

FreeNAS® 11.1-U4 User Guide

March 2018 Edition

FreeNAS® is © 2011-2018 iXsystems

FreeNAS® and the FreeNAS® logo are registered trademarks of iXsystems

FreeBSD® is a registered trademark of the FreeBSD Foundation

Written by users of the FreeNAS® network-attached storage operating system.

Version 11.1

Copyright © 2011-2018 iXsystems (<https://www.ixsystems.com/>)

Welcome	1
Typographic Conventions	2
1 Introduction	3
1.1 New Features in 11.1	3
1.2 Changes Since 11.1	5
1.2.1 U1	5
1.2.2 U2	6
1.2.3 U3	6
1.3 Path and Name Lengths	6
1.4 Hardware Recommendations	6
1.4.1 RAM	7
1.4.2 The Operating System Device	7
1.4.3 Storage Disks and Controllers	8
1.4.4 Network Interfaces	9
1.5 Getting Started with ZFS	10
2 Installing and Upgrading	11
2.1 Getting FreeNAS®	11
2.2 Preparing the Media	11
2.2.1 On FreeBSD or Linux	12
2.2.2 On Windows	12
2.2.3 On OS X	13
2.3 Performing the Installation	13
2.4 Installation Troubleshooting	21
2.5 Upgrading	22
2.5.1 Caveats	22
2.5.2 Initial Preparation	22
2.5.3 Upgrading Using the ISO	23
2.5.4 Upgrading From the GUI	26
2.5.5 If Something Goes Wrong	26
2.5.6 Upgrading a ZFS Pool	27
2.6 Virtualization	28
2.6.1 VirtualBox	29
2.6.2 VMware ESXi	37
3 Booting	44
3.1 Obtaining an IP Address	45
3.2 Logging In	46
3.3 Initial Configuration	47
4 Account	48
4.1 Groups	48
4.2 Users	51

5	System	55
5.1	Information	55
5.2	General	56
5.3	Boot	59
5.3.1	Mirroring the Boot Device	62
5.4	Advanced	63
5.4.1	Autotune	65
5.5	Email	65
5.6	System Dataset	67
5.7	Tunables	68
5.8	Update	70
5.8.1	Preparing for Updates	70
5.8.2	Updates and Trains	70
5.8.3	Checking for Updates	72
5.8.4	Applying Updates	72
5.8.5	Manual Updates	73
5.9	Cloud Credentials	73
5.10	Alert Services	74
5.10.1	How it Works	75
5.11	CAs	75
5.12	Certificates	78
5.13	Support	81
6	Tasks	84
6.1	Cloud Sync	84
6.1.1	Cloud Sync Example	87
6.2	Cron Jobs	89
6.3	Init/Shutdown Scripts	91
6.4	Rsync Tasks	92
6.4.1	Rsync Module Mode	96
6.4.2	Rsync over SSH Mode	96
6.5	S.M.A.R.T. Tests	99
7	Network	102
7.1	Global Configuration	102
7.2	Interfaces	104
7.3	IPMI	106
7.4	Link Aggregations	108
7.4.1	LACP, MPIO, NFS, and ESXi	108
7.4.2	Creating a Link Aggregation	109
7.5	Network Summary	112
7.6	Static Routes	112
7.7	VLANs	113
8	Storage	115
8.1	Volumes	115
8.1.1	Volume Manager	115
	Encryption	117
	Encryption Performance	118
	Manual Setup	118
	Extending a ZFS Volume	119
8.1.2	Change Permissions	120
8.1.3	Create Dataset	122
	Deduplication	124
	Compression	125
8.1.4	Create zvol	125
8.1.5	Import Disk	127
8.1.6	Import Volume	127

Importing an Encrypted Pool	128
8.1.7 View Disks	129
8.1.8 Volumes	131
Managing Encrypted Volumes	133
Additional Controls for Encrypted Volumes	134
8.1.9 View Multipaths	136
8.1.10 Replacing a Failed Drive	137
Replacing an Encrypted Drive	139
Removing a Log or Cache Device	139
8.1.11 Replacing Drives to Grow a ZFS Pool	139
8.1.12 Hot Spares	140
8.2 Periodic Snapshot Tasks	140
8.3 Replication Tasks	142
8.3.1 Examples: Common Configuration	142
Alpha (Source)	142
Beta (Destination)	143
8.3.2 Example: FreeNAS [®] to FreeNAS [®] Semi-Automatic Setup	143
8.3.3 Example: FreeNAS [®] to FreeNAS [®] Dedicated User Replication	145
8.3.4 Example: FreeNAS [®] to FreeNAS [®] or Other Systems, Manual Setup	146
Encryption Keys	146
8.3.5 Replication Options	149
8.3.6 Replication Encryption	150
8.3.7 Limiting Replication Times	150
8.3.8 Troubleshooting Replication	150
SSH	150
Compression	151
Manual Testing	151
8.4 Resilver Priority	151
8.5 Scrubs	152
8.6 Snapshots	155
8.7 VMware-Snapshot	157
9 Directory Services	158
9.1 Active Directory	158
9.1.1 Troubleshooting Tips	162
9.1.2 If the System Will not Join the Domain	163
9.2 LDAP	163
9.3 NIS	165
9.4 Kerberos Realms	167
9.5 Kerberos Keytabs	167
9.6 Kerberos Settings	168
10 Sharing	169
10.1 Apple (AFP) Shares	170
10.1.1 Creating AFP Guest Shares	172
10.1.2 Creating Authenticated and Time Machine Shares	173
10.2 Unix (NFS) Shares	177
10.2.1 Example Configuration	181
10.2.2 Connecting to the Share	181
From BSD or Linux	181
From Microsoft	182
From Mac OS X	182
10.2.3 Troubleshooting NFS	183
10.3 WebDAV Shares	184
10.4 Windows (SMB) Shares	185
10.4.1 Configuring Unauthenticated Access	190
10.4.2 Configuring Authenticated Access Without a Domain Controller	191

10.4.3	Configuring Shadow Copies	194
10.5	Block (iSCSI)	195
10.5.1	Target Global Configuration	196
10.5.2	Portals	197
10.5.3	Initiators	199
10.5.4	Authorized Accesses	200
10.5.5	Targets	202
10.5.6	Extents	203
10.5.7	Target/Extents	206
10.5.8	Connecting to iSCSI	207
10.5.9	Growing LUNs	207
	Zvol Based LUN	208
	File Extent Based LUN	208
11	Services	209
11.1	Control Services	209
11.2	AFP	211
11.2.1	Troubleshooting AFP	213
11.3	Domain Controller	213
11.3.1	Samba Domain Controller Backup	215
11.4	Dynamic DNS	215
11.5	FTP	216
11.5.1	Anonymous FTP	219
11.5.2	FTP in chroot	220
11.5.3	Encrypting FTP	221
11.5.4	Troubleshooting FTP	221
11.6	iSCSI	221
11.7	LLDP	222
11.8	Netdata	222
11.9	NFS	223
11.10	Rsync	225
11.10.1	Configure Rsyncd	225
11.10.2	Rsync Modules	226
11.11	S3	227
11.12	S.M.A.R.T.	228
11.13	SMB	230
11.13.1	Troubleshooting SMB	233
11.14	SNMP	234
11.15	SSH	236
11.15.1	SCP Only	237
11.15.2	Troubleshooting SSH	238
11.16	TFTP	238
11.17	UPS	239
11.17.1	Multiple Computers with One UPS	242
11.18	WebDAV	242
12	Plugins	244
12.1	Installing Plugins	244
12.2	Updating Plugins	248
12.3	Uploading Plugins	248
12.4	Deleting Plugins	248
12.5	Available Plugins	249
13	Jails	251
13.1	Jails Configuration	252
13.2	Adding Jails	253
13.2.1	Managing Jails	255
	Accessing a Jail Using SSH	257

Add Storage	257
13.2.2 Installing FreeBSD Packages	260
13.2.3 Compiling FreeBSD Ports	261
13.2.4 Starting Installed Software	263
13.3 Managing Jail Templates	264
13.4 Using iocage	267
13.4.1 Managing iocage Jails	267
14 Virtual Machines	270
14.1 Creating VMs	270
14.2 Adding Devices to a VM	271
14.2.1 Network Interfaces	272
14.2.2 Disk Devices	273
14.2.3 Raw Files	273
14.2.4 CD-ROM Devices	274
14.2.5 VNC Interface	274
14.2.6 Virtual Serial Ports	275
14.3 Running VMs	276
14.4 Deleting VMs	276
14.5 Docker/Rancher VM	277
14.5.1 Rancher VM Requirements	277
14.5.2 Create the Rancher VM	277
14.5.3 Start the Rancher VM	279
14.5.4 Installing the Rancher Server	279
15 Reporting	280
16 Wizard	282
17 Display System Processes	289
18 Shell	290
19 Log Out	292
20 Reboot	293
21 Shutdown	294
22 Support Icon	295
23 Guide	296
24 Alert	297
25 Support Resources	299
25.1 Website and Social Media	299
25.2 Forums	299
25.3 IRC	300
25.4 Videos	300
25.5 Professional Support	301
26 Command Line Utilities	302
26.1 Iperf	302
26.2 Netperf	305
26.3 IOzone	306
26.4 arcstat	308
26.5 tw_cli	313
26.6 MegaCli	315
26.7 freenas-debug	315

26.8	tmux	316
26.9	Dmidecode	316
27	Contributing to FreeNAS®	318
27.1	Translation	318
28	ZFS Primer	322
29	VAAI	326
29.1	VAAI for iSCSI	326
30	Using the API	327
30.1	A Simple API Example	328
30.2	A More Complex Example	329
	Index	331

Welcome

This Guide covers the installation and use of FreeNAS® 11.1.

The FreeNAS® User Guide is a work in progress and relies on the contributions of many individuals. If you are interested in helping us to improve the Guide, read the instructions in the [README](https://github.com/freenas/freenas-docs/blob/master/README.md) (<https://github.com/freenas/freenas-docs/blob/master/README.md>). IRC Freenode users are welcome to join the #freenas channel where you will find other FreeNAS® users.

The FreeNAS® User Guide is freely available for sharing and redistribution under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/3.0/) (<https://creativecommons.org/licenses/by/3.0/>). This means that you have permission to copy, distribute, translate, and adapt the work as long as you attribute iXsystems as the original source of the Guide.

FreeNAS® and the FreeNAS® logo are registered trademarks of iXsystems.

Active Directory® is a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries.

Apple, Mac and Mac OS are trademarks of Apple Inc., registered in the U.S. and other countries.

Broadcom is a trademark of Broadcom Corporation.

Chelsio® is a registered trademark of Chelsio Communications.

Cisco® is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

Django® is a registered trademark of Django Software Foundation.

Facebook® is a registered trademark of Facebook Inc.

FreeBSD® and the FreeBSD® logo are registered trademarks of the FreeBSD Foundation®.

Intel, the Intel logo, Pentium Inside, and Pentium are trademarks of Intel Corporation in the U.S. and/or other countries.

LinkedIn® is a registered trademark of LinkedIn Corporation.

Linux® is a registered trademark of Linus Torvalds.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Twitter is a trademark of Twitter, Inc. in the United States and other countries.

UNIX® is a registered trademark of The Open Group.

VirtualBox® is a registered trademark of Oracle.

VMware® is a registered trademark of VMware, Inc.

Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.

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

Typographic Conventions

Typographic Conventions

The FreeNAS® 11.1 User Guide uses these typographic conventions:

Table 1: Text Format Examples

Item	Visual Example
Graphical elements: buttons, icons, fields, columns, and boxes	Click the Import CA button.
Menu selections	Select <code>System</code> → <code>Information</code> .
Commands	Use the <code>scp</code> command.
File names and volume and dataset names	Locate the <code>/etc/rc.conf</code> file.
Keyboard keys	Press the <code>Enter</code> key.
Important points	This is important.
Values entered into fields, or device names	Enter 127.0.0.1 in the address field.

INTRODUCTION

FreeNAS® is an embedded open source network-attached storage (NAS) operating system based on FreeBSD and released under a [2-clause BSD license](https://opensource.org/licenses/BSD-2-Clause) (<https://opensource.org/licenses/BSD-2-Clause>). A NAS has an operating system optimized for file storage and sharing.

FreeNAS® provides a browser-based, graphical configuration interface. The built-in networking protocols provide storage access to multiple operating systems. A plugin system is provided for extending the built-in features by installing additional software.

1.1 New Features in 11.1

FreeNAS® 11.1 is a feature release, which includes several new significant features, many improvements and bug fixes to existing features, and version updates to the operating system, base applications, and drivers. Users are encouraged to [Update](#) (page 70) to this release in order to take advantage of these improvements and bug fixes.

These base applications and drivers have been updated or added:

- The base operating system has been updated to FreeBSD 11.1-STABLE. This brings in many new [features and drivers](#) (<https://www.freebsd.org/releases/11.1R/relnotes.html>). Improvements have been made to the [em\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=en>), [ixl\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=ixl>), [ixgbe\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=ixgbe>), and [mps\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=mps>) drivers. Additionally, the [netmap\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=netmap>) kernel module has been added to the build as some NIC drivers depend upon it.
- There have been many improvements to OpenZFS, including a significant speed improvement when listing a large number of snapshots and deleting multiple snapshots or large files.
- The algorithm used for scrubs and resilvers has received many improvements which are most noticeable on fragmented pools.
- Samba has been patched to address [these security vulnerabilities](#) (<https://www.samba.org/samba/history/samba-4.7.3>).
- The Dojo Toolkit has been updated to version 1.12.2.
- OpenVPN has been updated to version [2.4.3](#) (<https://github.com/OpenVPN/openvpn/blob/release/2.4/Changes.rst#version-243>).
- [Iperf version 3.2](#) (<http://software.es.net/iperf/>) has been added. To use this version, specify `iperf3` instead of `iperf`.
- `locage` has been updated to version 0.9.10.
- The new middleware now uses Python `asyncio` which simplifies asynchronous code and makes it more readable.
- The SNMP MIB has many improvements, including the ability to send SNMP traps for new alerts.
- The system now sends an email when a scrub finishes.
- `mmv` (<https://packages.debian.org/unstable/utils/mmv>) has been added. It can be used from the command line to safely move or copy multiple files using patterns, without any unexpected deletion of files due to target name collisions.

- [s3cmd](http://s3tools.org/s3cmd) (<http://s3tools.org/s3cmd>) has been added back as a CLI alternative to [S3](#) (page 227).
- The [zfs-stats](http://www.vx.sk/zfs-stats/) (<http://www.vx.sk/zfs-stats/>) CLI utility has been added. Type `zfs-stats` to see command usage.
- The hardware watchdog has been reenabled for recent firmware versions of AsrockRack C2750D4I. The BMC bug which required the watchdog to be disabled is resolved with the 00.35.00 or newer BMC firmware version.
- The system issues an alert if the system reboots itself.

These major features are new in this version:

- Scrubs can be paused and resumed from the command line. Scrub pause state and progress are periodically synced to disk. If the system is restarted or the pool is exported during a paused scrub, the scrub remains paused until it is resumed. When resumed, the scrub picks up from the place where it was last checkpointed to disk. Paused scrubs can be resumed with `zpool scrub`. Scrubs can be paused manually with `zpool scrub -p`. A future version of FreeNAS® will add a button to the UI to resume or pause a scrub.
- [Cloud Credentials](#) (page 73) has been added to [System](#) (page 55). This can be used to provide a secure connection to a cloud service providers. Supported services include Amazon S3, Azure Blob Storage, Backblaze B2, and Google Cloud Storage.
- [Cloud Sync](#) (page 84) has been added to [Tasks](#) (page 84) and can be used to synchronize files or directories to remote cloud storage providers with a specified transfer mode.
- The Server Side Encryption drop-down menu appears on [Tasks](#) → [Cloud Sync](#) → [Add Cloud Sync](#) when an S3 provider is selected.
- [Resilver Priority](#) (page 151) has been added to [Storage](#) (page 115). This provides the ability to run resilvering at a higher priority at configurable times and days of the week.
- The [Netdata](#) (page 222) real-time performance and monitoring system has been added to [Services](#) (page 209).
- [Virtual Machines](#) (page 270) have received significant improvements, including:
 - support for non-US keyboards.
 - the ability to restart a VM and to clone a VM.
 - the ability to specify the NIC used by the VM as well as the MAC address for the VM NIC. These options can be set with [VMs](#) → [Devices](#) → [Network Interface](#).
 - the ability to specify the sector size used by the emulated disk has been added to [VMs](#) → [Devices](#) → [Disk](#).
 - the ability to edit the VNC screen resolution, select the IP address to bind to, set the VNC password, and select the option to use the Web version of VNC. These options can be set with [VMs](#) → [Devices](#) → [VNC](#).

These screens have changed:

- The Change E-mail button has been removed from [Account](#) → [Users](#).
- Each device in a mirrored boot pool now displays a Detach button in [System](#) → [Boot](#) → [Status](#). This can be used to remove a device from the boot pool.
- The Enable Console Menu in [System](#) → [Advanced](#) has been renamed to Show Text Console Without Password Prompt.
- The Report CPU usage in percentage checkbox has been added to [System](#) → [Advanced](#).
- The FreeNAS-11-Nightlies-SDK train has been added and the FreeNAS-9.3-STABLE train has been removed from [System](#) → [Update](#).
- The Send Test Alert button has been added to [System](#) → [Alert Services](#) → [Edit](#).
- The Subject Alternate Names field has been added to [System](#) → [CAs](#) → [Create Internal CA](#), [System](#) → [CAs](#) → [Create Intermediate CA](#), [System](#) → [Certificates](#) → [Create Internal Certificate](#), and [System](#) → [Certificates](#) → [Create Certificate Signing Request](#) screens.
- The Sign CSR button has been added to [System](#) → [CAs](#).

- The ability to edit an existing certificate's Name, Certificate, and Private Key fields has been added to `System` → `Certificates` → `View`.
- An Enabled checkbox has been added to `Tasks` → `Init/Shutdown Scripts`.
- The Additional domains field has been added to `Network` → `Global Configuration`. This allows up to six additional DNS search domains with the caveat that adding more domains may negatively impact DNS lookup time.
- The Identify Light button has been added to `Network` → `IPMI` to make it easier to identify a system in a rack by flashing its IPMI LED light.
- The Priority Code Point (CoS) field has been added to `Network` → `VLANs` → `Add VLAN`. This can be useful in data-center environments to classify storage traffic on a given VLAN interface using IEEE 802.1p Class of Service (COS).
- The Read-Only drop-down menu has been added to `Storage` → `Datasets` → `Add Dataset` → `Advanced Mode`.
- The Promote Dataset button has been added to `Storage` → `Volumes`.
- The Replication column has been removed from `Storage` → `Snapshots`.
- The Time Machine Quota checkbox has been added to `Sharing` → `Apple (AFP) Shares` → `Add Apple (AFP) Share`.
- The Access Based Share Enumeration checkbox has been added to `Sharing` → `SMB (Windows) Shares` → `Add SMB (Windows) Share`.
- The Home Share Time Machine checkbox has been added to `Services` → `AFP`.
- The CheckIP Server SSL, CheckIP Server, CheckIP Path, and Use SSL fields have been added to `Services` → `DDNS`. The Forced update period and Auxiliary parameters fields have been removed. In addition, several dozen DDNS providers have been added to the Provider drop-down menu.
- The Certificate drop-down menu has been added to `Services` → `S3` in order to configure encrypted S3 connections.
- The Server minimum protocol and Server maximum protocol fields have been removed from `Services` → `SMB`.
- The Log Level drop-down menu has been added to `Services` → `SNMP`. It defaults to the Error log level.
- The No Communication Warning Time field has been added to `Services` → `UPS`. This can be used to configure the frequency of email notifications during the loss of UPS communications.
- The No Authentication choice has been added to the `Services` → `WebDAV` → `HTTP Authentication` drop-down menu.

1.2 Changes Since 11.1

FreeNAS® uses a “rolling release” model instead of point releases. The [Update](#) (page 70) mechanism makes it easy to keep up-to-date with the latest security fixes, bug fixes, and new features. Some updates affect the user interface, so this section lists any functional changes that have occurred since 11.1 was released.

Note: The screenshots in this documentation assume that the system has been fully updated to the latest STABLE version of FreeNAS® 11.1-U4. If a screen on the system is not the same as shown in this guide, make sure that all updates have been applied.

1.2.1 U1

- RancherOS has been updated to version [1.1.3](https://github.com/rancher/os/releases/tag/v1.1.3) (<https://github.com/rancher/os/releases/tag/v1.1.3>).
- Smartmontools has been updated to [6.6](https://www.smartmontools.org/browser/tags/RELEASE_6_6/smartmontools/NEWS) (https://www.smartmontools.org/browser/tags/RELEASE_6_6/smartmontools/NEWS).
- The Factory Restore button in `System` → `General` has been renamed to Reset Configuration to Defaults.

1.2.2 U2

- System Update has been removed from the [Console Setup Menu](#) (page 44).
- System -> Information shows the system serial number on systems supplied or certified by iXsystems.
- ZFS Sync write synchronization settings have been added to [Create Dataset](#) (page 122) and [Create zvol](#) (page 125).
- Midnight Commander has been updated to version 4.8.20 (<http://midnight-commander.org/wiki/NEWS-4.8.20>).

1.2.3 U3

- Samba has been patched to address [CVE-2018-1050](https://www.samba.org/samba/security/CVE-2018-1050.html) (<https://www.samba.org/samba/security/CVE-2018-1050.html>) and [CVE-2018-1057](https://wiki.samba.org/index.php/CVE-2018-1057) (<https://wiki.samba.org/index.php/CVE-2018-1057>).

1.3 Path and Name Lengths

Names of files, directories, and devices are subject to some limits imposed by the FreeBSD operating system. The limits shown here are for names using plain-text characters that each occupy one byte of space. Some UTF-8 characters take more than a single byte of space, and using those characters reduces these limits proportionally. System overhead can also reduce the length of these limits by one or more bytes.

Table 1.1: Path and Name Lengths

Type	Maximum Length	Description
File Paths	1024 bytes	Total file path length (PATH_MAX). The full path includes directory separator slash characters, subdirectory names, and the name of the file itself. For example, the path <code>/mnt/tank/mydataset/mydirectory/myfile.txt</code> is 42 bytes long. Using very long file or directory names can be problematic. A complete path with long directory and file names can exceed the 1024-byte limit, preventing direct access to that file until the directory names or filename are shortened or the file is moved into a directory with a shorter total path length.
File and Directory Names	255 bytes	Individual directory or file name length (NAME_MAX).
Mounted Filesystem Paths	88 bytes	Mounted filesystem path length (MNAMELEN). Longer paths can prevent a device from being mounted.
Device Filesystem Paths	63 bytes	devfs(8) (https://www.freebsd.org/cgi/man.cgi?query=devfs&sektion=8) device path lengths (SPECNAMELEN). Longer paths can prevent a device from being created.

1.4 Hardware Recommendations

FreeNAS® 11.1 is based on FreeBSD 11.1 and supports the same hardware found in the [FreeBSD Hardware Compatibility List](http://www.freebsd.org/releases/11.1R/hardware.html) (<http://www.freebsd.org/releases/11.1R/hardware.html>). Supported processors are listed in section 2.1 [amd64](https://www.freebsd.org/releases/11.1R/hardware.html#proc) (<https://www.freebsd.org/releases/11.1R/hardware.html#proc>). FreeNAS® is only available for 64-bit processors. This architecture is called amd64 by AMD and Intel 64 by Intel.

Note: FreeNAS® boots from a GPT partition. This means that the system BIOS must be able to boot using either the legacy BIOS firmware interface or EFI.

Actual hardware requirements vary depending on the usage of the FreeNAS® system. This section provides some starter guidelines. The [FreeNAS® Hardware Forum](https://forums.freenas.org/index.php?forums/hardware.18/) (<https://forums.freenas.org/index.php?forums/hardware.18/>) has performance tips from FreeNAS® users and is a place to post questions regarding the hardware best suited to meet specific requirements. [Hardware Recommendations](https://forums.freenas.org/index.php?resources/hardware-recommendations-guide.12/) (<https://forums.freenas.org/index.php?resources/hardware-recommendations-guide.12/>) gives detailed recommendations for system components, with the [FreeNAS® Quick Hardware Guide](https://forums.freenas.org/index.php?resources/freenas-quick-hardware-guide.7/) (<https://forums.freenas.org/index.php?resources/freenas-quick-hardware-guide.7/>) providing short lists of components for various configurations. [Building, Burn-In, and Testing your FreeNAS® system](https://forums.freenas.org/index.php?threads/building-burn-in-and-testing-your-freenas-system.17750/) (<https://forums.freenas.org/index.php?threads/building-burn-in-and-testing-your-freenas-system.17750/>) has detailed instructions on testing new hardware.

1.4.1 RAM

The best way to get the most out of a FreeNAS® system is to install as much RAM as possible. More RAM allows ZFS to provide better performance. The [FreeNAS® Forums](https://forums.freenas.org/index.php) (<https://forums.freenas.org/index.php>) provide anecdotal evidence from users on how much performance can be gained by adding more RAM.

General guidelines for RAM:

- A minimum of 8 GB of RAM is required.
Additional features require additional RAM, and large amounts of storage require more RAM for cache. An old, somewhat overstated guideline is 1 GB of RAM per terabyte of disk capacity.
- To use Active Directory with many users, add an additional 2 GB of RAM for the winbind internal cache.
- For iSCSI, install at least 16 GB of RAM if performance is not critical, or at least 32 GB of RAM if good performance is a requirement.
- [Jails](#) (page 251) are very memory-efficient, but can still use memory that would otherwise be available for ZFS. If the system will be running many jails, or a few resource-intensive jails, adding 1 to 4 additional gigabytes of RAM can be helpful. This memory is shared by the host and will be used for ZFS when not being used by jails.
- [Virtual Machines](#) (page 270) require additional RAM beyond any amounts listed here. Memory used by virtual machines is not available to the host while the VM is running, and is not included in the amounts described above. For example, a system that will be running two VMs that each need 1 GB of RAM requires an additional 2 GB of RAM.
- When installing FreeNAS® on a headless system, disable the shared memory settings for the video card in the BIOS.
- For ZFS deduplication, ensure the system has at least 5 GB of RAM per terabyte of storage to be deduplicated.

If the hardware supports it, install ECC RAM. While more expensive, ECC RAM is highly recommended as it prevents in-flight corruption of data before the error-correcting properties of ZFS come into play, thus providing consistency for the checksumming and parity calculations performed by ZFS. If your data is important, use ECC RAM. This [Case Study](http://research.cs.wisc.edu/adsl/Publications/zfs-corruption-fast10.pdf) (<http://research.cs.wisc.edu/adsl/Publications/zfs-corruption-fast10.pdf>) describes the risks associated with memory corruption.

Do not use FreeNAS® to store data without at least 8 GB of RAM. Many users expect FreeNAS® to function with less memory, just at reduced performance. The bottom line is that these minimums are based on feedback from many users. Requests for help in the forums or IRC are sometimes ignored when the installed system does not have at least 8 GB of RAM because of the abundance of information that FreeNAS® may not behave properly with less memory.

1.4.2 The Operating System Device

The FreeNAS® operating system is installed to at least one device that is separate from the storage disks. The device can be a SSD, USB memory stick, or DOM (Disk on Module). Installation to a hard drive is discouraged as that drive is then not available for data storage.

Note: To write the installation file to a USB stick, two USB ports are needed, each with an inserted USB device. One USB stick contains the installer, while the other USB stick is the destination for the FreeNAS® installation. Be careful to select the correct USB device for the FreeNAS® installation. FreeNAS® cannot be installed onto the same device that contains the

installer. After installation, remove the installer USB stick. It might also be necessary to adjust the BIOS configuration to boot from the new FreeNAS® boot device.

When determining the type and size of the target device where FreeNAS® is to be installed, keep these points in mind:

- The absolute bare minimum size is 8 GB. That does not provide much room. The recommended minimum is 16 GB. This provides room for the operating system and several boot environments created by updates. More space provides room for more boot environments and 32 GB or more is preferred.
- SSDs (Solid State Disks) are fast and reliable, and make very good FreeNAS® operating system devices. Their one disadvantage is that they require a disk connection which might be needed for storage disks.

Even a relatively large SSD (120 or 128 GB) is useful as a boot device. While it might appear that the unused space is wasted, that space is instead used internally by the SSD for wear leveling. This makes the SSD last longer and provides greater reliability.
- When planning to add your own boot environments, budget about 1 GB of storage per boot environment. Consider deleting older boot environments after making sure they are no longer needed. Boot environments can be created and deleted using `System → Boot`.
- Use quality, name-brand USB sticks, as ZFS will quickly reveal errors on cheap, poorly-made sticks.
- For a more reliable boot disk, use two identical devices and select them both during the installation. This will create a mirrored boot device.

Note: Current versions of FreeNAS® run directly from the operating system device. Early versions of FreeNAS® ran from RAM, but that has not been the case for years.

1.4.3 Storage Disks and Controllers

The [Disk section](http://www.freebsd.org/releases/11.1R/hardware.html#DISK) (<http://www.freebsd.org/releases/11.1R/hardware.html#DISK>) of the FreeBSD Hardware List lists the supported disk controllers. In addition, support for 3ware 6 Gbps RAID controllers has been added along with the CLI utility `tw_cli` for managing 3ware RAID controllers.

FreeNAS® supports hot pluggable drives. Using this feature requires enabling AHCI in the BIOS.

Reliable disk alerting and immediate reporting of a failed drive can be obtained by using an HBA such as an Broadcom MegaRAID controller or a 3Ware twa-compatible controller.

Note: Upgrading the firmware of Broadcom SAS HBAs to the latest version is recommended.

Some Highpoint RAID controllers do not support pass-through of S.M.A.R.T. data or other disk information, potentially including disk serial numbers. It is best to use a different disk controller with FreeNAS®.

Note: The system is configured to prefer the `mrsas(4)` (<https://www.freebsd.org/cgi/man.cgi?query=mrsas>) driver for controller cards like the Dell PERC H330 and H730 which are supported by several drivers. Although not recommended, the `mfi(4)` (<https://www.freebsd.org/cgi/man.cgi?query=mfi>) driver can be used instead by removing the loader `Tunable` (page 68): `hw.mfi.mrsas_enable` or setting the Value to 0.

Suggestions for testing disks before adding them to a RAID array can be found in this [forum post](https://forums.freenas.org/index.php?threads/checking-new-hdds-in-raid.12082/#post-55936) (<https://forums.freenas.org/index.php?threads/checking-new-hdds-in-raid.12082/#post-55936>). Additionally, `badblocks` (<https://linux.die.net/man/8/badblocks>) is installed with FreeNAS® for testing disks.

If the budget allows optimization of the disk subsystem, consider the read/write needs and RAID requirements:

- For steady, non-contiguous writes, use disks with low seek times. Examples are 10K or 15K SAS drives which cost about \$1/GB. An example configuration would be six 600 GB 15K SAS drives in a RAID 10 which would yield 1.8 TB of usable space, or eight 600 GB 15K SAS drives in a RAID 10 which would yield 2.4 TB of usable space.

For ZFS, [Disk Space Requirements for ZFS Storage Pools](http://docs.oracle.com/cd/E19253-01/819-5461/6n7ht6r12/index.html) (<http://docs.oracle.com/cd/E19253-01/819-5461/6n7ht6r12/index.html>) recommends a minimum of 16 GB of disk space. Due to the way that ZFS creates swap, it is not possible to format less than 3 GB of space with ZFS. However, on a drive that is below the minimum recommended size, a fair amount of storage space is lost to swap: for example, on a 4 GB drive, 2 GB will be reserved for swap.

Users new to ZFS who are purchasing hardware should read through [ZFS Storage Pools Recommendations](https://web.archive.org/web/20161028084224/http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide#ZFS_Storage_Pools) (https://web.archive.org/web/20161028084224/http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide#ZFS_Storage_Pools) first.

ZFS vdevs, groups of disks that act like a single device, can be created using disks of different sizes. However, the capacity available on each disk is limited to the same capacity as the smallest disk in the group. For example, a vdev with one 2 TB and two 4 TB disks will only be able to use 2 TB of space on each disk. In general, use disks that are the same size for the best space usage and performance.

The [ZFS Drive Size and Cost Comparison spreadsheet](https://forums.freenas.org/index.php?threads/zfs-drive-size-and-cost-comparison-spreadsheet.38092/) (<https://forums.freenas.org/index.php?threads/zfs-drive-size-and-cost-comparison-spreadsheet.38092/>) is available to compare usable space provided by different quantities and sizes of disks.

1.4.4 Network Interfaces

The [Ethernet section](http://www.freebsd.org/releases/11.1R/hardware.html#ethernet) (<http://www.freebsd.org/releases/11.1R/hardware.html#ethernet>) of the FreeBSD Hardware Notes indicates which interfaces are supported by each driver. While many interfaces are supported, FreeNAS® users have seen the best performance from Intel and Chelsio interfaces, so consider these brands when purchasing a new NIC. Realtek cards often perform poorly under CPU load as interfaces with these chipsets do not provide their own processors.

At a minimum, a GigE interface is recommended. While GigE interfaces and switches are affordable for home use, modern disks can easily saturate their 110 MB/s throughput. For higher network throughput, multiple GigE cards can be bonded together using the LACP type of [Link Aggregations](#) (page 108). The Ethernet switch must support LACP, which means a more expensive managed switch is required.

When network performance is a requirement and there is some money to spend, use 10 GigE interfaces and a managed switch. Managed switches with support for LACP and jumbo frames are preferred, as both can be used to increase network throughput. Refer to the [10 Gig Networking Primer](https://forums.freenas.org/index.php?threads/10-gig-networking-primer.25749/) (<https://forums.freenas.org/index.php?threads/10-gig-networking-primer.25749/>) for more information.

Note: At present, these are not supported: InfiniBand, FibreChannel over Ethernet, or wireless interfaces.

Both hardware and the type of shares can affect network performance. On the same hardware, SMB is slower than FTP or NFS because Samba is [single-threaded](#) (<https://www.samba.org/samba/docs/man/Samba-Developers-Guide/architecture.html>). So a fast CPU can help with SMB performance.

Wake on LAN (WOL) support depends on the FreeBSD driver for the interface. If the driver supports WOL, it can be enabled using [ifconfig\(8\)](#) (<http://www.freebsd.org/cgi/man.cgi?query=ifconfig>). To determine if WOL is supported on a particular interface, use the interface name with the following command. In this example, the capabilities line indicates that WOL is supported for the re0 interface:

```
ifconfig -m re0
re0: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
    options=42098<VLAN_MTU,VLAN_HWTAGGING,VLAN_HWCSUM,WOL_MAGIC,VLAN_HWTSO>
    capabilities=5399b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM, TSO4, WOL_UCAST, WOL_MCAST, ↵
↵ WOL_MAGIC, VLAN_HWFILTER, VLAN_H  WTSO>
```

If WOL support is shown but not working for a particular interface, create a bug report using the instructions in [Support](#) (page 81).

1.5 Getting Started with ZFS

Readers new to ZFS should take a moment to read the [ZFS Primer](#) (page 322).

INSTALLING AND UPGRADING

Please note that the FreeNAS® operating system must be installed on a separate device from the drives which hold the storage data. In other words, with only one disk drive, the FreeNAS® graphical interface is available, but there is no place to store any data. And storing data is, after all, the whole point of a NAS system. Home users experimenting with FreeNAS® can install FreeNAS® on an inexpensive USB thumb drive and use the computer's disks for storage.

This section describes:

- [Getting FreeNAS®](#) (page 11)
- [Preparing the Media](#) (page 11)
- [Performing the Installation](#) (page 13)
- [Installation Troubleshooting](#) (page 21)
- [Upgrading](#) (page 22)
- [Virtualization](#) (page 28)

2.1 Getting FreeNAS®

The latest STABLE version of FreeNAS® 11.1 can be downloaded from <https://download.freenas.org/stable/>.

Note: FreeNAS® requires 64-bit hardware.

The download page contains an .iso file. This is a bootable installer that can be written to either a CD or USB flash as described in [Preparing the Media](#) (page 11).

The .iso file has an associated `sha256.txt` file which is used to verify the integrity of the downloaded file. The command to verify the checksum varies by operating system:

- on a BSD system use the command `sha256 name_of_file`
- on a Linux system use the command `sha256sum name_of_file`
- on a Mac system use the command `shasum -a 256 name_of_file`
- Windows or Mac users can install additional utilities like [HashCalc](http://www.slavasoft.com/hashcalc/) (<http://www.slavasoft.com/hashcalc/>) or [HashTab](http://implbits.com/products/hashtab/) (<http://implbits.com/products/hashtab/>)

The value produced by running the command must match the value shown in the `sha256.txt` file. Checksum values that do not match indicate a corrupted installer file that should not be used.

2.2 Preparing the Media

The FreeNAS® installer can run from either a CD or a USB memory stick.

A CD burning utility is needed to write the `.iso` file to a CD.

The `.iso` file can also be written to a USB memory stick. The method used to write the file depends on the operating system. Examples for several common operating systems are shown below.

Note: To install from a USB stick to another USB stick, two USB ports are needed, each with an inserted USB device. One USB stick contains the installer. The other USB stick is the destination for the FreeNAS® installation. Take care to select the correct USB device for the FreeNAS® installation. It is not possible to install FreeNAS® onto the same USB stick containing the installer. After installation, remove the installer USB stick. It might also be necessary to adjust the BIOS configuration to boot from the new FreeNAS® USB stick.

Make sure that the boot device order in the BIOS is set to boot from the device containing the FreeNAS® installer media, then boot the system to start the installation.

2.2.1 On FreeBSD or Linux

On a FreeBSD or Linux system, the `dd` command can be used to write the `.iso` file to an inserted USB thumb drive. This example demonstrates writing the image to the first USB device (`/dev/da0`) connected to a FreeBSD system. Substitute the filename of the `.iso` file and the device name of the device to write.

Warning: The `dd` command is very powerful and can destroy any existing data on the specified device. Make absolutely sure of the device name to write to and do not mistype the device name when using `dd`! If you are uncomfortable using this command, write the `.iso` file to a CD instead.

```
dd if=FreeNAS-11.0-RELEASE.iso of=/dev/da0 bs=64k
6117+0 records in
6117+0 records out
400883712 bytes transferred in 88.706398 secs (4519220 bytes/sec)
```

When using the `dd` command:

- `if=` refers to the input file, or the name of the file to write to the device.
- `of=` refers to the output file; in this case, the device name of the flash card or removable USB drive. Note that USB device numbers are dynamic, and the target device might be `da1` or `da2` or another name depending on which devices are attached. Before attaching the target USB drive, use `ls /dev/da*`. Then attach the target USB drive, wait ten seconds, and run `ls /dev/da*` again to see the new device name and number of the target USB drive. On Linux, use `/dev/sdX`, where `X` refers to the letter of the USB device.
- `bs=` refers to the block size, the amount of data to write at a time. The larger 64K block size shown here helps speed up writes to the USB drive.

2.2.2 On Windows

Microsoft provides the USB/DVD Download Tool to create a USB bootable image from an `.iso` file. Follow [these instructions](https://www.microsoft.com/en-us/download/windows-usb-dvd-download-tool) (<https://www.microsoft.com/en-us/download/windows-usb-dvd-download-tool>), but enter the name of the downloaded `.iso` into the SOURCE FILE box.

[Image Writer](https://launchpad.net/win32-image-writer/) (<https://launchpad.net/win32-image-writer/>) and [Rufus](http://rufus.akeo.ie/) (<http://rufus.akeo.ie/>) are alternate programs for writing images to USB sticks on a computer running Windows. When using Rufus, check Create a bootable disk using and select DD Image from the drop-down menu.

2.2.3 On OS X

Insert the USB thumb drive. In the Finder, go to Applications → Utilities → Disk Utility. Unmount any mounted partitions on the USB thumb drive. Check that the USB thumb drive has only one partition, or partition table errors will be shown on boot. If needed, use Disk Utility to set up one partition on the USB drive. Selecting Free space when creating the partition works fine.

Determine the device name of the inserted USB thumb drive. From TERMINAL, navigate to the Desktop, then type this command:

```
diskutil list
/dev/disk0

#      TYPE NAME              SIZE               IDENTIFIER
0:      GUID_partition_scheme   *500.1 GB         disk0
1:      EFI                     209.7 MB          disk0s1
2:      Apple_HFS Macintosh HD   499.2 GB          disk0s2
3:      Apple_Boot Recovery HD   650.0 MB          disk0s3

/dev/disk1

#      TYPE NAME              SIZE               IDENTIFIER
0:      FDisk_partition_scheme  *8.0 GB           disk1
1:      DOS_FAT_32 UNTITLED      8.0 GB           disk1s1
```

This shows which devices are available to the system. Locate the target USB stick and record the path. If you are not sure which path is the correct one for the USB stick, remove the device, run the command again, and compare the difference. Once sure of the device name, navigate to the Desktop from TERMINAL, unmount the USB stick, and use the `dd` command to write the image to the USB stick. In this example, the USB thumb drive is `/dev/disk1`. It is first unmounted. The `dd` command is used to write the image to the faster “raw” version of the device (note the extra `r` in `/dev/rdisk1`). When running these commands, substitute the name of the installation file and the correct path to the USB thumb drive.

```
diskutil unmountDisk /dev/disk1
Unmount of all volumes on disk1 was successful

dd if=FreeNAS-11.0-RELEASE.iso of=/dev/rdisk1 bs=64k
```

Note: If the error “Resource busy” is shown when the `dd` command is run, go to Applications → Utilities → Disk Utility, find the USB thumb drive, and click on its partitions to make sure all of them are unmounted. If the error “dd: /dev/disk1: Permission denied” is shown, run the `dd` command by typing `sudo dd if=FreeNAS-11.0-RELEASE.iso of=/dev/rdisk1 bs=64k`. This will prompt for your password.

The `dd` command can take some minutes to complete. Wait until the prompt returns and a message is displayed with information about how long it took to write the image to the USB drive.

2.3 Performing the Installation

With the installation media inserted, boot the system from that media. The FreeNAS® installer GRUB menu is displayed as is shown in [Figure 2.1](#).

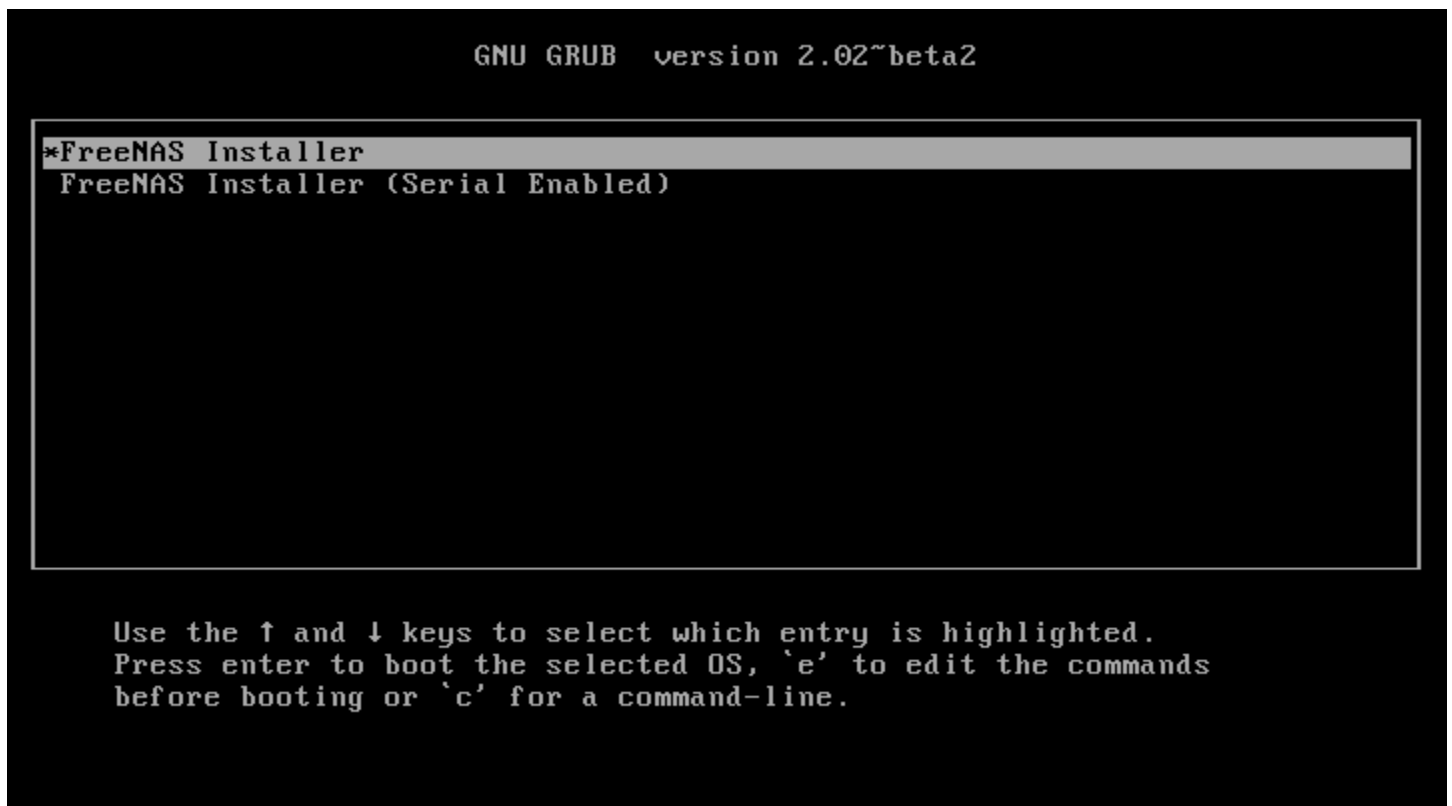


Fig. 2.1: Grub Menu

Tip: The Serial Enabled option is useful on systems which do not have a keyboard or monitor, but are accessed through a serial port, Serial over LAN, or [IPMI](#) (page 106).

Note: If the installer does not boot, verify that the installation device is listed first in the boot order in the BIOS. When booting from a CD, some motherboards may require connecting the CD device to SATA0 (the first connector) to boot from CD. If the installer stalls during bootup, double-check the SHA256 hash of the `.iso` file. If the hash does not match, re-download the file. If the hash is correct, burn the CD again at a lower speed or write the file to a different USB stick.

The installer will start automatically after a few seconds, or an option can be chosen by moving the highlight bar to it with the up and down arrow keys and pressing `Enter`. After booting, the installer menu is displayed as shown in [Figure 2.2](#).




Fig. 2.2: Installer Menu

Press `Enter` to select the default option, 1 Install/Upgrade. The next menu, shown in [Figure 2.3](#), lists all available drives. This includes any inserted USB thumb drives, which have names beginning with `da`.

Note: A minimum of 8 GB of RAM is required and the installer will present a warning message if less than 8 GB is detected.

In this example, the user is performing a test installation using VirtualBox and has created a 16 GB virtual disk to hold the operating system.

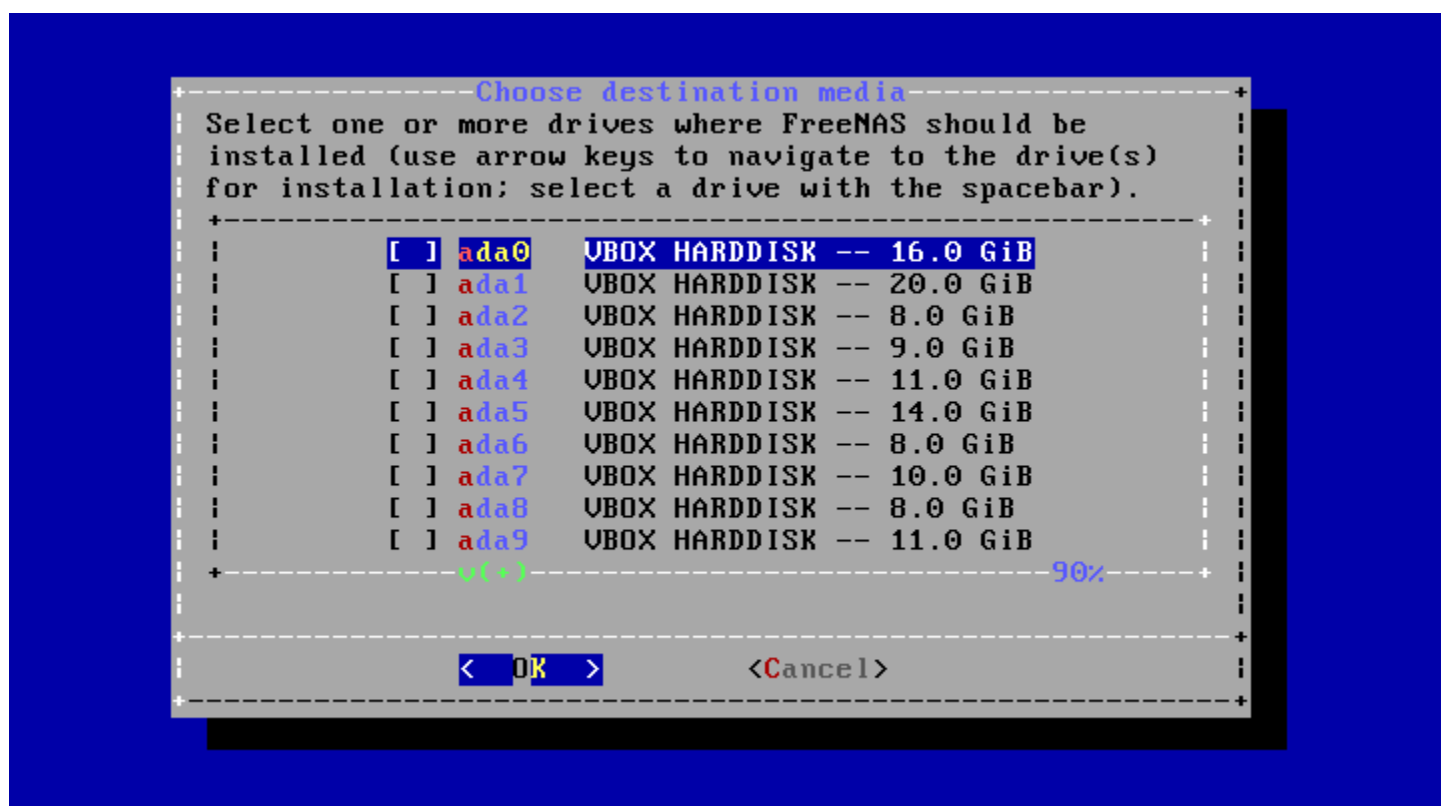


Fig. 2.3: Selecting the Install Drive

Use the arrow keys to highlight the destination USB drive, SSD, DOM (Disk on Module), or virtual disk. Press the `spacebar` to select it. To mirror the boot device, move to the second device and press `spacebar` to select it also. After making these selections, press `Enter`. The warning shown in [Figure 2.4](#) is displayed, a reminder not to install the operating system on a drive that is meant for storage. Press `Enter` to continue on to the screen shown in [Figure 2.6](#).

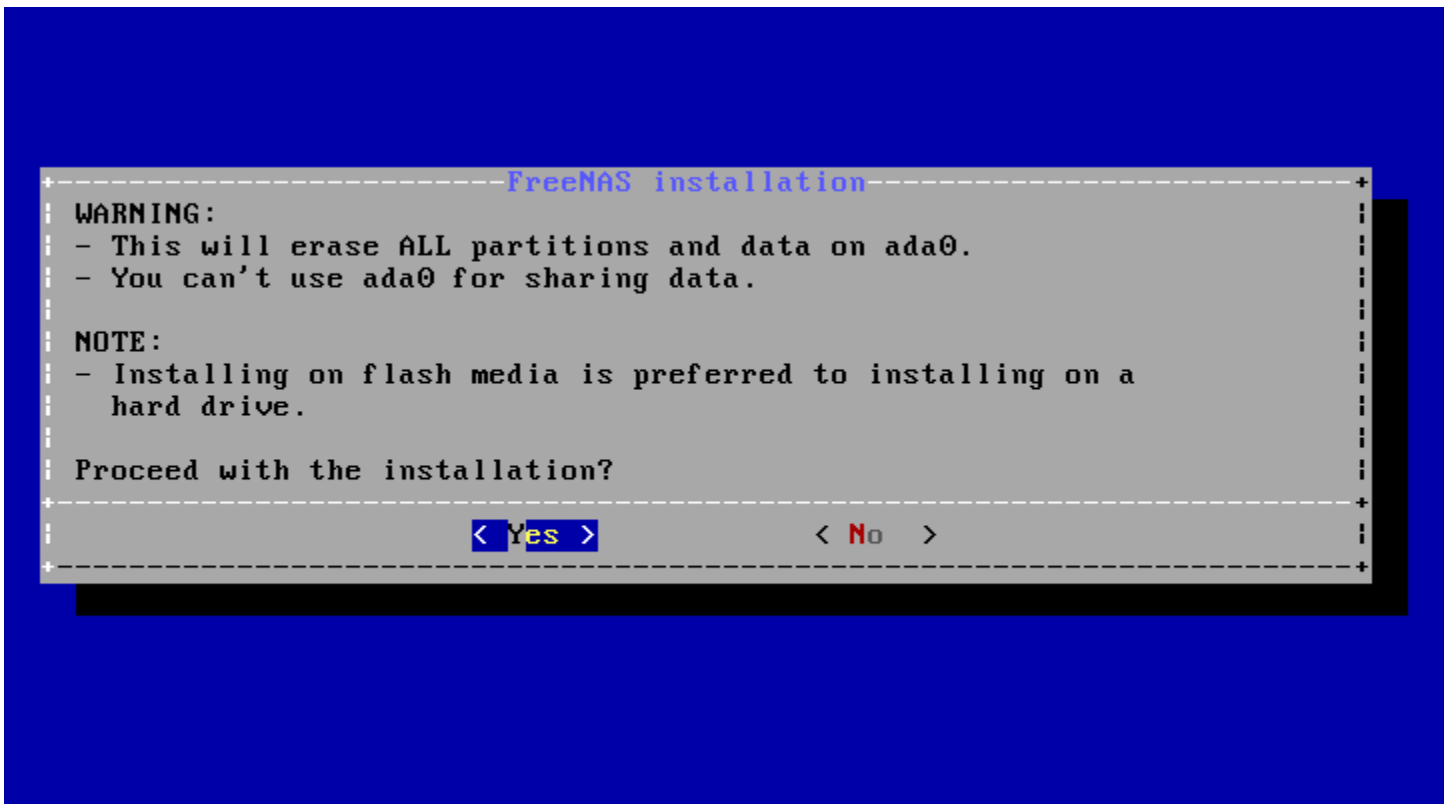


Fig. 2.4: Installation Warning

Note: A minimum of 8 GB of space on the boot device is required. However, 32 GB is recommended to provide room for future additions and boot environments. When using mirrored boot devices, it is best to use devices of the same size. If the device sizes are different, the mirror is limited to the size of the smaller device.

The installer recognizes existing installations of previous versions of FreeNAS® 8.x or 9.x. When an existing installation is present, the menu shown in [Figure 2.5](#) is displayed. To overwrite an existing installation, use the arrows to move to Fresh Install and press `Enter` twice to continue to the screen shown in [Figure 2.6](#).




Fig. 2.5: Performing a Fresh Install

The screen shown in [Figure 2.6](#) prompts for the root password which is used to log in to the administrative graphical interface.




Fig. 2.6: Set the Root Password

Setting a password is mandatory and the password cannot be blank. Since this password provides access to the administrative GUI, it should be hard to guess. Enter the password, press the down arrow key, and confirm the password. Then press `Enter` to continue with the installation.

Note: For security reasons, the SSH service and root SSH logins are disabled by default. Unless these are set, the only way to access a shell as root is to gain physical access to the console menu or to access the web shell within the administrative GUI. This means that the FreeNAS® system should be kept physically secure and that the administrative GUI should be behind a properly configured firewall and protected by a secure password.

FreeNAS® can be configured to boot with the standard BIOS boot mechanism or UEFI booting as shown [Figure 2.7](#). BIOS booting is recommended for legacy and enterprise hardware. UEFI is used on newer consumer motherboards.




Fig. 2.7: Choose UEFI or BIOS Booting

Note: Most UEFI systems can also boot in BIOS mode if CSM (Compatibility Support Module) is enabled in the UEFI setup screens.

The message in [Figure 2.8](#) is shown after the installation is complete.




Fig. 2.8: Installation Complete

Press `Enter` to return to the first menu, shown in [Figure 2.1](#). Highlight 3 Reboot System and press `Enter`. If booting from CD, remove the CDROM. As the system reboots, make sure that the device where FreeNAS® was installed is listed as the first boot entry in the BIOS so the system will boot from it. FreeNAS® boots into the Console Setup menu described in [Booting](#) (page 44).

2.4 Installation Troubleshooting

If the system does not boot into FreeNAS®, there are several things that can be checked to resolve the situation.

Check the system BIOS and see if there is an option to change the USB emulation from CD/DVD/floppy to hard drive. If it still will not boot, check to see if the card/drive is UDMA compliant.

If the system BIOS does not support EFI with BIOS emulation, see if it has an option to boot using legacy BIOS mode.

When the system starts to boot but hangs with this repeated error message:

```
run_interrupt_driven_hooks: still waiting after 60 seconds for xpt_config
```

go into the system BIOS and look for an onboard device configuration for a 1394 Controller. If present, disable that device and try booting again.

If the system starts to boot but hangs at a `mountroot>` prompt, follow the instructions in [Workaround/Semi-Fix for Mountroot Issues with 9.3](https://forums.freenas.org/index.php?threads/workaround-semi-fix-for-mountroot-issues-with-9-3.26071/) (<https://forums.freenas.org/index.php?threads/workaround-semi-fix-for-mountroot-issues-with-9-3.26071/>).

If the burned image fails to boot and the image was burned using a Windows system, wipe the USB stick before trying a second burn using a utility such as [Active@ KillDisk](http://how-to-erase-hard-drive.com/) (<http://how-to-erase-hard-drive.com/>). Otherwise, the second burn attempt will fail as Windows does not understand the partition which was written from the image file. Be very careful to specify the correct USB stick when using a wipe utility!

2.5 Upgrading

FreeNAS® provides flexibility for keeping the operating system up-to-date:

1. Upgrades to major releases, for example from version 9.3 to 9.10, can still be performed using either an ISO or the graphical administrative interface. Unless the Release Notes for the new major release indicate that the current version requires an ISO upgrade, either upgrade method can be used.
2. Minor releases have been replaced with signed updates. This means that it is not necessary to wait for a minor release to update the system with a system update or newer versions of drivers and features. It is also no longer necessary to manually download an upgrade file and its associated checksum to update the system.
3. The updater automatically creates a boot environment, making updates a low-risk operation. Boot environments provide the option to return to the previous version of the operating system by rebooting the system and selecting the previous boot environment from the boot menu.

This section describes how to perform an upgrade from an earlier version of FreeNAS® to 11.1. After 11.1 has been installed, use the instructions in [Update](#) (page 70) to keep the system updated.

2.5.1 Caveats

Be aware of these caveats before attempting an upgrade to 11.1:

- Warning: upgrading the ZFS pool can make it impossible to go back to a previous version. For this reason, the update process does not automatically upgrade the ZFS pool, though the [Alert](#) (page 297) system shows when newer feature flags are available for a pool. Unless a new feature flag is needed, it is safe to leave the pool at the current version and uncheck the alert. If the pool is upgraded, it will not be possible to boot into a previous version that does not support the newer feature flags.
- The [Wizard](#) (page 282) does not recognize an encrypted ZFS pool. If the ZFS pool is GELI-encrypted and the [Wizard](#) (page 282) starts after the upgrade, cancel the [Wizard](#) (page 282) and use the instructions in [Importing an Encrypted Pool](#) (page 128) to import the encrypted volume. The [Wizard](#) (page 282) can be run afterward for post-configuration. It will then recognize that the volume has been imported and not prompt to reformat the disks.
- Upgrading the firmware of Broadcom SAS HBAs to the latest version is recommended.
- If upgrading from 9.3.x, please read the [FAQ: Updating from 9.3 to 9.10](https://forums.freenas.org/index.php?threads/faq-updating-from-9-3-to-9-10.54260/) (<https://forums.freenas.org/index.php?threads/faq-updating-from-9-3-to-9-10.54260/>) first.
- Upgrades from FreeNAS® 0.7x are not supported. The system has no way to import configuration settings from 0.7x versions of FreeNAS®. The configuration must be manually recreated. If supported, the FreeNAS® 0.7x volumes or disks must be manually imported.
- Upgrades on 32-bit hardware are not supported. However, if the system is currently running a 32-bit version of FreeNAS® and the hardware supports 64-bit, the system can be upgraded. Any archived reporting graphs will be lost during the upgrade.
- UFS is no longer supported. If your data currently resides on one UFS-formatted disk, create a ZFS volume using other disks after the upgrade, then use the instructions in [Import Disk](#) (page 127) to mount the UFS-formatted disk and copy the data to the ZFS volume. With only one disk, back up its data to another system or media before the upgrade, format the disk as ZFS after the upgrade, then restore the backup. If the data currently resides on a UFS RAID of disks, it is not possible to directly import that data to the ZFS volume. Instead, back up the data before the upgrade, create a ZFS volume after the upgrade, then restore the data from the backup.
- The VMware Tools VMXNET3 drivers are no longer supported. Configure and use the [vmx\(4\)](https://www.freebsd.org/cgi/man.cgi?query=vmx) (<https://www.freebsd.org/cgi/man.cgi?query=vmx>) driver instead.

2.5.2 Initial Preparation

Before upgrading the operating system, perform the following steps:

1. Back up the FreeNAS® configuration in `System → General → Save Config`.
2. If any volumes are encrypted, make sure that you have set the passphrase and have a copy of the encryption key and the latest recovery key. After the upgrade is complete, use the instructions in [Importing an Encrypted Pool](#) (page 128) to import the encrypted volume.
3. Warn users that the FreeNAS® shares will be unavailable during the upgrade; you should schedule the upgrade for a time that will least impact users.
4. Stop all services in `Services → Control Services`.

2.5.3 Upgrading Using the ISO

To perform an upgrade using this method, [download](http://download.freenas.org/latest/) (<http://download.freenas.org/latest/>) the `.iso` to the computer that will be used to prepare the installation media. Burn the downloaded `.iso` file to a CD or USB thumb drive using the instructions in [Preparing the Media](#) (page 11).

Insert the prepared media into the system and boot from it. After the media finishes booting into the installation menu, press `Enter` to select the default option of 1 Install/Upgrade. The installer presents a screen showing all available drives.

Warning: All drives are shown, including boot drives and storage drives. Only choose boot drives when upgrading. Choosing the wrong drives to upgrade or install will cause loss of data. If unsure about which drives contain the FreeNAS® operating system, reboot and remove the install media. In the FreeNAS® GUI, use `System → Boot` to identify the boot drives. More than one drive is shown when a mirror has been used.

Move to the drive where FreeNAS® is installed and press the `Spacebar` to mark it with a star. If a mirror has been used for the operating system, mark all of the drives where the FreeNAS® operating system is installed. Press `Enter` when done.

The installer recognizes earlier versions of FreeNAS® installed on the boot drive or drives and presents the message shown in [Figure 2.9](#).




Fig. 2.9: Upgrading a FreeNAS® Installation

Note: If you choose a Fresh Install, the backup of your configuration data must be restored using `System → General → Upload Config` after booting into the new operating system.

To perform an upgrade, press `Enter` to accept the default of Upgrade Install. Again, the installer will remind you that the operating system should be installed on a disk that is not used for storage.

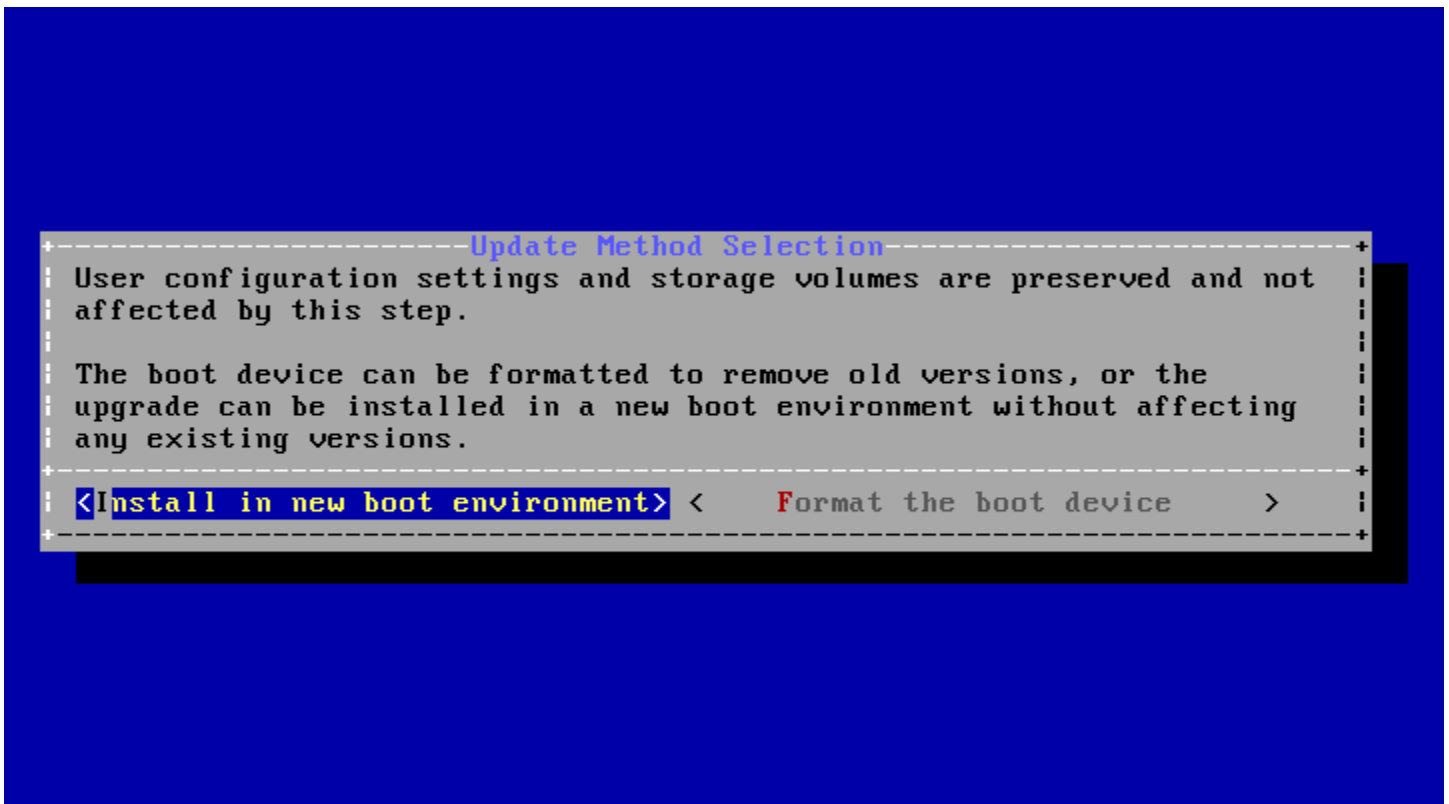


Fig. 2.10: Install in New Boot Environment or Format

The updated system can be installed in a new boot environment, or the entire boot device can be formatted to start fresh. Installing into a new boot environment preserves the old code, allowing a roll-back to previous versions if necessary. Formatting the boot device is usually not necessary but can reclaim space. User data and settings are preserved when installing to a new boot environment and also when formatting the boot device. Move the highlight to one of the options and press `Enter` to start the upgrade.

The installer unpacks the new image and displays the menu shown in [Figure 2.11](#). The database file that is preserved and migrated contains your FreeNAS® configuration settings.




Fig. 2.11: Preserve and Migrate Settings

Press `Enter`. FreeNAS® indicates that the upgrade is complete and a reboot is required. Press `OK`, highlight `3 Reboot System`, then press `Enter` to reboot the system. If the upgrade installer was booted from CD, remove the CD.

During the reboot there may be a conversion of the previous configuration database to the new version of the database. This happens during the “Applying database schema changes” line in the reboot cycle. This conversion can take a long time to finish, sometimes fifteen minutes or more, and might have to reboot the system again afterwards. Please be patient and the system will start normally. If database errors are shown but the graphical administrative interface is accessible, go to `Settings → General` and use the `Upload Config` button to upload the configuration that you saved before starting the upgrade.

2.5.4 Upgrading From the GUI

To perform an upgrade using this method, go to `System → Update`.

After the update is complete, you will temporarily lose your connection as the FreeNAS® system reboots into the new version of the operating system. The FreeNAS® system will normally receive the same IP address from the DHCP server. Refresh your browser after a moment to see if you can access the system.

2.5.5 If Something Goes Wrong

If an update fails, an alert is issued and the details are written to `/data/update.failed`.

To return to a previous version of the operating system, physical or IPMI access to the FreeNAS® console is needed. Reboot the system and watch for the boot menu. In the example shown in [Figure 2.12](#), the first boot menu entry, FreeNAS (default), refers to the initial installation, before the update was applied. The second boot entry, FreeNAS-1415259326, refers to the current version of the operating system, after the update was applied. This second entry is highlighted and begins with a star, indicating that this is the environment the system will boot unless another entry is manually selected. Both entries include a date and timestamp showing when that boot environment was created.

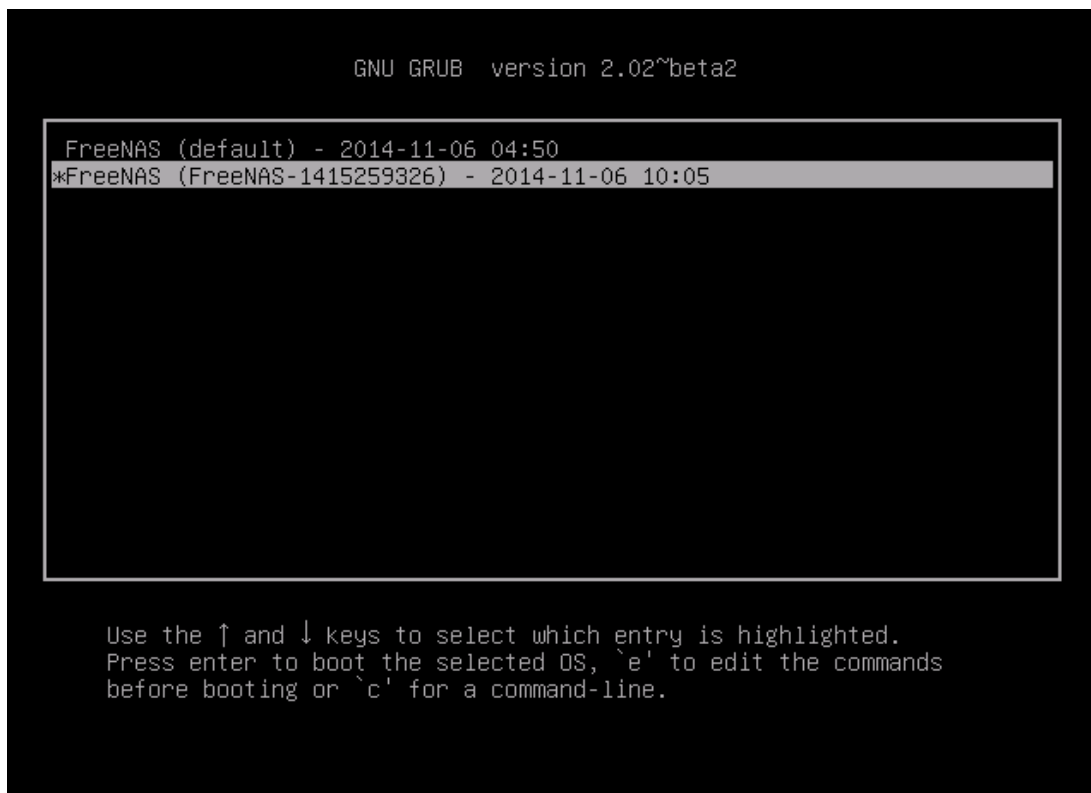


Fig. 2.12: Boot Menu

To boot into the previous version of the operating system, use the up or down arrow to select it and press `Enter`.

If a boot device fails and the system no longer boots, don't panic. The data is still on the disks and there is still a copy of the saved configuration. The system can be recovered with a few steps:

1. Perform a fresh installation on a new boot device.
2. Import the volumes in `Storage` → `Auto Import Volume`.
3. Restore the configuration in `System` → `General` → `Upload Config`.

Note: It is not possible to restore a saved configuration that is newer than the installed version. For example, if you reboot into an older version of the operating system, you cannot restore a configuration that was created in a later version.

2.5.6 Upgrading a ZFS Pool

In FreeNAS®, ZFS pools can be upgraded from the graphical administrative interface.

Before upgrading an existing ZFS pool, be aware of these caveats first:

- the pool upgrade is a one-way street, meaning that if you change your mind you cannot go back to an earlier ZFS version or downgrade to an earlier version of the software that does not support those feature flags.
- before performing any operation that may affect the data on a storage disk, always back up all data first and verify the integrity of the backup. While it is unlikely that the pool upgrade will affect the data, it is always better to be safe than sorry.
- upgrading a ZFS pool is optional. Do not upgrade the pool if the the possibility of reverting to an earlier version of FreeNAS® or repurposing the disks in another operating system that supports ZFS is desired. It is not necessary to

upgrade the pool unless newer ZFS feature flags are required. If a pool is upgraded to the latest feature flags, it will not be possible to import that pool into another operating system that does not yet support those feature flags.

To perform the ZFS pool upgrade, go to `Storage → Volumes → View Volumes` and highlight the volume (ZFS pool) to upgrade. Click the Upgrade button as shown in [Figure 2.13](#).

Note: If the Upgrade button does not appear, the pool is already at the latest feature flags and does not need to be upgraded.

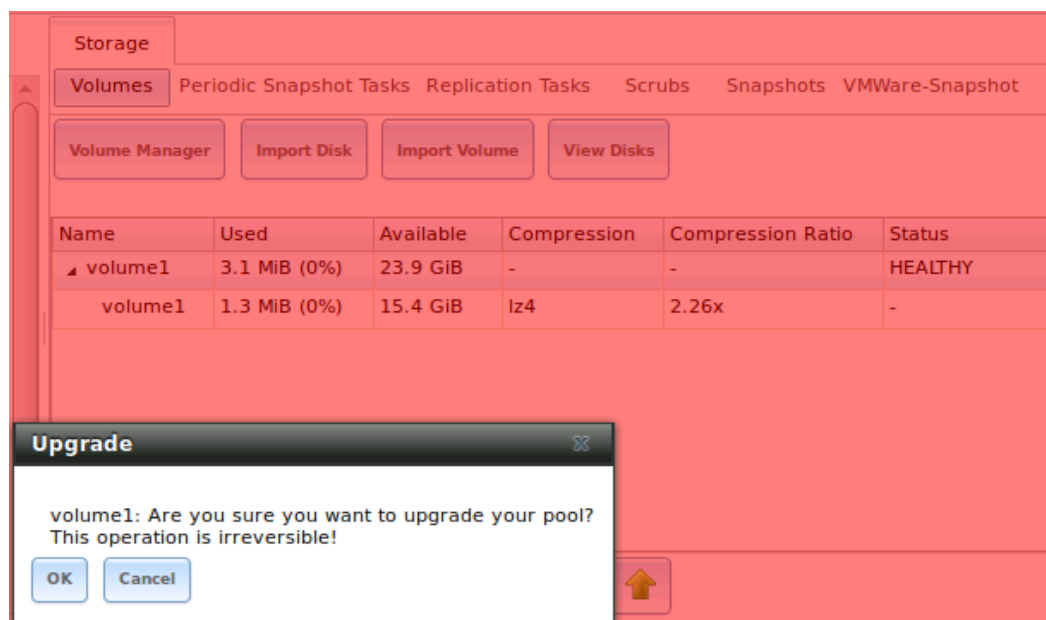


Fig. 2.13: Upgrading a ZFS Pool

The warning serves as a reminder that a pool upgrade is not reversible. Click OK to proceed with the upgrade.

The upgrade itself only takes a few seconds and is non-disruptive. It is not necessary to stop any sharing services to upgrade the pool. However, it is best to upgrade when the pool is not being heavily used. The upgrade process will suspend I/O for a short period, but is nearly instantaneous on a quiet pool.

2.6 Virtualization

FreeNAS® can be run inside a virtual environment for development, experimentation, and educational purposes. Please note that running FreeNAS® in production as a virtual machine is **not recommended** (<https://forums.freenas.org/index.php?threads/please-do-not-run-freenas-in-production-as-a-virtual-machine.12484/>). If you decide to use FreeNAS® within a virtual environment, **read this post first** (<https://forums.freenas.org/index.php?threads/absolutely-must-virtualize-freenas-a-guide-to-not-completely-losing-your-data.12714/>) as it contains useful guidelines for minimizing the risk of losing data.

To install or run FreeNAS® within a virtual environment, create a virtual machine that meets these minimum requirements:

- at least 8192 MB (8 GB) base memory size
- a virtual disk at least 8 GB in size to hold the operating system and boot environments
- at least one additional virtual disk at least 4 GB in size to be used as data storage
- a bridged network adapter

This section demonstrates how to create and access a virtual machine within VirtualBox and VMware ESXi environments.

2.6.1 VirtualBox

VirtualBox (<https://www.virtualbox.org/>) is an open source virtualization program originally created by Sun Microsystems. VirtualBox runs on Windows, BSD, Linux, Macintosh, and OpenSolaris. It can be configured to use a downloaded FreeNAS® .iso file, and makes a good testing environment for practicing configurations or learning how to use the features provided by FreeNAS®.

To create the virtual machine, start VirtualBox and click the New button, shown in [Figure 2.14](#), to start the new virtual machine wizard.

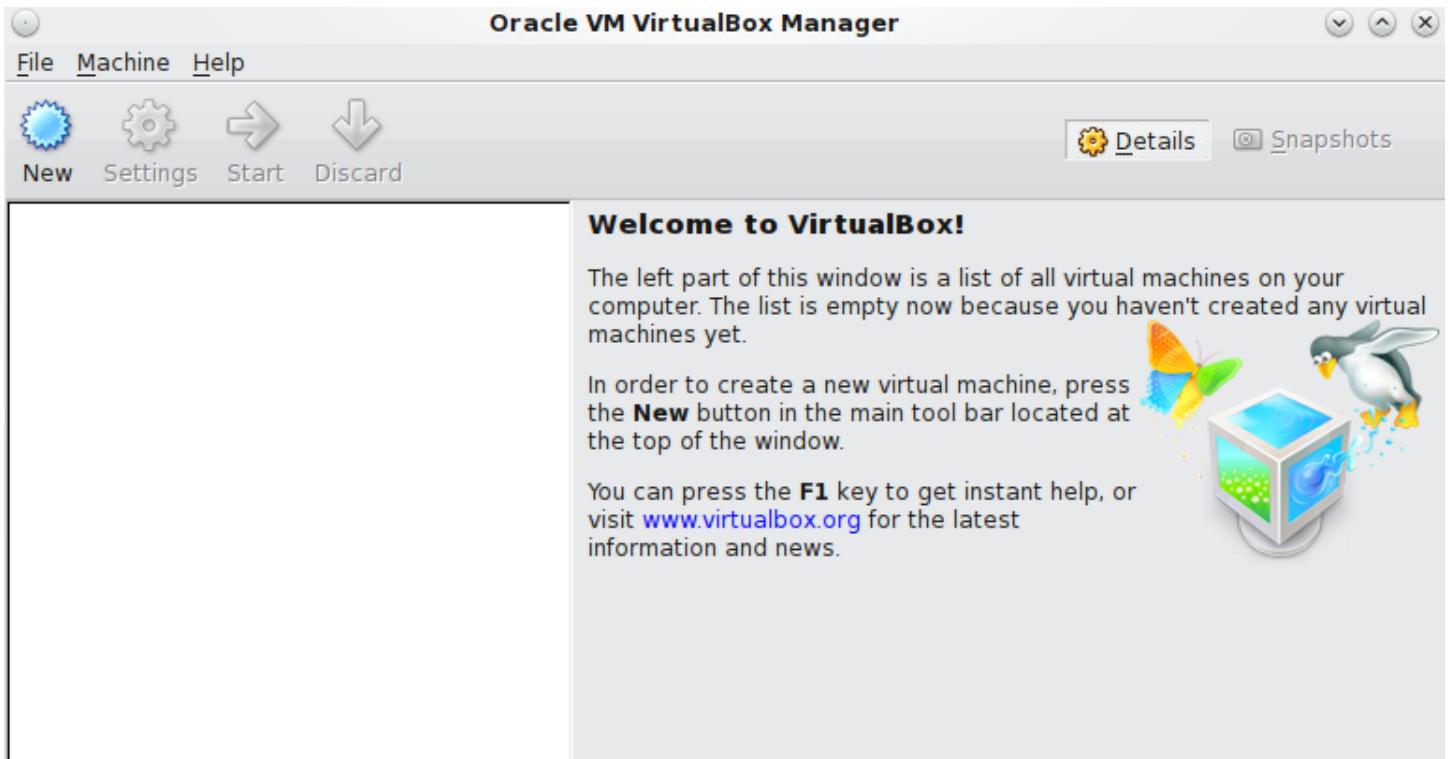


Fig. 2.14: Initial VirtualBox Screen

Click the Next button to see the screen in [Figure 2.15](#). Enter a name for the virtual machine, click the Operating System drop-down menu and select BSD, and select FreeBSD (64-bit) from the Version dropdown.

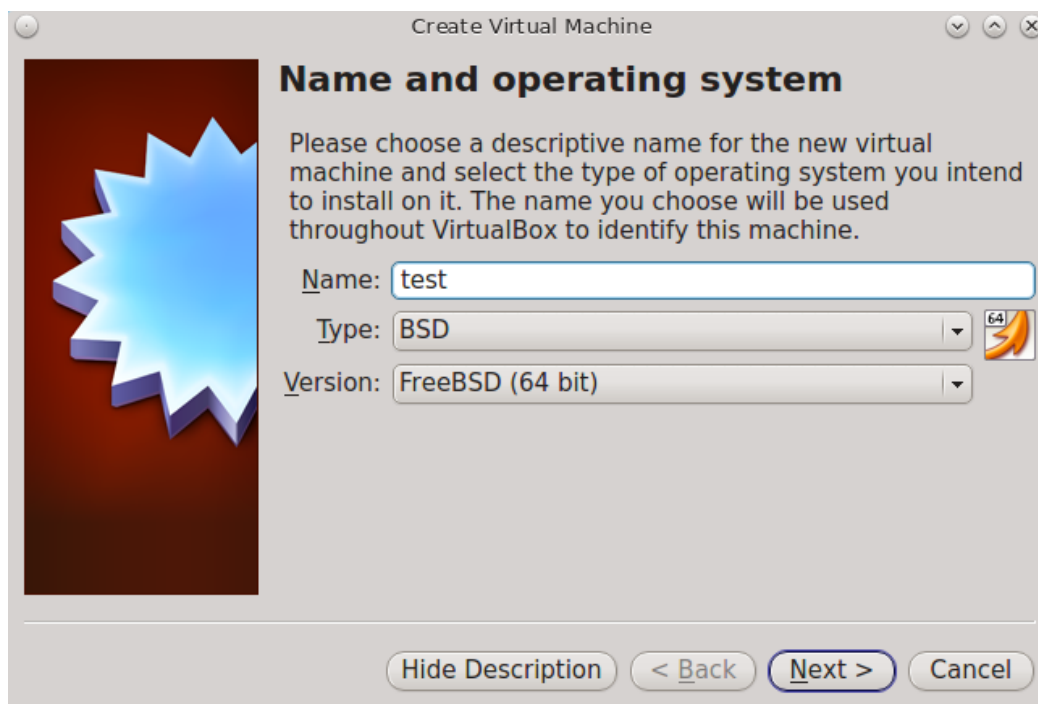


Fig. 2.15: Type in a Name and Select the Operating System for the New Virtual Machine

Click Next to see the screen in [Figure 2.16](#). The base memory size must be changed to at least 8192 MB. When finished, click Next to see the screen in [Figure 2.17](#).




Fig. 2.16: Select the Amount of Memory Reserved for the Virtual Machine

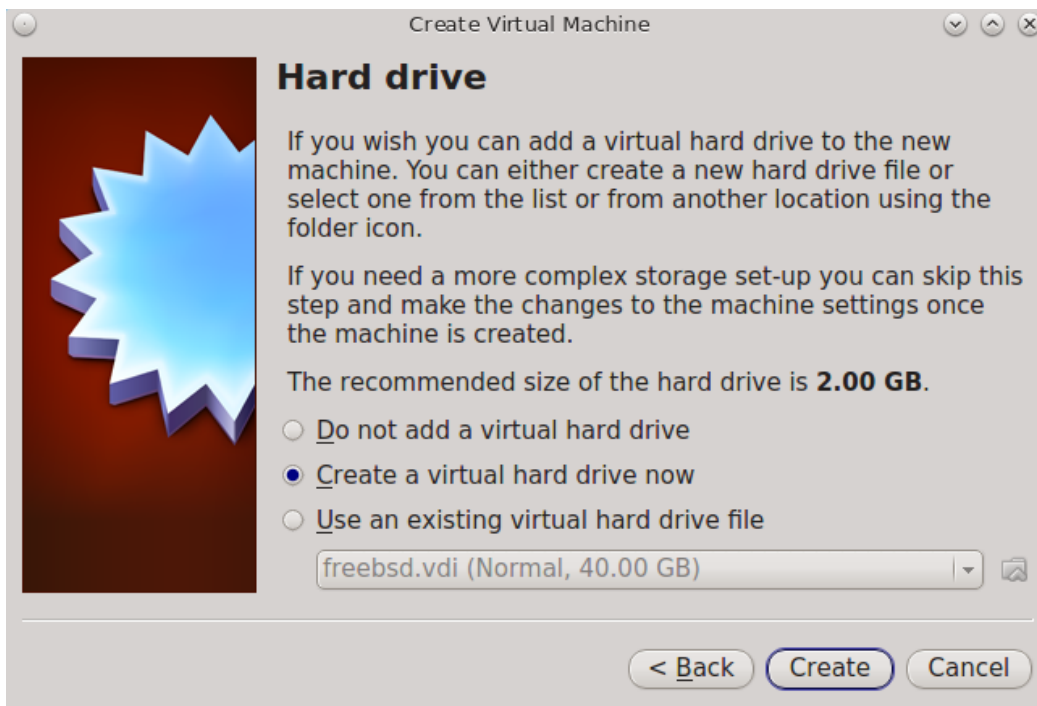


Fig. 2.17: Select Existing or Create a New Virtual Hard Drive

Click Create to launch the Create Virtual Hard Drive Wizard shown in [Figure 2.18](#).



Fig. 2.18: Create New Virtual Hard Drive Wizard

Select VDI and click the Next button to see the screen in [Figure 2.19](#).

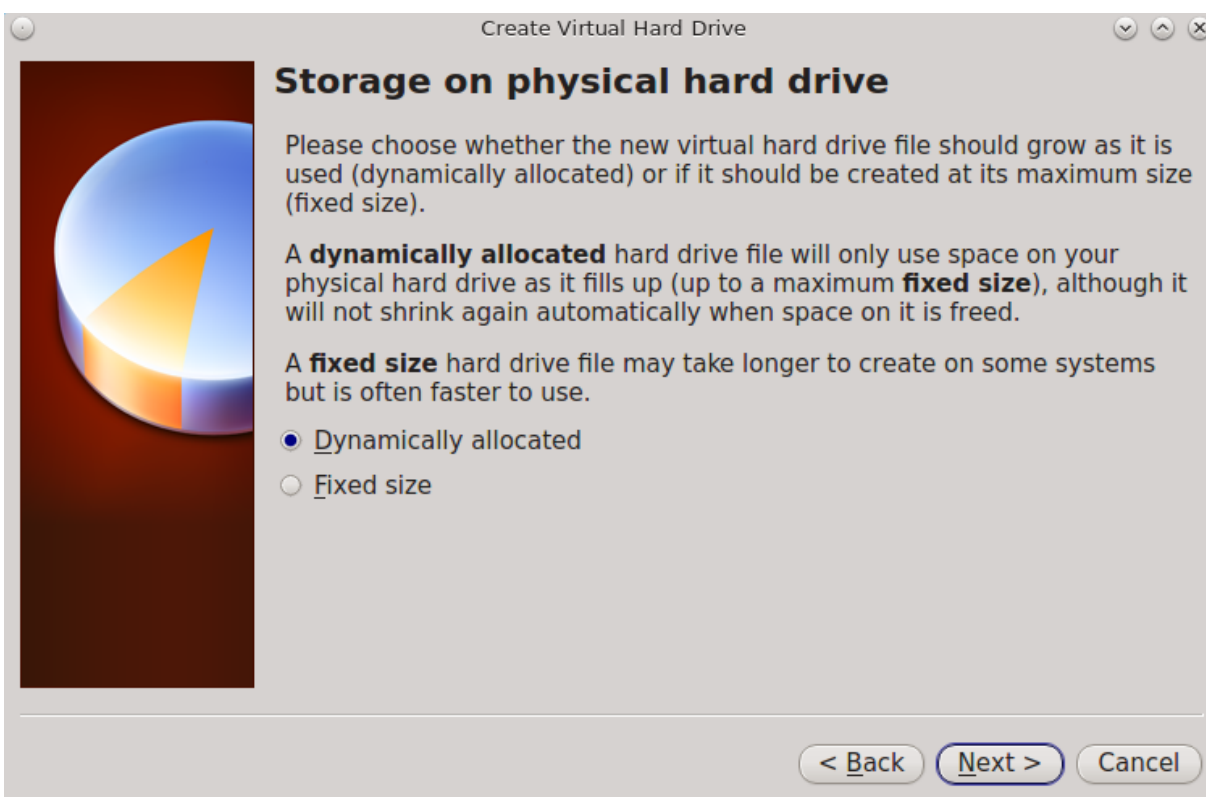


Fig. 2.19: Select Storage Type for Virtual Disk

Choose either Dynamically allocated or Fixed-size storage. The first option uses disk space as needed until it reaches the maximum size that is set in the next screen. The second option creates a disk the full amount of disk space, whether it is used or not. Choose the first option to conserve disk space; otherwise, choose the second option, as it allows VirtualBox to run slightly faster. After selecting Next, the screen in [Figure 2.20](#) is shown.

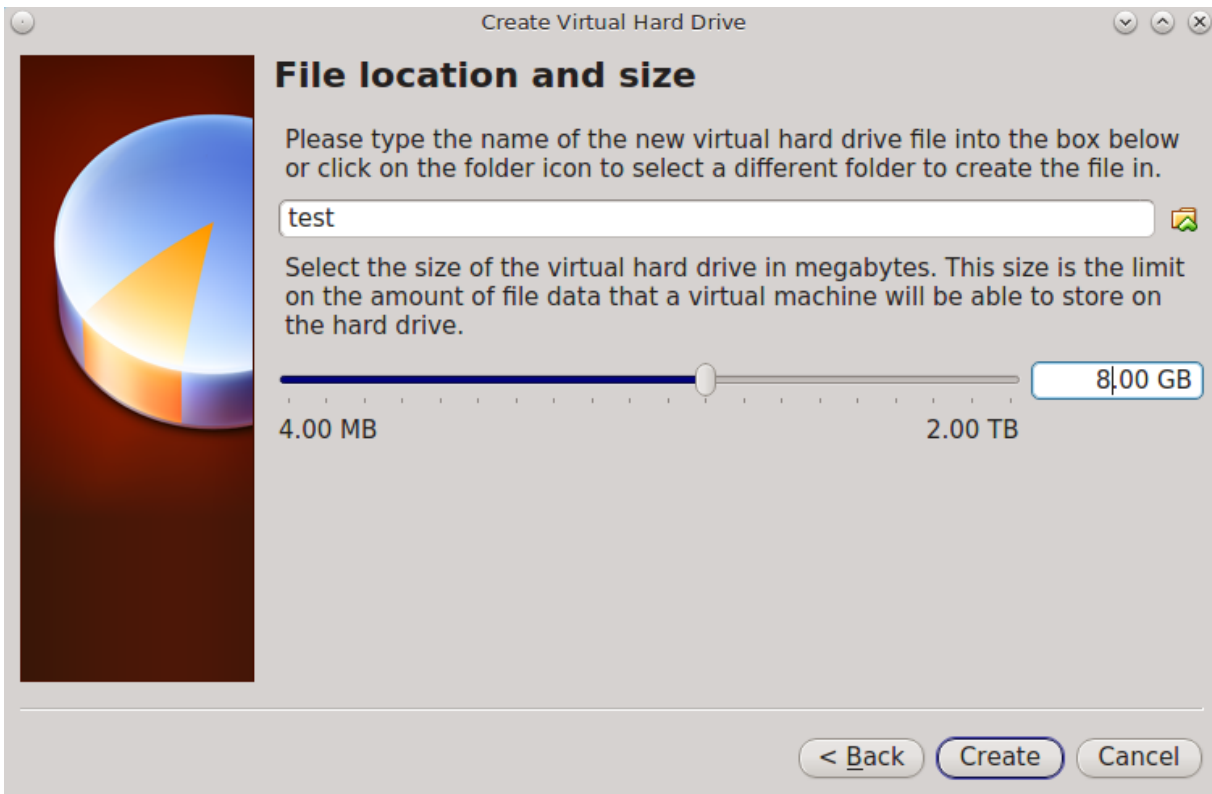


Fig. 2.20: Select File Name and Size of Virtual Disk

This screen is used to set the size (or upper limit) of the virtual disk. Increase the default size to 8 GB. Use the folder icon to browse to a directory on disk with sufficient space to hold the virtual disk files. Remember that there will be a system disk of at least 8 GB and at least one data storage disk of at least 4 GB.

After making a selection and pressing Next, a summary of the configuration options chosen is shown. Use the Back button to return to a previous screen if any values need to be modified. Otherwise, click Finish to complete the wizard. The new virtual machine is listed in the left frame, as shown in [Figure 2.21](#).

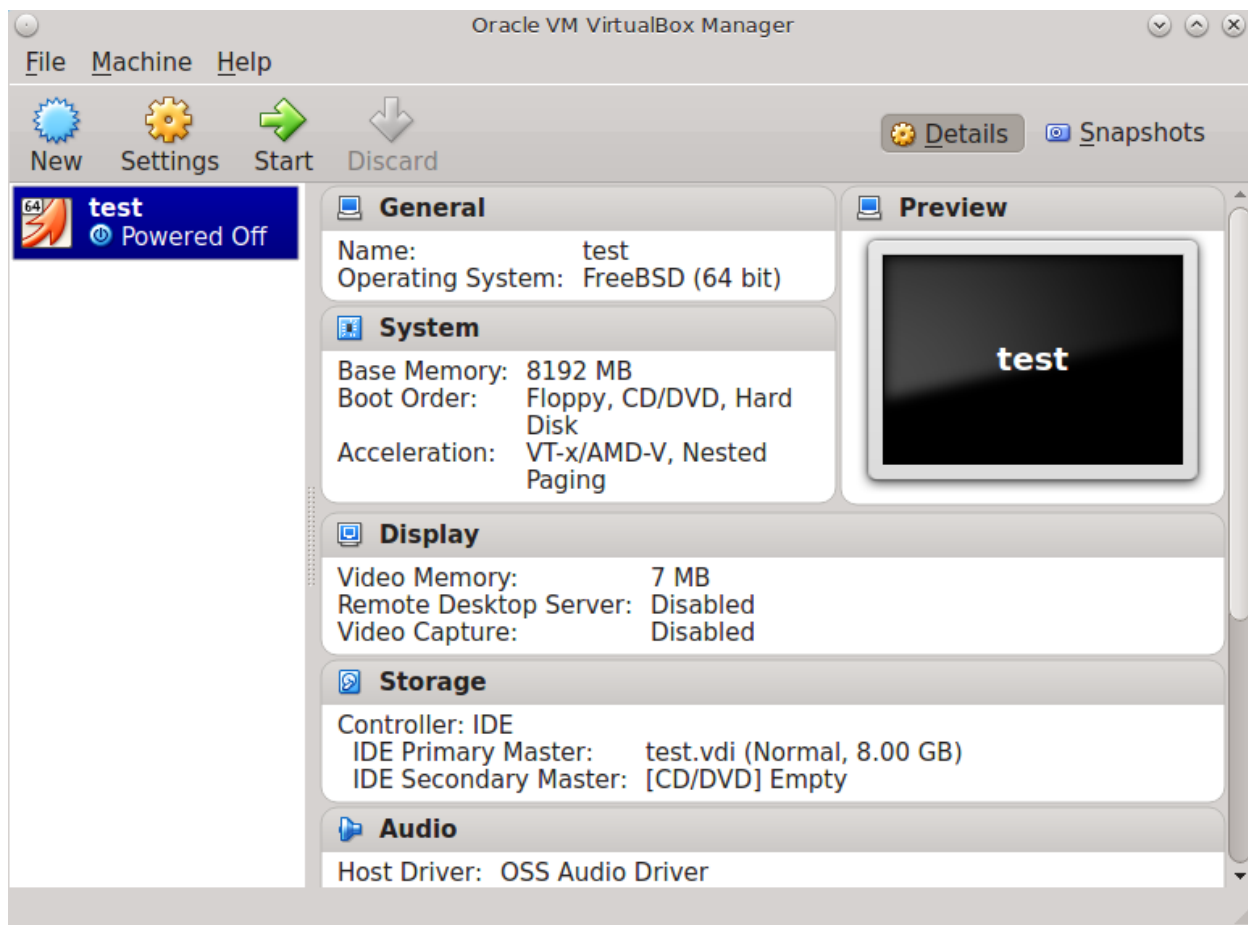


Fig. 2.21: The New Virtual Machine

Create the virtual disks to be used for storage. Click the Storage hyperlink in the right frame to access the storage screen seen in Figure 2.22.

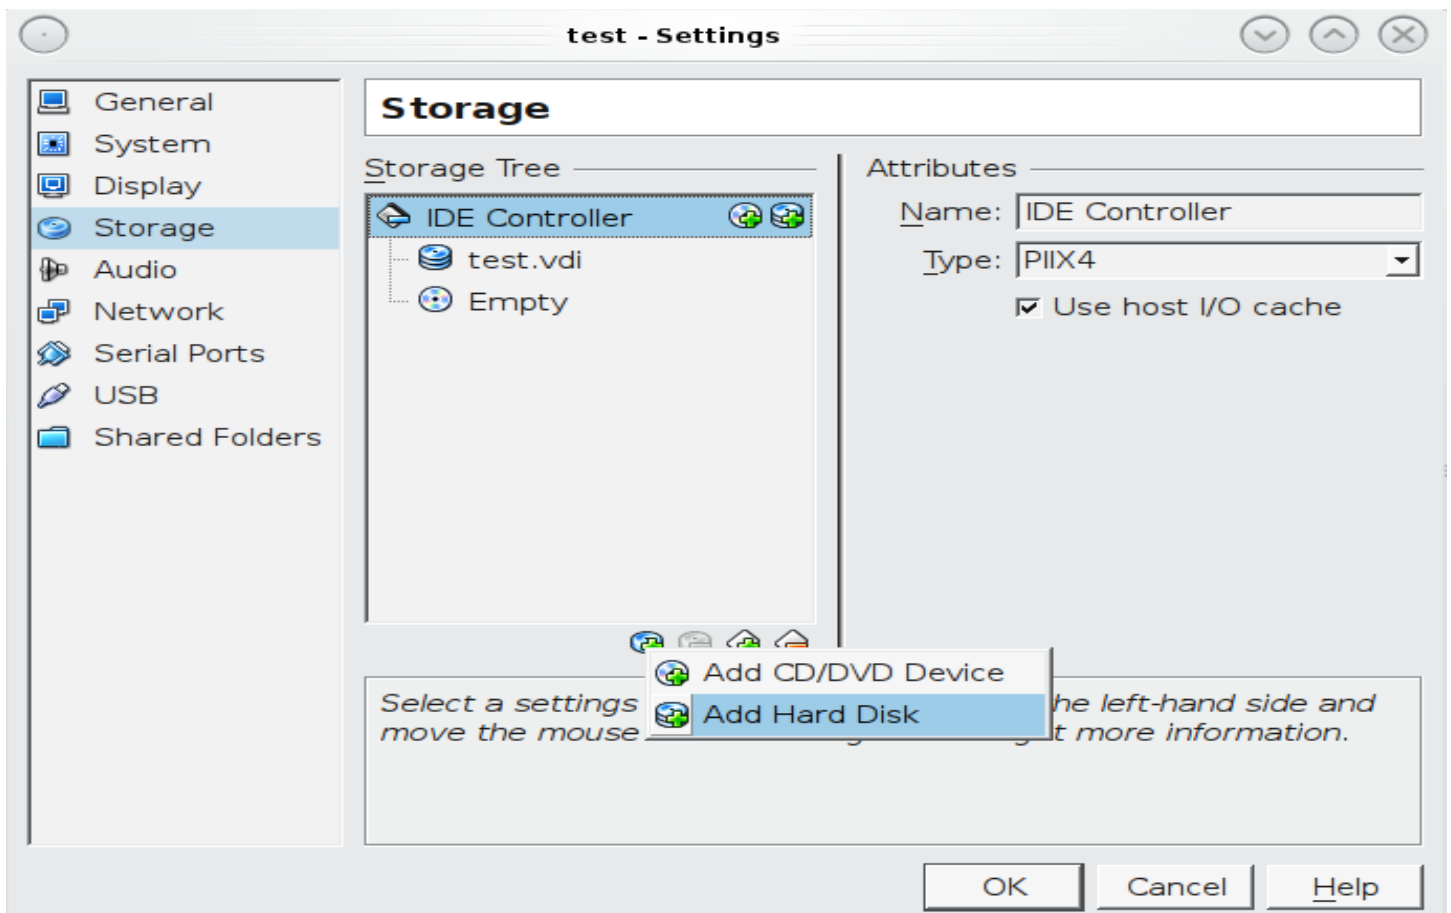


Fig. 2.22: Storage Settings of the Virtual Machine

Click the Add Attachment button, select Add Hard Disk from the pop-up menu, then click the Create New Disk button. This launches the Create New Virtual Hard Drive Wizard (seen in [Figure 2.18](#) and [2.19](#)). This disk will be used for storage, so create a size appropriate to your needs, making sure that it is at least 4 GB. To practice with RAID configurations, create as many virtual disks as needed. Two disks can be created on each IDE controller. For additional disks, click the Add Controller button to create another controller for attaching additional disks.

Create a device for the installation media. Highlight the word “Empty”, then click the CD icon as shown in [Figure 2.23](#).

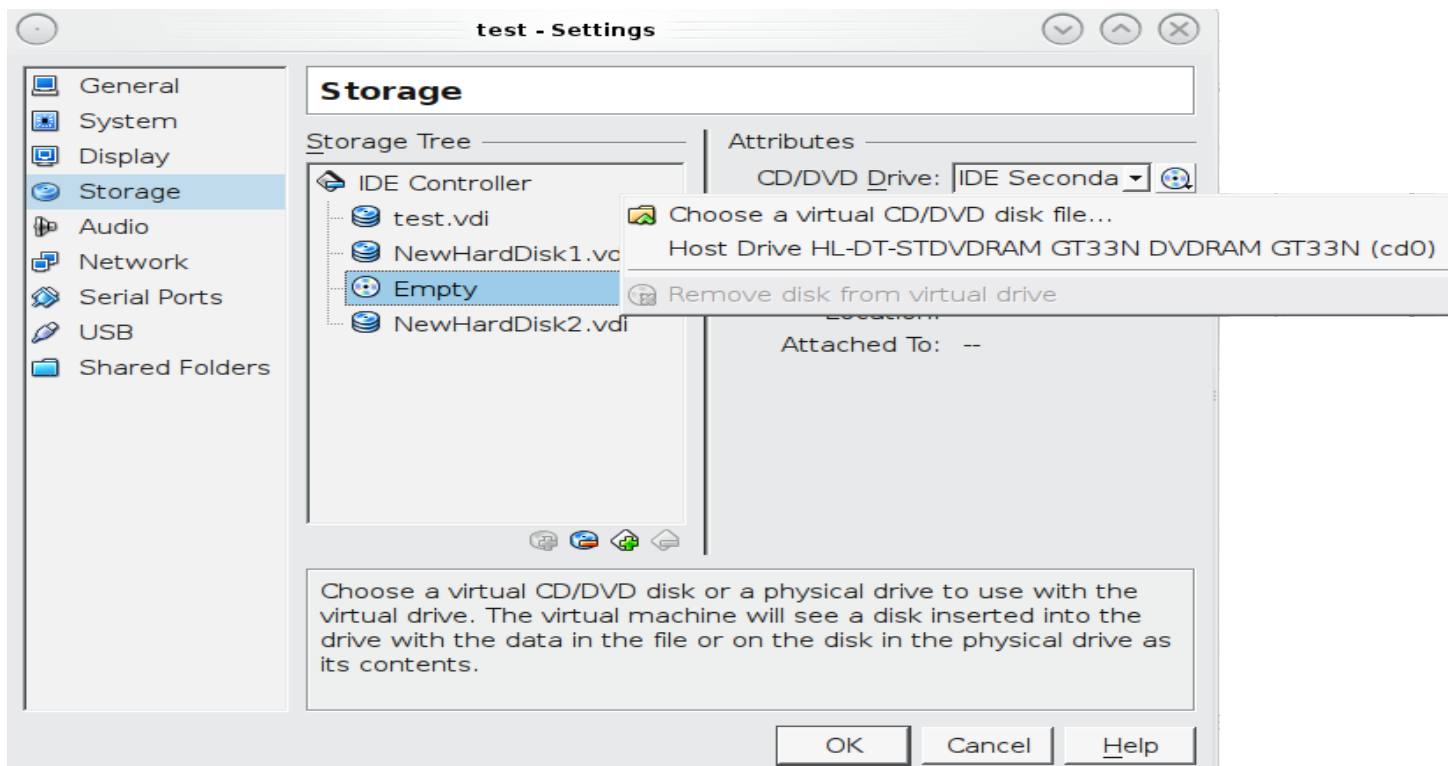


Fig. 2.23: Configuring ISO Installation Media

Click **Choose a virtual CD/DVD disk file...** to browse to the location of the `.iso` file. If the `.iso` was burned to CD, select the detected Host Drive.

Depending on the extensions available in the host CPU, it might not be possible to boot the VM from `.iso`. If “your CPU does not support long mode” is shown when trying to boot the `.iso`, the host CPU either does not have the required extension or AMD-V/VT-x is disabled in the system BIOS.

Note: If you receive a kernel panic when booting into the ISO, stop the virtual machine. Then, go to **System** and check the box **Enable IO APIC**.

To configure the network adapter, go to **Settings** → **Network**. In the **Attached to** drop-down menu select **Bridged Adapter**, then choose the name of the physical interface from the **Name** drop-down menu. In the example shown in [Figure 2.24](#), the Intel Pro/1000 Ethernet card is attached to the network and has a device name of `em0`.

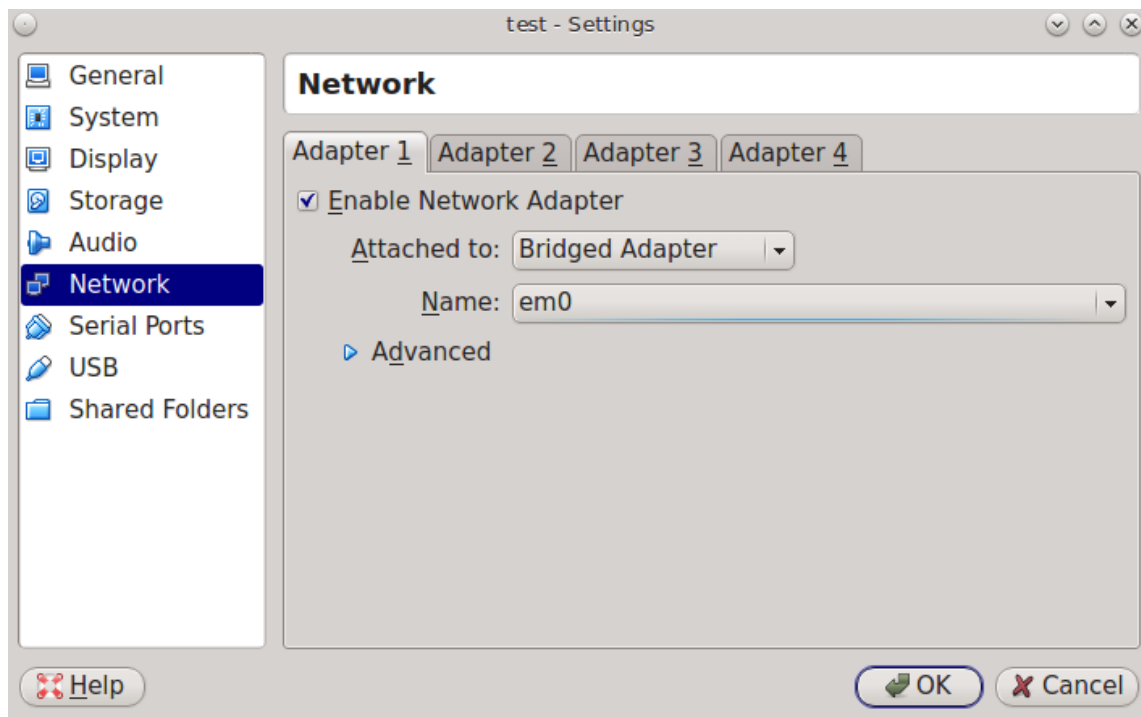


Fig. 2.24: Configuring a Bridged Adapter in VirtualBox

After configuration is complete, click the Start arrow and install FreeNAS® as described in [Performing the Installation](#) (page 13). Once FreeNAS® is installed, press **F12** when the VM starts to boot to access the boot menu. Select the primary hard disk as the boot option. You can permanently boot from disk by removing the CD/DVD device in Storage or by unchecking CD/DVD-ROM in the Boot Order section of System.

2.6.2 VMware ESXi

Before using ESXi, read [this post](https://forums.freenas.org/index.php?threads/sync-writes-or-why-is-my-esxi-nfs-so-slow-and-why-is-iscsi-faster.12506/) (<https://forums.freenas.org/index.php?threads/sync-writes-or-why-is-my-esxi-nfs-so-slow-and-why-is-iscsi-faster.12506/>) for an explanation of why iSCSI will be faster than NFS.

ESXi is a bare-metal hypervisor architecture created by VMware Inc. Commercial and free versions of the VMware vSphere Hypervisor operating system (ESXi) are available from the [VMware website](http://www.vmware.com/products/esxi-and-esx/overview) (<http://www.vmware.com/products/esxi-and-esx/overview>). After the operating system is installed on supported hardware, use a web browser to connect to its IP address. The welcome screen provides a link to download the VMware vSphere client which is used to create and manage virtual machines.

Once the VMware vSphere client is installed, use it to connect to the ESXi server. To create a new virtual machine, click **File** → **New** → **Virtual Machine**. The New Virtual Machine Wizard will launch as shown in [Figure 2.25](#).

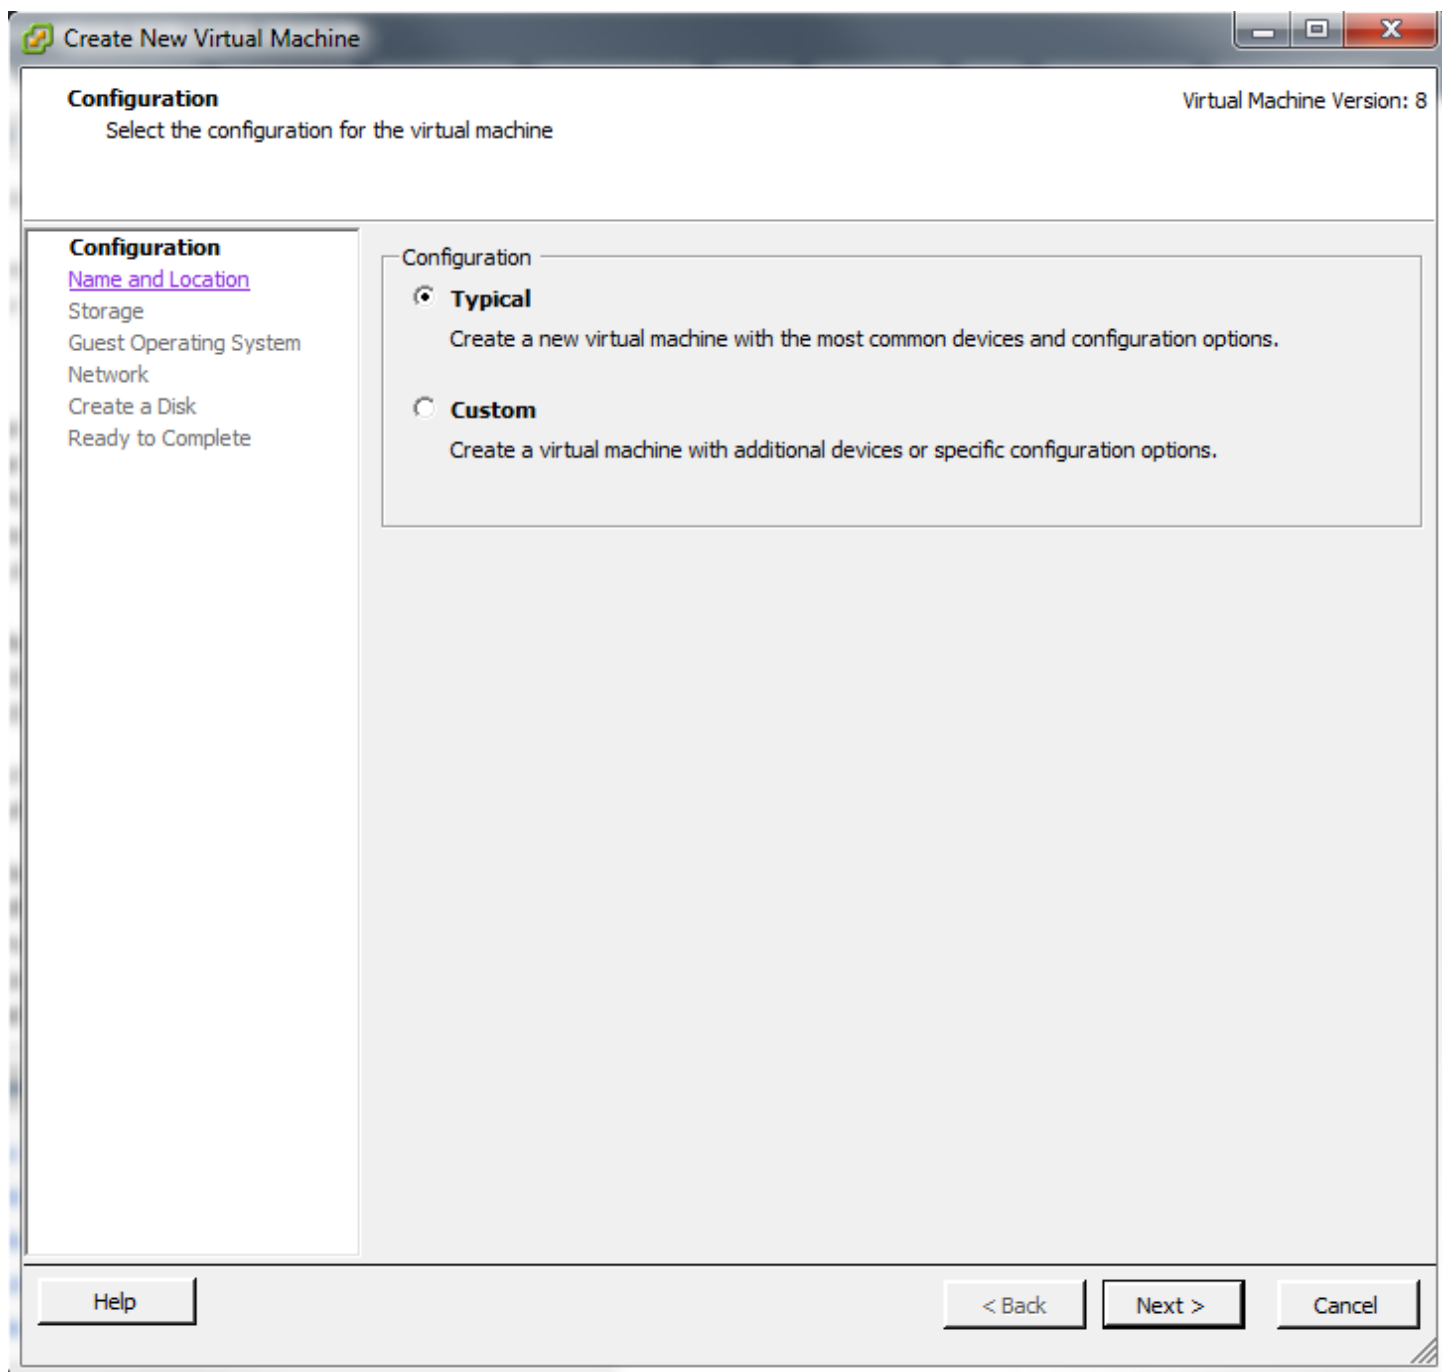


Fig. 2.25: New Virtual Machine Wizard

Click Next and enter a name for the virtual machine. Click Next and highlight a datastore. An example is shown in [Figure 2.26](#). Click Next. In the screen shown in [Figure 2.27](#), click Other, then select a FreeBSD 64-bit architecture.

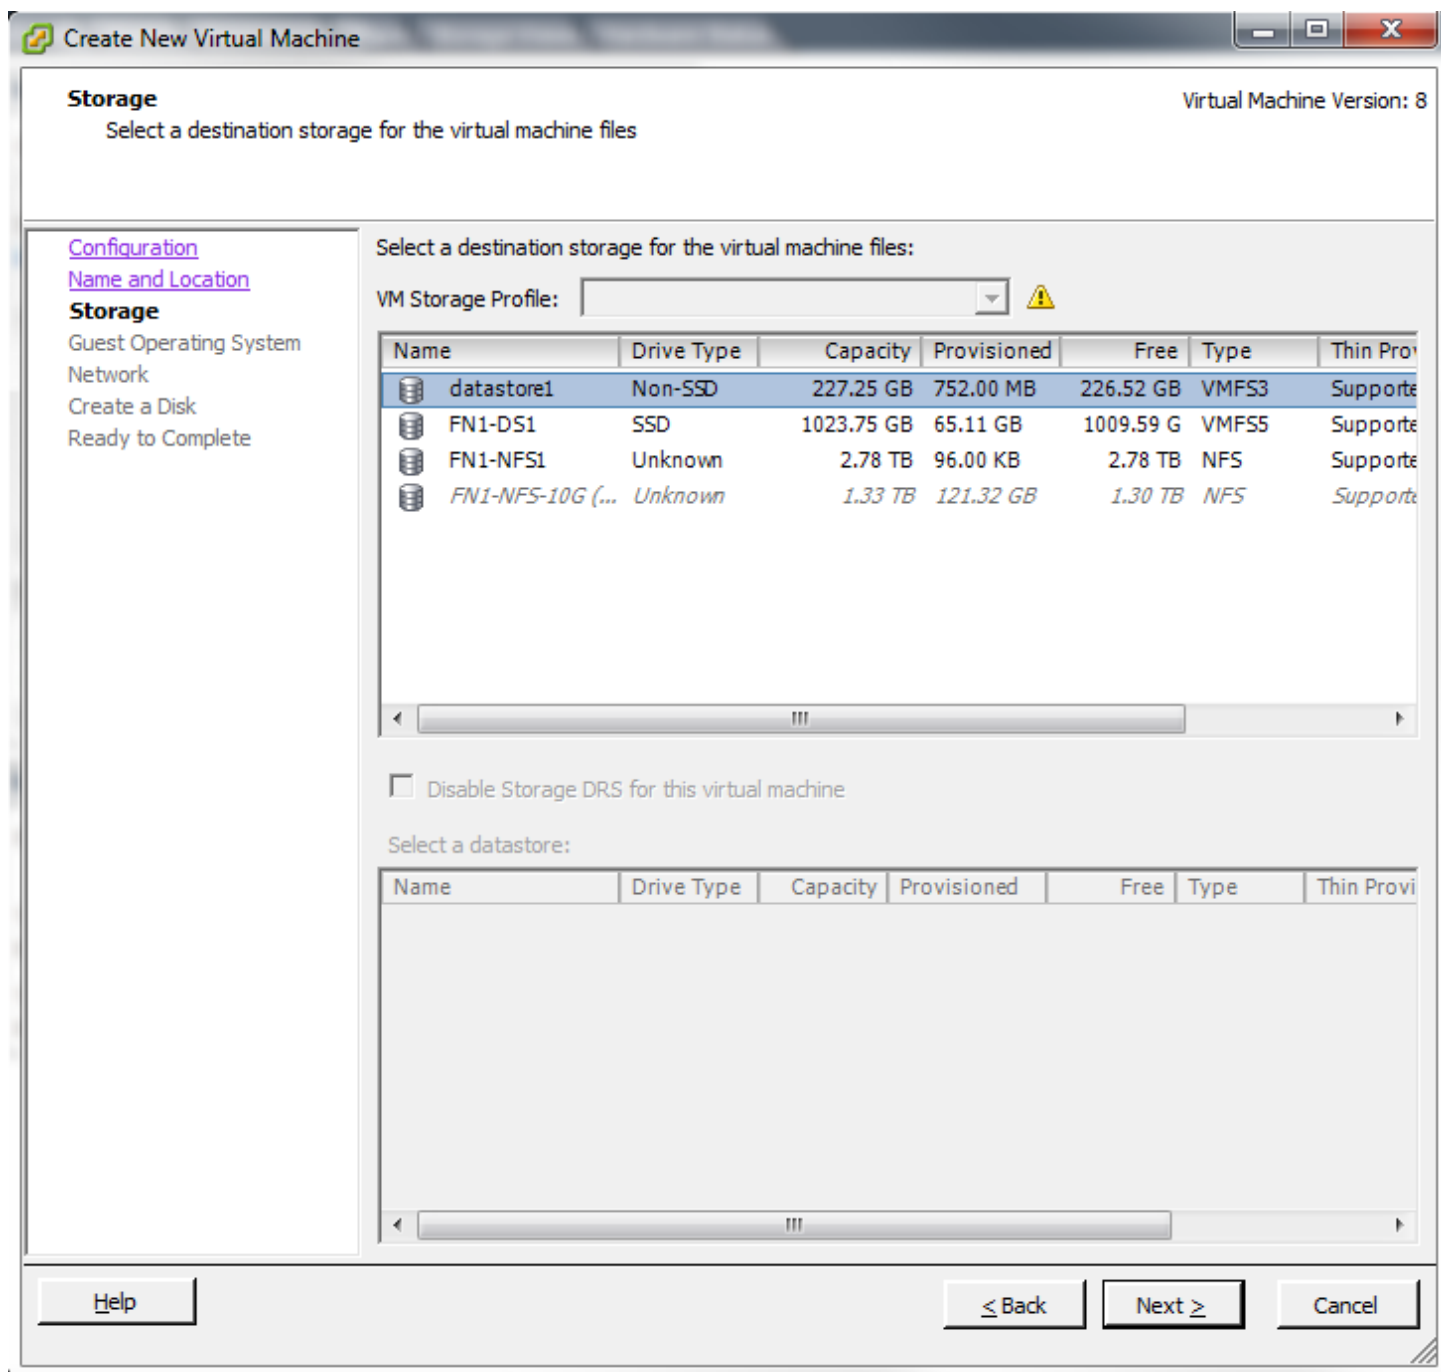


Fig. 2.26: Select Datastore

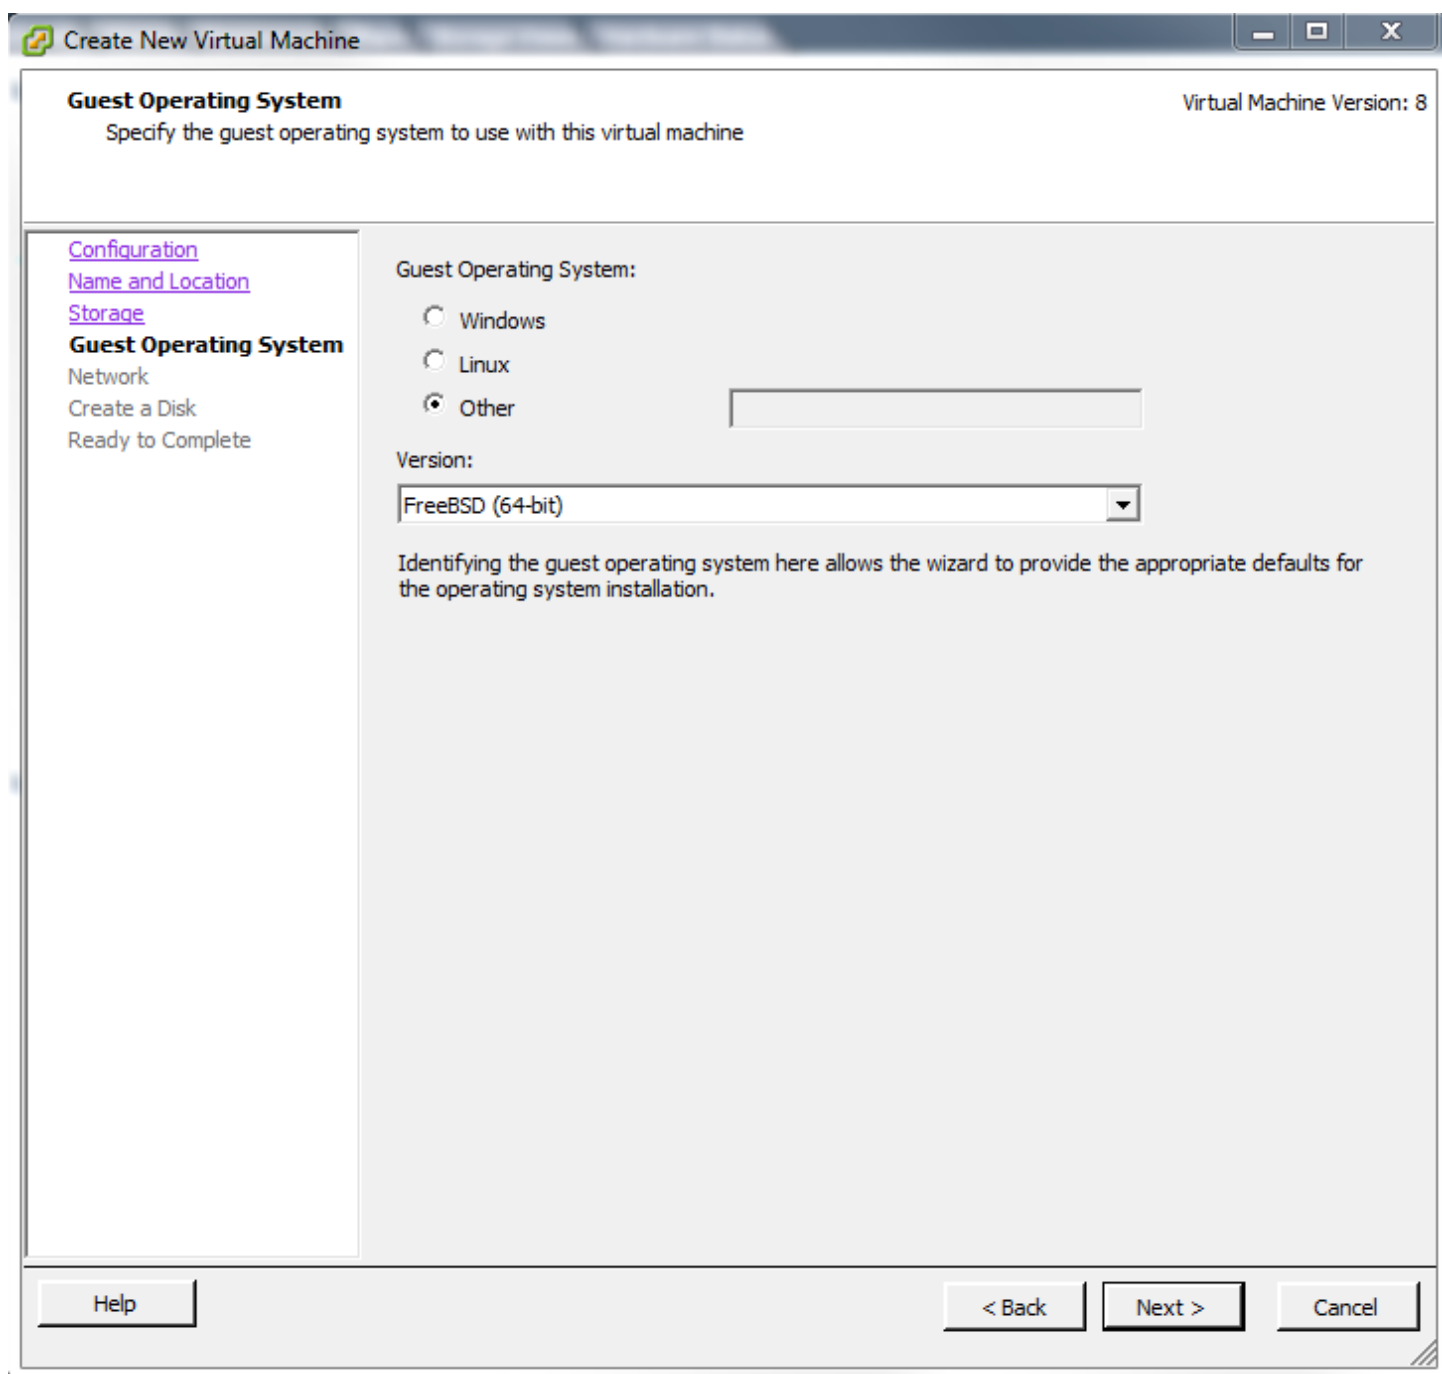


Fig. 2.27: Select Operating System

Click Next and create a virtual disk file of 8 GB to hold the FreeNAS® operating system, as shown in [Figure 2.28](#).

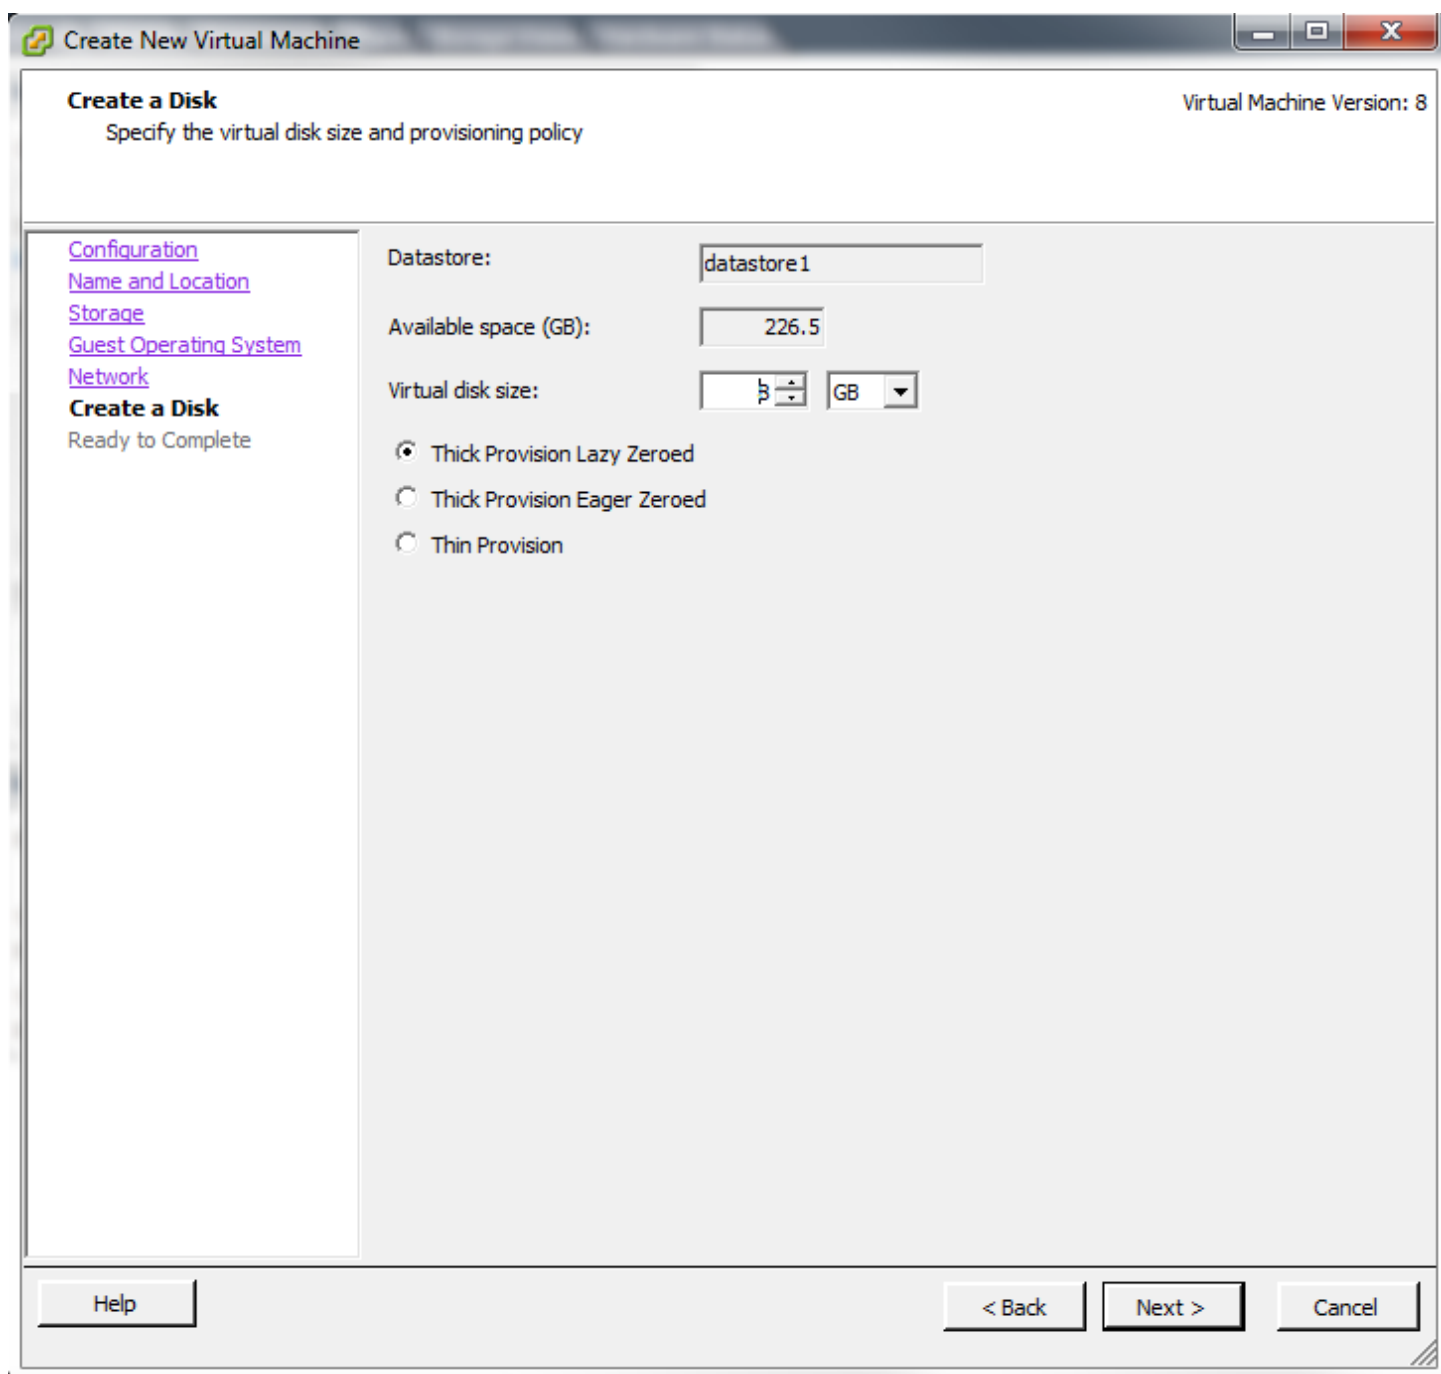


Fig. 2.28: Create Disk for the Operating System

Click Next and Finish. The new virtual machine is listed in the left frame. Right-click the virtual machine and select Edit Settings to access the screen shown in [Figure 2.29](#).

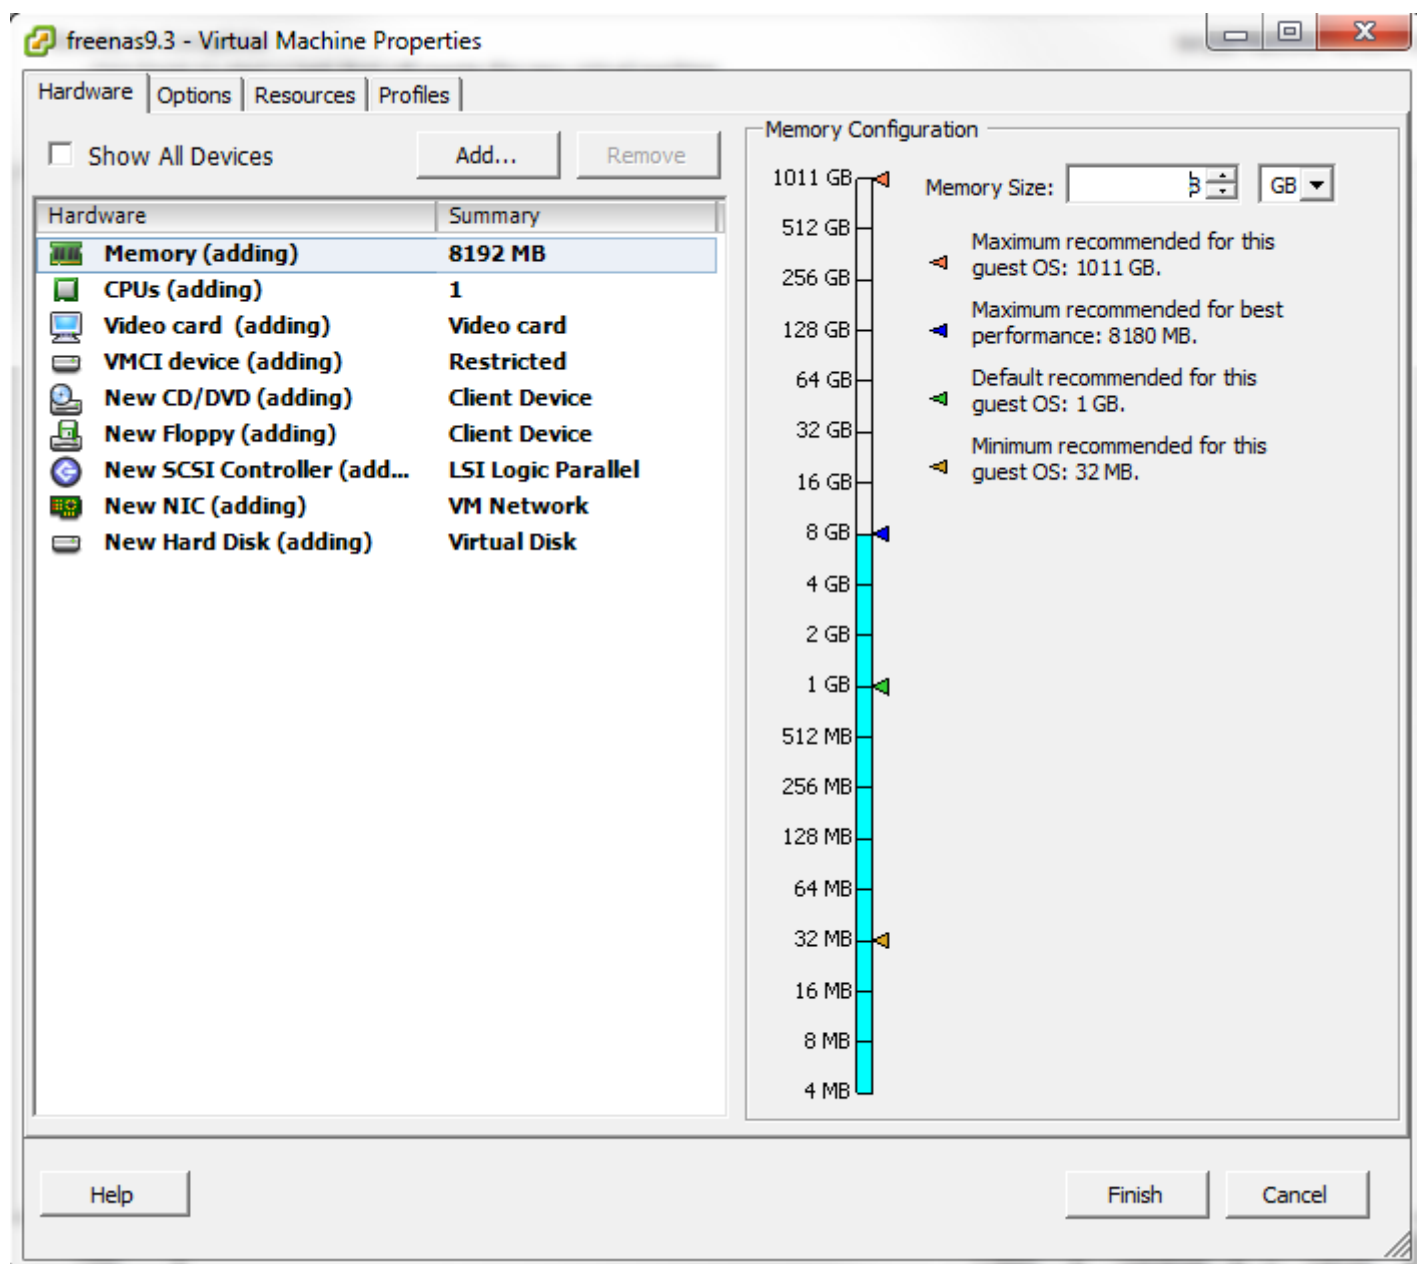


Fig. 2.29: Virtual Machine Settings

Increase the Memory Configuration to at least 8192 MB.

To create a storage disk, click **Hard disk 1** → **Add**. In the **Device Type** menu, highlight **Hard Disk** and click **Next**. Select **Create a new virtual disk** and click **Next**. In the screen shown in Figure 2.30, select the size of the disk. To dynamically allocate space as needed, check the box **Allocate and commit space on demand (Thin Provisioning)**. Click **Next**, then **Next**, then **Finish** to create the disk. Repeat to create the amount of storage disks needed to meet your requirements.

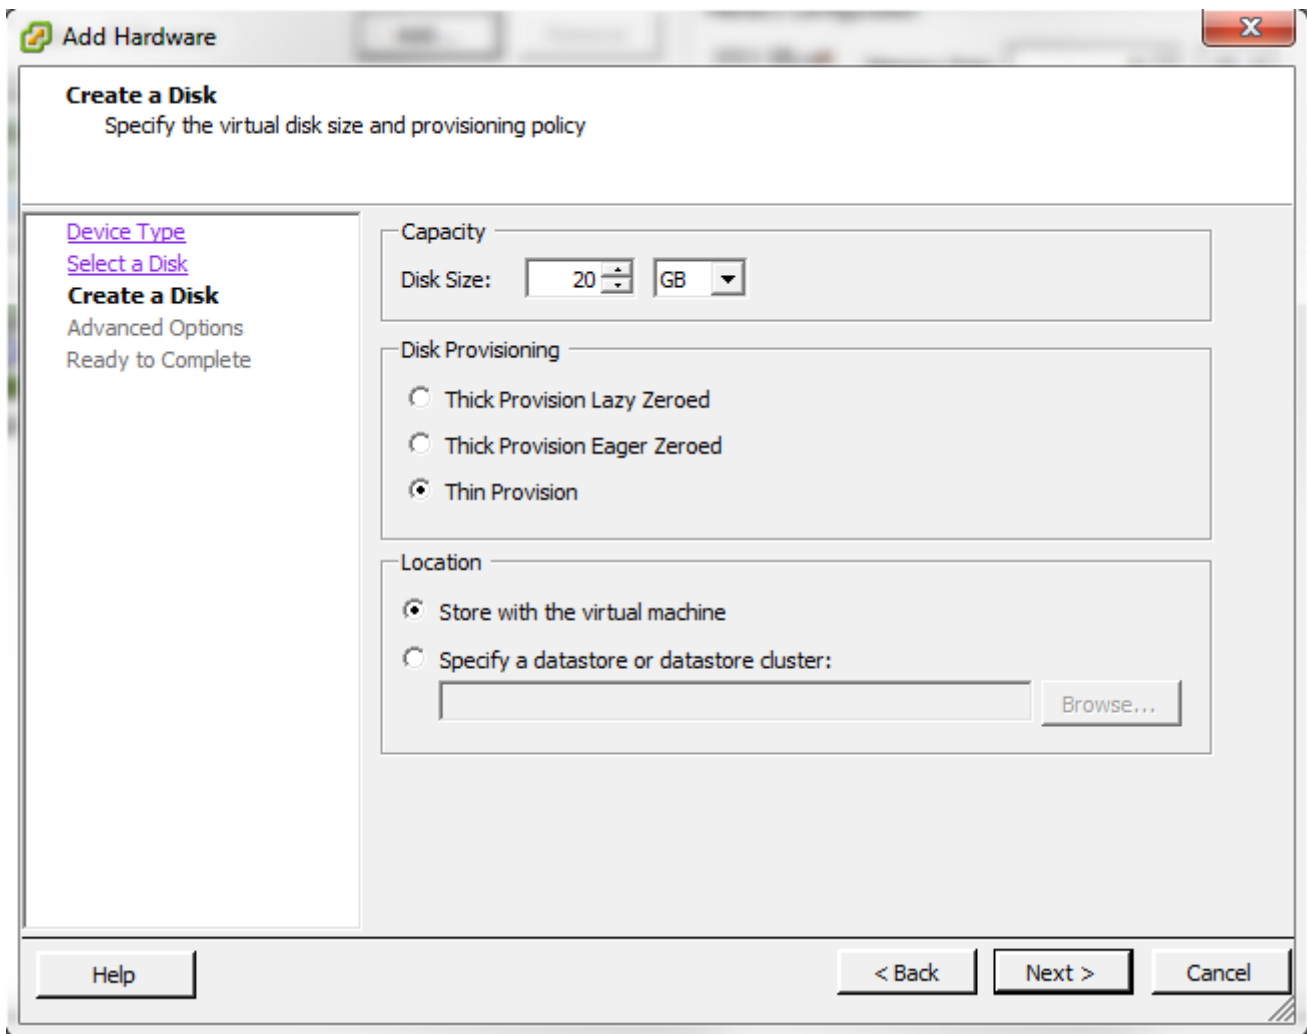


Fig. 2.30: Creating a Storage Disk

For ESX 5.0, Workstation 8.0, or Fusion 4.0 or higher, additional configuration is needed so that the virtual HPET setting does not prevent the virtual machine from booting.

If you are running ESX, while in Edit Settings, click `Options` → `Advanced` → `General` → `Configuration Parameters`. Change `hpet0.present` from `true` to `false`, then click OK twice to save the setting.


For Workstation or Player, while in Edit Settings, click `Options` → `Advanced` → `File Locations`. Locate the path for the Configuration file named `filename.vmx`. Open that file in a text editor, change `hpet0.present` from `true` to `false`, and save the change.

BOOTING

The Console Setup menu, shown in [Figure 3.1](#), appears at the end of the boot process. If the FreeNAS® system has a keyboard and monitor, this Console Setup menu can be used to administer the system.

Note: When connecting to the FreeNAS® system with SSH or the web [Shell](#) (page 290), the Console Setup menu is not shown by default. It can be started by the root user or another user with root permissions by typing `/etc/netcli`.

The Console Setup menu can be disabled by unchecking Enable Console Menu in `System → Settings → Advanced`.



```
Console setup
-----
1) Configure Network Interfaces
2) Configure Link Aggregation
3) Configure VLAN Interface
4) Configure Default Route
5) Configure Static Routes
6) Configure DNS
7) Reset Root Password
8) Reset Configuration to Defaults
9) Shell
10) Reboot
11) Shut Down

The web user interface is at:

http://10.0.0.102

Enter an option from 1-11: █
```

Fig. 3.1: Console Setup Menu

The menu provides these options:

- 1) Configure Network Interfaces provides a configuration wizard to set up the system's network interfaces.
- 2) Configure Link Aggregation is for creating or deleting link aggregations.
- 3) Configure VLAN Interface is used to create or delete VLAN interfaces.
- 4) Configure Default Route is used to set the IPv4 or IPv6 default gateway. When prompted, enter the IP address of the default gateway.
- 5) Configure Static Routes prompts for the destination network and gateway IP address. Re-enter this option for each static route needed.

- 6) Configure DNS prompts for the name of the DNS domain and the IP address of the first DNS server. When adding multiple DNS servers, press `Enter` to enter the next one. Press `Enter` twice to leave this option.
- 7) Reset Root Password is used to reset a lost or forgotten root password. Select this option and follow the prompts to set the password.
- 8) Reset Configuration to Defaults Caution! This option deletes all of the configuration settings made in the administrative GUI and is used to reset a FreeNAS® system back to defaults. Before selecting this option, make a full backup of all data and make sure all encryption keys and passphrases are known! After this option is selected, the configuration is reset to defaults and the system reboots. `Storage → Volumes → Import Volume` can be used to re-import volumes.
- 9) Shell starts a shell for running FreeBSD commands. To leave the shell, type `exit`.
- 10) Reboot reboots the system.
- 11) Shut Down shuts down the system.

Note: The numbering and quantity of options on this menu can change due to software updates, service agreements, or other factors. Please carefully check the menu before selecting an option, and keep this in mind when writing local procedures.

3.1 Obtaining an IP Address

During boot, FreeNAS® automatically attempts to connect to a DHCP server from all live network interfaces. If it successfully receives an IP address, the address is displayed so it can be used to access the graphical user interface. The example in [Figure 3.1](#) shows a FreeNAS® system that is accessible at `http://192.168.1.119`.

Some FreeNAS® systems are set up without a monitor, making it challenging to determine which IP address has been assigned. On networks that support Multicast DNS (mDNS), the hostname and domain can be entered into the address bar of a browser. By default, this value is `freenas.local`.

If the FreeNAS® server is not connected to a network with a DHCP server, use the console network configuration menu to manually configure the interface as shown here. In this example, the FreeNAS® system has one network interface, `em0`.

```
Enter an option from 1-12: 1
1) em0
Select an interface (q to quit): 1
Reset network configuration (y/n) n
Configure interface for DHCP? (y/n) n
Configure IPv4? (y/n) y
Interface name:      (press enter, the name can be blank)
Several input formats are supported
Example 1 CIDR Notation:
    192.168.1.1/24
Example 2 IP and Netmask separate:
    IP: 192.168.1.1
    Netmask: 255.255.255.0, or /24 or 24
IPv4 Address: 192.168.1.108/24
Saving interface configuration: Ok
Configure IPv6? (y/n) n
Restarting network: ok


...

The web user interface is at
http://192.168.1.108
```

After the system has an IP address, enter that address into a graphical web browser from a computer connected to the same network as the FreeNAS® system.

3.2 Logging In

The password for the root user is requested as shown in Figure 3.2.



A login dialog box titled "Welcome to FreeNAS® 11". It contains a "Username:" label with a text input field containing "root", a "Password:" label with an empty text input field, a "Log In" button, and a link that says "Try the Beta UI!". The iXsystems logo is in the bottom right corner.

Fig. 3.2: Enter the Root Password

Note: The FreeNAS® UI is in the process of being rewritten in Angular, with a new, asynchronous middleware. To see a preview of the new UI, click the Try the Beta UI! link in the login box. Note that the new UI is not expected to be feature complete until version 11.2. Until then, this Guide will only demonstrate the classic UI.

Enter the password chosen during the installation. The administrative GUI is displayed as shown in Figure 3.3.

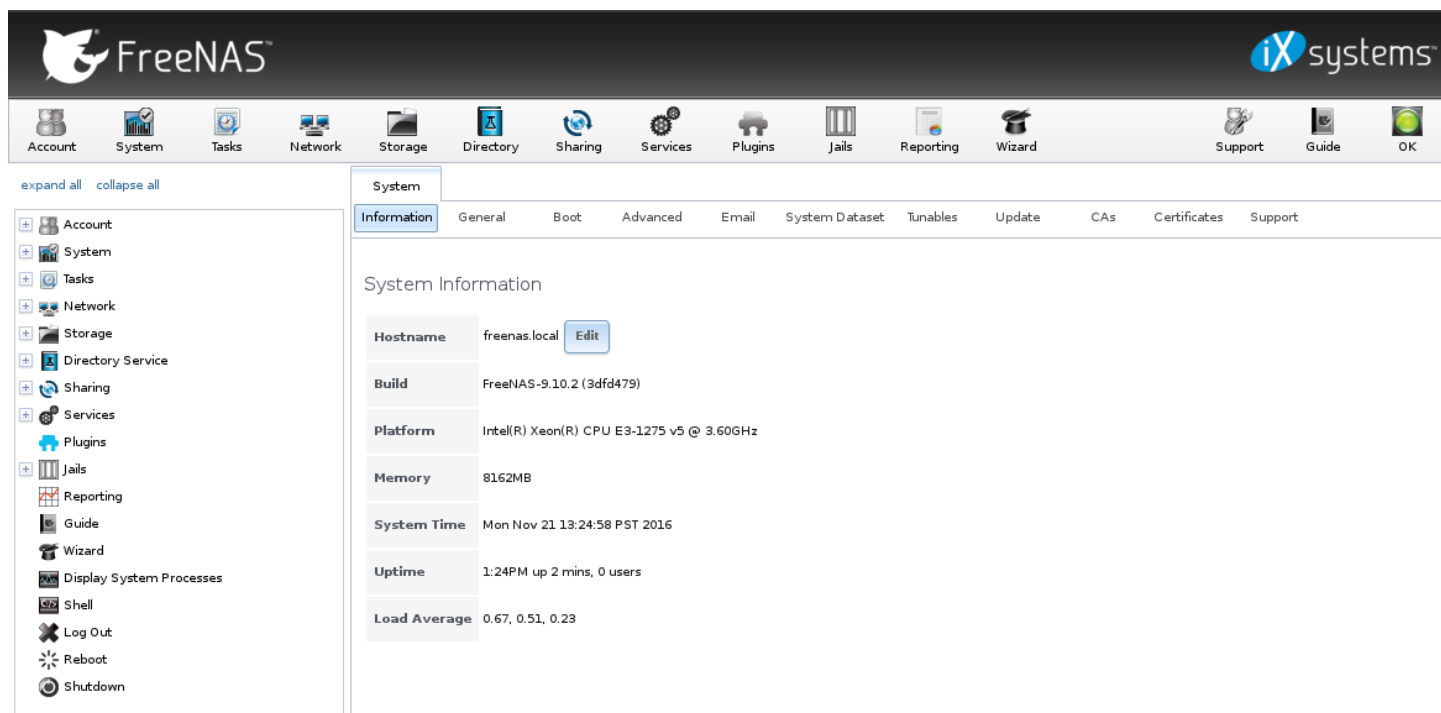


Fig. 3.3: FreeNAS® Graphical Configuration Menu

If the FreeNAS® system does not respond to the IP address or mDNS name entered in a browser:

- If proxy settings are enabled in the browser configuration, disable them and try connecting again.
- If the page does not load, check whether the FreeNAS® system's IP address responds to a `ping` from another computer on the same network. If the FreeNAS® IP address is in a private IP address range, it can only be accessed from within

that private network.

- If the user interface loads but is unresponsive or seems to be missing menu items, try a different web browser. IE9 has known issues and does not display the graphical administrative interface correctly if compatibility mode is turned on. [Firefox](https://www.mozilla.org/en-US/firefox/all/) (<https://www.mozilla.org/en-US/firefox/all/>) is recommended.
- If An error occurred! messages are shown when attempting to configure an item in the GUI, make sure that the browser is set to allow cookies from the FreeNAS® system.

This [blog post](http://fortysomethinggeek.blogspot.com/2012/10/ipad-iphone-connect-with-freenas-or-any.html) (<http://fortysomethinggeek.blogspot.com/2012/10/ipad-iphone-connect-with-freenas-or-any.html>) describes some applications which can be used to access the FreeNAS® system from an iPad or iPhone.

3.3 Initial Configuration

The first time the FreeNAS® GUI is accessed, the [Wizard](#) (page 282) starts automatically to help configure the FreeNAS® device quickly and easily.

ACCOUNT

The Account Configuration section of the administrative GUI describes how to manually create and manage users and groups. This section contains these entries:

- [Groups](#) (page 48): used to manage UNIX-style groups on the FreeNAS® system.
- [Users](#) (page 51): used to manage UNIX-style accounts on the FreeNAS® system.

Each entry is described in more detail in this section.

4.1 Groups

The Groups interface provides management of UNIX-style groups on the FreeNAS® system.

Note: If a directory service is running on the network, it is not necessary to recreate the network's users or groups. Instead, import the existing account information into FreeNAS®. Refer to [Directory Services](#) (page 158) for details.

This section describes how to create a group and assign user accounts to it. The next section, [Users](#) (page 51), describes creating user accounts.

Click `Groups` → `View Groups` to see a screen like [Figure 4.1](#).

Account

Groups

Users

Add Group

Group ID	Group Name	Built-in Group	Permit Sudo
0	wheel	true	false
1	daemon	true	false
2	kmem	true	false
3	sys	true	false
4	tty	true	false
5	operator	true	false
6	mail	true	false
7	bin	true	false
8	news	true	false
9	man	true	false
13	games	true	false
14	ftp	true	false
20	staff	true	false
22	sshd	true	false
25	smmsp	true	false
26	mailnull	true	false
31	guest	true	false
53	bind	true	false

Members

Fig. 4.1: Group Management

All groups that came with the operating system will be listed. Each group has an entry indicating the group ID, group name, whether or not it is a built-in group which was installed with FreeNAS®, and whether or not the group members are allowed to use `sudo`. Clicking a group entry causes a Members button to appear. Click the button to view and modify the group membership.

The Add Group button opens the screen shown in Figure 4.2. Table 4.1 summarizes the available options when creating a group.

Add Group

Group ID:

1001

Group Name:

Permit Sudo:

☐

Allow repeated GIDs:

☐

OK

Cancel

Fig. 4.2: Creating a New Group

Table 4.1: Group Creation Options

Setting	Value	Description
Group ID	string	the next available group ID will be suggested for you; by convention, UNIX groups containing user accounts have an ID greater than 1000 and groups required by a service have an ID equal to the default port number used by the service (e.g. the sshd group has an ID of 22)
Group Name	string	mandatory
Permit Sudo	checkbox	if checked, members of the group have permission to use <code>sudo</code> (http://www.sudo.ws/); when using sudo, a user will be prompted for their own password
Allow repeated GIDs	checkbox	allows multiple groups to share the same group id (GID); this is useful when a GID is already associated with the UNIX permissions for existing data

After a group and users are created, users can be made members of a group. Highlight the group where users will be assigned, then click the Members button. Highlight the user in the Member users list (which shows all user accounts on the system) and click >> to move that user to the right frame. The user accounts which appear in the right frame are added as members of the group.

In the example shown in Figure 4.3, the data1 group has been created and the user1 user account has been created with a primary group of user1. The Members button for the data1 group has been selected and user1 has been added as a member of the group.

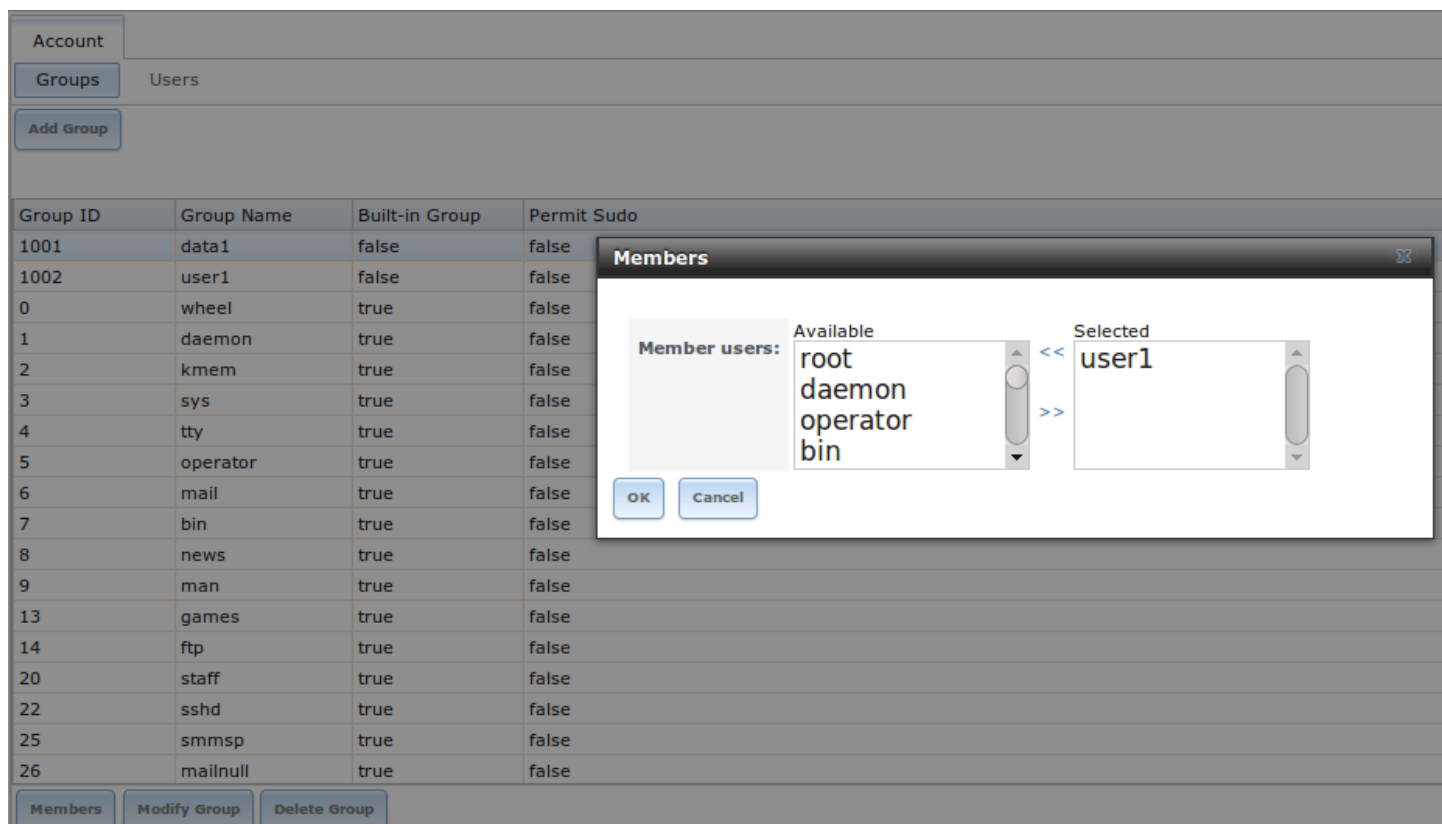


Fig. 4.3: Assigning a User to a Group

The Delete Group button deletes a group. The pop-up message asks whether all members of that group should also be deleted. Note that the built-in groups do not provide a Delete Group button.

4.2 Users

FreeNAS® supports users, groups, and permissions, allowing great flexibility in configuring which users have access to the data stored on FreeNAS®. To assign permissions to shares, one of the following must be done:

1. Create a guest account that all users will use or create a user account for every user in the network where the name of each account is the same as a logon name used on a computer. For example, if a Windows system has a login name of bobsmith, create a user account with the name bobsmith on FreeNAS®. A common strategy is to create groups with different sets of permissions on shares, then assign users to those groups.
2. If your network uses a directory service, import the existing account information using the instructions in [Directory Services](#) (page 158).

Account → Users → View Users provides a listing of all of the system accounts that were installed with the FreeNAS® operating system, as shown in [Figure 4.4](#).

Account											
Groups											
Users											
Add User											
User ID	Username	Primary Group ID	Home Directory	Shell	Full Name	Built-in User	E-mail	Disable password login	Lock user	Permit Sudo	Microsoft Account
0	root	0	/root	/bin/csh	root	true		false	false	false	false
1	daemon	1	/root	/usr/sbin/nologin	Owner of many system processes	true		false	false	false	false
2	operator	5	/	/usr/sbin/nologin	System &	true		false	false	false	false
3	bin	7	/	/usr/sbin/nologin	Binaries Commands and Source	true		false	false	false	false
4	tty	65533	/	/usr/sbin/nologin	Tty Sandbox	true		false	false	false	false
5	kmem	2	/	/usr/sbin/nologin	KMem Sandbox	true		false	false	false	false
7	games	13	/	/usr/sbin/nologin	Games pseudo-user	true		false	false	false	false
8	news	8	/	/usr/sbin/nologin	News Subsystem	true		false	false	false	false
9	man	9	/usr/share/man	/usr/sbin/nologin	Mister Man Pages	true		false	false	false	false
14	ftp	14	/nonexistent	/bin/csh		true		false	false	false	false
22	sshd	22	/var/empty	/usr/sbin/nologin	Secure Shell Daemon	true		false	false	false	false
25	smmsp	25	/var/spool/clientmqueue	/usr/sbin/nologin	Sendmail Submission User	true		false	false	false	false
26	mailnull	26	/var/spool/mqueue	/usr/sbin/nologin	Sendmail Default User	true		false	false	false	false
53	bind	53	/	/usr/sbin/nologin	Bind Sandbox	true		false	false	false	false
62	proxy	62	/nonexistent	/usr/sbin/nologin	Packet Filter pseudo-user	true		false	false	false	false
64	_pflogd	64	/var/empty	/usr/sbin/nologin	pflogd privsep user	true		false	false	false	false
65	_dhcp	65	/var/empty	/usr/sbin/nologin	dhcp programs	true		false	false	false	false
66	uucp	66	/var/spool/uucppublic	/usr/local/libexec/uucp/uucico	UUCP pseudo-user	true		false	false	false	false
68	pop	6	/nonexistent	/usr/sbin/nologin	Post Office Owner	true		false	false	false	false
78	auditdistd	77	/var/empty	/usr/sbin/nologin	Auditdistd unprivileged user	true		false	false	false	false
79	ladvd	78	/var/empty	/usr/sbin/nologin	ladvd user	true		false	false	false	false
80	www	80	/nonexistent	/usr/sbin/nologin	World Wide Web Owner	true		false	false	false	false
Modify User											

Fig. 4.4: Managing User Accounts

Each account entry indicates the user ID, username, primary group ID, home directory, default shell, full name, whether it is a built-in user that came with the FreeNAS® installation, the email address, whether logins are disabled, whether the user account is locked, whether the user is allowed to use `sudo`, and if the user connects from a Windows 8 or higher system. To reorder the list, click the desired column name. An arrow indicates which column controls the view sort order. Click the arrow to reverse the sort order.

Click a user account to cause these buttons to appear:

- **Modify User:** used to modify the account's settings, as listed in [Table 4.2](#).
- **Change E-mail:** used to change the email address associated with the account.

Note: It is important to set the email address for the built-in root user account as important system messages are sent to the root user. For security reasons, password logins are disabled for the root account and changing this setting is highly discouraged.

Except for the root user, the accounts that come with FreeNAS® are system accounts. Each system account is used by a service and should not be used as a login account. For this reason, the default shell on system accounts is `nologin(8)` (<http://www.freebsd.org/cgi/man.cgi?query=nologin>). For security reasons, and to prevent breakage of system services, do not modify the system accounts.

The Add User button opens the screen shown in [Figure 4.5](#). Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in `System → Advanced`. [Table 4.2](#) summarizes the options which are available when user accounts are created or modified.

Warning: When using [Active Directory](#) (page 158), Windows user passwords must be set from within Windows.

The screenshot shows the 'Add User' web interface. The form is titled 'Add User' and contains the following fields and controls:

- User ID:** Text input field containing 'h001'.
- Username:** Text input field.
- Create a new primary group for the user:** Checked checkbox.
- Primary Group:** Dropdown menu showing '-----'.
- Create Home Directory In:** Text input field containing '/nonexistent', with a 'Browse' button to its right.
- Shell:** Dropdown menu showing 'csh'.
- Full Name:** Text input field.
- E-mail:** Text input field.
- Password:** Text input field.
- Password confirmation:** Text input field, with an information icon to its right.
- Disable password login:** Unchecked checkbox.
- Lock user:** Unchecked checkbox.

Fig. 4.5: Adding or Editing a User Account

Table 4.2: User Account Configuration

Setting	Value	Advanced Mode	Description
User ID	integer		grayed out if user already created; when creating an account, the next numeric ID will be suggested; by convention, user accounts have an ID greater than 1000 and system accounts have an ID equal to the default port number used by the service
Username	string		grayed out if user already created; maximum 16 characters though a maximum of 8 is recommended for interoperability; cannot begin with a hyphen, if a \$ is used it can only be the last character, and it cannot contain a space, tab, or the characters , : + & # % ^ & () ! @ ~ * ? < > =
Create a new primary group	checkbox		by default, a primary group with the same name as the user will be created; uncheck this box to select a different primary group name
Primary Group	drop-down menu		must uncheck Create a new primary group to access this menu; for security reasons, FreeBSD will not give a user <code>su</code> permissions if wheel is their primary group; to give a user <code>su</code> access, add them to the wheel group in Auxiliary groups
Create Home Directory In	browse button		browse to the name of an existing volume or dataset that the user will be assigned permission to access
Home Directory Mode	checkboxes	✓	sets default Unix permissions of user's home directory; read-only for built-in users
Shell	drop-down menu		select shell to use for local and SSH logins; see Table 4.3 for an overview of available shells
Full Name	string		mandatory, may contain spaces
E-mail	string		email address associated with the account
Password	string		mandatory unless check box Disable password login; cannot contain a ?
Password confirmation	string		must match the value of Password
Disable password login	checkbox		when checked, disables password logins and authentication to SMB shares; to undo this setting, set a password for the user using the Modify User button for the user in View Users; checking this box grays out Lock user and Permit Sudo, which are mutually exclusive
Lock user	checkbox		a checked box prevents user from logging in until the account is unlocked (box is unchecked); checking this box will gray out Disable password login which is mutually exclusive
Permit Sudo	checkbox		if checked, members of the group have permission to use <code>sudo</code> (http://www.sudo.ws/); when using sudo, a user will be prompted for their own password
Microsoft Account	checkbox		check this box if the user will be connecting from a Windows 8 or higher system
SSH Public Key	string		paste the user's public SSH key to be used for key-based authentication (do not paste the private key!)
Auxiliary groups	mouse selection		highlight the groups to which the user is to be added; click the >> button to add the user to the highlighted groups

Note: Some fields cannot be changed for built-in users and will be grayed out.

Table 4.3: Available Shells

Shell	Description
netcli.sh	user is shown the Console Setup menu (Figure 3.1) on connection, even if it is disabled in <code>System → Advanced → Enable Console Menu</code> ; the user must be root or have root permissions (effective user ID 0, like toor)
csch	C shell (https://en.wikipedia.org/wiki/C_shell)
sh	Bourne shell (https://en.wikipedia.org/wiki/Bourne_shell)
tcsh	Enhanced C shell (https://en.wikipedia.org/wiki/Tcsh)
nologin	use when creating a system account or to create a user account that can authenticate with shares but which cannot login to the FreeNAS system using <code>ssh</code>
bash	Bourne Again shell (https://en.wikipedia.org/wiki/Bash_%28Unix_shell%29)
ksh93	Korn shell (http://www.kornshell.com/)
mksh	mirBSD Korn shell (https://www.mirbsd.org/mksh.htm)
rbash	Restricted bash (http://www.gnu.org/software/bash/manual/html_node/The-Restricted-Shell.html)
rzsh	Restricted zsh (http://www.csse.uwa.edu.au/programming/linux/zsh-doc/zsh_14.html)
scponly	select scponly (https://github.com/scponly/scponly/wiki) to restrict the user's SSH usage to only the <code>scp</code> and <code>sftp</code> commands
zsh	Z shell (http://www.zsh.org/)
git-shell	restricted git shell (http://git-scm.com/docs/git-shell)

Built-in user accounts needed by the system cannot be removed. A Remove User button appears for custom users that have been added by the system administrator. If the user to be removed is the last user in a custom group, an option is presented to delete the group as well.

SYSTEM

The System section of the administrative GUI contains these entries:

- [Information](#) (page 55) provides general FreeNAS® system information such as hostname, operating system version, platform, and uptime
- [General](#) (page 56) configures general settings such as HTTPS access, the language, and the timezone
- [Boot](#) (page 59) creates, renames, and deletes boot environments
- [Advanced](#) (page 63) configures advanced settings such as the serial console, swap space, and console messages
- [Email](#) (page 65) configures the email address to receive notifications
- [System Dataset](#) (page 67) configures the location where logs and reporting graphs are stored
- [Tunables](#) (page 68) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- [Update](#) (page 70) performs upgrades and checks for system updates
- [Cloud Credentials](#) (page 73) is used to enter connection credentials for remote cloud service providers
- [Alert Services](#) (page 74) configures services used to notify the administrator about system events.
- [CAs](#) (page 75): import or create internal or intermediate CAs (Certificate Authorities)
- [Certificates](#) (page 78): import existing certificates or create self-signed certificates
- [Support](#) (page 81): report a bug or request a new feature.

Each of these is described in more detail in this section.

5.1 Information

`System` → `Information` displays general information about the FreeNAS® system. An example is seen in [Figure 5.1](#).

The information includes the hostname, the build version, type of CPU (platform), the amount of memory, the current system time, the system's uptime, the number of users connected at the console or by serial, telnet, or SSH connections, and the current load average. On systems supplied or certified by iXsystems, an additional Serial Number field showing the hardware serial number is displayed.

To change the system's hostname, click the Edit button, type in the new hostname, and click OK. The hostname must include the domain name. If the network does not use a domain name, add `.local` after the hostname.

System

InformationGeneralBootAdvancedEmailSystem DatasetTunablesUpdateCloud CredentialsAlert ServicesCAsCertificatesSupport

System Information


Hostname	freenas.local	Edit
Build	FreeNAS-11.1-U2	
Platform	Intel(R) Atom(TM) CPU C2750 @ 2.40GHz	
Memory	16319MB	
System Time	Wed, 14 Feb 2018 14:15:03 -0500	
Uptime	2:15PM up 1:52, 0 users	
Load Average	0.46, 0.50, 0.46	
System Serial		
System Product	FREENAS-MINI-2.0	

Fig. 5.1: System Information Tab

5.2 General

System → General is shown in Figure 5.2.

The screenshot shows the 'General' configuration tab for the FreeNAS system. The settings are as follows:

- Protocol:** HTTP
- Certificate:** (empty dropdown)
- WebGUI IPv4 Address:** 0.0.0.0
- WebGUI IPv6 Address:** ::
- WebGUI HTTP Port:** 80
- WebGUI HTTPS Port:** 443
- WebGUI HTTP -> HTTPS Redirect:** ☒ (with info icon)
- Language (Require UI reload):** English
- Console Keyboard Map:** (empty dropdown)
- Timezone:** America/Los_Angeles
- Syslog level:** Info (with info icon)
- Syslog server:** (empty text field with info icon)

At the bottom of the form are five buttons: **Save**, **Reset Configuration to Defaults**, **Save Config**, **Upload Config**, and **HTTP Servers**.

Fig. 5.2: General Screen

Table 5.1 summarizes the settings that can be configured using the General tab:

Table 5.1: General Configuration Settings

Setting	Value	Description
Protocol	drop-down menu	protocol to use when connecting to the administrative GUI from a browser; if modified from the default of HTTP to HTTPS or to HTTP+HTTPS, select the certificate to use in Certificate; if you do not have a certificate, first create a CA (in CAs (page 75)), then the certificate itself (in Certificates (page 78))
Certificate	drop-down menu	required for HTTPS; browse to the location of the certificate to use for encrypted connections
WebGUI IPv4 Address	drop-down menu	choose from a list of recent IP addresses to limit the one to use when accessing the administrative GUI; the built-in HTTP server will automatically bind to the wildcard address of 0.0.0.0 (any address) and will issue an alert if the specified address becomes unavailable
WebGUI IPv6 Address	drop-down menu	choose from a list of recent IPv6 addresses to limit the one to use when accessing the administrative GUI; the built-in HTTP server will automatically bind to any address and will issue an alert if the specified address becomes unavailable
WebGUI HTTP Port	integer	allows configuring a non-standard port for accessing the administrative GUI over HTTP; changing this setting might also require changing a Firefox configuration setting (http://www.redbrick.dcu.ie/%7Ed_fens/articles/Firefox:_This_Address_is_Restricted)
WebGUI HTTPS Port	integer	allows configuring a non-standard port for accessing the administrative GUI over HTTPS

Continued on next page

Table 5.1 – continued from previous page

Setting	Value	Description
WebGUI HTTP → HTTPS Redirect	checkbox	when this box is checked, HTTP connections are automatically redirected to HTTPS if HTTPS is selected in Protocol, otherwise such connections will fail
Language	drop-down menu	select the localization from the drop-down menu and reload the browser; view the status of localization at pootle.freenas.org (http://pootle.freenas.org/)
Console Keyboard Map	drop-down menu	select the keyboard layout
Timezone	drop-down menu	select the timezone from the drop-down menu
Syslog level	drop-down menu	when Syslog server is defined, only logs matching this level are sent
Syslog server	string	IP address_or_hostname:optional_port_number of remote syslog server to send logs to; once set, log entries are written to both the console and the remote server

After making any changes, click the Save button.

This screen also contains these buttons:

Reset Configuration to Defaults: reset the configuration database to the default base version. This does not delete user SSH keys or any other data stored in a user's home directory. Since configuration changes stored in the configuration database are erased, this option is useful when a mistake has been made or to return a test system to the original configuration.

Save Config: save a backup copy of the current configuration database in the format hostname-version-architecture to the computer accessing the administrative interface. Saving the configuration after making any configuration changes is highly recommended. FreeNAS® automatically backs up the configuration database to the system dataset every morning at 3:45. However, this backup does not occur if the system is shut down at that time. If the system dataset is stored on the boot pool and the boot pool becomes unavailable, the backup will also not be available. The location of the system dataset can be viewed or set using `System → System Dataset`.

Note: [SSH](#) (page 236) keys are not stored in the configuration database and must be backed up separately.

There are two types of passwords. User account passwords for the base operating system are stored as hashed values, do not need to be encrypted to be secure, and are saved in the system configuration backup. Other passwords, like iSCSI CHAP passwords or Active Directory bind credentials, are stored in an encrypted form to prevent them from being visible as plain text in the saved system configuration. The key or seed for this encryption is normally stored only on the boot device. When Save Config is chosen, a dialog gives the option to Export Password Secret Seed with the saved configuration, allowing the configuration file to be restored to a different boot device where the decryption seed is not already present. Configuration backups containing the seed must be physically secured to prevent decryption of passwords and unauthorized access.

Warning: The Export Password Secret Seed option is off by default and should only be used when making a configuration backup that will be stored securely. After moving a configuration to new hardware, media containing a configuration backup with a decryption seed should be securely erased before reuse.

Upload Config: allows browsing to the location of a previously saved configuration file to restore that configuration. The screen turns red as an indication that the system will need to reboot to load the restored configuration.

NTP Servers: The network time protocol (NTP) is used to synchronize the time on the computers in a network. Accurate time is necessary for the successful operation of time sensitive applications such as Active Directory or other directory services. By default, FreeNAS® is pre-configured to use three public NTP servers. If your network is using a directory service, ensure that the FreeNAS® system and the server running the directory service have been configured to use the same NTP servers.

Available NTP servers can be found at <https://support.ntp.org/bin/view/Servers/NTPPoolServers>. For time accuracy, choose NTP servers that are geographically close to the FreeNAS® system's physical location.

NTP servers are added by clicking on NTP Servers → Add NTP Server to open the screen shown in Figure 5.3. Table 5.2 summarizes the options available when adding an NTP server. [ntp.conf\(5\)](http://www.freebsd.org/cgi/man.cgi?query=ntp.conf) (<http://www.freebsd.org/cgi/man.cgi?query=ntp.conf>) explains these options in more detail.




Fig. 5.3: Add an NTP Server

Table 5.2: NTP Servers Configuration Options

Setting	Value	Description
Address	string	name of NTP server
Burst	checkbox	recommended when Max. Poll is greater than 10; only use on your own servers i.e. do not use with a public NTP server
IBurst	checkbox	speeds the initial synchronization (seconds instead of minutes)
Prefer	checkbox	should only be used for NTP servers that are known to be highly accurate, such as those with time monitoring hardware
Min. Poll	integer	power of 2 in seconds; cannot be lower than 4 or higher than Max. Poll
Max. Poll	integer	power of 2 in seconds; cannot be higher than 17 or lower than Min. Poll
Force	checkbox	forces the addition of the NTP server, even if it is currently unreachable

5.3 Boot

FreeNAS® supports a ZFS feature known as multiple boot environments. With multiple boot environments, the process of updating the operating system becomes a low-risk operation. The updater automatically creates a snapshot of the current boot environment and adds it to the boot menu before applying the update. If the update fails, reboot the system and select the previous boot environment from the boot menu to instruct the system to go back to that system state.

Note: Boot environments are separate from the configuration database. Boot environments are a snapshot of the operating system at a specified time. When a FreeNAS® system boots, it loads the specified boot environment, or operating system, then reads the configuration database in order to load the current configuration values. If the intent is to make configuration changes rather than operating system changes, make a backup of the configuration database first using **System** → **General** → **Save Config**.

As seen in [Figure 5.4](#), two boot environments are created when FreeNAS® is installed. The system will boot into the default boot environment and users can make their changes and update from this version. The other boot environment, named Initial-Install can be booted into if the system needs to be returned to a pristine, non-configured version of the installation.

If the [Wizard](#) (page 282) was used, a third boot environment called `Wizard-date` is also created, indicating the date and time the [Wizard](#) (page 282) was run.

System			
Information	General	Boot	Advanced Email System Dataset Tunables Update Cloud Credentials Alert Services CAs Certificates Support
Create	Scrub Boot	Status	Boot Volume Condition: HEALTHY Size: 7.9 GiB Used: 949.1 MiB (11%)
7 Automatic scrub interval (in days)			
Name	Active	Created	Keep
default	On Reboot, Now	2017-10-16 02:16:00	No
Initial-Install		2017-10-16 02:22:00	No

Fig. 5.4: Viewing Boot Environments

Each boot environment entry contains this information:

- Name: the name of the boot entry as it will appear in the boot menu.
- Active: indicates which entry will boot by default if the user does not select another entry in the boot menu.
- Created: indicates the date and time the boot entry was created.
- Keep: indicates whether or not this boot environment can be pruned if an update does not have enough space to proceed. Click the entry's Keep button if that boot environment should not be automatically pruned.


Highlight an entry to view the configuration buttons for it. These configuration buttons are shown:

- Rename: used to change the name of the boot environment.
- Keep/Unkeep: used to toggle whether or not the updater can prune (automatically delete) this boot environment if there is not enough space to proceed with the update.
- Clone: used to create a copy of the highlighted boot environment.
- Delete: used to delete the highlighted entry, which also removes that entry from the boot menu. Since you cannot delete an entry that has been activated, this button will not appear for the active boot environment. If you need to delete an entry that is currently activated, first activate another entry, which will clear the On reboot field of the currently activated entry. Note that this button will not be displayed for the default boot environment as this entry is needed in order to return the system to the original installation state.
- Activate: only appears on entries which are not currently set to Active. Changes the selected entry to the default boot entry on next boot. Its status changes to On Reboot and the current Active entry changes from On Reboot, Now to Now, indicating that it was used on the last boot but will not be used on the next boot.

The buttons above the boot entries can be used to:

- Create: a manual boot environment. A pop-up menu prompts for entry of a Name for the boot environment. When entering the name, only alphanumeric characters, underscores, and dashes are allowed.
- Scrub Boot: can be used to perform a manual scrub of the boot devices. By default, the boot device is scrubbed every 7 days. To change the default interval, change the number in the Automatic scrub interval (in days) field. The date and results of the last scrub are also listed in this screen. The condition of the boot device should be listed as HEALTHY.

- Status: click this button to see the status of the boot devices. [Figure 5.5](#), shows only one boot device, which is ONLINE.



Boot Status				
Name	Read	Write	Checksum	Status
▲ freenas-boot	0	0	0	ONLINE
▲ stripe	0	0	0	ONLINE
da0p2	0	0	0	ONLINE



Replace

Fig. 5.5: Viewing the Status of the Boot Device

If the system has a mirrored boot pool, there will be a Detach button in addition to the Replace button. To remove a device from the boot pool, highlight the device and click its Detach button. Alternately, if one of the boot devices has an OFFLINE Status, click the device to replace, then click Replace to rebuild the boot mirror.

Note that you cannot replace the boot device if it is the only boot device as it contains the operating system itself.

[Figure 5.6](#) shows a sample boot menu.

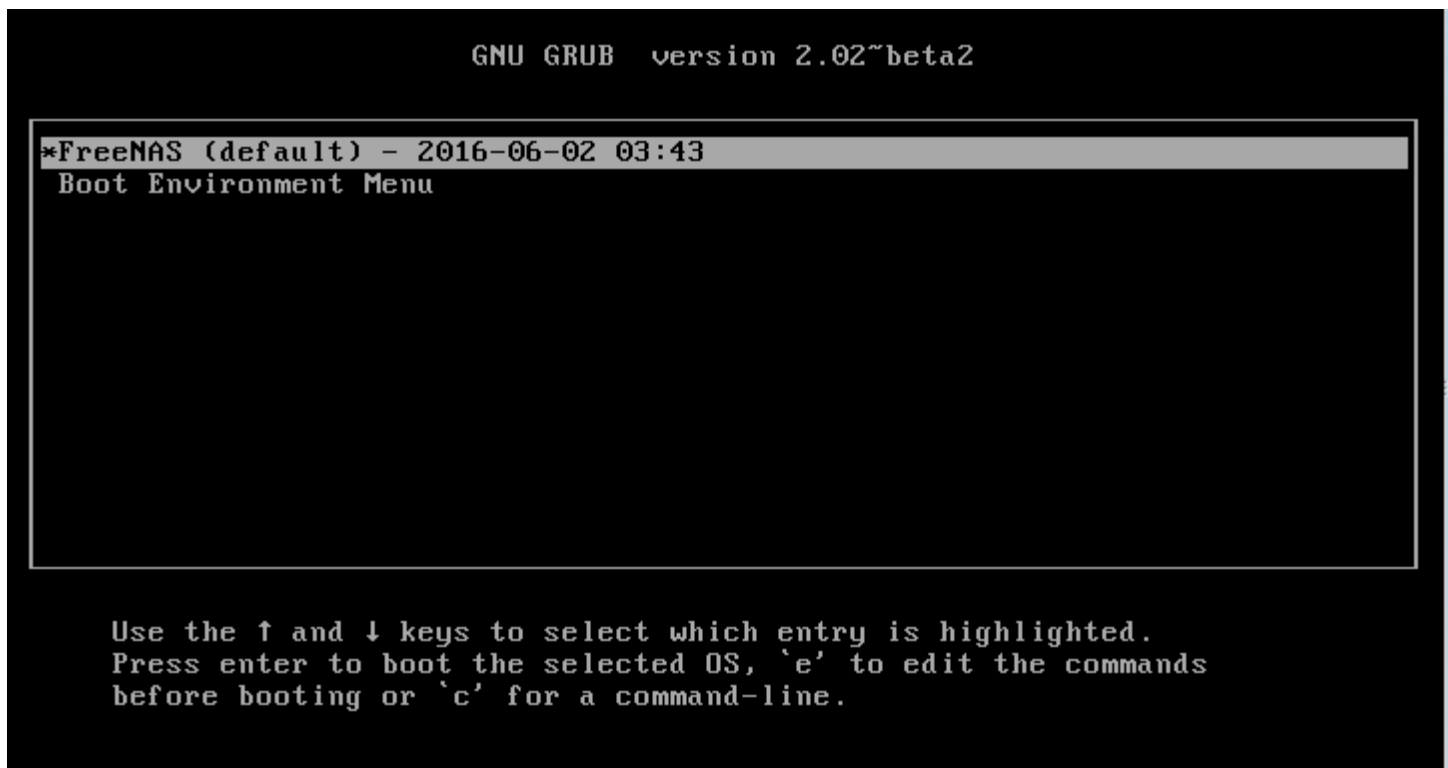


Fig. 5.6: Boot Environments in Boot Menu

The first entry is the active boot environment, or the one that the system has been configured to boot into. To boot into a different boot environment, press the `spacebar` to pause this screen, use the down arrow to select `Boot Environment Menu`, and press `Enter`. A menu displays the other available boot environments. Use the up/down arrows to select the desired boot environment and press `Enter` to boot into it. To always boot into that boot environment, go to `System → Boot`, highlight that entry, and click the `Activate` button.

5.3.1 Mirroring the Boot Device

If the system is currently booting from one device, you can add another device to create a mirrored boot device. This way, if one device fails, the system still has a copy of the boot file system and can be configured to boot from the remaining device in the mirror.

Note: When adding another boot device for a mirror, the new device must have at least the same capacity as the existing boot device. Larger capacity devices can be added, but the mirror will only have the capacity of the smallest device. Different models of devices which advertise the same nominal size are not necessarily the same actual size. For this reason, adding another of the same model of boot device is recommended.

In the example shown in [Figure 5.7](#), the user has clicked `System → Boot → Status` to display the current status of the boot device. The example indicates that there is currently one device, `ada0p2`, its status is `ONLINE`, and it is currently the only boot device as indicated by the word `stripe`. To create a mirrored boot device, click either the entry called `freenas-boot` or `stripe`, then click the `Attach` button. If another device is available, it appears in the `Member disk` drop-down menu. Select the desired device.

The `Use all disk space` checkbox gives control of how much of the new device is made available to ZFS. The default is unchecked, and the new device will be partitioned to the same size as the existing device. If either device in the mirror fails, it can be replaced with another of the same size as the original boot device.

When `Use all disk space` is checked, the entire capacity of the new device is used. If the original boot device fails and is removed, the boot mirror will consist of just the newer drive, and will grow to whatever capacity it provides. However, new

devices added to this mirror must now be as large as the new capacity.
Click Attach Disk to attach the new disk to the mirror.

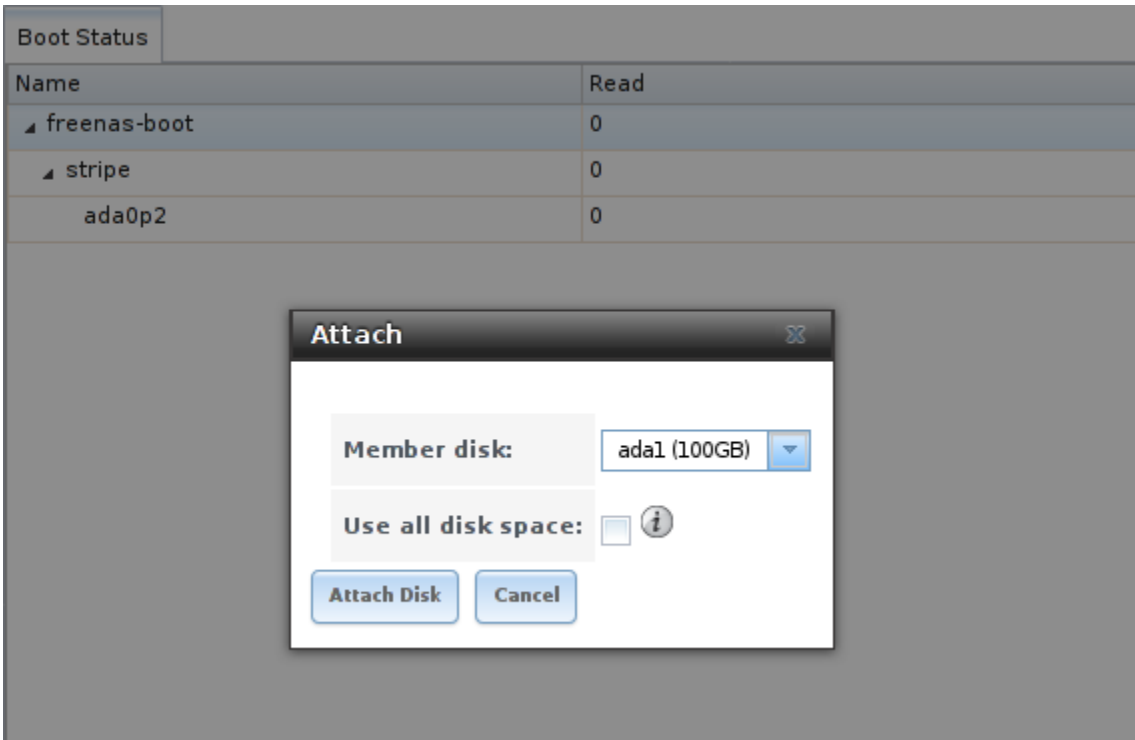


Fig. 5.7: Mirroring a Boot Device

After the mirror is created, the Status screen indicates that it is now a mirror. The number of devices in the mirror are shown, as seen in the example in [Figure 5.8](#).

Boot Status				
Name	Read	Write	Checksum	Status
▲ freenas-boot	0	0	0	ONLINE
▲ mirror-0	0	0	0	ONLINE
ada1p2	0	0	0	ONLINE
ada0p2	0	0	0	ONLINE

Fig. 5.8: Viewing the Status of a Mirrored Boot Device

5.4 Advanced

System → Advanced is shown in [Figure 5.9](#). The configurable settings are summarized in [Table 5.3](#).

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Alert Services
CAs
Certificates
Support

Enable Console Menu:
☒

Use Serial Console:
☐

Serial Port Address:
0x2f8 ⓘ

Serial Port Speed:
9600 ⓘ

Enable screen saver:
☐

Enable powerd (Power Saving Daemon):
☐

Swap size on each drive in GiB, affects new disks only. Setting this to 0 disables swap creation completely (STRONGLY DISCOURAGED).
2

Show console messages in the footer:
☐

Show tracebacks in case of fatal errors:
☒

Show advanced fields by default:
☐ ⓘ

Enable autotune:
☐ ⓘ

Enable debug kernel:
☐ ⓘ

Enable automatic upload of kernel crash dumps and daily telemetry:
☒

MOTD banner:
Welcome to FreeNAS

Periodic Notification User:
root ⓘ

Remote Graphite Server Hostname:
 ⓘ

Use FQDN for logging:
☐

Save
Save Debug

Fig. 5.9: Advanced Screen

Table 5.3: Advanced Configuration Settings

Setting	Value	Description
Show Text Console Without Password Prompt	checkbox	unchecking this box replaces the console menu shown in Figure 3.1 with a login prompt
Use Serial Console	checkbox	do not check this box if the serial port is disabled
Serial Port Address	string	serial port address in hex
Serial Port Speed	drop-down menu	select the speed used by the serial port
Enable screen saver	checkbox	enable or disable the console screen saver
Enable powerd (Power Saving Daemon)	checkbox	powerd(8) (http://www.freebsd.org/cgi/man.cgi?query=powerd) monitors the system state and sets the CPU frequency accordingly
Swap size	non-zero integer representing GB	by default, all data disks are created with this amount of swap; this setting does not affect log or cache devices as they are created without swap

Continued on next page

Table 5.3 – continued from previous page

Setting	Value	Description
Show console messages in the footer	checkbox	display console messages in real time at bottom of browser; click the console to bring up a scrollable screen; check the Stop refresh box in the scrollable screen to pause updating and uncheck the box to continue to watch the messages as they occur
Show tracebacks in case of fatal errors	checkbox	provides a pop-up of diagnostic information when a fatal error occurs
Show advanced fields by default	checkbox	several GUI menus provide an Advanced Mode button to access additional features; enabling this shows these features by default
Enable autotune	checkbox	enables Autotune (page 65) which attempts to optimize the system depending upon the hardware which is installed
Enable debug kernel	checkbox	when checked, next boot uses a debug version of the kernel
Enable automatic upload of kernel crash dumps and daily telemetry	checkbox	when checked, kernel crash dumps and telemetry (some system stats, collectedd RRDs, and select syslog messages) are automatically sent to the development team for diagnosis
MOTD banner	string	message to be shown when a user logs in with SSH
Periodic Notification User	drop-down menu	user to receive security output emails; this output runs nightly but only sends an email when the system reboots or encounters an error
Report CPU usage in percentage	checkbox	when checked, CPU usage is reported as percentages in Reporting (page 280)
Remote Graphite Server hostname	string	IP address or hostname of a remote server running Graphite (http://graphite.wikidot.com/)
Use FQDN for logging	checkbox	when checked, include the Fully-Qualified Domain Name in logs to precisely identify systems with similar hostnames

Click the Save button after making any changes.

This tab also contains this button:

Save Debug: used to generate a text file of diagnostic information. After the debug data is collected, the system prompts for a location to save the generated ASCII text file.

5.4.1 Autotune

FreeNAS® provides an autotune script which optimizes the system depending on the installed hardware. For example, if a ZFS volume exists on a system with limited RAM, the autotune script automatically adjusts some ZFS sysctl values in an attempt to minimize ZFS memory starvation issues. It should only be used as a temporary measure on a system that hangs until the underlying hardware issue is addressed by adding more RAM. Autotune will always slow such a system, as it caps the ARC.

The Enable autotune checkbox in *System* → *Advanced* is unchecked by default. Check this box to run the autotuner at boot time. If you would like the script to run immediately, the system must be rebooted.

If the autotune script adjusts any settings, the changed values appear in *System* → *Tunables*. These values can be modified and overridden. Note that deleting tunables that were created by autotune only affects the current session, as autotune-set tunables are recreated at boot.

When attempting to increase the performance of the FreeNAS® system, and particularly when the current hardware may be limiting performance, try enabling autotune.

For those who wish to see which checks are performed, the autotune script is located in `/usr/local/bin/autotune`.

5.5 Email

An automatic script sends a nightly email to the root user account containing important information such as the health of the disks. [Alert](#) (page 297) events are also emailed to the root user account. Problems with [Scrubs](#) (page 152) are reported

separately in an email sent at 03:00AM.

Note: [S.M.A.R.T.](#) (page 228) reports are mailed separately to the address configured in that service.

The administrator typically does not read email directly on the FreeNAS® system. Instead, these emails are usually sent to an external email address where they can be read more conveniently. It is important to configure the system so it can send these emails to the administrator's remote email account so they are aware of problems or status changes.

The first step is to set the remote address where email will be sent. Select **Account** → **Users**, click on **root** to highlight that user, then click **Modify User**. In the **E-mail** field, enter the email address on the remote system where email is to be sent, like **admin@example.com**. Click **OK** to save the settings.

Additional configuration is performed with **System** → **Email**, shown in [Figure 5.10](#).

The screenshot shows the 'System' configuration page with the 'Email' tab selected. The configuration fields are as follows:

- From email:** root@freenas.local
- Outgoing mail server:** (empty)
- Port to connect to:** 25
- TLS/SSL:** Plain
- Use SMTP Authentication:** ☐
- Username:** (empty)
- Password:** (empty)
- Password confirmation:** (empty)

A hint at the bottom states: "HINT: Test e-mails are sent to root user. To configure it use Account -> Users -> View Users -> root -> Change E-mail". At the bottom left are 'Save' and 'Send Test Mail' buttons.

Fig. 5.10: Email Screen

Table 5.4: Email Configuration Settings

Setting	Value	Description
From email	string	the envelope From address shown in the email; this can be set to assist with filtering mail on the receiving system
Outgoing mail server	string or IP address	hostname or IP address of SMTP server to use for sending this email
Port to connect to	integer	SMTP port number, typically 25, 465 (secure SMTP), or 587 (submission)
TLS/SSL	drop-down menu	encryption type; choices are Plain, SSL, or TLS
Use SMTP Authentication	checkbox	enable/disable SMTP AUTH (http://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN SASL; if checked, enter the required Username and Password
Username	string	enter the username if the SMTP server requires authentication
Password	string	enter the password if the SMTP server requires authentication
Password Confirmation	string	enter the same password again for confirmation

Click the **Send Test Mail** button to verify that the configured email settings are working. If the test email fails, double-check

the destination email address by clicking the Change E-mail button for the root account in [Account](#) → [Users](#) → [View Users](#). Test mail cannot be sent unless the root email address has been set.

Configuring email for TLS/SSL email providers is described in [Are you having trouble getting FreeNAS to email you in Gmail?](#) (<https://forums.freenas.org/index.php?threads/are-you-having-trouble-getting-freenas-to-email-you-in-gmail.22517/>).

5.6 System Dataset

[System](#) → [System Dataset](#), shown in [Figure 5.11](#), is used to select the pool which will contain the persistent system dataset. The system dataset stores debugging core files and Samba4 metadata such as the user/group cache and share level permissions. If the FreeNAS® system is configured to be a Domain Controller, all of the domain controller state is stored there as well, including domain controller users and groups.

Note: When the system dataset is moved, a new dataset is created and set active. The old dataset is intentionally not deleted by the system because the move might be transient or the information in the old dataset might be useful for later recovery.

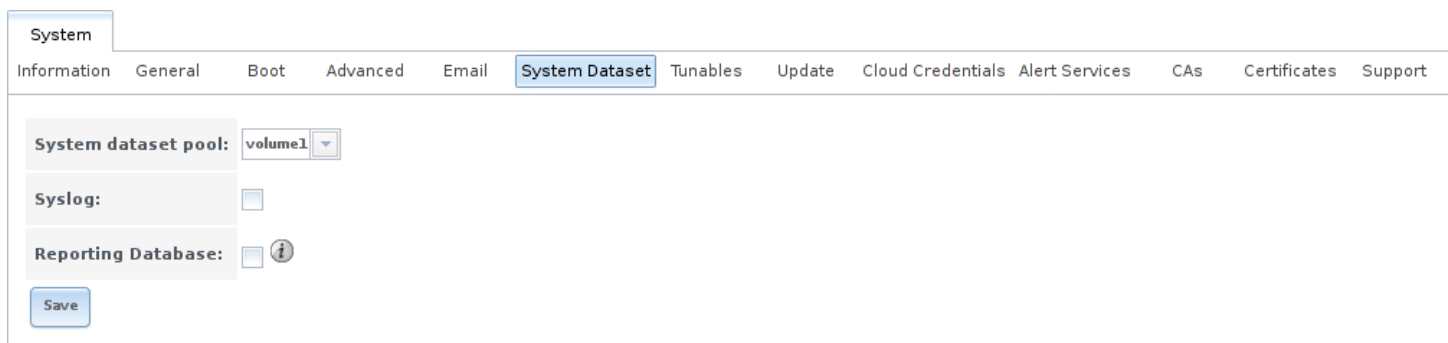


Fig. 5.11: System Dataset Screen

Note: Encrypted volumes are not displayed in the System dataset pool drop-down menu.

The system dataset can optionally be configured to also store the system log and [Reporting](#) (page 280) information. If there are lots of log entries or reporting information, moving these to the system dataset will prevent `/var/` on the device holding the operating system from filling up as `/var/` has limited space.

Use the drop-down menu to select the ZFS volume (pool) to contain the system dataset. Whenever the location of the system dataset is changed, a pop-up warning indicates that the SMB service must be restarted, causing a temporary outage of any active SMB connections.

To store the system log on the system dataset, check the Syslog box.

To store the reporting information on the system dataset, check the Reporting Database box. When this box is not checked, a RAM disk is created to prevent reporting information from filling up `/var`.

Click the Save button to save changes.

If the pool storing the system dataset is changed at a later time, FreeNAS® migrates the existing data in the system dataset to the new location.

Note: Depending on configuration, the system dataset can occupy a large amount of space and receive frequent writes. Do not put the system dataset on a flash drive or other media with limited space or write life.

5.7 Tunables

System → Tunables can be used to manage the following:

- 1. FreeBSD sysctls: a `sysctl(8)` (<http://www.freebsd.org/cgi/man.cgi?query=sysctl>) makes changes to the FreeBSD kernel running on a FreeNAS® system and can be used to tune the system.
- 2. FreeBSD loaders: a loader is only loaded when a FreeBSD-based system boots and can be used to pass a parameter to the kernel or to load an additional kernel module such as a FreeBSD hardware driver.
- 3. FreeBSD rc.conf options: `rc.conf(5)` (<https://www.freebsd.org/cgi/man.cgi?query=rc.conf&manpath=FreeBSD+11.0-RELEASE>) is used to pass system configuration options to the system startup scripts as the system boots. Since FreeNAS® has been optimized for storage, not all of the services mentioned in `rc.conf(5)` are available for configuration. Note that in FreeNAS®, customized `rc.conf` options are stored in `/tmp/rc.conf.freenas`.

Warning: Adding a `sysctl`, loader, or `rc.conf` option is an advanced feature. A `sysctl` immediately affects the kernel running the FreeNAS® system and a loader could adversely affect the ability of the FreeNAS® system to successfully boot. Do not create a tunable on a production system unless you understand and have tested the ramifications of that change.

Since `sysctl`, loader, and `rc.conf` values are specific to the kernel parameter to be tuned, the driver to be loaded, or the service to configure, descriptions and suggested values can be found in the man page for the specific driver and in many sections of the [FreeBSD Handbook](http://www.freebsd.org/handbook) (<http://www.freebsd.org/handbook>).

To add a loader, `sysctl`, or `rc.conf` option, go to System → Tunables → Add Tunable, to access the screen shown in [Figure 5.12](#).

Add Tunable

Variable:

Value:

Type:

Loader

Comment:

Enabled:

☒

OK

Cancel

Fig. 5.12: Adding a Tunable

Table 5.5 summarizes the options when adding a tunable.

Table 5.5: Adding a Tunable

Setting	Value	Description
Variable	string	typically the name of the <code>sysctl</code> or driver to load, as indicated by its man page
Continued on next page		

Table 5.5 – continued from previous page

Setting	Value	Description
Value	integer or string	value to associate with Variable; typically this is set to YES to enable the sysctl or driver specified by the "Variable"
Type	drop-down menu	choices are Loader, rc.conf, or Sysctl
Comment	string	optional, but a useful reminder for the reason behind adding this tunable
Enabled	checkbox	uncheck if you would like to disable the tunable without deleting it

Note: As soon as a Sysctl is added or edited, the running kernel changes that variable to the value specified. However, when a Loader or rc.conf value is changed, it does not take effect until the system is rebooted. Regardless of the type of tunable, changes persist at each boot and across upgrades unless the tunable is deleted or its Enabled checkbox is unchecked.

Any added tunables are listed in **System** → **Tunables**. To change the value of an existing tunable, click its Edit button. To remove a tunable, click its Delete button.

Restarting the FreeNAS® system after making sysctl changes is recommended. Some sysctls only take effect at system startup, and restarting the system guarantees that the setting values correspond with what is being used by the running system.

The GUI does not display the sysctls that are pre-set when FreeNAS® is installed. FreeNAS® 11.1 ships with the following sysctls set:

```
kern.metadelay=3
kern.dirdelay=4
kern.filedelay=5
kern.coredump=1
kern.sugid_coredump=1
vfs.timestamp_precision=3
net.link.lagg.lacp.default_strict_mode=0
vfs.zfs.min_auto_ashift=12
```

Do not add or edit these default sysctls as doing so may render the system unusable.

The GUI does not display the loaders that are pre-set when FreeNAS® is installed. FreeNAS® 11.1 ships with these loaders set:

```
autoboot_delay="2"
loader_logo="freenas"
loader_menu_title="Welcome to FreeNAS"
loader_brand="freenas-brand"
loader_version=" "
kern.cam.boot_delay="30000"
debug.debugger_on_panic=1
debug.ddb.textdump.pending=1
hw.hptrr.attach_generic=0
vfs.mountroot.timeout="30"
ispfw_load="YES"
freenas_sysctl_load="YES"
hint.isp.0.role=2
hint.isp.1.role=2
hint.isp.2.role=2
hint.isp.3.role=2
hint.isp.0.topology="nport-only"
hint.isp.1.topology="nport-only"
hint.isp.2.topology="nport-only"
hint.isp.3.topology="nport-only"
module_path="/boot/kernel;/boot/modules;/usr/local/modules"
net.inet6.ip6.auto_linklocal="0"
vfs.zfs.vol.mode=2
```

```
kern.geom.label.disk_ident.enable="0"
hint.ahciem.0.disabled="1"
hint.ahciem.1.disabled="1"
kern.msgbufsize="524288"
hw.mfi.mrsas_enable="1"
hw.usb.no_shutdown_wait=1
hw.cxgbe.toecaps_allowed=0
hw.cxgbe.rdmacaps_allowed=0
hw.cxgbe.iscsicaps_allowed=0
vfs.nfsd.fha.write=0
vfs.nfsd.fha.max_nfsds_per_fh=32
```

Do not add or edit the default tunables as doing so might make the system unusable.

The ZFS version used in 11.1 deprecates these tunables:

```
vfs.zfs.write_limit_override
vfs.zfs.write_limit_inflated
vfs.zfs.write_limit_max
vfs.zfs.write_limit_min
vfs.zfs.write_limit_shift
vfs.zfs.no_write_throttle
```

After upgrading from an earlier version of FreeNAS®, these tunables are automatically deleted. Please do not manually add them back.

5.8 Update

FreeNAS® has an integrated update system to make it easy to keep up to date.

5.8.1 Preparing for Updates

It is best to perform updates at times the FreeNAS® system is idle, with no clients connected and no scrubs or other disk activity going on. A reboot is required after most updates, so they are often planned for scheduled maintenance times to avoid disrupting user activities.

The update process will not proceed unless there is enough free space in the boot pool for the new update files. If a space warning is shown, use [Boot](#) (page 59) to remove unneeded boot environments.

5.8.2 Updates and Trains

FreeNAS® is updated with signed update files. This provides flexibility in deciding when to upgrade the system with patches, new drivers, or new features. It also allows “test driving” an upcoming release. Combined with boot environments, new features or system patches can be tested while still being able to revert to a previous version of the operating system (see [If Something Goes Wrong](#) (page 26)). Digital signing of update files eliminates the need to manually download both an upgrade file and the associated checksum to verify file integrity.

[Figure 5.13](#) shows an example of the `System → Update` screen.

Fig. 5.13: Update Options

By default, the system automatically checks for updates and issues an alert when a new update becomes available. The automatic check can be disabled by unchecking Automatically check for updates.

This screen lists the URL of the official update server in case that information is needed in a network with outbound firewall restrictions. It also shows which software branch, or train, is being tracked for updates.

Several trains are available for updates.

Caution: Only Production trains are recommended for regular usage. Other trains are made available for pre-production testing and updates to legacy versions. Pre-production testing trains are provided only to permit testing of new versions before switching to a new branch. Before using a non-production train, be prepared to experience bugs or problems. Testers are encouraged to submit bug reports at <https://redmine.ixsystems.com/projects/freenas/issues>.

These trains are available:

For Production Use

- FreeNAS-11-STABLE (Recommended)

After testing, new fixes and features are added to this train. Selecting this train and applying any pending updates is recommended.

For Pre-Production Testing

- FreeNAS-11-Nightlies: Do not use this train in production. It is the experimental branch for future versions and is meant only for testers and developers.
- FreeNAS-11-Nightlies-SDK: Do not use this train in production. This train is meant only for developers. It is similar to FreeNAS-11-Nightlies but with extra development and debugging utilities added.

- FreeNAS-HEAD-Nightlies: Do not use this train in production. This train is meant only for developers and contains the source that will eventually become FreeNAS® version 12.

Legacy Versions

- FreeNAS-9.10-STABLE

Maintenance-only updates to the older version of FreeNAS®. Upgrading to FreeNAS-11-STABLE is recommended to ensure that the system receives bug fixes and new features.

To change the train, use the drop-down menu to make a different selection.

Note: The train selector does not allow downgrades. For example, the STABLE train cannot be selected while booted into a Nightly boot environment, or a 9.10 train cannot be selected while booted into a 11 boot environment. To go back to an earlier version after testing or running a more recent version, reboot and select a boot environment for that earlier version. This screen can then be used to check for updates that train.

This screen also shows the URL of the official update server. That information can be required when using a network with outbound firewall restrictions.

The Verify Install button verifies that the operating system files in the current installation do not have any inconsistencies. If any problems are found, a pop-up menu lists the files with checksum mismatches or permission errors.

5.8.3 Checking for Updates

Checking for updates by making sure the desired train is selected and clicking the Check Now button. Any available updates are listed. In the example shown in [Figure 5.14](#), the numbers which begin with a # represent the issue number from [the issue tracker](https://redmine.ixsystems.com/projects/freenas/issues) (<https://redmine.ixsystems.com/projects/freenas/issues>). Numbers which do not begin with a # represent a git commit. Click the ChangeLog link to open the log of changes in a web browser. Click the ReleaseNotes link to open the Release Notes in the browser.

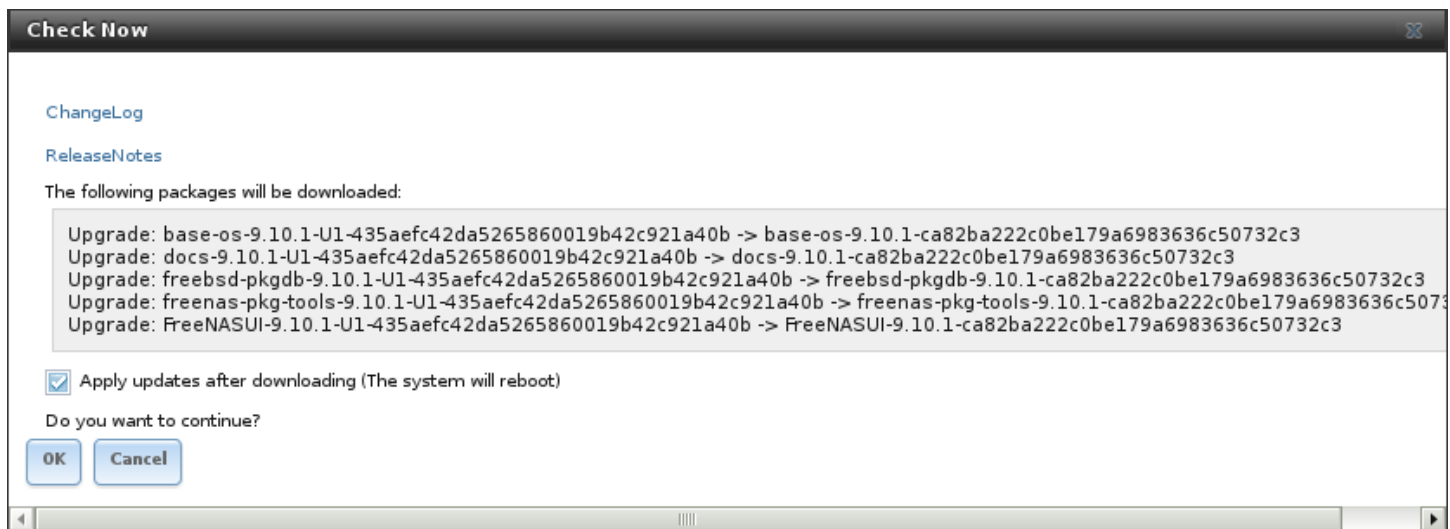


Fig. 5.14: Reviewing Updates

5.8.4 Applying Updates

Make sure the system is in a low-usage state as described above in [Preparing for Updates](#) (page 70).

Click the OK button to download and apply the updates. Be aware that some updates automatically reboot the system after they are applied.

Warning: Each update creates a boot environment. If the update process needs more space, it attempts to remove old boot environments. Boot environments marked with the Keep attribute as shown in [Boot](#) (page 59) will not be removed. If space for a new boot environment is not available, the upgrade fails. Space on the boot device can be manually freed using `System → Boot`. Review the boot environments and remove the Keep attribute or delete any boot environments that are no longer needed.

Updates can also be downloaded and applied later. To do so, uncheck the Apply updates after downloading box before pressing OK. In this case, this screen closes after updates are downloaded. Downloaded updates are listed in the Pending Updates section of the screen shown in [Figure 5.13](#). When ready to apply the previously downloaded updates, click the Apply Pending Updates button. Remember that the system might reboot after the updates are applied.

Warning: After updates have completed, reboot the system. Configuration changes made after an update but before that final reboot will not be saved.

5.8.5 Manual Updates

Updates can be manually downloaded as a file. These updates are then applied with the Manual Update button. After obtaining the update file, click Manual Update and choose a location to temporarily store the file on the FreeNAS® system. Use the file browser to locate the update file, then click Apply Update to apply it.


Manual update files can be identified by their filenames, which end in `-manual-update-unsigned.tar`.

Manual updates cannot be used to upgrade from older major versions.

5.9 Cloud Credentials

FreeNAS® can use cloud services for features like [Cloud Sync](#) (page 84). The credentials to provide secure connections with cloud services are entered here. Amazon S3, Azure Blob Storage, Backblaze B2, and Google Cloud Storage are supported.

Select `System → Cloud Credentials → Add Cloud Credential` to display the dialog shown in [Figure 5.15](#).



The screenshot shows a dialog box titled "Add Cloud Credential". It has a close button (X) in the top right corner. Inside the dialog, there are two main input areas. The first is labeled "Account Name:" followed by a text input field. The second is labeled "Provider:" followed by a dropdown menu showing a dashed line and a downward arrow. At the bottom of the dialog, there are two buttons: "OK" and "Cancel".

Fig. 5.15: Adding Cloud Credentials

Enter a descriptive name for the cloud credential in the Account Name field, then select a provider. The remaining options vary by provider, and are shown in [Table 5.6](#).

Table 5.6: Cloud Credential Options

Provider	Setting	Description
Amazon S3	Access Key, Secret Key	paste the Amazon account access key and secret key in the fields
Azure Blob Storage	Account Name, Account Key	enter the Azure Blob Storage account name and key in the fields
Backblaze B2	Account ID, Application Key	enter the Backblaze account ID and paste the application in the fields
Google Cloud Storage	JSON Server Account Key	browse to the location of the saved Google Cloud Storage key and select it

Additional fields are displayed after Provider is selected. For Amazon S3, Access Key and Secret Key are shown. These values can be found on the Amazon AWS website by clicking on the account name, then My Security Credentials and Access Keys (Access Key ID and Secret Access Key). Copy the Access Key value to the FreeNAS® Cloud Credential Access Key field, then enter the Secret Key value saved when the key pair was created. If the Secret Key value is not known, a new key pair can be created on the same Amazon screen.

5.10 Alert Services

FreeNAS® can use a number of methods to notify the administrator of system events that require attention. These events are system [Alerts](#) (page 297) marked WARN or CRITICAL.

Currently available alert services:

- [AWS-SNS](https://aws.amazon.com/sns/) (https://aws.amazon.com/sns/)
- [Hipchat](https://www.hipchat.com/) (https://www.hipchat.com/)
- [InfluxDB](https://www.influxdata.com/) (https://www.influxdata.com/)
- [Slack](https://slack.com/) (https://slack.com/)
- [Mattermost](https://about.mattermost.com/) (https://about.mattermost.com/)
- [OpsGenie](https://www.opsgenie.com/) (https://www.opsgenie.com/)
- [PagerDuty](https://www.pagerduty.com/) (https://www.pagerduty.com/)
- [VictorOps](https://victorops.com/) (https://victorops.com/)

Warning: These alert services might use a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before using their alert service. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Alert Services feature.

Select **System** → **Alert Services** to go to the Alert Services screen. Click **Add Service** to display the dialog shown in [Figure 5.16](#).

Fig. 5.16: Add Alert Service

The Service Name drop-down menu is used to pick a specific alert service. The fields shown in the rest of the dialog change to those required by that service. Enter the required information, check the Enabled checkbox, then click OK to save the settings.

System alerts marked WARN or CRITICAL are sent to each alert service that has been configured and enabled.

Alert services can be deleted from this list by clicking them and then clicking the Delete button at the bottom of the window. To disable an alert service temporarily, click Edit and remove the checkmark from the Enabled checkbox.

Note: To send a test alert, highlight an alert entry, click Edit, and click the Send Test Alert button.

5.10.1 How it Works

A nas-health service is registered with Consul. This service runs `/usr/local/etc/consul-checks/freenas_health.sh` periodically, currently every two minutes. If an alert marked WARNING or CRITICAL is found, the nas-health service is marked as “unhealthy”, triggering `consul-alerts` to notify configured alert services.

5.11 CAs

FreeNAS® can act as a Certificate Authority (CA). When encrypting SSL or TLS connections to the FreeNAS® system, either import an existing certificate, or create a CA on the FreeNAS® system, then create a certificate. This certificate will appear in the drop-down menus for services that support SSL or TLS.

For secure LDAP, the public key of an existing CA can be imported with Import CA, or a new CA created on the FreeNAS® system and used on the LDAP server also.

Figure 5.17 shows the screen after clicking `System → CAs`.




Fig. 5.17: Initial CA Screen

If your organization already has a CA, the CA's certificate and key can be imported. Click the Import CA button to open the configuration screen shown in Figure 5.18. The configurable options are summarized in Table 5.7.

The screenshot shows the 'Import CA' configuration window. It contains several input fields: 'Identifier' (with a tooltip: 'Internal identifier of the certificate. Only alphanumeric, "_" and "-" are allowed.'), 'Certificate' (highlighted with a red border), 'Private Key', 'Passphrase', 'Confirm Passphrase', and 'Serial'. Each field has an information icon (i) to its right. At the bottom are 'OK' and 'Cancel' buttons.

Fig. 5.18: Importing a CA

Table 5.7: Importing a CA Options

Setting	Value	Description
Identifier	string	mandatory; enter a descriptive name for the CA using only alphanumeric, underscore (_), and dash (-) characters
Certificate	string	mandatory; paste in the certificate for the CA
Private Key	string	if there is a private key associated with the Certificate, paste it here
Passphrase	string	if the Private Key is protected by a passphrase, enter it here and repeat it in the "Confirm Passphrase" field
Serial	string	mandatory; enter the serial number for the certificate

To instead create a new CA, first decide if it will be the only CA which will sign certificates for internal use or if the CA will be part of a [certificate chain](https://en.wikipedia.org/wiki/Root_certificate) (https://en.wikipedia.org/wiki/Root_certificate).

To create a CA for internal use only, click the Create Internal CA button which will open the screen shown in Figure 5.19.

Fig. 5.19: Creating an Internal CA

The configurable options are described in [Table 5.8](#). When completing the fields for the certificate authority, supply the information for your organization.

Table 5.8: Internal CA Options

Setting	Value	Description
Identifier	string	required; enter a descriptive name for the CA using only alphanumeric, underscore (<code>_</code>), and dash (<code>-</code>) characters
Key Length	drop-down menu	for security reasons, a minimum of 2048 is recommended
Digest Algorithm	drop-down menu	the default is acceptable unless your organization requires a different algorithm
Lifetime	integer	in days
Country	drop-down menu	select the country for the organization
State	string	required; enter the state or province of the organization
Locality	string	required; enter the location of the organization
Organization	string	required; enter the name of the company or organization
Email Address	string	required; enter the email address for the person responsible for the CA

Continued on next page

Table 5.8 – continued from previous page

Setting	Value	Description
Common Name	string	required; enter the fully-qualified hostname (FQDN) of the system; the Common Name must be unique within a certificate chain
Subject Alternate Names	string	newer browsers look for the values in this field to match the domain to the certificate; use a space to separate domain names

To instead create an intermediate CA which is part of a certificate chain, click the Create Intermediate CA button. This screen adds one more option to the screen shown in [Figure 5.19](#):

- **Signing Certificate Authority:** this drop-down menu is used to specify the root CA in the certificate chain. This CA must first be imported or created.

Any CAs that you import or create will be added as entries in *System* → *CAs*. The columns in this screen indicate the name of the CA, whether it is an internal CA, whether the issuer is self-signed, the number of certificates that have been issued by the CA, the distinguished name of the CA, the date and time the CA was created, and the date and time the CA expires.

Clicking the entry for a CA causes these buttons to become available:

- **Sign CSR:** used to sign internal Certificate Signing Requests created using *System* → *Certificates* → *Create Certificate Signing Request*.
- **Export Certificate:** prompts to browse to the location to save a copy of the CA's X.509 certificate on the computer being used to access the FreeNAS® system.
- **Export Private Key:** prompts to browse to the location to save a copy of the CA's private key on the computer being used to access the FreeNAS® system. This option only appears if the CA has a private key.
- **Delete:** prompts for confirmation before deleting the CA.

5.12 Certificates

FreeNAS® can import existing certificates, create new certificates, and issue certificate signing requests so that created certificates can be signed by the CA which was previously imported or created in [CAs](#) (page 75).

[Figure 5.20](#) shows the initial screen after clicking *System* → *Certificates*.




Fig. 5.20: Initial Certificates Screen

To import an existing certificate, click the Import Certificate button to open the configuration screen shown in [Figure 5.21](#). When importing a certificate chain, paste the primary certificate, followed by any intermediate certificates, followed by the root CA certificate.

The configurable options are summarized in [Table 5.9](#).

Import Certificate

Identifier:

Internal identifier of the certificate. Only alphanumeric, "_" and "-" are allowed.

Certificate:

i

Private Key:

i

Passphrase:

i

Confirm Passphrase:

OK

Cancel

Fig. 5.21: Importing a Certificate

Table 5.9: Certificate Import Options

Setting	Value	Description
Identifier	string	required; enter a descriptive name for the certificate using only alphanumeric, underscore (_), and dash (-) characters
Certificate	string	required; paste the contents of the certificate
Private Key	string	required; paste the private key associated with the certificate
Passphrase	string	if the private key is protected by a passphrase, enter it here and repeat it in the Confirm Passphrase field

To instead create a new self-signed certificate, click the Create Internal Certificate button to see the screen shown in [Figure 5.22](#). The configurable options are summarized in [Table 5.10](#). When completing the fields for the certificate authority, use the information for your organization. Since this is a self-signed certificate, use the CA that was imported or created with [CAs](#) (page 75) as the signing authority.

Create Internal Certificate

Signing Certificate Authority:

Identifier:

Key length:

2048

Digest Algorithm:

SHA256

Lifetime:

3,650

Country:

United States

State:

Locality:

Organization:

Email Address:

Common Name:

Subject Alternate Names:

OK

Cancel

Fig. 5.22: Creating a New Certificate

Table 5.10: Certificate Creation Options

Setting	Value	Description
Signing Certificate Authority	drop-down menu	required; select the CA which was previously imported or created using CAs (page 75)
Identifier	string	required; enter a descriptive name for the certificate using only alphanumeric, underscore (<code>_</code>), and dash (<code>-</code>) characters
Key Length	drop-down menu	for security reasons, a minimum of 2048 is recommended
Digest Algorithm	drop-down menu	the default is acceptable unless your organization requires a different algorithm
Lifetime	integer	in days
Country	drop-down menu	select the country for the organization
State	string	required; enter the state or province for the organization
Locality	string	required; enter the location for the organization
Organization	string	required; enter the name of the company or organization

Continued on next page

Table 5.10 – continued from previous page

Setting	Value	Description
Email Address	string	required; enter the email address for the person responsible for the CA
Common Name	string	required; enter the fully-qualified hostname (FQDN) of the system; the Common Name must be unique within a certificate chain
Subject Alternate Names	string	newer browsers look for the values in this field to match the domain to the certificate; use a space to separate domain names

If you need to use a certificate that is signed by an external CA, such as Verisign, instead create a certificate signing request. To do so, click the Create Certificate Signing Request button. A screen like the one in [Figure 5.22](#) opens, but without the Signing Certificate Authority field.

Certificates that are imported, self-signed, or for which a certificate signing request is created are added as entries to `System` → `Certificates`. In the example shown in [Figure 5.23](#), a self-signed certificate and a certificate signing request have been created for the fictional organization My Company. The self-signed certificate was issued by the internal CA named My Company and the administrator has not yet sent the certificate signing request to Verisign so that it can be signed. Once that certificate is signed and returned by the external CA, it should be imported using the Import Certificate button so that is available as a configurable option for encrypting connections.

Name	Issuer	Distinguished Name	From	Until
FreeNAS_Internal_Certificate	FreeNAS_Internal_CA	/C=US/ST=CA/L=Silicon Valley/O=iXsystems /CN= /emailAddress=	Thu Apr 27 18:40:59 2017	Sun Apr 25 18:40:59 2027

Fig. 5.23: Managing Certificates

Clicking an entry activates these configuration buttons:

- View: use this option to view or edit the contents of an existing certificate. These fields can be edited: Identifier (name), Certificate, and Private Key.
- Export Certificate saves a copy of the certificate or certificate signing request to the system being used to access the FreeNAS® system. For a certificate signing request, send the exported certificate to the external signing authority so that it can be signed.
- Export Private Key saves a copy of the private key associated with the certificate or certificate signing request to the system being used to access the FreeNAS® system.
- Delete is used to delete a certificate or certificate signing request.

5.13 Support

The FreeNAS® Support tab, shown in [Figure 5.24](#), provides a built-in ticketing system for generating bug reports and feature requests.

System

Information General Boot Advanced Email System Dataset Tunables Update Cloud Credentials Alert Services CAs Certificates **Support**

Before filing a bug report or feature request, search <http://bugs.freenas.org> to ensure the issue has not already been reported. If it has, add a comment to the existing issue instead of creating a new one. For enterprise-grade storage solutions and support, please visit <http://www.ixsystems.com/storage/>.

If you do not have an account, please [register](#).

Username

Password

Type

Category

Attach Debug Info ☒

Subject

Description

Attachments

No file selected.

Fig. 5.24: Support Tab

This screen provides a built-in interface to the FreeNAS® issue tracker located at <https://redmine.ixsystems.com/projects/freenas/issues>. If you have not yet used the FreeNAS® bug tracker, you must first go to that website, click the Register link, fill out the form, and reply to the registration email. This will create a username and password which can be used to create bug reports and receive notifications as the reports are actioned.

Before creating a bug report or feature request, ensure that an existing report does not already exist at <https://redmine.ixsystems.com/projects/freenas/issues>. If you find a similar issue that is not yet marked as closed or resolved, add a comment to that issue if you have new information to provide that can assist in resolving the issue. If you find a similar issue that is marked as closed or resolved, you can create a new issue and refer to the earlier issue number.

Note: If you are not updated to the latest version of STABLE, do that first to see if it resolves your issue.

To generate a report using the built-in Support screen, complete the following fields:

- Username: enter the login name created when registering at <https://redmine.ixsystems.com/projects/freenas/issues>.
- Password: enter the password associated with the registered login name.
- Type: select Bug when reporting an issue or Feature when requesting a new feature.
- Category: this drop-down menu is empty a registered “Username” and “Password” are entered. An error message is displayed if either value is incorrect. After the Username and Password are validated, possible categories are populated to the drop-down menu. Select the one that best describes the bug or feature being reported.
- Attach Debug Info: it is recommended to leave this box checked so that an overview of the system’s hardware, build string, and configuration is automatically generated and included with the ticket.
- Subject: enter a descriptive title for the ticket. A good Subject makes it easy for you and other users to find similar reports.

- Description: enter a one- to three-paragraph summary of the issue that describes the problem, and if applicable, what steps can be taken to reproduce it.
- Attachments: this is the only optional field. It is useful for including configuration files or screenshots of any errors or tracebacks.

Once you have finished completing the fields, click the Submit button to automatically generate and upload the report to <https://redmine.ixsystems.com/projects/freenas/issues>. A pop-up menu provides a clickable URL so to view status or add additional information to the report.

TASKS

The Tasks section of the administrative GUI is used to configure repetitive tasks:

- [Cloud Sync](#) (page 84) schedules data synchronization to cloud providers
- [Cron Jobs](#) (page 89) schedules a command or script to automatically execute at a specified time
- [Init/Shutdown Scripts](#) (page 91) configures a command or script to automatically execute during system startup or shutdown
- [Rsync Tasks](#) (page 92) schedules data synchronization to another system
- [S.M.A.R.T. Tests](#) (page 99) schedules disk tests

Each of these tasks is described in more detail in this section.

Note: By default, [Scrubs](#) (page 152) are run once a month by an automatically-created task. [S.M.A.R.T. Tests](#) (page 99) and [Periodic Snapshot Tasks](#) (page 140) must be set up manually.

6.1 Cloud Sync

Files or directories can be synchronized to remote cloud storage providers with the Cloud Sync feature.

Warning: This Cloud Sync task might go to a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before creating any Cloud Sync task. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Cloud Sync feature.

Selecting [Tasks](#) → [Cloud Sync](#) shows the screen in [Figure 6.1](#). This screen shows a single cloud sync called backup-acctg that “pushes” a file to cloud storage. The last run finished with a status of SUCCESS.

Existing cloud syncs can be run manually, edited, or deleted with the buttons that appear when a single cloud sync line is selected by clicking with the mouse.

Tasks										
Cloud Sync										
Cron Jobs Init/Shutdown Scripts Rsync Tasks S.M.A.R.T. Tests										
Add Cloud Sync										
Description	Direction	Path	Status	Minute	Hour	Day of month	Month	Day of week	Credential	Enabled
backup-acctg	PUSH	/mnt/volume1 /smb-storage /accounting-backup.bin	SUCCESS	00	Every hour	Everyday	Every month	Everyday	S3 Storage	true
<div> <div>Edit</div> <div>Delete</div> <div>Run Now</div> </div>										

Fig. 6.1: Cloud Sync Status

[Cloud Credentials](#) (page 73) must be defined before a cloud sync is created. One set of credentials can be used for more than one cloud sync. For example, a single set of credentials for Amazon S3 can be used for separate cloud syncs that push different sets of files or directories.

A cloud storage area must also exist. With Amazon S3, these are called buckets. The bucket must be created before a sync task can be created.

After the credentials and receiving bucket have been created, a cloud sync task is created with `Tasks → Cloud Sync → Add Cloud Sync`. The Add Cloud Sync dialog is shown in [Figure 6.2](#).

Add Cloud Sync

Warning: This Cloud Sync task might go to a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before creating any Cloud Sync task. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Cloud Sync feature.

Description:

Direction:

Push

Provider:

Credential

Path:

Browse

Transfer Mode:

Sync

Minute:

Every N minute

Each selected minute

00

01

02

03

04

05

06

07

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

Fig. 6.2: Adding a Cloud Sync

Table 6.1 shows the configuration options for Cloud Syncs.

Table 6.1: Cloud Sync Options

Setting	Value Type	Description
Description	string	a descriptive name for this Cloud Sync
Direction	string	Push to send data to cloud storage, or Pull to pull data from the cloud storage
Provider	drop-down menu	select the cloud storage provider; the list of providers is defined by Cloud Credentials (page 73)
Amazon S3 Buckets	drop-down menu	only appears when an S3 credential is the Provider; select the bucket to use

Continued on next page

Table 6.1 – continued from previous page

Setting	Value Type	Description
Folder	string	only appears when an S3 credential is the Provider; input the name of the folder to sync to
Server Side Encryption	drop-down menu	only appears when an S3 credential is the Provider; choices are None (no encryption) or AES-256 (encrypted)
Path	browse button	select the directories or files to be sent for Push syncs or the destinations for Pull syncs
Transfer Mode	drop-down menu	Sync (default): make files on destination system identical to those on the source; files removed from the source are removed from the destination (like <code>rsync --delete</code>) Copy: copy files from the source to the destination, skipping files that are identical (like <code>rsync</code>) Move: copy files from the source to the destination, deleting files from the source after the copy (like <code>mv</code>)
Minute	slider or minute selections	select Every N minutes and use the slider to choose a value, or select Each selected minute and choose specific minutes
Hour	slider or hour selections	select Every N hours and use the slider to choose a value, or select Each selected hour and choose specific hours
Day of month	slider or day of month selections	select Every N days of month and use the slider to choose a value, or select Each selected day of month and choose specific days
Month	checkboxes	months when the Cloud Sync runs
Day of week	checkboxes	days of the week when the Cloud Sync runs
Enabled	checkbox	uncheck to temporarily disable this Cloud Sync

Take care when choosing a Direction. Most of the time, Push will be used to send data to the cloud storage. Pull retrieves data from cloud storage, but be careful: files retrieved from cloud storage will overwrite local files with the same names in the destination directory.

Provider is the name of the cloud storage provider. These providers are defined by entering credentials in [Cloud Credentials](#) (page 73).

After the Provider is chosen, a list of available cloud storage areas from that provider is shown. With Amazon AWS, this is a drop-down with names of existing buckets. Choose a bucket, and a folder inside that bucket if desired.

Path is the path to the directories or files on the FreeNAS® system. On Push jobs, this is the source location for files sent to cloud storage. On Pull jobs, the Path is where the retrieved files are written. Again, be cautious about the destination of Pull jobs to avoid overwriting existing files.

The Minute, Hour, Days of month, `gui-label:Months`, and Days of week fields permit creating a flexible schedule of when the cloud synchronization takes place.


Finally, the Enabled field makes it possible temporarily disable a cloud sync job without deleting it.

6.1.1 Cloud Sync Example

This example shows a Push cloud sync which writes an accounting department backup file from the FreeNAS® system to Amazon S3 storage.

Before the new cloud sync was added, a bucket called `cloudsync-bucket` was created with the Amazon S3 web console for storing data from the FreeNAS® system.

`System → Cloud Credentials → Add Cloud Credential` is used to enter the credentials for storage on an Amazon AWS account. The credential is given the name `S3 Storage`, as shown in [Figure 6.3](#):



The screenshot shows a dialog box titled "Add Cloud Credential". It has a close button in the top right corner. The dialog contains four input fields, each with a label on the left and a text input on the right. The first field is "Account Name" with the value "S3 Storage". The second field is "Provider" with a dropdown menu showing "Amazon S3". The third field is "Access Key" with the value "XYZZYXSQUAWKASQUEEPS". The fourth field is "Secret Key" with the value "VutRWwPQEos+TetQEWE5si". At the bottom of the dialog are two buttons: "OK" and "Cancel".

Fig. 6.3: Example: Adding Cloud Credentials

The local data to be sent to the cloud is a single file called `accounting-backup.bin` on the `smb-storage` dataset. A cloud sync job is created with `Tasks → Cloud Sync → Add Cloud Sync`. The Description is set to `backup-acctg` to describe the job. This data is being sent to cloud storage, so this is a Push. The provider comes from the cloud credentials defined in the previous step, and the destination bucket `cloudsync-bucket` has been chosen.

The Path to the data file is selected.

The remaining fields are for setting a schedule. The default is to send the data to cloud storage once an hour, every day. The options provide great versatility in configuring when a cloud sync runs, anywhere from once a minute to once a year.

The Enabled field is checked by default, so this cloud sync will run at the next scheduled time.

The completed dialog is shown in [Figure 6.4](#):

Add Cloud Sync

Description: backup-acctg

Direction: Push

Provider: S3 Storage

Credential: cloudsync-bucket

Folder:

Path: /mnt/volume1/smb-storage/ Close

Minute: Every N minute | Each selected minute

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27	28	29
30	31	32	33	34	35	36	37	38	39
40	41	42	43	44	45	46	47	48	49
50	51	52	53	54	55	56	57	58	59

Hour: Every N hour | Each selected hour

1

Day of month: Every N day of month | Each selected day of month

1

Month:

- ☒ January
- ☒ February
- ☒ March
- ☒ April
- ☒ May
- ☒ June
- ☒ July
- ☒ August
- ☒ September
- ☒ October
- ☒ November
- ☒ December

Day of week:

- ☒ Monday
- ☒ Tuesday
- ☒ Wednesday
- ☒ Thursday
- ☒ Friday
- ☒ Saturday
- ☒ Sunday

Enabled: ☒

OK Cancel

Fig. 6.4: Example: Adding a Cloud Sync

6.2 Cron Jobs

cron(8) (<http://www.freebsd.org/cgi/man.cgi?query=cron>) is a daemon that runs a command or script on a regular schedule as a specified user.

Figure 6.5 shows the screen that opens after clicking **Tasks** → **Cron Jobs** → **Add Cron Job**.

Add Cron Job

User:

The user to run the command

Command:

Short description:

Minute:

Every N minute

Each selected minute

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27	28	29
30	31	32	33	34	35	36	37	38	39
40	41	42	43	44	45	46	47	48	49
50	51	52	53	54	55	56	57	58	59

Hour:

Every N hour

Each selected hour

1

Day of month:

Every N day of month

Each selected day of month

1

Month:

☒ January

Fig. 6.5: Creating a Cron Job

Table 6.2 lists the configurable options for a cron job.

Table 6.2: Cron Job Options

Setting	Value	Description
User	drop-down menu	make sure the selected user has permission to run the specified command or script
Command	string	the full path to the command or script to be run; if it is a script, test it at the command line first to make sure that it works as expected
Short description	string	optional
Minute	slider or minute selections	with the slider, the cron job occurs every N minutes; with minute selections, the cron job occurs at the highlighted minutes
Hour	slider or hour selections	with the slider, the cron job occurs every N hours; with hour selections, the cron job occurs at the highlighted hours
Day of month	slider or month selections	with the slider, cron job occurs every N days; with day selections, cron job occurs on the highlighted days each month
Month	checkboxes	cron job occurs on the selected months
Day of week	checkboxes	cron job occurs on the selected days
Redirect Stdout	checkbox	disables emailing standard output to the root user account
Redirect Stderr	checkbox	disables emailing errors to the root user account
Enabled	checkbox	uncheck disable the cron job without deleting it

Cron jobs are shown in View Cron Jobs. Highlight a cron job entry to display buttons to Edit, Delete, or Run Now.

Note: % symbols are automatically escaped and should not be prefixed with backslashes. For example, use `date '+%Y-%m-%d'` in a cron job to generate a filename based on the date.

6.3 Init/Shutdown Scripts

FreeNAS® provides the ability to schedule commands or scripts to run at system startup or shutdown.

Figure 6.6 shows the screen that opens after clicking **Tasks** → **Init/Shutdown Scripts** → **Add Init/Shutdown Script**. Table 6.3 summarizes the options.

Scheduled commands must be in the default path. The full path to the command can also be included in the entry. The path can be tested by typing `which commandname`. If the command is not found, it is not in the path.

When scheduling a script, make sure that the script is executable and has been fully tested to ensure it achieves the desired results.




Fig. 6.6: Add an Init/Shutdown Script

Table 6.3: Options When Adding an Init/Shutdown Script

Setting	Value	Description
Type	drop-down menu	select from Command (for an executable) or Script (for an executable script)
Command	string	if Command is selected, enter the command plus any desired options; if Script is selected, browse to the location of the script
When	drop-down menu	select when the command/script will run; choices are Pre Init (very early in boot process before filesystems are mounted), Post Init (towards end of boot process before FreeNAS services are started), or Shutdown
Enabled	checkbox	uncheck to disable the task

6.4 Rsync Tasks

Rsync (<http://www.samba.org/ftp/rsync/rsync.html>) is a utility that copies specified data from one system to another over a network. Once the initial data is copied, rsync reduces the amount of data sent over the network by sending only the differences between the source and destination files. Rsync can be used for backups, mirroring data on multiple systems, or for copying files between systems.

Rsync is most effective when only a relatively small amount of the data has changed. There are also [some limitations when using Rsync with Windows files](https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/) (<https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/>). For large amounts of data, data that has many changes from the previous copy, or Windows files, [Replication Tasks](#) (page 142) are often the faster and better solution.

Rsync is single-threaded, so gains little from multiple processor cores. To see whether rsync is currently running, use `pgrep rsync` from the [Shell](#) (page 290).

Both ends of an rsync connection must be configured:

- the rsync server: this system pulls (receives) the data. This system is referred to as PULL in the configuration examples.
- the rsync client: this system pushes (sends) the data. This system is referred to as PUSH in the configuration examples.

FreeNAS® can be configured as either an rsync client or an rsync server. The opposite end of the connection can be another FreeNAS® system or any other system running rsync. In FreeNAS® terminology, an rsync task defines which data is synchronized between the two systems. To synchronize data between two FreeNAS® systems, create the rsync task on the rsync client.

FreeNAS® supports two modes of rsync operation:

- rsync module mode: exports a directory tree, and its configured settings, as a symbolic name over an unencrypted connection. This mode requires that at least one module be defined on the rsync server. It can be defined in the FreeNAS® GUI under *Services* → *Rsync* → *Rsync Modules*. In other operating systems, the module is defined in *rsyncd.conf(5)* (<http://www.samba.org/ftp/rsync/rsyncd.conf.html>).
- rsync over SSH: synchronizes over an encrypted connection. Requires the configuration of SSH user and host public keys.

This section summarizes the options when creating an Rsync Task. It then provides a configuration example between two FreeNAS® systems for each mode of rsync operation.

Note: If there is a firewall between the two systems or if the other system has a built-in firewall, make sure that TCP port 873 is allowed.

Figure 6.7 shows the screen that appears after selecting *Tasks* → *Rsync Tasks* → *Add Rsync Task*. **Table 6.4** summarizes the options that can be configured when creating an rsync task.

Add Rsync Task

Path:

Browse

User:

i

Remote Host:

i

Rsync mode:

Rsync module

Remote Module Name:

i

Direction:

Push

i

Short description:

Minute:

Every N minute

Each selected minute

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27	28	29
30	31	32	33	34	35	36	37	38	39
40	41	42	43	44	45	46	47	48	49
50	51	52	53	54	55	56	57	58	59

i

Hour:

Every N hour

Each selected hour

1

i

Day of month:

Every N day of month

Each selected day of month

1

i

Fig. 6.7: Adding an Rsync Task

Table 6.4: Rsync Configuration Options

Setting	Value	Description
Path	browse button	browse to the path that to be copied; note that a path length greater than 255 characters will fail
User	drop-down menu	specified user must have permission to write to the specified directory on the remote system; due to a limitation in FreeBSD, the user name cannot contain spaces or exceed 17 characters
Remote Host	string	IP address or hostname of the remote system that will store the copy; use the format <code>username@remote_host</code> if the username differs on the remote host
Remote SSH Port	integer	only available in Rsync over SSH mode; allows specifying an SSH port other than the default of 22
Rsync mode	drop-down menu	choices are Rsync module or Rsync over SSH
Remote Module Name	string	only appears when using Rsync module mode, at least one module must be defined in <code>rsyncd.conf(5)</code> (http://www.samba.org/ftp/rsync/rsyncd.conf.html) of rsync server or in the Rsync Modules of another system
Remote Path	string	only appears when using Rsync over SSH mode, enter the existing path on the remote host to sync with (e.g. <code>/mnt/volume</code>); note that maximum path length is 255 characters
Validate Remote Path	checkbox	if the Remote Path does not yet exist, check this box to have it automatically created
Direction	drop-down menu	choices are Push or Pull; default is to push to a remote host
Short Description	string	optional
Minute	slider or minute selections	if use the slider, sync occurs every N minutes; if use minute selections, sync occurs at the highlighted minutes
Hour	slider or hour selections	if use the slider, sync occurs every N hours; if use hour selections, sync occurs at the highlighted hours
Day of month	slider or day selections	if use the slider, sync occurs every N days; if use day selections, sync occurs on the highlighted days
Month	checkboxes	task occurs on the selected months
Day of week	checkboxes	task occurs on the selected days of the week
Recursive	checkbox	if checked, copy will include all subdirectories of the specified volume
Times	checkbox	preserve modification times of files
Compress	checkbox	recommended on slow connections as reduces size of data to be transmitted
Archive	checkbox	equivalent to <code>-rlptgoD</code> (recursive, copy symlinks as symlinks, preserve permissions, preserve modification times, preserve group, preserve owner (super-user only), and preserve device files (super-user only) and special files)
Delete	checkbox	delete files in destination directory that do not exist in sending directory
Quiet	checkbox	suppresses informational messages from the remote server
Preserve permissions	checkbox	preserves original file permissions; useful if User is set to root
Preserve extended attributes	checkbox	both systems must support extended attributes (http://en.wikipedia.org/wiki/Xattr)
Delay Updates	checkbox	when checked, the temporary file from each updated file is saved to a holding directory until the end of the transfer, when all transferred files are renamed into place
Extra options	string	<code>rsync(1)</code> (http://rsync.samba.org/ftp/rsync/rsync.html) options not covered by the GUI; if the <code>*</code> character is used, it must be escaped with a backslash (<code>* .txt</code>) or used inside single quotes (<code>'* .txt'</code>)

Continued on next page

Table 6.4 – continued from previous page

Setting	Value	Description
Enabled	checkbox	uncheck to disable the rsync task without deleting it; note that when the Rsync (page 225) service is OFF, the rsync task will continue to look for the server unless this checkbox is unchecked

If the rsync server requires password authentication, enter `--password-file=/PATHTO/FILENAME` in the Extra options box, replacing `/PATHTO/FILENAME` with the appropriate path to the file containing the password.

Created rsync tasks will be listed in View Rsync Tasks. Highlight the entry for an rsync task to display buttons for Edit, Delete, or Run Now.

6.4.1 Rsync Module Mode

This configuration example configures rsync module mode between the two following FreeNAS® systems:

- 192.168.2.2 has existing data in `/mnt/local/images`. It will be the rsync client, meaning that an rsync task needs to be defined. It will be referred to as PUSH.
- 192.168.2.6 has an existing volume named `/mnt/remote`. It will be the rsync server, meaning that it will receive the contents of `/mnt/local/images`. An rsync module needs to be defined on this system and the rsyncd service needs to be started. It will be referred to as PULL.

On PUSH, an rsync task is defined in `Tasks → Rsync Tasks → Add Rsync Task`. In this example:

- the Path points to `/usr/local/images`, the directory to be copied
- the Remote Host points to 192.168.2.6, the IP address of the rsync server
- the Rsync Mode is Rsync module
- the Remote Module Name is backups; this will need to be defined on the rsync server
- the Direction is Push
- the rsync is scheduled to occur every 15 minutes
- the User is set to root so it has permission to write anywhere
- the Preserve Permissions checkbox is checked so that the original permissions are not overwritten by the root user

On PULL, an rsync module is defined in `Services → Rsync Modules → Add Rsync Module`. In this example:

- the Module Name is backups; this needs to match the setting on the rsync client
- the Path is `/mnt/remote`; a directory called `images` will be created to hold the contents of `/usr/local/images`
- the User is set to root so it has permission to write anywhere
- Hosts allow is set to 192.168.2.2, the IP address of the rsync client

Descriptions of the configurable options can be found in Rsync Modules.

To finish the configuration, start the rsync service on PULL in `Services → Control Services`. If the rsync is successful, the contents of `/mnt/local/images/` will be mirrored to `/mnt/remote/images/`.

6.4.2 Rsync over SSH Mode

SSH replication mode does not require the creation of an rsync module or for the rsync service to be running on the rsync server. It does require SSH to be configured before creating the rsync task:

- a public/private key pair for the rsync user account (typically root) must be generated on PUSH and the public key copied to the same user account on PULL
- to mitigate the risk of man-in-the-middle attacks, the public host key of PULL must be copied to PUSH

- the SSH service must be running on PULL

To create the public/private key pair for the rsync user account, open [Shell](#) (page 290) on PUSH and run `ssh-keygen`. This example generates an RSA type public/private key pair for the root user. When creating the key pair, do not enter the passphrase as the key is meant to be used for an automated task.

```
ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
f5:b0:06:d1:33:e4:95:cf:04:aa:bb:6e:a4:b7:2b:df root@freenas.local
The key's randomart image is:
+--[ RSA 2048 ]-----+
|      .o. oo      |
|      o+o. .      |
|      . =o +      |
|      + +   o      |
|      S o .        |
|      .o           |
|      o.           |
|      o oo         |
|      **oE         |
|-----|
|
|-----|
```

FreeNAS® supports RSA keys for SSH. When creating the key, use `-t rsa` to specify this type of key.

Note: If a different user account is used for the rsync task, use the `su -` command after mounting the filesystem but before generating the key. For example, if the rsync task is configured to use the `user1` user account, use this command to become that user:

```
su - user1
```

Next, view and copy the contents of the generated public key:

```
more .ssh/id_rsa.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAC1lBEXRgw1W8y8k+lXP1VR3xsmVSjtsoyIzV/PlQPo
SrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4h
dcD7Y5mvU3MAEeDClt02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmasti00qmDDcp/kO
xT+S6DFNDBY6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+l9RaEXMRuTyQgqJB/rsRcmJX5fApd
DmNfwrRSxLjDvUzfywnjFHlKk/+TQIT1gg1QQaj21PJD9pnDVF0AiJrWyWnR root@freenas.local
```

Go to PULL and paste (or append) the copied key into the SSH Public Key field of Account → Users → View Users → root → Modify User, or the username of the specified rsync user account. The paste for the above example is shown in [Figure 6.8](#). When pasting the key, ensure that it is pasted as one long line and, if necessary, remove any extra spaces representing line breaks.

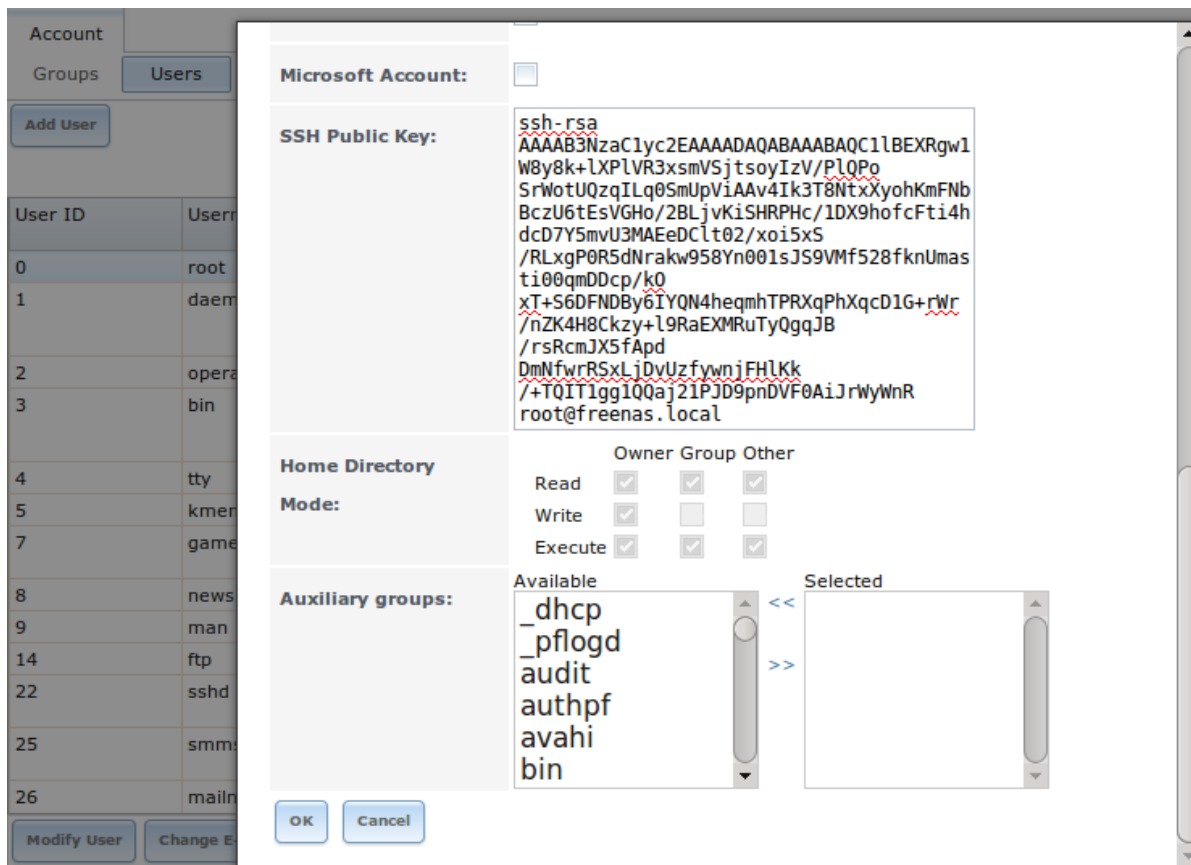


Fig. 6.8: Pasting the User's SSH Public Key

While on PULL, verify that the SSH service is running in **Services** → **Control Services** and start it if it is not.

Next, copy the host key of PULL using Shell on PUSH. The following command copies the RSA host key of the PULL server used in our previous example. Be sure to include the double bracket >> to prevent overwriting any existing entries in the `known_hosts` file:

```
ssh-keyscan -t rsa 192.168.2.6 >> /root/.ssh/known_hosts
```

Note: If PUSH is a Linux system, use this command to copy the RSA key to the Linux system:

```
cat ~/.ssh/id_rsa.pub | ssh user@192.168.2.6 'cat >> .ssh/authorized_keys'
```

The rsync task can now be created on PUSH. To configure rsync SSH mode using the systems in our previous example, the configuration is as follows:

- the Path points to `/mnt/local/images`, the directory to be copied
- the Remote Host points to `192.168.2.6`, the IP address of the rsync server
- the Rsync Mode is Rsync over SSH
- the rsync is scheduled to occur every 15 minutes
- the User is set to root so it has permission to write anywhere; the public key for this user must be generated on PUSH and copied to PULL
- the Preserve Permissions checkbox is checked so that the original permissions are not overwritten by the root user

Save the rsync task and the rsync will automatically occur according to the schedule. In this example, the contents of `/mnt/local/images/` will automatically appear in `/mnt/remote/images/` after 15 minutes. If the content does not appear, use Shell on PULL to read `/var/log/messages`. If the message indicates a `n` (newline character) in the key, remove the space in the pasted key—it will be after the character that appears just before the `n` in the error message.

6.5 S.M.A.R.T. Tests

S.M.A.R.T. (<http://en.wikipedia.org/wiki/S.M.A.R.T.>) (Self-Monitoring, Analysis and Reporting Technology) is a monitoring system for computer hard disk drives to detect and report on various indicators of reliability. When a failure is anticipated by S.M.A.R.T., the drive should be replaced. Most modern ATA, IDE, and SCSI-3 hard drives support S.M.A.R.T.—refer to the drive documentation for confirmation.

Figure 6.9 shows the configuration screen that appears after selecting `Tasks → S.M.A.R.T. Tests → Add S.M.A.R.T. Test`. Tests are listed under `View S.M.A.R.T. Tests`. After creating tests, check the configuration in `Services → S.M.A.R.T.`, then click the slider to ON for the S.M.A.R.T. service in `Services → Control Services`. The S.M.A.R.T. service will not start if there are no volumes.

Note: To prevent problems, do not enable the S.M.A.R.T. service if the disks are controlled by a RAID controller. It is the job of the controller to monitor S.M.A.R.T. and mark drives as Predictive Failure when they trip.

Add S.M.A.R.T. Test

Disks:

ada0
ada1
ada2
ada3

Type:

Short description:

Hour:

Every N hour

Each selected hour

1

Day of month:

Every N day of month

Each selected day of month

1

Month:

☒

January

☒

February

☒

March

☒

April

☒

May

☒

June

☒

July

☒

August

☒

September☒☒☒

Fig. 6.9: Adding a S.M.A.R.T. Test

Table 6.5 summarizes the configurable options when creating a S.M.A.R.T. test.

Table 6.5: S.M.A.R.T. Test Options

Setting	Value	Description
Disks	list	highlight disks to monitor
Type	drop-down menu	select type of test to run; see smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) for a description of each type of test (note that some test types will degrade performance or take disks offline; do not schedule S.M.A.R.T. tests at the same time as a scrub or during a resilver operation)
Short description	string	optional
Hour	slider or hour selections	if use the slider, test occurs every N hours; if use hour selections, test occurs at the highlighted hours
Day of month	slider or day selections	if use the slider, test occurs every N days; if use day selections, test occurs on the highlighted days
Month	checkboxes	select the months for the test to occur
Day of week	checkboxes	select the days of the week for the test to occur

An example configuration is to schedule a Short Self-Test once a week and a Long Self-Test once a month. These tests should not have a performance impact, as the disks prioritize normal I/O over the tests. If a disk fails a test, even if the overall status is Passed, start to think about replacing that disk.

Warning: Some S.M.A.R.T. tests cause heavy disk activity and can drastically reduce disk performance. Do not schedule S.M.A.R.T. tests to run at the same time as scrub or resilver operations or during other periods of intense disk activity.

Which tests will run and when can be verified by typing `smartd -q showtests` within [Shell](#) (page 290).

The results of a test can be checked from [Shell](#) (page 290) by specifying the name of the drive. For example, to see the results for disk `ada0`, type:

```
smartctl -l selftest /dev/ada0
```

If an email address is entered in the Email to report field of [Services](#) → S.M.A.R.T., the system will send email to that address when a test fails.

NETWORK

The Network section of the administrative GUI contains these components for viewing and configuring network settings on the FreeNAS® system:

- [Global Configuration](#) (page 102): general network settings.
- [Interfaces](#) (page 104): settings for each network interface.
- [IPMI](#) (page 106): settings controlling connection to the appliance through the hardware side-band management interface if the graphical user interface becomes unavailable.
- [Link Aggregations](#) (page 108): settings for network link aggregation and link failover.
- [Network Summary](#) (page 112): display an overview of the current network settings.
- [Static Routes](#) (page 112): add static routes.
- [VLANs](#) (page 113): configure IEEE 802.1q tagging for virtual LANs.

Each of these is described in more detail in this section.

7.1 Global Configuration

Network → Global Configuration, shown in [Figure 7.1](#), is for general network settings that are not unique to any particular network interface.

Network

Global Configuration
Interfaces
Link Aggregations
Network Summary
Static Routes
VLANs

Hostname:
freenas

Domain:
local

Additional domains:
i

IPv4 Default Gateway:

IPv6 Default Gateway:

Nameserver 1:

Nameserver 2:

Nameserver 3:

HTTP Proxy:

Enable netwait feature:
☐ *i*

Netwait IP list:
i

Host name data base:
i

Save

Fig. 7.1: Global Network Configuration

Table 7.1 summarizes the settings on the Global Configuration tab. Hostname and domain fields are pre-filled as shown in Figure 7.1, but can be changed to meet requirements of the local network.

Table 7.1: Global Configuration Settings

Setting	Value	Description
Hostname	string	system host name
Domain	string	system domain name
Additional domains	string	can enter up to 6 space delimited search domains; adding multiple domains may result in slower DNS lookups
IPv4 Default Gateway	IP address	typically not set (see Note below); if set, used instead of default gateway provided by DHCP

Continued on next page

Table 7.1 – continued from previous page

Setting	Value	Description
IPv6 Default Gateway	IP address	typically not set (see Note below)
Nameserver 1	IP address	primary DNS server (typically in Windows domain)
Nameserver 2	IP address	secondary DNS server
Nameserver 3	IP address	tertiary DNS server
HTTP Proxy	string	enter the proxy information for the network in the format <code>http://my.proxy.server:3128</code> or <code>http://user:password@my.proxy.server:3128</code>
Enable netwait feature	checkbox	if enabled, network services are not started at boot until the interface is able to ping the addresses listed in Netwait IP list
Netwait IP list	string	if Enable netwait feature is checked, list of IP addresses to ping; otherwise, ping the default gateway
Host name database	string	used to add one entry per line which will be appended to <code>/etc/hosts</code> ; use the format <code>IP_address space hostname</code> where multiple hostnames can be used if separated by a space

When Active Directory is being used, set the IP address of the realm's DNS server in the Nameserver 1 field.

If the network does not have a DNS server, or NFS, SSH, or FTP users are receiving "reverse DNS" or timeout errors, add an entry for the IP address of the FreeNAS® system in the Host name database field.

Note: In many cases, a FreeNAS® configuration does not include default gateway information as a way to make it more difficult for a remote attacker to communicate with the server. While this is a reasonable precaution, such a configuration does not restrict inbound traffic from sources within the local network. However, omitting a default gateway will prevent the FreeNAS® system from communicating with DNS servers, time servers, and mail servers that are located outside of the local network. In this case, it is recommended to add [Static Routes](#) (page 112) to be able to reach external DNS, NTP, and mail servers which are configured with static IP addresses. When a gateway to the Internet is added, make sure that the FreeNAS® system is protected by a properly configured firewall.

7.2 Interfaces

`Network → Interfaces` shows which interfaces have been manually configured and allows adding or editing a manually configured interface.

Note: Typically, the interface used to access the FreeNAS® administrative GUI is configured by DHCP. This interface does not appear in this screen, even though it is already dynamically configured and in use.

[Figure 7.2](#) shows the screen that opens on clicking `Interfaces → Add Interface`. [Table 7.2](#) summarizes the configuration options shown when adding an interface or editing an already configured interface. Note that if any changes to this screen require a network restart, the screen will turn red when the OK button is clicked and a pop-up message will point out that network connectivity to the FreeNAS® system will be interrupted while the changes are applied.

Fig. 7.2: Adding or Editing an Interface

Table 7.2: Interface Configuration Settings

Setting	Value	Description
NIC	drop-down menu	the FreeBSD device name of the interface; a read-only field when editing an interface
Interface Name	string	description of interface
DHCP	checkbox	requires static IPv4 or IPv6 configuration if unchecked; only one interface can be configured for DHCP
IPv4 Address	IP address	enter a static IP address if DHCP is unchecked
IPv4 Netmask	drop-down menu	enter a netmask if DHCP is unchecked
Auto configure IPv6	checkbox	only one interface can be configured for this option; if unchecked, manual configuration is required to use IPv6
IPv6 Address	IPv6 address	must be unique on network
IPv6 Prefix Length	drop-down menu	match the prefix used on network
Options	string	additional parameters from ifconfig(8) (http://www.freebsd.org/cgi/man.cgi?query=ifconfig), separate multiple parameters with a space; for example: <code>mtu 9000</code> increases the MTU for interfaces which support jumbo frames (but see this note (page 112) about MTU and lagg interfaces)

This screen also provides for the configuration of IP aliases, making it possible for a single interface to have multiple IP addresses. To set multiple aliases, click the Add extra alias link for each alias. Aliases are deleted by clicking the interface in the tree, clicking the Edit button, checking the Delete checkbox below the alias, then clicking the OK button.

Warning: Aliases are deleted by checking the Delete checkbox in the alias area, then clicking OK for the interface. Do not click the Delete button at the bottom of this screen, which deletes the entire interface.

Multiple interfaces cannot be members of the same subnet. See [Multiple network interfaces on a single subnet](https://forums.freenas.org/index.php?threads/multiple-network-interfaces-on-a-single-subnet.20204/) (<https://forums.freenas.org/index.php?threads/multiple-network-interfaces-on-a-single-subnet.20204/>) for more information. Check the subnet mask if an error is shown when setting the IP addresses on multiple interfaces.

This screen will not allow an interface's IPv4 and IPv6 addresses to both be set as primary addresses. An error is shown if both the IPv4 address and IPv6 address fields are filled in. Instead, set only one of these address fields and create an alias for the other address.

7.3 IPMI

Beginning with version 9.2.1, FreeNAS® provides a graphical screen for configuring an IPMI interface. This screen will only appear if the system hardware includes a Baseboard Management Controller (BMC).

IPMI provides side-band management if the graphical administrative interface becomes unresponsive. This allows for a few vital functions, such as checking the log, accessing the BIOS setup, and powering on the system without requiring physical access to the system. IPMI can also be used to allow another person remote access to the system to assist with a configuration or troubleshooting issue. Before configuring IPMI, ensure that the management interface is physically connected to the network. The IPMI device may share the primary Ethernet interface, or it may be a dedicated separate IPMI interface.

Warning: It is recommended to first ensure that the IPMI has been patched against the Remote Management Vulnerability before enabling IPMI. This [article](http://www.ixsystems.com/whats-new/how-to-fix-the-ipmi-remote-management-vulnerability/) (<http://www.ixsystems.com/whats-new/how-to-fix-the-ipmi-remote-management-vulnerability/>) provides more information about the vulnerability and how to fix it.

Note: Some IPMI implementations require updates to work with newer versions of Java. See [PSA: Java 8 Update 131 breaks ASRock's IPMI Virtual console](https://forums.freenas.org/index.php?threads/psa-java-8-update-131-breaks-asrock-ipmi-virtual-console.53911/) (<https://forums.freenas.org/index.php?threads/psa-java-8-update-131-breaks-asrock-ipmi-virtual-console.53911/>) for more information.

IPMI is configured from `Network` → `IPMI`. The IPMI configuration screen, shown in [Figure 7.3](#), provides a shortcut to the most basic IPMI configuration. Those already familiar with IPMI management tools can use them instead. [Table 7.3](#) summarizes the options available when configuring IPMI with the FreeNAS® GUI.

The screenshot shows the 'Network' configuration page with the 'IPMI' sub-tab selected. The form contains the following fields and controls:

- Channel:** A dropdown menu with '1' selected.
- Password:** A text input field.
- Password confirmation:** A text input field with an information icon to its right.
- DHCP:** A checked checkbox.
- IPv4 Address:** A text input field with a dotted pattern.
- IPv4 Netmask:** A dropdown menu with '/16 (255.255.0.0)' selected.
- IPv4 Default Gateway:** A text input field with a dotted pattern.
- VLAN ID:** A text input field.
- Buttons:** 'OK', 'Cancel', and 'Identify Light' at the bottom.

Fig. 7.3: IPMI Configuration

Table 7.3: IPMI Options

Setting	Value	Description
Channel	drop-down menu	select the channel to use
Password	string	enter the password used to connect to the IPMI interface from a web browser
DHCP	checkbox	if left unchecked, the following three fields must be set
IPv4 Address	string	IP address used to connect to the IPMI web GUI
IPv4 Netmask	drop-down menu	subnet mask associated with the IP address
IPv4 Default Gateway	string	default gateway associated with the IP address
VLAN ID	string	enter the VLAN identifier if the IPMI out-of-band management interface is not on the same VLAN as management networking

The Identify Light button can be used to identify a system in a multi-system rack by flashing its IPMI LED light. Clicking this button will present a pop-up with a menu of times, ranging from 15 seconds to 4 minutes, to flash the LED light.

After configuration, the IPMI interface is accessed using a web browser and the IP address specified in the configuration. The management interface prompts for a username and the configured password. Refer to the IPMI device's documentation to determine the default administrative username.

After logging in to the management interface, the default administrative username can be changed, and additional users created. The appearance of the IPMI utility and the functions that are available vary depending on the hardware.

A command-line utility called `ipmitool` is available to control many features of the IPMI interface. See [How To: Change IPMI Sensor Thresholds using ipmitool](https://forums.freenas.org/index.php?resources/how-to-change-ipmi-sensor-thresholds-using-ipmitool.35/) (<https://forums.freenas.org/index.php?resources/how-to-change-ipmi-sensor-thresholds-using-ipmitool.35/>) for some examples.

7.4 Link Aggregations

FreeNAS® uses FreeBSD's `lagg(4)` (<http://www.freebsd.org/cgi/man.cgi?query=lagg>) interface to provide link aggregation and link failover. The `lagg` interface allows aggregation of multiple network interfaces into a single virtual `lagg` interface, providing fault-tolerance and high-speed multi-link throughput. The aggregation protocols supported by `lagg` determine which ports are used for outgoing traffic and whether a specific port accepts incoming traffic. The link state of the `lagg` interface is used to validate whether the port is active.

Aggregation works best on switches supporting LACP, which distributes traffic bi-directionally while responding to failure of individual links. FreeNAS® also supports active/passive failover between pairs of links. The LACP and load-balance modes select the output interface using a hash that includes the Ethernet source and destination address, VLAN tag (if available), IP source and destination address, and flow label (IPv6 only). The benefit can only be observed when multiple clients are transferring files from the NAS. The flow entering into the NAS depends on the Ethernet switch load-balance algorithm.

The `lagg` driver currently supports several aggregation protocols, although only Failover is recommended on network switches that do not support LACP:

Failover: the default protocol. Sends traffic only through the active port. If the master port becomes unavailable, the next active port is used. The first interface added is the master port; any interfaces added after that are used as failover devices. By default, received traffic is only accepted when received through the active port. This constraint can be relaxed, which is useful for certain bridged network setups, by creating a tunable with a Variable of `net.link.lagg.failover_rx_all`, a Value of a non-zero integer, and a Type of `Sysctl` in `System` → `Tunables` → `Add Tunable`.

LACP: supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker Protocol. LACP negotiates a set of aggregable links with the peer into one or more link aggregated groups (LAGs). Each LAG is composed of ports of the same speed, set to full-duplex operation. Traffic is balanced across the ports in the LAG with the greatest total speed; in most cases there will only be one LAG which contains all ports. In the event of changes in physical connectivity, link aggregation will quickly converge to a new configuration. LACP must be configured on the switch, and LACP does not support mixing interfaces of different speeds. Only interfaces that use the same driver, like two `igb` ports, are recommended for LACP. Using LACP for iSCSI is not recommended, as iSCSI has built-in multipath features which are more efficient.

Load Balance: balances outgoing traffic across the active ports based on hashed protocol header information and accepts incoming traffic from any active port. This is a static setup and does not negotiate aggregation with the peer or exchange frames to monitor the link. The hash includes the Ethernet source and destination address, VLAN tag (if available), and IP source and destination address. Requires a switch which supports IEEE 802.3ad static link aggregation.

Round Robin: distributes outgoing traffic using a round-robin scheduler through all active ports and accepts incoming traffic from any active port. This mode can cause unordered packet arrival at the client. This has a side effect of limiting throughput as reordering packets can be CPU intensive on the client. Requires a switch which supports IEEE 802.3ad static link aggregation.

None: this protocol disables any traffic without disabling the `lagg` interface itself.

Note: When using LACP, verify that the switch is configured for active LACP. Passive LACP is not supported.

7.4.1 LACP, MPIO, NFS, and ESXi

LACP bonds Ethernet connections to improve bandwidth. For example, four physical interfaces can be used to create one mega interface. However, it cannot increase the bandwidth for a single conversation. It is designed to increase bandwidth when multiple clients are simultaneously accessing the same system. It also assumes that quality Ethernet hardware is used and it will not make much difference when using inferior Ethernet chipsets such as a Realtek.

LACP reads the sender and receiver IP addresses and, if they are deemed to belong to the same TCP connection, always sends the packet over the same interface to ensure that TCP does not need to reorder packets. This makes LACP ideal for load balancing many simultaneous TCP connections, but does nothing for increasing the speed over one TCP connection.

MPIO operates at the iSCSI protocol level. For example, if four IP addresses are created and there are four simultaneous TCP connections, MPIO will send the data over all available links. When configuring MPIO, make sure that the IP addresses

on the interfaces are configured to be on separate subnets with non-overlapping netmasks, or configure static routes to do point-to-point communication. Otherwise, all packets will pass through one interface.

LACP and other forms of link aggregation generally do not work well with virtualization solutions. In a virtualized environment, consider the use of iSCSI MPIO through the creation of an iSCSI Portal with at least two network cards on different networks. This allows an iSCSI initiator to recognize multiple links to a target, using them for increased bandwidth or redundancy. This [how-to](https://fojta.wordpress.com/2010/04/13/iscsi-and-esxi-multipathing-and-jumbo-frames/) (<https://fojta.wordpress.com/2010/04/13/iscsi-and-esxi-multipathing-and-jumbo-frames/>) contains instructions for configuring MPIO on ESXi.

NFS does not understand MPIO. Therefore, one fast interface is needed, since creating an iSCSI portal will not improve bandwidth when using NFS. LACP does not work well to increase the bandwidth for point-to-point NFS (one server and one client). LACP is a good solution for link redundancy or for one server and many clients.

7.4.2 Creating a Link Aggregation

Before creating a link aggregation, double-check that no interfaces have been manually configured in `Network → Interfaces → View Interfaces`.

If any manually-configured interfaces exist, delete them as lagg creation fails if any interfaces are manually configured.

Warning: Creating or editing link aggregations can disconnect clients using the FreeNAS® computer. Please verify that clients have saved their work and are not connected through the affected networks before making changes.

Figure 7.4 shows the configuration options when adding a lagg interface using `Network → Link Aggregations → Create Link Aggregation`.

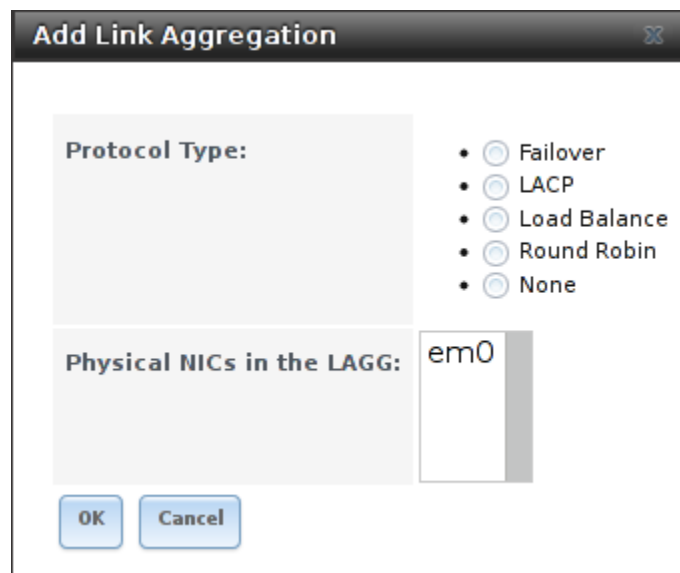


Fig. 7.4: Creating a lagg Interface

Note: If interfaces are installed but do not appear in the Physical NICs list, check that a FreeBSD driver for the interface exists [here](http://www.freebsd.org/releases/11.0R/hardware.html#ETHERNET) (<http://www.freebsd.org/releases/11.0R/hardware.html#ETHERNET>).

To create a link aggregation, select the desired Protocol Type. LACP is preferred. If the network switch does not support LACP, choose Failover. Highlight the interfaces to associate with the lagg device, and click the OK button.

Once the lagg device has been created, click its entry to enable its Edit, Delete, and Edit Members buttons.

Clicking the Edit button for a lagg opens the configuration screen shown in Figure 7.5. Table 7.4 describes the options in this screen.

If the network interface used to connect to the FreeNAS® web GUI is a member of the lagg, the network connection will be lost when the new lagg is created. The switch settings might also require changes to communicate through the new lagg interface.

The IP address of the new lagg can be set with DHCP or manually from the console setup menu. If the IP address is set manually, it might also be necessary to enter a default gateway to allow access to the GUI from the new lagg interface.

The screenshot shows the 'Edit' configuration screen for a Link Aggregation (lagg) interface. The interface is titled 'Edit' and is part of the 'Link Aggregation' section. It contains several input fields and checkboxes for configuring the lagg interface. The background shows a table of interfaces with 'lagg0 (none: em0)' selected.

Interface	Protocol Type
lagg0 (none: em0)	none

Buttons: Add Link Aggregation, Edit, Delete, Edit Members

Configuration fields:

- NIC: lagg0
- Interface Name: lagg0
- DHCP: ☐
- IPv4 Address:
- IPv4 Netmask:
- Auto configure IPv6: ☐
- IPv6 Address:
- IPv6 Prefix Length:
- Options:
- Alias:
- IPv4 Address:

Fig. 7.5: Editing a lagg

Table 7.4: Configurable Options for a lagg

Setting	Value	Description
NIC	string	read-only; automatically assigned the next available numeric ID
Interface Name	string	by default same as device (NIC) name, can be changed to a more descriptive value
DHCP	checkbox	check if the lagg device will get IP address info from DHCP server
IPv4 Address	string	enter a static IP address if DHCP is left unchecked
IPv4 Netmask	drop-down menu	enter a netmask if DHCP is left unchecked
Auto configure IPv6	checkbox	check only if DHCP server available to provide IPv6 address info
IPv6 Address	string	optional
IPv6 Prefix Length	drop-down menu	required if an IPv6 address is entered

Continued on next page

Table 7.4 – continued from previous page

Setting	Value	Description
Options	string	additional <code>ifconfig(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ifconfig) options

This screen also allows the configuration of an alias for the lagg interface. Multiple aliases can be added with the Add extra Alias link.

Click the Edit Members button, click the entry for a member, then click its Edit button to see the configuration screen shown in Figure 7.6. The configurable options are summarized in Table 7.5.

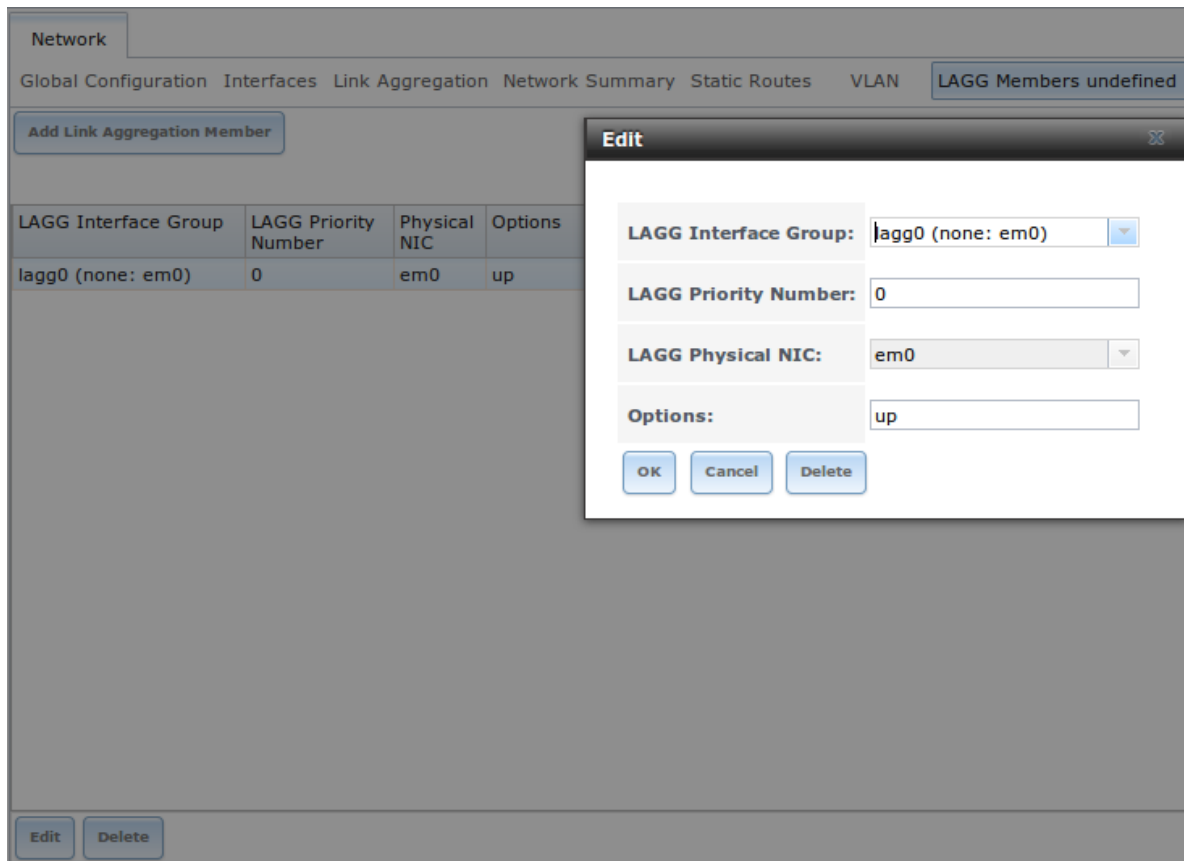


Fig. 7.6: Editing a Member Interface

Table 7.5: Configuring a Member Interface

Setting	Value	Description
LAGG Interface group	drop-down menu	select the member interface to configure
LAGG Priority Number	integer	order of selected interface within the lagg; configure a failover to set the master interface to 0 and the other interfaces to 1, 2, etc.
LAGG Physical NIC	drop-down menu	physical interface of the selected member
Options	string	additional parameters from <code>ifconfig(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ifconfig)

Options can be set at the lagg level using the Edit button, or at the individual parent interface level using the Edit Members button. Changes are typically made at the lagg level (Figure 7.5) as each interface member will inherit from the lagg. To

configure at the interface level (Figure 7.6) instead, the configuration must be repeated for each interface within the lagg. Some options can only be set on the parent interfaces and are inherited by the lagg interface. For example, to set the MTU on a lagg, use Edit Members to set the MTU for each parent interface.

If the MTU settings on the lagg member interfaces are not identical, the smallest value is used for the MTU of the entire lagg.

Note: A reboot is required after changing the MTU to create a jumbo frame lagg.

To see if the link aggregation is properly load balancing, run this command from Shell (page 290):

```
systat -ifstat
```

More information about this command can be found at [systat\(1\)](http://www.freebsd.org/cgi/man.cgi?query=systat) (<http://www.freebsd.org/cgi/man.cgi?query=systat>).

7.5 Network Summary

Network → Network Summary shows a quick summary of the addressing information of every configured interface. For each interface name, the configured IPv4 and IPv6 addresses, DNS servers, and default gateway are displayed.

7.6 Static Routes

No static routes are defined on a default FreeNAS® system. If a static route is required to reach portions of the network, add the route with Network → Static Routes → Add Static Route, shown in Figure 7.7.

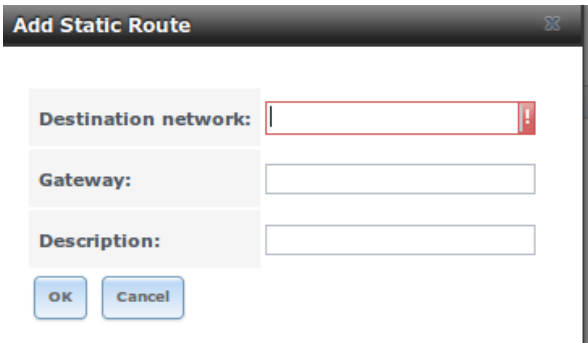


Fig. 7.7: Adding a Static Route

The available options are summarized in Table 7.6.

Table 7.6: Static Route Options

Setting	Value	Description
Destination net-work	integer	use the format A.B.C.D/E where E is the CIDR mask
Gateway	integer	enter the IP address of the gateway
Description	string	optional

Added static routes are shown in View Static Routes. Click a route’s entry to access the Edit and Delete buttons.

7.7 VLANs

FreeNAS® uses FreeBSD's [vlan\(4\)](http://www.freebsd.org/cgi/man.cgi?query=vlan) (<http://www.freebsd.org/cgi/man.cgi?query=vlan>) interface to demultiplex frames with IEEE 802.1q tags. This allows nodes on different VLANs to communicate through a layer 3 switch or router. A vlan interface must be assigned a parent interface and a numeric VLAN tag. A single parent can be assigned to multiple vlan interfaces provided they have different tags.

Note: VLAN tagging is the only 802.1q feature that is implemented. Additionally, not all Ethernet interfaces support full VLAN processing. See the **HARDWARE** section of [vlan\(4\)](http://www.freebsd.org/cgi/man.cgi?query=vlan) (<http://www.freebsd.org/cgi/man.cgi?query=vlan>) for details.

Click **Network** → **VLANs** → **Add VLAN**, to see the screen shown in [Figure 7.8](#).

Fig. 7.8: Adding a VLAN

[Table 7.7](#) summarizes the configurable fields.

Table 7.7: Adding a VLAN

Setting	Value	Description
Virtual Interface	string	use the format <code>vlanX</code> where X is a number representing a vlan interface not currently being used as a parent
Parent Interface	drop-down menu	usually an Ethernet card connected to a properly configured switch port; note that newly created Link Aggregations (page 108) will not appear in the drop-down until the system is rebooted
VLAN Tag	integer	number between 1 and 4095 which matches a numeric tag set up in the switched network
Priority Code Point	drop-down menu	available 802.1p Class of Service ranges from Best Effort (default) to Network Control (highest)
Description	string	optional

The parent interface of a VLAN must be up, but it can have an IP address or it can be unconfigured, depending upon the requirements of the VLAN configuration. This makes it difficult for the GUI to do the right thing without trampling the configuration. To remedy this, after adding the VLAN, go to **Network** → **Interfaces** → **Add Interface**. Select the parent interface from the NIC drop-down menu and in the Options field, type `up`. This will bring up the parent interface. If

an IP address is required, it can be configured using the rest of the options in the Add Interface screen.

Warning: Creating a vlan will cause network connectivity to be interrupted. Accordingly, the GUI will provide a warning and an opportunity to cancel the vlan creation.

STORAGE

The Storage section of the graphical interface allows configuration of these options:

- [Volumes](#) (page 115) creates and manages storage volumes.
- [Periodic Snapshot Tasks](#) (page 140) schedules automatic creation of filesystem snapshots.
- [Replication Tasks](#) (page 142) automate the replication of snapshots to a remote system.
- [Resilver Priority](#) (page 151) controls the priority of resilvers.
- [Scrubs](#) (page 152) schedules scrubs as part of ongoing disk maintenance.
- [Snapshots](#) (page 155) manages local snapshots.
- [VMware-Snapshot](#) (page 157) coordinates ZFS snapshots with a VMware datastore.

8.1 Volumes

The Volumes section of the FreeNAS® graphical interface can be used to format volumes, attach a disk to copy data onto an existing volume, or import a ZFS volume. It can also be used to create ZFS datasets and zvols and to manage their permissions.

Note: In ZFS terminology, groups of storage devices managed by ZFS are referred to as a pool. The FreeNAS® graphical interface uses the term volume to refer to a ZFS pool.

Proper storage design is important for any NAS. Please read through this entire chapter before configuring storage disks. Features are described to help make it clear which are beneficial for particular uses, and caveats or hardware restrictions which limit usefulness.

8.1.1 Volume Manager

The Volume Manager is used to add disks to a ZFS pool. Any old data on added disks is overwritten, so save it elsewhere before reusing a disk. Please see the [ZFS Primer](#) (page 322) for information on ZFS redundancy with multiple disks before using Volume Manager. It is important to realize that different layouts of virtual devices (vdevs) affect which operations can be performed on that volume later. For example, drives can be added to a mirror to increase redundancy, but that is not possible with RAIDZ arrays.

Selecting `Storage → Volumes → Volume Manager` opens a screen like the example shown in [Figure 8.1](#).

The screenshot shows the 'Volume Manager' window with the following fields and controls:

- Volume Name:** A text input field.
- Volume to extend:** A dropdown menu with a dashed line icon.
- Encryption:** A checkbox.
- Available disks:** A box containing a '+' button and the text '1 - 10.7 GB (3 drives, show)'.
- Volume layout (Estimated capacity: 0 B):** A section containing:
 - A dropdown menu.
 - A slider bar with 15 numbered positions (1-15).
 - Text: '0x1x0 B' and 'Capacity: 0 B'.
 - An 'Add Extra Device' button.
- Buttons at the bottom:** 'Add Volume' (with red text 'Existing data will be cleared'), 'Cancel', and 'Manual setup'.

Fig. 8.1: Creating a ZFS Pool Using Volume Manager

Table 8.1 summarizes the configuration options of this screen.

Table 8.1: ZFS Volume Creation Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html); choosing a name that will stick out in the logs (e.g. not a generic term like <code>data</code> or <code>freenas</code>) is recommended
Volume to extend	drop-down menu	extend an existing ZFS pool; see Extending a ZFS Volume (page 119) for more details
Encryption	checkbox	see the warnings in Encryption (page 117) before enabling encryption
Available disks	display	display the number and size of available disks; hover over show to list the available device names; click the + to add all of the disks to the pool
Volume layout	drag and drop	click and drag the icon to select the desired number of disks for a vdev; when at least one disk is selected, the layouts supported by the selected number of disks are added to the drop-down menu
Add Extra Device	button	configure multiple vdevs or add log or cache devices during pool creation
Manual setup	button	create a pool manually (not recommended); see Manual Setup (page 118) for details

Drag the slider to select the desired number of disks. Volume Manager displays the resulting storage capacity, taking reserved swap space into account. To change the layout or the number of disks, drag the slider to the desired volume layout. The Volume layout drop-down menu can also be clicked if a different level of redundancy is required.

Note: For performance and capacity reasons, this screen does not allow creating a volume from disks of differing sizes. While it is not recommended, it is possible to create a volume of differently-sized disks with the Manual setup button. Follow the instructions in [Manual Setup](#) (page 118).

Volume Manager only allows choosing a configuration if enough disks have been selected to create that configuration. These layouts are supported:

- Stripe: requires at least one disk
- Mirror: requires at least two disks
- RAIDZ1: requires at least three disks
- RAIDZ2: requires at least four disks
- RAIDZ3: requires at least five disks
- log device: requires at least one dedicated device, a fast, low-latency, power-protected SSD is recommended
- cache device: requires at least one dedicated device, SSD is recommended

When more than five disks are used, consideration must be given to the optimal layout for the best performance and scalability. An overview of the recommended disk group sizes as well as more information about log and cache devices can be found in the [ZFS Primer](#) (page 322).

The Add Volume button warns that existing data will be cleared. In other words, creating a new volume reformats the selected disks. To preserve existing data, click the Cancel button and refer to [Import Disk](#) (page 127) and [Import Volume](#) (page 127) to see if the existing format is supported. If so, perform that action instead. If the current storage format is not supported, it is necessary to back up the data to external media, format the disks, then restore the data to the new volume.

Depending on the size and number of disks, the type of controller, and whether encryption is selected, creating the volume may take some time. After the volume is created, the screen refreshes and the new volume is listed in the tree under *Storage* → *Volumes*. Click the + next to the volume name to access [Change Permissions](#) (page 120), [Create Dataset](#) (page 122), and [Create zvol](#) (page 125) options for that volume.

Encryption

Note: The encryption facility used by FreeNAS® is designed to protect against physical theft of the disks. It is not designed to protect against unauthorized software access. Ensure that only authorized users have access to the administrative GUI and that proper permissions are set on shares if sensitive data is stored on the system.

FreeNAS® supports [GELI](http://www.freebsd.org/cgi/man.cgi?query=geli) (<http://www.freebsd.org/cgi/man.cgi?query=geli>) full disk encryption for ZFS volumes. It is important to understand the details when considering whether encryption is right for your FreeNAS® system:

- FreeNAS® encryption is different from the encryption used in Oracle's proprietary, non-open source version of ZFS.
- In FreeNAS®, entire disks are encrypted, not individual filesystems. Encrypted devices are created from the underlying drives, then the volume (pool) is created on top of the encrypted devices. Data is encrypted as it is written and decrypted as it is read.
- This type of encryption is primarily useful for users storing sensitive data but wanting the ability to remove disks from the pool without having to first wipe the disk contents.
- The FreeNAS® encryption design is only suitable for safe disposal of disks independent of the encryption key. As long as the key and the disks are intact, the system is vulnerable to being decrypted. The key should be protected by a strong passphrase and any backups of the key should be securely stored.
- If the encryption key is lost, the data on the disks is inaccessible. Always back up the key!

- Encryption keys are per ZFS volume (pool). Each pool has a separate encryption key. Technical details about how encryption keys are used, stored, and managed within FreeNAS® are described in this [forum post](https://forums.freenas.org/index.php?threads/recover-encryption-key.16593/#post-85497) (<https://forums.freenas.org/index.php?threads/recover-encryption-key.16593/#post-85497>).
- Data in memory, including ARC, is not encrypted. ZFS data on disk, including ZIL and SLOG, are encrypted if the underlying disks are encrypted. Swap data on disk is always encrypted.

Warning: Data stored in Cache (L2ARC) drives is not encrypted. Do not use Cache (L2ARC) with encrypted volumes.

- At present, there is no one-step way to encrypt an existing, unencrypted volume. Instead, the data must be backed up, the existing pool destroyed, a new encrypted volume created, and the backup restored to the new volume.
- Hybrid pools are not supported. Added vdevs must match the existing encryption scheme. The [Volume Manager](#) (page 115) automatically encrypts a new vdev being added to an existing encrypted pool.

To create an encrypted volume, check the Encryption box shown in [Figure 8.1](#). A pop-up message shows a reminder that it is extremely important to make a backup of the key. Without the key, the data on the disks is inaccessible. See [Managing Encrypted Volumes](#) (page 133) for instructions.

Encryption Performance

Performance depends upon the number of disks encrypted. The more drives in an encrypted volume, the more encryption and decryption overhead, and the greater the impact on performance. Encrypted volumes composed of more than eight drives can suffer severe performance penalties. If encryption is desired, please benchmark such volumes before using them in production.

Note: Processors with support for the [AES-NI](https://en.wikipedia.org/wiki/AES-NI#Supporting_CPUs) (https://en.wikipedia.org/wiki/AES-NI#Supporting_CPUs) instruction set are strongly recommended. These processors can handle encryption of a small number of disks with negligible performance impact. They also retain performance better as the number of disks increases. Older processors without the AES-NI instructions see significant performance impact with even a single encrypted disk. This [forum post](https://forums.freenas.org/index.php?threads/encryption-performance-benchmarks.12157/) (<https://forums.freenas.org/index.php?threads/encryption-performance-benchmarks.12157/>) compares the performance of various processors.

Manual Setup

The Manual Setup button shown in [Figure 8.1](#) can be used to create a ZFS volume manually. While this is not recommended, it can, for example, be used to create a non-optimal volume containing disks of different sizes.

Note: The usable space of each disk in a volume is limited to the size of the smallest disk in the volume. Because of this, creating volumes with disks of the same size through the Volume Manager is recommended.

[Figure 8.2](#) shows the Manual Setup screen. [Table 8.2](#) shows the available options.

Manual Setup

Volume name:

Encryption: ☐

Member disks (0):

- ada1 (21.5 GB)
- ada2 (21.5 GB)
- ada3 (21.5 GB)
- ada4 (21.5 GB)
- ada5 (21.5 GB)

Deduplication:

ZFS Extra:

Disk	None	Log	Cache	Spare
ada1	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
ada2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
ada3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
ada4	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
ada5	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Existing data will be cleared

Fig. 8.2: Manually Creating a ZFS Volume

Note: Because of the disadvantages of creating volumes with disks of different sizes, the displayed list of disks is sorted by size.

Table 8.2: Manual Setup Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (http://docs.oracle.com/cd/E19082-01/817-2271/gbcpt/index.html); choose a name that will stand out in the logs (e.g. not <code>data</code> or <code>freenas</code>)
Encryption	checkbox	see the warnings in Encryption (page 117) before using encryption
Member disks	list	highlight desired number of disks from list of available disks
Deduplication	drop-down menu	choices are Off, Verify, and On; carefully consider the section on Deduplication (page 124) before changing this setting
ZFS Extra	bullet selection	specify disk usage: storage (None), a log device, a cache device, or a spare

Extending a ZFS Volume

The Volume to extend drop-down menu in `Storage → Volumes → Volume Manager`, shown in [Figure 8.1](#), is used to add disks to an existing ZFS volume to increase capacity. This menu is empty if there are no ZFS volumes yet.

If more than one disk is added, the arrangement of the new disks into stripes, mirrors, or RAIDZ vdevs can be specified. Mirrors and RAIDZ arrays provide redundancy for data protection if an individual drive fails.

Note: If the existing volume is encrypted, a warning message shows a reminder that extending a volume resets the passphrase and recovery key. After extending the volume, immediately recreate both using the instructions in [Managing Encrypted Volumes](#) (page 133).

After an existing volume has been selected from the drop-down menu, drag and drop the desired disks and select the desired

volume layout. For example, disks can be added to increase the capacity of the volume.

When adding disks to increase the capacity of a volume, ZFS supports the addition of virtual devices, or vdevs, to an existing ZFS pool. A vdev can be a single disk, a stripe, a mirror, a RAIDZ1, RAIDZ2, or a RAIDZ3. After a vdev is created, more drives cannot be added to that vdev. However, a new vdev can be striped with another of the same type of existing vdev to increase the overall size of the volume. Extending a volume often involves striping similar vdevs. Here are some examples:

- to extend a ZFS stripe, add one or more disks. Since there is no redundancy, disks do not have to be added in the same quantity as the existing stripe.
- to extend a ZFS mirror, add the same number of drives. The resulting striped mirror is a RAID 10. For example, if ten new drives are available, a mirror of two drives could be created initially, then extended by creating another mirror of two drives, and repeating three more times until all ten drives have been added.
- to extend a three drive RAIDZ1, add three additional drives. The result is a RAIDZ+0, similar to RAID 50 on a hardware controller.
- to extend a RAIDZ2 requires a minimum of four additional drives. The result is a RAIDZ2+0, similar to RAID 60 on a hardware controller.

If an attempt is made to add a non-matching number of disks to the existing vdev, an error message appears, indicating the number of disks that are required. Select the correct number of disks to continue.

Adding L2ARC or SLOG Devices

Storage → Volumes → Volume Manager (see [Figure 8.1](#)) is also used to add L2ARC or SLOG SSDs to improve specific types of volume performance. This is described in more detail in the [ZFS Primer](#) (page 322).

After the SSDs have been physically installed, click the Volume Manager button and choose the volume from the Volume to extend drop-down menu. Click the + next to the SSD in the Available disks list. In the Volume layout drop-down menu, select Cache (L2ARC) to add a cache device, or Log (ZIL) to add a log device. Finally, click Extend Volume to add the SSD.

8.1.2 Change Permissions

Setting permissions is an important aspect of configuring volumes. The graphical administrative interface is meant to set the initial permissions for a volume or dataset in order to make it available as a share. Once a share is available, the client operating system should be used to fine-tune the permissions of the files and directories that are created by the client.

The chapter on [Sharing](#) (page 169) contains configuration examples for several types of permission scenarios. This section provides an overview of the screen that is used to set permissions.

Note: For users and groups to be available, they must either be first created using the instructions in [Account](#) (page 48) or imported from a directory service using the instructions in [Directory Services](#) (page 158). If more than 50 users or groups are available, the drop-down menus described in this section will automatically truncate their display to 50 for performance reasons. In this case, start to type in the desired user or group name so that the display narrows its search to matching results.

After a volume or dataset is created, it is listed by its mount point name in Storage → Volumes. Clicking the Change Permissions icon for a specific volume/dataset displays the screen shown in [Figure 8.3](#). [Table 8.3](#) summarizes the options in this screen.

Change Permissions

Change permission

Change permission on /mnt/volume1 to:

Apply Owner (user): ☒

Owner (user): root

Apply Owner (group): ☒

Owner (group): wheel

Apply Mode: ☒

Mode:

	Owner	Group	Other
Read	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Write	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Execute	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Permission Type:

- ☒ Unix
- ☐ Mac
- ☐ Windows

Set permission recursively: ☐

Fig. 8.3: Changing Permissions on a Volume or Dataset

Table 8.3: Options When Changing Permissions

Setting	Value	Description
Apply Owner (user)	checkbox	uncheck to prevent new permission change from being applied to Owner (user), see Note below
Owner (user)	drop-down menu	user to control the volume/dataset; users which were manually created or imported from a directory service will appear in the drop-down menu
Apply Owner (group)	checkbox	uncheck to prevent new permission change from being applied to Owner (group), see Note below
Owner (group)	drop-down menu	group to control the volume/dataset; groups which were manually created or imported from a directory service will appear in the drop-down menu
Apply Mode	checkbox	uncheck to prevent new permission change from being applied to Mode, see Note below
Mode	checkboxes	only applies to the Unix or Mac "Permission Type" so will be grayed out if Windows is selected
Permission Type	bullet selection	choices are Unix, Mac or Windows; select the type which matches the type of client accessing the volume/dataset
Set permission recursively	checkbox	if checked, permissions will also apply to subdirectories of the volume/dataset; if data already exists on the volume/dataset, change the permissions on the client side to prevent a performance lag

Note: The Apply Owner (user), Apply Owner (group), and Apply Mode checkboxes allow fine-tuning of the change permissions behavior. By default, all boxes are checked and FreeNAS® resets the owner, group, and mode when the Change button is clicked. These checkboxes allow choosing which settings to change. For example, to change just the Owner (group) setting, uncheck the boxes Apply Owner (user) and Apply Mode.

The Windows Permission Type is used for SMB shares or when the FreeNAS® system is a member of an Active Directory domain. This adds ACLs to traditional Unix permissions. When the Windows Permission Type is set, ACLs are set to Windows defaults for new files and directories. A Windows client can be used to further fine-tune permissions as needed.

The Unix Permission Type is usually used with NFS shares. These permissions are compatible with most network clients and generally work well with a mix of operating systems or clients. However, Unix permissions do not support Windows ACLs and should not be used with SMB shares.

The Mac Permission Type is used with AFP shares.

After a volume or dataset has been set to Windows, it cannot be changed to Unix permissions because that would remove extended permissions provided by Windows ACLs.

8.1.3 Create Dataset

An existing ZFS volume can be divided into datasets. Permissions, compression, deduplication, and quotas can be set on a per-dataset basis, allowing more granular control over access to storage data. Like a folder or directory, permissions can be set on dataset. Datasets are also similar to filesystems in that properties such as quotas and compression can be set, and snapshots created.

Note: ZFS provides thick provisioning using quotas and thin provisioning using reserved space.

Selecting an existing ZFS volume in the tree and clicking Create Dataset shows the screen in [Figure 8.4](#).

Create Dataset

Create ZFS dataset in volume1

Dataset Name:

Comments:

Sync:

Inherit (standard)

Compression level:

Inherit (lz4)

Share type:

UNIX

Enable atime:

☒ Inherit (on)
 ☐ On
 ☐ Off

ZFS Deduplication:

Enabling dedup can drastically reduce performance and affect the ability to access data. Compression usually offers similar space savings with much lower performance impact and overhead.

Inherit (off)

Case Sensitivity:

Sensitive

Add Dataset

Cancel

Advanced Mode

Fig. 8.4: Creating a ZFS Dataset

Table 8.4 shows the options available when creating a dataset. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button, or configure the system to always display advanced settings by checking the box Show advanced fields by default in *System* → *Advanced*. Most attributes, except for the Dataset Name, Case Sensitivity, and Record Size, can be changed after dataset creation by highlighting the dataset name and clicking the Edit Options button in *Storage* → *Volumes*.

Table 8.4: ZFS Dataset Options

Setting	Value	Description
Dataset Name	string	mandatory; enter a unique name for the dataset
Comments	string	short comments or user notes about this dataset
Sync	drop-down menu	data write synchronization: Inherit inherits the sync settings from the parent dataset; Standard uses the sync settings that have been requested by the client software; Always always waits for data writes to complete; Disabled never waits for writes to complete

Continued on next page

Table 8.4 – continued from previous page

Setting	Value	Description
Compression Level	drop-down menu	see the section on Compression (page 125) for a description of the available algorithms
Share type	drop-down menu	select the type of share that will be used on the dataset; choices are UNIX for an NFS share, Windows for a SMB share, or Mac for an AFP share
Enable atime	Inherit, On, or Off	controls whether the access time for files is updated when they are read; setting this property to Off avoids producing log traffic when reading files and can result in significant performance gains
Quota for this dataset	integer	only available in Advanced Mode; default of 0 disables quotas; specifying a value means to use no more than the specified size and is suitable for user datasets to prevent users from hogging available space
Quota for this dataset and all children	integer	only available in Advanced Mode; a specified value applies to both this dataset and any child datasets
Reserved space for this dataset	integer	only available in Advanced Mode; default of 0 is unlimited; specifying a value means to keep at least this much space free and is suitable for datasets containing logs which could take up all available free space
Reserved space for this dataset and all children	integer	only available in Advanced Mode; a specified value applies to both this dataset and any child datasets
ZFS Deduplication	drop-down menu	read the section on Deduplication (page 124) before making a change to this setting
Read-Only	drop-down menu	only available in Advanced Mode; choices are Inherit (off), On, or Off
Record Size	drop-down menu	only available in Advanced Mode; while ZFS automatically adapts the record size dynamically to adapt to data, if the data has a fixed size (e.g. a database), matching that size may result in better performance
Case Sensitivity	drop-down menu	choices are sensitive (default, assumes filenames are case sensitive), insensitive (assumes filenames are not case sensitive), or mixed (understands both types of filenames)

After a dataset is created, you can click on that dataset and select Create Dataset, thus creating a nested dataset, or a dataset within a dataset. A zvol can also be created within a dataset. When creating datasets, double-check that you are using the Create Dataset option for the intended volume or dataset. If you get confused when creating a dataset on a volume, click all existing datasets to close them—the remaining Create Dataset will be for the volume.

Deduplication

Deduplication is the process of ZFS transparently reusing a single copy of duplicated data to save space. Depending on the amount of duplicate data, deduplication can improve storage capacity, as less data is written and stored. However, deduplication is RAM intensive. A general rule of thumb is 5 GB of RAM per terabyte of deduplicated storage. In most cases, compression provides storage gains comparable to deduplication with less impact on performance.

In FreeNAS®, deduplication can be enabled during dataset creation. Be forewarned that there is no way to undedup the data within a dataset once deduplication is enabled, as disabling deduplication has NO EFFECT on existing data. The more data written to a deduplicated dataset, the more RAM it requires. When the system starts storing the DDTs (dedup tables) on disk because they no longer fit into RAM, performance craters. Further, importing an unclean pool can require between 3-5 GB of RAM per terabyte of deduped data, and if the system does not have the needed RAM, it will panic. The only solution is to add more RAM or recreate the pool. Think carefully before enabling dedup! This [article](http://constantin.glez.de/blog/2011/07/zfs-dedupe-or-not-dedupe) (<http://constantin.glez.de/blog/2011/07/zfs-dedupe-or-not-dedupe>) provides a good description of the value versus cost considerations for deduplication.

Unless a lot of RAM and a lot of duplicate data is available, do not change the default deduplication setting of “Off”. For performance reasons, consider using compression rather than turning this option on.

If deduplication is changed to On, duplicate data blocks are removed synchronously. The result is that only unique data is stored and common components are shared among files. If deduplication is changed to Verify, ZFS will do a byte-to-

byte comparison when two blocks have the same signature to make sure that the block contents are identical. Since hash collisions are extremely rare, Verify is usually not worth the performance hit.

Note: After deduplication is enabled, the only way to disable it is to use the `zfs set dedup=off dataset_name` command from [Shell](#) (page 290). However, any data that has already been deduplicated will not be un-deduplicated. Only newly stored data after the property change will not be deduplicated. The only way to remove existing deduplicated data is to copy all of the data off of the dataset, set the property to off, then copy the data back in again. Alternately, create a new dataset with ZFS Deduplication left disabled, copy the data to the new dataset, and destroy the original dataset.

Tip: Deduplication is often considered when using a group of very similar virtual machine images. However, other features of ZFS can provide dedup-like functionality more efficiently. For example, create a dataset for a standard VM, then clone that dataset for other VMs. Only the difference between each created VM and the main dataset are saved, giving the effect of deduplication without the overhead.

Compression

When selecting a compression type, you need to balance performance with the amount of disk space saved by compression. Compression is transparent to the client and applications as ZFS automatically compresses data as it is written to a compressed dataset or zvol and automatically decompresses that data as it is read. These compression algorithms are supported:

- **lz4:** recommended compression method as it allows compressed datasets to operate at near real-time speed. This algorithm only compresses the files that will benefit from compression. By default, ZFS pools made using FreeNAS® 9.2.1 or higher use this compression method, meaning that this algorithm is used if the Compression level is left at Inherit when creating a dataset or zvol.
- **gzip:** varies from levels 1 to 9 where gzip fastest (level 1) gives the least compression and gzip maximum (level 9) provides the best compression but is discouraged due to its performance impact.
- **zle:** fast but simple algorithm to eliminate runs of zeroes.
- **lzjb:** provides decent data compression, but is considered deprecated as lz4 provides much better performance.

If you select Off as the Compression level when creating a dataset or zvol, compression will not be used on the dataset/zvol. This is not recommended as using lz4 has a negligible performance impact and allows for more storage capacity.

8.1.4 Create zvol

A zvol is a feature of ZFS that creates a raw block device over ZFS. This allows you to use a zvol as an [iSCSI](#) (page 221) device extent.

To create a zvol, select an existing ZFS volume or dataset from the tree then click Create zvol to open the screen shown in [Figure 8.5](#).

Create zvol

Create zvol on volume1

zvol name:

Comments:

Size for this zvol:

i

Force size:

☐

i

Sync:

Inherit (standard)

Compression level:

Inherit (lz4)

ZFS Deduplication:

Enabling dedup can drastically reduce performance and affect the ability to access data. Compression usually offers similar space savings with much lower performance impact and overhead.

Inherit (off)

Sparse volume:

☐

i

Add zvol

Cancel

Advanced Mode

Fig. 8.5: Creating a Zvol

The configuration options are described in [Table 8.5](#). Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking Show advanced fields by default in `System → Advanced`.

Table 8.5: zvol Configuration Options

Setting	Value	Description
zvol Name	string	mandatory; enter a name for the zvol; note that there is a 63-character limit on device path names in devfs, so using long zvol names can prevent accessing zvols as devices; for example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent
Comments	string	short comments or user notes about this zvol
Size for this zvol	integer	specify size and value such as 10Gib; if the size is more than 80% of the available capacity, the creation will fail with an “out of space” error unless Force size is checked
Force size	checkbox	by default, the system will not let you create a zvol if that operation will bring the pool to over 80% capacity; while NOT recommended, checking this box will force the creation of the zvol in this situation

Continued on next page

Table 8.5 – continued from previous page

Setting	Value	Description
Compression level	drop-down menu	see the section on Compression (page 125) for a description of the available algorithms
Sparse volume	checkbox	used to provide thin provisioning; use with caution for when this option is selected, writes will fail when the pool is low on space
Block size	drop-down menu	only available in Advanced Mode and by default is based on the number of disks in pool; can be set to match the block size of the filesystem which will be formatted onto the iSCSI target

8.1.5 Import Disk

The **Volume → Import Disk** screen, shown in [Figure 8.6](#), is used to import a single disk that has been formatted with the UFS, NTFS, MSDOS, or EXT2 filesystem. The import is meant to be a temporary measure to copy the data from a disk to an existing ZFS dataset. Only one disk can be imported at a time.

Note: Imports of EXT3 or EXT4 filesystems are possible in some cases, although neither is fully supported. EXT3 journaling is not supported, so those filesystems must have an external fsck utility, like the one provided by [E2fsprogs utilities](http://e2fsprogs.sourceforge.net/) (<http://e2fsprogs.sourceforge.net/>), run on them before import. EXT4 filesystems with extended attributes or inodes greater than 128 bytes are not supported. EXT4 filesystems with EXT3 journaling must have an fsck run on them before import, as described above.

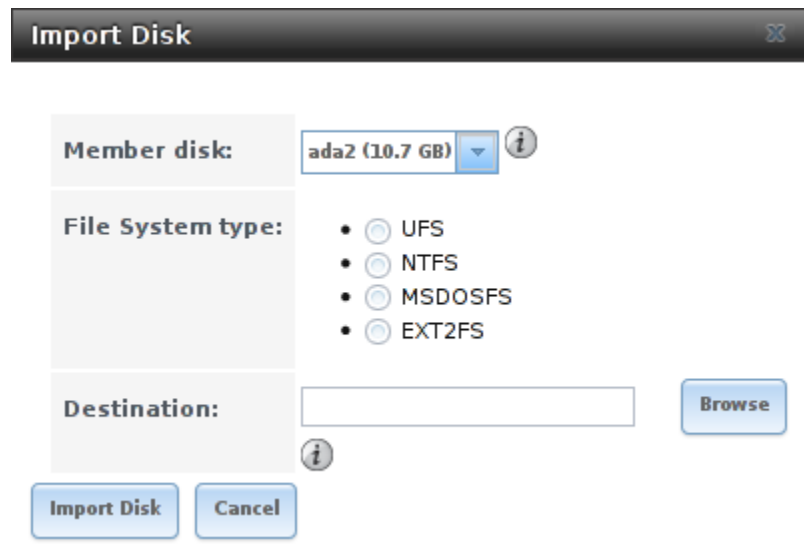


Fig. 8.6: Importing a Disk

Use the drop-down menu to select the disk to import, select the type of filesystem on the disk, and browse to the ZFS dataset that will hold the copied data. When you click **Import Volume**, the disk is mounted, its contents are copied to the specified ZFS dataset, and the disk is unmounted after the copy operation completes.

8.1.6 Import Volume

If you click **Storage → Volumes → Import Volume**, you can configure FreeNAS® to use an existing ZFS pool. This action is typically performed when an existing FreeNAS® system is re-installed. Since the operating system is separate from the storage disks, a new installation does not affect the data on the disks. However, the new operating system needs to be configured to use the existing volume.

Figure 8.7 shows the initial pop-up window that appears when you import a volume.




Fig. 8.7: Initial Import Volume Screen

If you are importing an unencrypted ZFS pool, select No: Skip to import to open the screen shown in Figure 8.8.



Fig. 8.8: Importing a Non-Encrypted Volume

Existing volumes should be available for selection from the drop-down menu. In the example shown in Figure 8.8, the FreeNAS® system has an existing, unencrypted ZFS pool. Once the volume is selected, click the OK button to import the volume.

If an existing ZFS pool does not show in the drop-down menu, run `zpool import` from [Shell](#) (page 290) to import the pool.

If you plan to physically install ZFS formatted disks from another system, be sure to export the drives on that system to prevent an “in use by another machine” error during the import.

If you suspect that your hardware is not being detected, run `camcontrol devlist` from [Shell](#) (page 290). If the disk does not appear in the output, check to see if the controller driver is supported or if it needs to be loaded using [Tunables](#) (page 68).

Importing an Encrypted Pool

Disks in existing GELI-encrypted ZFS pools must be decrypted before importing the pool. In the Import Volume dialog shown in Figure 8.7, select Yes: Decrypt disks. The screen shown in Figure 8.9 is then displayed.

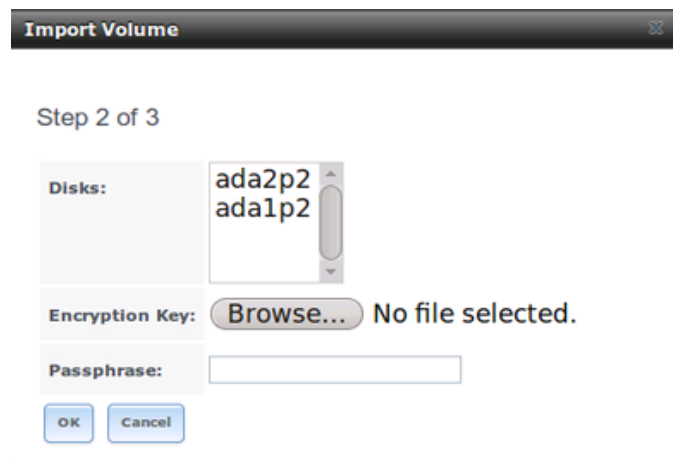


Fig. 8.9: Decrypting Disks Before Importing a ZFS Pool

Select the disks in the encrypted pool, browse to the location of the saved encryption key, enter the passphrase associated with the key, then click OK to decrypt the disks.

Note: The encryption key is required to decrypt the pool. If the pool cannot be decrypted, it cannot be re-imported after a failed upgrade or lost configuration. This means that it is very important to save a copy of the key and to remember the passphrase that was configured for the key. Refer to [Managing Encrypted Volumes](#) (page 133) for instructions on how to manage the keys for encrypted volumes.

After the pool is decrypted, it appears in the drop-down menu of [Figure 8.8](#). Click the OK button to finish the volume import.

Note: For security reasons, GELI keys for encrypted volumes are not saved in a configuration backup file. When FreeNAS® has been installed to a new device and a saved configuration file restored to it, the GELI keys for encrypted disks will not be present, and the system will not request them. To correct this, export the encrypted volume with Detach Volume, making sure that the checkboxes which clear data are not selected (there are no check marks on Mark the disks as new (destroy data) or Also delete the share's configuration). Then import the volume again. During the import, the GELI keys can be entered as described above.

8.1.7 View Disks

Storage → Volumes → View Disks shows all of the disks recognized by the FreeNAS® system. An example is shown in [Figure 8.10](#).

View Disks									
Name	Serial	Disk Size	Description	Transfer Mode	HDD Standby	Advanced Power Management	Acoustic Level	Enable S.M.A.R.T.	S.M.A.R.T. extra options
ada0	JP2940HZ3SNPDC	1.0 TB		Auto	Always On	Disabled	Disabled	true	
ada1	JP2940HZ3SN61C	1.0 TB		Auto	Always On	Disabled	Disabled	true	
ada2	JP2940HZ3SNPVC	1.0 TB		Auto	Always On	Disabled	Disabled	true	
ada3	JP2940HZ3SK5VC	1.0 TB		Auto	Always On	Disabled	Disabled	true	
da1		0		Auto	Always On	Disabled	Disabled	true	
da2		0		Auto	Always On	Disabled	Disabled	true	

Edit
Wipe

Fig. 8.10: Viewing Disks

The current configuration of each device is displayed. Click a disk entry and the Edit button to change its configuration. The configurable options are described in [Table 8.6](#).

Table 8.6: Disk Options

Setting	Value	Description
Name	string	read-only value showing FreeBSD device name for disk
Serial	string	read-only value showing the disk's serial number
Description	string	optional
HDD Standby	drop-down menu	indicates the time of inactivity (in minutes) before the drive enters standby mode in order to conserve energy; this forum post (https://forums.freenas.org/index.php?threads/how-to-find-out-if-a-drive-is-spinning-down-properly.2068/) demonstrates how to determine if a drive has spun down
Advanced Power Management	drop-down menu	default is Disabled, can select a power management profile from the menu
Acoustic Level	drop-down menu	default is Disabled; can be modified for disks that understand AAM (https://en.wikipedia.org/wiki/Automatic_acoustic_management)
Enable S.M.A.R.T.	checkbox	enabled by default if the disk supports S.M.A.R.T.; unchecking this box will disable any configured S.M.A.R.T. Tests (page 99) for the disk
S.M.A.R.T. extra options	string	additional smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools) options

Note: If a disk's serial number is not displayed in this screen, use the `smartctl` command from [Shell](#) (page 290). For

example, to determine the serial number of disk `ada0`, type `smartctl -a /dev/ada0 | grep Serial`.

The Wipe function is provided for when an unused disk is to be discarded.

Warning: Make certain that all data has been backed up and that the disk is no longer in use. Triple-check that the correct disk is being selected to be wiped, as recovering data from a wiped disk is usually impossible. If there is any doubt, physically remove the disk, verify that all data is still present on the FreeNAS® system, and wipe the disk in a separate computer.

Clicking Wipe offers several choices. Quick erases only the partitioning information on a disk, making it easy to reuse but without clearing other old data. For more security, Full with zeros overwrites the entire disk with zeros, while Full with random data overwrites the entire disk with random binary data.

Quick wipes take only a few seconds. A Full with zeros wipe of a large disk can take several hours, and a Full with random data takes longer. A progress bar is displayed during the wipe to track status.

8.1.8 Volumes

`Storage` → `Volumes` is used to view and further configure existing ZFS pools, datasets, and zvols. The example shown in [Figure 8.11](#) shows one ZFS pool (volume1) with two datasets (the one automatically created with the pool, volume1, and dataset1) and one zvol (zvol1).

Note that in this example, there are two datasets named volume1. The first represents the ZFS pool and its Used and Available entries reflect the total size of the pool, including disk parity. The second represents the implicit or root dataset and its Used and Available entries indicate the amount of disk space available for storage.

Buttons are provided for quick access to Volume Manager, Import Disk, Import Volume, and View Disks. If the system has multipath-capable hardware, an extra button will be added, View Multipaths. For each entry, the columns indicate the Name, how much disk space is Used, how much disk space is Available, the type of Compression, the Compression Ratio, the Status, whether it is mounted as read-only, and any Comments entered for the volume.

Storage							
Volumes Periodic Snapshot Tasks Replication Tasks Resilver Priority Scrubs Snapshots VMware-Snapshot							
Volume Manager Import Disk Import Volume View Disks							
Name	Used	Available	Compression	Compression Ratio	Status	Readonly	Comments
▲ volume1	2.7 MiB (0%)	7.9 GiB	-	-	HEALTHY		
volume1	1.1 MiB (0%)	7.7 GiB	lz4	3.08x	-	inherit (off)	

Fig. 8.11: Viewing Volumes

Clicking the entry for a pool causes several buttons to appear at the bottom of the screen.

Detach Volume: allows exporting the pool or deleting the contents of the pool, depending upon the choice made in the screen shown in [Figure 8.12](#). The Detach Volume screen displays the current used space and indicates whether there are any shares, provides checkboxes to Mark the disks as new (destroy data) and to Also delete the share's configuration, and asks if you are sure about doing this. The browser window turns red to indicate that some choices will make the data inaccessible. When the box to mark the disks as new is left unchecked, the volume is exported. The data is not destroyed and the volume can be re-imported at a later time. When moving a ZFS pool from one system to another, perform this export action first as it flushes any unwritten data to disk, writes data to the disk indicating that the export was done, and removes all knowledge of the pool from the system.

When the box to mark the disks as new is checked, the pool and all the data in its datasets, zvols, and shares is destroyed and the individual disks are returned to their raw state. Desired data must be backed up to another disk or device before using this option.



Fig. 8.12: Detach or Delete a Volume

Scrub Volume: scrubs and scheduling them are described in more detail in [Scrubs](#) (page 152). This button allows manually initiating a scrub. Scrubs are I/O intensive and can negatively impact performance. Avoid initiating a scrub when the system is busy.

A Cancel button is provided to cancel a scrub. When a scrub is cancelled, it is abandoned. The next scrub to run starts from the beginning, not where the cancelled scrub left off.

The status of a running scrub or the statistics from the last completed scrub can be seen by clicking the Volume Status button.

Volume Status: as shown in the example in [Figure 8.13](#), this screen shows the device name and status of each disk in the ZFS pool as well as any read, write, or checksum errors. It also indicates the status of the latest ZFS scrub. Clicking the entry for a device causes buttons to appear to edit the device's options (shown in [Figure 8.14](#)), offline or online the device, or replace the device (as described in [Replacing a Failed Drive](#) (page 137)).

Upgrade: used to upgrade the pool to the latest ZFS features, as described in [Upgrading a ZFS Pool](#) (page 27). This button does not appear if the pool is running the latest version of feature flags.

Volume Status

Scrub
Status: Completed
Errors: 0 Repaired: 0 Date: Mon Oct 16 13:10:08 2017

Name	Read	Write	Checksum	Status
volume1	0	0	0	ONLINE
raidz1-0	0	0	0	ONLINE
ada3p2	0	0	0	ONLINE
ada2p2	0	0	0	ONLINE
ada1p2	0	0	0	ONLINE

Fig. 8.13: Volume Status

Selecting a disk in Volume Status and clicking its Edit Disk button shows the screen in [Figure 8.14](#). [Table 8.6](#) summarizes the configurable options.

Name:	ada0
Serial:	JP2940HZ3SNPDC
Description:	
HDD Standby:	Always On
Advanced Power Management:	Disabled
Acoustic Level:	Disabled
Enable S.M.A.R.T.	<input checked="" type="checkbox"/>
S.M.A.R.T. extra options:	

OK Cancel

Fig. 8.14: Editing a Disk

Note: Versions of FreeNAS® prior to 8.3.1 required a reboot to apply changes to the HDD Standby, Advanced Power Management, and Acoustic Level settings. As of 8.3.1, changes to these settings are applied immediately.

Clicking a dataset in **Storage** → **Volumes** causes buttons to appear at the bottom of the screen, providing these options:

Change Permissions: edit the dataset's permissions as described in [Change Permissions](#) (page 120).

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use [Periodic Snapshot Tasks](#) (page 140).

Promote Dataset: only applies to clones. When a clone is promoted, the origin filesystem becomes a clone of the clone making it possible to destroy the filesystem that the clone was created from. Otherwise, a clone cannot be destroyed while the origin filesystem exists.

Destroy Dataset: clicking the Destroy Dataset button causes the browser window to turn red to indicate that this is a destructive action. The Destroy Dataset screen forces you to check the box I'm aware this will destroy all child datasets and snapshots within this dataset before it will perform this action.

Edit Options: edit the volume's properties described in [Table 8.4](#). Note that it will not allow changing the dataset's name.

Create Dataset: used to create a child dataset within this dataset.

Create zvol: create a child zvol within this dataset.

Clicking a zvol in **Storage** → **Volumes** causes icons to appear at the bottom of the screen: Create Snapshot, Edit zvol, and Destroy zvol. Similar to datasets, a zvol's name cannot be changed, and destroying a zvol requires confirmation.

Managing Encrypted Volumes

FreeNAS® generates and stores a randomized encryption key whenever a new encrypted volume is created. This key is required to read and decrypt any data on the volume.

Encryption keys can also be downloaded as a safety measure, to allow decryption on a different system in the event of failure, or to allow the locally stored key to be deleted for extra security. Encryption keys can also be optionally protected with a passphrase for additional security. The combination of encryption key location and whether a passphrase is used provide several different security scenarios:

- Key stored locally, no passphrase: the encrypted volume is decrypted and accessible when the system running. Protects “data at rest” only.
- Key stored locally, with passphrase: the encrypted volume is not accessible until the passphrase is entered by the FreeNAS® administrator.
- Key not stored locally: the encrypted volume is not accessible until the FreeNAS® administrator provides the key. If a passphrase is set on the key, it must also be entered before the encrypted volume can be accessed (two factor authentication (https://en.wikipedia.org/wiki/Multi-factor_authentication)).

Encrypted data cannot be accessed when the disks are removed or the system has been shut down. On a running system, encrypted data cannot be accessed when the volume is locked (see below) and the key is not available. If the key is protected with a passphrase, both the key and passphrase are required for decryption.

Encryption applies to a volume, not individual users. When a volume is unlocked, data is accessible to all users with permissions to access it.

Note: GELI (<http://www.freebsd.org/cgi/man.cgi?query=geli>) uses two randomized encryption keys for each disk. The first has been discussed here. The second, the disk’s “master key”, is encrypted and stored in the on-disk GELI metadata. Loss of a disk master key due to disk corruption is equivalent to any other disk failure, and in a redundant pool, other disks will contain accessible copies of the uncorrupted data. While it is possible to separately back up disk master keys, it is usually not necessary or useful.

Additional Controls for Encrypted Volumes

If the Encryption box is checked during the creation of a pool, additional buttons appear in the entry for the volume in Storage → Volumes. An example is shown in Figure 8.15.




Fig. 8.15: Encryption Icons Associated with an Encrypted Volume

These additional encryption buttons are used to:

Create/Change Passphrase: set and confirm a passphrase associated with the GELI encryption key. The desired passphrase is entered and repeated for verification. A red warning is a reminder to Remember to add a new recovery key as this action invalidates the previous recovery key. Unlike a password, a passphrase can contain spaces and is typically a series of words. A good passphrase is easy to remember (like the line to a song or piece of literature) but hard to guess (people who know you should not be able to guess the passphrase). Remember this passphrase. An encrypted volume cannot be reimported without it. In other words, if the passphrase is forgotten, the data on the volume can become inaccessible if it becomes necessary to reimport the pool. Protect this passphrase, as anyone who knows it could reimport the encrypted volume, thwarting the reason for encrypting the disks in the first place.


A dialog box titled "Create Passphrase" with a close button in the top right corner. Inside the dialog, there is a red warning message: "Remember to add a new recovery key as this action invalidates the previous recovery key". Below the warning, there are two input fields. The first is labeled "Passphrase:" and the second is labeled "Confirm Passphrase:". Both fields contain a series of dots, indicating that the text is masked. At the bottom of the dialog, there are two buttons: "OK" and "Cancel".

Fig. 8.16: Add or Change a Passphrase to an Encrypted Volume

After the passphrase is set, the name of this button changes to Change Passphrase. After setting or changing the passphrase, it is important to immediately create a new recovery key by clicking the Add recovery key button. This way, if the passphrase is forgotten, the associated recovery key can be used instead.

Encrypted volumes with a passphrase display an additional lock button:



Fig. 8.17: Lock Button

These encrypted volumes can be locked. The data is not accessible until the volume is unlocked by supplying the passphrase or encryption key, and the button changes to an unlock button:



Fig. 8.18: Unlock Button

To unlock the volume, click the unlock button to display the Unlock dialog:




Fig. 8.19: Unlock Locked Volume

Unlock the volume by entering a passphrase or using the Browse button to load the recovery key. If both a passphrase and a recovery key are entered, only the passphrase is used. By default, the services listed will restart when the volume is unlocked. This allows them to see the new volume and share or access data on it. Individual services can be prevented from restarting by unchecking them. However, a service that is not restarted might not be able to access the unlocked volume.

Download Key: download a backup copy of the GELI encryption key. The encryption key is saved to the client system, not on the FreeNAS® system. The FreeNAS® administrative password must be entered, then the directory in which to store the key is chosen. Since the GELI encryption key is separate from the FreeNAS® configuration database, it is highly recommended to make a backup of the key. If the key is ever lost or destroyed and there is no backup key, the data on the disks is inaccessible.

Encryption Re-key: generate a new GELI encryption key. Typically this is only performed when the administrator suspects that the current key may be compromised. This action also removes the current passphrase.

Add recovery key: generate a new recovery key. This screen prompts for the FreeNAS® administrative password and then the directory in which to save the key. Note that the recovery key is saved to the client system, not on the FreeNAS® system. This recovery key can be used if the passphrase is forgotten. Always immediately add a recovery key whenever the passphrase is changed.

Remove recovery key: Typically this is only performed when the administrator suspects that the current recovery key may be compromised. Immediately create a new passphrase and recovery key.

Note: The passphrase, recovery key, and encryption key must be protected. Do not reveal the passphrase to others. On the system containing the downloaded keys, take care that the system and its backups are protected. Anyone who has the keys has the ability to re-import the disks if they are discarded or stolen.

Warning: If a re-key fails on a multi-disk system, an alert is generated. Do not ignore this alert as doing so may result in the loss of data.

8.1.9 View Multipaths

FreeNAS® uses `gmultipath(8)` (<http://www.freebsd.org/cgi/man.cgi?query=gmultipath>) to provide `multipath I/O` (https://en.wikipedia.org/wiki/Multipath_I/O) support on systems containing hardware that is capable of multipath.

An example would be a dual SAS expander backplane in the chassis or an external JBOD.

Multipath hardware adds fault tolerance to a NAS as the data is still available even if one disk I/O path has a failure.

FreeNAS® automatically detects active/active and active/passive multipath-capable hardware. Any multipath-capable devices that are detected will be placed in multipath units with the parent devices hidden. The configuration will be displayed in *Storage → Volumes → View Multipaths*. Note that this option is not displayed in the *Storage → Volumes* tree on systems that do not contain multipath-capable hardware.

8.1.10 Replacing a Failed Drive

With any form of redundant RAID, failed drives must be replaced as soon as possible to repair the degraded state of the RAID. Depending on the hardware's capabilities, it might be necessary to reboot to replace the failed drive. Hardware that supports AHCI does not require a reboot.

Note: Striping (RAID0) does not provide redundancy. If a disk in a stripe fails, the volume will be destroyed and must be recreated and the data restored from backup.

Note: If the volume is encrypted with GELI, refer to [Replacing an Encrypted Drive](#) (page 139) before proceeding.

Before physically removing the failed device, go to *Storage → Volumes*. Select the volume's name. At the bottom of the interface are several icons, one of which is Volume Status. Click the Volume Status icon and locate the failed disk. Then perform these steps:

1. Click the disk's entry, then its Offline button to change the disk status to OFFLINE. This step is needed to properly remove the device from the ZFS pool and to prevent swap issues. If the hardware supports hot-pluggable disks, click the disk's Offline button and pull the disk, then skip to step 3. If there is no Offline button but only a Replace button, the disk is already offlined and this step can be skipped.

Note: If the process of changing the disk's status to OFFLINE fails with a "disk offline failed - no valid replicas" message, the ZFS volume must be scrubbed first with the Scrub Volume button in *Storage → Volumes*. After the scrub completes, try to Offline the disk again before proceeding.

2. If the hardware is not AHCI capable, shut down the system to physically replace the disk. When finished, return to the GUI and locate the OFFLINE disk.
3. After the disk has been replaced and is showing as OFFLINE, click the disk again and then click its Replace button. Select the replacement disk from the drop-down menu and click the Replace Disk button. After clicking the Replace Disk button, the ZFS pool begins resilvering.
4. After the drive replacement process is complete, re-add the replaced disk in the [S.M.A.R.T. Tests](#) (page 99) screen.

In the example shown in [Figure 8.20](#), a failed disk is being replaced by disk `ada5` in the volume named `volume1`.




Fig. 8.20: Replacing a Failed Disk

After the resilver is complete, Volume Status shows a Completed resilver status and indicates any errors. [Figure 8.21](#) indicates that the disk replacement was successful in this example.

Note: A disk that is failing but has not completely failed can be replaced in place, without first removing it. Whether this is a good idea depends on the overall condition of the failing disk. A disk with a few newly-bad blocks that is otherwise functional can be left in place during the replacement to provide data redundancy. A drive that is experiencing continuous errors can actually slow down the replacement. In extreme cases, a disk with serious problems might spend so much time retrying failures that it could prevent the replacement resilvering from completing before another drive fails.

Volume Status				
Resilver				
Status: Completed				
Errors: 0 Date: Fri Aug 29 11:22:39 2014				
Name	Read	Write	Checksum	Status
▲ volume1	0	0	0	ONLINE
▲ raidz1-0	0	0	0	ONLINE
ada4p2	0	0	0	ONLINE
ada3p2	0	0	0	ONLINE
ada2p2	0	0	0	ONLINE
ada5p2	0	0	0	ONLINE

Fig. 8.21: Disk Replacement is Complete

Replacing an Encrypted Drive

If the ZFS pool is encrypted, additional steps are needed when replacing a failed drive.

First, make sure that a passphrase has been set using the instructions in [Encryption](#) (page 117) before attempting to replace the failed drive. Then, follow the steps 1 and 2 as described above. During step 3, you will be prompted to input and confirm the passphrase for the pool. Enter this information then click the Replace Disk button. Wait until the resilvering is complete.

Next, restore the encryption keys to the pool. If the following additional steps are not performed before the next reboot, access to the pool might be permanently lost.

1. Highlight the pool that contains the disk that was just replaced and click the Encryption Re-key button in the GUI. Entry of the root password will be required.
2. Highlight the pool that contains the disk you just replaced and click Create Passphrase and enter the new passphrase. The old passphrase can be reused if desired.
3. Highlight the pool that contains the disk you just replaced and click the Download Key button to save the new encryption key. Since the old key will no longer function, any old keys can be safely discarded.
4. Highlight the pool that contains the disk that was just replaced and click the Add Recovery Key button to save the new recovery key. The old recovery key will no longer function, so it can be safely discarded.

Removing a Log or Cache Device

Added log or cache devices appear in `Storage → Volumes → Volume Status`. Clicking the device enables its Replace and Remove buttons.

Log and cache devices can be safely removed or replaced with these buttons. Both types of devices improve performance, and throughput can be impacted by their removal.

8.1.11 Replacing Drives to Grow a ZFS Pool

The recommended method for expanding the size of a ZFS pool is to pre-plan the number of disks in a vdev and to stripe additional vdevs using [Volume Manager](#) (page 115) as additional capacity is needed.

However, this is not an option if there are no open drive ports and a SAS/SATA HBA card cannot be added. In this case, one disk at a time can be replaced with a larger disk, waiting for the resilvering process to incorporate the new disk into the pool, then repeating with another disk until all of the original disks have been replaced.

The safest way to perform this is to use a spare drive port or an eSATA port and a hard drive dock. The process follows these steps:

1. Shut down the system.
2. Install one new disk.
3. Start up the system.
4. Go to `Storage` → `Volumes`, select the pool to expand and click the `Volume Status` button. Select a disk and click the `Replace` button. Choose the new disk as the replacement.
5. The status of the resilver process can be viewed by running `zpool status`. When the new disk has resilvered, the old one will be automatically offlined. The system is then shut down to physically remove the replaced disk. One advantage of this approach is that there is no loss of redundancy during the resilver.

If a spare drive port is not available, a drive can be replaced with a larger one using the instructions in [Replacing a Failed Drive](#) (page 137). This process is slow and places the system in a degraded state. Since a failure at this point could be disastrous, do not attempt this method unless the system has a reliable backup. Replace one drive at a time and wait for the resilver process to complete on the replaced drive before replacing the next drive. After all the drives are replaced and the final resilver completes, the added space will appear in the pool.

8.1.12 Hot Spares

ZFS provides the ability to have “hot” spares. These are drives that are connected to a volume, but not in use. If the volume experiences the failure of a data drive, the system uses the hot spare as a temporary replacement. If the failed drive is replaced with a new drive, the hot spare drive is no longer needed and reverts to being a hot spare. If the failed drive is instead removed from the volume, the spare is promoted to a full member of the volume.

Hot spares can be added to a volume during or after creation. On FreeNAS®, hot spare actions are implemented by `zfsd(8)` (<https://www.freebsd.org/cgi/man.cgi?query=zfsd>).

8.2 Periodic Snapshot Tasks

A periodic snapshot task allows scheduling the creation of read-only versions of ZFS volumes and datasets at a given point in time. Snapshots can be created quickly and, if little data changes, new snapshots take up very little space. For example, a snapshot where no files have changed takes 0 MB of storage, but as changes are made to files, the snapshot size changes to reflect the size of the changes.

Snapshots provide a clever way of keeping a history of files, providing a way to recover an older copy or even a deleted file. For this reason, many administrators take snapshots often (perhaps every fifteen minutes), store them for a period of time (possibly a month), and store them on another system (typically using [Replication Tasks](#) (page 142)). Such a strategy allows the administrator to roll the system back to a specific point in time. If there is a catastrophic loss, an off-site snapshot can be used to restore the system up to the time of the last snapshot.

An existing ZFS volume is required before creating a snapshot. Creating a volume is described in [Volume Manager](#) (page 115).

To create a periodic snapshot task, click `Storage` → `Periodic Snapshot Tasks` → `Add Periodic Snapshot` which opens the screen shown in [Figure 8.22](#). [Table 8.7](#) summarizes the fields in this screen.

Note: If only a one-time snapshot is needed, instead use `Storage` → `Volumes` and click the `Create Snapshot` button for the volume or dataset to snapshot.

Periodic Snapshots

Volume/Dataset: volume1

Recursive: ☐

Snapshot Lifetime 2 **Week(s)** ⓘ

Begin: 09:00:00 ⓘ

End: 18:00:00 ⓘ

Interval: 1 hour ⓘ

Weekday:

- ☒ Monday
- ☒ Tuesday
- ☒ Wednesday
- ☒ Thursday
- ☒ Friday
- ☐ Saturday
- ☐ Sunday

Enabled: ☒

OK **Cancel**

Fig. 8.22: Creating a Periodic Snapshot

Table 8.7: Options When Creating a Periodic Snapshot

Setting	Value	Description
Volume/Dataset	drop-down menu	select an existing ZFS volume, dataset, or zvol
Recursive	checkbox	select this box to take separate snapshots of the volume/dataset and each of its child datasets; if unchecked, a single snapshot is taken of only the specified volume/dataset, but not any child datasets
Snapshot Life-time	integer and drop-down menu	length of time to retain the snapshot on this system; if the snapshot is replicated, it is not removed from the receiving system when the lifetime expires
Begin	drop-down menu	do not create snapshots before this time of day
End	drop-down menu	do not create snapshots after this time of day
Interval	drop-down menu	how often to take snapshot between Begin and End times
Weekday	checkboxes	which days of the week to take snapshots
Enabled	checkbox	uncheck to disable the scheduled snapshot task without deleting it

If the Recursive box is checked, child datasets of this dataset are included in the snapshot and there is no need to create snapshots for each child dataset. The downside is that there is no way to exclude particular child datasets from a recursive snapshot.

Click the OK button to save the task. Entries for each task are shown in View Periodic Snapshot Tasks. Click an entry to

display Edit and Delete buttons for it.

8.3 Replication Tasks

Replication is the duplication of snapshots from one FreeNAS® system to another computer. When a new snapshot is created on the source computer, it is automatically replicated to the destination computer. Replication is typically used to keep a copy of files on a separate system, with that system sometimes being at a different physical location.

The basic configuration requires a source system with the original data and a destination system where the data will be replicated. The destination system is prepared to receive replicated data, a [periodic snapshot](#) (page 140) of the data on the source system is created, and then a replication task is created. As snapshots are automatically created on the source computer, they are automatically replicated to the destination computer.

Note: Replicated data is not visible on the receiving system until the replication task completes.

Note: The target dataset on the receiving system is automatically created in read-only mode to protect the data. To mount or browse the data on the receiving system, create a clone of the snapshot and use the clone. Clones are created in read/write mode, making it possible to browse or mount them. See [Snapshots](#) (page 155) for more information on creating clones.

8.3.1 Examples: Common Configuration

The examples shown here use the same setup of source and destination computers.

Alpha (Source)

Alpha is the source computer with the data to be replicated. It is at IP address 10.0.0.102. A [volume](#) (page 115) named alphavol has already been created, and a [dataset](#) (page 122) named alphadata has been created on that volume. This dataset contains the files which will be snapshotted and replicated onto Beta.

This new dataset has been created for this example, but a new dataset is not required. Most users will already have datasets containing the data they wish to replicate.

Create a periodic snapshot of the source dataset by selecting `Storage → Periodic Snapshot Tasks`. Click the alphavol/alphadata dataset to highlight it. Create a [periodic snapshot](#) (page 140) of it by clicking Periodic Snapshot Tasks, then Add Periodic Snapshot as shown in [Figure 8.23](#).

This example creates a snapshot of the alphavol/alphadata dataset every two hours from Monday through Friday between the hours of 9:00 and 18:00 (6:00 PM). Snapshots are automatically deleted after their chosen lifetime of two weeks expires.

Fig. 8.23: Create a Periodic Snapshot for Replication

Beta (Destination)

Beta is the destination computer where the replicated data will be copied. It is at IP address 10.0.0.118. A [volume](#) (page 115) named betavol has already been created.

Snapshots are transferred with [SSH](#) (page 236). To allow incoming connections, this service is enabled on Beta. The service is not required for outgoing connections, and so does not need to be enabled on Alpha.

8.3.2 Example: FreeNAS® to FreeNAS® Semi-Automatic Setup

FreeNAS® offers a special semi-automatic setup mode that simplifies setting up replication. Create the replication task on Alpha by clicking Replication Tasks and Add Replication. alphavol/alphadata is selected as the dataset to replicate. betavol is the destination volume where alphadata snapshots are replicated. The Setup mode dropdown is set to Semi-automatic as shown in [Figure 8.24](#). The IP address of Beta is entered in the Remote hostname field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP → HTTPS Redirect has been enabled in [System](#) → [General](#) on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

Add Replication

Volume/Dataset:

alphavol/alphadata

Remote ZFS Volume/Dataset:

betavol

Recursively replicate child dataset's snapshots:

☐

Delete stale snapshots on remote system:

☐

Replication Stream Compression:

lz4 (fastest)

Limit (kB/s):

0

Begin:

00:00:00

End:

23:59:00

Enabled:

☒

Setup mode:

Semi-automatic

This method only works with remote version greater or equal than 9.10.2

Remote hostname:

10.0.0.118

Remote HTTP/HTTPS Port:

80

Remote HTTPS:

☐

Remote Auth Token:

On the remote host go to Storage -> Replication Tasks, click the Temporary Auth Token button and paste the resulting value in to this field.

Dedicated User Enabled:

☐

Dedicated User:

Encryption Cipher:

Standard

OK

Cancel

Fig. 8.24: Add Replication Dialog, Semi-Automatic

The Remote Auth Token field expects a special token from the Beta computer. On Beta, choose `Storage` → `Replication Tasks`, then click `Temporary Auth Token`. A dialog showing the temporary authorization token is shown as in [Figure 8.25](#).

Highlight the temporary authorization token string with the mouse and copy it.




Fig. 8.25: Temporary Authentication Token on Destination

On the Alpha system, paste the copied temporary authorization token string into the Remote Auth Token field as shown in Figure 8.26.




Fig. 8.26: Temporary Authentication Token Pasted to Source

Finally, click the OK button to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See [Limiting Replication Times](#) (page 150) for information about restricting when replication is allowed to run.

Note: The temporary authorization token is only valid for a few minutes. If a Token is invalid message is shown, get a new temporary authorization token from the destination system, clear the Remote Auth Token field, and paste in the new one.

8.3.3 Example: FreeNAS® to FreeNAS® Dedicated User Replication

A dedicated user can be used for replications rather than the root user. This example shows the process using the semi-automatic replication setup between two FreeNAS® systems with a dedicated user named repluser. SSH key authentication is used to allow the user to log in remotely without a password.

In this example, the periodic snapshot task has not been created yet. If the periodic snapshot shown in the [example configuration](#) (page 142) has already been created, go to **Storage → Periodic Snapshot Tasks**, click on the task to select it, and click Delete to remove it before continuing.

On Alpha, select **Account → Users**. Click the Add User. Enter repluser for Username, enter /mnt/alphavol/repluser in the Create Home Directory In field, enter Replication Dedicated User for the Full Name, and set the Disable password login checkbox. Leave the other fields at their default values, but note the User ID number. Click OK to create the user.

On Beta, the same dedicated user must be created as was created on the sending computer. Select **Account → Users**. Click the Add User. Enter the User ID number from Alpha, repluser for Username, enter /mnt/betavol/repluser in the Create Home Directory In field, enter Replication Dedicated User for the Full Name, and set the Disable password login checkbox. Leave the other fields at their default values. Click OK to create the user.

A dataset with the same name as the original must be created on the destination computer, Beta. Select **Storage → Volumes**, click on betavol, then click the Create Dataset icon at the bottom. Enter alphadata as the Dataset Name, then click Add Dataset.

The replication user must be given permissions to the destination dataset. Still on Beta, open a [Shell](#) (page 290) and enter this command:

```
zfs allow -ldu repluser create,destroy,diff,mount,readonly,receive,release,send,userprop betavol/
↪ alphadata
```

The destination dataset must also be set to read-only. Enter this command in the [Shell](#) (page 290):

```
zfs set readonly=on betavol/alphadata
```

Close the [Shell](#) (page 290) by typing `exit` and pressing `Enter`.

The replication user must also be able to mount datasets. Still on Beta, go to `System` → `Tunables`. Click `Add Tunable`. Enter `vfs.usermount` for the Variable, 1 for the Value, and choose `Sysctl` from the Type drop-down. Click `OK` to save the tunable settings.

Back on Alpha, create a periodic snapshot of the source dataset by selecting `Storage` → `Periodic Snapshot Tasks`. Click the `alphavol/alphadata` dataset to highlight it. Create a [periodic snapshot](#) (page 140) of it by clicking `Periodic Snapshot Tasks`, then `Add Periodic Snapshot` as shown in [Figure 8.23](#).

Still on Alpha, create the replication task by clicking `Replication Tasks` and `Add Replication`. `alphavol/alphadata` is selected as the dataset to replicate. `betavol/alphadata` is the destination volume and dataset where `alphadata` snapshots are replicated.

The Setup mode dropdown is set to `Semi-automatic` as shown in [Figure 8.24](#). The IP address of Beta is entered in the Remote hostname field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP → HTTPS Redirect has been enabled in `System` → `General` on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

The Remote Auth Token field expects a special token from the Beta computer. On Beta, choose `Storage` → `Replication Tasks`, then click `Temporary Auth Token`. A dialog showing the temporary authorization token is shown as in [Figure 8.25](#).

Highlight the temporary authorization token string with the mouse and copy it.

On the Alpha system, paste the copied temporary authorization token string into the Remote Auth Token field as shown in [Figure 8.26](#).

Set the `Dedicated User` checkbox. Choose `repluser` in the `Dedicated User` drop-down.

Click the `OK` button to create the replication task.

Note: The temporary authorization token is only valid for a few minutes. If a Token is invalid message is shown, get a new temporary authorization token from the destination system, clear the Remote Auth Token field, and paste in the new one.

Replication will begin when the periodic snapshot task runs.

Additional replications can use the same dedicated user that has already been set up. The permissions and read only settings made through the [Shell](#) (page 290) must be set on each new destination dataset.

8.3.4 Example: FreeNAS® to FreeNAS® or Other Systems, Manual Setup

This example uses the same basic configuration of source and destination computers shown above, but the destination computer is not required to be a FreeNAS® system. Other operating systems can receive the replication if they support SSH, ZFS, and the same features that are in use on the source system. The details of creating volumes and datasets, enabling SSH, and copying encryption keys will vary when the destination computer is not a FreeNAS® system.

Encryption Keys

A public encryption key must be copied from Alpha to Beta to allow a secure connection without a password prompt. On Alpha, select `Storage` → `Replication Tasks` → `View Public Key`, producing the window shown in [Figure 8.27](#). Use the mouse to highlight the key data shown in the window, then copy it.



Fig. 8.27: Copy the Replication Key

On Beta, select **Account** → **Users** → **View Users**. Click the root account to select it, then click **Modify User**. Paste the copied key into the **SSH Public Key** field and click **OK** as shown in Figure 8.28.




Fig. 8.28: Paste the Replication Key

Back on Alpha, create the replication task by clicking **Replication Tasks** and **Add Replication**. `alphavol/alphadata` is selected as the dataset to replicate. The destination volume is `betavol`. The `alphadata` dataset and snapshots are replicated there. The IP address of Beta is entered in the **Remote hostname** field as shown in Figure 8.29. A hostname can be entered here if local DNS resolves for that hostname.

Click the **SSH Key Scan** button to retrieve the SSH host keys from Beta and fill the **Remote hostkey** field. Finally, click **OK** to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See [Limiting Replication Times](#) (page 150) for information about restricting when replication is allowed to run.

Add Replication

Volume/Dataset:

alphavol/alphadata

Remote ZFS Volume/Dataset:

betavol

Recursively replicate child dataset's snapshots:

Delete stale snapshots on remote system:

Replication Stream Compression:

lz4 (fastest)

Limit (kB/s):

0

Begin:

00:00:00

End:

23:59:00

Enabled:

Setup mode:

Manual

Remote hostname:

10.0.0.118

Remote port:

22

Dedicated User Enabled:

Dedicated User:

Encryption Cipher:

Standard

Remote hostkey:

10.0.0.118 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQCA4WnS+kfJa
CDL1SnPWEqHwuVjE0k8pl+kU8JlS8yyfOALP1/aB
c82DdZoNGwtJjn14xTyxA1XJKXio1YYkTnTiLj7M
R+S905HLt+vwSUhkfs3EdD8/oOCFmeiw
/00dzjT9oiCrqqnHiL+dySqBjAE0yfoQyTGfzbsy
FYG9BZ6aLSzA+oEd7i+aJlE++n6oRCENUCopeFGF
m9gADtWwETiHxJkY292JRqhY02k7JrhyzYPSLZvL
Yy3mw0bSG1Xjf8D2xGgxs7qdi ai3r6aKl+TRA4Bi
/d8GxVAKwzJPgv
/K/aWiibmaUcVBavUbM60yaRFg9uuhn43HYMHbJa
4fE/r1
10.0.0.118 ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlz
dHAyNTYAAABBBBANGLOmMyTZl/FplaScYX
/8S/b3nvXibX
/levDCDwJecuD1ASWY5Xx+Wp8YkraJzLv9bonf1w
yc2fCL4gzFs0Ag=
10.0.0.118 ssh-ed25519
AAAAC3NzaC1lZDI1NTE5AAAAIOZtUTtc59hv90WH
7nDoD4li3GdRKaZR/V70gzT8t7GE

OK

Cancel

SSH Key Scan

8.3.5 Replication Options

Table 8.8 describes the options in the replication task dialog.

Table 8.8: Replication Task Options

Setting	Value	Description
Volume/Dataset	drop-down menu	ZFS volume or dataset on the source computer containing the snapshots to be replicated; the drop-down menu is empty if a snapshot does not already exist
Remote ZFS Volume/Dataset	string	ZFS volume on the remote or destination computer which will store the snapshots; if the destination dataset is not present, it will be created; <code>/mnt/</code> is assumed, do not include it in the path
Recursively replicate child dataset's snapshots	checkbox	when checked, also replicate snapshots of datasets that are children of the main dataset
Delete stale snapshots	checkbox	when checked, delete previous snapshots on the remote or destination computer which are no longer present on the source computer
Replication Stream Compression	drop-down menu	choices are lz4 (fastest), pigz (all rounder), plzip (best compression), or Off (no compression); selecting a compression algorithm can reduce the size of the data being replicated
Limit (kB/s)	integer	limit replication speed to the specified value in kilobytes/second; default of 0 is unlimited
Begin	drop-down menu	replication is not allowed to start before this time; times entered in the Begin and End fields set when replication can occur
End	drop-down menu	replication must start by this time; once started, replication will continue until it is finished
Enabled	checkbox	uncheck to disable the scheduled replication task without deleting it
Setup mode	drop-down menu	Manual or Semi-automatic
Remote hostname	string	IP address or DNS name of remote computer where replication is sent
Remote port	string	the port used by the SSH server on the remote or destination computer
Dedicated User Enabled	checkbox	allow a user account other than root to be used for replication
Dedicated User	drop-down menu	only available if Dedicated User Enabled is checked; select the user account to be used for replication
Encryption Cipher	drop-down menu	Standard, Fast, or Disabled
Remote hostkey	string	use the SSH Key Scan button to retrieve the public host key of the remote or destination computer and populate this field with that key

The replication task runs after a new periodic snapshot is created. The periodic snapshot and any new manual snapshots of the same dataset are replicated onto the destination computer.

When multiple replications have been created, replication tasks run serially, one after another. Completion time depends on the number and size of snapshots and the bandwidth available between the source and destination computers.

The first time a replication runs, it must duplicate data structures from the source to the destination computer. This can take much longer to complete than subsequent replications, which only send differences in data.

Warning: Snapshots record incremental changes in data. If the receiving system does not have at least one snapshot that can be used as a basis for the incremental changes in the snapshots from the sending system, there is no way to identify only the data that has changed. In this situation, the snapshots in the receiving system target dataset are removed so a complete initial copy of the new replicated data can be created.

Selecting `Storage → Replication Tasks` displays [Figure 8.30](#), the list of replication tasks. The Last snapshot sent to remote side column shows the name of the last snapshot that was successfully replicated, and Status shows the current

status of each replication task. The display is updated every five seconds, always showing the latest status.

Volume/Dataset	Last snapshot sent to remote side	Remote Hostname	Status	Remote ZFS Volume/Dataset	Delete stale snapshots on remote system	Replication Stream Compression	Limit (kB/s)	Begin	End	Enabled
volume1/smb-storage	auto-20170116.0950	beta	Succeeded	betavol	true	lz4	0	00:00:00	23:59:00	true

Fig. 8.30: Replication Task List

Note: The encryption key that was copied from the source computer (Alpha) to the destination computer (Beta) is an RSA public key located in the `/data/ssh/replication.pub` file on the source computer. The host public key used to identify the destination computer (Beta) is from the `/etc/ssh/ssh_host_rsa_key.pub` file on the destination computer.

8.3.6 Replication Encryption

The default Encryption Cipher Standard setting provides good security. Fast is less secure than Standard but can give reasonable transfer rates for devices with limited cryptographic speed. For networks where the entire path between source and destination computers is trusted, the Disabled option can be chosen to send replicated data without encryption.

8.3.7 Limiting Replication Times

The Begin and End times in a replication task make it possible to restrict when replication is allowed. These times can be set to only allow replication after business hours, or at other times when disk or network activity will not slow down other operations like snapshots or [Scrubs](#) (page 152). The default settings allow replication to occur at any time.

These times control when replication task are allowed to start, but will not stop a replication task that is already running. Once a replication task has begun, it will run until finished.

8.3.8 Troubleshooting Replication

Replication depends on SSH, disks, network, compression, and encryption to work. A failure or misconfiguration of any of these can prevent successful replication.

SSH

[SSH](#) (page 236) must be able to connect from the source system to the destination system with an encryption key. This can be tested from [Shell](#) (page 290) by making an [SSH](#) (page 236) connection from the source system to the destination system. From the previous example, this is a connection from Alpha to Beta at 10.0.0.118. Start the [Shell](#) (page 290) on the source machine (Alpha), then enter this command:

```
ssh -vv -i /data/ssh/replication 10.0.0.118
```

On the first connection, the system might say

```
No matching host key fingerprint found in DNS.
Are you sure you want to continue connecting (yes/no)?
```

Verify that this is the correct destination computer from the preceding information on the screen and type `yes`. At this point, an [SSH](#) (page 236) shell connection is open to the destination system, Beta.

If a password is requested, SSH authentication is not working. See [Figure 8.27](#) above. This key value must be present in the `/root/.ssh/authorized_keys` file on Beta, the destination computer. The `/var/log/auth.log` file can show diagnostic errors for login problems on the destination computer also.

Compression

Matching compression and decompression programs must be available on both the source and destination computers. This is not a problem when both computers are running FreeNAS®, but other operating systems might not have `lz4`, `pigz`, or `plzip` compression programs installed by default. An easy way to diagnose the problem is to set Replication Stream Compression to Off. If the replication runs, select the preferred compression method and check `/var/log/debug.log` on the FreeNAS® system for errors.

Manual Testing

On Alpha, the source computer, the `/var/log/messages` file can also show helpful messages to locate the problem.

On the source computer, Alpha, open a [Shell](#) (page 290) and manually send a single snapshot to the destination computer, Beta. The snapshot used in this example is named `auto-20161206.1110-2w`. As before, it is located in the `alphavol/alphadata` dataset. A `@` symbol separates the name of the dataset from the name of the snapshot in the command.

```
zfs send alphavol/alphadata@auto-20161206.1110-2w | ssh -i /data/ssh/replication 10.0.0.118 zfs recv ↵
↵betavol
```

If a snapshot of that name already exists on the destination computer, the system will refuse to overwrite it with the new snapshot. The existing snapshot on the destination computer can be deleted by opening a [Shell](#) (page 290) on Beta and running this command:

```
zfs destroy -R betavol/alphadata@auto-20161206.1110-2w
```

Then send the snapshot manually again. Snapshots on the destination system, Beta, can be listed from the [Shell](#) (page 290) with `zfs list -t snapshot` or by going to `Storage → Snapshots`.

Error messages here can indicate any remaining problems.

8.4 Resilver Priority

Resilvering, or the process of copying data to a replacement disk, is best completed as quickly as possible. Increasing the priority of resilvers can help them to complete more quickly. The Resilver Priority tab makes it possible to increase the priority of resilvering at times where the additional I/O or CPU usage will not affect normal usage. Select `Storage → Resilver Priority` to display the screen shown in [Figure 8.31](#). [Table 8.9](#) describes the fields on this screen.

Storage

Volumes Periodic Snapshot Tasks Replication Tasks **Resilver Priority** Scrubs Snapshots VMware-Snapshot

Enabled: ☐

Begin higher priority resilvering at this time: 6:00 PM ▼

End higher priority resilvering at this time: 9:00 AM ▼

Weekday:

- ☒ Monday
- ☒ Tuesday
- ☒ Wednesday
- ☒ Thursday
- ☒ Friday
- ☒ Saturday
- ☒ Sunday

Save

Fig. 8.31: Resilver Priority

Table 8.9: Resilver Priority Options

Setting	Value	Description
Enabled	checkbox	check to enable higher-priority resilvering
Begin higher priority resilvering at this time	drop-down	start time to begin higher-priority resilvering
End higher priority resilvering at this time	drop-down	end time to begin higher-priority resilvering
Weekday	checkboxes	use higher-priority resilvering on these days of the week

8.5 Scrubs

A scrub is the process of ZFS scanning through the data on a volume. Scrubs help to identify data integrity problems, detect silent data corruptions caused by transient hardware issues, and provide early alerts of impending disk failures. FreeNAS® makes it easy to schedule periodic automatic scrubs.

Each volume should be scrubbed at least once a month. Bit errors in critical data can be detected by ZFS, but only when that data is read. Scheduled scrubs can find bit errors in rarely-read data. The amount of time needed for a scrub is proportional to the quantity of data on the volume. Typical scrubs take several hours or longer.

The scrub process is I/O intensive and can negatively impact performance. Schedule scrubs for evenings or weekends to minimize impact to users. Make certain that scrubs and other disk-intensive activity like [S.M.A.R.T. Tests](#) (page 99) are scheduled to run on different days to avoid disk contention and extreme performance impacts.

Scrubs only check used disk space. To check unused disk space, schedule [S.M.A.R.T. Tests](#) (page 99) of Type Long Self-Test to run once or twice a month.

Scrubs are scheduled and managed with `Storage → Scrubs`.

When a volume is created, a ZFS scrub is automatically scheduled. An entry with the same volume name is added to `Storage` → `Scrubs`. A summary of this entry can be viewed with `Storage` → `Scrubs` → `View Scrubs`. [Figure 8.32](#) displays the default settings for the volume named `volume1`. In this example, the entry has been highlighted and the `Edit` button clicked to display the `Edit` screen. [Table 8.10](#) summarizes the options in this screen.

Edit

Volume: volume1

Threshold days: 35

Description:

Minute:

Every N minute

Each selected minute

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27	28	29
30	31	32	33	34	35	36	37	38	39
40	41	42	43	44	45	46	47	48	49
50	51	52	53	54	55	56	57	58	59

Hour:

Every N hour

Each selected hour

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23						

Day of month:

Every N day of month

Each selected day of month

1

Month:

January

Fig. 8.32: Viewing a Volume's Default Scrub Settings

Table 8.10: ZFS Scrub Options

Setting	Value	Description
Volume	drop-down menu	volume to be scrubbed
Threshold days	integer	prevent scrub from running for this number of days after a scrub has completed, regardless of the calendar schedule; the default is a multiple of 7 to ensure that the scrub always occurs on the same day of the week
Description	string	optional text description of scrub
Minute	slider or minute selections	if the slider is used, a scrub occurs every N minutes; if specific minutes are chosen, a scrub runs only at the selected minute values
Hour	slider or hour selections	if the slider is used, a scrub occurs every N hours; if specific hours are chosen, a scrub runs only at the selected hour values
Day of Month	slider or month selections	if the slider is used, a scrub occurs every N days; if specific days of the month are chosen, a scrub runs only on the selected days of the selected months
Month	checkboxes	a scrub occurs on the selected months
Day of week	checkboxes	a scrub occurs on the selected days; the default is Sunday to least impact users; note that this field and the Day of Month field are ORed together: setting Day of Month to 01,15 and Day of week to Thursday will cause scrubs to run on the 1st and 15th days of the month, but also on any Thursday
Enabled	checkbox	uncheck to disable the scheduled scrub without deleting it

Review the default selections and, if necessary, modify them to meet the needs of the environment. Note that the Threshold field is used to prevent scrubs from running too often, and overrides the schedule chosen in the other fields.

Scheduled crubs can be deleted with the Delete button, but this is not recommended. Scrubs can provide an early indication of disk issues before a disk failure. If a scrub is too intensive for the hardware, consider temporarily unchecking the Enabled button for the scrub until the hardware can be upgraded.

8.6 Snapshots

The Snapshots tab is used to review the listing of available snapshots. An example is shown in [Figure 8.33](#).

Note: If snapshots do not appear, check that the current time configured in [Periodic Snapshot Tasks](#) (page 140) does not conflict with the Begin, End, and Interval settings. If the snapshot was attempted but failed, an entry is added to `/var/log/messages`. This log file can be viewed in [Shell](#) (page 290).

Storage

VolumesPeriodic Snapshot TasksReplication TasksResilver PriorityScrubsSnapshotsVMware-Snapshot











	Volume/Dataset	Snapshot Name	Used	Refer	Available Actions
 ...	No filter applied				
<input type="checkbox"/>	volume1	auto-20171018.0840-2w	0	88.0 KiB	 
<input type="checkbox"/>	volume1	auto-20171018.0850-2w	0	88.0 KiB	 
<input type="checkbox"/>	volume1	auto-20171018.0900-2w	0	88.0 KiB	 
<input type="checkbox"/>	volume1	auto-20171018.0910-2w	0	88.0 KiB	  

Fig. 8.33: Viewing Available Snapshots

The listing includes the name of the volume or dataset, the name of each snapshot, and the amount of used and referenced data.

Used is the amount of space consumed by this dataset and all of its descendants. This value is checked against the dataset's quota and reservation. The space used does not include the dataset's reservation, but does take into account the reservations of any descendent datasets. The amount of space that a dataset consumes from its parent, as well as the amount of space that are freed if this dataset is recursively destroyed, is the greater of its space used and its reservation. When a snapshot is created, the space is initially shared between the snapshot and the filesystem, and possibly with previous snapshots. As the filesystem changes, space that was previously shared becomes unique to the snapshot, and is counted in the snapshot's space used. Additionally, deleting snapshots can increase the amount of space unique to (and used by) other snapshots. The amount of space used, available, or referenced does not take into account pending changes. While pending changes are generally accounted for within a few seconds, disk changes do not necessarily guarantee that the space usage information is updated immediately.

Tip: Space used by individual snapshots can be seen by running `zfs list -t snapshot` from [Shell](#) (page 290).

Refer indicates the amount of data accessible by this dataset, which may or may not be shared with other datasets in the pool. When a snapshot or clone is created, it initially references the same amount of space as the file system or snapshot it was created from, since its contents are identical.

Snapshots have icons on the right side for several actions.

Clone Snapshot prompts for the name of the clone to create. A clone is a writable copy of the snapshot. Since a clone is actually a dataset which can be mounted, it appears in the Volumes tab rather than the Snapshots tab. By default, `-clone` is added to the name of a snapshot when a clone is created.

Destroy Snapshot a pop-up message asks for confirmation. Child clones must be destroyed before their parent snapshot can be destroyed. While creating a snapshot is instantaneous, deleting a snapshot can be I/O intensive and can take a long time, especially when deduplication is enabled. In order to delete a block in a snapshot, ZFS has to walk all the allocated blocks to see if that block is used anywhere else; if it is not, it can be freed.

The most recent snapshot also has a Rollback Snapshot icon. Clicking the icon asks for confirmation before rolling back to this snapshot state. Confirming by clicking Yes causes any files that have changed since the snapshot was taken to be reverted back to their state at the time of the snapshot.

Note: Rollback is a potentially dangerous operation and causes any configured replication tasks to fail as the replication system uses the existing snapshot when doing an incremental backup. To restore the data within a snapshot, the recommended steps are:

1. Clone the desired snapshot.
2. Share the clone with the share type or service running on the FreeNAS® system.
3. After users have recovered the needed data, destroy the clone in the Active Volumes tab.

This approach does not destroy any on-disk data and has no impact on replication.

A range of snapshots can be selected with the mouse. Click on the checkbox in the left column of the first snapshot, then press and hold `Shift` and click on the checkbox for the end snapshot. This can be used to select a range of obsolete snapshots to be deleted with the Destroy icon at the bottom. Be cautious and careful when deleting ranges of snapshots.

Periodic snapshots can be configured to appear as shadow copies in newer versions of Windows Explorer, as described in [Configuring Shadow Copies](#) (page 194). Users can access the files in the shadow copy using Explorer without requiring any interaction with the FreeNAS® graphical administrative interface.

The ZFS Snapshots screen allows the creation of filters to view snapshots by selected criteria. To create a filter, click the Define filter icon (near the text No filter applied). When creating a filter:

- Select the column or leave the default of Any Column.
- Select the condition. Possible conditions are: contains (default), is, starts with, ends with, does not contain, is not, does not start with, does not end with, and is empty.
- Enter a value that meets your view criteria.

- Click the Filter button to save the filter and exit the define filter screen. Alternately, click the + button to add another filter.

When creating multiple filters, select the filter to use before leaving the define filter screen. After a filter is selected, the No filter applied text changes to Clear filter. Clicking Clear filter produces a pop-up message indicates that this removes the filter and all available snapshots are listed.

8.7 VMware-Snapshot

Storage → VMware-Snapshot allows you to coordinate ZFS snapshots when using FreeNAS® as a VMware datastore. Once this type of snapshot is created, FreeNAS® will automatically snapshot any running VMware virtual machines before taking a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore. The temporary VMware snapshots are then deleted on the VMware side but still exist in the ZFS snapshot and can be used as stable resurrection points in that snapshot. These coordinated snapshots will be listed in [Snapshots](#) (page 155).

Figure 8.34 shows the menu for adding a VMware snapshot and Table 8.11 summarizes the available options.

Fig. 8.34: Adding a VMware Snapshot

Table 8.11: VMware Snapshot Options

Setting	Value	Description
Hostname	string	IP address or hostname of VMware host; when clustering, this is the vCenter server for the cluster
Username	string	user on VMware host with enough permission to snapshot virtual machines
Password	string	password associated with Username
ZFS Filesystem	drop-down menu	the filesystem to snapshot
Datastore	drop-down menu	after entering the Hostname, Username, and Password, click Fetch Datastores to populate the menu and select the datastore with which to synchronize

DIRECTORY SERVICES

FreeNAS® supports integration with these directory services:

- [Active Directory](#) (page 158) (for Windows 2000 and higher networks)
- [LDAP](#) (page 163)
- [NIS](#) (page 165)

It also supports [Kerberos Realms](#) (page 167), [Kerberos Keytabs](#) (page 167), and the ability to add additional parameters to [Kerberos Settings](#) (page 168).

This section summarizes each of these services and their available configurations within the FreeNAS® GUI.

9.1 Active Directory

Active Directory (AD) is a service for sharing resources in a Windows network. AD can be configured on a Windows server that is running Windows Server 2000 or higher or on a Unix-like operating system that is running [Samba version 4](#) (https://wiki.samba.org/index.php/Samba4/HOWTO#Provisioning_The_Samba_Active_Directory). Since AD provides authentication and authorization services for the users in a network, it is not necessary to recreate these user accounts on the FreeNAS® system. Instead, configure the Active Directory service so that it can import the account information and imported users can be authorized to access the SMB shares on the FreeNAS® system.

Many changes and improvements have been made to Active Directory support within FreeNAS®. It is strongly recommended to update the system to the latest FreeNAS® 11.1 before attempting Active Directory integration.

Before configuring the Active Directory service, ensure name resolution is properly configured by `ping`ing the domain name of the Active Directory domain controller from [Shell](#) (page 290) on the FreeNAS® system. If the `ping` fails, check the DNS server and default gateway settings in `Network → Global Configuration` on the FreeNAS® system.

Next, add a DNS record for the FreeNAS® system on the Windows server and verify that the hostname of the FreeNAS® system can be pinged from the domain controller.

Active Directory relies on Kerberos, which is a time sensitive protocol. The time on both the FreeNAS® system and the Active Directory Domain Controller cannot be out of sync by more than a few minutes. The best way to ensure that the same time is running on both systems is to configure both systems to:

- use the same NTP server (set in `System → NTP Servers` on the FreeNAS® system)
- have the same timezone
- be set to either localtime or universal time at the BIOS level

[Figure 9.1](#) shows the screen that appears when `Directory Service → Active Directory` is chosen. [Table 9.1](#) describes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box `Show advanced fields` by default in `System → Advanced`.

The screenshot shows the 'Directory Service' configuration page in FreeNAS. The 'Active Directory' sub-tab is selected. The form includes the following fields and controls:

- Domain Name (DNS/Realm-Name):** A text input field with an information icon.
- Domain Account Name:** A text input field with an information icon.
- Domain Account Password:** A text input field with an information icon.
- AD check connectivity frequency (seconds):** A text input field with the value '60' and an information icon.
- How many recovery attempts:** A text input field with the value '10' and an information icon.
- Enable Monitoring:** A checkbox that is currently unchecked, with an information icon.
- Enable:** A checkbox that is currently unchecked.

At the bottom of the form are three buttons: 'Save', 'Advanced Mode', and 'Rebuild Directory Service Cache'.

Fig. 9.1: Configuring Active Directory

Table 9.1: Active Directory Configuration Options

Setting	Value	Advanced Mode	Description
Domain Name (DNS/Realm-Name)	string		name of Active Directory domain (example.com) or child domain (sales.example.com); this setting is mandatory and the GUI will refuse to save the settings if the domain controller for the specified domain cannot be found
Domain Account Name	string		name of the Active Directory administrator account; this setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this account name
Domain Account Password	string		password for the Active Directory administrator account; this setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this password
AD check connectivity frequency (seconds)	integer		how often to verify that Active Directory services are active
How many recovery attempts	integer		number of times to attempt reconnecting to the Active Directory server; tries forever when set to 0
Enable Monitoring	checkbox		restart Active Directory automatically if the service is disconnected
Encryption Mode	drop-down menu	✓	choices are Off, SSL, or TLS
Certificate	drop-down menu	✓	select the certificate of the Active Directory server if SSL connections are used; if a certificate does not exist yet, create a CA (page 75), then create a certificate on the Active Directory server and import it to the FreeNAS® system with Certificates (page 78)
Verbose logging	checkbox	✓	when checked, logs attempts to join the domain to <code>/var/log/messages</code>

Continued on next page

Table 9.1 – continued from previous page

Setting	Value	Advanced Mode	Description
UNIX extensions	checkbox	✓	only check this box if the AD server has been explicitly configured to map permissions for UNIX users; checking this box provides persistent UIDs and GIDs, otherwise, users/groups are mapped to the UID/GUID range configured in Samba
Allow Trusted Domains	checkbox	✓	should only be enabled if network has active do-main/forest trusts (https://technet.microsoft.com/en-us/library/cc757352(WS.10).aspx) and you need to manage files on multiple domains; use with caution as it will generate more winbindd traffic, slowing down the ability to filter through user/group information
Use Default Domain	checkbox	✓	when unchecked, the domain name is prepended to the username; if Allow Trusted Domains is checked and multiple domains use the same usernames, uncheck this box to prevent name collisions
Allow DNS updates	checkbox	✓	when unchecked, disables Samba from doing DNS updates when joining a domain
Disable Active Directory user/group cache	checkbox	✓	when checked, disables caching AD users and groups; useful if you cannot bind to a domain with a large number of users or groups
User Base	string	✓	distinguished name (DN) of the user container in Active Directory
Group Base	string	✓	distinguished name (DN) of the group container in Active Directory
Site Name	string	✓	the relative distinguished name of the site object in Active Directory
Domain Controller	string	✓	will automatically be added to the SRV record for the domain and, when multiple controllers are specified, FreeNAS® selects the closest DC which responds; use the short form of the FQDN (Example: sampleserver)
Global Catalog Server	string	✓	if the hostname of the global catalog server to use is specified, make sure it is resolvable
Kerberos Realm	drop-down menu	✓	select the realm created using the instructions in Kerberos Realms (page 167)
Kerberos Principal	drop-down menu	✓	browse to the location of the keytab created using the instructions in Kerberos Keytabs (page 167)
AD timeout	integer	✓	in seconds, increase if the AD service does not start after connecting to the domain
DNS timeout	integer	✓	in seconds, increase if AD DNS queries timeout
Idmap backend	drop-down menu and Edit	✓	select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs; see Table 9.2 for a summary of the available backends; click the Edit link to configure that backend's editable options
Windbind NSS Info	drop-down menu	✓	defines the schema to use when querying AD for user/group info; rfc2307 uses the RFC2307 schema support included in Windows 2003 R2, sfu20 is for Services For Unix 3.0 or 3.5, and sfu is for Services For Unix 2.0
SASL wrapping	drop-down menu	✓	defines how LDAP traffic is transmitted; choices are plain (plain text), sign (signed only), or seal (signed and encrypted); Windows 2000 SP3 and higher can be configured to enforce signed LDAP connections
Enable	checkbox		Enable the Active Directory service

Continued on next page

Table 9.1 – continued from previous page

Setting	Value	Advanced Mode	Description
NetBIOS name	string	✓	limited to 15 characters; automatically populated with the system's original hostname; it must be different from the Workgroup name
NetBIOS alias	string	✓	limited to 15 characters

Table 9.2 summarizes the backends which are available in the Idmap backend drop-down menu. Each backend has its own [man page](http://samba.org.ru/samba/docs/man/manpages/) (<http://samba.org.ru/samba/docs/man/manpages/>) which gives implementation details. Since selecting the wrong backend will break Active Directory integration, a pop-up menu will appear whenever changes are made to this setting.

Table 9.2: ID Mapping Backends

Value	Description
ad	AD server uses RFC2307 or Services For Unix schema extensions; mappings must be provided in advance by adding the uidNumber attributes for users and gidNumber attributes for groups in the AD
autorid	similar to rid, but automatically configures the range to be used for each domain, so there is no need to specify a specific range for each domain in the forest; the only needed configuration is the range of UID/GIDs to use for user/group mappings and an optional size for the ranges
fruit	generate IDs the way Apple Mac OS X does, so UID and GID can be identical on all FreeNAS® servers on the network; for use in LDAP (page 163) environments where Apple's Open Directory is the authoritative LDAP server
ldap	stores and retrieves mapping tables in an LDAP directory service; default for LDAP directory service
nss	provides a simple means of ensuring that the SID for a Unix user is reported as the one assigned to the corresponding domain user
rfc2307	an AD server is required to provide the mapping between the name and SID and an LDAP server is required to provide the mapping between the name and the UID/GID
rid	default for AD; requires an explicit idmap configuration for each domain, using disjoint ranges where a writeable default idmap range should be defined, using a backend like tdb or ldap
script	stores mapping tables for clustered environments in the winbind_cache tdb
tdb	default backend used by winbindd for storing mapping tables
tdb2	substitute for tdb used by winbindd in clustered environments

Click the Rebuild Directory Service Cache button if a new Active Directory user needs immediate access to FreeNAS®. This occurs automatically once a day as a cron job.

Note: Active Directory places restrictions on which characters are allowed in Domain and NetBIOS names, a limits the length of those names to 15 characters. If there are problems connecting to the realm, [verify](https://support.microsoft.com/en-us/kb/909264) (<https://support.microsoft.com/en-us/kb/909264>) that your settings do not include any disallowed characters. Also, the Administrator account password cannot contain the \$ character. If a \$ exists in the domain administrator's password, `kinit` will report a "Password Incorrect" error and `ldap_bind` will report an "Invalid credentials (49)" error.

It can take a few minutes after configuring the Active Directory service for the AD information to be populated to the FreeNAS® system. Once populated, the AD users and groups will be available in the drop-down menus of the Permissions screen of a volume/dataset. For performance reasons, every available user may not show in the listing. However, it will autocomplete all applicable users when typing in a username.

The Active Directory users and groups that have been imported to the FreeNAS® system can be shown by using these commands from the FreeNAS® [Shell](#) (page 290). To view users:

```
wbinfo -u
```

To view groups:

```
wbinfo -g
```

In addition, `wbinfo -t` will test the connection and, if successful, will show a message similar to:

```
checking the trust secret for domain YOURDOMAIN via RPC calls succeeded
```

To manually check that a specified user can authenticate:

```
net ads join -S dcname -U username
```

If no users or groups are listed in the output, these commands can provide more troubleshooting information:

```
getent passwd  
getent group
```

If the `wbinfo` commands display the network users, but they do not show up in the drop-down menu of a Permissions screen, it may be because it is taking longer than the default ten seconds for the FreeNAS® system to join Active Directory. Try bumping up the value of AD timeout to 60 seconds.

Tip: To change a certificate, set the Encryption Mode to Off, then disable AD by unchecking Enable. Click the Save button. Select the new Certificate, set the Encryption Mode as desired, set the Enable checkbox to re-enable AD, and click the Save button to restart AD.

9.1.1 Troubleshooting Tips

When running AD in a 2003/2008 mixed domain, [refer to](https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) (https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) for instructions on how to prevent the secure channel key from becoming corrupt.

Active Directory uses DNS to determine the location of the domain controllers and global catalog servers in the network. Use the `host -t srv _ldap._tcp.domainname.com` command to determine the network's SRV records and, if necessary, change the weight and/or priority of the SRV record to reflect the fastest server. More information about SRV records can be found in the Technet article [How DNS Support for Active Directory Works](https://technet.microsoft.com/en-us/library/cc759550(Ws.10).aspx) (https://technet.microsoft.com/en-us/library/cc759550(Ws.10).aspx).

The realm that is used depends upon the priority in the SRV DNS record, meaning that DNS can override your Active Directory settings. When unable to connect to the correct realm, check the SRV records on the DNS server. [This article](http://www.informit.com/guides/content.aspx?g=security&seqNum=37&rll=1) (http://www.informit.com/guides/content.aspx?g=security&seqNum=37&rll=1) describes how to configure KDC discovery over DNS and provides some examples of records with differing priorities.

If the cache becomes out of sync due to an AD server being taken off and back online, resync the cache using `Directory Service → Active Directory → Rebuild Directory Service Cache`.

An expired password for the administrator account will cause kinit to fail, so ensure that the password is still valid. Also, double-check that the password on the AD account being used does not include any spaces or special symbols, and is not unusually long.

If the Windows server version is lower than 2008 R2, try creating a Computer entry on the Windows server's OU. When creating this entry, enter the FreeNAS® hostname in the name field. Make sure that it is under 15 characters and that it is the same name as the one set in the Hostname field in `Network → Global Configuration` and the NetBIOS Name in `Directory Service → Active Directory settings`. Make sure the hostname of the domain controller is set in the Domain Controller field of `Directory Service → Active Directory`.

9.1.2 If the System Will not Join the Domain

If the system will not join the Active Directory domain, run these commands in the order listed. If any of the commands fail or result in a traceback, create a bug report at <https://redmine.ixsystems.com/projects/freenas/issues> that includes the commands in the order in which they were run and the exact wording of the error message or traceback.

Start with these commands, where the `echo` commands should return a value of 0 and the `klist` command should show a Kerberos ticket:

```
sqlite3 /data/freenas-v1.db "update directoryservice_activedirectory set ad_enable=1;"
echo $?
service ix-kerberos start
service ix-nsswitch start
service ix-kinit start
service ix-kinit status
echo $?
klist
```

Next, only run these two commands if the Unix extensions box is checked in Advanced Mode and a keytab has been uploaded using [Kerberos Keytabs](#) (page 167):

```
service ix-sssd start
service sssd start
```

Finally, run these commands. Again, the `echo` command should return a 0:

```
python /usr/local/www/freenasUI/middleware/notifier.py start cifs
service ix-activedirectory start
service ix-activedirectory status
echo $?
python /usr/local/www/freenasUI/middleware/notifier.py restart cifs
service ix-pam start
service ix-cache start &
```

9.2 LDAP

FreeNAS® includes an [OpenLDAP](http://www.openldap.org/) (http://www.openldap.org/) client for accessing information from an LDAP server. An LDAP server provides directory services for finding network resources such as users and their associated permissions. Examples of LDAP servers include Microsoft Server (2000 and newer), Mac OS X Server, Novell eDirectory, and OpenLDAP running on a BSD or Linux system. If an LDAP server is running on your network, configure the FreeNAS® LDAP service so network users can authenticate to the LDAP server and have authorized access to the data stored on the FreeNAS® system.

Note: LDAP authentication for SMB shares is disabled unless the LDAP directory has been configured for and populated with Samba attributes. The most popular script for performing this task is [smbldap-tools](http://download.gna.org/smbldap-tools/) (http://download.gna.org/smbldap-tools/) and instructions for using it can be found at [The Linux Samba-OpenLDAP Howto](http://download.gna.org/smbldap-tools/docs/samba-ldap-howto/#htoc29) (http://download.gna.org/smbldap-tools/docs/samba-ldap-howto/#htoc29). In addition, the LDAP server must support SSL/TLS and the certificate for the LDAP server CA must be imported with `System → Certificates → Import Certificate`. Note that non-CA certificates are not supported at this time.

Tip: Apple's [Open Directory](https://manuals.info.apple.com/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) (https://manuals.info.apple.com/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) is an LDAP-compatible directory service into which FreeNAS® can be integrated. See [FreeNAS with Open Directory in Mac OS X environments](https://forums.freenas.org/index.php?threads/howto-freenas-with-open-directory-in-mac-os-x-environments.46493/) (https://forums.freenas.org/index.php?threads/howto-freenas-with-open-directory-in-mac-os-x-environments.46493/).

Figure 9.2 shows the LDAP Configuration screen that is seen after clicking `Directory Service → LDAP`.

The screenshot shows the 'Directory Service' configuration page in FreeNAS. The 'LDAP' tab is selected, with other tabs like 'Active Directory', 'NIS', 'Kerberos Realms', 'Kerberos Keytabs', and 'Kerberos Settings' visible. The configuration fields include:

- Hostname:** A text input field with an information icon.
- Base DN:** A text input field with an information icon.
- Bind DN:** A text input field with an information icon.
- Bind password:** A text input field with an information icon.
- Enable:** A checkbox that is currently unchecked.

 At the bottom, there are three buttons: 'Save', 'Advanced Mode', and 'Rebuild Directory Service Cache'.

Fig. 9.2: Configuring LDAP

Table 9.3 summarizes the available configuration options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in *System* → *Advanced*.

Those who are new to LDAP terminology should skim through the [OpenLDAP Software 2.4 Administrator's Guide](http://www.openldap.org/doc/admin24/) (<http://www.openldap.org/doc/admin24/>).

Table 9.3: LDAP Configuration Options

Setting	Value	Advanced Mode	Description
Hostname	string		hostname or IP address of LDAP server
Base DN	string		top level of the LDAP directory tree to be used when searching for resources (e.g. dc=test,dc=org)
Bind DN	string		name of administrative account on LDAP server (e.g. cn=Manager,dc=test,dc=org)
Bind password	string		password for Root bind DN
Allow Anonymous Binding	checkbox	✓	instructs LDAP server to not provide authentication and to allow read and write access to any client
User Suffix	string	✓	optional; can be added to name when user account added to LDAP directory (e.g. dept. or company name)
Group Suffix	string	✓	optional; can be added to name when group added to LDAP directory (e.g. dept. or company name)
Password Suffix	string	✓	optional; can be added to password when password added to LDAP directory
Machine Suffix	string	✓	optional; can be added to name when system added to LDAP directory (e.g. server, accounting)
SUDO Suffix	string	✓	use if LDAP-based users need superuser access

Continued on next page

Table 9.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Kerberos Realm	drop-down menu	✓	select the realm created using the instructions in Kerberos Realms (page 167)
Kerberos Principal	drop-down menu	✓	browse to the location of the principal in the keytab created as described in Kerberos Keytabs (page 167)
Encryption Mode	drop-down menu	✓	choices are Off, SSL, or TLS; note that either SSL or TLS and a Certificate must be selected in order for authentication to work
Certificate	drop-down menu	✓	select the certificate of the LDAP CA (required if authentication is used); the certificate for the LDAP server CA must first be imported with System → Certificates → Import Certificate
LDAP timeout	integer	✓	increase this value (in seconds) if obtaining a Kerberos ticket times out
DNS timeout	integer	✓	increase this value (in seconds) if DNS queries timeout
Idmap backend	drop-down menu and Edit	✓	select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs; see Table 9.2 for a summary of the available backends; click the Edit link to configure the backend's editable options
Samba Schema	checkbox	✓	only check this box if you need LDAP authentication for SMB shares and have already configured the LDAP server with Samba attributes
Auxiliary Parameters	string	✓	additional options for sssd.conf(5) (https://jhrozek.fedorapeople.org/sss/1.11.6/man/sss.conf.5.html)
Schema	drop-down menu	✓	if Samba Schema is checked, select the schema to use; choices are rfc2307 and rfc2307bis
Enable	checkbox		uncheck to disable the configuration without deleting it
NetBIOS Name	string	✓	limited to 15 characters; automatically populated with the system's original hostname; must be different from the Workgroup name
NetBIOS Alias	string	✓	limited to 15 characters

Click the Rebuild Directory Service Cache button after adding a user to LDAP who needs immediate access to FreeNAS®. Otherwise this occurs automatically once a day as a cron job.

Note: FreeNAS® automatically appends the root DN. This means that the scope and root DN should not be included when configuring the user, group, password, and machine suffixes.

LDAP users and groups appear in the drop-down menus of the Permissions screen of a volume/dataset after configuring the LDAP service. Type `getent passwd` from [Shell](#) (page 290) to verify that the users have been imported. Type `getent group` to verify that the groups have been imported.

If the users and groups are not listed, refer to [Common errors encountered when using OpenLDAP Software](#) (<http://www.openldap.org/doc/admin24/appendix-common-errors.html>) for common errors and how to fix them. When troubleshooting LDAP, open [Shell](#) (page 290) and look for error messages in `/var/log/auth.log`.

9.3 NIS

Network Information Service (NIS) is a service which maintains and distributes a central directory of Unix user and group information, hostnames, email aliases, and other text-based tables of information. If a NIS server is running on your network, the FreeNAS® system can be configured to import the users and groups from the NIS directory.

Note: In Windows Server 2016, Microsoft removed the Identity Management for Unix (IDMU) and NIS Server Role. See [Clarification regarding the status of Identity Management for Unix \(IDMU\) & NIS Server Role in Windows Server 2016 Technical Preview and beyond](https://blogs.technet.microsoft.com/activedirectoryua/2016/02/09/identity-management-for-unix-idmu-is-deprecated-in-windows-server/) (<https://blogs.technet.microsoft.com/activedirectoryua/2016/02/09/identity-management-for-unix-idmu-is-deprecated-in-windows-server/>).

Figure 9.3 shows the configuration screen which opens when you click `Directory Service` → `NIS`. Table 9.4 summarizes the configuration options.

Fig. 9.3: NIS Configuration

Table 9.4: NIS Configuration Options

Setting	Value	Description
NIS domain	string	name of NIS domain
NIS servers	string	comma delimited list of hostnames or IP addresses
Secure mode	checkbox	if checked, <code>ypbind(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ypbind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, <code>ypbind</code> will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet
Enable	checkbox	uncheck to disable the configuration without deleting it

Click the `Rebuild Directory Service Cache` button after adding a user to NIS who needs immediate access to FreeNAS®. Otherwise this occurs automatically once a day as a cron job.

9.4 Kerberos Realms

A default Kerberos realm is created for the local system in FreeNAS®. [Directory Service](#) → [Kerberos Realms](#) can be used to view and add Kerberos realms. If the network contains a KDC, click the Add kerberos realm button to add the Kerberos realm. This configuration screen is shown in [Figure 9.4](#).



Fig. 9.4: Adding a Kerberos Realm

[Table 9.5](#) summarizes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in [System](#) → [Advanced](#).

Table 9.5: Kerberos Realm Options

Setting	Value	Advanced Mode	Description
Realm	string		mandatory; name of the realm
KDC	string	✓	name of the Key Distribution Center
Admin Server	string	✓	server where all changes to the database are performed
Password Server	string	✓	server where all password changes are performed

9.5 Kerberos Keytabs

Kerberos keytabs are used to do Active Directory or LDAP joins without a password. This means that the password for the Active Directory or LDAP administrator account does not need to be saved into the FreeNAS® configuration database, which is a security risk in some environments.

When using a keytab, it is recommended to create and use a less privileged account for performing the required queries as the password for that account will be stored in the FreeNAS® configuration database. To create the keytab on a Windows system, use these commands:

```
ktpass.exe -out hostname.keytab host/ hostname@DOMAINNAME -ptype KRB5_NT_PRINCIPAL -mapuser_
↪DOMAIN\username -pass userpass

setspn -A host/ hostname@DOMAINNAME DOMAIN\username
```

where:

- hostname is the fully qualified hostname of the domain controller
- DOMAINNAME is the domain name in all caps
- DOMAIN is the pre-Windows 2000 short name for the domain
- username is the privileged account name

- userpass is the password associated with username

This will create a keytab with sufficient privileges to grant tickets.

After the keytab is generated, use `Directory Service → Kerberos Keytabs → Add kerberos keytab` to add it to the FreeNAS® system.

To instruct the Active Directory service to use the keytab, select the installed keytab using the drop-down Kerberos keytab menu in `Directory Service → Active Directory`. When using a keytab with Active Directory, make sure that the “username” and “userpass” in the keytab matches the “Domain Account Name” and “Domain Account Password” fields in `Directory Service → Active Directory`.

To instruct LDAP to use a principal from the keytab, select the principal from the drop-down Kerberos Principal menu in `Directory Service → LDAP`.

9.6 Kerberos Settings

To configure additional Kerberos parameters, use `Directory Service → Kerberos Settings`. Figure 9.5 shows the fields available:

- Appdefaults auxiliary parameters: contains settings used by some Kerberos applications. The available settings and their syntax are listed in the [\[appdefaults\] section of krb.conf\(5\)](http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#appdefaults) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#appdefaults).
- Libdefaults auxiliary parameters: contains settings used by the Kerberos library. The available settings and their syntax are listed in the [\[libdefaults\] section of krb.conf\(5\)](http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#libdefaults) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#libdefaults).

Fig. 9.5: Additional Kerberos Settings

SHARING

Shares are created to make part or all of a volume accessible to other computers on the network. The type of share to create depends on factors like which operating systems are being used by computers on the network, security requirements, and expectations for network transfer speeds.

FreeNAS® provides a [Wizard](#) (page 282) for creating shares. The [Wizard](#) (page 282) automatically creates the correct type of dataset and permissions for the type of share, sets the default permissions for the share type, and starts the service needed by the share. It is recommended to use the Wizard to create shares, fine-tune the share settings using the instructions in the rest of this chapter if needed, then fine-tune the default permissions from the client operating system to meet the requirements of the network.

Note: Shares are created to provide and control access to an area of storage. Before creating shares, it is recommended to make a list of the users that need access to storage data, which operating systems these users are using, whether all users should have the same permissions to the stored data, and whether these users should authenticate before accessing the data. This information can help determine which type of shares are needed, whether multiple datasets are needed to divide the storage into areas with different access and permissions, and how complex it will be to set up those permission requirements. Note that shares are used to provide access to data. When a share is deleted, it removes access to data but does not delete the data itself.

These types of shares and services are available:

- [AFP](#) (page 170): Apple File Protocol shares are often used when the client computers all run Mac OS X. Apple has slowly shifted to preferring [SMB](#) (page 185) for modern networks, although Time Machine still requires AFP.
- [Unix \(NFS\)](#) (page 177): Network File System shares are accessible from Mac OS X, Linux, BSD, and the professional and enterprise versions (but not the home editions) of Windows. This can be a good choice when the client computers do not all run the same operating system but NFS client software is available for all of them.
- [WebDAV](#) (page 184): WebDAV shares are accessible using an authenticated web browser (read-only) or [WebDAV client](#) (<https://en.wikipedia.org/wiki/WebDAV#Clients>) running on any operating system.
- [SMB](#) (page 185): Server Message Block shares, also known as Common Internet File System (CIFS) shares, are accessible by Windows, Mac OS X, Linux, and BSD computers. Access is slower than an NFS share due to the single-threaded design of Samba. SMB provides more configuration options than NFS and is a good choice on a network for Windows systems. However, it is a poor choice if the CPU on the FreeNAS® system is limited; if the CPU is maxed out, upgrade the CPU or consider another type of share.
- [Block \(iSCSI\)](#) (page 195): block or iSCSI shares appear as an unformatted disk to clients running iSCSI initiator software or a virtualization solution such as VMware. These are usually used as virtual drives.

Fast access from any operating system can be obtained by configuring the [FTP](#) (page 216) service instead of a share and using a cross-platform FTP file manager application such as [Filezilla](#) (<https://filezilla-project.org/>). Secure FTP can be configured if the data needs to be encrypted.

When data security is a concern and the network users are familiar with SSH command line utilities or [WinSCP](#) (<http://winscp.net/eng/index.php>), consider using the [SSH](#) (page 236) service instead of a share. It is slower than unencrypted FTP due to the encryption overhead, but the data passing through the network is encrypted.

Note: It is generally a mistake to share a volume or dataset with more than one share type or access method. Different types of shares and services use different file locking methods. For example, if the same volume is configured to use both NFS and FTP, NFS will lock a file for editing by an NFS user, but an FTP user can simultaneously edit or delete that file. This results in lost edits and confused users. Another example: if a volume is configured for both AFP and SMB, Windows users can be confused by the “extra” filenames used by Mac files and delete them. This corrupts the files on the AFP share. Pick the one type of share or service that makes the most sense for the types of clients accessing that volume, and use that single type of share or service. To support multiple types of shares, divide the volume into datasets and use one dataset per share.

This section demonstrates configuration and fine-tuning of AFP, NFS, SMB, WebDAV, and iSCSI shares. FTP and SSH configurations are described in [Services](#) (page 209).

10.1 Apple (AFP) Shares

FreeNAS® uses the [Netatalk](http://netatalk.sourceforge.net/) (<http://netatalk.sourceforge.net/>) AFP server to share data with Apple systems. This section describes the configuration screen for fine-tuning AFP shares created using the [Wizard](#) (page 282). It then provides configuration examples for using the [Wizard](#) (page 282) to create a guest share, configuring Time Machine to back up to a dataset on the FreeNAS® system, and for connecting to the share from a Mac OS X client.

To view the AFP share created by the Wizard, click [Sharing](#) → [Apple \(AFP\)](#) and highlight the name of the share. Click its [Edit](#) button to see the configuration options shown in [Figure 10.1](#). The values showing for these options will vary, depending upon the information given when the share was created.

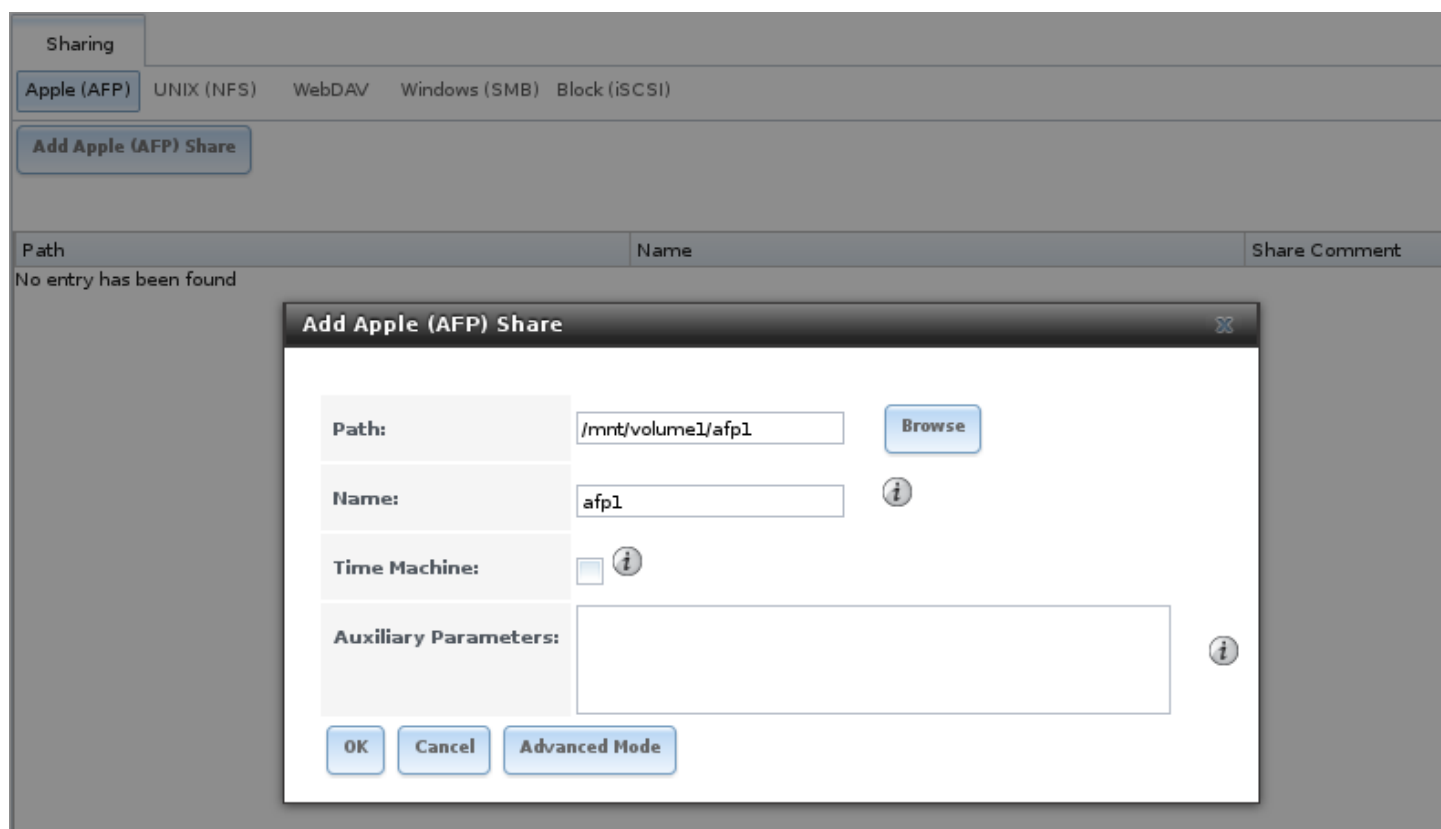


Fig. 10.1: Creating an AFP Share

Note: [Table 10.1](#) summarizes the options available to fine-tune an AFP share. These options should usually be left at the default settings. Changing them might cause unexpected behavior. Most settings are only available with Advanced Mode.

Do not change an advanced option without fully understanding the function of that option. Refer to [Setting up Netatalk](http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html) (<http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html>) for a more detailed explanation of these options.

Table 10.1: AFP Share Configuration Options

Setting	Value	Advanced Mode	Description
Path	browse button		browse to the volume/dataset to share; do not nest additional volumes, datasets, or symbolic links beneath this path because Netatalk does not fully support that
Name	string		volume name which appears in the Mac computer's connect to server dialog; limited to 27 characters and cannot contain a period
Share Comment	string	✓	optional comment
Allow List	string	✓	comma-delimited list of allowed users and/or groups where groupname begins with a @; note that adding an entry will deny any user/group that is not specified
Deny List	string	✓	comma-delimited list of denied users and/or groups where groupname begins with a @; note that adding an entry will allow all users/groups that are not specified
Read-only Access	string	✓	comma-delimited list of users and/or groups who only have read access where groupname begins with a @
Read-write Access	string	✓	comma-delimited list of users and/or groups who have read and write access where groupname begins with a @
Time Machine	checkbox		when checked, FreeNAS® advertises itself as a Time Machine disk so it can be found by Macs; due to a limitation in how the Mac deals with low-diskspace issues when multiple Macs share the same volume, checking Time Machine on multiple shares could result in intermittent failed backups
Time Machine Quota	checkbox		only appears when Time Machine is checked; when checked, each time machine backup on the share has its own quota
Zero Device Numbers	checkbox	✓	enable when the device number is not constant across a reboot
No Stat	checkbox	✓	if checked, AFP does not stat the volume path when enumerating the volumes list; useful for automounting or volumes created by a preexec script
AFP3 UNIX Privs	checkbox	✓	enable Unix privileges supported by OSX 10.5 and higher; do not enable this if the network contains Mac OS X 10.4 clients or lower as they do not support this feature
Default file permission	checkboxes	✓	only works with Unix ACLs; new files created on the share are set with the selected permissions
Default directory permission	checkboxes	✓	only works with Unix ACLs; new directories created on the share are set with the selected permissions
Default umask	integer	✓	umask used for newly created files, default is 000 (anyone can read, write, and execute)
Hosts Allow	string	✓	comma-, space-, or tab-delimited list of allowed hostnames or IP addresses
Hosts Deny	string	✓	comma-, space-, or tab-delimited list of denied hostnames or IP addresses
Auxiliary Parameters	string		additional afp.conf (http://netatalk.sourceforge.net/3.1/htmldocs/afp.conf.5.html) parameters not covered by other option fields

10.1.1 Creating AFP Guest Shares

AFP supports guest logins, meaning that Mac OS X users can access the AFP share without requiring their user accounts to first be created on or imported into the FreeNAS® system.

Note: When a guest share is created along with a share that requires authentication, AFP only maps users who log in as guest to the guest share. If a user logs in to the share that requires authentication, permissions on the guest share can prevent that user from writing to the guest share. The only way to allow both guest and authenticated users to write to a guest share is to set the permissions on the guest share to 777 or to add the authenticated users to a guest group and set the permissions to 77x.

Before creating a guest share, go to *Services* → *AFP* and make sure that the Guest Access box is checked.

To create the AFP guest share, click *Wizard*, then click the *Next* button twice to display the screen shown in [Figure 10.2](#). Complete these fields in this screen:

1. **Share name:** enter a name for the share that is identifiable but less than 27 characters long. This name cannot contain a period. In this example, the share is named `afp_guest`.
2. Click the button for *Mac OS X (AFP)*.
3. Click the *Ownership* button. Click the drop-down *User* menu and select *nobody*. Click the *Return* button to return to the previous screen.
4. Click the *Add* button. The share is not created until the button is clicked. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.

The screenshot shows a 'Wizard' window with a dark title bar. Inside, the 'Share name' field is filled with 'afp_guest'. Below this, the 'Purpose' section has four radio buttons: 'Windows (SMB)', 'Mac OS X (AFP)' (which is selected), 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' and 'Time Machine'. An 'Ownership' button is to the right of the 'Purpose' section. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. Below these is a list box labeled 'Name' which contains the entry 'afp_guest'. At the bottom of the window are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 10.2: Creating a Guest AFP Share

Click the Next button twice, then the Confirm button to create the share. The Wizard automatically creates a dataset for the share that contains the correct default permissions and starts the AFP service so the share is immediately available. The new share is also added as an entry to **Sharing** → **Apple** (AFP).

Mac OS X users can connect to the guest AFP share by clicking **Go** → **Connect to Server**. In the example shown in [Figure 10.3](#), the user has entered `afp://` followed by the IP address of the FreeNAS® system.

Click the Connect button. Once connected, Finder opens automatically. The name of the AFP share is displayed in the **SHARED** section in the left frame and the contents of any data saved in the share is displayed in the right frame.




Fig. 10.3: Connect to Server Dialogue

To disconnect from the volume, click the eject button in the Shared sidebar.

10.1.2 Creating Authenticated and Time Machine Shares

Mac OS X includes the Time Machine application which can be used to schedule automatic backups. In this configuration example, a Time Machine user will be configured to backup to an AFP share on a FreeNAS® system. It is recommended to create a separate Time Machine share for each user that will be using Time Machine to backup their Mac OS X system to FreeNAS®. The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

To use the Wizard to create an authenticated or Time Machine share, enter the following information, as seen in the example in [Figure 10.4](#).

1. Share name: enter a name for the share that is identifiable but less than 27 characters long. The name cannot contain a period. In this example, the share is named `backup_user1`.
2. Click the button for Mac OS X (AFP) and check the box for Time Machine.
3. Click the Ownership button. If the user already exists on the FreeNAS® system, click the drop-down User menu to select their user account. If the user does not yet exist on the FreeNAS® system, type their name into the User field and check the Create User checkbox. If the user will be a member of a group that already exists on the FreeNAS® system, click the drop-down Group menu to select the group name. To create a new group to be used by Time Machine users, enter the name in the Group field and check the Create Group checkbox. Otherwise, enter the same name as the user. In

the example shown in [Figure 10.5](#), both a new user1 user and a new tm_backups group will be created. Since a new user is being created, this screen prompts for the user password to be used when accessing the share. It also provides an opportunity to change the default permissions on the share. When finished, click Return to return to the screen shown in [Figure 10.4](#).

4. Click the Add button. Remember to do this or the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

To configure multiple authenticated or Time Machine shares, repeat for each user, giving each user their own Share name and Ownership. When finished, click the Next button twice, then the Confirm button to create the shares. The Wizard automatically creates a dataset for each share with the correct ownership and starts the AFP service so the shares are immediately available. The new shares are also added to *Sharing* → *Apple* (AFP).

The screenshot shows the 'Wizard' window with a dark title bar. Inside, the 'Share name' field contains 'backup_user1'. Below it, the 'Purpose' section has four radio buttons: 'Windows (SMB)', 'Mac OS X (AFP)' (selected), 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' (unchecked) and 'Time Machine' (checked). A 'Size' field is empty. An 'Ownership' button is to the right of the checkboxes. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. A list box labeled 'Name' contains 'backup_user1'. At the bottom are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 10.4: Creating a Time Machine Share

Wizard

User: user1 ☒ Create User ⓘ

User Password: ●●●●●●

Confirm User Password: ●●●●●●

Group: tm_backups ☒ Create Group ⓘ

Mode:

	Owner	Group	Other
Read	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Write	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Execute	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Return Cancel

Fig. 10.5: Creating an Authenticated User

At this point, it may be desirable to configure a quota for each Time Machine share, to restrict backups from using all of the available space on the FreeNAS® system. The first time Time Machine makes a backup, it will create a full backup after waiting two minutes. It will then create a one hour incremental backup for the next 24 hours, and then one backup each day, each week and each month. Since the oldest backups are deleted when a Time Machine share becomes full, make sure that the quota size is sufficient to hold the desired number of backups. Note that a default installation of Mac OS X is ~21 GB in size.

To configure a quota, go to **Storage** → **Volumes** and highlight the entry for the share. In the example shown in [Figure 10.6](#), the Time Machine share name is backup_user1. Click the **Edit Options** button for the share, then **Advanced Mode**. Enter a value in the **Quota** for this dataset field, then click **Edit Dataset** to save the change. In this example, the Time Machine share is restricted to 200 GB.

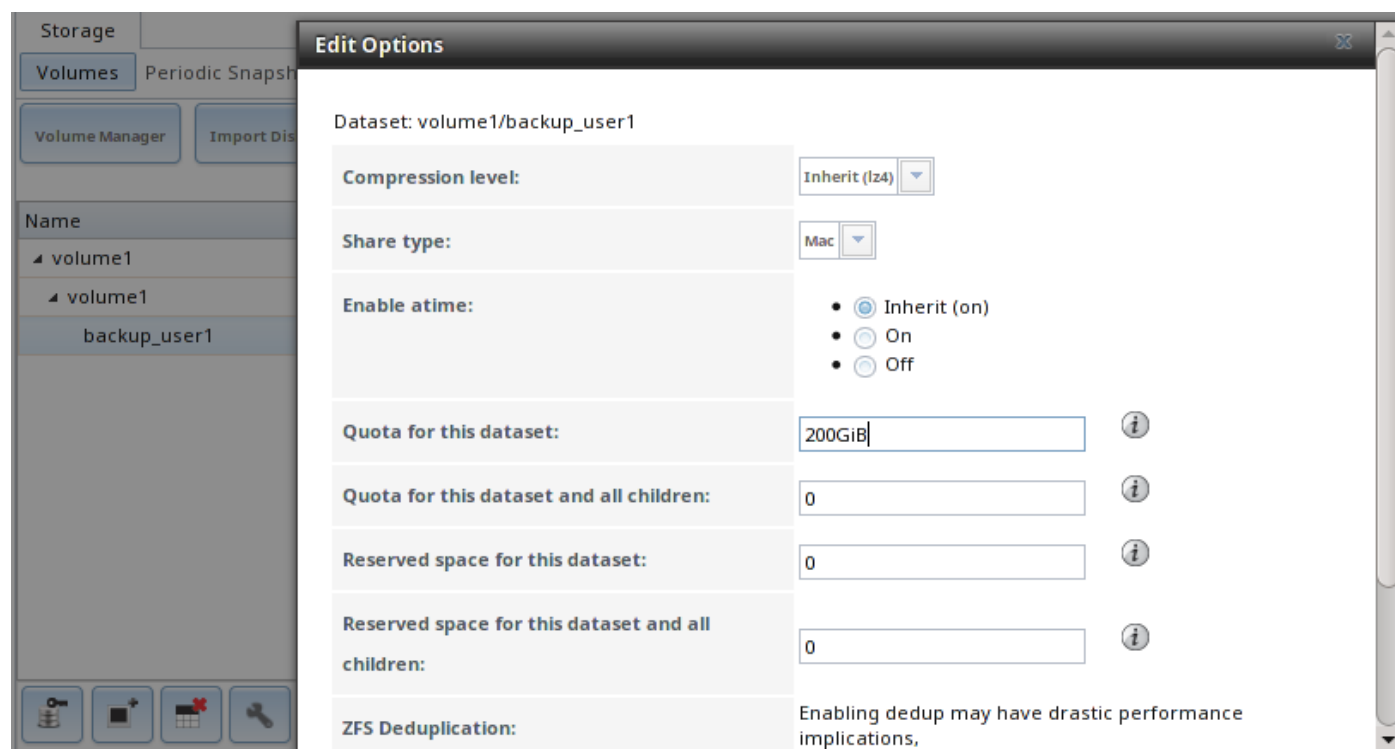


Fig. 10.6: Setting a Quota

Note: An alternative is to create a global quota using the instructions in [Set up Time Machine for multiple machines with OSX Server-Style Quotas](https://forums.freenas.org/index.php?threads/how-to-set-up-time-machine-for-multiple-machines-with-osx-server-style-quotas.47173/) (<https://forums.freenas.org/index.php?threads/how-to-set-up-time-machine-for-multiple-machines-with-osx-server-style-quotas.47173/>).

To configure Time Machine on the Mac OS X client, go to `System Preferences → Time Machine` which opens the screen shown in [Figure 10.7](#). Click ON and a pop-up menu shows the FreeNAS® system as a backup option. In our example, it is listed as backup_user1 on "freenas". Highlight the FreeNAS® system and click Use Backup Disk. A connection bar opens and prompts for the user account's password—in this example, the password that was set for the user1 account.



Fig. 10.7: Configuring Time Machine on Mac OS X Lion

If Time Machine could not complete the backup. The backup disk image could not be created (error 45) is shown when backing up to the FreeNAS® system, a sparsebundle image must be created using [these instructions](http://forum1.netgear.com/showthread.php?t=49482) (<http://forum1.netgear.com/showthread.php?t=49482>).

If Time Machine completed a verification of your backups. To improve reliability, Time Machine must create a new backup for you. is shown, follow the instructions in [this post](http://www.garth.org/archives/2011,08,27,169,fix-time-machine-sparsebundle-nas-based-backup-errors.html) (<http://www.garth.org/archives/2011,08,27,169,fix-time-machine-sparsebundle-nas-based-backup-errors.html>) to avoid making another backup or losing past backups.

10.2 Unix (NFS) Shares

FreeNAS® supports sharing over the Network File System (NFS). Clients use the `mount` command to mount the share. Once mounted, the NFS share appears as just another directory on the client system. Some Linux distros require the installation of additional software in order to mount an NFS share. On Windows systems, enable Services for NFS in the Ultimate or Enterprise editions or install an NFS client application.

Note: For performance reasons, iSCSI is preferred to NFS shares when FreeNAS® is installed on ESXi. When considering creating NFS shares on ESXi, read through the performance analysis at [Running ZFS over NFS as a VMware Store](http://blog.laspina.ca/ubiquitous/running-zfs-over-nfs-as-a-vmware-store) (<http://blog.laspina.ca/ubiquitous/running-zfs-over-nfs-as-a-vmware-store>).

To create an NFS share using the [Wizard](#) (page 282), click the Next button twice to display the screen shown in [Figure 10.8](#). Enter a Share name. Spaces are not allowed in these names. Click the button for Generic Unix (NFS), then click Add so the

share name appears in the Name frame. When finished, click the Next button twice, then the Confirm button to create the share. Creating an NFS share using the wizard automatically creates a new dataset for the share, starts the services required for NFS, and adds an entry in *Sharing* → *Unix (NFS) Shares*. Depending on your requirements, the IP addresses that are allowed to access the NFS share can be restricted, or the permissions adjusted.

Wizard

Share name:

Purpose

☐ Windows (SMB) ☐ Allow Guest

☐ Mac OS X (AFP) ☐ Time Machine

☒ Generic Unix (NFS)

☐ Block Storage (iSCSI) Size:

Ownership

Add Delete Update

Name
nfs_share1

Previous Next Exit

Fig. 10.8: NFS Share Wizard

NFS shares are edited by clicking *Sharing* → *Unix (NFS)*, highlighting the entry for the share, and clicking its Edit button. In the example shown in [Figure 10.9](#), the configuration screen is open for the `nfs_share1` share.

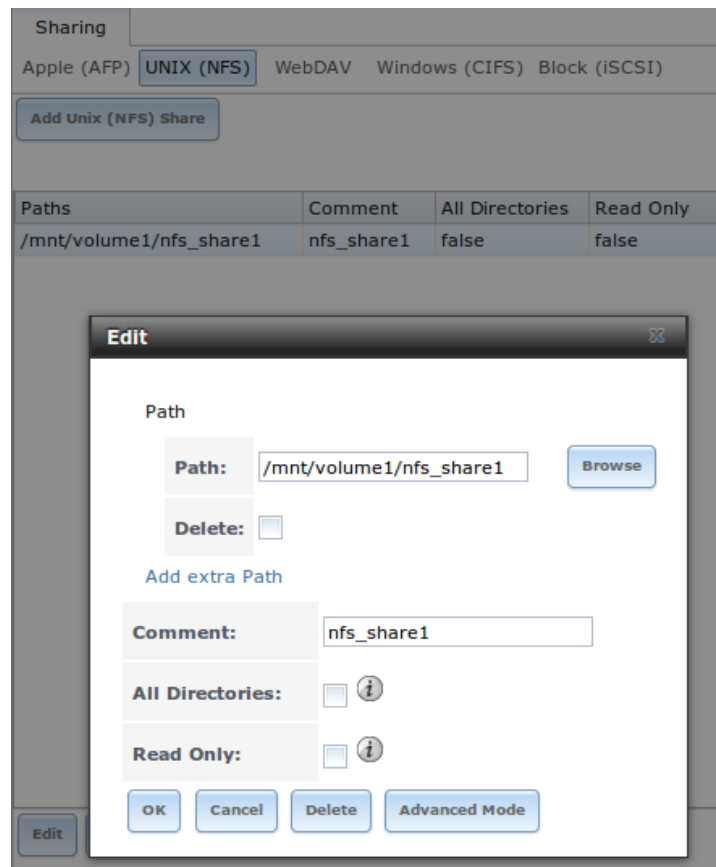


Fig. 10.9: NFS Share Settings

Table 10.2 summarizes the available configuration options in this screen. Some settings are only available by clicking the Advanced Mode button.

Table 10.2: NFS Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		browse to the volume or dataset to be shared; click Add extra path to select multiple paths
Comment	string		set the share name; if left empty, share name is the list of selected Path entries
Authorized networks	string	✓	list of allowed networks in network/mask CIDR notation, like 1.2.3.0/24, space-delimited; leave empty to allow all
Authorized IP addresses or hosts	string	✓	list of allowed IP addresses or hostnames, space-delimited; leave empty to allow all
All directories	checkbox		when checked, allow the client to mount any subdirectory within the Path
Read only	checkbox		prohibit writing to the share
Quiet	checkbox	✓	inhibit otherwise-useful syslog diagnostics to avoid some annoying error messages; see exports(5) (http://www.freebsd.org/cgi/man.cgi?query=exports) for examples
Maproot User	drop-down menu	✓	when a user is selected, the root user is limited to that user's permissions

Continued on next page

Table 10.2 – continued from previous page

Setting	Value	Advanced Mode	Description
Maproot Group	drop-down menu	✓	when a group is selected, the root user is also limited to that group's permissions
Mapall User	drop-down menu	✓	the specified user's permissions are used by all clients
Mapall Group	drop-down menu	✓	the specified group's permissions are used by all clients
Security	selection	✓	only appears if Enable NFSv4 is checked in <i>Services</i> → <i>NFS</i> ; choices are sys or these Kerberos options: krb5 (authentication only), krb5i (authentication and integrity), or krb5p (authentication and privacy); if multiple security mechanisms are added to the Selected column using the arrows, use the Up or Down buttons to list in order of preference

When creating NFS shares, keep these points in mind:

1. Clients will specify the Path when mounting the share.
2. The Maproot and Mapall options are exclusive, meaning only one can be used—the GUI does not allow both. The Mapall options supersede the Maproot options. To restrict only the root user's permissions, set the Maproot option. To restrict permissions of all users, set the Mapall options.
3. Each volume or dataset is considered to be its own filesystem and NFS is not able to cross filesystem boundaries.
4. The network or host must be unique per share and per filesystem or directory.
5. The All directories option can only be used once per share per filesystem.

To better understand these restrictions, consider a scenario where there are:

- two networks, 10.0.0.0/8 and 20.0.0.0/8
- a ZFS volume named `volume1` with 2 datasets named `dataset1` and `dataset2`
- `dataset1` contains a directory named `directory1`

Because of restriction #3, an error is shown when trying to create one NFS share like this:

- Authorized networks set to 10.0.0.0/8 20.0.0.0/8
- Path set to `/mnt/volume1/dataset1` and `/mnt/volume1/dataset1/directory1`

Instead, set a Path of `/mnt/volume1/dataset1` and check the All directories box.

That directory could also be restricted to one of the networks by creating two shares instead:

First NFS share:

- Authorized networks set to 10.0.0.0/8
- Path set to `/mnt/volume1/dataset1`

Second NFS share:

- Authorized networks set to 20.0.0.0/8
- Path set to `/mnt/volume1/dataset1/directory1`

Note that this requires the creation of two shares. It cannot be done with only one share.

10.2.1 Example Configuration

By default, the Mapall fields are not set. This means that when a user connects to the NFS share, the user has the permissions associated with their user account. This is a security risk if a user is able to connect as root as they will have complete access to the share.

A better option is to do this:

1. Specify the built-in nobody account to be used for NFS access.
2. In the Change Permissions screen of the volume/dataset that is being shared, change the owner and group to nobody and set the permissions according to your requirements.
3. Select nobody in the Mapall User and Mapall Group drop-down menus for the share in *Sharing* → *Unix (NFS) Shares*.

With this configuration, it does not matter which user account connects to the NFS share, as it will be mapped to the nobody user account and will only have the permissions that were specified on the volume/dataset. For example, even if the root user is able to connect, it will not gain root access to the share.

10.2.2 Connecting to the Share

The following examples share this configuration:

1. The FreeNAS® system is at IP address 192.168.2.2.
2. A dataset named `/mnt/volume1/nfs_share1` is created and the permissions set to the nobody user account and the nobody group.
3. An NFS share is created with these attributes:
 - Path: `/mnt/volume1/nfs_share1`
 - Authorized Networks: 192.168.2.0/24
 - All Directories checkbox is checked
 - MapAll User is set to nobody
 - MapAll Group is set to nobody

From BSD or Linux

NFS shares are mounted on BSD or Linux clients with this command executed as the superuser (root) or with `sudo`:

```
mount -t nfs 192.168.2.2:/mnt/volume1/nfs_share1 /mnt
```

- `-t nfs` specifies the filesystem type of the share
- 192.168.2.2 is the IP address of the FreeNAS® system
- `/mnt/volume/nfs_share1` is the name of the directory to be shared, a dataset in this case
- `/mnt` is the mountpoint on the client system. This must be an existing, empty directory. The data in the NFS share appears in this directory on the client computer.

A successful mounting of the share returns to the command prompt without any status or error messages.

Note: If this command fails on a Linux system, make sure that the [nfs-utils](http://sourceforge.net/projects/nfs/files/nfs-utils/) (<http://sourceforge.net/projects/nfs/files/nfs-utils/>) package is installed.

This configuration allows users on the client system to copy files to and from `/mnt` (the mount point). All files are owned by `nobody:nobody`. Changes to any files or directories in `/mnt` are written to the FreeNAS® system's `/mnt/volume1/nfs_share1` dataset.

Settings cannot be changed on the NFS share if it is mounted on any client computers. The `umount` command is used to unmount the share on BSD and Linux clients. Run it as the superuser or with `sudo` on each client computer:

```
umount /mnt
```

From Microsoft

Windows NFS client support varies with versions and releases. For best results, use [Windows \(SMB\) Shares](#) (page 185).

From Mac OS X

To mount the NFS volume from a Mac OS X client, click on `Go → Connect to Server`. In the Server Address field, enter `nfs://` followed by the IP address of the FreeNAS® system and the name of the volume/dataset being shared by NFS. The example shown in [Figure 10.10](#) continues with our example of `192.168.2.2:/mnt/volume1/nfs_share1`.

Finder opens automatically after connecting. The IP address of the FreeNAS® system is displayed in the SHARED section in the left frame and the contents of the share are displayed in the right frame. In the example shown in [Figure 10.11](#), `/mnt/data` has one folder named `images`. The user can now copy files to and from the share.

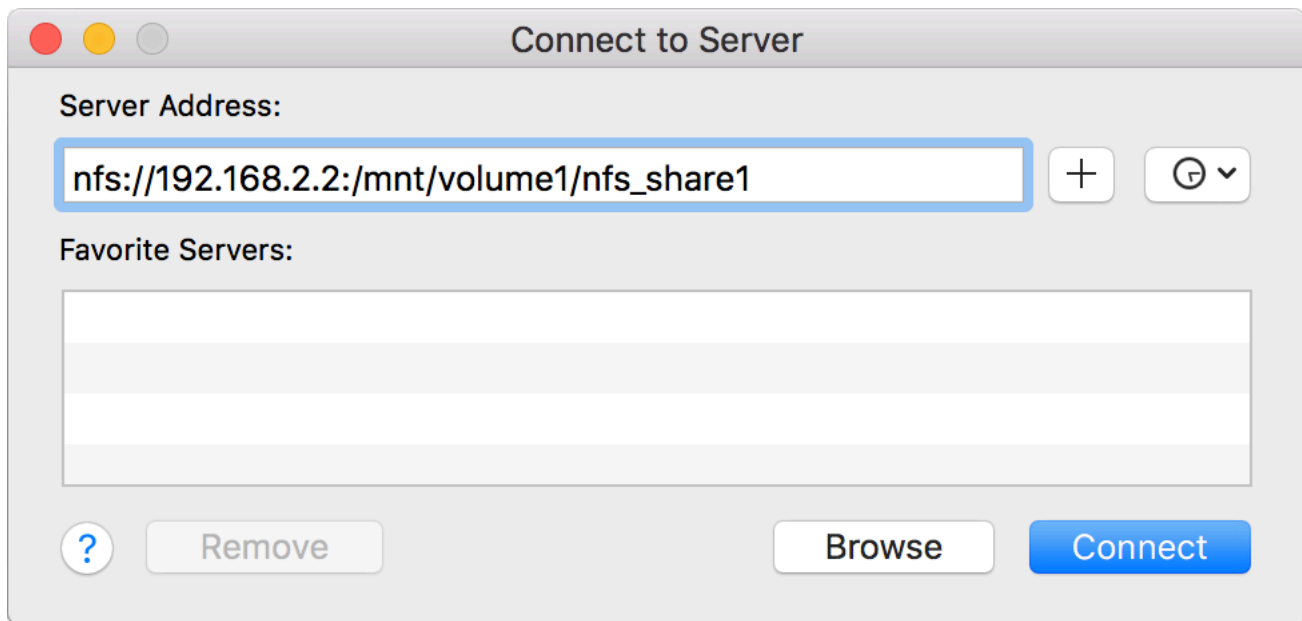


Fig. 10.10: Mounting the NFS Share from Mac OS X




Fig. 10.11: Viewing the NFS Share in Finder

10.2.3 Troubleshooting NFS

Some NFS clients do not support the NLM (Network Lock Manager) protocol used by NFS. This is the case if the client receives an error that all or part of the file may be locked when a file transfer is attempted. To resolve this error, add the option `-o nolock` when running the `mount` command on the client to allow write access to the NFS share.

If a “time out giving up” error is shown when trying to mount the share from a Linux system, make sure that the portmapper service is running on the Linux client. If portmapper is running and timeouts are still shown, force the use of TCP by including `-o tcp` in the `mount` command.

If a “RPC: Program not registered” error is shown, upgrade to the latest version of FreeNAS® and restart the NFS service after the upgrade to clear the NFS cache.

If clients see “reverse DNS” errors, add the FreeNAS® IP address in the Host name database field of `Network → Global Configuration`.

If clients receive timeout errors when trying to mount the share, add the client IP address and hostname to the Host name data base field in `Network → Global Configuration`.

Some older versions of NFS clients default to UDP instead of TCP and do not auto-negotiate for TCP. By default, FreeNAS®

uses TCP. To support UDP connections, go to *Services* → *NFS* and check the box *Serve UDP NFS clients*.

The `nfsstat -c` or `nfsstat -s` commands can be helpful to detect problems from the *Shell* (page 290). A high proportion of retries and timeouts compared to reads usually indicates network problems.

10.3 WebDAV Shares

In FreeNAS®, WebDAV shares can be created so that authenticated users can browse the contents of the specified volume, dataset, or directory from a web browser.

Configuring WebDAV shares is a two step process. First, create the WebDAV shares to specify which data can be accessed. Then, configure the WebDAV service by specifying the port, authentication type, and authentication password. Once the configuration is complete, the share can be accessed using a URL in the format:

```
protocol://IP_address:port_number/share_name
```

where:

- **protocol:** is either `http` or `https`, depending upon the Protocol configured in *Services* → *WebDAV*.
- **IP address:** is the IP address or hostname of the FreeNAS® system. Take care when configuring a public IP address to ensure that the network's firewall only allows access to authorized systems.
- **port_number:** is configured in *Services* → *WebDAV*. If the FreeNAS® system is to be accessed using a public IP address, consider changing the default port number and ensure that the network's firewall only allows access to authorized systems.
- **share_name:** is configured in *Sharing* → *WebDAV Shares*.

Entering the URL in a web browser brings up an authentication pop-up message. Enter a username of `webdav` and the password configured in *Services* → *WebDAV*.

Warning: At this time, only the `webdav` user is supported. For this reason, it is important to set a good password for this account and to only give the password to users which should have access to the WebDAV share.

To create a WebDAV share, click *Sharing* → *WebDAV Shares* → *Add WebDAV Share* which will open the screen shown in Figure 10.12.

Fig. 10.12: Adding a WebDAV Share

Table 10.3 summarizes the available options.

Table 10.3: WebDAV Share Options

Setting	Value	Description
Share Path Name	string	input a name for the share
Comment	string	optional
Path	browse button	browse to the volume/dataset to share
Read Only	checkbox	if checked, users cannot write to the share
Change User & Group Ownership	checkbox	if checked, automatically sets the share's contents to the webdav user and group

After clicking OK, a pop-up asks about enabling the service. Once the service starts, review the settings in *Services* → *WebDAV* as they are used to determine which URL is used to access the WebDAV share and whether or not authentication is required to access the share. These settings are described in [WebDAV](#) (page 242).

10.4 Windows (SMB) Shares

FreeNAS® uses [Samba](https://www.samba.org/) (<https://www.samba.org/>) to share volumes using Microsoft's SMB protocol. SMB is built into the Windows and Mac OS X operating systems and most Linux and BSD systems pre-install the Samba client in order to provide support for SMB. If your distro did not, install the Samba client using the distro's software repository.

The SMB protocol supports many different types of configuration scenarios, ranging from the very simple to quite complex. The complexity of the scenario depends upon the types and versions of the client operating systems that will connect to the share, whether the network has a Windows server, and whether Active Directory is being used. Depending on the authentication requirements, it might be necessary to create or import users and groups.

Samba supports server-side copy of files on the same share with clients from Windows 8 and higher. Copying between two different shares is not server-side. Windows 7 clients support server-side copying with [Robocopy](#) (<https://technet.microsoft.com/en-us/library/cc733145>).

This chapter starts by summarizing the available configuration options. It demonstrates some common configuration scenarios as well as offering some troubleshooting tips. It is recommended to first read through this entire chapter before creating any SMB shares to get a better idea of the configuration scenario that best meets your network's needs.

Tip: [SMB Tips and Tricks](https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/) (<https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/>) shows helpful hints for configuring and managing SMB networking. The [FreeNAS and Samba \(CIFS\) permissions](https://www.youtube.com/watch?v=RxggaE935PM) (<https://www.youtube.com/watch?v=RxggaE935PM>) and [Advanced Samba \(CIFS\) permissions on FreeNAS](https://www.youtube.com/watch?v=QhwOyLtArw0) (<https://www.youtube.com/watch?v=QhwOyLtArw0>) videos clarify setting up permissions on SMB shares. Another helpful reference is [Methods For Fine-Tuning Samba Permissions](https://forums.freenas.org/index.php?threads/methods-for-fine-tuning-samba-permissions.50739/) (<https://forums.freenas.org/index.php?threads/methods-for-fine-tuning-samba-permissions.50739/>).

Tip: Run `smbstatus` from the [Shell](#) (page 290) for a list of active connections and users.

Figure 10.13 shows the configuration screen that appears after clicking *Sharing* → *Windows (SMB Shares)* → *Add Windows (SMB) Share*.

Add Windows (SMB) Share

Path: **Browse**

Use as home share: ☐

Name:

Apply Default Permissions: ☒ ⓘ

Allow Guest Access: ☐ ⓘ

OK **Cancel** **Advanced Mode**

Fig. 10.13: Adding an SMB Share

Table 10.4 summarizes the options when creating a SMB share. Some settings are only available after clicking the Advanced Mode button. For simple sharing scenarios, Advanced Mode options are not needed. For more complex sharing scenarios, only change an Advanced Mode option after fully understanding the function of that option. [smb.conf\(5\)](https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>) provides more details for each configurable option.

Table 10.4: Options for a SMB Share

Setting	Value	Advanced Mode	Description
Path	browse button		select volume/dataset/directory to share
Use as home share	checkbox		check this box if the share is meant to hold user home directories; only one share can be the homes share
Name	string		mandatory; name of share
Comment	string	✓	optional description
Apply Default Permissions	checkbox		sets the ACLs to allow read/write for owner/group and read-only for others; should only be unchecked when creating a share on a system that already has custom ACLs set
Export Read Only	checkbox	✓	prohibits write access to the share
Browsable to Network Clients	checkbox	✓	when checked, users see the contents of /homes (including other home directories of other users) and when unchecked, users see only their own home directory
Export Recycle Bin	checkbox	✓	deleted files are moved to a hidden <code>.recycle</code> in the root folder of the share; the <code>.recycle</code> directory can be deleted to reclaim space and is automatically recreated when a file is deleted
Show Hidden Files	checkbox	✓	if enabled, the Windows hidden attribute is not set when file-names that begin with a dot (a Unix hidden file) are created; existing files are not affected

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Allow Guest Access	checkbox		if checked, a password is not required to connect to the share; connections with a bad password are rejected unless the user account does not exist, in which case it is mapped to the guest account and granted the permissions of the guest user defined in the SMB (page 230) service
Only Allow Guest Access	checkbox	✓	requires Allow guest access to also be checked; forces guest access for all connections
Access Based Share Enumeration	checkbox	✓	when checked, users can only see the shares they have permission to access; to change the default that grants Everyone access, use the computer management MMC on Windows or the sharesec command-line utility
Hosts Allow	string	✓	comma-, space-, or tab-delimited list of allowed hostnames or IP addresses
Hosts Deny	string	✓	comma-, space-, or tab-delimited list of denied hostnames or IP addresses; allowed hosts take precedence so can use ALL in this field and specify allowed hosts in Hosts Allow
VFS Objects	selection	✓	adds virtual file system modules to enhance functionality; Table 10.5 summarizes the available modules
Periodic Snapshot Task	drop-down menu	✓	used to configure directory shadow copies on a per-share basis; select the pre-configured periodic snapshot task to use for the share's shadow copies; periodic snapshot must be recursive
Auxiliary Parameters	string	✓	additional <code>smb4.conf</code> parameters not covered by other option fields

Note the following regarding some of the Advanced Mode settings:

- Hostname lookups add some time to accessing the SMB share. If you only use IP addresses, uncheck the Hostnames lookups box in `Services` → `SMB`.
- When the Browseable to Network Clients box is checked (the default), the share is visible through Windows File Explorer or through `net view`. When the Use as a home share box is checked, unchecking the Browseable to Network Clients box hides the share named homes so that only the dynamically generated share containing the authenticated user's home directory will be visible. By default, the homes share and the user's home directory are both visible. Users are not automatically granted read or write permissions on browsable shares. This option provides no real security because shares that are not visible in Windows File Explorer can still be accessed with a UNC path.
- If some files on a shared volume should be hidden and inaccessible to users, put a `veto files=` line in the Auxiliary Parameters field. The syntax for the veto files option and some examples can be found in the [smb.conf manual page](#) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>).

Samba disables NTLMv1 authentication by default for security. Standard configurations of Windows XP and some configurations of later clients like Windows 7 will not be able to connect with NTLMv1 disabled. [Security guidance for NTLMv1 and LM network authentication](#) (<https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-network-authentication>) has information about the security implications and ways to enable NTLMv2 on those clients. If changing the client configuration is not possible, NTLMv1 authentication can be enabled by checking the box NTLMv1 auth in `Services` → `SMB`.

[Table 10.5](#) provides an overview of the available VFS modules. Be sure to research each module before adding or deleting it from the Selected column of the VFS Objects field of the share. Some modules need additional configuration after they are added. Refer to [Stackable VFS modules](#) (<https://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/VFS.html>) and the [vfs_* man pages](#) (<https://www.samba.org/samba/docs/man/manpages/>) for more details.

Table 10.5: Available VFS Modules

Value	Description
acl_tdb	stores NTFS ACLs in a tdb file to enable full mapping of Windows ACLs
acl_xattr	stores NTFS ACLs in Extended Attributes (EAs) to enable the full mapping of Windows ACLs
aio_fork	enables async I/O
aio_pthread	implements async I/O in Samba vfs using a pthread pool instead of the internal Posix AIO interface
audit	logs share access, connects/disconnects, directory opens/creates/removes, and file opens/closes/renames/unlinks/chmods to syslog
cacheprime	primes the kernel file data cache
cap	translates filenames to and from the CAP encoding format, commonly used in Japanese language environments
catia	improves Mac interoperability by translating characters that are unsupported by Windows
commit	tracks the amount of data written to a file and synchronizes it to disk when a specified amount accumulates
crossrename	allows server side rename operations even if source and target are on different physical devices
default_quota	stores the default quotas that are reported to a windows client in the quota record of a user
dfs_samba4	distributed file system for providing an alternative name space, load balancing, and automatic failover
dirsort	sorts directory entries alphabetically before sending them to the client
expand_msdfs	enables support for Microsoft Distributed File System (DFS)
extd_audit	sends audit logs to both syslog and the Samba log files
fake_acls	stores file ownership and ACLs as extended attributes
fake_perms	allows roaming profile files and directories to be set as read-only
fruit	enhances OS X support by providing the SMB2 AAPL extension and Netatalk interoperability; automatically loads catia and streams_xattr but read the caveat in NOTE below table
full_audit	record selected client operations to the system log; if selected, a warning will indicate that Windows 10 clients may experience issues when transferring files to the NAS system when this module is enabled
linux_xfs_sgid	used to work around an old Linux XFS bug
media_harmony	allows Avid editorial workstations to share a network drive
netatalk	eases the co-existence of SMB and AFP shares
offline	marks all files in the share with the DOS offline attribute; this can prevent Windows Explorer from reading files just to make thumbnail images
posix_eadb	provides Extended Attributes (EAs) support so they can be used on filesystems which do not provide native support for EAs

Continued on next page

Table 10.5 – continued from previous page

Value	Description
preopen	useful for video streaming applications that want to read one file per frame
readahead	useful for Windows Vista clients reading data using Windows Explorer
readonly	marks a share as read-only for all clients connecting within the configured time period
shadow_copy	allows Microsoft shadow copy clients to browse shadow copies on Windows shares
shadow_copy_test	shadow copy testing
shell_snap	provides shell-script callouts for snapshot creation and deletion operations issued by remote clients using the File Server Remote VSS Protocol (FSRVP)
skel_opaque	implements dummy versions of all VFS modules (useful to VFS module developers)
skel_transparent	implements dummy passthrough functions of all VFS modules (useful to VFS module developers)
snapper	provides the ability for remote SMB clients to access shadow copies of FSRVP snapshots using Windows Explorer
streams_depot	experimental module to store alternate data streams in a central directory; the association with the primary file can be lost due to inode numbers changing when a directory is copied to a new location (see http://marc.info/?l=samba&m=132542069802160&w=2)
streams_xattr	enables storing of NTFS alternate data streams in the file system
syncops	ensures metadata operations are performed synchronously
time_audit	logs system calls that take longer than the number of defined milliseconds
unityed_media	allows multiple Avid clients to share a network drive
winmsa	emulate Microsoft's MoveSecurityAttributes=0 registry option, setting the ACL for file and directory hierarchies to inherit from the parent directory into which they are moved
worm	controls the writability of files and folders depending on their change time and an adjustable grace period
xattr_tdb	stores Extended Attributes (EAs) in a tdb file so they can be used on filesystems which do not provide support for EAs
zfs_space	correctly calculates ZFS space used by the share, including space used by ZFS snapshots, quotas, and reservations; enabled by default
zfsacl	provide ACL extensions for proper integration with ZFS; enabled by default

Note: Be careful when using multiple SMB shares, some with and some without fruit. OS X clients negotiate SMB2 AAPL protocol extensions on the first connection to the server, so mixing shares with and without fruit will globally disable AAPL if the first connection occurs without fruit. To resolve this, all OS X clients need to disconnect from all SMB shares and the first reconnection to the server has to be to a fruit-enabled share.

These VFS objects do not appear in the selection box:

- `recycle`: moves deleted files to the recycle directory instead of deleting them. Controlled by Export Recycle Bin in the [SMB share options](#) (page 186).
- `shadow_copy2`: a more recent implementation of `shadow_copy` with some additional features. `shadow_copy2` and the associated parameters are automatically added to the `smb4.conf` when a Periodic Snapshot Task is selected.

10.4.1 Configuring Unauthenticated Access

SMB supports guest logins, meaning that users can access the SMB share without needing to provide a username or password. This type of share is convenient as it is easy to configure, easy to access, and does not require any users to be configured on the FreeNAS® system. This type of configuration is also the least secure as anyone on the network can access the contents of the share. Additionally, since all access is as the guest user, even if the user inputs a username or password, there is no way to differentiate which users accessed or modified the data on the share. This type of configuration is best suited for small networks where quick and easy access to the share is more important than the security of the data on the share.

To configure an unauthenticated SMB share, click Wizard, then click the Next button twice to display the screen shown in [Figure 10.14](#). Complete the following fields in this screen:

1. Share name: enter a name for the share that is useful to you. In this example, the share is named `smb_insecure`.
2. Click the button for Windows (SMB) and check the box for Allow Guest.
3. Click the Ownership button. Click the drop-down User menu and select nobody. Click the Return button to return to the previous screen.
4. Click the Add button. If you forget to do this, the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

Wizard

Share name:

Purpose

☒ Windows (SMB) ☒ Allow Guest
☐ Mac OS X (AFP) ☐ Time Machine
☐ Generic Unix (NFS)
☐ Block Storage (iSCSI) Size:

Ownership

Add **Delete** **Update**

Name
smb_insecure

Previous **Next** **Exit**

Fig. 10.14: Creating an Unauthenticated SMB Share

Click the Next button twice, then the Confirm button to create the share. The Wizard automatically creates a dataset for the share and starts the SMB service so the share is immediately available. The new share is also be added to *Sharing* → *Windows (SMB)*.

Users can now access the share from any SMB client and will not be prompted for their username or password. For example, to access the share from a Windows system, open Explorer and click on Network. For this configuration example, a system named FREENAS appears with a share named insecure_smb. The user can copy data to and from the unauthenticated SMB share.

10.4.2 Configuring Authenticated Access Without a Domain Controller

Most configuration scenarios require each user to have their own user account and to authenticate before accessing the share. This allows the administrator to control access to data, provide appropriate permissions to that data, and to determine who accesses and modifies stored data. A Windows domain controller is not needed for authenticated SMB shares, which means that additional licensing costs are not required. However, since there is no domain controller to provide authentication for the network, each user account needs to be created on the FreeNAS® system. This type of configuration scenario is often used in home and small networks as it does not scale well if many users accounts are needed.

Before configuring this scenario, determine which users will need authenticated access. While not required for the configuration, it eases troubleshooting if the username and password that will be created on the FreeNAS® system matches that information on the client system. Next, determine if each user should have their own share to store their own data or if several users will be using the same share. The simpler configuration is to make one share per user as it does not require the creation of groups, adding the correct users to the groups, and ensuring that group permissions are set correctly.

To use the Wizard to create an authenticated SMB share, enter the following information, as shown in the example in [Figure 10.15](#).

1. Share name: enter a name for the share that is useful to you. In this example, the share is named smb_user1.
2. Click the button for Windows (SMB).
3. Click the Ownership button. To create the user account on the FreeNAS® system, type their name into the User field and check the Create User checkbox. The user's password is then entered and confirmed. If the user will not be sharing this share with other users, type their name into the Group field and click Create Group. If, however, the share will be used by several users, instead type in a group name and check the Create Group box. In the example shown in [Figure 10.16](#), user1 has been used for both the user and group name, meaning that this share will only be used by user1. When finished, click Return to return to the screen shown in [Figure 10.15](#).
4. Click the Add button. If you forget to do this, the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

If you wish to configure multiple authenticated shares, repeat for each user, giving each user their own Share name and Ownership. When finished, click Next twice, then Confirm to create the shares. The Wizard automatically creates a dataset with the correct ownership for each share and starts the SMB service so the shares are available immediately. The new shares are also added to `Sharing → Windows (SMB)`.

The screenshot shows the 'Wizard' window with a dark title bar. Inside, the 'Share name' field contains 'smb_user1'. Below it, the 'Purpose' section has four radio buttons: 'Windows (SMB)' (selected), 'Mac OS X (AFP)', 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' (checked) and 'Time Machine' (unchecked). A 'Size' field is next to the iSCSI option. An 'Ownership' button is to the right of the Purpose section. Below the Purpose section are three buttons: 'Add', 'Delete', and 'Update'. A list box labeled 'Name' contains the entry 'smb_user1'. At the bottom are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 10.15: Creating an Authenticated SMB Share

Wizard

User: user1 ☒ Create User ⓘ

User Password: ••••••••

Confirm User Password: ••••••••

Group: user1 ☒ Create Group ⓘ

Mode:

	Owner	Group	Other
Read	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Write	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Execute	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Fig. 10.16: Creating the User and Group

The authenticated share can now be tested from any SMB client. For example, to test an authenticated share from a Windows system, open Explorer and click on Network. For this configuration example, a system named FREENAS appears with a share named smb_user1. If you click on smb_user1, a Windows Security pop-up screen prompts for that user's username and password. Enter the values that were configured for that share, in this case user user1. After authentication, the user can copy data to and from the SMB share.

To prevent Windows Explorer from hanging when accessing the share, map the share as a network drive. To do this, right-click the share and select Map network drive.... Choose a drive letter from the drop-down menu and click the Finish button.

Note that Windows systems cache a user's credentials. This can cause issues when testing or accessing multiple authenticated shares as only one authentication is allowed at a time. If you are having problems authenticating to a share and are sure that you are entering the correct username and password, type cmd in the Search programs and files box and use the following command to see if you have already authenticated to a share. In this example, the user has already authenticated to the smb_user1 share:

```
net use
New connections will be remembered.

Status          Local    Remote                               Network
-----
OK              \\FREENAS\smb_user1 Microsoft Windows Network
The command completed successfully.
```

To clear the cache:

```
net use * /DELETE
You have these remote connections:
    \\FREENAS\smb_user1
Continuing will cancel the connections.

Do you want to continue this operation? <Y/N> [N]: y
```

An additional warning is shown if the share is currently open in Explorer:

```
There are open files and/or incomplete directory searches pending on the connection
to \\FREENAS\smb_user1.

Is it OK to continue disconnecting and force them closed? <Y/N> [N]: y
The command completed successfully.
```

The next time a share is accessed with Explorer, you will be prompted to authenticate.

10.4.3 Configuring Shadow Copies

Shadow Copies (https://en.wikipedia.org/wiki/Shadow_copy), also known as the Volume Shadow Copy Service (VSS) or Previous Versions, is a Microsoft service for creating volume snapshots. Shadow copies allow you to easily restore previous versions of files from within Windows Explorer. Shadow Copy support is built into Vista and Windows 7. Windows XP or 2000 users need to install the **Shadow Copy client** (<http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=16220>).

When you create a periodic snapshot task on a ZFS volume that is configured as a SMB share in FreeNAS®, it is automatically configured to support shadow copies.

Before using shadow copies with FreeNAS®, be aware of the following caveats:

- If the Windows system is not fully patched to the latest service pack, Shadow Copies may not work. If you are unable to see any previous versions of files to restore, use Windows Update to make sure that the system is fully up-to-date.
- Shadow copy support only works for ZFS pools or datasets. This means that the SMB share must be configured on a volume or dataset, not on a directory.
- Datasets are filesystems and shadow copies cannot traverse filesystems. If you want to be able to see the shadow copies in your child datasets, create separate shares for them.
- Shadow copies will not work with a manual snapshot, you must create a periodic snapshot task for the pool or dataset being shared by SMB or a recursive task for a parent dataset.
- The periodic snapshot task should be created and at least one snapshot should exist before creating the SMB share. If the SMB share was created first, restart the SMB service in `Services → Control Services`.
- Appropriate permissions must be configured on the volume/dataset being shared by SMB.
- Users cannot delete shadow copies on the Windows system due to the way Samba works. Instead, the administrator can remove snapshots from the FreeNAS® administrative GUI. The only way to disable shadow copies completely is to remove the periodic snapshot task and delete all snapshots associated with the SMB share.

To configure shadow copy support, use the instructions in [Configuring Authenticated Access Without a Domain Controller](#) (page 191) to create the desired number of shares. In this configuration example, a Windows 7 computer has two users: user1 and user2. For this example, two authenticated shares are created so that each user account has their own share. The first share is named user1 and the second share is named user2. Then:

1. Use `Storage → Periodic Snapshot Tasks → Add Periodic Snapshot` to create at least one periodic snapshot task. You can either create a snapshot task for each user's dataset, in this example the datasets `/mnt/volume1/user1` and `/mnt/volume1/user2`, or you can create one periodic snapshot task for the entire volume, in this case `/mnt/volume1`. Before continuing to the next step, confirm that at least one snapshot for each defined task is displayed in the `Storage → Snapshots` tab. When creating the schedule for the periodic snapshot tasks, keep in mind how often your users need to access modified files and during which days and time of day they are likely to make changes.
2. Go to `Sharing → Windows (SMB) Shares`. Highlight a share and click `Edit`, then `Advanced Mode`. Click the `Periodic Snapshot Task` drop-down menu and select the periodic snapshot task to use for that share. Repeat for each share being configured as a shadow copy. For this example, the share named `/mnt/volume1/user1` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/volume1/user1` dataset and the share named `/mnt/volume1/user2` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/volume1/user2` dataset.
3. Verify that the SMB service is set to ON in `Services → Control Services`.

Figure 10.17 provides an example of using shadow copies while logged in as user1 on the Windows system. In this example, the user right-clicked modified file and selected `Restore previous versions` from the menu. This particular file has three versions: the current version, plus two previous versions stored on the FreeNAS® system. The user can choose to open one of the previous versions, copy a previous version to the current folder, or restore one of the previous versions, overwriting the existing file on the Windows system.

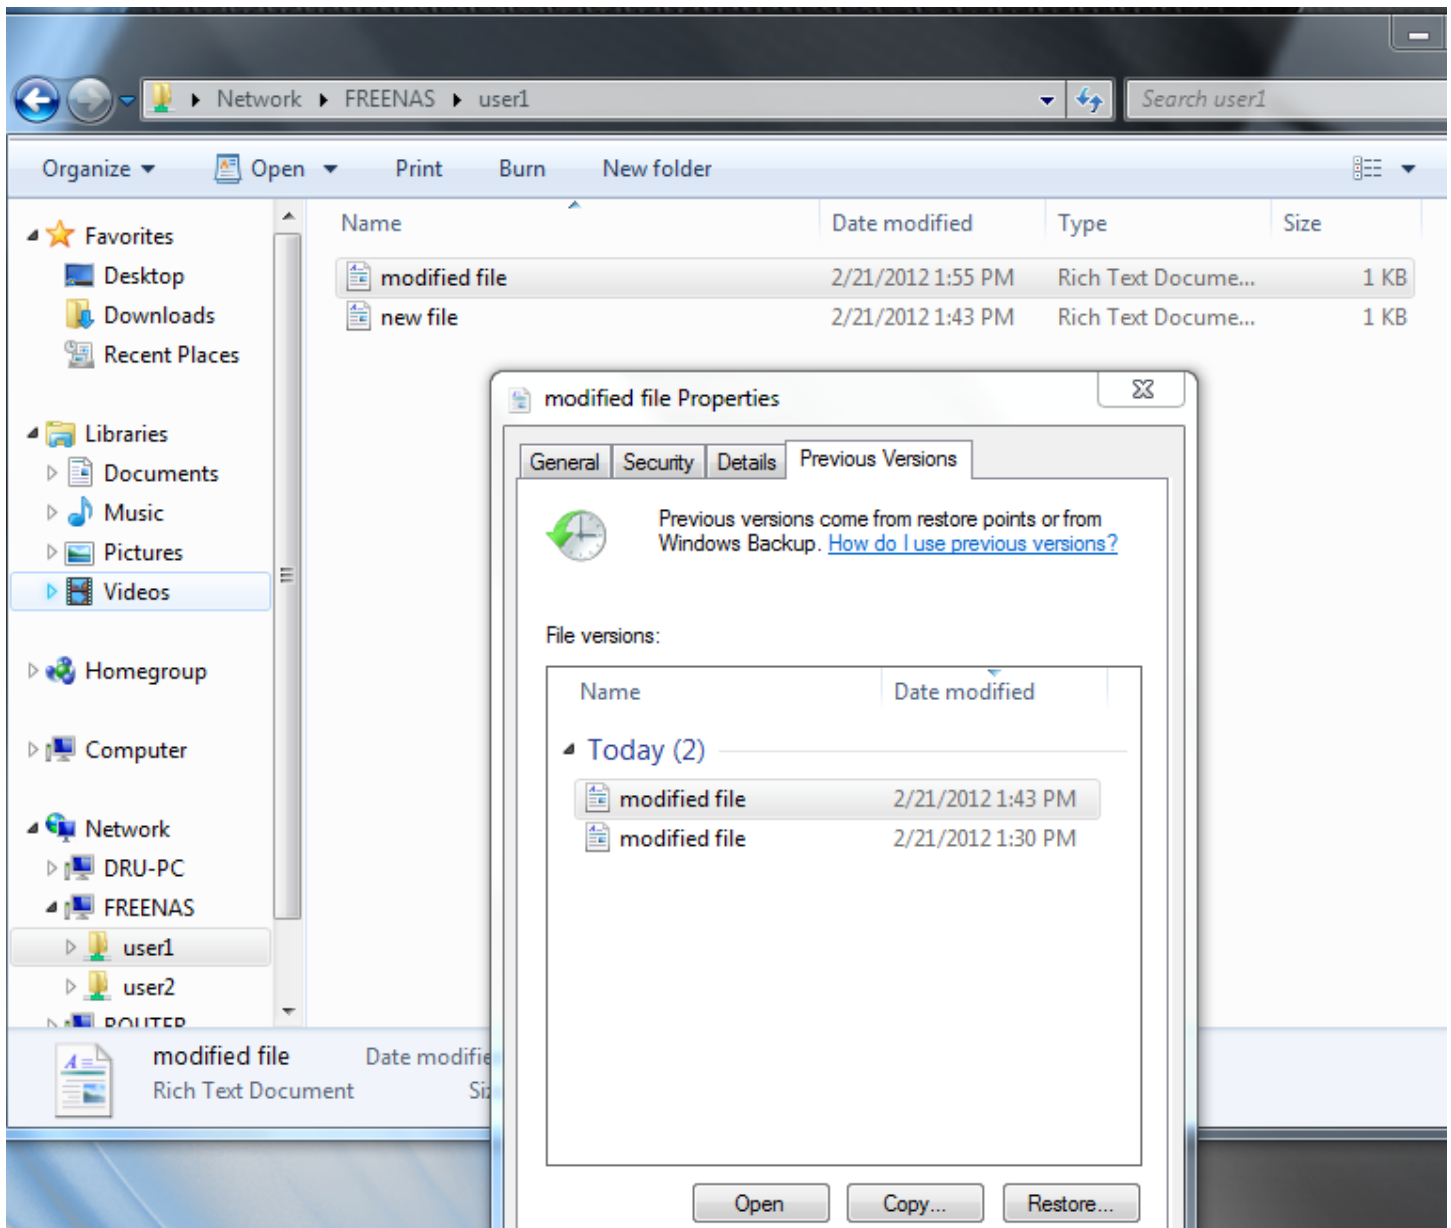


Fig. 10.17: Viewing Previous Versions within Explorer

10.5 Block (iSCSI)

iSCSI is a protocol standard for the consolidation of storage data. iSCSI allows FreeNAS® to act like a storage area network (SAN) over an existing Ethernet network. Specifically, it exports disk devices over an Ethernet network that iSCSI clients (called initiators) can attach to and mount. Traditional SANs operate over fibre channel networks which require a fibre channel infrastructure such as fibre channel HBAs, fibre channel switches, and discrete cabling. iSCSI can be used over an existing Ethernet network, although dedicated networks can be built for iSCSI traffic in an effort to boost performance. iSCSI also provides an advantage in an environment that uses Windows shell programs; these programs tend to filter "Network Location" but iSCSI mounts are not filtered.

Before configuring the iSCSI service, be familiar with this iSCSI terminology:

CHAP: an authentication method which uses a shared secret and three-way authentication to determine if a system is authorized to access the storage device and to periodically confirm that the session has not been hijacked by another system.

In iSCSI, the initiator (client) performs the CHAP authentication.

Mutual CHAP: a superset of CHAP in that both ends of the communication authenticate to each other.

Initiator: a client which has authorized access to the storage data on the FreeNAS® system. The client requires initiator software to initiate the connection to the iSCSI share.

Target: a storage resource on the FreeNAS® system. Every target has a unique name known as an iSCSI Qualified Name (IQN).

Internet Storage Name Service (iSNS): protocol for the automated discovery of iSCSI devices on a TCP/IP network.

Extent: the storage unit to be shared. It can either be a file or a device.

Portal: indicates which IP addresses and ports to listen on for connection requests.

LUN: Logical Unit Number representing a logical SCSI device. An initiator negotiates with a target to establish connectivity to a LUN. The result is an iSCSI connection that emulates a connection to a SCSI hard disk. Initiators treat iSCSI LUNs as if they were a raw SCSI or SATA hard drive. Rather than mounting remote directories, initiators format and directly manage filesystems on iSCSI LUNs. When configuring multiple iSCSI LUNs, create a new target for each LUN. Since iSCSI multiplexes a target with multiple LUNs over the same TCP connection, there can be TCP contention when more than one target accesses the same LUN. FreeNAS® supports up to 1024 LUNs.

In FreeNAS®, iSCSI is built into the kernel. This version of iSCSI supports [Microsoft Offloaded Data Transfer \(ODX\)](https://technet.microsoft.com/en-us/library/hh831628) (<https://technet.microsoft.com/en-us/library/hh831628>), meaning that file copies happen locally, rather than over the network. It also supports the [VAAI](#) (page 326) (vStorage APIs for Array Integration) primitives for efficient operation of storage tasks directly on the NAS. To take advantage of the VAAI primitives, create a zvol using the instructions in [Create zvol](#) (page 125) and use it to create a device extent, as described in [Extents](#) (page 203).

To configure iSCSI:

1. Review the target global configuration parameters.
2. Create at least one portal.
3. Determine which hosts are allowed to connect using iSCSI and create an initiator.
4. Decide if authentication will be used, and if so, whether it will be CHAP or mutual CHAP. If using authentication, create an authorized access.
5. Create a target.
6. Create either a device or a file extent to be used as storage.
7. Associate a target with an extent.
8. Start the iSCSI service in `Services → Control Services`.

The rest of this section describes these steps in more detail.

10.5.1 Target Global Configuration

`Sharing → Block (iSCSI) → Target Global Configuration`, shown in [Figure 10.18](#), contains settings that apply to all iSCSI shares. [Table 10.6](#) summarizes the settings that can be configured in the Target Global Configuration screen.

Some built-in values affect iSNS usage. Fetching of allowed initiators from iSNS is not implemented, so target ACLs must be configured manually. To make iSNS registration useful, iSCSI targets should have explicitly configured port IP addresses. This avoids initiators attempting to discover unconfigured target portal addresses like 0.0.0.0.

The iSNS registration period is 900 seconds. Registered Network Entities not updated during this period are unregistered. The timeout for iSNS requests is 5 seconds.

Sharing

Apple (AFP) UNIX (NFS) WebDAV Windows (SMB) **Block (iSCSI)**

Target Global Configuration Portals Initiators Authorized Access Targets Extents Associated Targets

Base Name: ⓘ

ISNS Servers: ⓘ

Pool Available Space Threshold (%): ⓘ

Save

Fig. 10.18: iSCSI Target Global Configuration Variables

Table 10.6: Target Global Configuration Settings

Setting	Value	Description
Base Name	string	see the “Constructing iSCSI names using the iqn. format” section of RFC 3721 (https://tools.ietf.org/html/rfc3721) if unfamiliar with this format
ISNS Servers	string	space delimited list of hostnames or IP addresses of ISNS servers with which to register the system’s iSCSI targets and portals
Pool Available Space Threshold	integer	enter the percentage of free space that should remain in the pool; when this percentage is reached, the system issues an alert, but only if zvols are used; see VAAI (page 326) Threshold Warning

10.5.2 Portals

A portal specifies the IP address and port number to be used for iSCSI connections. `Sharing → Block (iSCSI) → Portals → Add Portal` brings up the screen shown in [Figure 10.19](#).

[Table 10.19](#) summarizes the settings that can be configured when adding a portal. If you need to assign additional IP addresses to the portal, click the link `Add extra Portal IP`.

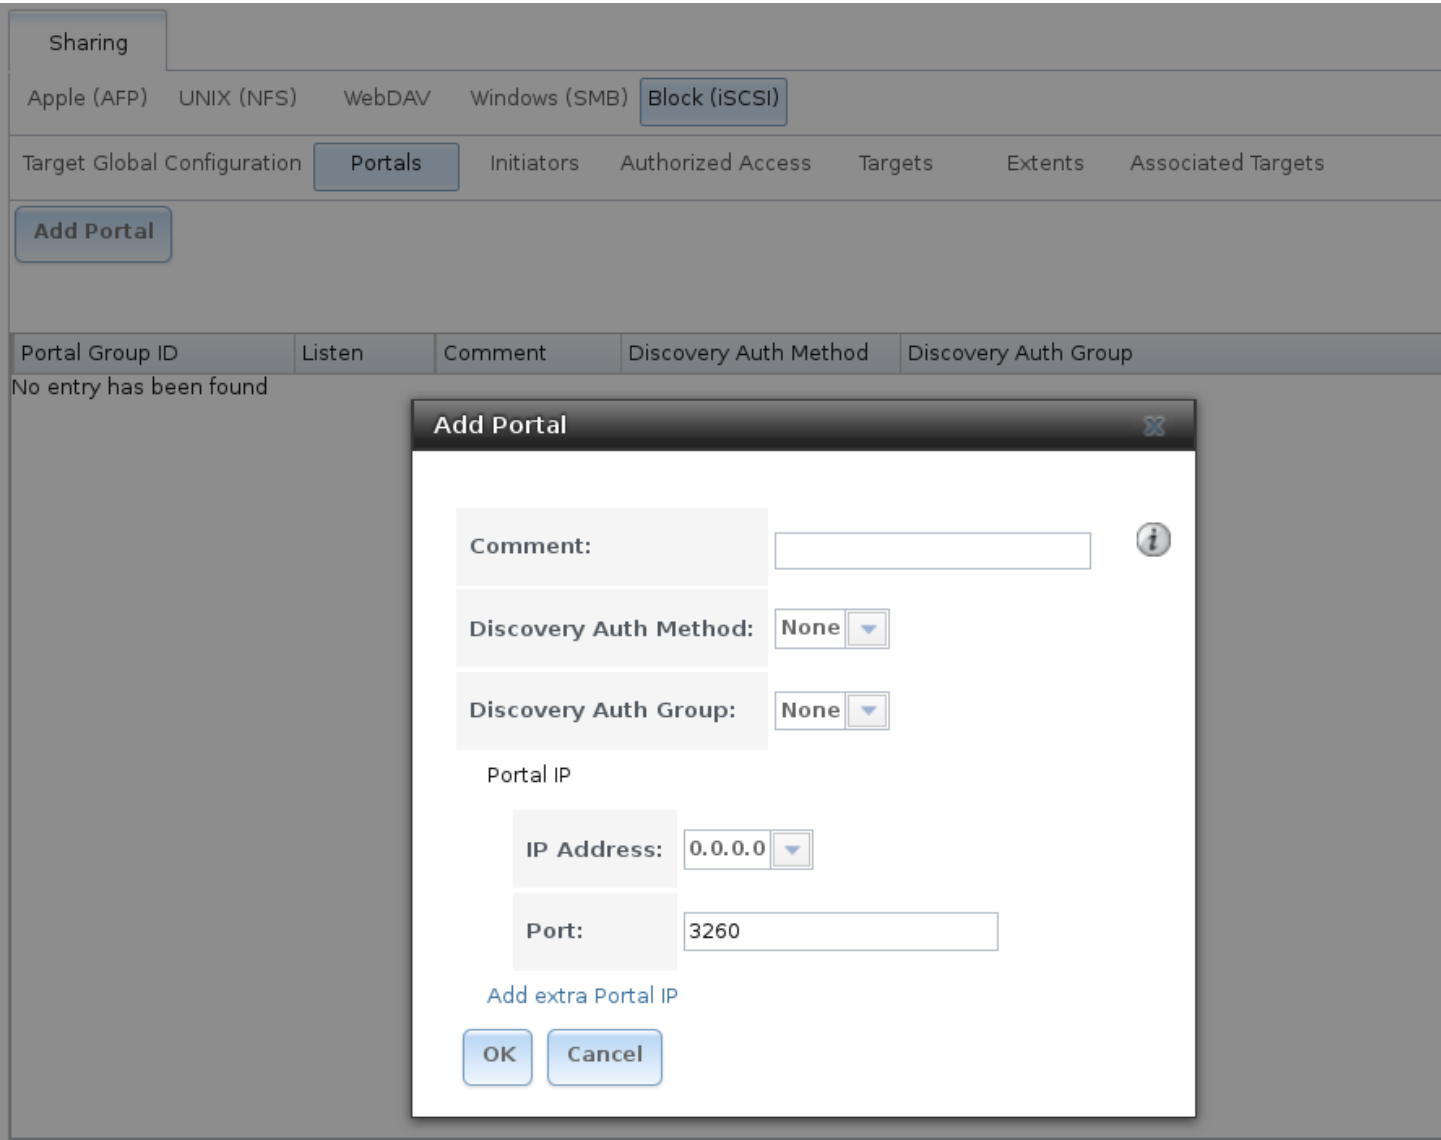


Fig. 10.19: Adding an iSCSI Portal

Table 10.7: Portal Configuration Settings

Setting	Value	Description
Comment	string	optional description; portals are automatically assigned a numeric group ID
Discovery Auth Method	drop-down menu	configures the authentication level required by the target for discovery of valid devices, where None will allow anonymous discovery while CHAP and Mutual CHAP require authentication
Discovery Auth Group	drop-down menu	select a user created in Authorized Access if the Discovery Auth Method is set to CHAP or Mutual CHAP
IP address	drop-down menu	select the IP address associated with an interface or the wildcard address of 0.0.0.0 (any interface)
Port	integer	TCP port used to access the iSCSI target; default is 3260

FreeNAS® systems with multiple IP addresses or interfaces can use a portal to provide services on different interfaces or subnets. This can be used to configure multi-path I/O (MPIO). MPIO is more efficient than a link aggregation.

If the FreeNAS® system has multiple configured interfaces, portals can also be used to provide network access control. For example, consider a system with four interfaces configured with the following addresses:

192.168.1.1/24

192.168.2.1/24

192.168.3.1/24

192.168.4.1/24

You could create a portal containing the first two IP addresses (group ID 1) and a portal containing the remaining two IP addresses (group ID 2). You could then create a target named A with a Portal Group ID of 1 and a second target named B with a Portal Group ID of 2. In this scenario, the iSCSI service would listen on all four interfaces, but connections to target A would be limited to the first two networks and connections to target B would be limited to the last two networks.

Another scenario would be to create a portal which includes every IP address except for the one used by a management interface. This would prevent iSCSI connections to the management interface.

10.5.3 Initiators

The next step is to configure authorized initiators, or the systems which are allowed to connect to the iSCSI targets on the FreeNAS® system. To configure which systems can connect, use **Sharing → Block (iSCSI) → Initiators → Add Initiator**, shown in [Figure 10.20](#).

Fig. 10.20: Adding an iSCSI Initiator

[Table 10.8](#) summarizes the settings that can be configured when adding an initiator.

Table 10.8: Initiator Configuration Settings

Setting	Value	Description
Initiators	string	use ALL keyword or a list of initiator hostnames separated by spaces
Authorized network	string	use ALL keyword or a network address with CIDR mask such as 192.168.2.0/24
Comment	string	optional description

In the example shown in [Figure 10.21](#), two groups have been created. Group 1 allows connections from any initiator on any network. Group 2 allows connections from any initiator on the 10.10.1.0/24 network. Click an initiator's entry to display its Edit and Delete buttons.

Note: Attempting to delete an initiator causes a warning that indicates if any targets or target/extent mappings depend upon the initiator. Confirming the delete causes these to be deleted also.

Sharing

Apple (AFP)UNIX (NFS)WebDAVWindows (SMB)Block (iSCSI)

Target Global ConfigurationPortalsInitiatorsAuthorized AccessTargetsExtentsAssociated Targets

Add Initiator

Group ID	Initiators	Authorized network	Comment
1	ALL	ALL	
2	ALL	10.10.1.0/24	

Fig. 10.21: Sample iSCSI Initiator Configuration

10.5.4 Authorized Accesses

If you will be using CHAP or mutual CHAP to provide authentication, you must create an authorized access in Sharing → Block (iSCSI) → Authorized Accesses → Add Authorized Access. This screen is shown in Figure 10.22.

Note: This screen sets login authentication. This is different from discovery authentication which is set in Target Global Configuration (page 196).

Add Authorized Access

Group ID: 1

User:

Secret:

Secret (Confirm):

Peer User:

Peer Secret:

Peer Secret (Confirm):

OK Cancel

Fig. 10.22: Adding an iSCSI Authorized Access

Table 10.9 summarizes the settings that can be configured when adding an authorized access:

Table 10.9: Authorized Access Configuration Settings

Setting	Value	Description
Group ID	integer	allows different groups to be configured with different authentication profiles; for instance, all users with a Group ID of 1 will inherit the authentication profile associated with Group 1
User	string	name of user account to create for CHAP authentication with the user on the remote system; many initiators default to using the initiator name as the user
Secret	string	password to be associated with User; the iSCSI standard requires that this be between 12 and 16 characters
Peer User	string	only input when configuring mutual CHAP; in most cases it will need to be the same value as User
Peer Secret	string	the mutual secret password which must be different than the Secret; required if Peer User is set

Note: CHAP does not work with GlobalSAN initiators on Mac OS X.

As authorized accesses are added, they will be listed under View Authorized Accesses. In the example shown in Figure 10.23, three users (test1, test2, and test3) and two groups (1 and 2) have been created, with group 1 consisting of one CHAP user and group 2 consisting of one mutual CHAP user and one CHAP user. Click an authorized access entry to display its Edit and Delete buttons.

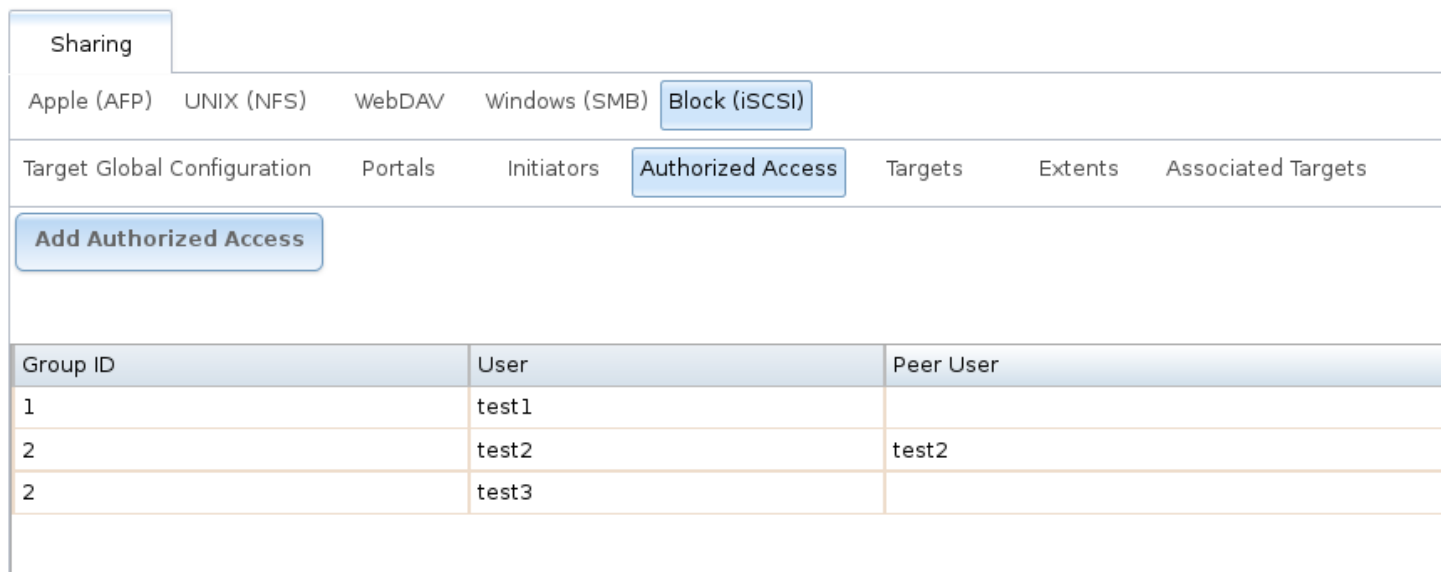


Fig. 10.23: Viewing Authorized Accesses

10.5.5 Targets

Next, create a Target using `Sharing → Block (iSCSI) → Targets → Add Target`, as shown in [Figure 10.24](#). A target combines a portal ID, allowed initiator ID, and an authentication method. [Table 10.10](#) summarizes the settings that can be configured when creating a Target.

Note: An iSCSI target creates a block device that may be accessible to multiple initiators. A clustered filesystem is required on the block device, such as VMFS used by VMware ESX/ESXi, in order for multiple initiators to mount the block device read/write. If a traditional filesystem such as EXT, XFS, FAT, NTFS, UFS, or ZFS is placed on the block device, care must be taken that only one initiator at a time has read/write access or the result will be filesystem corruption. If multiple clients need access to the same data on a non-clustered filesystem, use SMB or NFS instead of iSCSI, or create multiple iSCSI targets (one per client).

Fig. 10.24: Adding an iSCSI Target

Table 10.10: Target Settings

Setting	Value	Description
Target Name	string	required value; base name will be appended automatically if it does not start with iqn
Target Alias	string	optional user-friendly name
Portal Group ID	drop-down menu	leave empty or select number of existing portal to use
Initiator Group ID	drop-down menu	select which existing initiator group has access to the target
Auth Method	drop-down menu	choices are None, Auto, CHAP, or Mutual CHAP
Authentication Group number	drop-down menu	None or integer representing number of existing authorized access

10.5.6 Extents

iSCSI targets provide virtual access to resources on the FreeNAS® system. Extents are used to define resources to share with clients. There are two types of extents: device and file.

Device extents provide virtual storage access to zvols, zvol snapshots, or physical devices like a disk, an SSD, a hardware RAID volume, or a [HAST device](http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html) (http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html).

File extents provide virtual storage access to an individual file.

Tip: For typical use as storage for virtual machines where the virtualization software is the iSCSI initiator, device extents with zvols provide the best performance and most features. For other applications, device extents sharing a raw device can be appropriate. File extents do not have the performance or features of device extents, but do allow creating multiple extents

on a single filesystem.

Virtualized zvols support all the FreeNAS® [VAAI](#) (page 326) primitives and are recommended for use with virtualization software as the iSCSI initiator.

The ATS, WRITE SAME, XCOPY and STUN, primitives are supported by both file and device extents. The UNMAP primitive is supported by zvols and raw SSDs. The threshold warnings primitive is fully supported by zvols and partially supported by file extents.

Virtualizing a raw device like a single disk or hardware RAID volume limits performance to the abilities of the device. Because this bypasses ZFS, such devices do not benefit from ZFS caching or provide features like block checksums or snapshots.

Virtualizing a zvol adds the benefits of ZFS, such as read and write cache. Even if the client formats a device extent with a different filesystem, the data still resides on a ZFS volume and benefits from ZFS features like block checksums and snapshots.

Warning: For performance reasons and to avoid excessive fragmentation, keep the used space of the pool below 50% when using iSCSI. The capacity of an existing extent can be increased as shown in [Growing LUNs](#) (page 207).

To add an extent, go to `Sharing → Block (iSCSI) → Extents → Add Extent`. In the example shown in [Figure 10.25](#), the device extent is using the `export` zvol that was previously created from the `/mnt/volume1` volume.

[Table 10.11](#) summarizes the settings that can be configured when creating an extent. Note that file extent creation will fail if you do not append the name of the file to be created to the volume/dataset name.

Add Extent

Extent Name:

Extent Type: Device

Device: ada1 (10.0 GiB)

Serial:

Logical Block Size: 512

Disable Physical Block Size Reporting: ☐

Comment:

Enable TPC: ☒

Xen initiator compat mode: ☐

LUN RPM: SSD

Read-only: ☐

String identifier of the extent.

Fig. 10.25: Adding an iSCSI Extent

Table 10.11: Extent Configuration Settings

Setting	Value	Description
Extent Name	string	name of extent; if the Extent size is not 0, it cannot be an existing file within the volume/dataset
Extent Type	drop-down menu	select from File or Device
Device	drop-down menu	only appears if Device is selected; select the unformatted disk, controller, zvol, zvol snapshot, or HAST device
Serial	string	unique LUN ID; the default is generated from the system's MAC address
Path to the extent	browse button	only appears if File is selected; browse to an existing file and use 0 as the Extent size, or browse to the volume or dataset, click Close, append the Extent Name to the path, and specify a value in Extent size; extents cannot be created inside the jail root directory
Extent size	integer	only appears if File is selected; if the size is specified as 0, the file must already exist and the actual file size will be used; otherwise, specify the size of the file to create
Logical Block Size	drop-down menu	only override the default if the initiator requires a different block size

Continued on next page

Table 10.11 – continued from previous page

Setting	Value	Description
Disable Physical Block Size Reporting	checkbox	if the initiator does not support physical block size values over 4K (MS SQL), check this box
Available Space Threshold	string	only appears if File or a zvol is selected; when the specified percentage of free space is reached, the system issues an alert; see VAAI (page 326) Threshold Warning
Comment	string	optional
Enable TPC	checkbox	if checked, an initiator can bypass normal access control and access any scannable target; this allows <code>xcopy</code> operations otherwise blocked by access control
Xen initiator compat mode	checkbox	check this box when using Xen as the iSCSI initiator
LUN RPM	drop-down menu	do NOT change this setting when using Windows as the initiator; only needs to be changed in large environments where the number of systems using a specific RPM is needed for accurate reporting statistics
Read-only	checkbox	check this box to prevent the initiator from initializing this LUN

10.5.7 Target/Extents

The last step is associating an extent to a target within `Sharing → Block (iSCSI) → Associated Targets → Add Target/Extent`. This screen is shown in [Figure 10.26](#). Use the drop-down menus to select the existing target and extent. Click OK to add an entry for the LUN.

Fig. 10.26: Associating a Target With an Extent

[Table 10.12](#) summarizes the settings that can be configured when associating targets and extents.

Table 10.12: Target/Extents Configuration Settings

Setting	Value	Description
Target	drop-down menu	select the pre-created target
LUN ID	drop-down menu	select the value to use or type in a value between 1 and 1023; note that some initiators expect a value below 256
Extent	drop-down menu	select the pre-created extent

It is recommended to always associate extents to targets in a one-to-one manner, even though the GUI will allow multiple extents to be associated with the same target.

Note: Each LUN entry has Edit and Delete buttons for modifying the settings or deleting the LUN entirely. A verification popup appears when the Delete button is clicked. If an initiator has an active connection to the LUN, it is indicated in red text. It is recommended to clear initiator connections to a LUN before deleting it.

After iSCSI has been configured, remember to start it in `Services → Control Services`. Click the red OFF button next to iSCSI. After a second or so, it will change to a blue ON, indicating that the service has started.

10.5.8 Connecting to iSCSI

To access the iSCSI target, clients must use iSCSI initiator software.

An iSCSI Initiator client is pre-installed with Windows 7. A detailed how-to for this client can be found [here](http://www.windowsnetworking.com/articles-tutorials/windows-7/Connecting-Windows-7-iSCSI-SAN.html) (<http://www.windowsnetworking.com/articles-tutorials/windows-7/Connecting-Windows-7-iSCSI-SAN.html>). A client for Windows 2000, XP, and 2003 can be found [here](http://www.microsoft.com/en-us/download/details.aspx?id=18986) (<http://www.microsoft.com/en-us/download/details.aspx?id=18986>). This [how-to](http://blog.pluralsight.com/freenas-8-iscsi-target-windows-7) (<http://blog.pluralsight.com/freenas-8-iscsi-target-windows-7>) shows how to create an iSCSI target for a Windows 7 system.

Mac OS X does not include an initiator. [globalSAN](http://www.studionetworksolutions.com/globalsan-iscsi-initiator/) (<http://www.studionetworksolutions.com/globalsan-iscsi-initiator/>) is a commercial, easy-to-use Mac initiator.

BSD systems provide command line initiators: `iscontrol(8)` (<http://www.freebsd.org/cgi/man.cgi?query=iscontrol>) comes with FreeBSD versions 9.x and lower, `iscsictl(8)` (<https://www.freebsd.org/cgi/man.cgi?query=iscsictl>) comes with FreeBSD versions 10.0 and higher, `iscsi-initiator(8)` (<http://netbsd.gw.com/cgi-bin/man-cgi?iscsi-initiator++NetBSD-current>) comes with NetBSD, and `iscsid(8)` (<http://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/.man8/iscsid.8?query=iscsid>) comes with OpenBSD.

Some Linux distros provide the command line utility `iscsiadm` from [Open-iSCSI](http://www.open-iscsi.com/) (<http://www.open-iscsi.com/>). Use a web search to see if a package exists for your distribution should the command not exist on your Linux system.

If a LUN is added while `iscsiadm` is already connected, it will not see the new LUN until rescanned with `iscsiadm -m node -R`. Alternately, use `iscsiadm -m discovery -t st -p portal_IP` to find the new LUN and `iscsiadm -m node -T LUN_Name -l` to log into the LUN.

Instructions for connecting from a VMware ESXi Server can be found at [How to configure FreeNAS 8 for iSCSI and connect to ESXi\(i\)](http://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/) (<http://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/>). Note that the requirements for booting vSphere 4.x off iSCSI differ between ESX and ESXi. ESX requires a hardware iSCSI adapter while ESXi requires specific iSCSI boot firmware support. The magic is on the booting host side, meaning that there is no difference to the FreeNAS® configuration. See the [iSCSI SAN Configuration Guide](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) (http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) for details.

The VMware firewall only allows iSCSI connections on port 3260 by default. If a different port has been selected, outgoing connections to that port must be manually added to the firewall before those connections will work.

If the target can be seen but does not connect, check the Discovery Auth settings in Target Global Configuration.

If the LUN is not discovered by ESXi, make sure that promiscuous mode is set to Accept in the vSwitch.

10.5.9 Growing LUNs

The method used to grow the size of an existing iSCSI LUN depends on whether the LUN is backed by a file extent or a zvol. Both methods are described in this section.

Enlarging a LUN with one of the methods below gives it more unallocated space, but does not automatically resize filesystems or other data on the LUN. This is the same as binary-copying a smaller disk onto a larger one. More space is available on the new disk, but the partitions and filesystems on it must be expanded to use this new space. Resizing virtual disk images is usually done from virtual machine management software. Application software to resize filesystems is dependent on the type of filesystem and client, but is often run from within the virtual machine. For instance, consider a Windows VM with the last partition on the disk holding an NTFS filesystem. The LUN is expanded and the partition table edited to add the new

space to the last partition. The Windows disk manager must still be used to resize the NTFS filesystem on that last partition to use the new space.

Zvol Based LUN

To grow a zvol based LUN, go to **Storage** → **Volumes** → **View Volumes**, highlight the zvol to be grown, and click **Edit zvol**. In the example shown in [Figure 10.27](#), the current size of the zvol named zvol1 is 4GB.

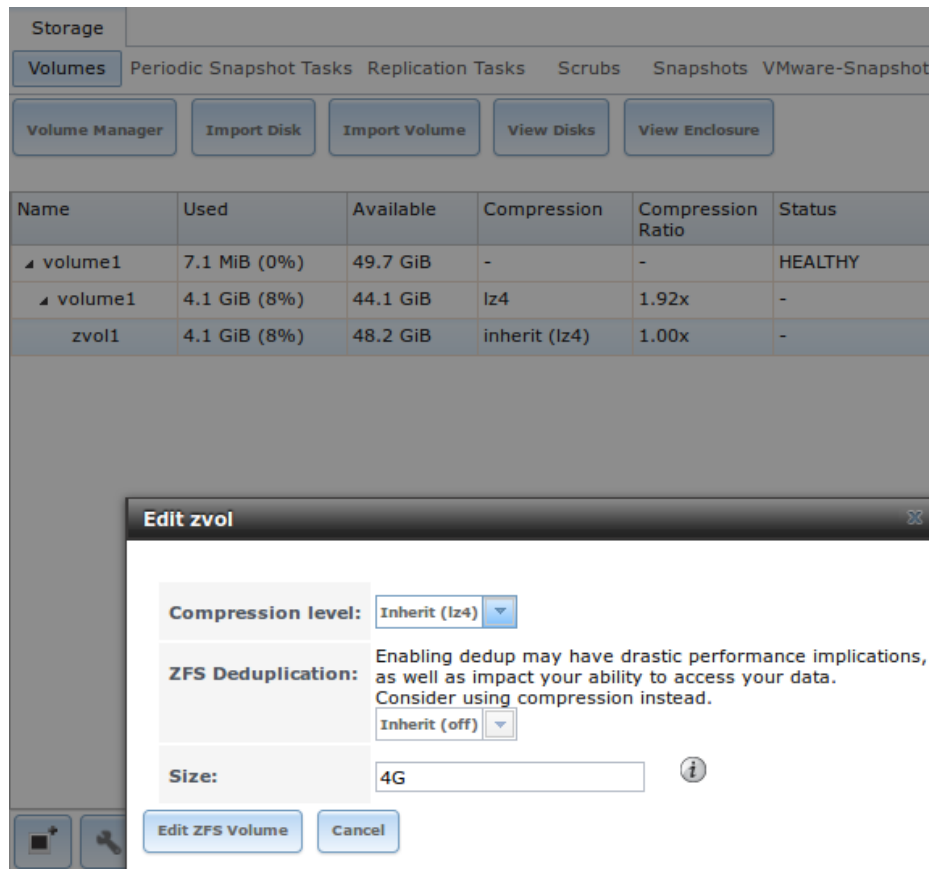


Fig. 10.27: Editing an Existing Zvol

Enter the new size for the zvol in the Size field and click **Edit ZFS Volume**. This menu closes and the new size for the zvol is immediately shown in the Used column of the View Volumes screen.

Note: The GUI does not allow reducing (shrinking) the size of the zvol, as doing so could result in loss of data. It also does not allow increasing the size of the zvol past 80% of the volume size.

File Extent Based LUN

To grow a file extent based LUN, go to **Services** → **iSCSI** → **File Extents** → **View File Extents** to determine the path of the file extent to grow. Open **Shell** to grow the extent. This example grows `/mnt/volume1/data` by 2 G:

```
truncate -s +2g /mnt/volume1/data
```

Go back to **Services** → **iSCSI** → **File Extents** → **View File Extents** and click the **Edit** button for the file extent. Set the size to 0 as this causes the iSCSI target to use the new size of the file.

SERVICES

The Services section of the GUI is where various services that ship with the FreeNAS® system are configured, started, or stopped. FreeNAS® includes these built-in services:

- [AFP](#) (page 211)
- [Domain Controller](#) (page 213)
- [Dynamic DNS](#) (page 215)
- [FTP](#) (page 216)
- [iSCSI](#) (page 221)
- [LLDP](#) (page 222)
- [Netdata](#) (page 222)
- [NFS](#) (page 223)
- [Rsync](#) (page 225)
- [S3](#) (page 227)
- [S.M.A.R.T.](#) (page 228)
- [SMB](#) (page 230)
- [SNMP](#) (page 234)
- [SSH](#) (page 236)
- [TFTP](#) (page 238)
- [UPS](#) (page 239)
- [WebDAV](#) (page 242)

This section demonstrates starting a FreeNAS® service and the available configuration options for each FreeNAS® service.

11.1 Control Services

`Services → Control Services`, shown in [Figure 11.1](#), shows which services are currently running and can start, stop, or configure them. The S.M.A.R.T. service is enabled by default, but only runs if the storage devices support [S.M.A.R.T. data](#) (<http://en.wikipedia.org/wiki/S.M.A.R.T.>) Other services default to off until started.

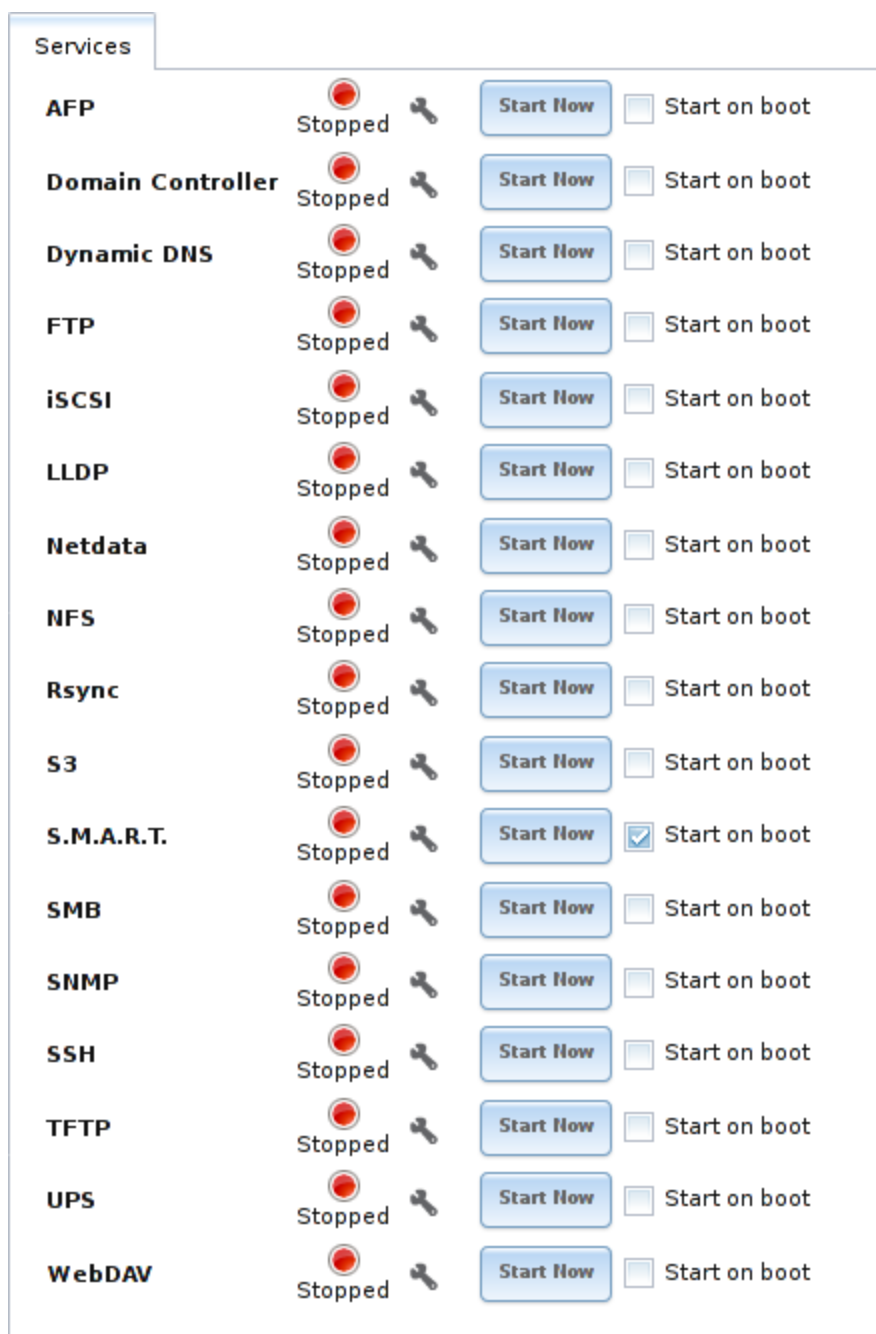


Fig. 11.1: Control Services

Stopped services show a red stop symbol and a Start Now button. Running services show a green light with a Stop Now button.

Tip: Using a proxy server can prevent the list of services from being displayed. If a proxy server is used, configure it to not proxy local network connections or websocket connections. VPN software can also cause problems. If the list of services is displayed when connecting on the local network but not when connecting through the VPN, check the VPN software configuration.

Services are configured by clicking the wrench icon or the name of the service in the Services section of the tree menu.

If a service does not start, go to `System` → `Advanced` and check the box `Show console messages` in the footer. Console

messages appear at the bottom of the browser. Clicking the console message area makes it into a pop-up window, allowing scrolling through or copying the messages. Watch these messages for errors when stopping or starting the problematic service.

To read the system logs for more information about a service failure, open [Shell](#) (page 290) and type `more /var/log/messages`.

11.2 AFP

The settings that are configured when creating AFP Shares in [Sharing](#) → [Apple \(AFP\) Shares](#) → [Add Apple \(AFP\) Share](#) are specific to each configured AFP Share. In contrast, global settings which apply to all AFP shares are configured in [Services](#) → [AFP](#).

[Figure 11.2](#) shows the available global AFP configuration options which are described in [Table 11.1](#).

Settings

Guest Access:

☐

Guest account:

nobody

Max. Connections:

50

Enable home directories:

☐

Home directories:

Browse

Home share name:

Home Share Time Machine:

☐

Database Path:

Browse

Global auxiliary parameters:

Map ACLs:

Rights

Chmod Request:

Preserve

Bind IP Addresses:

☐ 10.0.0.102

OK

Cancel

Fig. 11.2: Global AFP Configuration

Table 11.1: Global AFP Configuration Options

Setting	Value	Description
Guest Access	checkbox	if checked, clients will not be prompted to authenticate before accessing AFP shares
Guest account	drop-down menu	select account to use for guest access; the selected account must have permissions to the volume or dataset being shared
Max Connections	integer	maximum number of simultaneous connections
Continued on next page		

Table 11.1 – continued from previous page

Setting	Value	Description
Enable home directories	checkbox	if checked, any user home directories located under Home directories will be available over the share
Home directories	browse button	select the volume or dataset which contains user home directories
Home share name	string	overrides default home folder name with the specified value
Home Share Time Machine	checkbox	when checked, enables Time Machine Lock Stealing; Apple recommends that shares designated for Time Machine backups be used exclusively for Time Machine backups
Database Path	browse button	select the path to store the CNID databases used by AFP (default is the root of the volume); the path must be writable
Global auxiliary parameters	string	additional afp.conf(5) (http://netatalk.sourceforge.net/3.0/html/docs/afp.conf.5.html) parameters not covered elsewhere in this screen
Map ACLs	drop-down menu	choose mapping of effective permissions for authenticated users; Rights (default, Unix-style permissions), Mode (ACLs), or None
Chmod Request	drop-down menu	indicates how to handle ACLs; choices are Ignore, Preserve, or Simple
Bind IP Addresses	selection	specify the IP addresses to listen for FTP connections; highlight the desired IP addresses in the Available list and use the >> button to add to the Selected list

When configuring home directories, it is recommended to create a dataset to hold the home directories which contains a child dataset for each user. As an example, create a dataset named `volume1/homedirs` and browse to this dataset when configuring the Home directories field of the AFP service. Then, as you create each user, first create a child dataset for that user. For example, create a dataset named `volume1/homedirs/user1`. When you create the user1 user, browse to the `volume1/homedirs/user1` dataset in the Home Directory field of the Add New User screen.

11.2.1 Troubleshooting AFP

You can determine which users are connected to an AFP share by typing `afpusers`.

If Something wrong with the volume's CNID DB is shown, run this command from [Shell](#) (page 290), replacing the path to the problematic AFP share:

```
dbd -rf /path/to/share
```

This command may take a while, depending upon the size of the volume or dataset being shared. This command will wipe the CNID database and rebuild it from the CNIDs stored in the AppleDouble files.

11.3 Domain Controller

FreeNAS® can be configured to act either as the domain controller for a network or to join an existing [Active Directory](#) (page 158) network as a domain controller.

Note: This section demonstrates how to configure the FreeNAS® system to act as a domain controller. If the goal is to integrate with an existing [Active Directory](#) (page 158) network to access its authentication and authorization services, configure [Active Directory](#) (page 158) instead.

Be aware that configuring a domain controller is a complex process that requires a good understanding of how [Active Directory](#) (page 158) works. While `Services → Domain Controller` makes it easy to enter the needed settings into the administrative graphical interface, it is important to understand what those settings should be. Before beginning configuration, read through the [Samba AD DC HOWTO](https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO) (https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO). After FreeNAS®

is configured, use the RSAT utility from a Windows system to manage the domain controller. The Samba AD DC HOWTO includes instructions for installing and configuring RSAT.

Figure 11.3 shows the configuration screen for creating a domain controller and Table 11.2 summarizes the available options.

Fig. 11.3: Domain Controller Settings

Table 11.2: Domain Controller Configuration Options

Setting	Value	Description
Realm	string	capitalized DNS realm name
Domain	string	capitalized domain name
Server Role	drop-down menu	at this time, the only supported role is as the domain controller for a new domain
DNS Forwarder	string	IP address of DNS forwarder; required for recursive queries when SAMBA_INTERNAL is selected
Domain Forest Level	drop-down menu	choices are 2000, 2003, 2008, 2008_R2, 2012, or 2012_R2; refer to Understanding Active Directory Domain Services (AD DS) Functional Levels (https://technet.microsoft.com/en-us/library/understanding-active-directory-functional-levels(Ws.10).aspx) for details
Administrator password	string	password to be used for the Active Directory (page 158) administrator account
Kerberos Realm	drop-down menu	auto-populates with information from the Realm when the settings in this screen are saved

11.3.1 Samba Domain Controller Backup

A `samba_backup` script is available to back up Samba4 domain controller settings is available. From the [Shell](#) (page 290), run `/usr/local/bin/samba_backup --usage` to show the input options.

11.4 Dynamic DNS

Dynamic DNS (DDNS) is useful if the FreeNAS® system is connected to an ISP that periodically changes the IP address of the system. With dynamic DNS, the system can automatically associate its current IP address with a domain name, allowing you to access the FreeNAS® system even if the IP address changes. DDNS requires you to register with a DDNS service such as [DynDNS](http://dyn.com/dns/) (<http://dyn.com/dns/>).

[Figure 11.4](#) shows the DDNS configuration screen and [Table 11.3](#) summarizes the configuration options. The values for these fields are provided by the DDNS provider. After configuring DDNS, remember to start the DDNS service in [Services](#) → [Control Services](#).

Field	Type	Value	Info Icon
Provider	Dropdown	-----	No
CheckIP Server SSL	Checkbox	<input type="checkbox"/>	No
CheckIP Server	Text		Yes
CheckIP Path	Text		Yes
Use SSL	Checkbox	<input type="checkbox"/>	No
Domain name	Text		Yes
Username	Text	admin	No
Password	Text		No
Confirm Password	Text		No
Update Period	Text	300	Yes

Fig. 11.4: Configuring DDNS

Table 11.3: DDNS Configuration Options

Setting	Value	Description
Provider	drop-down menu	several providers are supported; if your provider is not listed, select Custom Provider and enter the information in the Custom Server and Custom Path fields
CheckIP Server SSL	string	when checked, HTTPS is used for the connection to the CheckIP Server
CheckIP Server	string	enter the name and port of the server that reports the external IP address, in the format server.name.org:port
CheckIP Path	string	enter the path that is requested by the CheckIP Server to determine the user's IP address
Use SSL		when checked, HTTPS is used for the connection to the server that updates the DNS record
Domain name	string	fully qualified domain name (e.g. yourname.dyndns.org)
Username	string	username used to logon to the provider and update the record
Password	string	password used to logon to the provider and update the record
Update period	integer	how often the IP is checked in seconds

When using `he.net`, enter the domain name for Username and enter the DDNS key generated for that domain's A entry at the `he.net` website for Password.

11.5 FTP

FreeNAS® uses the `proftpd` (<http://www.proftpd.org/>) FTP server to provide FTP services. Once the FTP service is configured and started, clients can browse and download data using a web browser or FTP client software. The advantage of FTP is that easy-to-use cross-platform utilities are available to manage uploads to and downloads from the FreeNAS® system. The disadvantage of FTP is that it is considered to be an insecure protocol, meaning that it should not be used to transfer sensitive files. If you are concerned about sensitive data, see [Encrypting FTP](#).

This section provides an overview of the FTP configuration options. It then provides examples for configuring anonymous FTP, specified user access within a chroot environment, encrypting FTP connections, and troubleshooting tips.

[Figure 11.5](#) shows the configuration screen for `Services` → `FTP`. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box `Show advanced fields by default` in `System` → `Advanced`.

The screenshot shows the 'FTP Settings' window. The fields and their values are: Port: 21, Clients: 5, Connections: 2, Login Attempts: 1, Timeout: 600. There are checkboxes for 'Allow Root Login', 'Allow Anonymous Login', 'Allow Local User Login', and 'Allow Transfer Resumption', all of which are currently unchecked. There is a text field for 'Path' with a 'Browse' button next to it, and a 'Display Login' text field. Information icons (i) are present next to the Port, Clients, Connections, Login Attempts, and Timeout fields.

Fig. 11.5: Configuring FTP

Table 11.4 summarizes the available options when configuring the FTP server.

Table 11.4: FTP Configuration Options

Setting	Value	Advanced Mode	Description
Port	integer		port the FTP service listens on
Clients	integer		maximum number of simultaneous clients
Connections	integer		maximum number of connections per IP address where 0 means unlimited
Login Attempts	integer		maximum number of attempts before client is disconnected; increase this if users are prone to typos
Timeout	integer		maximum client idle time in seconds before client is disconnected
Allow Root Login	checkbox		discouraged as increases security risk
Allow Anonymous Login	checkbox		enables anonymous FTP logins with access to the directory specified in Path
Path	browse button		root directory for anonymous FTP connections
Allow Local User Login	checkbox		required if Anonymous Login is disabled
Display Login	string		message displayed to local login users after authentication; not displayed to anonymous login users
File Permission	checkboxes	✓	sets default permissions for newly created files
Directory Permission	checkboxes	✓	sets default permissions for newly created directories

Continued on next page

Table 11.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Enable FXP (https://en.wikipedia.org/wiki/File_eXchange_Protocol)	checkbox	✓	enables File eXchange Protocol which is discouraged as it makes the server vulnerable to FTP bounce attacks
Allow Transfer Re-umption	checkbox		allows FTP clients to resume interrupted transfers
Always Chroot	checkbox		a local user is only allowed access to their home directory unless the user is a member of group wheel
Require IDENT Authentication	checkbox	✓	will result in timeouts if <code>identd</code> is not running on the client
Perform Reverse DNS Lookups	checkbox		perform reverse DNS lookups on client IPs; can cause long delays if reverse DNS is not configured
Masquerade address	string		public IP address or hostname; set if FTP clients cannot connect through a NAT device
Minimum passive port	integer	✓	used by clients in PASV mode, default of 0 means any port above 1023
Maximum passive port	integer	✓	used by clients in PASV mode, default of 0 means any port above 1023
Local user upload bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Local user download bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Anonymous user upload bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Anonymous user download bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Enable TLS	checkbox	✓	enables encrypted connections and requires a certificate to be created or imported using Certificates (page 78)
TLS policy	drop-down menu	✓	the selected policy defines whether the control channel, data channel, both channels, or neither channel of an FTP session must occur over SSL/TLS; the policies are described here (http://www.proftpd.org/docs/directives/linked/config_ref_TLSRequired.h)
TLS allow client renegotiations	checkbox	✓	checking this box is not recommended as it breaks several security measures; for this and the rest of the TLS fields, refer to mod_tls (http://www.proftpd.org/docs/contrib/mod_tls.html) for more details
TLS allow dot login	checkbox	✓	if checked, the user's home directory is checked for a <code>.tlslogin</code> file which contains one or more PEM-encoded certificates; if not found, the user is prompted for password authentication
TLS allow per user	checkbox	✓	if checked, the user's password may be sent unencrypted
TLS common name required	checkbox	✓	if checked, the common name in the certificate must match the FQDN of the host
TLS enable diagnostics	checkbox	✓	if checked when troubleshooting a connection, logs more verbosely
TLS export certificate data	checkbox	✓	if checked, exports the certificate environment variables
TLS no certificate request	checkbox	✓	try checking this box if the client cannot connect and it is suspected that the client software is not properly handling the server's certificate request
TLS no empty fragments	checkbox	✓	checking this box is not recommended as it bypasses a security mechanism
TLS no session reuse required	checkbox	✓	checking this box reduces the security of the connection, so only use it if the client does not understand reused SSL sessions

Continued on next page

Table 11.4 – continued from previous page

Setting	Value	Advanced Mode	Description
TLS export standard vars	checkbox	✓	if checked, sets several environment variables
TLS DNS name required	checkbox	✓	if checked, the client's DNS name must resolve to its IP address and the cert must contain the same DNS name
TLS IP address required	checkbox	✓	if checked, the client's certificate must contain the IP address that matches the IP address of the client
Certificate	drop-down menu		the SSL certificate to be used for TLS FTP connections; to create a certificate, use System → Certificates
Auxiliary parameters	string	✓	used to add proftpd(8) (http://linux.die.net/man/8/proftpd) parameters not covered elsewhere in this screen

This example demonstrates the auxiliary parameters that prevent all users from performing the FTP DELETE command:

```
<Limit DELE>
DenyAll
</Limit>
```

11.5.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the FreeNAS® system is not accessible from the Internet and everyone in your internal network needs easy access to the stored data. Anonymous FTP does not require you to create a user account for every user. In addition, passwords are not required so it is not necessary to manage changed passwords on the FreeNAS® system.

To configure anonymous FTP:

1. Give the built-in ftp user account permissions to the volume/dataset to be shared in [Storage](#) → [Volumes](#) as follows:
 - Owner(user): select the built-in ftp user from the drop-down menu
 - Owner(group): select the built-in ftp group from the drop-down menu
 - Mode: review that the permissions are appropriate for the share

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means that you always use Unix ACLs, even if Windows clients will be accessing FreeNAS® via FTP.

2. Configure anonymous FTP in [Services](#) → [FTP](#) by setting the following attributes:
 - check the box [Allow Anonymous Login](#)
 - Path: browse to the volume/dataset/directory to be shared
3. Start the FTP service in [Services](#) → [Control Services](#). Click the [Start Now](#) button next to FTP. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to [Stop Now](#).
4. Test the connection from a client using a utility such as [Filezilla](https://filezilla-project.org/) (<https://filezilla-project.org/>).

In the example shown in [Figure 11.6](#), the user has enter the following information into the Filezilla client:

- IP address of the FreeNAS® server: 192.168.1.113
- Username: anonymous
- Password: the email address of the user

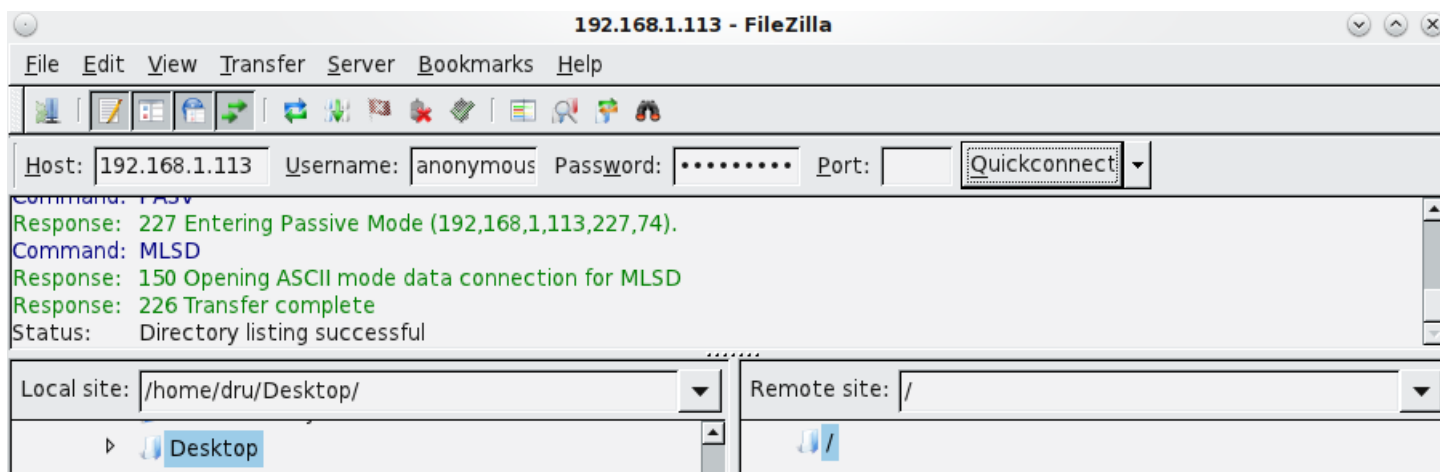


Fig. 11.6: Connecting Using Filezilla

The messages within the client indicate that the FTP connection is successful. The user can now navigate the contents of the root folder on the remote site—this is the volume/dataset that was specified in the FTP service configuration. The user can also transfer files between the local site (their system) and the remote site (the FreeNAS® system).

11.5.2 FTP in chroot

If you require your users to authenticate before accessing the data on the FreeNAS® system, you will need to either create a user account for each user or import existing user accounts using [Active Directory](#) (page 158) or LDAP. If you then create a ZFS dataset for each user, you can chroot each user so that they are limited to the contents of their own home directory. Datasets provide the added benefit of configuring a quota so that the size of the user's home directory is limited to the size of the quota.

To configure this scenario:

1. Create a ZFS dataset for each user in **Storage** → **Volumes**. Click an existing ZFS volume → **Create ZFS Dataset** and set an appropriate quota for each dataset. Repeat this process to create a dataset for every user that needs access to the FTP service.
2. If you are not using AD or LDAP, create a user account for each user in **Account** → **Users** → **Add User**. For each user, browse to the dataset created for that user in the **Home Directory** field. Repeat this process to create a user account for every user that needs access to the FTP service, making sure to assign each user their own dataset.
3. Set the permissions for each dataset in **Storage** → **Volumes**. Click the **Change Permissions** button for a dataset to assign a user account as Owner of that dataset and to set the desired permissions for that user. Repeat for each dataset.

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means that you always use Unix ACLs, even if Windows clients will be accessing FreeNAS® via FTP.

4. Configure FTP in **Services** → **FTP** with these attributes:
 - Path: browse to the parent volume containing the datasets
 - make sure the boxes for **Allow Anonymous Login** and **Allow Root Login** are unchecked
 - check the box **Allow Local User Login**
 - check the box **Always Chroot**

5. Start the FTP service in *Services* → *Control Services*. Click the Start Now button next to FTP. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to Stop Now.
6. Test the connection from a client using a utility such as Filezilla.

To test this configuration in Filezilla, use the IP address of the FreeNAS® system, the Username of a user that has been associated with a dataset, and the Password for that user. The messages should indicate that the authorization and the FTP connection are successful. The user can now navigate the contents of the root folder on the remote site—this time it is not the entire volume but the dataset that was created for that user. The user should be able to transfer files between the local site (their system) and the remote site (their dataset on the FreeNAS® system).

11.5.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

1. Import or create a certificate authority using the instructions in [CAs](#) (page 75). Then, import or create the certificate to use for encrypted connections using the instructions in [Certificates](#) (page 78).
2. In *Services* → *FTP*, check the box Enable TLS and select the certificate in the Certificate drop-down menu.
3. Specify secure FTP when accessing the FreeNAS® system. For example, in Filezilla enter `ftps://IP_address` (for an implicit connection) or `ftpes://IP_address` (for an explicit connection) as the Host when connecting. The first time a user connects, they will be presented with the certificate of the FreeNAS® system. Click OK to accept the certificate and negotiate an encrypted connection.
4. To force encrypted connections, select on for the TLS Policy.

11.5.4 Troubleshooting FTP

The FTP service will not start if it cannot resolve the system's hostname to an IP address using DNS. To see if the FTP service is running, open [Shell](#) (page 290) and issue the command:

```
sockstat -4p 21
```

If there is nothing listening on port 21, the FTP service is not running. To see the error message that occurs when FreeNAS® tries to start the FTP service, go to *System* → *Advanced*, check the box Show console messages in the footer and click Save. Next, go to *Services* → *Control Services* and switch the FTP service off, then back on. Watch the console messages at the bottom of the browser for errors.

If the error refers to DNS, either create an entry in the local DNS server with the FreeNAS® system's hostname and IP address or add an entry for the IP address of the FreeNAS® system in the Host name database field of *Network* → *Global Configuration*.

11.6 iSCSI

Refer to [Block \(iSCSI\)](#) (page 195) for instructions on configuring iSCSI. To start the iSCSI service, click its entry in *Services*.

Note: A warning message is shown if you stop the iSCSI service when initiators are connected. Type `ctladm islist` to determine the names of the connected initiators.

11.7 LLDP

The Link Layer Discovery Protocol (LLDP) is used by network devices to advertise their identity, capabilities, and neighbors on an Ethernet network. FreeNAS® uses the [ladvd](https://github.com/sspan/ladvd) (<https://github.com/sspan/ladvd>) LLDP implementation. If your network contains managed switches, configuring and starting the LLDP service will tell the FreeNAS® system to advertise itself on the network.

Figure 11.7 shows the LLDP configuration screen and Table 11.5 summarizes the configuration options for the LLDP service.




Fig. 11.7: Configuring LLDP

Table 11.5: LLDP Configuration Options

Setting	Value	Description
Interface De- scription	checkbox	when checked, receive mode is enabled and received peer information is saved in interface descriptions
Country Code	string	required for LLDP location support; enter a two-letter ISO 3166 country code
Location	string	optional; specify the physical location of the host

11.8 Netdata

Netdata is a real-time performance and monitoring system. It displays data as web dashboards.

Start the Netdata service from the [Services](#) (page 209) screen. Click the wrench icon to display the Netdata settings dialog shown in Figure 11.8.



Fig. 11.8: Netdata Settings Dialog

Click the Take me to the Netdata UI button to view the web dashboard as shown in Figure 11.9.



Fig. 11.9: Netdata Web Dashboard

More information on configuring and using Netdata is available at the [Netdata website](https://my-netdata.io/) (<https://my-netdata.io/>).

11.9 NFS

The settings that are configured when creating NFS Shares in **Sharing** → **Unix (NFS) Shares** → **Add Unix (NFS) Share** are specific to each configured NFS Share. In contrast, global settings which apply to all NFS shares are configured in **Services** → **NFS**.

Figure 11.10 shows the configuration screen and Table 11.6 summarizes the configuration options for the NFS service.

Settings

Number of servers:

4

Serve UDP NFS clients:

Bind IP Addresses:

10.0.0.142

Allow non-root mount:

Enable NFSv4:

NFSv3 ownership model for NFSv4:

Require Kerberos for NFSv4:

mountd(8) bind port:

rpc.statd(8) bind port:

rpc.lockd(8) bind port:

Support >16 groups:

Log mountd(8) requests:

Log rpc.statd(8) and rpc.lockd(8):

OK

Cancel

Fig. 11.10: Configuring NFS

Table 11.6: NFS Configuration Options

Setting	Value	Description
Number of servers	integer	the number of servers can be increased if NFS client responses are slow; to limit CPU context switching, keep this number less than or equal to the number of CPUs reported by <code>sysctl -n kern.smp.cpus</code> .
Serve UDP NFS clients	checkbox	check if NFS clients need to use UDP
Bind IP Ad-dresses	checkboxes	IP addresses to listen on for NFS requests; when unchecked, NFS listens on all available addresses

Continued on next page

Table 11.6 – continued from previous page

Setting	Value	Description
Allow non-root mount	checkbox	check this box only if the NFS client requires it
Enable NFSv4	checkbox	NFSv3 is the default, check this box to switch to NFSv4
NFSv3 ownership model for NFSv4	checkbox	grayed out unless Enable NFSv4 is checked and, in turn, grays out Support>16 groups which is incompatible; check this box if NFSv4 ACL support is needed without requiring the client and the server to sync users and groups
Require Kerberos for NFSv4	checkbox	when checked, NFS shares will fail if the Kerberos ticket is unavailable
mountd(8) bind port	integer	optional; specify port that mountd(8) (http://www.freebsd.org/cgi/man.cgi?query=mountd) binds to
rpc.statd(8) bind port	integer	optional; specify port that rpc.statd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.statd) binds to
rpc.lockd(8) bind port	integer	optional; specify port that rpc.lockd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) binds to
Support>16 groups	checkbox	check this box if any users are members of more than 16 groups (useful in AD environments); note that this assumes that group membership has been configured correctly on the NFS server
Log mountd(8) requests	checkbox	enable logging of mountd(8) (http://www.freebsd.org/cgi/man.cgi?query=mountd) requests by syslog
Log rpc.statd(8) and rpc.lockd(8)	checkbox	enable logging of rpc.statd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.statd) and rpc.lockd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) requests by syslog

Note: NFSv4 sets all ownership to nobody:nobody if user and group do not match on client and server.

11.10 Rsync

`Services` → `Rsync` is used to configure an rsync server when using rsync module mode. Refer to [Rsync Module Mode](#) (page 96) for a configuration example.

This section describes the configurable options for the `rsyncd` service and rsync modules.

11.10.1 Configure Rsyncd

[Figure 11.11](#) shows the rsyncd configuration screen which is accessed from `Services` → `Rsync` → `Configure Rsyncd`.

Configure Rsyncd

TCP Port

Auxiliary parameters

OK

Cancel

Fig. 11.11: Rsyncd Configuration

Table 11.7 summarizes the options that can be configured for the rsync daemon:

Table 11.7: Rsyncd Configuration Options

Setting	Value	Description
TCP Port	integer	port for <code>rsyncd</code> to listen on, default is 873
Auxiliary parameters	string	additional parameters from <code>rsyncd.conf(5)</code> (https://www.samba.org/ftp/rsync/rsyncd.conf.html)

11.10.2 Rsync Modules

Figure 11.12 shows the configuration screen that appears after clicking `Services → Rsync → Rsync Modules → Add Rsync Module`.

Table 11.8 summarizes the options that can be configured when creating a rsync module.

Add Rsync Module

Module name:

Comment:

Path: **Browse** ⓘ

Access Mode: **Read and Write** ⓘ

Maximum connections: ⓘ

User: **nobody** ⓘ

Group: **nobody** ⓘ

Hosts allow: ⓘ

Hosts deny:

Fig. 11.12: Adding an Rsync Module

Table 11.8: Rsync Module Configuration Options

Setting	Value	Description
Module name	string	mandatory; needs to match the setting on the rsync client
Comment	string	optional description
Path	browse button	volume/dataset to hold received data
Access Mode	drop-down menu	choices are Read and Write, Read-only, or Write-only
Maximum connections	integer	0 is unlimited
User	drop-down menu	select user that file transfers to and from that module should take place as
Group	drop-down menu	select group that file transfers to and from that module should take place as
Hosts allow	string	see rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html) for allowed formats
Hosts deny	string	see rsyncd.conf(5) for allowed formats
Auxiliary parameters	string	additional parameters from rsyncd.conf(5)

11.11 S3

S3 is a distributed or clustered filesystem protocol compatible with Amazon S3 cloud storage. The FreeNAS® S3 service uses [Minio](https://minio.io/) (<https://minio.io/>) to provide S3 storage hosted on the FreeNAS® system itself. Minio also provides features beyond the limits of the basic Amazon S3 specifications.

Figure 11.13 shows the S3 service configuration screen and Table 11.9 summarizes the configuration options. After configuring the S3 service, start it in *Services* → *Control Services*.

The screenshot shows the 'Settings' window for the S3 service. It includes the following fields and controls:

- IP Address:** A dropdown menu showing '0.0.0.0' with an information icon.
- Port:** A text input field containing '9000' with an information icon.
- Access key of 5 to 20 characters in length:** An empty text input field with an information icon.
- Secret key of 8 to 40 characters in length:** An empty text input field with an information icon.
- Confirm S3 Key:** An empty text input field.
- Disks:** An empty text input field with a 'Browse' button and an information icon.
- Certificate:** A dropdown menu with a dashed line and an information icon.
- Enable Browser:** A checkbox that is currently unchecked with an information icon.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom left.

Fig. 11.13: Configuring S3

Table 11.9: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	the IP address on which to run the S3 service; 0.0.0.0 sets the server to listen on all addresses
Port	string	TCP port on which to provide the S3 service (default 9000)
Access Key	string	the S3 user name
Secret Key	string	the password to be used by connecting S3 systems; must be at least 8 but no more than 40 characters long
Confirm S3 Key	string	re-enter the S3 password to confirm
Disks	string	S3 filesystem directory
Certificate	drop-down menu	the SSL certificate to be used for secure S3 connections; to create a certificate, use <i>System</i> → <i>Certificates</i>
Enable Browser	checkbox	Enable the web user interface for the S3 service

11.12 S.M.A.R.T.

S.M.A.R.T., or *Self-Monitoring, Analysis, and Reporting Technology* (<http://en.wikipedia.org/wiki/S.M.A.R.T.>), is an industry standard for disk monitoring and testing. Drives can be monitored for status and problems, and several types of self-tests can be run to check the drive health.

Tests run internally on the drive. Most tests can run at the same time as normal disk usage. However, a running test can greatly reduce drive performance, so they should be scheduled at times when the system is not busy or in normal use. It is very important to avoid scheduling disk-intensive tests at the same time. For example, do not schedule S.M.A.R.T. tests to run at the same time, or preferably, even on the same days as [Scrubs](#) (page 152).

Of particular interest in a NAS environment are the Short and Long S.M.A.R.T. tests. Details vary between drive manufacturers, but a Short test generally does some basic tests of a drive that takes a few minutes. The Long test scans the entire disk surface, and can take several hours on larger drives.

FreeNAS® uses the [smartd\(8\)](http://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in) (<http://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in>) service to monitor S.M.A.R.T. information. A complete configuration consists of:

1. Scheduling when S.M.A.R.T. tests are run in [Tasks](#) → S.M.A.R.T. Tests → Add S.M.A.R.T. Test.
2. Enabling or disabling S.M.A.R.T. for each disk member of a volume in [Volumes](#) → View Disks. This setting is enabled by default for disks that support S.M.A.R.T.
3. Checking the configuration of the S.M.A.R.T. service as described in this section.
4. Starting the S.M.A.R.T. service with [Services](#) → Control Services.

[Figure 11.14](#) shows the configuration screen that appears after clicking [Services](#) → S.M.A.R.T.

Fig. 11.14: S.M.A.R.T Configuration Options

Note: `smartd` wakes up at the configured Check Interval. It checks the times configured in [Tasks](#) → S.M.A.R.T. Tests to see whether tests should be run. Since the smallest time increment for a test is an hour (60 minutes), it does not make sense to set a Check Interval value higher than 60 minutes. For example, if the Check Interval is set to 120 minutes and the smart test to every hour, the test will only be run every two hours because `smartd` only wakes up every two hours.

[Table 11.10](#) summarizes the options in the S.M.A.R.T configuration screen.

Table 11.10: S.M.A.R.T Configuration Options

Setting	Value	Description
Check interval	integer	in minutes, how often <code>smartd</code> wakes up to check if any tests have been configured to run
Power mode	drop-down menu	tests are not performed if the system enters the specified power mode; choices are: Never, Sleep, Standby, or Idle

Continued on next page

Table 11.10 – continued from previous page

Setting	Value	Description
Difference	integer in degrees Celsius	default of 0 disables this check, otherwise reports if the temperature of a drive has changed by N degrees Celsius since last report
Informational	integer in degrees Celsius	default of 0 disables this check, otherwise will message with a log level of LOG_INFO if the temperature is higher than specified degrees in Celsius
Critical	integer in degrees Celsius	default of 0 disables this check, otherwise will message with a log level of LOG_CRIT and send an email if the temperature is higher than specified degrees in Celsius
Email to report	string	email address of person or alias to receive S.M.A.R.T. alerts

11.13 SMB

The settings that are configured when creating SMB Shares in *Sharing* → *Windows (SMB) Shares* → *Add Windows (SMB) Share* are specific to each configured SMB Share. In contrast, global settings which apply to all SMB shares are configured in *Services* → *SMB*.

Note: After starting the SMB service, it can take several minutes for the *master browser election* (<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/NetworkBrowsing.html#id2581357>) to occur and for the FreeNAS® system to become available in Windows Explorer.

Figure 11.15 shows the global SMB configuration options which are described in Table 11.11. This configuration screen is really a front-end to *smb4.conf* (<https://www.freebsd.org/cgi/man.cgi?query=smb4.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>).

Settings	
NetBIOS name:	<input type="text" value="freenas"/>
NetBIOS alias:	<input type="text"/>
Workgroup:	<input type="text" value="WORKGROUP"/> ⓘ
Description:	<input type="text" value="FreeNAS Server"/> ⓘ
DOS charset:	<input type="text" value="CP437"/> ▼
UNIX charset:	<input type="text" value="UTF-8"/> ▼
Log level:	<input type="text" value="Minimum"/> ▼
Use syslog only:	<input type="checkbox"/>
Local Master:	<input checked="" type="checkbox"/>
Domain logons:	<input type="checkbox"/>
Time Server for Domain:	<input checked="" type="checkbox"/>
Guest account:	<input type="text" value="nobody"/> ▼ ⓘ
File mask:	<input type="text"/> ⓘ
Directory mask:	<input type="text"/> ⓘ
Allow Empty Password:	<input type="checkbox"/>
Auxiliary parameters:	<input type="text"/> ⓘ
Unix Extensions:	<input checked="" type="checkbox"/> ⓘ
Zeroconf share discovery:	<input checked="" type="checkbox"/> ⓘ
Hostnames lookups:	<input checked="" type="checkbox"/> ⓘ
Allow execute always:	<input checked="" type="checkbox"/> ⓘ
Obey pam restrictions:	<input checked="" type="checkbox"/> ⓘ
NTLMv1 auth:	<input type="checkbox"/> ⓘ
Bind IP Addresses:	<input type="checkbox"/> 10.0.0.102 ⓘ

Table 11.11: Global SMB Configuration Options

Setting	Value	Description
NetBIOS Name	string	automatically populated with the system's original hostname; limited to 15 characters; it must be different from the Workgroup name
NetBIOS Alias	string	limited to 15 characters
Workgroup	string	must match Windows workgroup name; this setting is ignored if the Active Directory (page 158) or LDAP (page 163) service is running
Description	string	optional
DOS charset	drop-down menu	the character set Samba uses when communicating with DOS and Windows 9x/ME clients; default is CP437
UNIX charset	drop-down menu	default is UTF-8 which supports all characters in all languages
Log level	drop-down menu	choices are Minimum, Normal, or Debug
Use syslog only	checkbox	when checked, authentication failures are logged to <code>/var/log/messages</code> instead of the default of <code>/var/log/samba4/log.smbd</code>
Local Master	checkbox	determines whether or not the system participates in a browser election; should be disabled when network contains an AD or LDAP server and is not necessary if Vista or Windows 7 machines are present
Domain logons	checkbox	only check if need to provide the netlogin service for older Windows clients
Time Server for Domain	checkbox	determines whether or not the system advertises itself as a time server to Windows clients; should be disabled when network contains an AD or LDAP server
Guest Account	drop-down menu	account to be used for guest access; default is nobody; account must have permission to access the shared volume/dataset; if Guest Account user is deleted, resets to nobody
File mask	integer	overrides default file creation mask of 0666 which creates files with read and write access for everybody
Directory mask	integer	overrides default directory creation mask of 0777 which grants directory read, write and execute access for everybody
Allow Empty Password	checkbox	if checked, users can just press <code>Enter</code> when prompted for a password; requires that the username/password be the same as the Windows user account
Auxiliary parameters	string	<code>smb.conf</code> options not covered elsewhere in this screen; see the Samba Guide (http://www.oreilly.com/openbook/samba/book/appb_02.html) for additional settings
Unix Extensions	checkbox	allows non-Windows SMB clients to access symbolic links and hard links, has no effect on Windows clients
Zeroconf share discovery	checkbox	enable if Mac clients will be connecting to the SMB share
Hostname lookups	checkbox	allows using hostnames rather than IP addresses in the Hosts Allow or Hosts Deny fields of a SMB share; uncheck if IP addresses are used to avoid the delay of a host lookup
Allow execute always	checkbox	if checked, Samba will allow the user to execute a file, even if that user's permissions are not set to execute
Obey pam restrictions	checkbox	uncheck this box to allow cross-domain authentication, to allow users and groups to be managed on another forest, or to allow permissions to be delegated from Active Directory (page 158) users and groups to domain admins on another forest
NTLMv1 auth	checkbox	when checked, allow NTLMv1 authentication, required by Windows XP clients and sometimes by clients in later versions of Windows
Bind IP Addresses	checkboxes	check the IP addresses on which SMB should listen

Continued on next page

Table 11.11 – continued from previous page

Setting	Value	Description
Idmap Range Low	integer	the beginning UID/GID for which this system is authoritative; any UID/GID lower than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs
Idmap Range High	integer	the ending UID/GID for which this system is authoritative; any UID/GID higher than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs

Changes to SMB settings take effect immediately. Changes to share settings only take effect after the client and server negotiate a new session.

Note: Do not set the directory name cache size as an Auxiliary parameter. Due to differences in how Linux and BSD handle file descriptors, directory name caching is disabled on BSD systems to improve performance.

Note: [SMB](#) (page 230) cannot be disabled while [Active Directory](#) (page 158) is enabled.

11.13.1 Troubleshooting SMB

Do not connect to SMB shares as `root`, and do not add the root user in the SMB user database. There are security implications in attempting to do so, and Samba 4 and later take measures to prevent such actions. This can produce `auth_check_ntlm_password` and `FAILED with error NT_STATUS_WRONG_PASSWORD` errors.

Samba is single threaded, so CPU speed makes a big difference in SMB performance. A typical 2.5Ghz Intel quad core or greater should be capable of handling speeds in excess of Gb LAN while low power CPUs such as Intel Atoms and AMD C-30sE-350E-450 will not be able to achieve more than about 30-40MB/sec typically. Remember that other loads such as ZFS will also require CPU resources and may cause Samba performance to be less than optimal.

Samba's write cache parameter has been reported to improve write performance in some configurations and can be added to the Auxiliary parameters field. Use an integer value which is a multiple of `_SC_PAGESIZE` (typically 4096) to avoid memory fragmentation. This will increase Samba's memory requirements and should not be used on systems with limited RAM.

Windows automatically caches file sharing information. If changes are made to an SMB share or to the permissions of a volume/dataset being shared by SMB and the share becomes inaccessible, try logging out and back into the Windows system. Alternately, users can type `net use /delete` from the command line to clear their SMB sessions.

Windows also automatically caches login information. To require users to log in every time access is required, reduce the cache settings on the client computers.

Where possible, avoid using a mix of case in filenames as this can cause confusion for Windows users. [Representing and resolving filenames with Samba](http://www.oreilly.com/openbook/samba/book/ch05_04.html) (http://www.oreilly.com/openbook/samba/book/ch05_04.html) explains in more detail.

If a particular user cannot connect to a SMB share, make sure that their password does not contain the `?` character. If it does, have the user change the password and try again.

If permissions work for Windows users but not for OS X users, try disabling Unix Extensions and restarting the SMB service.

If the SMB service will not start, run this command from [Shell](#) (page 290) to see if there is an error in the configuration:

```
testparm /usr/local/etc/smb4.conf
```

If clients have problems connecting to the SMB share, go to `Services` → `SMB` and verify that Server maximum protocol is set to SMB2.

It is recommended to use a dataset for SMB sharing. When creating the dataset, make sure that the Share type is set to Windows.

Do not use `chmod` to attempt to fix the permissions on a SMB share as it destroys the Windows ACLs. The correct way to manage permissions on a SMB share is to manage the share security from a Windows system as either the owner of the share or a member of the group that owns the share. To do so, right-click on the share, click Properties and navigate to the Security tab. If you already destroyed the ACLs using `chmod`, `winacl` can be used to fix them. Type `winacl` from [Shell](#) (page 290) for usage instructions.

The [Common Errors](#) (<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/domain-member.html#id2573692>) section of the Samba documentation contains additional troubleshooting tips.

The Samba [Performance Tuning](#) (https://wiki.samba.org/index.php/Performance_Tuning) page describes options to improve performance.

Directory listing speed in folders with a large number of files is sometimes a problem. A few specific changes can help improve the performance. However, changing these settings can affect other usage. In general, the defaults are adequate. Do not change these settings unless there is a specific need.

- Hostname Lookups and Log Level can also have a performance penalty. When not needed, they can be disabled or reduced in the [global SMB service options](#) (page 232).
- Make Samba datasets case insensitive by setting Case Sensitivity to Insensitive when creating them. This ZFS property is only available when creating a dataset. It cannot be changed on an existing dataset. To convert such datasets, back up the data, create a new case-insensitive dataset, create an SMB share on it, set the share level auxiliary parameter `case sensitive = true`, then copy the data from the old one onto it. After the data has been checked and verified on the new share, the old one can be deleted.
- If present, remove options for extended attributes and DOS attributes in the share's [Auxiliary Parameters](#) (page 186).
- Disable as many VFS Objects as possible in the [share settings](#) (page 186). Many have performance overhead.

11.14 SNMP

SNMP (Simple Network Management Protocol) is used to monitor network-attached devices for conditions that warrant administrative attention. FreeNAS® uses [Net-SNMP](#) (<http://net-snmp.sourceforge.net/>) to provide SNMP. When you start the SNMP service, the following port will be enabled on the FreeNAS® system:

- UDP 161 (listens here for SNMP requests)

Available MIBS are located in `/usr/local/share/snmp/mibs`.

[Figure 11.16](#) shows the SNMP configuration screen. [Table 11.12](#) summarizes the configuration options.

Settings

Location: ⓘ

Contact: ⓘ

SNMP v3 Support: ☐

Community: ⓘ

Username:

Authentication Type: SHA ▾

Password:

Confirm Password:

Privacy Protocol: ----- ▾

Privacy Passphrase:

Confirm Privacy Passphrase:

Log Level: Error ▾

Auxiliary parameters: ⓘ

OK **Cancel**

Fig. 11.16: Configuring SNMP

Table 11.12: SNMP Configuration Options

Setting	Value	Description
Location	string	optional description of system's location
Contact	string	optional email address of administrator
SNMP v3 Support	checkbox	check this box to enable support for SNMP version 3
Community	string	default is public and should be changed for security reasons; can only contain alphanumeric characters, underscores, dashes, periods, and spaces; this value can be empty for SNMPv3 networks

Continued on next page

Table 11.12 – continued from previous page

Setting	Value	Description
Username	string	only applies if SNMP v3 Support is checked; specify the username to register with this service; refer to snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) for more information regarding the configuration of this setting as well as the Authentication Type, Password, Privacy Protocol, and “Privacy Passphrase” fields
Authentication Type	drop-down menu	only applies if SNMP v3 Support is checked; choices are MD5 or SHA
Password	string	only applies if SNMP v3 Support is checked; specify and confirm a password of at least eight characters
Privacy Protocol	drop-down menu	only applies if SNMP v3 Support is checked; choices are AES or DES
Privacy Passphrase	string	if not specified, Password is used
Log Level	drop-down menu	choices range from the least log entries (Emergency) to the most (Debug)
Auxiliary Parameters	string	additional snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) options not covered in this screen, one per line

11.15 SSH

Secure Shell (SSH) allows for files to be transferred securely over an encrypted network. If you configure your FreeNAS® system as an SSH server, the users in your network will need to use [SSH client software](https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) (https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) to transfer files with SSH.

This section shows the FreeNAS® SSH configuration options, demonstrates an example configuration that restricts users to their home directory, and provides some troubleshooting tips.

Figure 11.17 shows the **Services** → **SSH** configuration screen. After configuring SSH, remember to start it in **Services** → **Control Services**.

Fig. 11.17: SSH Configuration

Table 11.13 summarizes the configuration options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button, or configure the system to always display these settings by checking the box **Show advanced fields by default** in **System** → **Advanced**.

Table 11.13: SSH Configuration Options

Setting	Value	Advanced Mode	Description
Bind Interfaces	selection	✓	by default, SSH listens on all interfaces unless specific interfaces are highlighted in the Available field and added to the Selected field
TCP Port	integer		port to open for SSH connection requests; 22 by default
Login as Root with password	checkbox		for security reasons, root logins are discouraged and disabled by default if enabled, password must be set for root user in View Users
Allow Password Authentication	checkbox		if unchecked, key-based authentication for all users is required; requires additional setup (http://the.earth.li/%7Esgtatham/putty/0.55/html/doc/Chapter8.html) on both the SSH client and server
Allow Kerberos Authentication	checkbox		before checking this box, ensure that Kerberos Realms (page 167) and Kerberos Keytabs (page 167) have been configured and that the FreeNAS® system can communicate with the KDC
Allow TCP Port Forwarding	checkbox		allows users to bypass firewall restrictions using SSH's port forwarding feature (http://www.symantec.com/connect/articles/ssh-port-forwarding)
Compress Connections	checkbox		may reduce latency over slow networks
SFTP Log Level	drop-down menu	✓	select the syslog(3) (http://www.freebsd.org/cgi/man.cgi?query=syslog) level of the SFTP server
SFTP Log Facility	drop-down menu	✓	select the syslog(3) (http://www.freebsd.org/cgi/man.cgi?query=syslog) facility of the SFTP server
Extra Options	string	✓	additional sshd_config(5) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options not covered in this screen, one per line; these options are case-sensitive and misspellings may prevent the SSH service from starting

A few [sshd_config\(5\)](#) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options that are useful to enter in the Extra Options field include:

- increase the ClientAliveInterval if SSH connections tend to drop
- ClientMaxStartup defaults to 10; increase this value if you need more concurrent SSH connections

11.15.1 SCP Only

When you configure SSH, authenticated users with a user account created using **Account → Users → Add User** can use the `ssh` command to login to the FreeNAS® system over the network. A user's home directory will be the volume/dataset specified in the Home Directory field of their FreeNAS® user account. While the SSH login will default to the user's home directory, users are able to navigate outside of their home directory, which can pose a security risk.

It is possible to allow users to use the `scp` and `sftp` commands to transfer files between their local computer and their home directory on the FreeNAS® system, while restricting them from logging into the system using `ssh`. To configure this scenario, go to **Account → Users → View Users**, select the user and click **Modify User**, and change the user's Shell to `scponly`. Repeat for each user that needs restricted SSH access.

Test the configuration from another system by running the `sftp`, `ssh`, and `scp` commands as the user. The `sftp` and `scp` commands should work but the `ssh` should fail.

Note: Some utilities such as WinSCP and Filezilla can bypass the `scponly` shell. This section assumes that users are accessing the system using the command line versions of `scp` and `sftp`.

11.15.2 Troubleshooting SSH

When adding any Extra Options, be aware that the keywords listed in `sshd_config(5)` (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) are case sensitive. This means that your configuration will fail to do what you intended if you do not match the upper and lowercase letters of the keyword.

If your clients are receiving “reverse DNS” or timeout errors, add an entry for the IP address of the FreeNAS® system in the Host name database field of **Network** → **Global Configuration**.

When configuring SSH, always test your configuration as an SSH user account to ensure that the user is limited to what you have configured and that they have permission to transfer files within the intended directories. If the user account is experiencing problems, the SSH error messages are usually pretty specific to what the problem is. Type the following command within **Shell** (page 290) to read these messages as they occur:

```
tail -f /var/log/messages
```

Additional messages regarding authentication errors may be found in `/var/log/auth.log`.

11.16 TFTP

Trivial File Transfer Protocol (TFTP) is a light-weight version of FTP usually used to transfer configuration or boot files between machines, such as routers, in a local environment. TFTP provides an extremely limited set of commands and provides no authentication.

If the FreeNAS® system will be used to store images and configuration files for the network’s devices, configure and start the TFTP service. Starting the TFTP service will open UDP port 69.

Figure 11.18 shows the TFTP configuration screen and Table 11.14 summarizes the available options:

Fig. 11.18: TFTP Configuration

Table 11.14: TFTP Configuration Options

Setting	Value	Description
Directory	browse button	browse to an existing directory to be used for storage; some devices require a specific directory name, refer to the device's documentation for details
Allow New Files	checkbox	enable if network devices need to send files to the system (for example, to back up their configuration)
Port	integer	UDP port to listen for TFTP requests, 69 by default
Username	drop-down menu	account used for tftp requests; must have permission to the Directory
Umask	integer	umask for newly created files, default is 022 (everyone can read, nobody can write); some devices require a less strict umask
Extra options	string	additional tftpd(8) (http://www.freebsd.org/cgi/man.cgi?query=tftpd) options not shown in this screen, one per line

11.17 UPS

FreeNAS® uses [NUT](http://www.networkupstools.org/) (<http://www.networkupstools.org/>) (Network UPS Tools) to provide UPS support. If the FreeNAS® system is connected to a UPS device, configure the UPS service then start it in *Services* → *Control Services*.

[Figure 11.19](#) shows the UPS configuration screen:

Settings

UPS Mode:	Master	
Identifier:	ups	i
Driver:	-----	i
Port:		
Auxiliary parameters (ups.conf):	i	
Auxiliary parameters (upsd.conf):	i	
Description:		
Shutdown mode:	UPS goes on battery	
Shutdown timer:	30	i
Shutdown Command:	/sbin/shutdown -p now	i
No Communication Warning Time:		i
Monitor User:	upsmon	
Monitor Password:	fixmepass	
Extra users (upsd.users):		
Remote Monitor:	<input type="checkbox"/>	
Send Email Status Updates:	<input type="checkbox"/>	
To email:		i
Email Subject:	UPS report generated by %h	i
Power Off UPS:	<input type="checkbox"/>	i

Fig. 11.19: UPS Configuration Screen

Table 11.15 summarizes the options in the UPS Configuration screen.

Table 11.15: UPS Configuration Options

Setting	Value	Description
UPS Mode	drop-down menu	select from Master or Slave
Identifier	string	can contain alphanumeric, period, comma, hyphen, and underscore characters
Driver	drop-down menu	supported UPS devices are listed at http://www.networkupstools.org/stable-hcl.html
Port	drop-down menu	select the serial or USB port the UPS is plugged into (see NOTE below)
Auxiliary Parameters (ups.conf)	string	additional options from ups.conf(5) (http://www.networkupstools.org/docs/man/ups.conf.html)
Auxiliary Parameters (upsd.conf)	string	additional options from upsd.conf(5) (http://www.networkupstools.org/docs/man/upsd.conf.html)
Description	string	optional
Shutdown mode	drop-down menu	choices are UPS goes on battery and UPS reaches low battery
Shutdown timer	integer	in seconds; will initiate shutdown after this many seconds after UPS enters UPS goes on battery, unless power is restored
Shutdown Command	string	the command to run to shut down the computer when battery power is low or shutdown timer runs out
No Communication Warning Time	string	the frequency, in seconds, of email notifications during the loss of UPS communications; the default is 300
Monitor User	string	default is upsmon
Monitor Password	string	default is known value fixmepass and should be changed; cannot contain a space or #
Extra users	string	defines the accounts that have administrative access; see upsd.users(5) (http://www.networkupstools.org/docs/man/upsd.users.html) for examples
Remote monitor	checkbox	if enabled, be aware that the default is to listen on all interfaces and to use the known values user upsmon and password fixmepass
Send Email Status Updates	checkbox	if checked, activates the To email field
To email	email address	if Send Email box checked, email address to receive status updates; separate multiple email addresses with a semicolon
Email Subject	string	subject line to be used in the email
Power Off UPS	checkbox	if checked, the UPS will also power off after shutting down the FreeNAS system

Note: For USB devices, the easiest way to determine the correct device name is to check the box Show console messages in System → Advanced. Plug in the USB device and look for a /dev/ugen or /dev/uhid device name in the console messages.

[upsc\(8\)](http://www.networkupstools.org/docs/man/upsc.html) (<http://www.networkupstools.org/docs/man/upsc.html>) can be used to get status variables from the UPS daemon such as the current charge and input voltage. It can be run from Shell using the following syntax. The man page gives some other usage examples.

```
upsc upsc@localhost
```

[upscmd\(8\)](http://www.networkupstools.org/docs/man/upscmd.html) (<http://www.networkupstools.org/docs/man/upscmd.html>) can be used to send commands directly to the UPS, assuming that the hardware supports the command being sent. Only users with administrative rights can use this command. These users are created in the Extra users field.

11.17.1 Multiple Computers with One UPS

A UPS with adequate capacity can be used to power multiple computers. One computer is connected to the UPS data port with a serial or USB cable. This master makes UPS status available on the network for other computers. These slave computers are powered by the UPS, but receive UPS status data from the master computer. See the [NUT User Manual](http://networkupstools.org/docs/user-manual.chunked/index.html) (<http://networkupstools.org/docs/user-manual.chunked/index.html>) and [NUT User Manual Pages](http://networkupstools.org/docs/man/index.html#User_man) (http://networkupstools.org/docs/man/index.html#User_man).

11.18 WebDAV

The WebDAV service can be configured to provide a file browser over a web connection. Before starting this service, you must create at least one WebDAV share using *Sharing* → *WebDAV Shares* → *Add WebDAV Share*. Refer to [WebDAV Shares](#) (page 184) for instructions on how to create a share and then how to connect to it once the service is configured and started.

The settings in the WebDAV service apply to all WebDAV shares. [Figure 11.20](#) shows the WebDAV configuration screen. [Table 11.16](#) summarizes the available options.

Fig. 11.20: WebDAV Configuration Screen

Table 11.16: WebDAV Configuration Options

Setting	Value	Description
Protocol	drop-down menu	choices are HTTP (connection always unencrypted), HTTPS (connection always encrypted), or HTTP+HTTPS (both types of connections allowed)
HTTP Port	string	only appears if the selected Protocol is HTTP or HTTP+HTTPS and is used to specify the port to be used for unencrypted connections; the default of 8080 should work, if you change it, do not use a port number already being used by another service
HTTPS Port	string	only appears if the selected Protocol is HTTPS or HTTP+HTTPS and is used to specify the port to be used for encrypted connections; the default of 8081 should work, if you change it, do not use a port number already being used by another service
Webdav SSL Certificate	drop-down menu	only appears if the selected Protocol is HTTPS or HTTP+HTTPS; select the SSL certificate to be used for encrypted connections; to create a certificate, use <i>System</i> → <i>Certificates</i>
HTTP Authentication	drop-down menu	choices are No Authentication, Basic Authentication (unencrypted), or Digest Authentication (encrypted)

Continued on next page

Table 11.16 – continued from previous page

Setting	Value	Description
Webdav Password	string	default is davtest; this should be changed as it is a known value

PLUGINS

FreeNAS® 8.2.0 introduced the ability to extend the built-in NAS services by providing a mechanism for installing additional software. This mechanism was known as the Plugins architecture and is based on [FreeBSD jails](https://en.wikipedia.org/wiki/Freebsd_jail) (https://en.wikipedia.org/wiki/Freebsd_jail) and [PC-BSD 9.x PBIs](http://wiki.pcbsd.org/index.php/AppCafe%C2%AE/9.2) (<http://wiki.pcbsd.org/index.php/AppCafe%C2%AE/9.2>). This allowed users to install and configure additional applications once they had created and configured a plugins jail.

FreeNAS® 9.x simplifies this procedure by providing two methods for software installation. The Plugins method, described in this section, is meant for users who prefer to browse for, install, and configure available software using the GUI. This method is very easy to use, but is limited in the amount of software that is available. Each application will automatically be installed into its own jail, meaning that this method may not be suitable for users who wish to run multiple applications within the same jail.

The Jails method provides much more control over software installation but assumes that the user is comfortable working from the command line and has a good understanding of networking basics and software installation on FreeBSD-based systems.

It is recommended that users skim through both the [Plugins](#) (page 244) and [Jails](#) (page 251) sections in order to become familiar with the features and limitations of each and to choose the method that best meets their software needs.

Note: Plugins created for FreeNAS® 9.3 or later are expected to work on the current release. Plugins created for earlier releases of FreeNAS® must be reinstalled.

12.1 Installing Plugins

A plugin is a self-contained application installer which has been designed to integrate into the FreeNAS® GUI. A plugin offers several advantages:

- the FreeNAS® GUI provides a browser for viewing the list of available plugins
- the FreeNAS® GUI provides buttons for installing, starting, managing, and deleting plugins
- if the plugin has configuration options, a screen will be added to the FreeNAS® GUI so that these options can be configured from the GUI

To install a plugin, click Plugins. As seen in [Figure 12.1](#), the list of available plugins will be displayed.



























Plugins		
Available Installed Configuration		
<div>Refresh</div> <div>Upload</div>		
Name	Description	Version
 bacula-sd	Network backup solution (server)	5.2.12_3
 CouchPotato	An automatic NZB and torrent downloader	autoupdate.9_5
 crashplan	Crashplan backs up data to remote servers, other computers, or hard drives	3.6.3_1
 Emby	A home media server built using other popular open source technologies	3.2.13.0
 firefly	Firefly (mt-daapd) is an open-source media server for the Roku SoundBridge and Apple iTunes	1696_8
 Headphones	Automatic music downloader for SABnzbd	9.3.1
 HTPC-Manager	A fully responsive interface to manage all your favorite software on your Htpc.	autoupdate.9_2
 LazyLibrarian	A program to follow authors and grab metadata for all your digital reading needs.	9.3.0
 Madsonic	A web-based media streamer and jukebox fork of Subsonic.	6.1
 Maraschino	A simple web interface to act as a nice overview/front page for an XBMC HTPC	9.3.1
 MineOS	A web interface to create and manage Minecraft server instances.	9.3.5
 Mylar	An automated Comic Book downloader (cbr/cbz) trying to follow in the lines of sickbeard and headphones.	9.3.2
 Nextcloud	Nextcloud is a system for the creation and management of personal cloud resources	10.0.1
 NZBHydra	A meta search for NZB indexers	autoupdate.9_1
 ownCloud	Owncloud is a system for the creation and management of personal cloud resources	9.1.2
 PlexMediaServer	The Plex Media Server component	1.5.5.3634
 Resilio	Distributed peer-to-peer file syncing application	2.4.4
 s3cmd	A plugin which allows you to backup a dataset to Amazon's S3 service	1.0.1_1
 SABnzbd	Open Source Binary Newsreader	2.0.0
 SickBeard	PVR for newsgroup users	9.3.1
 SickRage	Video File Manager for TV Shows	autoupdate.9_1
 Sonarr	Smart PVR for newsgroup and bittorrent users.	2.0.0.3732
 Subsonic	A free, web-based media streamer, providing ubiquitous access to your music.	6.0_2
 Syncthing	Open Source Continuous Replication / Cluster Synchronization Thing	0.14.3
 Transmission	A lightweight, yet powerful BitTorrent client	2.92
 XDM	eXtensible Download Manager. Plugin based media collection manager.	9.3.1

Fig. 12.1: Viewing the List of Available Plugins

Note: if the list of available plugins is not displayed, open [Shell](#) (page 290) and verify that the FreeNAS® system can ping an address on the Internet. If it cannot, you may have to add a default gateway address and/or DNS server address in [Network](#) → [Global Configuration](#).

Highlight the plugin you would like to install, click its Install button, then click OK. In the example shown in [Figure 12.2](#), SABnzbd is selected for installation.




Fig. 12.2: Installing a Plugin

The installation will take a few minutes as the system will first download and configure a jail to contain the installed software. It will then install the plugin and add it to the Installed tab as shown in [Figure 12.3](#).

Warning: Be patient and wait for the installation to finish. Navigating away from the installation before it is finished will cause problems with the installation.




Fig. 12.3: Viewing Installed PBIs

As seen in the example shown in [Figure 12.3](#), entries for the installed PBI will appear in the following locations:

- the Installed tab of Plugins
- the Plugins section of the tree
- the Jails section of the tree

The entry in the Installed tab of Plugins will display the plugin name and version, the name of the PBI that was installed, the name of the jail that was created, whether the application status is ON or OFF, and a button to delete the application and its associated jail. If a newer version of the application is available as a plugin, a button to update the application will also appear.

Note: The Service status of a plugin must be turned to ON before the installed application is available. Before starting the service, check to see if it has a configuration menu by clicking its entry in the Plugins section of the tree. If the application is configurable, this will open a screen that contains the available configuration options. Plugins which are not configurable will instead display a message with a hyperlink for accessing the software. However, that hyperlink does not work until the plugin is started.

Always review a plugin's configuration options before attempting to start it. Some plugins have options that need to be set before their service will successfully start. If you have never configured that application before, check the application's website to see what documentation is available. A link to the website for each available plugin can be found in [Available Plugins](#) (page 249).

If the application requires access to the data stored on the FreeNAS® system, click the entry for the associated jail in the Jails section of the tree and add a storage as described in [Add Storage](#) (page 257).

If you need to access the shell of the jail containing the application to complete or test your configuration, click the entry for the associated jail in the Jails section of the tree. You can then click its "shell" icon as described in [Managing Jails](#) (page 255).

Once the configuration is complete, click the red OFF button for the entry for the plugin. If the service starts successfully, it will change to a blue ON. If it fails to start, click the jail's Shell icon and type `tail /var/log/messages` to see if any errors were logged.

12.2 Updating Plugins

When a newer version of a plugin becomes available in the official repository, an Update button is added to the entry for the plugin in the Installed tab. In the example shown in [Figure 12.4](#), a newer version of Transmission is available.

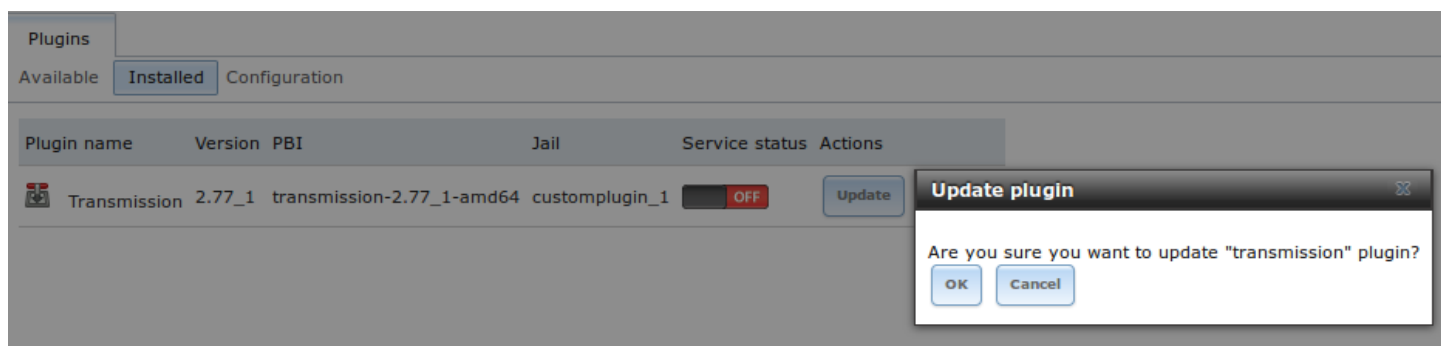


Fig. 12.4: Updating an Installed Plugin

Click the OK button to start the download and installation of the latest version of the plugin. Once the update is complete, the entry for the plugin will be refreshed to show the new version number and the Update button will disappear.

12.3 Uploading Plugins

The Available tab of Plugins contains an Upload button. This button allows installation of plugins that are not yet available in the official repository or which are still being tested. These plugins must be manually downloaded and should end in a .pbi extension. When downloading a plugin, make sure that it is 64-bit and that it was developed for 9.x. as 8.x and 10.x applications will not work on a 9.x FreeNAS® system.

Upload the new plugin with the Upload button. As seen in the example in [Figure 12.5](#), this prompts you to browse to the location of the plugin file. Select the file and click Upload to begin the installation.




Fig. 12.5: Installing a Previously Downloaded .pbi File

When the installation is complete, an entry for the plugin will be added to the Installed tab and its associated jail is listed under Jails. However, if it is not a FreeNAS® plugin, it will not be added to Plugins in the tree. In this case, any required jail configuration must be done from the command line of the jail's shell instead of from the GUI.

12.4 Deleting Plugins

When you install a plugin, an associated jail is created. If you decide to delete a plugin, the associated jail is also deleted as it is no longer required. Before deleting a plugin, make sure that you do not have any data or configuration in the jail that

you need to save. If you do, back up that data first, before deleting the plugin.

In the example shown in Figure 12.6, Sabnzbd has been installed and the user has clicked its Delete button. A pop-up message asks the user if they are sure that they want to delete. This is the one and only warning. If the user clicks Yes, the plugin and the associated jail are permanently deleted.




Fig. 12.6: Deleting an Installed Plugin

12.5 Available Plugins

These plugins are available for FreeNAS® 11.1:

- [bacula-sd \(storage daemon\)](http://bacula.org/) (<http://bacula.org/>)
- [CouchPotato](https://couchpota.to/) (<https://couchpota.to/>)
- [crashplan](http://www.code42.com/crashplan/) (<http://www.code42.com/crashplan/>)
- [Emby](http://emby.media/) (<http://emby.media/>)
- [firefly](https://en.wikipedia.org/wiki/Firefly_Media_Server) (https://en.wikipedia.org/wiki/Firefly_Media_Server)
- [Headphones](https://github.com/rembo10/headphones) (<https://github.com/rembo10/headphones>)
- [HTPC-Manager](http://htpc.io/) (<http://htpc.io/>)
- [LazyLibrarian](https://github.com/lazylibrarian/LazyLibrarian) (<https://github.com/lazylibrarian/LazyLibrarian>)
- [Madsonic](http://madsonic.org/) (<http://madsonic.org/>)
- [Maraschino](http://www.maraschinoproject.com/) (<http://www.maraschinoproject.com/>)
- [MineOS](http://minecraft.codeemo.com/) (<http://minecraft.codeemo.com/>)
- [Mylar](https://github.com/evilhero/mylar) (<https://github.com/evilhero/mylar>)
- [Nextcloud](https://nextcloud.com/) (<https://nextcloud.com/>)
- [NZBHydra](https://github.com/theotherp/nzbhydra) (<https://github.com/theotherp/nzbhydra>)
- [ownCloud](https://owncloud.org/) (<https://owncloud.org/>)
- [PlexMediaServer](https://plex.tv/) (<https://plex.tv/>)
- [Resilio](https://www.resilio.com/) (<https://www.resilio.com/>)
- [s3cmd](http://s3tools.org/s3cmd) (<http://s3tools.org/s3cmd>)
- [SABnzbd](http://sabnzbd.org/) (<http://sabnzbd.org/>)

- [SickBeard](http://sickbeard.com/) (<http://sickbeard.com/>)
- [SickRage](https://github.com/SiCKRAGETV/SickRage) (<https://github.com/SiCKRAGETV/SickRage>)
- [Sonarr](https://sonarr.tv/) (<https://sonarr.tv/>)
- [Subsonic](http://www.subsonic.org/pages/index.jsp) (<http://www.subsonic.org/pages/index.jsp>)
- [Syncthing](https://syncthing.net/) (<https://syncthing.net/>)
- [Transmission](http://www.transmissionbt.com/) (<http://www.transmissionbt.com/>)
- [XDM](https://github.com/lad1337/XDM) (<https://github.com/lad1337/XDM>)
- [XMRig](https://github.com/xmrig/xmrig) (<https://github.com/xmrig/xmrig>)

While the FreeNAS® Plugins system makes it easy to install software, it is still up to you to know how to configure and use the installed application. When in doubt, refer to the documentation for that application.

JAILS

The previous section described how to find, install, and configure software using [Plugins](#) (page 244).

This section describes how to use jails, which allow users who are comfortable with the command line to have more control over software installation and management. Any software installed using Jails must be managed from the command line of the jail. If you prefer to use a GUI to manage software, use [Plugins](#) (page 244) instead.

Note: The jails infrastructure is transitioning from the old warden backend to the new iocage backend. This transition process requires the middleware API calls to be rewritten for the new UI. It is expected that the transition will be complete with FreeNAS® version 11.2. Since jails created in the old UI use the warden backend, jails created in the new UI use the iocage backend, and both use different API versions, they are not compatible. While a migration script will be made available when the transition is complete, it will not be able to anticipate every configuration scenario for every application installed in jails. At that time, the recommendation will be to: create new jails using the new UI, copy over any existing configurations, and delete the old jail datasets once the new jails are working as expected.

FreeNAS® automatically creates a jail whenever a plugin is installed, but does not let the user install multiple plugins into the same jail. In contrast, using Jails allows users to create as many jails as needed and to customize the operating system and installed software within each jail.

By default, a [FreeBSD jail](https://en.wikipedia.org/wiki/Freebsd_jail) (https://en.wikipedia.org/wiki/Freebsd_jail) is created. This provides a very light-weight, operating system-level virtualization. Consider it as another independent instance of FreeBSD running on the same hardware, without all of the overhead usually associated with virtualization. The jail will install the FreeBSD software management utilities so FreeBSD ports can be compiled and FreeBSD packages can be installed from the command line of the jail.

It is important to understand that any users, groups, installed software, and configurations within a jail are isolated from both the FreeNAS® operating system and any other jails running on that system. During creation, the VIMAGE option can be selected to provide the jail with an independent networking stack. The jail can then do its own IP broadcasting, which is required by some applications.

Advanced users can also create custom templates to automate the creation of pre-installed and customized operating systems.

The ability to create multiple jails running different operating systems offers great flexibility regarding software management. For example, the administrator can choose to provide application separation by installing different applications in each jail, or to create one jail for all installed applications, or to mix and match how software is installed into each jail.

Note: Jails created with FreeNAS® 9.3 or later are expected to work with the current release. Jails created on older versions of FreeNAS® must be reinstalled due to ABI changes.

The rest of this section describes:

- [Jails Configuration](#) (page 252)
- [Adding Jails](#) (page 253)
- [Managing Jail Templates](#) (page 264)

- [Using iocage](#) (page 267)

13.1 Jails Configuration

Jails are stored in a volume or dataset. Using a separate dataset for the Jail Root is strongly recommended. The volume or dataset to be used must already exist or can be created with [Volume Manager](#) (page 115).

Note: The Jail Root volume or dataset cannot be created on a [Share](#) (page 169).

Begin global jail configuration by choosing [Jails](#) → [Configuration](#) to open the screen shown in [Figure 13.1](#). Jails are automatically installed into their own dataset under the specified path as they are created. For example, if the Jail Root is set to `/mnt/volume1/dataset1` and a jail named `jail1` is created, it is installed into its own dataset named `/mnt/volume1/dataset1/jail1`.

Fig. 13.1: Global Jail Configuration

Warning: If any [Plugins](#) (page 244) have already been installed, the Jail Root, IPv4 Network, IPv4 Network Start Address, and IPv4 Network End Address are automatically filled. Double-check that the pre-configured IP address values are appropriate for the jails and do not conflict with addresses used by other systems on the network.

[Table 13.1](#) summarizes the fields in this configuration screen. Refer to the text below the table for more details on how to properly configure the Jail Root and network settings. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in [System](#) → [Advanced](#).

Table 13.1: Jail Configuration Options

Setting	Value	Advanced Mode	Description
Jail Root	browse button		mandatory; jails cannot be added until this is set
IPv4 DHCP	checkbox		check this box if the network has a DHCP server
IPv4 Network	string	✓	format is IP address of network/CIDR mask
IPv4 Network Start Address	string	✓	enter the first IP address in the reserved range in the format host/CIDR mask
IPv4 Network End Address	string	✓	enter the last IP address in the reserved range in the format host/CIDR mask

Continued on next page

Table 13.1 – continued from previous page

Setting	Value	Advanced Mode	Description
IPv6 Autoconfigure	checkbox		check this box if the network has a DHCPv6 server and IPv6 will be used to access jails
IPv6 Network	string	✓	enter the network address for a properly configured IPv6 network
IPv6 Network Start Address	string	✓	enter the first IP address in the reserved range for a properly configured IPv6 network
IPv6 Network End Address	string	✓	enter the last IP address in the reserved range for a properly configured IPv6 network
Collection URL	string	✓	changing the default may break the ability to install jails

When selecting the Jail Root, ensure that the size of the selected volume or dataset is sufficient to hold the number of jails to be installed as well as any software, log files, and data to be stored within each jail. At a bare minimum, budget at least 2 GB per jail and do not select a dataset that is less than 2 GB in size.

Note: If you plan to add storage to a jail, be aware that the path size is limited to 88 characters. Make sure that the length of the volume name plus the dataset name plus the jail name does not exceed this limit.

If the network contains a DHCP server, it is recommended to check the box IPv4 DHCP (or IPv6 Autoconfigure, for a properly configured IPv6 network). This will prevent IP address conflicts on the network as the DHCP server will automatically assign the jail the next available lease and record the lease as in use.

If a static IP address is needed so that users always know the IP address of the jail, enter the start and end address for the IPv4 and/or IPv6 network. The range defined by the start and end addresses will be automatically assigned as jails are created. For example, if you plan to create 5 jails on the 192.168.1.0 network, enter a IPv4 Network Start Address of 192.168.1.100 and a IPv4 Network End Address of 192.168.1.104.

If you create a start and end range on a network that contains a DHCP server, it is very important that you also reserve those addresses on the DHCP server. Otherwise, the DHCP server will not be aware that those addresses are being used by jails and there will be IP address conflicts and weird networking errors on the network. When troubleshooting jails that do not install or which are unavailable, double-check that the IP address being used by the jail is not also being used by another jail or system in the network.

FreeNAS® will automatically detect and display the IPv4 Network to which the administrative interface is connected. This setting is important. The IP addresses used by the jails must be pingable from the FreeNAS® system for the jails and any installed software to be accessible. If the network topology requires changing the default value, a default gateway and possibly a static route need to be added to the specified network. After changing this value, ensure that the subnet mask value is correct, as an incorrect mask can make the IP network unreachable. When in doubt, keep the default setting for IPv4 Network. With VMware, make sure that the vswitch is set to “promiscuous mode”. With VirtualBox, make sure `Network -> Advanced -> Promiscuous Mode` is not set to “Deny”.

After clicking the Save button to save the configuration, the system is ready to create and manage jails as described in the rest of this chapter.

13.2 Adding Jails

To create a jail, click `Jails -> Add Jail` to access the screen shown in [Figure 13.2](#).

Note: the Add Jail menu item will not appear until after you configure `Jails -> Configuration`.




Fig. 13.2: Creating a Jail

By default, the only required value to create a jail is a name. FreeBSD jails are created by default.

Table 13.2 summarizes the available options. Most settings are only available in Advanced Mode and are not needed if the intent is to create a FreeBSD jail. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in *System* → *Advanced*.

Table 13.2: Jail Configuration Options

Setting	Value	Advanced Mode	Description
Jail Name	string		mandatory; can only contain letters, numbers, dashes, or the underscore character
Template	drop-down menu	✓	contains any created custom templates as described in Managing Jail Templates (page 264)
IPv4 DHCP	checkbox	✓	if unchecked, make sure that the defined address does not conflict with the DHCP server's pool of available addresses
IPv4 address	integer	✓	this and the other IPv4 settings are grayed out if IPv4 DHCP is checked; enter a unique IP address that is in the local network and not already used by anyother computer
IPv4 netmask	drop-down menu	✓	select the subnet mask associated with IPv4 address
IPv4 bridge address	integer	✓	grayed out unless VIMAGE is checked; see NOTE below
IPv4 bridge netmask	drop-down menu	✓	select the subnet mask associated with IPv4 bridge address; grayed out unless VIMAGE is checked
IPv4 default gateway	string	✓	grayed out unless VIMAGE is checked
IPv6 Autoconfigure	checkbox	✓	if unchecked, make sure that the defined address does not conflict with the DHCP server's pool of available addresses
IPv6 address	integer	✓	this and other IPv6 settings are grayed out if IPv6 Autoconfigure is checked; enter a unique IPv6 address that is in the local network and not already used by any other computer
IPv6 prefix length	drop-down menu	✓	select the prefix length associated with IPv6 address
IPv6 bridge address	integer	✓	grayed out unless VIMAGE is checked; see NOTE below
IPv6 bridge prefix length	drop-down menu	✓	grayed out unless VIMAGE is checked; select the prefix length associated with IPv6 address

Continued on next page

Table 13.2 – continued from previous page

Setting	Value	Advanced Mode	Description
IPv6 default gateway	string	✓	grayed out unless VIMAGE is checked; used to set the jail's default gateway IPv6 address
MAC	string	✓	grayed out unless VIMAGE is checked; if a static MAC address is entered, unique static MAC addresses must be entered for every jail created
NIC	drop-down menu	✓	grayed out if VIMAGE is checked; can be used to specify the interface to use for jail connections
Sysctls	string	✓	comma-delimited list of sysctls to set inside jail (like allow.sysvipc=1,allow.raw_sockets=1)
Autostart	checkbox	✓	uncheck if the jail will be started manually
VIMAGE	checkbox	✓	gives a jail its own virtualized network stack; requires promiscuous mode be enabled on the interface
NAT	checkbox	✓	grayed out for Linux jails or if VIMAGE is unchecked; enables Network Address Translation for the jail

Note: The IPv4 and IPv6 bridge interface is used to bridge the [epair\(4\)](http://www.freebsd.org/cgi/man.cgi?query=epair) (<http://www.freebsd.org/cgi/man.cgi?query=epair>) device, which is automatically created for each started jail, to a physical network device. The default network device is the one that is configured with a default gateway. So, if em0 is the FreeBSD name of the physical interface and three jails are running, these virtual interfaces are automatically created: bridge0, epair0a, epair1a, and epair2a. The physical interface em0 will be added to the bridge, as well as each epair device. The other half of the epair will be placed inside the jail and will be assigned the IP address specified for that jail. The bridge interface will be assigned an alias of the default gateway for that jail, if configured, or the bridge IP, if configured; either is correct.

The only time an IP address and mask are required for the bridge is when the jail will be on a different network than the FreeNAS® system. For example, if the FreeNAS® system is on the 10.0.0.0/24 network and the jail will be on the 192.168.0.0/24 network, set the IPv4 bridge address and IPv4 bridge netmask fields for the jail.

If both the VIMAGE and NAT boxes are unchecked, the jail must be configured with an IP address within the same network as the interface it is bound to, and that address will be assigned as an alias on that interface. To use a VIMAGE jail on the same subnet, uncheck NAT and configure an IP address within the same network. In both of these cases, configure only an IP address and do not configure a bridge or a gateway address.

After making selections, click the OK button. The jail is created and added to the Jails tab as well as in the tree menu under Jails. Jails start automatically. To prevent this, uncheck the Autostart box.

The first time a jail is added or used as a template, the GUI automatically downloads the necessary components from the internet. A progress bar indicates the status of the download and provides an estimated time for the process to complete. If it is unable to connect to the internet, jail creation fails.

Warning: Failure to download is often caused by the default gateway not being set, preventing internet access. See the Network [Global Configuration](#) (page 102) section for information on setting the default gateway.

After the first jail is created or a template has been used, subsequent jails will be added very quickly because the downloaded base for creating the jail has been saved to the Jail Root.

13.2.1 Managing Jails

Click Jails to view and configure the added jails. In the example shown in [Figure 13.3](#), the list entry for the jail named xdm_1 has been clicked to enable that jail's configuration options. The entry indicates the jail name, IP address, whether it will start automatically at system boot, if it is currently running, and jail type: standard for a FreeBSD jail, or pluginjail if it was installed using [Plugins](#) (page 244).




Fig. 13.3: Viewing Jails

From left to right, these configuration icons are available:

Edit Jail: edit the jail settings which were described in [Table 13.2](#).

After a jail has been created, the jail name and type cannot be changed, so these fields will be grayed out.

Note: To modify the IP address information for a jail, use the Edit Jail button instead of the associated networking commands from the command line of the jail.

Add Storage: configure the jail to access an area of storage as described in [Add Storage](#) (page 257).

Upload Plugin: manually upload a plugin previously downloaded from the [plugins repository](http://download.freenas.org/plugins/9/x64/) (<http://download.freenas.org/plugins/9/x64/>).

Start/Stop: this icon changes appearance depending on the current Status of the jail. When the jail is not running, the icon is green and clicking it starts the jail. When the jail is already running, the icon is red and clicking it stops the jail. A stopped jail and its applications are inaccessible until it is restarted.

Restart: restart the jail.

Shell: access a root command prompt to configure the selected jail from the command line. When finished, type `exit` to close the shell.

Delete: delete the jail and any periodic snapshots of it. The contents of the jail are entirely removed.

Warning: Back up data and programs in the jail before deleting it. There is no way to recover the contents of a jail after deletion.

Accessing a Jail Using SSH

`ssh` can be used to access a jail instead of the jail's Shell icon. This requires starting the `ssh` service and creating a user account for `ssh` access. Start by clicking the Shell icon for the desired jail.

Find the `sshd_enable=` line in the jail's `/etc/rc.conf` and set it to "YES":

```
sshd_enable="YES"
```

Then start the SSH daemon:

```
service sshd start
```

The first time the service runs, the jail's RSA key pair is generated and the key fingerprint and random art image displayed.

Add a user account by typing `adduser` and following the prompts. If the user needs superuser privileges, they must be added to the wheel group. For those users, enter `wheel` at this prompt:

```
Login group is user1. Invite user1 into other groups? []: wheel
```

After creating the user, set the root password so that the new user will be able to use the `su` command to gain superuser privilege. To set the password, type `passwd` then enter and confirm the desired password.

Finally, test from another system that the user can successfully `ssh` in and become the superuser. In this example, a user named `user1` uses `ssh` to access the jail at `192.168.2.3`. The first time the user logs in, they will be asked to verify the fingerprint of the host:

```
ssh user1@192.168.2.3
The authenticity of host '192.168.2.3 (192.168.2.3)' can't be established.
RSA key fingerprint is 6f:93:e5:36:4f:54:ed:4b:9c:c8:c2:71:89:c1:58:f0.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.2.3' (RSA) to the list of known hosts.
Password: type_password_here
```

Note: Each jail has its own user accounts and service configuration. These steps must be repeated for each jail that requires SSH access.

Add Storage

It is possible to give a FreeBSD jail access to an area of storage on the FreeNAS® system. This is useful for applications that store a large amount of data or if an application in a jail needs access to the data stored on the FreeNAS® system. One example is transmission, which stores torrents. The storage is added using the `mount_nullfs(8)` (http://www.freebsd.org/cgi/man.cgi?query=mount_nullfs) mechanism, which links data that resides outside of the jail as a storage area within the jail.

To add storage, click the Add Storage button for a highlighted jail's entry to open the screen shown in [Figure 13.4](#). This screen can also be accessed by expanding the jail name in the tree view and clicking `Storage` → `Add Storage`.

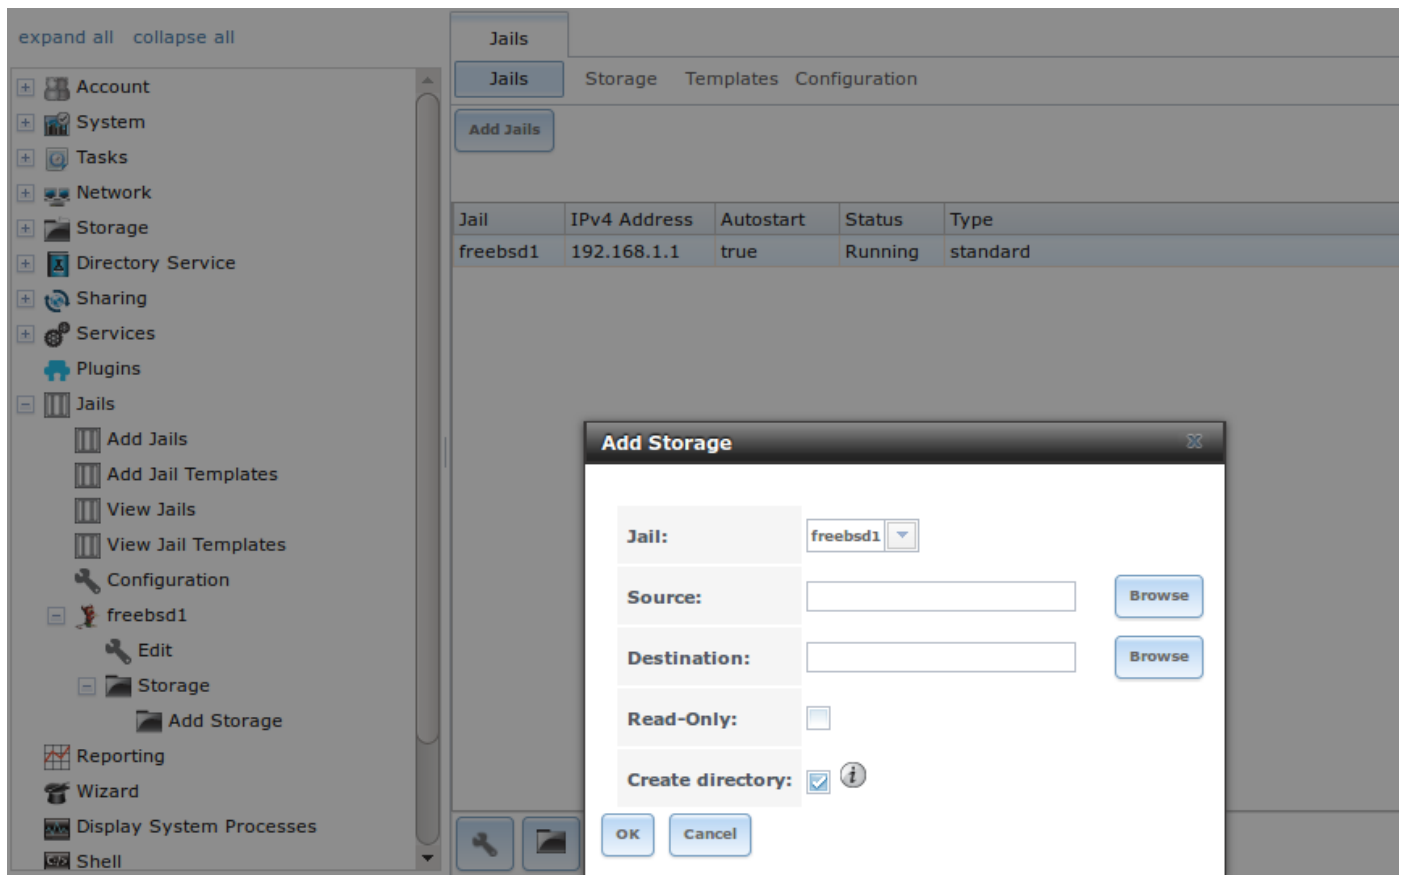


Fig. 13.4: Adding Storage to a Jail

Browse to the Source and Destination, where:

- **Source:** is the directory or dataset on the FreeNAS® system which will be accessed by the jail. This directory must reside outside of the volume or dataset being used by the jail. This is why it is recommended to create a separate dataset to store jails, so the dataset holding the jails is always separate from any datasets used for storage on the FreeNAS® system.
- **Destination:** select an existing, empty directory within the jail to link to the Source storage area. If that directory does not exist yet, enter the desired directory name and check the Create directory box.

Storage is typically added because the user and group account associated with an application installed inside of a jail needs to access data stored on the FreeNAS® system. Before selecting the Source, it is important to first ensure that the permissions of the selected directory or dataset grant permission to the user/group account inside of the jail. This is not the default, as the users and groups created inside of a jail are totally separate from the users and groups of the FreeNAS® system.

The workflow for adding storage usually goes like this:

1. Determine the name of the user and group account used by the application. For example, the installation of the transmission application automatically creates a user account named transmission and a group account also named transmission. When in doubt, check the files `/etc/passwd` (to find the user account) and `/etc/group` (to find the group account) inside the jail. Typically, the user and group names are similar to the application name. Also, the UID and GID are usually the same as the port number used by the service.

A media user and group (GID 8675309) are part of the base system. Having applications run as this group or user makes it possible to share storage between multiple applications in a single jail, between multiple jails, or even between the host and jails.

2. On the FreeNAS® system, create a user account and group account that match the user and group names used by the application in the jail.

3. Decide whether the jail should have access to existing data or if a new area of storage will be set aside for the jail to use.
4. If the jail will access existing data, edit the permissions of the volume or dataset so the user and group accounts have the desired read and write access. If multiple applications or jails are to have access to the same data, create a new group and add each needed user account to that group.
5. If an area of storage is being set aside for that jail or individual application, create a dataset. Edit the permissions of that dataset so the user and group account has the desired read and write access.
6. Use the Add Storage button of the jail and select the configured volume/dataset as the Source.

To prevent writes to the storage, check the box Read-Only.

By default, the Create directory box is checked. This means that the directory will automatically be created under the specified Destination path if the directory does not already exist.

After storage has been added or created, it appears in the tree under the specified jail. In the example shown in Figure 13.5, a dataset named `volume1/data` has been chosen as the Source as it contains the files stored on the FreeNAS® system. When the storage was created, the user browsed to `volume1/jails/freebsd1/usr/local` in the Destination field, then entered `test` as the directory. Since this directory did not already exist, it was created, because the Create directory box was left checked. The resulting storage was added to the `freenas1` entry in the tree as `/usr/local/test`. The user has clicked this `/usr/local/test` entry to access the Edit screen.




Fig. 13.5: Example Storage

Storage is normally mounted as it is created. To unmount the storage, uncheck the Mounted? box.

Note: A mounted dataset will not automatically mount any of its child datasets. While the child datasets may appear to be browsable inside the jail, any changes will not be visible. Since each dataset is considered to be its own filesystem, each child dataset must have its own mount point, so separate storage must be created for any child datasets which need to be mounted.

To delete the storage, click its Delete button.

Warning: It is important to realize that added storage is really just a pointer to the selected storage directory on the FreeNAS® system. It does not copy that data to the jail. Files that are deleted from the Destination directory in the jail are really deleted from the Source directory on the FreeNAS® system. However, removing the jail storage entry only removes the pointer, leaving the data intact but not accessible from the jail.

13.2.2 Installing FreeBSD Packages

The quickest and easiest way to install software inside the jail is to install a FreeBSD package. FreeBSD packages are pre-compiled. They contain all the binaries and a list of dependencies required for the software to run on a FreeBSD system.

A huge amount of software has been ported to FreeBSD, currently over 24,000 applications, and most of that software is available as a package. One way to find FreeBSD software is to use the search bar at [FreshPorts.org](http://www.freshports.org/) (<http://www.freshports.org/>).

After finding the name of the desired package, use the `pkg install` command to install it. For example, to install the audiotag package, use this command:

```
pkg install audiotag
```

When prompted, type `y` to complete the installation. The installation messages will indicate if the package and its dependencies successfully download and install.

Warning: Some older versions of FreeBSD used package systems which are now obsolete. Do not use commands from those obsolete package systems in a FreeNAS® jail, as they will cause inconsistencies in the jail's package management database. Use the current FreeBSD package system as shown in these examples.

A successful installation can be confirmed by querying the package database:

```
pkg info -f audiotag
audiotag-0.19_1
Name:          audiotag
Version:       0.19_1
Installed on:  Fri Nov 21 10:10:34 PST 2014
Origin:        audio/audiotag
Architecture:  freebsd:9:x86:64
Prefix:        /usr/local
Categories:    multimedia audio
Licenses:      GPLv2
Maintainer:    ports@FreeBSD.org
WWW:           http://github.com/Daenyth/audiotag
Comment:       Command-line tool for mass tagging/renaming of audio files
Options:
  DOCS:         on
  FLAC:         on
  ID3:          on
  MP4:          on
  VORBIS:       on
Annotations:
  repo_type:    binary
  repository:   FreeBSD
Flat size:     62.8KiB
Description:    Audiotag is a command-line tool for mass tagging/renaming of audio files
                 it supports the vorbis comment, id3 tags, and MP4 tags.
WWW:           http://github.com/Daenyth/audiotag
```

To show what was installed by the package:

```
pkg info -l audiotag
audiotag-0.19_1:
/usr/local/bin/audiotag
/usr/local/share/doc/audiotag/COPYING
/usr/local/share/doc/audiotag/ChangeLog
/usr/local/share/doc/audiotag/README
/usr/local/share/licenses/audiotag-0.19_1/GPLv2
/usr/local/share/licenses/audiotag-0.19_1/LICENSE
/usr/local/share/licenses/audiotag-0.19_1/catalog.mk
```

In FreeBSD, third-party software is always stored in `/usr/local` to differentiate it from the software that came with the operating system. Binaries are almost always located in a subdirectory called `bin` or `sbin` and configuration files in a subdirectory called `etc`.

13.2.3 Compiling FreeBSD Ports

Software is typically installed into FreeBSD jails using packages. But sometimes there are good reasons to compile a port instead. Compiling ports offers these advantages:

- Not every port has an available package. This is usually due to licensing restrictions or known, unaddressed security vulnerabilities.
- Sometimes the package is out-of-date and a feature is needed that only became available in the newer version.
- Some ports provide compile options that are not available in the pre-compiled package. These options are used to add or remove features or options.

Compiling a port has these disadvantages:

- It takes time. Depending upon the size of the application, the amount of dependencies, the speed of the CPU, the amount of RAM available, and the current load on the FreeNAS® system, the time needed can range from a few minutes to a few hours or even to a few days.

Note: If the port does not provide any compile options, it saves time and preserves the FreeNAS® system's resources to just use the `pkg install` command instead.

The [FreshPorts.org](http://www.freshports.org/) (<http://www.freshports.org/>) listing shows whether a port has any configurable compile options. [Figure 13.6](#) shows the Configuration Options for audiotag.

Port details

audiotag Command-line tool for mass tagging/renaming of audio files
0.19_1 [audio](#)

There is no maintainer for this port.
 Any concerns regarding this port should be directed to the FreeBSD Ports mailing list via ports@FreeBSD.org

Port Added: 15 Apr 2008 13:43:37
Also Listed In: [multimedia](#)
License: GPLv2+

Audiotag is a command-line tool for mass tagging/renaming of audio files
 it supports the vorbis comment, id3 tags, and MP4 tags.

WWW: <http://github.com/Daenyth/audiotag>
SVNWeb: [Homepage](#) : [PortsMon](#)

To install [the port](#): `cd /usr/ports/audio/audiotag/ && make install clean`
To add [the package](#): `pkg install audiotag`

PKGNAME: audiotag

distinfo:

```

SHA256 (audiotag-0.19.tar.bz2) = 7b6a2de751058a95755f0842b83f2b1d8b94e5cd7634cbe71d67257208bf4646
SIZE (audiotag-0.19.tar.bz2) = 15016

```

NOTE: FreshPorts displays only information on required and default dependencies. Optional dependencies are not covered.

Runtime dependencies:

1. flac : [audio/flac](#)
2. id3tag : [audio/id3lib](#)
3. AtomicParsley : [multimedia/atomicparsley](#)
4. vorbiscomment : [audio/vorbis-tools](#)
5. perl5<=5.20<5.21 : [lang/perl5.20](#)

There are no ports dependent upon this port

Configuration Options

```

==> The following configuration options are available for audiotag-0.19_1:
DOCS=on: Build and/or install documentation
FLAC=on: FLAC lossless audio codec support
ID3=on: ID3 tags support
MP4=on: MP4 media format support
VORBIS=on: Ogg Vorbis audio codec support
==> Use 'make config' to modify these settings

```

Search

Enter Keywords:

[more...](#)

Latest Vulnerabilities

file-roller	Oct 12
ImageMagick	Oct 12
ImageMagick-nox11	Oct 12
lftp	Oct 12
mupdf	Oct 12
openoffice-4	Oct 12
openoffice-devel	Oct 12
virtualbox-ose	Oct 12
zathura-pdf-mupdf	Oct 12
libvncserver	Oct 11
linux-c6-openssl*	Oct 11
openjpeg	Oct 11
openssl*	Oct 11
openssl-devel*	Oct 11

Fig. 13.6: Configuration Options for Audiotag

This port has five configurable options (DOCS, FLAC, ID3, MP4, and VORBIS) and each option is enabled (on) by default.

FreeBSD packages are always built using the default options. When compiling a port yourself, those options are presented in a menu, allowing the default values to be changed.

The Ports Collection must be installed in a jail before ports can be compiled. Inside the jail, use the `portsnap` utility. This command downloads the ports collection and extracts it to the jail's `/usr/ports/` directory:

```
portsnap fetch extract
```

Note: To install additional software at a later date, make sure the ports collection is updated with `portsnap fetch update`.

To compile a port, `cd` into a subdirectory of `/usr/ports/`. The entry for the port at FreshPorts provides the location to `cd` into and the `make` command to run. This example compiles and installs the audiotag port:

```
cd /usr/ports/audio/audiotag
make install clean
```

Since this port has configurable options, the first time this command is run, the configure screen shown in [Figure 13.7](#) is displayed:

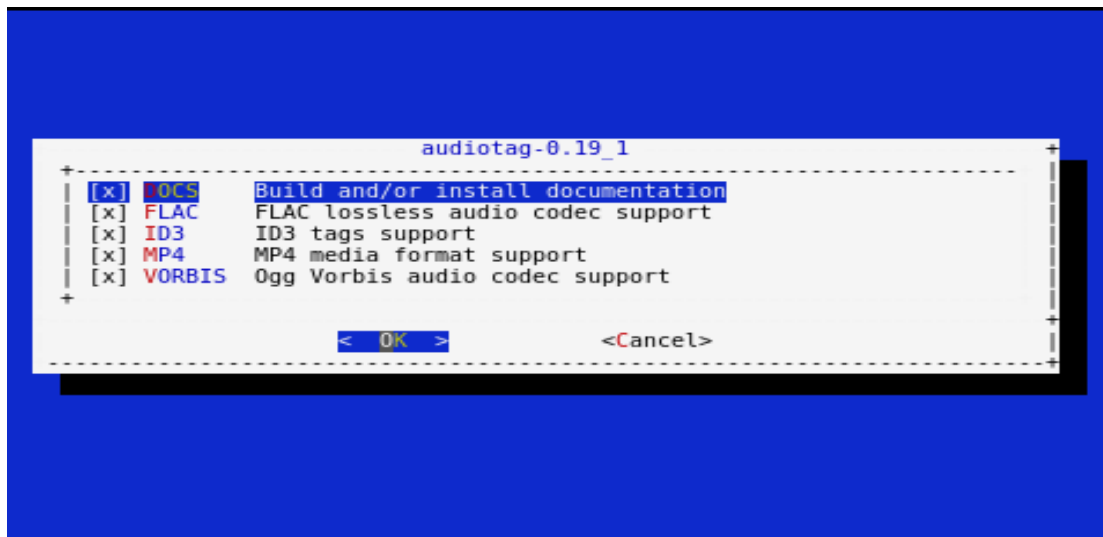


Fig. 13.7: Configuration Options for Audiotag Port

Use the arrow keys to select an option and press `spacebar` to toggle the value. When all the values are as desired, press `Enter`. The port will begin to compile and install.

Note: The configuration screen will not be shown again, even if the build is stopped and restarted. It can be redisplayed by typing `make config`. Change the settings, then rebuild with `make clean install clean`.

Many ports depend on other ports. Those other ports can also have configuration screens that will be shown before compiling begins. It is a good idea to keep an eye on the compile until it finishes and the command prompt returns.

When the port is installed, it is registered in the same package database that manages packages. The same `pkg info` command can be used to determine what was installed, as described in the previous section.

13.2.4 Starting Installed Software

After packages or ports are installed, they need to be configured and started. If you are familiar with the software, look for the configuration file in `/usr/local/etc` or a subdirectory of it. Many FreeBSD packages contain a sample configuration file as a reference. If you are unfamiliar with the software, you will need to spend some time at the software's website to learn which configuration options are available and which configuration files require editing.

Most FreeBSD packages that contain a startable service include a startup script which is automatically installed to `/usr/local/etc/rc.d/`. After the configuration is complete, the starting of the service can be tested by running the script with the `onestart` option. As an example, if `openvpn` is installed into the jail, these commands run its startup script and verify that the service started:

```
/usr/local/etc/rc.d/openvpn onestart
Starting openvpn.

/usr/local/etc/rc.d/openvpn onestatus
openvpn is running as pid 45560.

sockstat -4
USER COMMAND      PID    FD    PROTO  LOCAL ADDRESS  FOREIGN ADDRESS
root openvpn      48386   4     udp4   *:54789        *:*
```

If it produces an error:

```
/usr/local/etc/rc.d/openvpn onestart
Starting openvpn.
/usr/local/etc/rc.d/openvpn: WARNING: failed to start openvpn
```

Run `tail /var/log/messages` to see if any error messages hint at the problem. Most startup failures are related to a misconfiguration: either a typo or a missing option in a configuration file.

After verifying that the service starts and is working as intended, add a line to `/etc/rc.conf` to start the service automatically when the jail is started. The line to start a service always ends in `_enable="YES"` and typically starts with the name of the software. For example, this is the entry for the `openvpn` service:

```
openvpn_enable="YES"
```

When in doubt, the startup script shows the line to put in `/etc/rc.conf`. This is the description in `/usr/local/etc/rc.d/openvpn`:

```
# This script supports running multiple instances of openvpn.
# To run additional instances link this script to something like
# % ln -s openvpn openvpn_foo

# and define additional openvpn_foo_* variables in one of
# /etc/rc.conf, /etc/rc.conf.local or /etc/rc.conf.d /openvpn_foo

#
# Below NAME should be substituted with the name of this script. By default
# it is openvpn, so read as openvpn_enable. If you linked the script to
# openvpn_foo, then read as openvpn_foo_enable etc.
#
# The following variables are supported (defaults are shown).
# You can place them in any of
# /etc/rc.conf, /etc/rc.conf.local or /etc/rc.conf.d/NAME
#
# NAME_enable="NO"
# set to YES to enable openvpn
```

The startup script also indicates if any additional parameters are available:

```
# NAME_if=
# driver(s) to load, set to "tun", "tap" or "tun tap"
#
# it is OK to specify the if_ prefix.
#
# # optional:
# NAME_flags=
# additional command line arguments
# NAME_configfile="/usr/local/etc/openvpn/NAME.conf"
# --config file
# NAME_dir="/usr/local/etc/openvpn"
# --cd directory
```

13.3 Managing Jail Templates

FreeNAS® supports the ability to add custom templates to the Templates drop-down menu described in [Table 13.2](#).

To create a custom template, first install the desired operating system and configure it as needed. The installation can be either to an existing jail or on another system.

Next, create an `mtree` specification using this command, replacing `/path/to/jail` with the actual path to the jail:


```
mtree -c -p /path/to/jail -k sha256digest > file.mtree
```

After configuration is complete, create a tarball of the entire operating system to be used as a template. This tarball needs to be compressed with `gzip` and end in a `.tgz` extension. Be careful when creating the tarball as it is possible to end up in a recursive loop. In other words, the resulting tarball must be saved outside of the operating system being tarballed, such as to an external USB drive or network share. Alternately, create a temporary directory within the operating system and use the `-exclude` switch to `tar` to exclude this directory from the tarball. The exact `tar` command to use will vary, depending upon the operating system being used to create the tarball.

Save the generated `.mtree` and `.tgz` files to either an FTP share or an HTTP server. The FTP or HTTP URL is needed to add the template to the list of available templates.

To add the template, click **Jails** → **Templates** → **Add Jail Templates** which opens the screen shown in [Figure 13.8](#).


Fig. 13.8: Adding A Custom Jail Template

[Table 13.3](#) summarizes the fields in this screen.

Table 13.3: Jail Template Options

Setting	Value	Description
Name	string	value appears in the Name column of View Jail Templates
OS	drop-down menu	choices are FreeBSD or Linux
Architecture	drop-down menu	choices are x86 (32-bit) or x64 (64-bit)
URL	string	enter the full URL to the <code>.tgz</code> file, including the protocol (<code>ftp://</code> or <code>http://</code>)
Mtree	string	paste the mtree specification for the template

Added templates appear in **Jails** → **Templates**. An example is shown in [Figure 13.9](#).



The screenshot shows the 'Templates' tab in the FreeNAS interface. At the top, there are tabs for 'Jails', 'Storage', 'Templates', and 'Configuration'. Below the tabs is a button labeled 'Add Jail Template'. A table lists available templates with columns for Name, URL, and Instances.


Name	URL	Instances
freenas-pluginjail	http://download.freenas.org/jails/10/x64/freenas-pluginjail-10.3-RELEASE.tgz	0
freenas-standard	http://download.freenas.org/jails/10/x64/freenas-standard-10.3-RELEASE.tgz	0

Fig. 13.9: Viewing Available Templates

The listing contains these columns:

- Name: appears in the Template drop-down menu when adding a new jail.
- URL: when adding a new jail using this template, the template is downloaded from this location.
- Instances: indicates if the template has been used to create a jail. In this example, the template has not yet been used, so Instances shows as 0.

Click the entry for a template to access its Edit and Delete buttons. Clicking a template's Edit button opens the configuration screen shown in Figure 13.10.



The screenshot shows the 'Edit' dialog box for a template. The dialog has a title bar 'Edit' with a close button. It contains several fields: 'Name' (text input with 'customfreebsd'), 'OS' (dropdown menu with 'FreeBSD'), 'Architecture' (dropdown menu with 'x64'), 'URL' (text input with 'http://myurl.com'), and 'Mtree' (text input with an information icon). At the bottom are 'OK' and 'Cancel' buttons. The background shows the 'Templates' tab with the 'customfreebsd' template selected.

Fig. 13.10: Editing Template Options

Clicking a template's Delete button shows a warning message that prompts for confirmation of the deletion. Note that once

a template is deleted, it is removed from the Templates drop-down menu and will no longer be available for creating new jails.

13.4 Using iocage

Beginning with FreeNAS® 9.10.1, the **iocage** (<https://github.com/iocage/iocage>) command line utility is included for creating and managing jails. Click the Shell option to open the command line and begin using iocage.

Note: The jails infrastructure is transitioning from the old warden backend to the new iocage backend. This transition process requires the middleware API calls to be rewritten for the new UI. It is expected that the transition will be complete with FreeNAS® version 11.2. Since jails created in the old UI use the warden backend, jails created in the new UI use the iocage backend, and both use different API versions, they are not compatible. While a migration script will be made available when the transition is complete, it will not be able to anticipate every configuration scenario for every application installed in jails. At that time, the recommendation will be to: create new jails using the new UI, copy over any existing configurations, and delete the old jail datasets once the new jails are working as expected.

iocage has several options to help users:

- There is built-in help displayed by entering `iocage --help | more`. Each subcommand also has help, displayed by giving the subcommand name followed by the `--help` flag. For example, help for the `activate` subcommand displays with `iocage activate --help`.
- The iocage manual page is accessed by typing `man iocage`.
- The iocage project also has documentation available on [readthedocs.io](http://iocage.readthedocs.io/en/latest/index.html) (<http://iocage.readthedocs.io/en/latest/index.html>).

13.4.1 Managing iocage Jails

Creating a jail automatically starts the iocage configuration process for the FreeNAS® system. Jail properties can also be specified with the `iocage create` command.

In this example a new jail named `examplejail` is created. Additional properties are a manually designated IP address of 192.168.1.10, a netmask of /24 on the `em0` interface, and using the FreeBSD 11.1-RELEASE:

```
[root@freenas ~]# iocage create -n examplejail ip4_addr="em0|192.168.1.10/24" -r
11.1-RELEASE
...
examplejail successfully created!
```

Jail creation may take a few moments. After completion, start the new jail with `iocage start`:

```
[root@freenas ~]# iocage start examplejail
* Starting examplejail
+ Started OK
+ Starting services OK
```

To open the console in the started jail, use `iocage console`

```
[root@freenas ~]# iocage console examplejail
FreeBSD 11.1-STABLE (FreeNAS.amd64) #0 35e0ef284(freenas/11-stable): Wed Oct 18
17:44:36 UTC 2017

Welcome to FreeBSD!

Release Notes, Errata: https://www.FreeBSD.org/releases/
Security Advisories:  https://www.FreeBSD.org/security/
FreeBSD Handbook:     https://www.FreeBSD.org/handbook/
```

```
FreeBSD FAQ:           https://www.FreeBSD.org/faq/
Questions List: https://lists.FreeBSD.org/mailman/listinfo/freebsd-questions/
FreeBSD Forums:       https://forums.FreeBSD.org/
```

Documents installed with the system are in the `/usr/local/share/doc/freebsd/` directory, or can be installed later with: `pkg install en-freebsd-doc`
For other languages, replace "en" with a language code like `de` or `fr`.

```
Show the version of FreeBSD installed:  freebsd-version ; uname -a
Please include that output and any error messages when posting questions.
Introduction to manual pages:  man man
FreeBSD directory layout:      man hier
```

```
Edit /etc/motd to change this login announcement.
root@examplejail:~ #
```

Jails can be shut down with `iocage stop`:

```
[root@freenas ~]# iocage stop examplejail
* Stopping examplejail
  + Running prestop OK
  + Stopping services OK
  + Removing jail process OK
  + Running poststop OK
```

Jails are deleted with `iocage destroy`:

```
[root@freenas ~]# iocage stop examplejail
* Stopping examplejail2
  + Running prestop OK
  + Stopping services OK
  + Removing jail process OK
  + Running poststop OK
```

To adjust the properties of a jail, use `iocage set` and `iocage get`. All properties of a jail are viewed with `iocage get all`:

Tip: This example shows an abbreviated list of `examplejail`'s properties. The `iocage` manual page (`man iocage`) describes even more configurable properties for jails.

```
[root@freenas ~]# iocage get all examplejail | less
allow_mount:0
allow_mount_devfs:0
allow_sysvipc:0
available:readonly
basejail:no
boot:off
bpf:no
children_max:0
cloned_release:11.1-RELEASE
comment:none
compression:lz4
compressratio:readonly
coredumpsize:off
count:1
cpuset:off
cputime:off
datasize:off
dedup:off
```

```
defaultrouter:none  
defaultrouter6:none  
...
```

To adjust a jail property, use `iocage set`:

```
[root@freenas ~]# iocage set notes="This is a testing jail." examplejail  
Property: notes has been updated to This is a testing jail.
```

VIRTUAL MACHINES

A Virtual Machine (VM) is an environment on a host computer that can be used as if it were a separate physical computer. VMs can be used to run multiple operating systems simultaneously on a single computer. Operating systems running inside a VM see emulated virtual hardware rather than the actual hardware of the host computer. This provides more isolation than [jails](#) (page 251), although there is additional overhead. A portion of system RAM is assigned to each VM, and each VM uses a [zvol](#) (page 125) for storage. While a VM is running, these resources are not available to the host computer or other VMs.

FreeNAS® VMs use the [bhyve\(8\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=bhyve&manpath=FreeBSD+11.0-RELEASE+and+Ports>) virtual machine software. This type of virtualization requires an Intel processor with Extended Page Tables (EPT) or an AMD processor with Rapid Virtualization Indexing (RVI) or Nested Page Tables (NPT).

To verify that an Intel processor has the required features, use [Shell](#) (page 290) to run `grep VT-x /var/run/dmesg.boot`. If the EPT and UG features are shown, this processor can be used with bhyve.

To verify that an AMD processor has the required features, use [Shell](#) (page 290) to run `grep POPCNT /var/run/dmesg.boot`. If the output shows the POPCNT feature, this processor can be used with bhyve.

Note: By default, new VMs have the [bhyve\(8\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=bhyve&manpath=FreeBSD+11.0-RELEASE+and+Ports>) `-H` option is set. This causes the virtual CPU thread to yield when a HLT instruction is detected, and prevents idle VMs from consuming all of the host's CPU.

Note: AMD K10 “Kuma” processors include POPCNT but do not support NRIPS, which is required for use with bhyve. Production of these processors ceased in 2012 or 2013.

14.1 Creating VMs

Select VMs → Add VM for the Add VM dialog shown in [Figure 14.1](#):

Fig. 14.1: Add VM

VM configuration options are described in [Table 14.1](#).

Table 14.1: VM Options

Setting	Value	Description
VM Type	drop-down menu	type of VM, either Virtual Machine for a typical instance, or Docker VM for a special VM to run Docker
Name	string	a name to identify the VM
Description	string	a short description of the VM or its purpose
Virtual CPUs	integer	quantity of virtual CPUs allocated to the VM, up to 16; although these are virtual and not strictly related to host processor cores, the host CPU might limit the maximum number; the operating system used in the VM might also have operational or licensing restrictions on the number of CPUs allowed
Memory Size (MiB)	integer	megabytes of RAM allocated to the VM
Boot Method	drop-down menu	UEFI for newer operating systems, or UEFI-CSM (Compatibility Support Mode) for older operating systems that only understand BIOS booting
Autostart	checkbox	when checked, start the VM automatically on boot

14.2 Adding Devices to a VM

After creating the VM, click it to select it, then click Devices and Add Device to add virtual hardware to it:



Fig. 14.2: Add Devices to a VM

Select the name of the VM from the VM drop-down menu, then select the Type of device to add. These types are available:

- [Network Interfaces](#) (page 272)
- [Disk Devices](#) (page 273)
- [Raw Files](#) (page 273)
- [CD-ROMs](#) (page 274)
- [VNC Interface](#) (page 274)

Figure 14.3 shows the fields that appear when Network Interface is the selected Type.

14.2.1 Network Interfaces




Fig. 14.3: VM Network Interface Device

The default Adapter Type emulates an Intel E1000 (82545) Ethernet card for compatibility with most operating systems. This can be changed to VirtIO to provide better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.

If the system has multiple physical network interface cards, the Nic to attach drop-down menu can be used to specify which physical interface to associate with the VM.

By default, the VM receives an auto-generated random MAC address. To override the default with a custom value, enter the desired address into the MAC Address field.

14.2.2 Disk Devices

Zvols (page 125) are typically used as virtual hard drives. After [creating a zvol](#) (page 125), associate it with the VM by selecting Add device, choose the VM, select a Type of Disk, select the created zvol, then set the Mode. If a specific sector size is required, enter the number of bytes into Disk sectorsize. The default of 0 leaves the sector size unset.




Fig. 14.4: VM Disk Device

AHCI emulates an AHCI hard disk for best software compatibility. VirtIO uses paravirtualized drivers and can provide better performance, but requires the operating system installed in the VM to support VirtIO disk devices.

14.2.3 Raw Files

Raw Files are similar to [Zvol](#) (page 125) disk devices, but the disk image comes from a file. These are typically used with existing read-only binary images of drives, like an installer disk image file meant to be copied onto a USB stick.

After obtaining and copying the image file to the FreeNAS® system, select Add device, choose the VM, select a Type of Raw File, browse to the image file, then set the Mode. AHCI emulates an AHCI hard disk for best software compatibility. VirtIO uses paravirtualized drivers and can provide better performance, but requires the operating system installed in the VM to support VirtIO disk devices.

If a specific sector size is required, enter the number of bytes into Disk sectorsize. The default of 0 leaves the sector size unset.

The 'Add device' dialog box contains the following fields and controls:

- VM:** A dropdown menu with 'samplevm' selected.
- Type:** A dropdown menu with 'Raw File' selected.
- Mode:** A dropdown menu with 'AHCI' selected.
- Raw File:** A text input field containing '/mnt/volume1/FreeBSD-11.1.' and a 'Browse' button to its right.
- Disk sectorsize:** A text input field containing '0' and an information icon (i) to its right.
- At the bottom are 'OK' and 'Cancel' buttons.

Fig. 14.5: VM Raw File Disk Device

14.2.4 CD-ROM Devices

Adding a CD-ROM device makes it possible to boot the VM from a CD-ROM image, typically an installation CD. The image must be present on an accessible portion of the FreeNAS® storage. In this example, a FreeBSD installation image is shown:

The 'Edit' dialog box contains the following fields and controls:

- VM:** A dropdown menu with 'samplevm' selected.
- Type:** A dropdown menu with 'CD-ROM' selected.
- CD-ROM (ISO):** A text input field containing '/mnt/volume1/isos/FreeBSD-' and a 'Browse' button to its right.
- At the bottom are 'OK', 'Cancel', and 'Delete' buttons.

Fig. 14.6: VM CD-ROM Device

Note: VMs from other virtual machine systems can be recreated for use in FreeNAS®. Back up the original VM, then create a new FreeNAS® VM with virtual hardware as close as possible to the original VM. Binary-copy the disk image data into the `zvol` (page 125) created for the FreeNAS® VM with a tool that operates at the level of disk blocks, like `dd(1)` (<https://www.freebsd.org/cgi/man.cgi?query=dd>). For some VM systems, it is best to back up data, install the operating system from scratch in a new FreeNAS® VM, and restore the data into the new VM.

14.2.5 VNC Interface

VMs set to UEFI booting are also given a VNC (Virtual Network Computing) remote connection. A standard VNC (https://en.wikipedia.org/wiki/Virtual_Network_Computing) client can connect to the VM to provide screen output and key-

board and mouse input.

Figure 14.7 shows the fields that appear when VNC is the selected Type.

Fig. 14.7: VM VNC Device

The Resolution drop-down menu can be used to modify the default screen resolution used by the VNC session.

The VNC port can be set to 0, left empty for FreeNAS® to assign a port when the VM is started, or set to a fixed, preferred port number.

Select the IP address for VNC to listen on with the Bind to drop-down menu.

Check the Wait to boot checkbox to indicate that the VNC client should wait until the VM has booted before attempting the connection.

To automatically pass the VNC password, enter it into the Password field. Note that the password is limited to 8 characters.

To use the VNC web interface, check the VNC Web checkbox.

Tip: If a RealVNC 5.X Client shows the error `RFB protocol error: invalid message type`, disable the Adapt to network speed option and move the slider to Best quality. On later versions of RealVNC, select `File → Preferences`, click Expert, ProtocolVersion, then select 4.1 from the drop-down menu.

14.2.6 Virtual Serial Ports

VMs automatically include a virtual serial port.

- `/dev/nmdm1B` is assigned to the first VM
- `/dev/nmdm2B` is assigned to the second VM


And so on. These virtual serial ports allow connecting to the VM console from the [Shell](#) (page 290). To connect to the first VM:

```
cu -s 9600 -l /dev/nmdm1B
```

See [cu\(1\)](https://www.freebsd.org/cgi/man.cgi?query=cu) (<https://www.freebsd.org/cgi/man.cgi?query=cu>) for more information on operating `cu`.

14.3 Running VMs

Select **VMs** to see a list of configured VMs. Configuration and control buttons appear at the bottom of the screen when an individual VM is selected with a mouse click:



The screenshot shows a web interface for managing VMs. At the top, there are three buttons: 'VMs', 'VMs', and 'Add VM'. Below these is a table with the following data:

Name	Description	Info	Virtual CPUs	Memory Size (MiB)	Boot Method	Autostart
samplevm	A sample virtual machine	State: STOPPED VNC Port: 5901	1	2048	UEFI	false

Fig. 14.8: VM Configuration and Control Buttons

The name, description, running state, VNC port (if present), and other configuration values are shown. Click on an individual VM for additional options.

Some standard buttons are shown for all VMs:

- Edit changes VM settings.
- Delete [removes the VM](#) (page 276).
- Devices is used to add and remove devices to this VM.

When a VM is not running, these buttons are shown:

- Start starts the VM.
- Clone clones or copies the VM to a new VM. The new VM is given the same name as the original, with `_cloneN` appended.

When a VM is already running, these buttons are shown:

- Stop shuts down the VM.
- Power off immediately halts the VM, equivalent to disconnecting the power on a physical computer.
- Restart restarts the VM.
- Vnc via Web starts a web VNC connection to the VM. The VM must have a VNC device, and VNC Web enabled in that device.

14.4 Deleting VMs

When a VM is no longer needed, it can be deleted by clicking on the VM, then clicking Delete at the bottom of the screen. A dialog will show any related devices that will also be deleted and ask for confirmation.

Tip: [Zvols](#) (page 125) used in [disk devices](#) (page 273) and image files used in [raw file](#) (page 273) devices are not removed when a VM is deleted. These resources can be removed manually after it is determined that the data in them has been backed up or is no longer needed.

14.5 Docker/Rancher VM

Docker (<https://www.docker.com/what-docker>) is open source software for automating application deployment inside containers. A container provides a complete filesystem, runtime, system tools, and system libraries, so applications always see the same environment.

Rancher (<http://rancher.com/>) is a GUI tool for managing Docker containers.

FreeNAS® runs the Rancher GUI as a separate VM.

14.5.1 Rancher VM Requirements

20 GiB of storage space is required for the Rancher VM. For setup, the [SSH](#) (page 236) service must be enabled.

The Rancher VM requires 2 GiB of RAM while running.


14.5.2 Create the Rancher VM

Click VMs, then the Add VM button. Set the VM Type to Docker VM. Enter RancherUI for the name, Rancher UI VM for the Description, leave the number of Virtual CPUs at 1, and enter 2048 for the Memory Size. To have the Rancher VM start when the FreeNAS® system boots, check the Autostart checkbox. Click OK to create the virtual machine.

Fig. 14.9: Rancher VM Configuration

A location to store the disk image must now be chosen. In this example, a [dataset](#) (page 122) called vm-storage has already been created as a location to store VM data. Click VMs, then click on the RancherUI line to select it. Click on the Devices button to show the devices attached to that VM. Click on the RAW device to select it, then click the Edit button. In the Raw File field, browse to the dataset and select it. Then add a filename by typing /rancherui.img at the end of the path in the text box.

Set the Disk boot checkbox, enter a password for the `rancher` user in the Password field, then enter 20G in the Disk size field. Click OK to save the device.



The screenshot shows the 'Edit' dialog box for configuring a virtual machine. The dialog has a title bar with 'Edit' and a close button. The configuration fields are as follows:

- VM:** A dropdown menu showing 'RancherUI'.
- Type:** A dropdown menu showing 'Raw File'.
- Mode:** A dropdown menu showing 'AHCI'.
- Raw File:** A text field containing 'k/vm-storage/rancherui.img'. To the right of this field is a 'Close' button. Below the text field is a file browser showing a directory structure: a root directory with subdirectories 'mnt' and 'tank', and a 'vm-storage' directory selected at the bottom.
- Disk boot:** A checkbox that is checked.
- Password:** A text field with 10 black dots, indicating a masked password. To the right of the field is an information icon.
- Disk sectorsize:** A text field containing '0'. To the right of the field is an information icon.
- Disk size:** A text field containing '20G'. To the right of the field is an information icon.

At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Delete'.

Fig. 14.10: Rancher Image Storage

14.5.3 Start the Rancher VM

Click VMs, then click on the RancherUI line to select it. Click the Start button and then Yes to start the VM.

The first time the Rancher VM is started, it downloads the Rancher disk image file. How long this takes to complete depends on the speed of the network connection. A status dialog reports the progress of the download.

After the image is downloaded, the VM is started.

14.5.4 Installing the Rancher Server

Click VMs and locate the line for the RancherUI VM. The Info column shows the `Com Port` for the Rancher VM. In this example, `/dev/nmdm3B` is used.

Further setup of the Rancher VM is done from the command line. Use an SSH client to connect to the FreeNAS® server. Remember that this requires the [SSH](#) (page 236) service to be running. Depending on local configuration, it might also require changes to the setting of the service, like allowing root user login with a password.

At the FreeNAS® console prompt, connect to the Rancher VM with `cu` (<https://www.freebsd.org/cgi/man.cgi?query=cu>), replacing `/dev/nmdm3B` with the value from the RancherUI Info column:

```
cu -l /dev/nmdm3B
```

If the terminal does not show a `rancher login:` prompt, press `Enter`.

Enter `rancher` as the username, press `Enter`, then type the password that was entered when the raw file was created above and press `Enter` again. After logging in, a `[rancher@rancher ~]$` prompt is displayed.

Download and install the Rancher system with this command:

```
sudo docker run -d --restart=unless-stopped -p 8080:8080 rancher/server
```

Note: If the error `Cannot connect to the Docker daemon` is shown, run `sudo dockerd`. Then give the `sudo docker run` command above again.

Installation time varies with processor and network connection speed, but typically takes a few minutes. After the process finishes and a command prompt is shown, type this command:

```
ifconfig eth0 | grep 'inet addr'
```

The first value is the IP address of the Rancher server. Enter the IP address and port 8080 as the URL in a web browser. For example, if the IP address was `10.231.3.208`, enter `10.231.3.208:8080` as the URL in the web browser.

The Rancher server takes a few minutes to start. The web browser might show a connection error while the Rancher GUI is still starting. If the browser shows a `connection has timed out` or a similar error, wait one minute and try again.

In the Rancher GUI, click `Add a host` and enter the same IP address and port number. Click `Save` to save the information.

For more information on using Rancher, see the Rancher [Quick Start Guide](https://rancher.com/docs/rancher/v1.6/en/quick-start-guide/) (<https://rancher.com/docs/rancher/v1.6/en/quick-start-guide/>).

REPORTING

Reporting displays several graphs, as seen in the example in Figure 15.1. Click the tab for a device type to see its graphs.




Fig. 15.1: Reporting Graphs

FreeNAS® uses [collectd](https://collectd.org/) (<https://collectd.org/>) to provide reporting statistics. The resulting graphs are grouped into several tabs on the Reporting page:

- CPU
 - CPU (<https://collectd.org/wiki/index.php/Plugin:CPU>) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle.
- Disk
 - Disk (<https://collectd.org/wiki/index.php/Plugin:Disk>) shows statistics on I/O, percent busy, latency, operations per second, and pending I/O requests.
- Memory
 - Memory (<https://collectd.org/wiki/index.php/Plugin:Memory>) displays memory usage.
 - Swap (<https://collectd.org/wiki/index.php/Plugin:Swap>) displays the amount of free and used swap space.

- Network
 - [Interface](https://collectd.org/wiki/index.php/Plugin:Interface) (<https://collectd.org/wiki/index.php/Plugin:Interface>) shows received and transmitted traffic in bits per second for each configured interface.
- Partition
 - [Disk space](https://collectd.org/wiki/index.php/Plugin:DF) (<https://collectd.org/wiki/index.php/Plugin:DF>) displays free and used space for each volume and dataset. However, the disk space used by an individual zvol is not displayed as it is a block device.
- System
 - [Processes and Uptime](https://collectd.org/wiki/index.php/Plugin:Processes) (<https://collectd.org/wiki/index.php/Plugin:Processes>) displays the number of processes, grouped by state.
 - [Uptime](https://collectd.org/wiki/index.php/Plugin:Uptime) (<https://collectd.org/wiki/index.php/Plugin:Uptime>) keeps track of the system uptime, the average running time, and the maximum reached uptime.
- Target
 - Target shows bandwidth statistics for iSCSI ports.
- ZFS
 - [ZFS](https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) (https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) shows ARC size, hit ratio, and requests.

Reporting data is saved to permit viewing and monitoring usage trends over time. This data is preserved across system upgrades and restarts.

Data files are saved in `/var/db/collectd/rrd/`.

The reporting data file recording method is controlled by the `System → System Dataset Reporting database` checkbox. When unchecked, data files are recorded in a temporary filesystem and copied hourly to on-disk files.

When `System → System Dataset Reporting database` is checked, data files are written directly to the [System Dataset](#) (page 67).

Warning: Reporting data is frequently written and should not be stored on the boot pool or boot device.

Use the magnifier buttons next to each graph to increase or decrease the displayed time increment from 10 minutes, hourly, daily, weekly, or monthly. The << and >> buttons can be used to scroll through the output.

[Update on using Graphite with FreeNAS](http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html) (<http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html>) contains instructions for sending the collected information to a [Graphite](http://graphite.wikidot.com/) (<http://graphite.wikidot.com/>) server.

WIZARD

FreeNAS® provides a wizard which helps complete the steps needed to quickly configure FreeNAS® for serving data over a network. The wizard can be run at any time by clicking the Wizard icon.

Figure 16.1 shows the first wizard configuration screen.



Fig. 16.1: Configuration Wizard

Note: You can exit the wizard at any time by clicking the Exit button. However, exiting the wizard will not save any selections. The wizard can always be run again by clicking the Wizard icon. Alternately, the FreeNAS® GUI can be used to configure the system, as described in the rest of this Guide.

This screen can be used to change the default language, keyboard map, and timezone. After making your selections, click Next. The next screen depends on whether or not the storage disks have already been formatted into a ZFS pool.

Figure 16.2 shows the configuration screen that appears if the storage disks have not yet been formatted.

Fig. 16.2: Volume Creation Wizard

Note: The wizard will not recognize an encrypted ZFS pool. If your ZFS pool is GELI-encrypted, cancel the wizard and use the instructions in [Importing an Encrypted Pool](#) (page 128) to import the encrypted volume. You can then rerun the wizard afterwards, if you wish to use it for post-configuration, and it will recognize that the volume has been imported and will not prompt to reformat the disks.

Enter a name for the ZFS pool that conforms to these [naming conventions](http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html) (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html). It is recommended to choose a name that will stick out in the logs (e.g. not `data` or `freenas`).

Decide if the pool should provide disk redundancy, and if so, which type. The [ZFS Primer](#) (page 322) discusses RAIDZ redundancy in more detail. If you prefer to make a more complex configuration, click the Exit button to close the wizard and instead use [Volume Manager](#) (page 115).

These redundancy types are available:

- Automatic: automatically creates a mirrored, RAIDZ1, or RAIDZ2 pool, depending upon the number of disks. If you prefer to control the type of redundancy, select one of the other options.
- RAID 10: creates a striped mirror and requires a minimum of 4 disks.
- RAIDZ2: requires a minimum of 4 disks. Up to 2 disks can fail without data loss.
- RAIDZ1: requires a minimum of 3 disks. Up to 1 disk can fail without data loss.
- Stripe: requires a minimum of 1 disk. Provides no redundancy, meaning if any of the disks in the stripe fails, all data in the stripe is lost.

Once you have made your selection, click Next to continue.

If the disks have already been formatted with ZFS and the disks have not been encrypted, the next screen will instead prompt to import the volume, as shown in [Figure 16.3](#).



Fig. 16.3: Volume Import Screen

Select the existing volume from the drop-down menu and click Next to continue.
The next screen in the wizard is shown in [Figure 16.4](#).

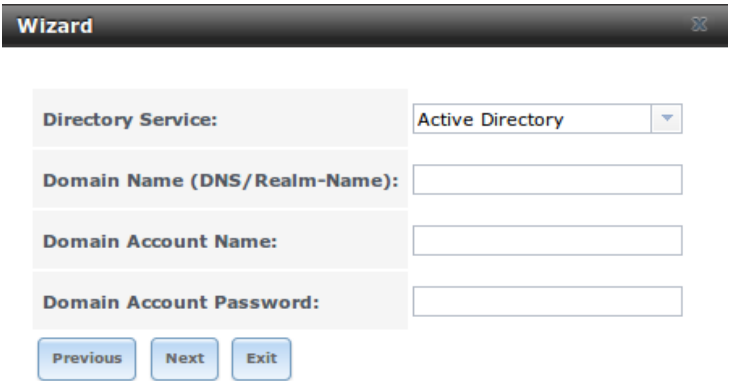


Fig. 16.4: Directory Service Selection

If the FreeNAS® system is on a network that does not contain an Active Directory, LDAP, or NIS server, click Next to skip to the next screen.

However, if the FreeNAS® system is on a network containing an Active Directory, LDAP, or NIS server and you wish to import the users and groups from that server, select the type of directory service in the Directory Service drop-down menu. The rest of the fields in this screen will vary, depending upon which directory service is selected. Available configuration options for each directory service are summarized in [Tables 16.1 through 16.3](#).

Note: Additional configuration options are available for each directory service. The wizard can be used to set the initial values required to connect to that directory service. You can then review the other available options in [Directory Services](#) (page 158) to determine if additional configuration is required.

Table 16.1: Active Directory Options

Setting	Value	Description
Domain Name	string	name of Active Directory domain (e.g. example.com) or child domain (e.g. sales.example.com)
Domain Account Name	string	name of the Active Directory administrator account
Domain Account Password	string	password for the Active Directory administrator account

Table 16.2: LDAP Options

Setting	Value	Description
Hostname	string	hostname or IP address of LDAP server
Continued on next page		

Table 16.2 – continued from previous page

Setting	Value	Description
Base DN	string	top level of the LDAP directory tree to be used when searching for resources (e.g. dc=test,dc=org)
Bind DN	string	name of administrative account on LDAP server (e.g. cn=Manager,dc=test,dc=org)
Base password	string	password for

Table 16.3: NIS Options

Setting	Value	Description
NIS domain	string	name of NIS domain
NIS servers	string	comma delimited list of hostnames or IP addresses
Secure mode	checkbox	if checked, ypbind(8) (http://www.freebsd.org/cgi/man.cgi?query=ybind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, ybind will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet

The next configuration screen, shown in [Figure 16.5](#), is used to create network shares.

Fig. 16.5: Network Shares

FreeNAS® supports several types of shares for providing storage data to the clients in a network. The initial wizard can be

used to quickly make shares using default permissions which should “just work” for common scenarios. For more complex scenarios, refer to the section on [Sharing](#) (page 169).

To create a share using the wizard, enter a name for the share, then select the Purpose of the share:

- **Windows (SMB):** this type of share can be accessed by any operating system using a SMB client. Check the box for Allow Guest to allow users to access the share without a password. SMB shares created with the wizard can be fine-tuned afterward with [Windows \(SMB\) Shares](#) (page 185).
- **Mac OS X (AFP):** this type of share can be accessed by Mac OS X users. Check the box for Time Machine if Mac users will be using the FreeNAS® system as a backup device. AFP shares created with the wizard can be fine-tuned afterward with [Apple \(AFP\) Shares](#) (page 170).
- **Generic Unix (NFS):** this type of share can be accessed by any operating system using a NFS client. NFS shares created using the wizard can be fine-tuned afterward with [Unix \(NFS\) Shares](#) (page 177).
- **Block Storage (iSCSI):** this type of share can be accessed by any operating system using iSCSI initiator software. Enter the size of the block storage to create in the format 20G (for 20 GB). iSCSI shares created with the wizard can be fine-tuned afterward with [iSCSI](#) (page 221).

After selecting the Purpose, click the Ownership button to see the screen shown in [Figure 16.6](#).

	Owner	Group	Other
Read	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Write	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Execute	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Fig. 16.6: Share Permissions

The default permissions for the share are displayed. To create a user or group, enter the desired name, then check the Create User box to create that user and the Create Group box to create the group. Check or uncheck the boxes in the Mode section to set the initial access permissions for the share. When finished, click the Return button to return to the share creation screen. Click the Add button to finish creating that share, which will then appear in the Name frame.

The Delete button can be used to remove the share highlighted in the Name frame. To edit a share, highlight it, make the change, then press the Update button.

When finished making shares, click the Next button to advance to the screen shown in [Figure 16.7](#).

Fig. 16.7: Miscellaneous Settings

This screen can be used to configure these settings:

- **Console messages:** check this box if you would like to view system messages at the bottom of the graphical administrative interface. This can be handy when troubleshooting a service that will not start. When using the console message view, if you click the console messages area, it will pop-up as a window, allowing you to scroll through the output and to copy its contents.
- **Root E-mail:** FreeNAS® provides an “Alert” icon in the upper right corner to provide a visual indication of events that warrant administrative attention. The alert system automatically emails the root user account whenever an alert is issued. It is important to enter the email address of the person to receive these alerts and other administrative emails. The rest of the email settings in this screen should also be reviewed and edited as necessary. Before leaving this screen, click the “Send Test Mail” button to ensure that email notifications are working correctly.
- **From email:** the from email address to use when sending email notifications.
- **Outgoing mail server:** hostname or IP address of SMTP server.
- **Port to connect to:** port number used by the SMTP server.
- **TLS/SSL:** encryption type used by the SMTP server.
- **Use SMTP Authentication:** check this box if the SMTP server requires authentication.
- **Username:** enter the username if the SMTP server requires authentication.
- **Password:** enter the password if the SMTP server requires authentication.

When finished, click Next. A message will indicate that the wizard is ready to perform all of the saved actions. To make changes, click the Return to Wizard button to review your edits. If you click the Exit without saving button, none of your selections will be saved. To save your edits, click the Confirm button. A status bar will indicate when the wizard has completed applying the new settings.

In addition to the settings that you specify, the wizard will automatically enable [S.M.A.R.T. Tests](#) (page 99), create a boot environment, and add the new boot environment to the boot menu. If you also wish to save a backup of the configuration database to the system being used to access the administrative graphical interface, go to *System* → *General*, click the


Save Config button, and browse to the directory where the configuration will be saved. Always back up your configuration after making any configuration changes.

The rest of this Guide describes the FreeNAS® graphical interface in more detail. The layout of this Guide follows the order of the menu items in the tree located in the left frame of the graphical interface.

Note: It is important to use the GUI (or the Console Setup menu) for all configuration changes. FreeNAS® uses a configuration database to store its settings. While it is possible to use the command line to modify your configuration, changes made at the command line are not written to the configuration database. This means that any changes made at the command line will not persist after a reboot and will be overwritten by the values in the configuration database during an upgrade.

DISPLAY SYSTEM PROCESSES

Clicking Display System Processes opens a screen showing the output of `top(1)` (<http://www.freebsd.org/cgi/man.cgi?query=top>). An example is shown in Figure 17.1.



```

Running Processes

last pid: 4533; load averages: 0.04, 0.04, 0.00 up 0+01:17:36 06:26:29
21 processes: 1 running, 20 sleeping

Mem: 103M Active, 118M Inact, 224M Wired, 3220K Cache, 152M Buf, 7375M Free
ARC: 2543K Total, 1052K MFU, 1126K MRU, 16K Anon, 90K Header, 258K Other
Swap: 8192M Total, 8192M Free

  PID USERNAME   THR PRI NICE   SIZE    RES STATE  C  TIME  WCPU COMMAND
 2014 root         6  20   0   382M    138M usem   3   0:07  0.00% python2.7
 2586 root         1  52   0   147M   51312K ttyin  3   0:01  0.00% python2.7
 3942 root         7  20   0    122M   13920K uwait   3   0:00  0.00% collectd
 1742 root         1  20   0  22216K   3852K select  2   0:00  0.00% ntpd
 3387 www          1  20   0  26028K   5540K kqread  1   0:00  0.00% nginx
 1557 root         1  20   0  12044K   1724K select  2   0:00  0.00% syslogd
 2200 root         1  52   0  14128K   1808K nanslp  1   0:00  0.00% cron
 2442 root         1  20   0  14128K   1852K select  0   0:00  0.00% rpcbind
 1290 root         1  20   0  10376K   4400K select  3   0:00  0.00% devd
 2088 root         1  20   0  26028K   5028K pause   2   0:00  0.00% nginx
 4533 root         1  20   0  16556K   2184K CPU3    3   0:00  0.00% top
 2591 root         1  52   0  12044K   1620K ttyin   3   0:00  0.00% getty
 2446 root         1  23   0  12040K   1912K select  0   0:00  0.00% mountd
 2587 root         1  52   0  12044K   1620K ttyin   0   0:00  0.00% getty
 2589 root         1  52   0  12044K   1620K ttyin   0   0:00  0.00% getty
 2593 root         1  52   0  12044K   1620K ttyin   0   0:00  0.00% getty
 2588 root         1  52   0  12044K   1620K ttyin   3   0:00  0.00% getty
 2592 root         1  52   0  12044K   1620K ttyin   2   0:00  0.00% getty
  
```

Fig. 17.1: System Processes Running on FreeNAS®

The display will automatically refresh itself. Click the X in the upper right corner to close the display. Note that the display is read-only, meaning that you will not be able to issue a `kill` command within it.

SHELL

Beginning with version 8.2.0, the FreeNAS® GUI provides a web shell, making it convenient to run command line tools from the web browser as the root user. The link to Shell is the fourth entry from the bottom of the menu tree. In [Figure 18.1](#), the link has been clicked and Shell is open.




Fig. 18.1: Web Shell

The prompt indicates that the current user is root, the hostname is freenas, and the current working directory is ~ (root's home directory).

To change the size of the shell, click the 80x25 drop-down menu and select a different size.

To copy text from shell, highlight the text, right-click, and select Copy from the right-click menu. To paste into the shell, click the Paste button, paste the text into the box that opens, and click the OK button to complete the paste operation.

Shell provides history (use your up arrow to see previously entered commands and press `Enter` to repeat the currently displayed command) and tab completion (type a few letters and press `tab` to complete a command name or filename in the

current directory). When you are finished using Shell, type `exit` to leave the session.

While you are in Shell, you will not have access to any of the other GUI menus. If you need to have access to a prompt while using the GUI menus, use [tmux](#) (page 316) instead as it supports multiple shell sessions and the detachment and reattachment of sessions.

Note: Not all of Shell's features render correctly in Chrome. Firefox is the recommended browser for using Shell.

Most FreeBSD command line utilities are available in Shell. Additional troubleshooting utilities that are provided by FreeNAS® are described in [Command Line Utilities](#) (page 302).

LOG OUT

Click the Log Out entry in the FreeNAS® GUI to log out.

After logging out, a message appears with a link to log back in. When logging back in, the root password is required.

REBOOT

Clicking the Reboot entry in the tree shows the warning message in [Figure 20.1](#). The browser screen color changes to red to indicate that this option will negatively impact current users of the FreeNAS® system.

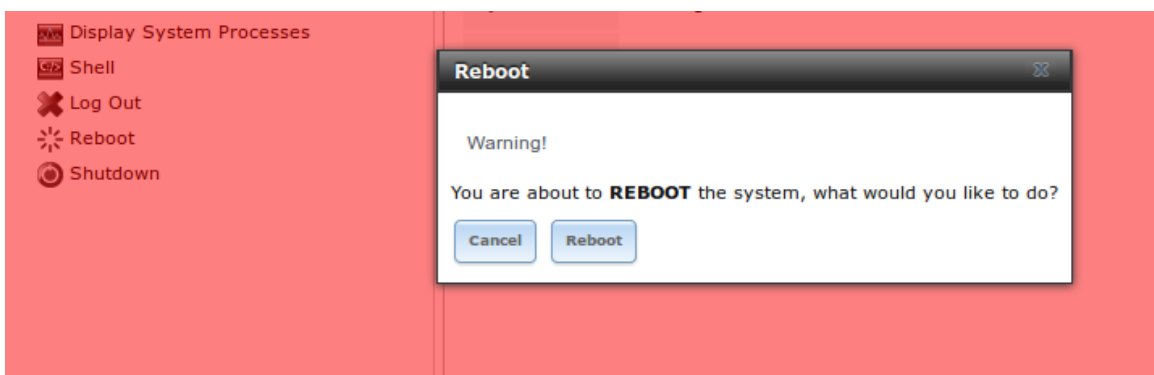


Fig. 20.1: Reboot Warning Message

If a scrub or resilver is in progress when a reboot is requested, an additional warning will ask you to make sure that you wish to proceed. In this case, it is recommended to Cancel the reboot request and to periodically run `zpool status` from Shell until it is verified that the scrub or resilver process is complete. Once complete, the reboot request can be re-issued.

Click the Cancel button to cancel the reboot request. Otherwise, click the Reboot button to reboot the system. Rebooting the system disconnects all clients, including the web administration GUI. The URL in the web browser changes to add `/system/reboot/` to the end of the IP address. Wait a few minutes for the system to boot, then use your browser's "back" button to return to the FreeNAS® system's IP address. If all went well, the GUI login screen will appear. If the login screen does not appear, physical access to the FreeNAS® system's monitor and keyboard is needed to determine what problem is preventing the system from resuming normal operation.

SHUTDOWN

Clicking the Shutdown entry in the tree opens the warning message shown in [Figure 21.1](#). The browser window color changes to red to indicate that this command will negatively impact current users of the FreeNAS® system.



Fig. 21.1: Shutdown Warning Message

If a scrub or resilver is in progress when a shutdown is requested, an additional warning will ask if you wish to proceed. In this case, it is recommended to Cancel the shutdown request and to periodically run `zpool status` from [Shell](#) (page 290) until the scrub or resilver process is complete. Once complete, the shutdown request can be re-issued.

Click the Cancel button to cancel the shutdown request. Otherwise, click the Shutdown button to halt the system. Shutting down the system disconnects all clients, including the web administration GUI, and powers off the FreeNAS® system. Physical access to the FreeNAS® system will be needed to turn it back on.

SUPPORT ICON

The Support icon, the third icon from the left in the top menubar, provides a shortcut to *System* → *Support*. This screen can be used to create a support ticket. Refer to [Support](#) (page 81) for detailed usage instructions.

GUIDE

The Documentation icon, the second icon from the left in the top menubar, provides a built-in browser to the FreeNAS® User Guide (this documentation).

ALERT

FreeNAS® provides an alert system to provide a visual warning of any conditions that require administrative attention. The Alert button in the far right corner flashes red when there is an outstanding alert. In the example alert shown in [Figure 24.1](#), the system is warning that the S.M.A.R.T. service is not running.



Fig. 24.1: Example Alert Message

Informational messages have a green OK, warning messages flash yellow, and messages requiring attention are listed as a red CRITICAL. CRITICAL messages are also emailed to the root user account. If you are aware of a critical condition but wish to remove the flashing alert until it can be dealt with, uncheck the box next to that message.

Behind the scenes, an alert daemon checks for various alert conditions, such as volume and disk status, and writes the current conditions to `/var/tmp/alert`. The daemon retrieves the current alert status every minute and changes the solid green alert icon to flashing red when a new alert is detected.

Current alerts can also be viewed from the Shell option of the Console Setup Menu ([Figure 3.1](#)) or from the Web Shell ([Figure 18.1](#)) by running `alertcli.py`.

Some of the conditions that trigger an alert include:

- used space on a volume, dataset, or zvol goes over 80%; the alert goes red at 95%
- new OpenZFS feature flags are available for the pool; this alert can be unchecked if a pool upgrade is not desired at present
- a new update is available
- the system reboots itself
- non-optimal multipath states are detected
- ZFS pool status changes from HEALTHY
- a S.M.A.R.T. error occurs
- the system dataset does not reside on the boot pool
- the system is unable to bind to the WebGUI IPv4 Address set in `System → General`
- the system can not find an IP address configured on an iSCSI portal
- a periodic snapshot or replication task fails
- a VMware login or a [VMware-Snapshot](#) (page 157) task fails

- deleting a VMware snapshot fails
- a Certificate Authority or certificate is invalid or malformed
- an update failed, or the system needs to reboot to complete a successful update
- a re-key operation fails on an encrypted pool
- LDAP failed to bind to the domain
- any member interfaces of a lagg interface are not active
- the status of an Avago MegaRAID SAS controller has changed; [mfiutil\(8\)](http://www.freebsd.org/cgi/man.cgi?query=mfiutil) (<http://www.freebsd.org/cgi/man.cgi?query=mfiutil>) is included for managing these devices

SUPPORT RESOURCES

FreeNAS® has a large installation base and an active user community. This means that many usage questions have already been answered and the details are available on the Internet. If you get stuck using FreeNAS®, spend a few moments searching the Internet for the word FreeNAS with some keywords that describe the error message or the function you are trying to implement.

The rest of this section discusses the following resources which are available to FreeNAS® users:

- [Website and Social Media](#) (page 299)
- [Forums](#) (page 299)
- [IRC](#) (page 300)
- [Videos](#) (page 300)
- [Professional Support](#) (page 301)

25.1 Website and Social Media

The [FreeNAS® website](http://www.freenas.org/) (<http://www.freenas.org/>) contains links to all of the available documentation, support, and social media resources. Major announcements are also posted to the main page.

Users are welcome to network on the FreeNAS® social media sites:

- [LinkedIn](http://www.linkedin.com/groups/FreeNAS8-3903140) (<http://www.linkedin.com/groups/FreeNAS8-3903140>)
- [Google+](https://plus.google.com/110373675402281849911/posts) (<https://plus.google.com/110373675402281849911/posts>)
- [Facebook FreeNAS Community](https://www.facebook.com/freenascommunity) (<https://www.facebook.com/freenascommunity>)
- [Facebook FreeNAS Consortium \(please request to be added\)](https://www.facebook.com/groups/1707686686200221) (<https://www.facebook.com/groups/1707686686200221>)
- [Twitter](https://twitter.com/freenasteam) (<https://twitter.com/freenasteam>)

25.2 Forums

The [FreeNAS Forums](https://forums.freenas.org/index.php) (<https://forums.freenas.org/index.php>) are an active online resource where people can ask questions, receive help, and share findings with other FreeNAS® users. New users are encouraged to post a brief message about themselves and how they use or are considering using FreeNAS® in the [Introductions](https://forums.freenas.org/index.php?forums/introductions.25/) (<https://forums.freenas.org/index.php?forums/introductions.25/>) forum.

The [Resources](https://forums.freenas.org/index.php?resources/) (<https://forums.freenas.org/index.php?resources/>) section contains categorized, user-contributed guides on many aspects of building and using FreeNAS® systems.

Language-specific categories are available under International, allowing FreeNAS® users to interact with each other in their native language:

- [Chinese](https://forums.freenas.org/index.php?forums/chinese-%E4%B8%AD%E6%96%87.60/) (<https://forums.freenas.org/index.php?forums/chinese-%E4%B8%AD%E6%96%87.60/>)

- Dutch - Nederlands (<https://forums.freenas.org/index.php?forums/dutch-nederlands.35/>)
- French - Français (<https://forums.freenas.org/index.php?forums/french-francais.29/>)
- German - Deutsch (<https://forums.freenas.org/index.php?forums/german-deutsch.31/>)
- Italian - Italiano (<https://forums.freenas.org/index.php?forums/italian-italiano.30/>)
- Portuguese - Português (<https://forums.freenas.org/index.php?forums/portuguese-portugu%C3%AAs.44/>)
- Romanian - Română (<https://forums.freenas.org/index.php?forums/romanian-rom%C3%A2n%C4%83.53/>)
- Russian - Русский (

To join the forums, create an account with the Sign Up Now! link.

Before asking a question on the forum, check the [Resources](https://forums.freenas.org/index.php?resources/) (https://forums.freenas.org/index.php?resources/) to see if the information is already there. Then see the [Forum Rules](https://forums.freenas.org/index.php?threads/updated-forum-rules-4-11-17.45124/) (https://forums.freenas.org/index.php?threads/updated-forum-rules-4-11-17.45124/) for guidelines on posting your hardware information and how to ask questions that will get a response.

25.3 IRC

To ask a question in real time, you can use the #freenas channel on IRC [Freenode](http://freenode.net/) (<http://freenode.net/>). Depending on the time of day and your time zone, FreeNAS® developers or other users may be available to provide assistance. If no one answers right away, remain on the channel, as other users tend to read the channel history to answer questions as time permits.

Typically, an IRC [client](http://en.wikipedia.org/wiki/Comparison_of_Internet_Relay_Chat_clients) (http://en.wikipedia.org/wiki/Comparison_of_Internet_Relay_Chat_clients) is used to access the #freenas IRC channel. Alternately, use [webchat](http://webchat.freenode.net/?channels=freenas) (<http://webchat.freenode.net/?channels=freenas>) from a web browser.

To get the most out of the IRC channel, keep these points in mind:

- Do not ask “can anyone help me?”. Just ask the question. If someone knows the answer, they will try to help.
- Do not ask a question and then leave. Users who know the answer cannot help you if you disappear.
- Do not take it personally if no one answers or demand that someone answers your question. Maybe no one who knows the answer is available, maybe your question is really difficult, or maybe it is a question that has already been answered many times in the other support resources. Try asking again in a few hours or research the other resources to see if you have missed anything.
- Do not post error messages in the channel as the IRC software will probably kick you out. Instead, use a pasting service such as [pastebin](http://www.pastebin.com/) (<http://www.pastebin.com/>) and paste the resulting URL into the IRC discussion.

25.4 Videos

A series of instructional videos are available for FreeNAS®:

- [Install Murmur \(Mumble server\) on FreeNAS/FreeBSD \(https://www.youtube.com/watch?v=aAeZRNfarJc\)](https://www.youtube.com/watch?v=aAeZRNfarJc)
- [FreeNAS® 9.10 - Certificate Authority & SSL Certificates \(https://www.youtube.com/watch?v=OT1Le5VQIc0\)](https://www.youtube.com/watch?v=OT1Le5VQIc0)
- [How to Update FreeNAS® 9.10 \(https://www.youtube.com/watch?v=2nvb90AhgL8\)](https://www.youtube.com/watch?v=2nvb90AhgL8)
- [FreeNAS® 9.10 LAGG & VLAN Overview \(https://www.youtube.com/watch?v=wqSH_uQSArQ\)](https://www.youtube.com/watch?v=wqSH_uQSArQ)
- [FreeNAS 9.10 and Samba \(SMB\) Permissions \(https://www.youtube.com/watch?v=RxxgaE935PM\)](https://www.youtube.com/watch?v=RxxgaE935PM)
- [FreeNAS® 11 - What's New \(https://www.youtube.com/watch?v=-uJ_7eG88zk\)](https://www.youtube.com/watch?v=-uJ_7eG88zk)

- [FreeNAS® 11 - How to Install](https://www.youtube.com/watch?v=R3f-Sr6y-c4) (<https://www.youtube.com/watch?v=R3f-Sr6y-c4>)

25.5 Professional Support

In addition to the freely available community resources, professional support may be available through iXsystem's network of third-party consultants. Submit a support inquiry using the form at <https://www.ixsystems.com/freenas-commercial-support/>.

COMMAND LINE UTILITIES

Several command line utilities which are provided with FreeNAS® are demonstrated in this section.

The following utilities can be used for benchmarking and performance testing:

- [Iperf](#) (page 302): used for measuring maximum TCP and UDP bandwidth performance
- [Netperf](#) (page 305): a tool for measuring network performance
- [IOzone](#) (page 306): filesystem benchmark utility used to perform a broad filesystem analysis
- [arcstat](#) (page 308): used to gather ZFS ARC statistics

The following utilities are specific to RAID controllers:

- [tw_cli](#) (page 313): used to monitor and maintain 3ware RAID controllers
- [MegaCli](#) (page 315): used to configure and manage Broadcom MegaRAID SAS family of RAID controllers

This section also describes these utilities:

- [freenas-debug](#) (page 315): the backend used to dump FreeNAS® debugging information
- [tmux](#) (page 316): a terminal multiplexer similar to GNU screen
- [Dmidecode](#) (page 316): reports information about system hardware as described in the system's BIOS

26.1 Iperf

Iperf is a utility for measuring maximum TCP and UDP bandwidth performance. It can be used to chart network throughput over time. For example, you can use it to test the speed of different types of shares to determine which type best performs on your network.

FreeNAS® includes the Iperf server. To perform network testing, you will need to install an Iperf client on a desktop system that has network access to the FreeNAS® system. This section will demonstrate how to use the [xjperf GUI client](#) (<http://code.google.com/p/xjperf/downloads/detail?name=jperf-2.0.2.zip>) as it works on Windows, Mac OS X, Linux, and BSD systems.

Since this client is Java-based, the appropriate [JRE](#) (<http://www.oracle.com/technetwork/java/javase/downloads/index.html>) must be installed on the client computer.

Linux and BSD users will need to install the iperf package using their operating system's package management system.

To start xjperf on Windows: unzip the downloaded file, start Command Prompt in Run as administrator mode, `cd` to the unzipped folder, and run `jperf.bat`.

To start xjperf on Mac OS X, Linux, or BSD, unzip the downloaded file, `cd` to the unzipped directory, type `chmod u+x jperf.sh`, and run `./jperf.sh`.

Once the client is ready, you need to start the Iperf server on FreeNAS®.

Note: Beginning with FreeNAS® version 11.1, both [iperf2](https://sourceforge.net/projects/iperf2/) (<https://sourceforge.net/projects/iperf2/>) and [iperf3](http://software.es.net/iperf/) (<http://software.es.net/iperf/>) are pre-installed. To use iperf2, use `iperf`. To use iperf3, instead type `iperf3`. The examples below are for iperf2.

To see the available server options, open Shell and type:

```
iperf --help | more
```

or:

```
iperf3 --help | more
```

For example, to perform a TCP test and start the server in daemon mode (to get the prompt back), type:

```
iperf -sD
-----
Server listening on TCP port 5001
TCP window size: 64.0 KByte (default)
-----
Running Iperf Server as a daemon
The Iperf daemon process ID: 4842
```

Note: If you close [Shell](#) (page 290), the daemon process will stop. Have your environment set up (e.g. shares configured and started) before starting the iperf process.

From your desktop, open the client. Enter the IP of address of the FreeNAS® system, specify the running time for the test under *Application layer options* → *Transmit* (the default test time is 10 seconds), and click the *Run Iperf!* button. [Figure 26.1](#) shows an example of the client running on a Windows system while an SFTP transfer is occurring on the network.




Fig. 26.1: Viewing Bandwidth Statistics Using xjperf

Depending upon the traffic being tested (e.g. the type of share running on your network), you may need to test UDP instead of TCP. To start the iperf server in UDP mode, use `iperf -sDu` as the `u` specifies UDP; the startup message should indicate that the server is listening for UDP datagrams. If you are not sure if the traffic that you wish to test is UDP or TCP, run this command to determine which services are running on the FreeNAS® system:

```
sockstat -4 | more
```

USER	COMMAND	PID	FD	PROTO	LOCAL ADDRESS	FOREIGN ADDRESS
root	iperf	4870	6	udp4	*:5001	*:*
root	iperf	4842	6	tcp4	*:5001	*:*
www	nginx	4827	3	tcp4	127.0.0.1:15956	127.0.0.1:9042
www	nginx	4827	5	tcp4	192.168.2.11:80	192.168.2.26:56964
www	nginx	4827	7	tcp4	*:80	*:*
root	sshd	3852	5	tcp4	*:22	*:*
root	python	2503	5	udp4	*:*	*:*
root	mountd	2363	7	udp4	*:812	*:*
root	mountd	2363	8	tcp4	*:812	*:*
root	rpcbind	2359	9	udp4	*:111	*:*
root	rpcbind	2359	10	udp4	*:886	*:*
root	rpcbind	2359	11	tcp4	*:111	*:*
root	nginx	2044	7	tcp4	*:80	*:*
root	python	2029	3	udp4	*:*	*:*
root	python	2029	4	tcp4	127.0.0.1:9042	*:*


```

root    python    2029    7    tcp4          127.0.0.1:9042    127.0.0.1:15956
root    ntpd        1548    20   udp4          *:123             *:
root    ntpd        1548    22   udp4          192.168.2.11:123*:
root    ntpd        1548    25   udp4          127.0.0.1:123    *:
root    syslogd    1089    6    udp4          127.0.0.1:514    *:

```

When you are finished testing, either type `killall iperf` or close Shell to terminate the iperf server process.

26.2 Netperf

Netperf is a benchmarking utility that can be used to measure the performance of unidirectional throughput and end-to-end latency.

Before you can use the `netperf` command, you must start its server process using this command:

```

netserver
Starting netserver with host 'IN(6)ADDR_ANY' port '12865' and family AF_UNSPEC

```

The following command will display the available options for performing tests with the `netperf` command. The [Netperf Manual](http://www.netperf.org/svn/netperf2/tags/netperf-2.6.0/doc/netperf.html) (<http://www.netperf.org/svn/netperf2/tags/netperf-2.6.0/doc/netperf.html>) describes each option in more detail and explains how to perform many types of tests. It is the best reference for understanding how each test works and how to interpret your results. When you are finished with your tests, type `killall netserver` to stop the server process.

```

netperf -h |more
Usage: netperf [global options] -- [test options]
Global options:
  -a send,recv      Set the local send,recv buffer alignment
  -A send,recv      Set the remote send,recv buffer alignment
  -B brandstr        Specify a string to be emitted with brief output
  -c [cpu_rate]      Report local CPU usage
  -C [cpu_rate]      Report remote CPU usage
  -d                Increase debugging output
  -D [secs,units] *  Display interim results at least every secs seconds
                    using units as the initial guess for units per second
  -f G|M|K|g|m|k     Set the output units
  -F fill_file        Pre-fill buffers with data from fill_file
  -h                Display this text
  -H name|ip,fam *   Specify the target machine and/or local ip and family
  -i max,min          Specify the max and min number of iterations (15,1)
  -I lvl[,intvl]      Specify confidence level (95 or 99) (99)
                    and confidence interval in percentage (10)
  -j                Keep additional timing statistics
  -l testlen          Specify test duration (>0 secs) (<0 bytes|trans)
  -L name|ip,fam *   Specify the local ip|name and address family
  -o send,recv        Set the local send,recv buffer offsets
  -O send,recv        Set the remote send,recv buffer offset
  -n numcpu           Set the number of processors for CPU util
  -N                Establish no control connection, do 'send' side only
  -p port,lport*      Specify netserver port number and/or local port
  -P 0|1             Don't/Do display test headers
  -r                Allow confidence to be hit on result only
  -s seconds          Wait seconds between test setup and test start
  -S                Set SO_KEEPALIVE on the data connection
  -t testname         Specify test to perform
  -T lcpu,rcpu        Request netperf/netserver be bound to local/remote cpu
  -v verbosity        Specify the verbosity level
  -W send,recv        Set the number of send,recv buffers
  -v level            Set the verbosity level (default 1, min 0)
  -V                Display the netperf version and exit

```

For those options taking two parms, at least one must be specified; specifying one value without a comma will set both parms to that value, specifying a value with a leading comma will set just the second parm, a value with a trailing comma will set just the first. To set each parm to unique values, specify both and separate them with a comma.

For these options taking two parms, specifying one value with no comma will only set the first parms and will leave the second at the default value. To set the second value it must be preceded with a comma or be a comma-separated pair. This is to retain previous netperf behaviour.

26.3 IOzone

IOzone is a disk and filesystem benchmarking tool. It can be used to test file I/O performance for the following operations: read, write, re-read, re-write, read backwards, read strided, fread, fwrite, random read, pread, mmap, aio_read, and aio_write.

FreeNAS® ships with IOzone, meaning that it can be run from Shell. When using IOzone on FreeNAS®, `cd` to a directory in a volume that you have permission to write to, otherwise you will get an error about being unable to write the temporary file.

Before using IOzone, read through the [IOzone documentation PDF](http://www.iozone.org/docs/IOzone_msword_98.pdf) (http://www.iozone.org/docs/IOzone_msword_98.pdf) as it describes the tests, the many command line switches, and how to interpret your results.

If you have never used this tool before, these resources provide good starting points on which tests to run, when to run them, and how to interpret the results:

- [How To Measure Linux Filesystem I/O Performance With iozone](http://www.cyberciti.biz/tips/linux-filesystem-benchmarking-with-iozone.html) (<http://www.cyberciti.biz/tips/linux-filesystem-benchmarking-with-iozone.html>)
- [Analyzing NFS Client Performance with IOzone](http://www.iozone.org/docs/NFSClientPerf_revised.pdf) (http://www.iozone.org/docs/NFSClientPerf_revised.pdf)
- [10 iozone Examples for Disk I/O Performance Measurement on Linux](http://www.thegeekstuff.com/2011/05/iozone-examples/) (<http://www.thegeekstuff.com/2011/05/iozone-examples/>)

You can receive a summary of the available switches by typing the following command. As you can see from the number of options, IOzone is comprehensive and it may take some time to learn how to use the tests effectively.

Starting with version 9.2.1, FreeNAS® enables compression on newly created ZFS pools by default. Since IOzone creates test data that is compressible, this can skew test results. To configure IOzone to generate incompressible test data, include the options `-tw 1 -ty 1 -c 1`.

Alternatively, consider temporarily disabling compression on the ZFS pool or dataset when running IOzone benchmarks.

Note: If you prefer to visualize the collected data, scripts are available to render IOzone's output in [Gnuplot](http://www.gnuplot.info/) (<http://www.gnuplot.info/>).

```
iozone -h | more
iozone: help mode
Usage: iozone[-s filesize_Kb] [-r record_size_Kb] [-f [path]filename] [-h]
        [-i test] [-E] [-p] [-a] [-A] [-z] [-Z] [-m] [-M] [-t children]
        [-l min_number_procs] [-u max_number_procs] [-v] [-R] [-x] [-o]
        [-d microseconds] [-F path1 path2...] [-V pattern] [-j stride]
        [-T] [-C] [-B] [-D] [-G] [-I] [-H depth] [-k depth] [-U mount_point]
        [-S cache_size] [-O] [-L cacheline_size] [-K] [-g maxfilesize_Kb]
        [-n minfilesize_Kb] [-N] [-Q] [-P start_cpu] [-e] [-c] [-b Excel.xls]
        [-J milliseconds] [-X write_telemetry_filename] [-w] [-W]
        [-Y read_telemetry_filename] [-y minrecsize_Kb] [-q maxrecsize_Kb]
        [-+u] [-+m cluster_filename] [-+d] [-+x multiplier] [-+p # ]
        [-+r] [-+t] [-+X] [-+Z] [-+w percent dedupable] [-+y percent_interior_dedup]
        [-+C percent_dedup_within]
-a Auto mode
-A Auto2 mode
-b Filename Create Excel worksheet file
```

```

-B Use mmap() files
-c Include close in the timing calculations
-C Show bytes transferred by each child in throughput testing
-d # Microsecond delay out of barrier
-D Use msync(MS_ASYNC) on mmap files
-e Include flush (fsync,fflush) in the timing calculations
-E Run extension tests
-f filename to use
-F filenames for each process/thread in throughput test
-g # Set maximum file size (in Kbytes) for auto mode (or #m or #g)
-G Use msync(MS_SYNC) on mmap files
-h help
-H # Use POSIX async I/O with # async operations
-i # Test to run (0=write/rewrite, 1=read/re-read, 2=random-read/write
    3=Read-backwards, 4=Re-write-record, 5=stride-read, 6=fwrite/re-fwrite
    7=fread/Re-fread, 8=random_mix, 9=pwrite/Re-pwrite, 10=pread/Re-pread
    11=pwritev/Re-pwritev, 12=preadv/Re-preadv)
-I Use VxFS VX_DIRECT, O_DIRECT, or O_DIRECTIO for all file operations
-j # Set stride of file accesses to (# * record size)
-J # milliseconds of compute cycle before each I/O operation
-k # Use POSIX async I/O (no bcopy) with # async operations
-K Create jitter in the access pattern for readers
-l # Lower limit on number of processes to run
-L # Set processor cache line size to value (in bytes)
-m Use multiple buffers
-M Report uname -a output
-n # Set minimum file size (in Kbytes) for auto mode (or #m or #g)
-N Report results in microseconds per operation
-o Writes are synch (O_SYNC)
-O Give results in ops/sec.
-p Purge on
-P # Bind processes/threads to processors, starting with this cpu
-q # Set maximum record size (in Kbytes) for auto mode (or #m or #g)
-Q Create offset/latency files
-r # record size in Kb
    or -r #k .. size in Kb
    or -r #m .. size in Mb
    or -r #g .. size in Gb
-R Generate Excel report
-s # file size in Kb
    or -s #k .. size in Kb
    or -s #m .. size in Mb
    or -s #g .. size in Gb
-S # Set processor cache size to value (in Kbytes)
-t # Number of threads or processes to use in throughput test
-T Use POSIX pthreads for throughput tests
-u # Upper limit on number of processes to run
-U Mount point to remount between tests
-v version information
-V # Verify data pattern write/read
-w Do not unlink temporary file
-W Lock file when reading or writing
-x Turn off stone-walling
-X filename Write telemetry file. Contains lines with (offset reclen compute_time) in ascii
-y # Set minimum record size (in Kbytes) for auto mode (or #m or #g)
-Y filename Read telemetry file. Contains lines with (offset reclen compute_time) in ascii
-z Used in conjunction with -a to test all possible record sizes
-Z Enable mixing of mmap I/O and file I/O
-+E Use existing non-Iozone file for read-only testing
-+K Sony special. Manual control of test 8.
-+m Cluster_filename Enable Cluster testing

```

```

-d File I/O diagnostic mode. (To troubleshoot a broken file I/O subsystem)
-u Enable CPU utilization output (Experimental)
-x # Multiplier to use for incrementing file and record sizes
-p # Percentage of mix to be reads
-r Enable O_RSYNC|O_SYNC for all testing.
-t Enable network performance test. Requires -m
-n No retests selected.
-k Use constant aggregate data set size.
-q Delay in seconds between tests.
-l Enable record locking mode.
-L Enable record locking mode, with shared file.
-B Sequential mixed workload.
-A # Enable madvise. 0 = normal, 1=random, 2=sequential 3=dontneed, 4=willneed
-N Do not truncate existing files on sequential writes.
-S # Dedup-able data is limited to sharing within each numerically identified file set
-V Enable shared file. No locking.
-X Enable short circuit mode for filesystem testing ONLY
  ALL Results are NOT valid in this mode.
-Z Enable old data set compatibility mode. WARNING.. Published
  hacks may invalidate these results and generate bogus, high values for results.
-w ## Percent of dedup-able data in buffers.
-y ## Percent of dedup-able within & across files in buffers.
-C ## Percent of dedup-able within & not across files in buffers.
-H Hostname Hostname of the PIT server.
-P Service Service of the PIT server.
-z Enable latency histogram logging.

```

26.4 arcstat

Arcstat is a script that prints out ZFS ARC (https://en.wikipedia.org/wiki/Adaptive_replacement_cache) statistics. Originally it was a perl script created by Sun. That perl script was ported to FreeBSD and was then ported as a Python script for use on FreeNAS®.

Watching ARC hits/misses and percentages will provide an indication of how well your ZFS pool is fetching from the ARC rather than using disk I/O. Ideally, you want as many things fetching from cache as possible. Keep your load in mind as you review the stats. For random reads, expect a miss and having to go to disk to fetch the data. For cached reads, expect it to pull out of the cache and have a hit.

Like all cache systems, the ARC takes time to fill with data. This means that it will have a lot of misses until the pool has been in use for a while. If there continues to be lots of misses and high disk I/O on cached reads, there is cause to investigate further and tune the system.

The [FreeBSD ZFS Tuning Guide](https://wiki.FreeBSD.org/ZFSTuningGuide) (<https://wiki.FreeBSD.org/ZFSTuningGuide>) provides some suggestions for commonly tuned `sysctl` values. It should be noted that performance tuning is more of an art than a science and that any changes you make will probably require several iterations of tune and test. Be aware that what needs to be tuned will vary depending upon the type of workload and that what works for one person's network may not benefit yours.

In particular, the value of pre-fetching depends upon the amount of memory and the type of workload, as seen in these two examples:

- [Understanding ZFS: Prefetch](http://www.cuddletech.com/blog/pivot/entry.php?id=1040) (<http://www.cuddletech.com/blog/pivot/entry.php?id=1040>)
- [ZFS prefetch algorithm can cause performance drawbacks](http://southbrain.com/south/2008/04/the-nightmare-comes-slowly-zfs.html) (<http://southbrain.com/south/2008/04/the-nightmare-comes-slowly-zfs.html>)

FreeNAS® provides two command line scripts which can be manually run from [Shell](#) (page 290):

- `arc_summary.py`: provides a summary of the statistics
- `arcstat.py`: used to watch the statistics in real time

The advantage of these scripts is that they can be used to provide real time (right now) information, whereas the current GUI reporting mechanism is designed to only provide graphs charted over time.

This [forum post](https://forums.freenas.org/index.php?threads/benchmarking-zfs.7928/) (<https://forums.freenas.org/index.php?threads/benchmarking-zfs.7928/>) demonstrates some examples of using these scripts with hints on how to interpret the results.

To view the help for arcstat.py:

```
arcstat.py -h
Usage: arcstat [-hvx] [-f fields] [-o file] [-s string] [interval [count]]
-h: Print this help message
-v: List all possible field headers and definitions
-x: Print extended stats
-f: Specify specific fields to print (see -v)
-o: Redirect output to the specified file
-s: Override default field separator with custom character or string

Examples:
arcstat -o /tmp/a.log 2 10
arcstat -s "," -o /tmp/a.log 2 10
arcstat -v
arcstat -f time,hit%,dh%,ph%,mh% 1
```

To view ARC statistics in real time, specify an interval and a count. This command will display every 1 second for a count of five.

```
arcstat.py 1 5
      time  read  miss  miss%  dmis  dm%  pmis  pm%  mmis  mm%  arcsz  c
06:19:03    7    0    0    0    0    0    0    0    0    0  153M  6.6G
06:19:04   257    0    0    0    0    0    0    0    0    0  153M  6.6G
06:19:05   193    0    0    0    0    0    0    0    0    0  153M  6.6G
06:19:06   193    0    0    0    0    0    0    0    0    0  153M  6.6G
06:19:07   255    0    0    0    0    0    0    0    0    0  153M  6.6G
```

Table 26.1 briefly describes the columns in the output.

Table 26.1: arcstat Column Descriptions

Column	Description
read	total ARC accesses/second
miss	ARC misses/second
miss%	ARC miss percentage
dmis	demand data misses/second
dm%	demand data miss percentage
pmis	prefetch misses per second
pm%	prefetch miss percentage
mmis	metadata misses/second
mm%	metadata miss percentage
arcsz	arc size
c	arc target size

To receive a summary of statistics, use:

```
arcsummary.py
System Memory:
  2.36%  93.40  MiB Active,    8.95%  353.43  MiB Inact
  8.38%  330.89  MiB Wired,    0.15%   5.90  MiB Cache
 80.16%   3.09  GiB Free,    0.00%   0  Bytes Gap
Real Installed:                4.00  GiB
Real Available:                99.31%  3.97  GiB
```

```

Real Managed:          97.10%  3.86   GiB
Logical Total:         4.00   GiB
Logical Used:          13.93%  570.77 MiB
Logical Free:          86.07%  3.44   GiB
Kernel Memory:        87.62   MiB
  Data:                69.91%  61.25   MiB
  Text:                30.09%  26.37   MiB
Kernel Memory Map:    3.86   GiB
  Size:                5.11%  201.70 MiB
  Free:               94.89%  3.66   GiB
ARC Summary: (HEALTHY)
  Storage pool Version: 5000
  Filesystem Version: 5
  Memory Throttle Count: 0
ARC Misc:
  Deleted: 8
  Mutex Misses: 0
  Evict Skips: 0
ARC Size: 5.83% 170.45 MiB
  Target Size: (Adaptive) 100.00% 2.86 GiB
  Min Size (Hard Limit): 12.50% 365.69 MiB
  Max Size (High Water): 8:1 2.86 GiB
ARC Size Breakdown:
  Recently Used Cache Size: 50.00% 1.43 GiB
  Frequently Used Cache Size: 50.00% 1.43 GiB
ARC Hash Breakdown:
  Elements Max: 5.90k
  Elements Current: 100.00% 5.90k
  Collisions: 72
  Chain Max: 1
  Chains: 23
ARC Total accesses: 954.06k
  Cache Hit Ratio: 99.18% 946.25k
  Cache Miss Ratio: 0.82% 7.81k
  Actual Hit Ratio: 98.84% 943.00k
  Data Demand Efficiency: 99.20% 458.77k
  CACHE HITS BY CACHE LIST:
    Anonymously Used: 0.34% 3.25k
    Most Recently Used: 3.73% 35.33k
    Most Frequently Used: 95.92% 907.67k
    Most Recently Used Ghost: 0.00% 0
    Most Frequently Used Ghost: 0.00% 0
  CACHE HITS BY DATA TYPE:
    Demand Data: 48.10% 455.10k
    Prefetch Data: 0.00% 0
    Demand Metadata: 51.56% 487.90k
    Prefetch Metadata: 0.34% 3.25k
  CACHE MISSES BY DATA TYPE:
    Demand Data: 46.93% 3.66k
    Prefetch Data: 0.00% 0
    Demand Metadata: 49.76% 3.88k
    Prefetch Metadata: 3.30% 258
ZFS Tunable (sysctl):
  kern.maxusers 590
  vm.kmem_size 4141375488
  vm.kmem_size_scale 1
  vm.kmem_size_min 0
  vm.kmem_size_max 1319413950874
  vfs.zfs.vol.unmap_enabled 1
  vfs.zfs.vol.mode 2
  vfs.zfs.sync_pass_rewrite 2

```

vfs.zfs.sync_pass_dont_compress	5
vfs.zfs.sync_pass_deferred_free	2
vfs.zfs.zio.exclude_metadata	0
vfs.zfs.zio.use_uma	1
vfs.zfs.cache_flush_disable	0
vfs.zfs.zil_replay_disable	0
vfs.zfs.version.zpl	5
vfs.zfs.version.spa	5000
vfs.zfs.version.acl	1
vfs.zfs.version.ioctl	5
vfs.zfs.debug	0
vfs.zfs.super_owner	0
vfs.zfs.min_auto_ashift	9
vfs.zfs.max_auto_ashift	13
vfs.zfs.vdev.write_gap_limit	4096
vfs.zfs.vdev.read_gap_limit	32768
vfs.zfs.vdev.aggregation_limit	131072
vfs.zfs.vdev.trim_max_active	64
vfs.zfs.vdev.trim_min_active	1
vfs.zfs.vdev.scrub_max_active	2
vfs.zfs.vdev.scrub_min_active	1
vfs.zfs.vdev.async_write_max_active	10
vfs.zfs.vdev.async_write_min_active	1
vfs.zfs.vdev.async_read_max_active	3
vfs.zfs.vdev.async_read_min_active	1
vfs.zfs.vdev.sync_write_max_active	10
vfs.zfs.vdev.sync_write_min_active	10
vfs.zfs.vdev.sync_read_max_active	10
vfs.zfs.vdev.sync_read_min_active	10
vfs.zfs.vdev.max_active	1000
vfs.zfs.vdev.async_write_active_max_dirty_percent	60
vfs.zfs.vdev.async_write_active_min_dirty_percent	30
vfs.zfs.vdev.mirror.non_rotating_seek_incl	1
vfs.zfs.vdev.mirror.non_rotating_inc	0
vfs.zfs.vdev.mirror.rotating_seek_offset	1048576
vfs.zfs.vdev.mirror.rotating_seek_inc	5
vfs.zfs.vdev.mirror.rotating_inc	0
vfs.zfs.vdev.trim_on_init	1
vfs.zfs.vdev.larger_ashift_minimal	0
vfs.zfs.vdev.bio_delete_disable	0
vfs.zfs.vdev.bio_flush_disable	0
vfs.zfs.vdev.cache.bshift	16
vfs.zfs.vdev.cache.size	0
vfs.zfs.vdev.cache.max	16384
vfs.zfs.vdev.metaslabs_per_vdev	200
vfs.zfs.vdev.trim_max_pending	10000
vfs.zfs.txg.timeout	5
vfs.zfs.trim.enabled	1
vfs.zfs.trim.max_interval	1
vfs.zfs.trim.timeout	30
vfs.zfs.trim.txg_delay	32
vfs.zfs.space_map_blksize	4096
vfs.zfs.spa_slop_shift	5
vfs.zfs.spa_asize_inflation	24
vfs.zfs.deadman_enabled	1
vfs.zfs.deadman_checktime_ms	5000
vfs.zfs.deadman_synctime_ms	1000000
vfs.zfs.recover	0
vfs.zfs.spa_load_verify_data	1
vfs.zfs.spa_load_verify_metadata	1
vfs.zfs.spa_load_verify_maxinflight	10000

```

vfs.zfs.check_hostid 1
vfs.zfs.mg_fragmentation_threshold 85
vfs.zfs.mg_noalloc_threshold 0
vfs.zfs.condense_pct 200
vfs.zfs.metaslab.bias_enabled 1
vfs.zfs.metaslab.lba_weighting_enabled 1
vfs.zfs.metaslab.fragmentation_factor_enabled 1
vfs.zfs.metaslab.preload_enabled 1
vfs.zfs.metaslab.preload_limit 3
vfs.zfs.metaslab.unload_delay 8
vfs.zfs.metaslab.load_pct 50
vfs.zfs.metaslab.min_alloc_size 33554432
vfs.zfs.metaslab.df_free_pct 4
vfs.zfs.metaslab.df_alloc_threshold 131072
vfs.zfs.metaslab.debug_unload 0
vfs.zfs.metaslab.debug_load 0
vfs.zfs.metaslab.fragmentation_threshold 70
vfs.zfs.metaslab.gang_bang 16777217
vfs.zfs.free_bpobj_enabled 1
vfs.zfs.free_max_blocks 18446744073709551615
vfs.zfs.no_scrub_prefetch 0
vfs.zfs.no_scrub_io 0
vfs.zfs.resilver_min_time_ms 3000
vfs.zfs.free_min_time_ms 1000
vfs.zfs.scan_min_time_ms 1000
vfs.zfs.scan_idle 50
vfs.zfs.scrub_delay 4
vfs.zfs.resilver_delay 2
vfs.zfs.top_maxinflight 32
vfs.zfs.delay_scale 500000
vfs.zfs.delay_min_dirty_percent 60
vfs.zfs.dirty_data_sync 67108864
vfs.zfs.dirty_data_max_percent 10
vfs.zfs.dirty_data_max_max 4294967296
vfs.zfs.dirty_data_max 426512793
vfs.zfs.max_records_size 1048576
vfs.zfs.zfetch.array_rd_sz 1048576
vfs.zfs.zfetch.max_distance 8388608
vfs.zfs.zfetch.min_sec_reap 2
vfs.zfs.zfetch.max_streams 8
vfs.zfs.prefetch_disable 1
vfs.zfs.mdcomp_disable 0
vfs.zfs.nopwrite_enabled 1
vfs.zfs.dedup.prefetch 1
vfs.zfs.l2c_only_size 0
vfs.zfs.mfu_ghost_data_lsize 0
vfs.zfs.mfu_ghost_metadata_lsize 0
vfs.zfs.mfu_ghost_size 0
vfs.zfs.mfu_data_lsize 26300416
vfs.zfs.mfu_metadata_lsize 1780736
vfs.zfs.mfu_size 29428736
vfs.zfs.mru_ghost_data_lsize 0
vfs.zfs.mru_ghost_metadata_lsize 0
vfs.zfs.mru_ghost_size 0
vfs.zfs.mru_data_lsize 122090496
vfs.zfs.mru_metadata_lsize 2235904
vfs.zfs.mru_size 139389440
vfs.zfs.anon_data_lsize 0
vfs.zfs.anon_metadata_lsize 0
vfs.zfs.anon_size 163840
vfs.zfs.l2arc_norw 1

```


vfs.zfs.l2arc_feed_again	1
vfs.zfs.l2arc_noprefetch	1
vfs.zfs.l2arc_feed_min_ms	200
vfs.zfs.l2arc_feed_secs	1
vfs.zfs.l2arc_headroom	2
vfs.zfs.l2arc_write_boost	8388608
vfs.zfs.l2arc_write_max	8388608
vfs.zfs.arc_meta_limit	766908416
vfs.zfs.arc_free_target	7062
vfs.zfs.arc_shrink_shift	7
vfs.zfs.arc_average_blocksize	8192
vfs.zfs.arc_min	383454208
vfs.zfs.arc_max	3067633664

When reading the tunable values, 0 means no, 1 typically means yes, and any other number represents a value. To receive a brief description of a “sysctl” value, use `sysctl -d`. For example:

```
sysctl -d vfs.zfs.zio.use_uma
vfs.zfs.zio.use_uma: Use uma(9) for ZIO allocations
```

The ZFS tunables require a fair understanding of how ZFS works, meaning that you will be reading man pages and searching for the meaning of acronyms you are unfamiliar with. Do not change a tunable’s value without researching it first. If the tunable takes a numeric value (rather than 0 for no or 1 for yes), do not make one up. Instead, research examples of beneficial values that match your workload.

If you decide to change any of the ZFS tunables, continue to monitor the system to determine the effect of the change. It is recommended that you test your changes first at the command line using `sysctl`. For example, to disable pre-fetch (i.e. change disable to 1 or yes):

```
sysctl vfs.zfs.prefetch_disable=1
vfs.zfs.prefetch_disable: 0 -> 1
```

The output will indicate the old value followed by the new value. If the change is not beneficial, change it back to the original value. If the change turns out to be beneficial, you can make it permanent by creating a `sysctl` using the instructions in [Tunables](#) (page 68).

26.5 tw_cli

FreeNAS® includes the `tw_cli` command line utility for providing controller, logical unit, and drive management for AMCC/3ware ATA RAID Controllers. The supported models are listed in the man pages for the [twe\(4\)](#) (<http://www.freebsd.org/cgi/man.cgi?query=twe>) and [twa\(4\)](#) (<http://www.freebsd.org/cgi/man.cgi?query=twa>) drivers.

Before using this command, read its [man page](#) (http://www.cyberciti.biz/files/tw_cli.8.html) as it describes the terminology and provides some usage examples.

If you type `tw_cli` in Shell, the prompt will change, indicating that you have entered interactive mode where you can run all sorts of maintenance commands on the controller and its arrays.

Alternately, you can specify one command to run. For example, to view the disks in the array:

```
tw_cli /c0 show
```

Unit	UnitType	Status	%RCmpl	%V/I/M	Stripe	Size (GB)	Cache	AVrfy
u0	RAID-6	OK	-	-	256K	5587.88	RiW	ON
u1	SPARE	OK	-	-	-	931.505	-	OFF
u2	RAID-10	OK	-	-	256K	1862.62	RiW	ON

VPort	Status	Unit	Size	Type	Phy	Encl-Slot	Model

p8	OK	u0	931.51	GB	SAS	-	/c0/e0/slt0	SEAGATE	ST31000640SS
p9	OK	u0	931.51	GB	SAS	-	/c0/e0/slt1	SEAGATE	ST31000640SS
p10	OK	u0	931.51	GB	SAS	-	/c0/e0/slt2	SEAGATE	ST31000640SS
p11	OK	u0	931.51	GB	SAS	-	/c0/e0/slt3	SEAGATE	ST31000640SS
p12	OK	u0	931.51	GB	SAS	-	/c0/e0/slt4	SEAGATE	ST31000640SS
p13	OK	u0	931.51	GB	SAS	-	/c0/e0/slt5	SEAGATE	ST31000640SS
p14	OK	u0	931.51	GB	SAS	-	/c0/e0/slt6	SEAGATE	ST31000640SS
p15	OK	u0	931.51	GB	SAS	-	/c0/e0/slt7	SEAGATE	ST31000640SS
p16	OK	u1	931.51	GB	SAS	-	/c0/e0/slt8	SEAGATE	ST31000640SS
p17	OK	u2	931.51	GB	SATA	-	/c0/e0/slt9	ST31000340NS	
p18	OK	u2	931.51	GB	SATA	-	/c0/e0/slt10	ST31000340NS	
p19	OK	u2	931.51	GB	SATA	-	/c0/e0/slt11	ST31000340NS	
p20	OK	u2	931.51	GB	SATA	-	/c0/e0/slt15	ST31000340NS	

Name	OnlineState	BBUReady	Status	Volt	Temp	Hours	LastCapTest
bbu	On	Yes	OK	OK	OK	212	03-Jan-2012

Or, to review the event log:

tw_cli /c0 show events					
Ctl	Date	Severity	AEN Message		
c0	[Thu Feb 23 2012 14:01:15]	INFO	Battery charging started		
c0	[Thu Feb 23 2012 14:03:02]	INFO	Battery charging completed		
c0	[Sat Feb 25 2012 00:02:18]	INFO	Verify started: unit=0		
c0	[Sat Feb 25 2012 00:02:18]	INFO	Verify started: unit=2,subunit=0		
c0	[Sat Feb 25 2012 00:02:18]	INFO	Verify started: unit=2,subunit=1		
c0	[Sat Feb 25 2012 03:49:35]	INFO	Verify completed: unit=2,subunit=0		
c0	[Sat Feb 25 2012 03:51:39]	INFO	Verify completed: unit=2,subunit=1		
c0	[Sat Feb 25 2012 21:55:59]	INFO	Verify completed: unit=0		
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery health check started		
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery health check completed		
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery charging started		
c0	[Thu Mar 01 2012 13:53:03]	INFO	Battery charging completed		
c0	[Sat Mar 03 2012 00:01:24]	INFO	Verify started: unit=0		
c0	[Sat Mar 03 2012 00:01:24]	INFO	Verify started: unit=2,subunit=0		
c0	[Sat Mar 03 2012 00:01:24]	INFO	Verify started: unit=2,subunit=1		
c0	[Sat Mar 03 2012 04:04:27]	INFO	Verify completed: unit=2,subunit=0		
c0	[Sat Mar 03 2012 04:06:25]	INFO	Verify completed: unit=2,subunit=1		
c0	[Sat Mar 03 2012 16:22:05]	INFO	Verify completed: unit=0		
c0	[Thu Mar 08 2012 13:41:39]	INFO	Battery charging started		
c0	[Thu Mar 08 2012 13:43:42]	INFO	Battery charging completed		
c0	[Sat Mar 10 2012 00:01:30]	INFO	Verify started: unit=0		
c0	[Sat Mar 10 2012 00:01:30]	INFO	Verify started: unit=2,subunit=0		
c0	[Sat Mar 10 2012 00:01:30]	INFO	Verify started: unit=2,subunit=1		
c0	[Sat Mar 10 2012 05:06:38]	INFO	Verify completed: unit=2,subunit=0		
c0	[Sat Mar 10 2012 05:08:57]	INFO	Verify completed: unit=2,subunit=1		
c0	[Sat Mar 10 2012 15:58:15]	INFO	Verify completed: unit=0		

If you add some disks to the array and they are not showing up in the GUI, try running this command:

```
tw_cli /c0 rescan
```

Use the drives to create units and export them to the operating system. When finished, run `camcontrol rescan all` and they should now be available in the FreeNAS® GUI.

This [forum post](https://forums.freenas.org/index.php?threads/3ware-drive-monitoring.13835/) (<https://forums.freenas.org/index.php?threads/3ware-drive-monitoring.13835/>) contains a handy wrapper script that will notify you of errors.

26.6 MegaCli

MegaCli is the command line interface for the Broadcom :MegaRAID SAS family of RAID controllers. FreeNAS® also includes the [mfiutil\(8\)](http://www.freebsd.org/cgi/man.cgi?query=mfiutil) (<http://www.freebsd.org/cgi/man.cgi?query=mfiutil>) utility which can be used to configure and manage connected storage devices.

The MegaCli command is quite complex with several dozen options. The commands demonstrated in the [Emergency Cheat Sheet](http://tools.rapidsoft.de/perc/perc-cheat-sheet.html) (<http://tools.rapidsoft.de/perc/perc-cheat-sheet.html>) can get you started.

26.7 freenas-debug

The FreeNAS® GUI provides an option to save debugging information to a text file using **System → Advanced → Save Debug**. This debugging information is created by the `freenas-debug` command line utility and a copy of the information is saved to `/var/tmp/fndebug`.

This command can be run manually from [Shell](#) (page 290) to gather specific debugging information. To see a usage explanation listing all options, run the command without any options:

```
freenas-debug
Usage: /usr/local/bin/freenas-debug <options>
Where options are:
  -e Email debug log to this comma-delimited list of email addresses
  -A Dump all debug information

  -a Dump Active Directory Configuration
  -f Dump AFP Configuration
  -c Dump (AD|LDAP) Cache
  -B Dump System Configuration Database
  -D Dump Domain Controller Configuration
  -d Dump DTrace Scripts
  -g Dump GEOM Configuration
  -G Dump Grub Configuration
  -h Dump Hardware Configuration
  -I Dump IPMI Configuration
  -i Dump iSCSI Configuration
  -j Dump Jail Information
  -l Dump LDAP Configuration
  -T Loader Configuration Information
  -n Dump Network Configuration
  -N Dump NFS Configuration
  -S Dump SMART Information
  -C Dump SMB Configuration
  -s Dump SSL Configuration
  -y Dump Sysctl Configuration
  -t Dump System Information
  -v Dump Boot System File Verification Status and Inconsistencies
  -z Dump ZFS Configuration
```

Individual tests can be run alone. For example, when troubleshooting an Active Directory configuration, use:

```
freenas-debug -a
```

To collect the output of every module, use `-A`:

```
freenas-debug -A
```

26.8 tmux

`tmux` is a terminal multiplexer which enables a number of `:terminals` to be created, accessed, and controlled from a single `:screen`. `tmux` is an alternative to GNU `screen`. Similar to `screen`, `tmux` can be detached from a screen and continue running in the background, then later reattached. Unlike [Shell](#) (page 290), `tmux` allows you to have access to a command prompt while still providing access to the graphical administration screens.

To start a session, simply type `tmux`. As seen in [Figure 26.2](#), a new session with a single window opens with a status line at the bottom of the screen. This line shows information on the current session and is used to enter interactive commands.



Fig. 26.2: `tmux` Session

To create a second window, press `Ctrl+b` then `"`. To close a window, type `exit` within the window.

`tmux(1)` (<http://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/./man1/tmux.1?query=tmux>) lists all of the key bindings and commands for interacting with `tmux` windows and sessions.

If you close [Shell](#) (page 290) while `tmux` is running, it will detach its session. The next time you open `Shell`, run `tmux attach` to return to the previous session. To leave the `tmux` session entirely, type `exit`. If you have multiple windows running, you will need to `exit` out of each first.

These resources provide more information about using `tmux`:

- [A tmux Crash Course](https://robots.thoughtbot.com/a-tmux-crash-course) (<https://robots.thoughtbot.com/a-tmux-crash-course>)
- [TMUX - The Terminal Multiplexer](http://blog.hawkhost.com/2010/06/28/tmux-the-terminal-multiplexer/) (<http://blog.hawkhost.com/2010/06/28/tmux-the-terminal-multiplexer/>)

26.9 Dmidecode

`Dmidecode` reports hardware information as reported by the system BIOS. `Dmidecode` does not scan the hardware, it only reports what the BIOS told it to. A sample output can be seen [here](http://www.nongnu.org/dmidecode/sample/dmidecode.txt) (<http://www.nongnu.org/dmidecode/sample/dmidecode.txt>).

To view the BIOS report, type the command with no arguments:

dmidecode | more

[dmidecode\(8\)](http://linux.die.net/man/8/dmidecode) (<http://linux.die.net/man/8/dmidecode>) describes the supported strings and types.

CONTRIBUTING TO FREENAS®

FreeNAS® is an open source community, relying on the input and expertise of its users to help grow and improve FreeNAS®. When you take time to assist the community, your contributions benefit everyone who uses FreeNAS®.

This section describes some areas of participation to get you started. It is by no means an exhaustive list. If you have an idea that you think would benefit the FreeNAS® community, bring it up on one of the resources mentioned in [Support Resources](#) (page 299).

This section demonstrates how you can:

- [Help with Translation](#) (page 318)

27.1 Translation

Not everyone speaks English, and having a complete translation of the user interface into native languages can make FreeNAS® much more useful to communities around the world.

FreeNAS® uses [Weblate](https://weblate.org/) (<https://weblate.org/>) to manage the translation of text shown in the FreeNAS® graphical administrative interface. Weblate provides an easy-to-use web-based editor and commenting system, making it possible for individuals to assist with translation or comment on existing translations.

To see the status of translations, open <http://weblate.trueos.org/projects/freenas/>, as shown in [Figure 27.1](#).

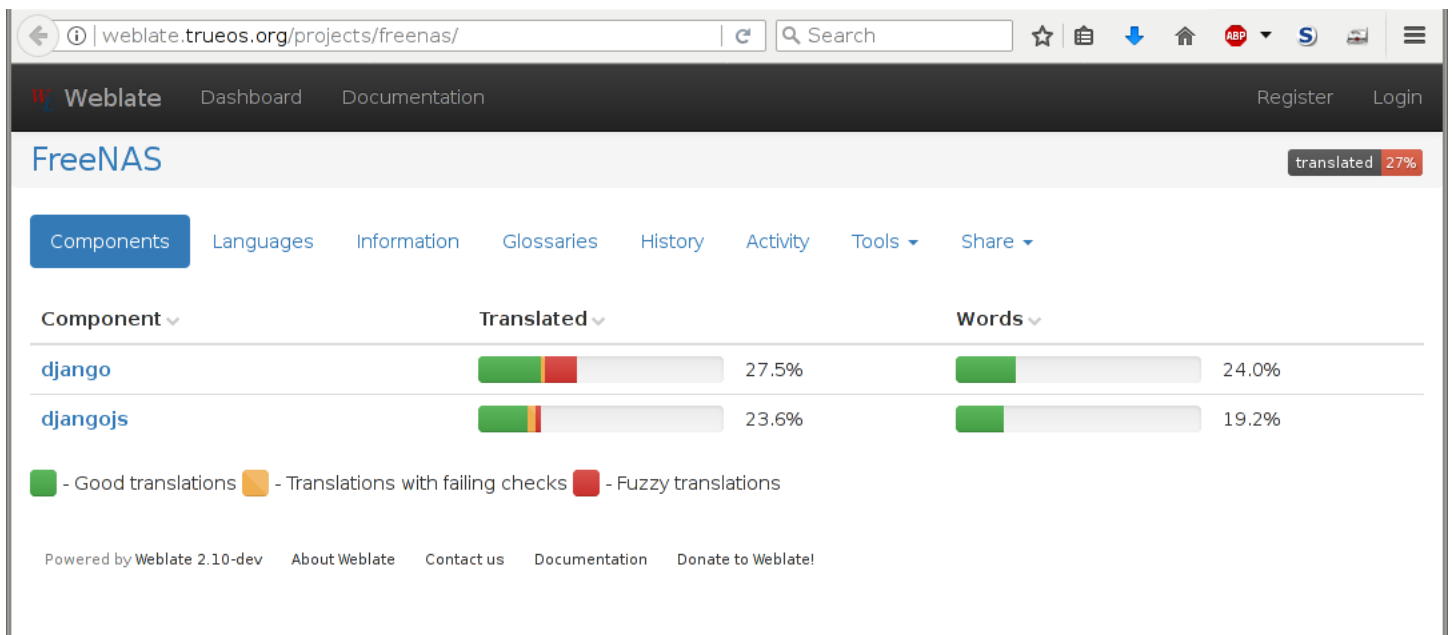


Fig. 27.1: FreeNAS® Translation System

To assist with translating FreeNAS®, create an account by clicking the Register button. Enter the information requested, then a confirmation email will be sent. Follow the link in the email to set a password and complete the account creation. The Dashboard screen is shown after logging in:

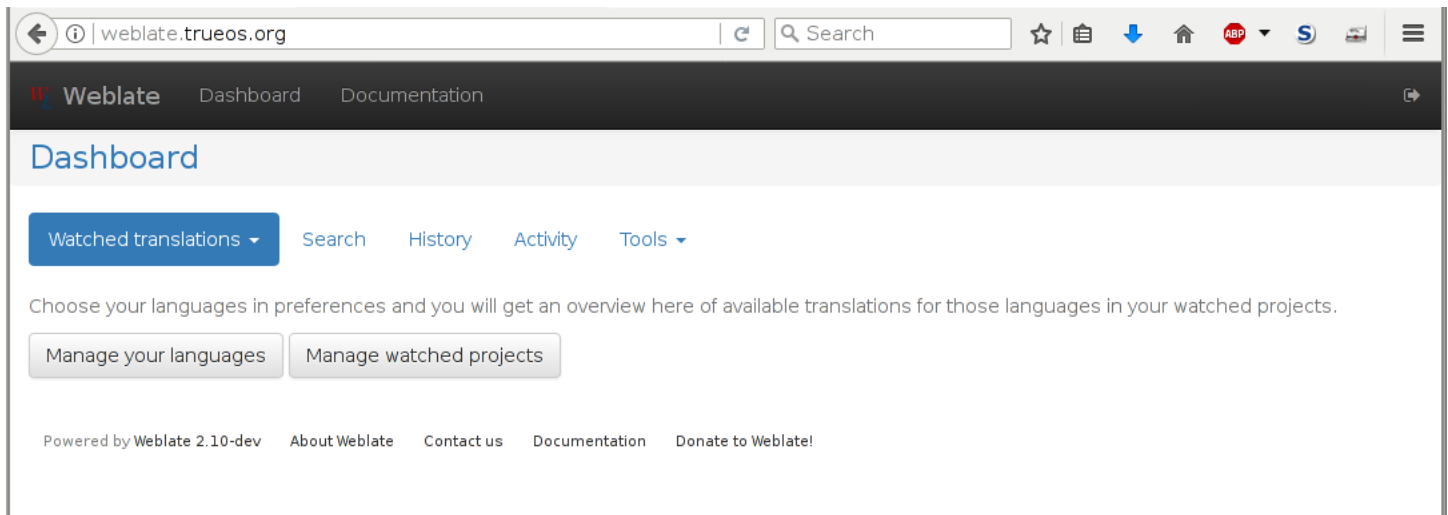


Fig. 27.2: Weblate Dashboard

Click Manage your languages to choose languages for translation. Select languages, then click Save. Click the Dashboard link at the top of the screen to go back to the dashboard, then choose Your languages from the drop-down menu:

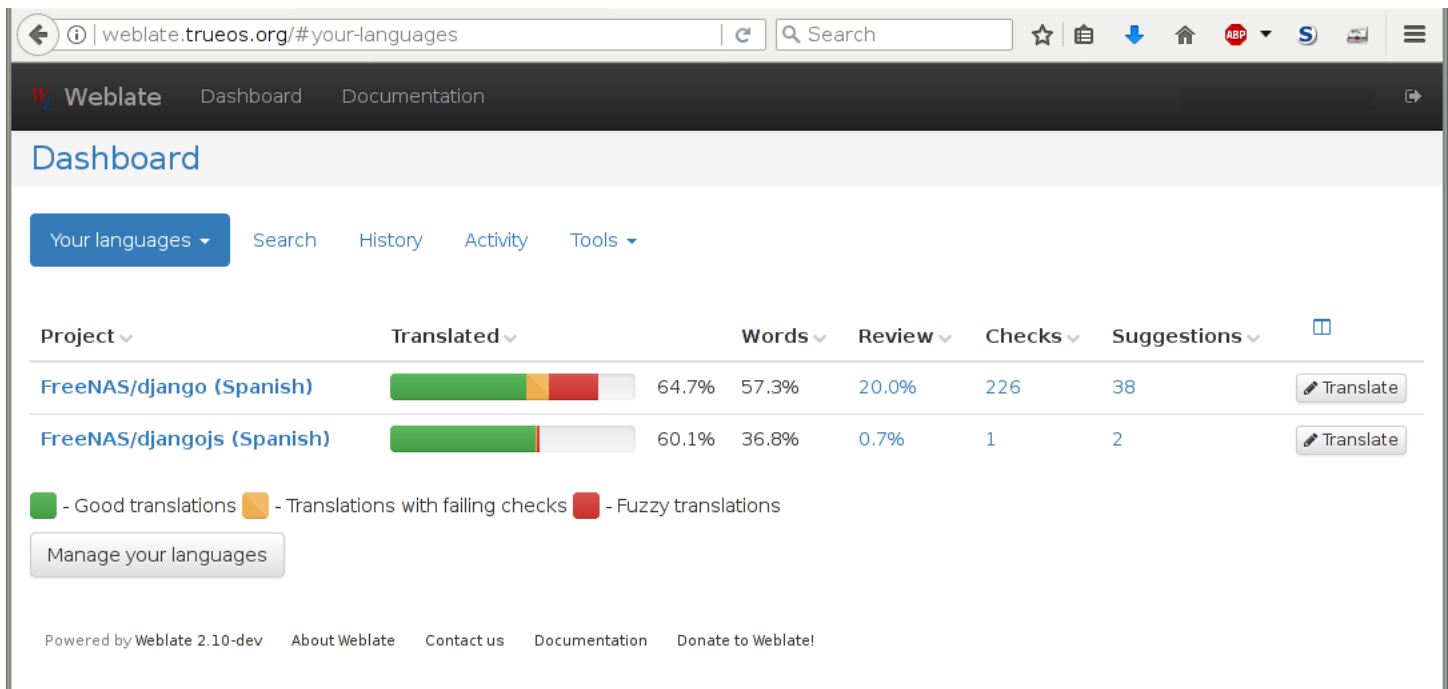


Fig. 27.3: Selected Languages

Projects are a collection of text to be translated. In this example, the Django and DjangoJS projects have both been partially translated into Spanish. Click one of the entries under Project to help translate that project.

The Overview screen shows the current translation status along with categories of translatable strings:

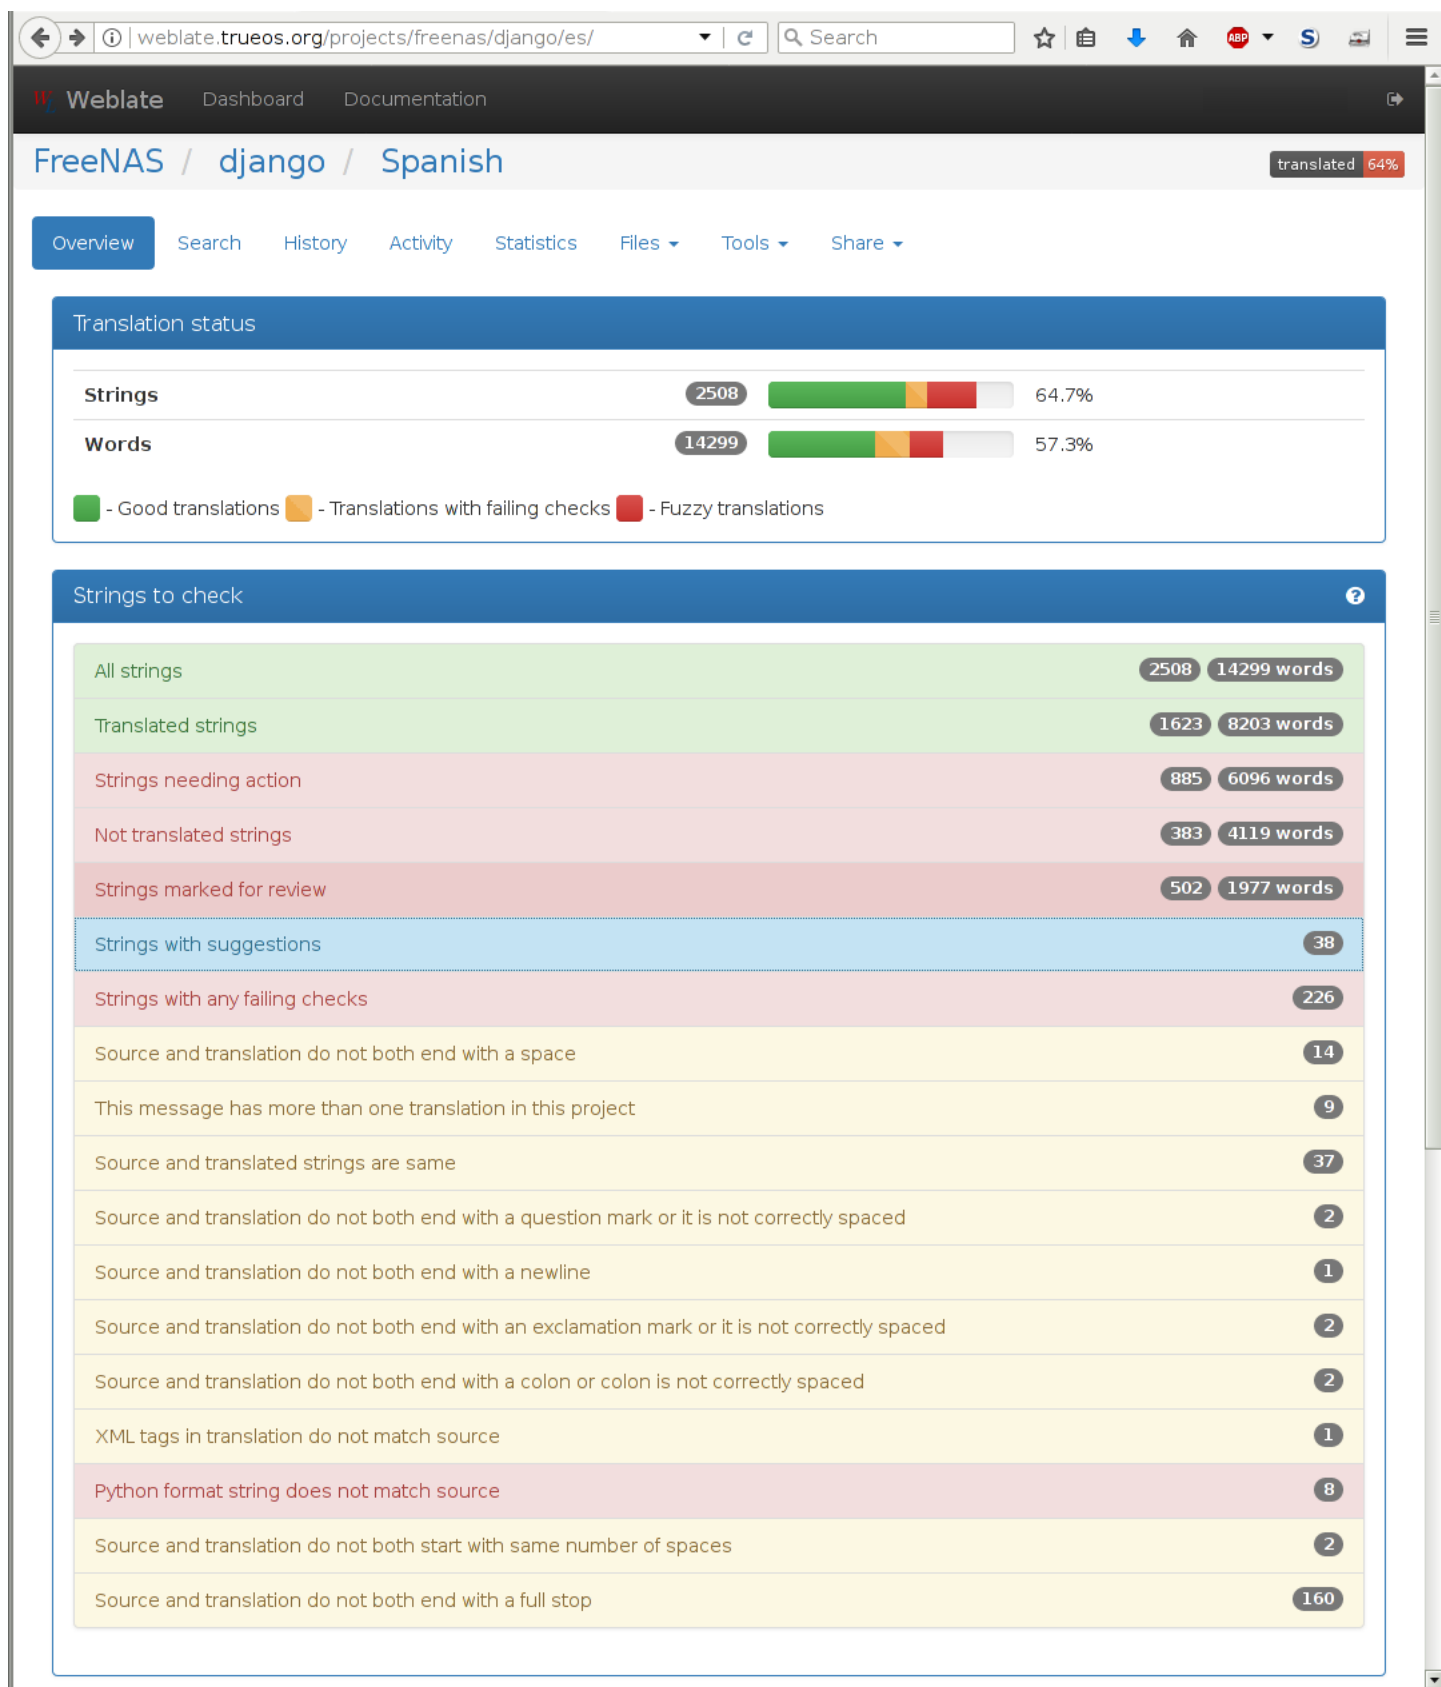


Fig. 27.4: Translation Overview

Click on a category of string, like Strings needing action, to see the translation screen:

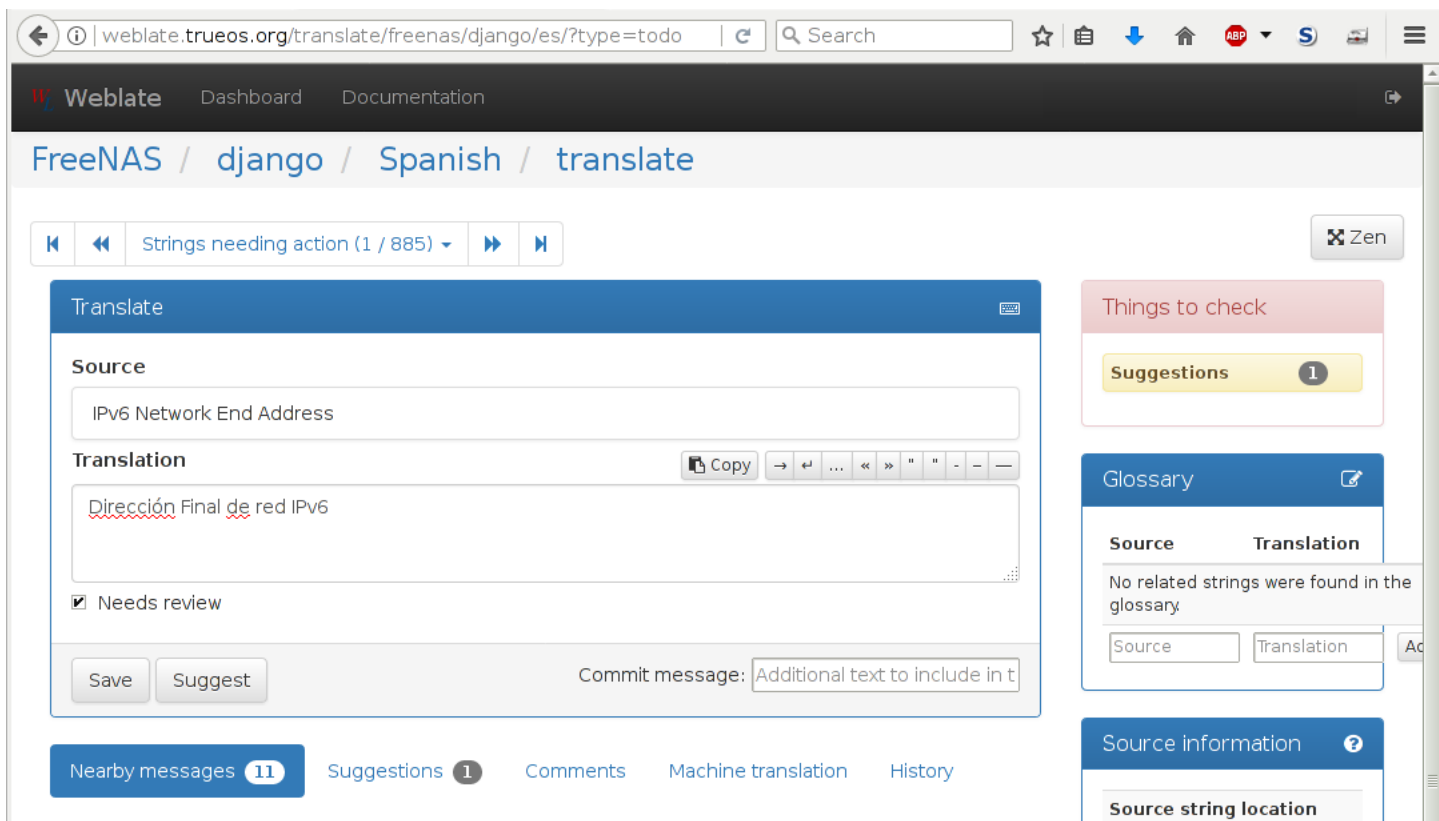


Fig. 27.5: Translate Strings

Enter translations here, clicking Save to save the work. The controls at the top of the screen can be used to skip forward and back in the list of strings to be translated. Click Dashboard at the top of the screen to return to the Dashboard.

All assistance with translations helps to benefit the FreeNAS® community. Thank you!

ZFS PRIMER

ZFS is an advanced, modern filesystem that was specifically designed to provide features not available in traditional UNIX filesystems. It was originally developed at Sun with the intent to open source the filesystem so that it could be ported to other operating systems. After the Oracle acquisition of Sun, some of the original ZFS engineers founded [OpenZFS](http://open-zfs.org/wiki/Main_Page) (http://open-zfs.org/wiki/Main_Page) to provide continued, collaborative development of the open source version. To differentiate itself from Oracle ZFS version numbers, OpenZFS uses feature flags. Feature flags are used to tag features with unique names in order to provide portability between OpenZFS implementations running on different platforms, as long as all of the feature flags enabled on the ZFS pool are supported by both platforms. FreeNAS® uses OpenZFS and each new version of FreeNAS® keeps up-to-date with the latest feature flags and OpenZFS bug fixes.

Here is an overview of the features provided by ZFS:

ZFS is a transactional, Copy-On-Write (COW) (https://en.wikipedia.org/wiki/ZFS#Copy-on-write_transactional_model) filesystem. For each write request, a copy is made of the associated disk blocks and all changes are made to the copy rather than to the original blocks. When the write is complete, all block pointers are changed to point to the new copy. This means that ZFS always writes to free space, most writes are sequential, and old versions of files are not unlinked until a complete new version has been written successfully. ZFS has direct access to disks and bundles multiple read and write requests into transactions. Most filesystems cannot do this, as they only have access to disk blocks. A transaction either completes or fails, meaning there will never be a [write-hole](https://blogs.oracle.com/bonwick/entry/raid_z) (https://blogs.oracle.com/bonwick/entry/raid_z) and a filesystem checker utility is not necessary. Because of the transactional design, as additional storage capacity is added, it becomes immediately available for writes. To rebalance the data, one can copy it to re-write the existing data across all available disks. As a 128-bit filesystem, the maximum filesystem or file size is 16 exabytes.

ZFS was designed to be a self-healing filesystem. As ZFS writes data, it creates a checksum for each disk block it writes. As ZFS reads data, it validates the checksum for each disk block it reads. Media errors or “bit rot” can cause data to change, and the checksum no longer matches. When ZFS identifies a disk block checksum error on a pool that is mirrored or uses RAIDZ, it replaces the corrupted data with the correct data. Since some disk blocks are rarely read, regular scrubs should be scheduled so that ZFS can read all of the data blocks to validate their checksums and correct any corrupted blocks. While multiple disks are required in order to provide redundancy and data correction, ZFS will still provide data corruption detection to a system with one disk. FreeNAS® automatically schedules a monthly scrub for each ZFS pool and the results of the scrub are displayed by selecting the [Volume](#) (page 115) then clicking the Volume Status button. Checking scrub results can provide an early indication of potential disk problems.

Unlike traditional UNIX filesystems, it is not necessary to define partition sizes when filesystems are created. Instead, a group of disks, known as a vdev, are built into a ZFS pool. Filesystems are created from the pool as needed. As more capacity is needed, identical vdevs can be striped into the pool. In FreeNAS®, [Volume Manager](#) (page 115) can be used to create or extend ZFS pools. After a pool is created, it can be divided into dynamically-sized datasets or fixed-size zvols as needed. Datasets can be used to optimize storage for the type of data being stored as permissions and properties such as quotas and compression can be set on a per-dataset level. A zvol is essentially a raw, virtual block device which can be used for applications that need raw-device semantics such as iSCSI device extents.

ZFS supports real-time data compression. Compression happens when a block is written to disk, but only if the written data will benefit from compression. When a compressed block is accessed, it is automatically decompressed. Since compression happens at the block level, not the file level, it is transparent to any applications accessing the compressed data. By default, ZFS pools made using FreeNAS® version 9.2.1 or later will use the recommended LZ4 compression algorithm.

ZFS provides low-cost, instantaneous snapshots of the specified pool, dataset, or zvol. Due to COW, snapshots initially take no additional space. The size of a snapshot increases over time as changes to the files in the snapshot are written to disk.

Snapshots can be used to provide a copy of data at the point in time the snapshot was created. When a file is deleted, its disk blocks are added to the free list; however, the blocks for that file in any existing snapshots are not added to the free list until all referencing snapshots are removed. This makes snapshots a clever way to keep a history of files, useful for recovering an older copy of a file or a deleted file. For this reason, many administrators take snapshots often (e.g., every 15 minutes), store them for a period of time (e.g., for a month), and store them on another system. Such a strategy allows the administrator to roll the system back to a specific time. If there is a catastrophic loss, an off-site snapshot can restore the system up to the last snapshot interval (e.g., within 15 minutes of the data loss). Snapshots are stored locally but can also be replicated to a remote ZFS pool. During replication, ZFS does not do a byte-for-byte copy but instead converts a snapshot into a stream of data. This design means that the ZFS pool on the receiving end does not need to be identical and can use a different RAIDZ level, volume size, or compression settings.

ZFS boot environments provide a method for recovering from a failed upgrade. In FreeNAS®, a snapshot of the dataset the operating system resides on is automatically taken before an upgrade or a system update. This saved boot environment is automatically added to the GRUB boot loader. Should the upgrade or configuration change fail, simply reboot and select the previous boot environment from the boot menu. Users can also create their own boot environments in `System` → `Boot` as needed, for example before making configuration changes. This way, the system can be rebooted into a snapshot of the system that did not include the new configuration changes.

ZFS provides a write cache in RAM as well as a ZFS Intent Log (ZIL (https://blogs.oracle.com/realneel/entry/the_zfs_intent_log)). The ZIL is a storage area that temporarily holds *synchronous* writes until they are written to the ZFS pool (<https://pthree.org/2013/04/19/zfs-administration-appendix-a-visualizing-the-zfs-intent-log/>). Adding a fast (low-latency), power-protected SSD as a SLOG (Separate Log) device permits much higher performance. This is a necessity for NFS over ESXi, and highly recommended for database servers or other applications that depend on synchronous writes. More detail on SLOG benefits and usage is available in these blog and forum posts:

- [The ZFS ZIL and SLOG Demystified](http://www.freenas.org/blog/zfs-zil-and-slog-demystified/) (<http://www.freenas.org/blog/zfs-zil-and-slog-demystified/>)
- [Some insights into SLOG/ZIL with ZFS on FreeNAS®](https://forums.freenas.org/index.php?threads/some-insights-into-slog-zil-with-zfs-on-freenas.13633/) (<https://forums.freenas.org/index.php?threads/some-insights-into-slog-zil-with-zfs-on-freenas.13633/>)
- [ZFS Intent Log](http://nex7.blogspot.com/2013/04/zfs-intent-log.html) (<http://nex7.blogspot.com/2013/04/zfs-intent-log.html>)

Synchronous writes are relatively rare with SMB, AFP, and iSCSI, and adding a SLOG to improve performance of these protocols only makes sense in special cases. The `zilstat` utility can be run from `Shell` (page 290) to determine if the system will benefit from a SLOG. See [this website](http://www.richardelling.com/Home/scripts-and-programs-1/zilstat) (<http://www.richardelling.com/Home/scripts-and-programs-1/zilstat>) for usage information.

ZFS currently uses 16 GB of space for SLOG. Larger SSDs can be installed, but the extra space will not be used. SLOG devices cannot be shared between pools. Each pool requires a separate SLOG device. Bandwidth and throughput limitations require that a SLOG device must only be used for this single purpose. Do not attempt to add other caching functions on the same SSD, or performance will suffer.

In mission-critical systems, a mirrored SLOG device is highly recommended. Mirrored SLOG devices are required for ZFS pools at ZFS version 19 or earlier. ZFS pool version can be checked from the `Shell` (page 290) with `zpool get version poolname`. A version value of - means the ZFS pool is version 5000 (also known as Feature Flags) or later.

ZFS provides a read cache in RAM, known as the ARC, which reduces read latency. FreeNAS® adds ARC stats to `top(1)` (<http://www.freebsd.org/cgi/man.cgi?query=top>) and includes the `arc_summary.py` and `arcstat.py` tools for monitoring the efficiency of the ARC. If an SSD is dedicated as a cache device, it is known as an L2ARC (<https://blogs.oracle.com/brendan/entry/test>). Additional read data is cached here, which can increase random read performance. L2ARC does not reduce the need for sufficient RAM. In fact, L2ARC needs RAM to function. If there is not enough RAM for an adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of random reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 32 GB of RAM, and the size of an L2ARC should not exceed ten times the amount of RAM. In some cases, it may be more efficient to have two separate pools: one on SSDs for active data, and another on hard drives for rarely used content. After adding an L2ARC device, monitor its effectiveness using tools such as `arcstat`. To increase the size of an existing L2ARC, stripe another cache device with it. The GUI will always stripe L2ARC, not mirror it, as the contents of L2ARC

are recreated at boot. Failure of an individual SSD from an L2ARC pool will not affect the integrity of the pool, but may have an impact on read performance, depending on the workload and the ratio of dataset size to cache size. Note that dedicated L2ARC devices cannot be shared between ZFS pools.

ZFS was designed to provide redundancy while addressing some of the inherent limitations of hardware RAID such as the write-hole and corrupt data written over time before the hardware controller provides an alert. ZFS provides three levels of redundancy, known as RAIDZ, where the number after the RAIDZ indicates how many disks per vdev can be lost without losing data. ZFS also supports mirrors, with no restrictions on the number of disks in the mirror. ZFS was designed for commodity disks so no RAID controller is needed. While ZFS can also be used with a RAID controller, it is recommended that the controller be put into JBOD mode so that ZFS has full control of the disks.

When determining the type of ZFS redundancy to use, consider whether the goal is to maximize disk space or performance:

- RAIDZ1 maximizes disk space and generally performs well when data is written and read in large chunks (128K or more).
- RAIDZ2 offers better data availability and significantly better mean time to data loss (MTTDL) than RAIDZ1.
- A mirror consumes more disk space but generally performs better with small random reads. For better performance, a mirror is strongly favored over any RAIDZ, particularly for large, uncacheable, random read loads.
- Using more than 12 disks per vdev is not recommended. The recommended number of disks per vdev is between 3 and 9. With more disks, use multiple vdevs.
- Some older ZFS documentation recommends that a certain number of disks is needed for each type of RAIDZ in order to achieve optimal performance. On systems using LZ4 compression, which is the default for FreeNAS® 9.2.1 and higher, this is no longer true. See [ZFS RAIDZ stripe width, or: How I Learned to Stop Worrying and Love RAIDZ](http://blog.delphix.com/matt/2014/06/06/zfs-stripe-width/) (<http://blog.delphix.com/matt/2014/06/06/zfs-stripe-width/>) for details.

These resources can also help determine the RAID configuration best suited to your storage needs:

- [Getting the Most out of ZFS Pools](https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/) (<https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/>)
- [A Closer Look at ZFS, Vdevs and Performance](http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance) (<http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance>)

Warning: RAID AND DISK REDUNDANCY ARE NOT A SUBSTITUTE FOR A RELIABLE BACKUP STRATEGY. BAD THINGS HAPPEN AND A GOOD BACKUP STRATEGY IS STILL REQUIRED TO PROTECT VALUABLE DATA. See [Periodic Snapshot Tasks](#) (page 140) and [Replication Tasks](#) (page 142) to use replicated ZFS snapshots as part of a backup strategy.

ZFS manages devices. When an individual drive in a mirror or RAIDZ fails and is replaced by the user, ZFS adds the replacement device to the vdev and copies redundant data to it in a process called resilvering. Hardware RAID controllers usually have no way of knowing which blocks were in use and must copy every block to the new device. ZFS only copies blocks that are in use, reducing the time it takes to rebuild the vdev. Resilvering is also interruptable. After an interruption, resilvering resumes where it left off rather than starting from the beginning.

While ZFS provides many benefits, there are some caveats:

- At 90% capacity, ZFS switches from performance- to space-based optimization, which has massive performance implications. For maximum write performance and to prevent problems with drive replacement, add more capacity before a pool reaches 80%. If you are using iSCSI, it is recommended to not let the pool go over 50% capacity to prevent fragmentation issues.
- When considering the number of disks to use per vdev, consider the size of the disks and the amount of time required for resilvering, which is the process of rebuilding the vdev. The larger the size of the vdev, the longer the resilvering time. When replacing a disk in a RAIDZ, it is possible that another disk will fail before the resilvering process completes. If the number of failed disks exceeds the number allowed per vdev for the type of RAIDZ, the data in the pool will be lost. For this reason, RAIDZ1 is not recommended for drives over 1 TB in size.
- It is recommended to use drives of equal sizes when creating a vdev. While ZFS can create a vdev using disks of differing sizes, its capacity will be limited by the size of the smallest disk.

For those new to ZFS, the [Wikipedia entry on ZFS](https://en.wikipedia.org/wiki/Zfs) (<https://en.wikipedia.org/wiki/Zfs>) provides an excellent starting point to learn more about its features. These resources are also useful for reference:

- [FreeBSD ZFS Tuning Guide](https://wiki.FreeBSD.org/ZFSTuningGuide) (<https://wiki.FreeBSD.org/ZFSTuningGuide>)
- [ZFS Administration Guide](http://docs.oracle.com/cd/E19253-01/819-5461/index.html) (<http://docs.oracle.com/cd/E19253-01/819-5461/index.html>)
- [Becoming a ZFS Ninja \(video\)](https://blogs.oracle.com/video/entry/becoming_a_zfs_ninja) (https://blogs.oracle.com/video/entry/becoming_a_zfs_ninja)
- [Slideshow explaining VDev, zpool, ZIL and L2ARC and other newbie mistakes!](https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-and-other-newbie-mistakes!) (<https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-and-other-newbie-mistakes!>)
- [A Crash Course on ZFS](http://www.bsdnow.tv/tutorials/zfs) (<http://www.bsdnow.tv/tutorials/zfs>)
- [ZFS: The Last Word in File Systems - Part 1 \(video\)](https://www.youtube.com/watch?v=uT2i2ryhCio) (<https://www.youtube.com/watch?v=uT2i2ryhCio>)
- [The Zettabyte Filesystem](https://www.youtube.com/watch?v=ptY6-K78McY) (<https://www.youtube.com/watch?v=ptY6-K78McY>)

VMware's vStorage APIs for Array Integration, or VAAI, allows storage tasks such as large data moves to be offloaded from the virtualization hardware to the storage array. These operations are performed locally on the NAS without transferring bulk data over the network.

29.1 VAAI for iSCSI

VAAI for iSCSI supports these operations:

- Atomic Test and Set (ATS) allows multiple initiators to synchronize LUN access in a fine-grained manner rather than locking the whole LUN and preventing other hosts from accessing the same LUN simultaneously.
- Clone Blocks (XCOPY) copies disk blocks on the NAS. Copies occur locally rather than over the network. The operation is similar to [Microsoft ODX](https://technet.microsoft.com/en-us/library/hh831628) (<https://technet.microsoft.com/en-us/library/hh831628>).
- LUN Reporting allows a hypervisor to query the NAS to determine whether a LUN is using thin provisioning.
- Stun pauses running virtual machines when a volume runs out of space. The space issue can then be fixed and the virtual machines can continue rather than reporting write errors.
- Threshold Warning the system reports a warning when a configurable capacity is reached. In FreeNAS®, this threshold can be configured at the pool level when using zvols (see [Table 10.6](#)) or at the extent level (see [Table 10.11](#)) for both file- and device-based extents. Typically, the warning is set at the pool level, unless file extents are used, in which case it must be set at the extent level.
- Unmap informs FreeNAS® that the space occupied by deleted files should be freed. Without unmap, the NAS is unaware of freed space created when the initiator deletes files. For this feature to work, the initiator must support the unmap command.
- Zero Blocks or Write Same zeros out disk regions. When allocating virtual machines with thick provisioning, the zero write is done locally, rather than over the network. This makes virtual machine creation and any other zeroing of disk regions much quicker.

USING THE API

A [REST](https://en.wikipedia.org/wiki/Representational_state_transfer) (https://en.wikipedia.org/wiki/Representational_state_transfer) API is provided to be used as an alternate mechanism for remotely controlling a FreeNAS® system.

REST provides an easy-to-read, HTTP implementation of functions, known as resources, which are available beneath a specified base URL. Each resource is manipulated using the HTTP methods defined in [RFC 2616](https://tools.ietf.org/html/rfc2616.html) (<https://tools.ietf.org/html/rfc2616.html>), such as GET, PUT, POST, or DELETE.

As shown in [Figure 30.1](#), an online version of the API is available at api.freenas.org (<http://api.freenas.org>).




Fig. 30.1: API Documentation

The rest of this section shows code examples to illustrate the use of the API.

Note: Beginning with FreeNAS® 9.10.2, a new API has been added. The old API is still present for compatibility. Documentation for the new API is available on the FreeNAS® system at the `/api/docs/` URL. For example, if the FreeNAS® system is at IP address 192.168.1.119, enter `http://192.168.1.119/api/docs/` in a browser to see the API documentation.

30.1 A Simple API Example

The [api directory of the FreeNAS® github repository](https://github.com/freenas/freenas/tree/master/examples/api) (<https://github.com/freenas/freenas/tree/master/examples/api>) contains some API usage examples. This section provides a walk-through of the `newuser.py` script, shown below, as it provides a simple example that creates a user.

A FreeNAS® system running at least version 9.2.0 is required when creating a customized script based on this example. To test the scripts directly on the FreeNAS® system, create a user account and select an existing volume or dataset for the user's Home Directory. After creating the user, start the SSH service using **Services → Control Services**. That user will now be able to `ssh` to the IP address of the FreeNAS® system to create and run scripts. Alternately, scripts can be tested on any system with the required software installed as shown in the previous section.

To customize this script, copy the contents of this example into a filename that ends in `.py`. The text that is highlighted in red below can be modified in the new version to match the needs of the user being created. The text in black should not be changed. After saving changes, run the script by typing `python scriptname.py`. If all goes well, the new user account will appear in **Account → Users → View Users** in the FreeNAS® GUI.

Here is the example script with an explanation of the line numbers below it.

```

1 import json
2 import requests
3 r = requests.post(
4     'https://freenas.mydomain/api/v1.0/account/users/',
5     auth=('root', 'freenas'),
6     headers={'Content-Type': 'application/json'},
7     verify=False,
8     data=json.dumps({
9         'bsdusr_uid': '1100',
10        'bsdusr_username': 'myuser',
11        'bsdusr_mode': '755',
12        'bsdusr_creategroup': 'True',
13        'bsdusr_password': '12345',
14        'bsdusr_shell': '/usr/local/bin/bash',
15        'bsdusr_full_name': 'Full Name',
16        'bsdusr_email': 'name@provider.com',
17    })
18 )
19 print r.text

```

Where:

Lines 1-2: import the Python modules used to make HTTP requests and handle data in JSON format.

Line 4: replace `freenas.mydomain` with the Hostname value in **System → System Information**. Note that the script will fail if the machine running it is not able to resolve that hostname. Change `https` to `http` to use HTTP rather than HTTPS to access the FreeNAS® system.

Line 5: replace `freenas` with the password used to access the FreeNAS® system.

Line 7: if you are using HTTPS and want to force validation of the SSL certificate, change `False` to `True`.

Lines 8-16: set the values for the user being created. The [Users resource](http://api.freenas.org/resources/account.html#users) (<http://api.freenas.org/resources/account.html#users>) describes this in more detail. Allowed parameters are listed in the JSON Parameters section of that resource. Since this resource creates a FreeBSD user, the values entered must be valid for a FreeBSD user account. [Table 30.1](#) summarizes acceptable values. This resource uses JSON, so the boolean values are `True` or `False`.

Table 30.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
<code>bsdusr_username</code>	string	maximum 32 characters, though a maximum of 8 is recommended for interoperability; can include numerals but cannot include a space
Continued on next page		

Table 30.1 – continued from previous page

JSON Parameter	Type	Description
bsdusr_full_name	string	may contain spaces and uppercase characters
bsdusr_password	string	can include a mix of upper and lowercase letters, characters, and numbers
bsdusr_uid	integer	by convention, user accounts have an ID greater than 1000 with a maximum allowable value of 65,535
bsdusr_group	integer	if bsdusr_creategroup is set to False, specify the numeric ID of the group to create
bsdusr_creategroup	boolean	if set to True, a primary group with the same numeric ID as bsdusr_uid will be created automatically
bsdusr_mode	string	sets default numeric UNIX permissions of user's home directory
bsdusr_shell	string	specify full path to a UNIX shell that is installed on the system
bsdusr_password_disabled	boolean	if set to True, user is not allowed to log in
bsdusr_locked	boolean	if set to True, user is not allowed to log in
bsdusr_sudo	boolean	if set to True, <code>sudo</code> is enabled for the user
bsdusr_sshpubkey	string	contents of SSH authorized keys file

Note: When using boolean values, JSON returns raw lowercase values but Python uses uppercase values. So use True or False in Python scripts even though the example JSON responses in the API documentation are displayed as true or false.

30.2 A More Complex Example

This section provides a walk-through of a more complex example found in the `startup.py` script. Use the searchbar within the API documentation to quickly locate the JSON parameters used here. This example defines a class and several methods to create a ZFS volume, create a ZFS dataset, share the dataset over CIFS, and enable the CIFS service. Responses from some methods are used as parameters in other methods. In addition to the import lines seen in the previous example, two additional Python modules are imported to provide parsing functions for command line arguments:

```
import argparse
import sys
```

It then creates a Startup class which is started with the hostname, username, and password provided by the user via the command line:

```
1 class Startup(object):
2     def __init__(self, hostname, user, secret):
3         self._hostname = hostname
4         self._user = user
5         self._secret = secret
6         self._ep = 'http://%s/api/v1.0' % hostname
7     def request(self, resource, method='GET', data=None):
8         if data is None:
9             data = ''
10        r = requests.request(
11            method,
12            '%s/%s/' % (self._ep, resource),
13            data=json.dumps(data),
14            headers={'Content-Type': 'application/json'},
15            auth=(self._user, self._secret),
16        )
17    if r.ok:
18        try:
19            return r.json()
20        except:
```

```

21         return r.text
22     raise ValueError(r)

```

A `get_disks` method is defined to get all the disks in the system as a `disk_name` response. The `create_pool` method uses this information to create a ZFS pool named `tank` which is created as a `stripe`. The `volume_name` and `layout` JSON parameters are described in the “Storage Volume” resource of the API documentation.

```

1  def _get_disks(self):
2      disks = self.request('storage/disk')
3      return [disk['disk_name'] for disk in disks]
4
5  def create_pool(self):
6      disks = self._get_disks()
7      self.request('storage/volume', method='POST', data={
8          'volume_name': 'tank',
9          'layout': [
10             {'vdevtype': 'stripe', 'disks': disks},
11         ],
12     })

```

The `create_dataset` method is defined which creates a dataset named `MyShare`:

```

1  def create_dataset(self):
2      self.request('storage/volume/tank/datasets', method='POST', data={
3          'name': 'MyShare',
4      })

```

The `create_cifs_share` method is used to share `/mnt/tank/MyShare` with guest-only access enabled. The `cifs_name`, `cifs_path`, `cifs_guestonly` JSON parameters, as well as the other allowable parameters, are described in the “Sharing CIFS” resource of the API documentation.

```

1  def create_cifs_share(self):
2      self.request('sharing/cifs', method='POST', data={
3          'cifs_name': 'My Test Share',
4          'cifs_path': '/mnt/tank/MyShare',
5          'cifs_guestonly': True
6      })

```

Finally, the `service_start` method enables the CIFS service. The `srv_enable` JSON parameter is described in the Services resource.

```

1  def service_start(self, name):
2      self.request('services/services/%s' % name, method='PUT', data={
3          'srv_enable': True,
4      })
5

```

Symbols

802.1Q, 112

A

Add Group, 49

Add Jail, 253

Add User, 52

Adding Devices to a VM, 271

AFP, 170, 211

Alert, 296

Alert Services, 74

API, 326

Apple Filing Protocol, 170, 211

arcstat, 308

Autotune, 65

B

Boot Environments, 59

Burn ISO, 11

C

CA, 75

Certificate Authority, 75

Certificates, 78

Checksum, 11

CIFS, 185, 230

Cloud Credentials, 73

Cloud Sync, 84

Compression, 125

Create Dataset, 122

Create Group, 49

Create Jail, 253

Create User, 52

Creating VMs, 270

Cron Jobs, 89

D

DC, 213

DDNS, 215

Deduplication, 124

Delete Group, 50

Delete User, 54

Deleting VMs, 276

Dell PERC H330, 8

Dell PERC H730, 8

Dmidecode, 316

Domain Controller, 213

Download, 11

Dynamic DNS, 215

E

Email, 65

Encryption, 117

EtherChannel, 107

F

File Transfer Protocol, 216

Forums, 299

freenas-debug, 315

FTP, 216

G

Getting FreeNAS\ :sup:'@', 11

Groups, 48

Guide, 295

H

Hardware Recommendations, 6

Highpoint RAID, 8

Hot Spares, 140

I

Install, 13

Internet Small Computer System Interface, 195

iocage, 267

IOzone, 306

Iperf, 302

IRC, 300

iSCSI, 195

ISO, 11

J

Jails, 250

L

LACP, 107

LAGG, 107

Link Aggregation, 107

Link Layer Discovery Protocol, 221

LLDP, 221

Localize, 318

Log Out, 291

M

MegaCli, 314

Minio, 227

Mirroring the Boot Device, 62

Multiple Boot Environments, 59

N

Netdata, 222

Netperf, 305

Network File System, 177, 223

Network Settings, 101

New Group, 49

New Jail, 253

New User, 52

NFS, 177, 223

P

Path and Name Lengths, 6

Periodic Snapshot, 140

Plugin, 243

Professional Support, 301

R

Reboot, 292

Remove Group, 50

Remove User, 54

Replace Failed Drive, 137

Replication, 142

Reporting, 279

RFC

 RFC 2616, 327

 RFC 3721, 197

Route, 112

Rsync, 225

Rsync Tasks, 92

Running VMs, 276

S

S.M.A.R.T., 228

S.M.A.R.T. Tests, 99

S3, 227

Samba, 185, 230

SCP, 237

Scrub, 152

Secure Copy, 237

Secure Shell, 236

Services, 208

Shadow Copies, 194

Shell, 289

Shutdown, 293

Simple Network Management Protocol, 234

SMB, 185, 230

Snapshot, 140

Snapshots, 155

SNMP, 234

Spares, 140

SSH, 236

Start Service, 209

Static Route, 112

Stop Service, 209

Support, 81, 294

System Dataset, 67

T

Tasks, 83

TFTP, 238

Time Machine, 173

tmux, 315

Translate, 318

Translation, 318

Trivial File Transfer Protocol, 238

Trunking, 112

Tunables, 67

tw_cli, 313

U

Uninterruptible Power Supply, 239

Upgrade, 21

Upgrade ZFS Pool, 27

UPS, 239

USB Stick, 11

Users, 50

V

VAAI, 325

VAAI for iSCSI, 326

Virtualization, 28

VLAN, 112

VM, 28

VMs, 269

VMware Snapshot, 157

Volumes, 115

W

WebDAV, 184, 242

Windows File Share, 230

Windows Shares, 185

Wizard, 281

Z

ZVOL, 125