

FreeNAS® 11.2-RELEASE User Guide

December 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 11.2

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

CONTENTS

Welcome	8
Typographic Conventions	10
1 Introduction	11
1.1 New Features in 11.2	11
1.2 Path and Name Lengths	14
1.3 Hardware Recommendations	15
1.3.1 RAM	15
1.3.2 The Operating System Device	16
1.3.3 Storage Disks and Controllers	17
1.3.4 Network Interfaces	17
1.4 Getting Started with ZFS	18
2 Installing and Upgrading	19
2.1 Getting FreeNAS®	19
2.2 Preparing the Media	19
2.2.1 On FreeBSD or Linux	20
2.2.2 On Windows	20
2.2.3 On macOS	20
2.3 Performing the Installation	21
2.4 Installation Troubleshooting	29
2.5 Upgrading	30
2.5.1 Caveats	30
2.5.2 Initial Preparation	30
2.5.3 Upgrading Using the ISO	31
2.5.4 Upgrading From the Web Interface	33
2.5.5 If Something Goes Wrong	33
2.5.6 Upgrading a ZFS Pool	35
2.6 Virtualization	37
2.6.1 VirtualBox	37
2.6.2 VMware ESXi	48
3 Booting	55
3.1 Obtaining an IP Address	56
3.2 Logging In	57
4 Settings	59
4.1 Edit root Account	59
4.2 Change Password	59
4.3 Preferences	59
4.3.1 Web Interface Preferences	59
4.3.2 Themes	60
4.3.2.1 Theme Selector	60
4.3.2.2 Create New Themes	61

4.4	About	63
4.5	Legacy Web Interface	63
5	Accounts	64
5.1	Groups	64
5.2	Users	67
6	System	71
6.1	General	71
6.2	NTP Servers	73
6.3	Boot Environments	75
6.3.1	Mirroring the Boot Device	78
6.4	Advanced	79
6.4.1	Autotune	81
6.4.2	Self-Encrypting Drives	81
6.5	Email	83
6.6	System Dataset	84
6.7	Alert Services	85
6.8	Alert Settings	87
6.9	Cloud Credentials	88
6.10	Tunables	91
6.11	Update	94
6.11.1	Preparing for Updates	94
6.11.2	Updates and Trains	94
6.11.3	Checking for Updates	95
6.11.4	Saving the Configuration File	96
6.11.5	Applying Updates	97
6.11.6	Manual Updates	97
6.12	CAs	97
6.13	Certificates	101
6.14	Support	106
7	Tasks	108
7.1	Cron Jobs	108
7.2	Init/Shutdown Scripts	110
7.3	Rsync Tasks	111
7.3.1	Rsync Module Mode	113
7.3.2	Rsync over SSH Mode	114
7.4	S.M.A.R.T. Tests	116
7.5	Periodic Snapshot Tasks	118
7.6	Replication Tasks	120
7.6.1	Examples: Common Configuration	120
7.6.1.1	<i>Alpha</i> (Source)	120
7.6.1.2	<i>Beta</i> (Destination)	121
7.6.2	Example: FreeNAS® to FreeNAS® Semi-Automatic Setup	121
7.6.3	Example: FreeNAS® to FreeNAS® Dedicated User Replication	124
7.6.4	Example: FreeNAS® to FreeNAS® or Other Systems, Manual Setup	125
7.6.4.1	Encryption Keys	126
7.6.5	Replication Options	128
7.6.6	Replication Encryption	130
7.6.7	Limiting Replication Times	130
7.6.8	Troubleshooting Replication	130
7.6.8.1	SSH	131
7.6.8.2	Compression	131
7.6.8.3	Manual Testing	131
7.7	Resilver Priority	131
7.8	Scrub Tasks	132
7.9	Cloud Sync Tasks	134

7.9.1	Cloud Sync Example	137
8	Network	140
8.1	Global Configuration	140
8.2	Interfaces	142
8.3	IPMI	144
8.4	Link Aggregations	146
8.4.1	LACP, MPIO, NFS, and ESXi	146
8.4.2	Creating a Link Aggregation	147
8.4.3	Link Aggregation Options	150
8.5	Static Routes	151
8.6	VLANs	152
9	Storage	154
9.1	Pools	154
9.1.1	Creating Pools	154
9.1.2	Managing Encrypted Pools	157
9.1.3	Adding Cache or Log Devices	160
9.1.4	Removing Cache or Log Devices	160
9.1.5	Adding Spare Devices	160
9.1.6	Extending a Pool	161
9.1.7	Detaching a Pool	161
9.1.8	Importing a Pool	163
9.1.9	Viewing Pool Scrub Status	165
9.1.10	Adding Datasets	166
9.1.10.1	Deduplication	168
9.1.10.2	Compression	168
9.1.11	Adding Zvols	169
9.1.12	Setting Permissions	170
9.1.13	View Multipaths	172
9.2	Snapshots	172
9.2.1	Browsing a Snapshot Collection	174
9.3	VMware-Snapshots	175
9.4	Disks	176
9.4.1	Replacing a Failed Disk	178
9.4.1.1	Replacing an Encrypted Disk	180
9.4.1.2	Removing a Log or Cache Device	180
9.4.2	Replacing Disks to Grow a Pool	181
9.5	Importing a Disk	181
10	Directory Services	183
10.1	Active Directory	183
10.1.1	Troubleshooting Tips	187
10.1.2	If the System Does not Join the Domain	187
10.2	LDAP	188
10.3	NIS	191
10.4	Kerberos Realms	192
10.5	Kerberos Keytabs	193
10.6	Kerberos Settings	194
11	Sharing	195
11.1	Apple (AFP) Shares	196
11.1.1	Creating AFP Guest Shares	198
11.1.2	Creating Authenticated and Time Machine Shares	200
11.2	Unix (NFS) Shares	203
11.2.1	Example Configuration	207
11.2.2	Connecting to the Share	207
11.2.2.1	From BSD or Linux	207

11.2.2.2 From Microsoft	208
11.2.2.3 From macOS	208
11.2.3 Troubleshooting NFS	209
11.3 WebDAV Shares	210
11.4 Windows (SMB) Shares	211
11.4.1 Configuring Unauthenticated Access	217
11.4.2 Configuring Authenticated Access Without a Domain Controller	218
11.4.3 Configuring Shadow Copies	221
11.5 Block (iSCSI)	223
11.5.1 Target Global Configuration	224
11.5.2 Portals	225
11.5.3 Initiators	226
11.5.4 Authorized Accesses	228
11.5.5 Targets	229
11.5.6 Extents	231
11.5.7 Associated Targets	233
11.5.8 Connecting to iSCSI	234
11.5.9 Growing LUNs	235
11.5.9.1 Zvol Based LUN	235
11.5.9.2 File Extent Based LUN	236
12 Services	237
12.1 Configure Services	237
12.2 AFP	238
12.2.1 Troubleshooting AFP	239
12.3 Domain Controller	240
12.3.1 Samba Domain Controller Backup	241
12.4 Dynamic DNS	242
12.5 FTP	243
12.5.1 Anonymous FTP	246
12.5.2 FTP in chroot	247
12.5.3 Encrypting FTP	248
12.5.4 Troubleshooting FTP	248
12.6 iSCSI	248
12.7 LLDP	248
12.8 Netdata	249
12.9 NFS	250
12.10 Rsync	252
12.10.1 Configure Rsyncd	252
12.10.2 Rsync Modules	253
12.11 S3	255
12.12 S.M.A.R.T.	256
12.13 SMB	258
12.13.1 Troubleshooting SMB	260
12.14 SNMP	261
12.15 SSH	263
12.15.1 SCP Only	264
12.15.2 Troubleshooting SSH	265
12.16 TFTP	265
12.17 UPS	267
12.17.1 Multiple Computers with One UPS	269
12.18 WebDAV	269
13 Plugins	271
13.1 Install	271
13.2 Update	275
13.3 Delete	275

13.4	Create a Plugin	275
13.4.1	Test a Plugin	279
13.5	Official Plugins	279
14	Jails	281
14.1	Jail Storage	281
14.2	Creating Jails	282
14.2.1	Jail Wizard	282
14.2.2	Advanced Jail Creation	283
14.3	Managing Jails	289
14.3.1	Jail Updates and Upgrades	291
14.3.2	Accessing a Jail Using SSH	291
14.3.3	Additional Storage	293
14.4	Jail Software	296
14.4.1	Installing FreeBSD Packages	296
14.4.2	Compiling FreeBSD Ports	297
14.4.3	Starting Installed Software	300
14.5	Using iocage	302
14.5.1	Managing iocage Jails	302
15	Reporting	305
16	Virtual Machines	307
16.1	Creating VMs	307
16.2	Adding Devices to a VM	310
16.2.1	CD-ROM Devices	311
16.2.2	NIC (Network Interfaces)	312
16.2.3	Disk Devices	313
16.2.4	Raw Files	314
16.2.5	VNC Interface	315
16.2.6	Virtual Serial Ports	317
16.3	Running VMs	317
16.4	Deleting VMs	319
16.5	Docker Hosts	319
16.5.1	Docker Host Requirements	319
16.5.2	Creating Docker Host	319
16.5.3	Start the Docker Host	322
16.5.4	SSH into the Docker Host	322
16.5.5	Installing and Configuring the Rancher Server	323
17	Display System Processes	324
18	Shell	325
19	Log Out, Restart, or Shut Down	327
19.1	Log Out	327
19.2	Restart	327
19.3	Shut Down	328
20	Alert	330
21	Support Resources	332
21.1	Website and Social Media	332
21.2	Forums	332
21.3	IRC	333
21.4	Videos	333
21.5	Professional Support	333

22 Contributing to FreeNAS®	334
22.1 Translation	334
22.1.1 Translate with GitHub	335
22.1.2 Download and Translate Offline	336
22.1.3 Translation Pull Requests	337
23 Command Line Utilities	338
23.1 lperf	338
23.2 Netperf	341
23.3 IOzone	342
23.4 arcstat	344
23.5 tw_cli	349
23.6 MegaCli	350
23.7 freenas-debug	351
23.8 tmux	351
23.9 Dmidecode	352
23.10 Midnight Commander	353
24 ZFS Primer	354
25 VAAI	358
25.1 VAAI for iSCSI	358
26 Using the API	359
26.1 A Simple API Example	360
26.2 A More Complex Example	361
Index	363

Welcome

This Guide covers the installation and use of FreeNAS® 11.2.

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® 11.2 User Guide uses these typographic conventions:

Table 1: Text Format Examples

Item	Visual Example
Graphical elements: buttons, icons, fields, columns, and boxes	Click the <i>Import CA</i> button.
Menu selections	Select <i>System</i> → <i>Information</i> .
Commands	Use the <code>scp</code> command.
File names and pool 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 <code>127.0.0.1</code> in the address field.

Table 2: FreeNAS® Icons

Icon	Usage
<i>ADD</i>	Add a new item.
 (Settings)	Show a settings menu.
 (Options)	Show an Options menu.
 (Power)	Show a power options menu.
 (Show)	Reveal characters in a password field.
 (Hide)	Hide characters in a password field.
 (Configure)	Edit settings.
 (Launch)	Launch a service.
 (Start)	Start jails.
 (Stop)	Stop jails.
 (Update)	Update jails.
 (Delete)	Delete jails.

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 11.2

FreeNAS® 11.2 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 94) to this release in order to take advantage of these improvements and bug fixes.

These major features are new in this version:

- The login screen defaults to the new, Angular-based UI. Users who wish to continue to use the classic UI can select “Legacy UI” in the login screen.
- Beginning with this release, the screenshots that appear in the [published version of the Guide](http://doc.freenas.org/11.2/freenas.html) (<http://doc.freenas.org/11.2/freenas.html>) and in the *Guide* option within the new UI are for the new UI. However, users who click on the *Guide* icon while logged into the classic UI will continue to see screenshots for the old UI. The availability of both versions of the Guide is to assist users as they become familiar with the new UI during the transition period before the classic UI is deprecated in a future release.
- The rewrite from the old API to the new middlewared continues. Once the rewrite is complete, api.freenas.org (<http://api.freenas.org/>) will be deprecated and replaced by the new API documentation. In the mean time, to see the API documentation for the new middleware, log into the new UI, click on the URL for the FreeNAS system in your browser’s location bar, and add `/api/docs` to the end of that URL.
- The boot loader has changed from GRUB to the native FreeBSD boot loader. This should resolve several issues that some users experienced with GRUB. GRUB was introduced as a temporary solution until the FreeBSD boot loader had full support for boot environments, which it now has.
- The [Plugins](#) (page 271) and [Jails](#) (page 281) backend has switched from `warden` to `iocage` and `warden` will no longer receive bug fixes. The new UI will automatically use `iocage` to create and manage [Plugins](#) (page 271) and [Jails](#) (page 281). Users are encouraged to recreate any existing [Plugins](#) (page 271) and [Jails](#) (page 281) using the new UI to ensure that they are running the latest supported application versions.
- [Plugins](#) (page 271) have switched to FreeBSD 11.2-RELEASE and all Plugins have been rebuilt for this version.
- [Virtual Machines](#) (page 307) are more crash-resistant. When a guest is started, the amount of available memory is checked and an initialization error will occur if there is insufficient system resources. There is an option to overcommit memory to the guest when it is started, but this is not recommended for normal use. When a guest is stopped, its resources are returned to the system. In addition, the UEFI boot menu fix allows Linux kernels 4.15 and higher to boot properly.
- [Cloud Sync Tasks](#) (page 134) provides configuration options to encrypt data before it is transmitted and to keep it in the encrypted format while stored on the cloud. The filenames can also be encrypted.

- Preliminary support has been added for *Self-Encrypting Drives* (page 81) (SEDs).

This software has been added or updated:

- The base operating system is the STABLE branch of [FreeBSD 11.2](https://www.freebsd.org/releases/11.2R/announce.html) (<https://www.freebsd.org/releases/11.2R/announce.html>), which brings in many updated drivers and bug fixes. This branch has been patched to include the FreeBSD security advisories up to [FreeBSD-SA-18:13.nfs](https://www.freebsd.org/security/advisories/FreeBSD-SA-18:13.nfs) (<https://www.freebsd.org/security/advisories/FreeBSD-SA-18:13.nfs.asc>).
- OpenZFS is up-to-date with Illumos and slightly ahead due to support for sorted scrubs which were ported from ZFS on Linux. Notable improvements include channel programs, data disk removal, more resilient volume import, the ability to import a pool with missing vdevs, pool checkpoints, improved compressed ARC performance, and ZIL batching. As part of this change, the default ZFS indirect block size is reduced to 32 KiB from 128 KiB. Note that many of these improvements need further testing so have not yet been integrated into the UI.
- The IPsec kernel module has been added. It can be manually loaded with `kldload ipsec`.
- Support for eMMC flash storage has been added.
- The `em` (<https://www.freebsd.org/cgi/man.cgi?query=em&apropos=0&sektion=4>), `igb` (<https://www.freebsd.org/cgi/man.cgi?query=igb&apropos=0&sektion=4>), `ixgbe` (<https://www.freebsd.org/cgi/man.cgi?query=ixgbe&apropos=0&sektion=4>) and `ixl` (<https://www.freebsd.org/cgi/man.cgi?query=ixl&apropos=0&sektion=4>) Intel drivers have been patched to resolve a performance degradation issue that occurs when the MTU is set to 9000 (9k jumbo clusters). Before configuring 9k jumbo clusters for `cxgbe` (<https://www.freebsd.org/cgi/man.cgi?query=cxgbe&apropos=0&sektion=4>) create a *Tunables* (page 91) with a *Variable* of `hw.cxgbe.largest_rx_cluster`, a *Type of Loader*, and a *Value* of 4096. The `cxgbe` (<https://www.freebsd.org/cgi/man.cgi?query=cxgbe&apropos=0&sektion=4>) driver does not support jumbo clusters and should not use an MTU greater than 4096.
- The `bnxt` (<https://www.freebsd.org/cgi/man.cgi?query=bnxt>) driver has been added which provides support for Broadcom NetXtreme-C and NetXtreme-E Ethernet drivers.
- The `vt terminal` (<https://www.freebsd.org/cgi/man.cgi?query=vt&sektion=4&manpath=FreeBSD+11.2-RELEASE+and+Ports>) is now used by default and the `syscons` terminal is removed from the kernel.
- `ncdu` (<https://dev.yorhel.nl/ncdu>) has been added to the base system. This CLI utility can be used to analyze disk usage from the console or an SSH session.
- `drm-next-kmod` (<https://www.freshports.org/graphics/drm-next-kmod/>) has been added to the base system, adding support for UTF-8 fonts to the console for Intel graphic cards.
- Netatalk has been updated to the 3.1.12 development version which addresses known issues with Time Machine time-outs.
- Samba 4.7 has been patched to address the latest round of *security vulnerabilities* (https://www.samba.org/samba/latest_news.html#4.9.3).
- `rsync` has been updated to *version 3.1.3* (<https://download.samba.org/pub/rsync/src/rsync-3.1.3-NEWS>).
- `rclone` has been updated to *version 1.44* (<https://rclone.org/changelog/#v1-44-2018-10-15>).
- Minio has been updated to *version 2018-04-04T05* (<https://github.com/minio/minio/releases/tag/RELEASE.2018-04-04T05-20-54Z>).
- Netdata has been updated to *version 1.10.1* (<https://github.com/firehol/netdata/releases/tag/v1.10.0>).
- `iocage` has been synced with upstream as of October 3, providing many bug fixes and improved IPv6 support.
- RancherOS has been updated to *version 1.4.2* (<https://github.com/rancher/os/releases/tag/v1.4.2>).
- `zsh` (<http://www.zsh.org/>) is the root shell for new installations. Upgrades will continue to use the `csh` shell as the default root shell.
- `xattr` (<https://github.com/xattr/xattr>) has been added to the base system and can be used to modify file extended attributes from the command line. Type `xattr -h` to view the available options.
- `convmv` (<https://www.j3e.de/linux/convmv/man/>) has been added to the base system and can be used to convert the encoding of filenames from the command line. Type `convmv` to view the available options.
- The `cloneacl` CLI utility has been added. It can be used to quickly clone a complex ACL recursively to or from an existing share. Type `cloneacl` for usage instructions.

- These switches have been added to [freenas-debug](#) (page 351): `-M` for dumping SATADOM info and `-Z` to delete old debug information. The `-G` switch has been removed as the system no longer uses GRUB. The `-J` switch has been removed and the `-j` switch has been reworked to show iocage jail information instead of Warden.
- These switches have been added to [arcstat](#) (page 344): `-a` for displaying all available statistics and `-p` for displaying raw numbers without suffixes.

These screen options have changed:

- The *ATA Security User*, *SED Password*, and *Reset SED Password* fields have been added to *System* → *Advanced*.
- The *Enable Console Screensaver* field has been removed from *System* → *Advanced*.
- The *Enable automatic upload of kernel crash dumps and daily telemetry* checkbox has been removed from *System* → *Advanced*.
- The *Enable Power Saving Daemon* option has been removed from *System* → *Advanced*.
- *Alert Settings* has been added to *System* and can be used to list the available alert conditions and to configure the notification frequency on a per-alert basis.
- These [Cloud Credentials](#) (page 88) have been added to *System* → *Cloud Credentials*: Amazon Cloud Drive, Box, Dropbox, FTP, Google Drive, HTTP, Hubic, Mega, Microsoft OneDrive, pCloud, SFTP, WebDAV, and Yandex.
- The *Team Drive ID* field has been added to *System* → *Cloud Credentials* → *Add* and appears when *Google Drive* is the *Provider*.
- The *Endpoint URL* has been added to *System* → *Cloud Credentials* → *Add Cloud Credential* but only appears when *Amazon S3* is selected as the *Provider*. This can be used to configure a connection to another S3-compatible service, such as Wasabi.
- *Drive Account Type* and *Drive ID* has been added to *System* → *Cloud Credentials* → *Add Cloud Credential*. These fields appear when *Microsoft OneDrive* is selected as the *Provider*.
- The *Automatically check for new updates* option in *System* → *Update* has been renamed to *Check for Updates Daily and Download if Available*.
- The *Train* selector in *System* → *Update* has been changed so that only allowable trains are displayed in the drop-down menu. Each train option has an expanded description.
- There is now an option to add a prompt to save a copy of the system configuration and include the *Password Secret Seed* before doing a system upgrade. This popup can be enabled by going to *i18n* (Settings) → *Preferences* and unsetting *Enable "Save Configuration" Dialog Before Upgrade*.
- The *Container*, *Remote encryption*, *Filename encryption*, *Encryption password*, and *Encryption salt* fields have been added to *Tasks* → *Cloud Sync Tasks* → *Add Cloud Sync*.
- The *NIC* and *Interface Name* fields in *Network* → *Interfaces* → *Add Interface* are preconfigured with the web interface NIC settings when configuring the first interface. A warning is shown when a user attempts to configure a different interface before the web interface NIC.
- The *Block size* field in *Storage* → *Pools* → *Add Zvol* → *ADVANCED MODE* no longer allows choosing sizes smaller than 4K. This is to prevent performance issues from setting a block size that is too small for efficient processing.
- The *Exec* field has been added to *Storage* → *Pools* → *Add Dataset* → *ADVANCED MODE*. The *Record Size* field no longer allows choosing sizes smaller than 4K. This is to prevent performance issues from setting a block size that is too small for efficient processing.
- The *Password for SED* column has been added to *Storage* → *Disks*.
- The *MSDOSFS locale* drop-down menu has been added to *Storage* → *Import Disk*.
- The *User Base* and *Group Base* fields have been removed from *Directory Services* → *Active Directory* → *Advanced Mode*.
- The *Enable home directories*, *Home directories*, *Home share name*, and *Home Share Time Machine* fields have been removed from *Services* → *AFP* and the *Time Machine Quota* field has been removed from *Sharing* → *Apple (AFP) Shares*. These fields have been replaced by *Sharing* → *Apple (AFP) Shares* → *Use as home share*.
- The *Umask* field in *Services* → *TFTP* has changed to a *File Permissions* selector.

- The *Hostname* field has been added to *Services* → *UPS*. This field replaces the *Port* field when a *UPS Driver* with `snmp` is selected.
- The BitTorrent Sync plugin has been renamed to Resilio Sync.
- Disk temperature graphs have been added to *Reporting* → *Disk*. This category has been reworked to allow the user to choose the devices and metrics before graphs are displayed.
- Uptime graphs have been removed from the *Reporting* → *System* tab.
- *Virtual Machines* → *Device* add and edit forms now have a *Device Order* field to set boot priority for VM devices.

Note: To keep lists aligned when using zoom in Firefox, ensure *View* → *Zoom* → *Zoom Text Only* is not set.

1.2 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 (<i>PATH_MAX</i>). 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 (<i>NAME_MAX</i>).
Mounted Filesystem Paths	88 bytes	Mounted filesystem path length (<i>MNAMELEN</i>). Longer paths can prevent a device from being mounted.
Device Filesystem Paths	63 bytes	<code>devfs(8)</code> (https://www.freebsd.org/cgi/man.cgi?query=devfs) device path lengths (<i>SPECMNAMELEN</i>). Longer paths can prevent a device from being created.

Note: 88 bytes is equal to 88 ASCII characters. The number of characters varies when using Unicode.

Warning: If the mounted path length for a snapshot exceeds 88 bytes, the data in the snapshot is safe but inaccessible. When the mounted path length of the snapshot is less than the 88 byte limit, the data will be accessible again.

The 88 byte limit affects automatic and manual snapshot mounts in slightly different ways:

- **Automatic mount:** ZFS temporarily mounts a snapshot whenever a user attempts to view or search the files within the snapshot. The mountpoint used will be in the hidden directory `.zfs/snapshot/name` within the same ZFS dataset. For example, the snapshot `mypool/dataset/snap1@snap2` is mounted at `/mnt/mypool/dataset/.zfs/snapshot/snap2/`. If the length of this path exceeds 88 bytes the snapshot will not be automatically mounted by ZFS and the snapshot contents will not be visible or searchable. This can be resolved by renaming the ZFS pool or dataset containing the snapshot to shorter names (`mypool` or `dataset`), or by shortening the second part of the snapshot name (`snap2`),

so that the total mounted path length does not exceed 88 bytes. ZFS will automatically perform any necessary unmount or remount of the file system as part of the rename operation. After renaming, the snapshot data will be visible and searchable again.

- **Manual mount:** The same example snapshot is mounted manually from the [Shell](#) (page 325) with `mount -t zfs mypool/dataset/snap1@snap2 /mnt/mymountpoint`. The path `/mnt/mountpoint/` must not exceed 88 bytes, and the length of the snapshot name is irrelevant. When renaming a manual mountpoint, any object mounted on the mountpoint must be manually unmounted with the `umount` command before renaming the mountpoint. It can be remounted afterwards.

Note: A snapshot that cannot be mounted automatically by ZFS can still be mounted manually from the [Shell](#) (page 325) with a shorter mountpoint path. This makes it possible to mount and access snapshots that cannot be accessed automatically in other ways, such as from the web interface or from features such as “File History” or “Versions”.

1.3 Hardware Recommendations

FreeNAS® 11.2 is based on FreeBSD 11.2 and supports the same hardware found in the [FreeBSD Hardware Compatibility List](#) (<https://www.freebsd.org/releases/11.2R/hardware.html>). Supported processors are listed in section 2.1 [amd64](#) (<https://www.freebsd.org/releases/11.2R/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/>) 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/>) gives detailed recommendations for system components, with the [FreeNAS® Quick Hardware Guide](#) (<https://forums.freenas.org/index.php?resources/freenas%C2%AE-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/>) has detailed instructions on testing new hardware.

1.3.1 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>) provide anecdotal evidence from users on how much performance can be gained by adding more RAM.

General guidelines for RAM:

- **A minimum of 8 GiB 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 GiB of RAM per terabyte of disk capacity.

- To use Active Directory with many users, add an additional 2 GiB of RAM for the winbind internal cache.
- For iSCSI, install at least 16 GiB of RAM if performance is not critical, or at least 32 GiB of RAM if good performance is a requirement.
- [Jails](#) (page 281) 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 307) 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 GiB of RAM requires an additional 2 GiB 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 GiB 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 GiB 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 GiB of RAM because of the abundance of information that FreeNAS® may not behave properly with less memory.

1.3.2 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 GiB. That does not provide much room. The *recommended* minimum is 16 GiB. This provides room for the operating system and several boot environments created by updates. More space provides room for more boot environments and 32 GiB 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 GiB) 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 GiB 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.

1.3.3 Storage Disks and Controllers

The [Disk section](https://www.freebsd.org/releases/11.2R/hardware.html#disk) (<https://www.freebsd.org/releases/11.2R/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) (<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) (<https://www.freebsd.org/cgi/man.cgi?query=mfi>) driver can be used instead by removing the loader [Tunable](#) (page 91): `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) (<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/GiB. An example configuration would be six 600 GiB 15K SAS drives in a RAID 10 which would yield 1.8 TiB of usable space, or eight 600 GiB 15K SAS drives in a RAID 10 which would yield 2.4 TiB of usable space.

For ZFS, [Disk Space Requirements for ZFS Storage Pools](https://docs.oracle.com/cd/E19253-01/819-5461/6n7ht6r12/index.html) (<https://docs.oracle.com/cd/E19253-01/819-5461/6n7ht6r12/index.html>) recommends a minimum of 16 GiB of disk space. FreeNAS® allocates 2 GiB of swap space on each drive. Combined with ZFS space requirements, this means that **it is not possible to format drives smaller than 3 GiB**. Drives larger than 3 GiB but smaller than the minimum recommended capacity might be usable but lose a significant portion of storage space to swap allocation. For example, a 4 GiB drive only has 2 GiB of available space after swap allocation.

New ZFS user 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_Recommendations) (https://web.archive.org/web/20161028084224/http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide#ZFS_Storage_Pools_Recommendations) 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 TiB and two 4 TiB disks will only be able to use 2 TiB 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.

1.3.4 Network Interfaces

The [Ethernet section](https://www.freebsd.org/releases/11.2R/hardware.html#ethernet) (<https://www.freebsd.org/releases/11.2R/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 MiB/s throughput. For higher network throughput, multiple GigE cards can be bonded

together using the LACP type of [Link Aggregations](#) (page 146). 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/old/Samba3-Developers-Guide/architecture.html) (https://www.samba.org/samba/docs/old/Samba3-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)` (https://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 106).

1.4 Getting Started with ZFS

Readers new to ZFS should take a moment to read the [ZFS Primer](#) (page 354).

INSTALLING AND UPGRADING

The FreeNAS® operating system has to be installed on a separate device from the drives which hold the storage data. With only one disk drive, the FreeNAS® web 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 disks for storage.

This section describes:

- *Getting FreeNAS®* (page 19)
- *Preparing the Media* (page 19)
- *Performing the Installation* (page 21)
- *Installation Troubleshooting* (page 29)
- *Upgrading* (page 30)
- *Virtualization* (page 37)

2.1 Getting FreeNAS®

The latest STABLE version of FreeNAS® 11.2 is available for download from <https://download.freenas.org/latest/>.

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 19).

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.

Ensure 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 is used to write the `.iso` file to an inserted USB thumb drive.

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`! This command can be avoided by writing the `.iso` file to a CD instead.

This example demonstrates writing the image to the first USB device connected to a FreeBSD system. This first device usually reports as `/dev/da0`. Replace `FreeNAS-RELEASE.iso` with the filename of the downloaded FreeNAS® ISO file. Replace `/dev/da0` with the device name of the device to write.

```
dd if=FreeNAS-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 macOS

Insert the USB thumb drive. In 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. To determine the correct path 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, replace `FreeNAS-RELEASE.iso` with the name of the FreeNAS® ISO and `/dev/rdisk1` with the correct path to the USB thumb drive:

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

dd if=FreeNAS-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 a “Permission denied” error is shown, use `sudo` for elevated rights: `sudo dd if=FreeNAS-11.0-RELEASE.iso of=/dev/rdisk1 bs=64k`. This will prompt for the 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 boot menu is displayed as is shown in [Figure 2.1](#).

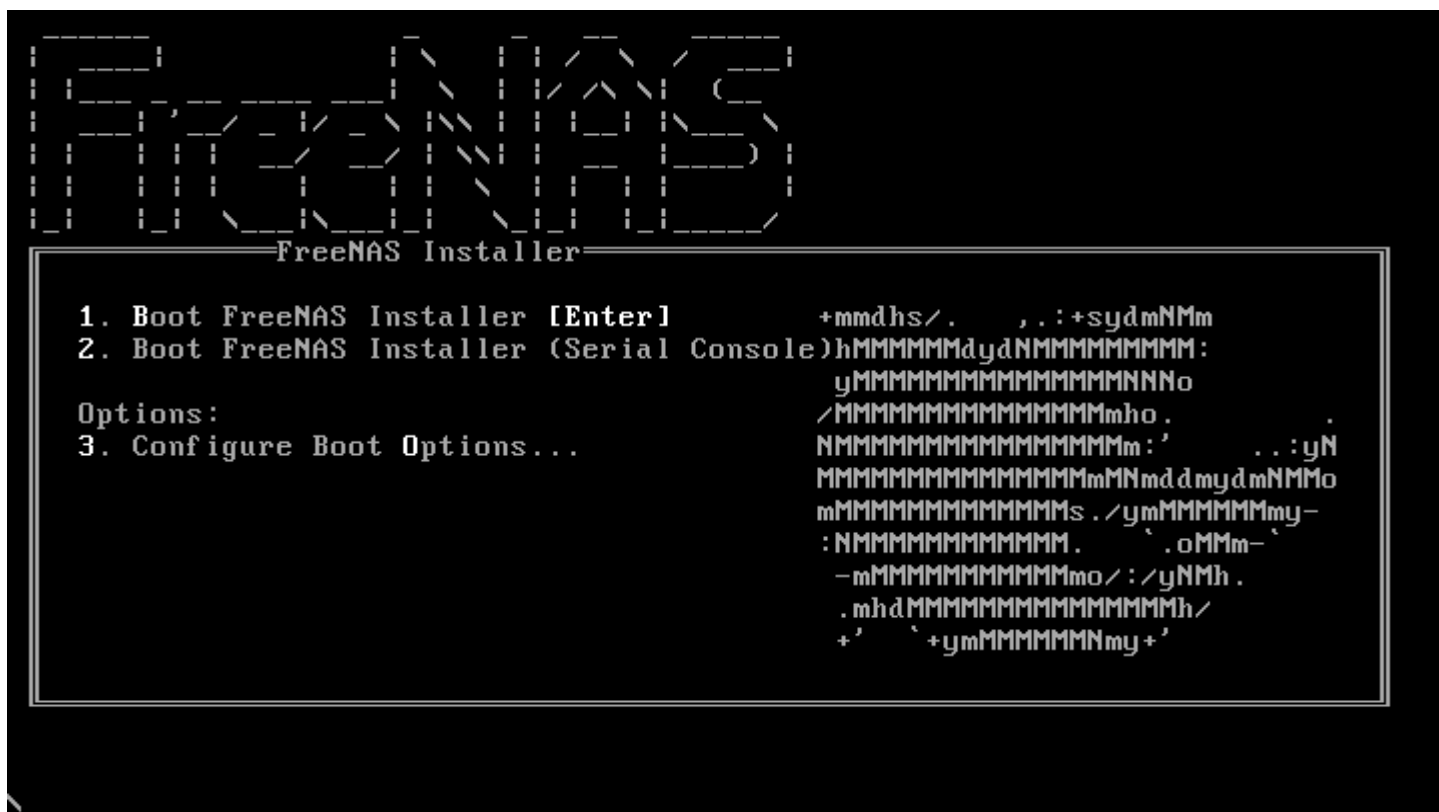


Fig. 2.1: Installer Boot Menu

The FreeNAS® installer automatically boots into the default option after ten seconds. If needed, choose another boot option by pressing the **Spacebar** to stop the timer and then enter the number of the desired option.

Tip: The *Serial Console* 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 144).

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.

Once the installer has finished 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 GiB of RAM is required and the installer will present a warning message if less than 8 GiB is detected.

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

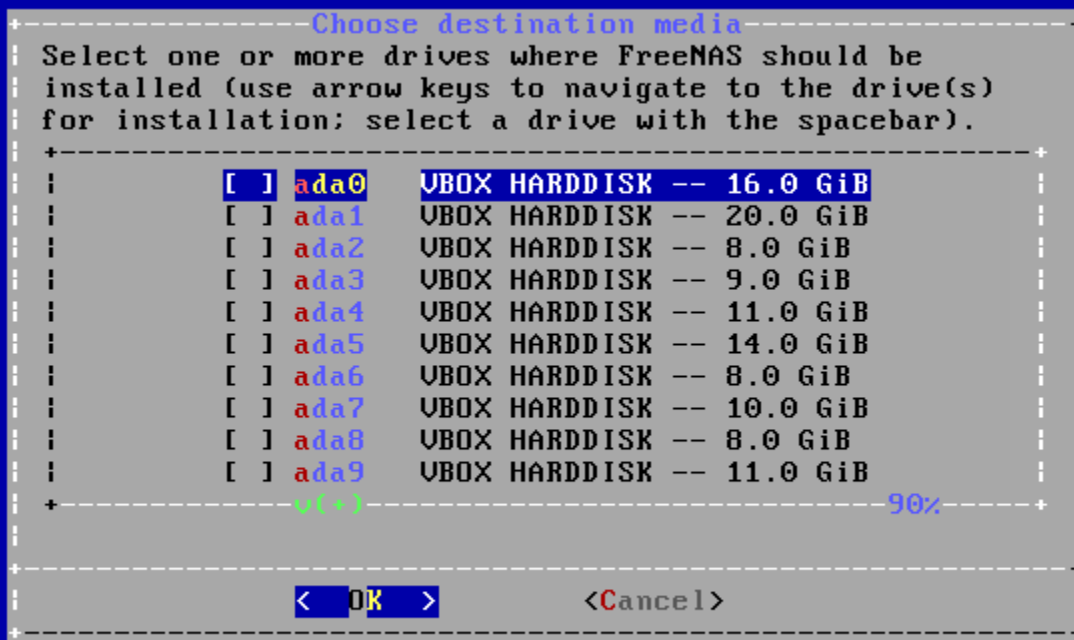


Fig. 2.3: Selecting the Install Drive

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 GiB of space on the boot device is required. However, 32 GiB 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®. 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 web 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 web interface, it needs to be hard to guess. Enter the password, press the down arrow key, and confirm the password. Then press `Enter` to continue with the installation. Choosing *Cancel* skips setting a root password during the installation, but the web interface will require setting a root password when logging in for the first time.

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 web interface. This means that the FreeNAS® system needs to be kept physically secure and that the web interface needs to 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 [Installer Menu](#) (page 23). 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 55) after waiting five seconds in the *boot menu* (page 34). Press the `Spacebar` to stop the timer and use the boot menu.

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/>).

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/>). 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 web 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 11.2. After 11.2 has been installed, use the instructions in [Update](#) (page 94) to keep the system updated.

2.5.1 Caveats

Be aware of these caveats **before** attempting an upgrade to 11.2:

- **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 330) 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.
- Upgrading the firmware of Broadcom SAS HBAs to the latest version is recommended.
- If upgrading from 9.3.x, 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/>) 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 pools 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 not supported.** If the data currently resides on **one** UFS-formatted disk, create a ZFS pool using **other** disks after the upgrade, then use the instructions in [Importing a Disk](#) (page 181) to mount the UFS-formatted disk and copy the data to the ZFS pool. 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 pool. Instead, back up the data before the upgrade, create a ZFS pool after the upgrade, then restore the data from the backup.
- **The VMware Tools VMXNET3 drivers are not supported.** Configure and use the [vmx\(4\)](#) (<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 pools are encrypted, **remember** to set a passphrase and download a copy of the encryption key and the latest recovery key. After the upgrade is complete, use the instructions in [Importing a Pool](#) (page 163) to import the encrypted pools.

3. Warn users that the FreeNAS® shares will be unavailable during the upgrade; it is recommended to schedule the upgrade for a time that will least impact users.
4. Stop all services in *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 19).

Insert the prepared media into the system and boot from it. The installer waits ten seconds in the *installer boot menu* (page 22) before booting the default option. If needed, press the `Spacebar` to stop the timer and choose another boot option. 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® web interface, 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: A FreeNAS® system can be upgraded by backing up the existing configuration data, performing a fresh install, then restoring the configuration. Back up the configuration with *System → General → Save Config*. Choose *Fresh Install* during the installation. After booting into the new install, restore the previous configuration data with *System → General → Upload Config*.

To perform an upgrade, press `Enter` to accept the default of *Upgrade Install*. Again, the installer will display a reminder 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 can 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 can cause the system to reboot again. The system will start normally afterwards. If database errors are shown but the web interface is accessible, go to *Settings* → *General* and use the *UPLOAD CONFIG* button to upload the configuration that was saved before starting the upgrade.

2.5.4 Upgrading From the Web Interface

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

The connection is lost temporarily when the update is complete. It returns after the FreeNAS® system reboots into the new version of the operating system. The FreeNAS® system normally receives the same IP address from the DHCP server. Refresh the browser after a moment to see if the system is accessible.

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:



Fig. 2.12: Boot Menu

FreeNAS® waits five seconds before booting into the default boot environment. Press the **Spacebar** to stop the automatic boot timer. Press **4** to display the available boot environments and press **3** as needed to scroll through multiple pages.

In the example shown in [Figure 2.13](#), the first entry in *Boot Environments* is `11.2-MASTER-201807250900`. This is the current version of the operating system, after the update was applied. Since it is the first entry, it is the default selection.

The next entry is `Initial-Install`. This is the original boot environment created when FreeNAS® was first installed. Since there are no other entries between the initial installation and the first entry, only one update has been applied to this system since its initial installation.

To boot into another version of the operating system, enter the number of the boot environment to set it as *Active*. Press **Backspace** to return to the *Boot Menu* (page 34) and press **Enter** to boot into the chosen *Active* boot environment.

If a boot device fails and the system no longer boots, do not 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 pools in *Storage* → *Auto Import Pool*.
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 a reboot into an older version of the operating system is performed, a configuration created in a later version cannot be restored.

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 the end user has a specific need for the newer ZFS feature flags. 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* → *Pools* and click  (Settings) to upgrade. Click the *Upgrade Pool* button as shown in [Figure 2.14](#).

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

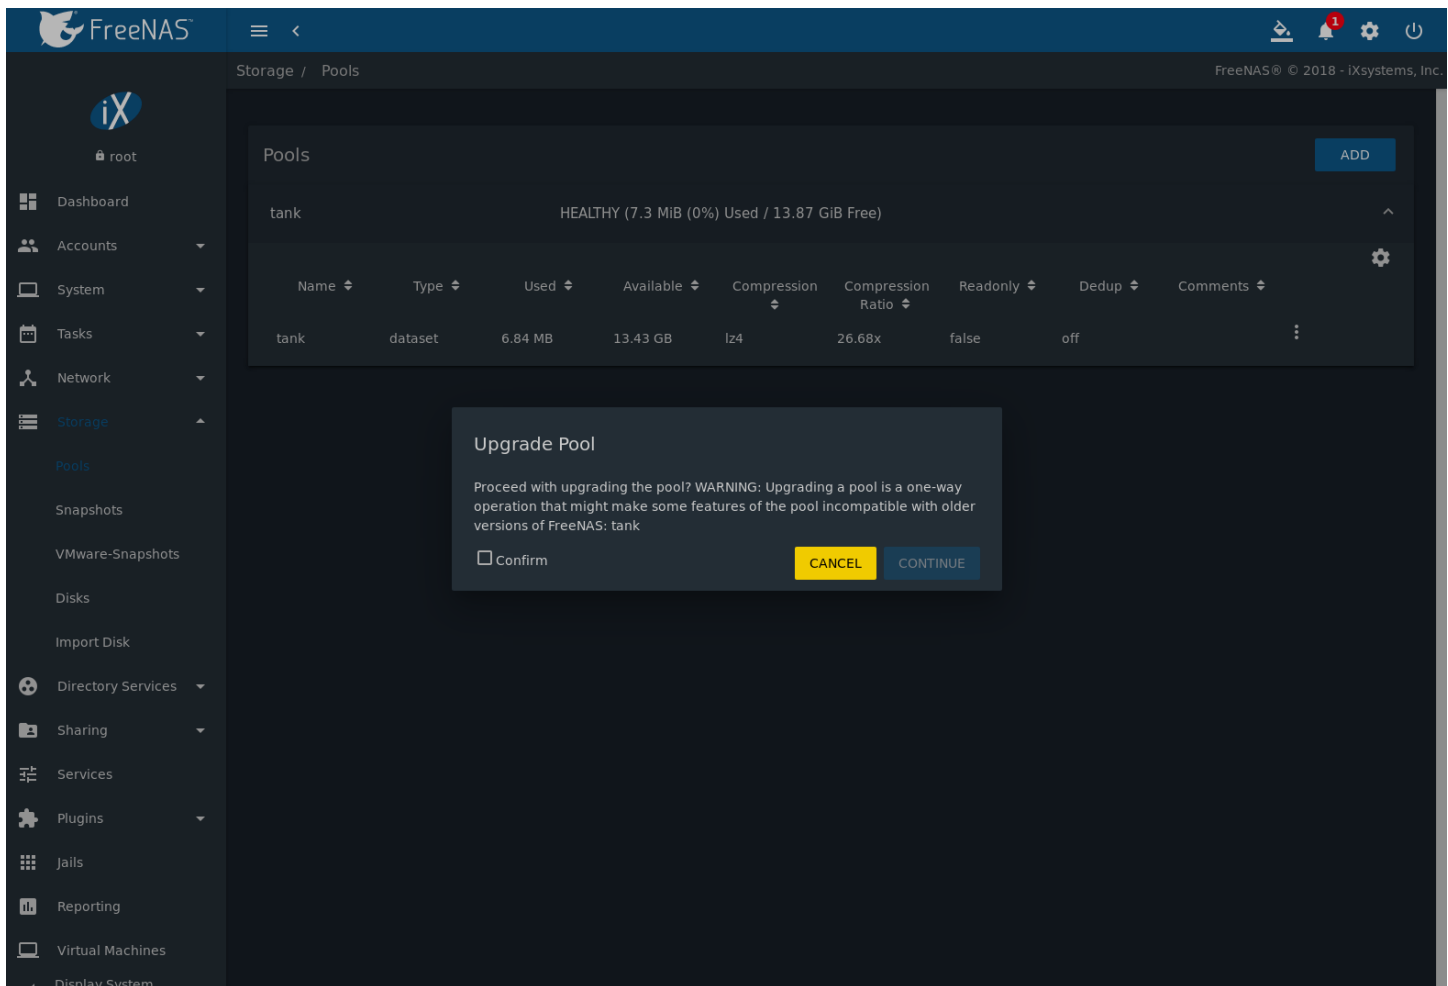


Fig. 2.14: Upgrading a Pool

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. 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/) (https://forums.freenas.org/index.php?threads/please-do-not-run-freenas-in-production-as-a-virtual-machine.12484/). When using 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/) (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 MiB (8 GiB) base memory size
- a virtual disk **at least 8 GiB in size** to hold the operating system and boot environments
- at least one additional virtual disk **at least 4 GiB 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/) (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 2.15](#), to start the new virtual machine wizard.


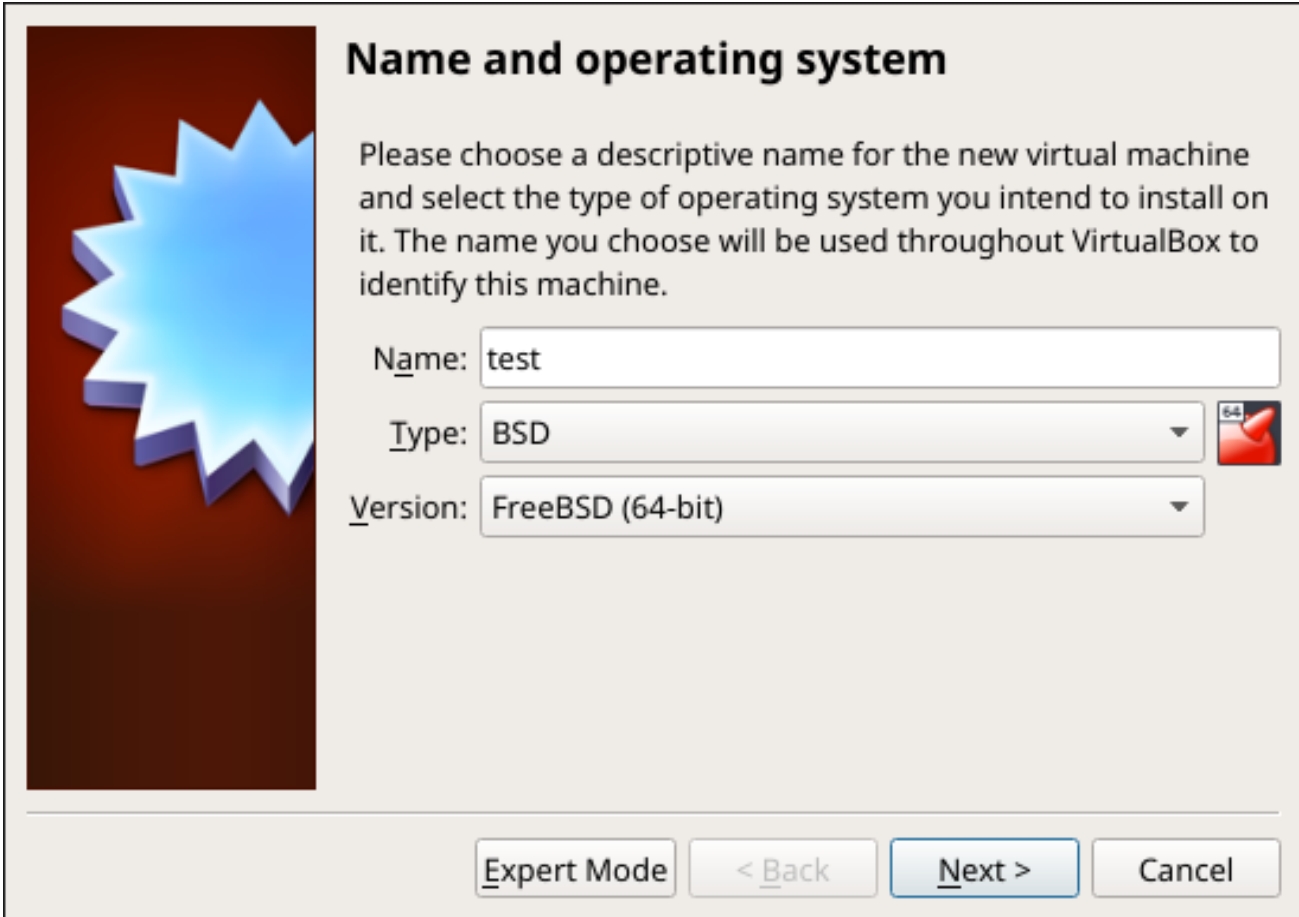


Fig. 2.15: Initial VirtualBox Screen


Click the *Next* button to see the screen in [Figure 2.16](#). 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.



Name and operating system

Please choose a descriptive name for the new virtual machine and select the type of operating system you intend to install on it. The name you choose will be used throughout VirtualBox to identify this machine.

Name:

Type: 

Version:

Fig. 2.16: Enter Name and Operating System for the New Virtual Machine

Click *Next* to see the screen in [Figure 2.17](#). The base memory size must be changed to **at least 8192 MiB**. When finished, click *Next* to see the screen in [Figure 2.18](#).

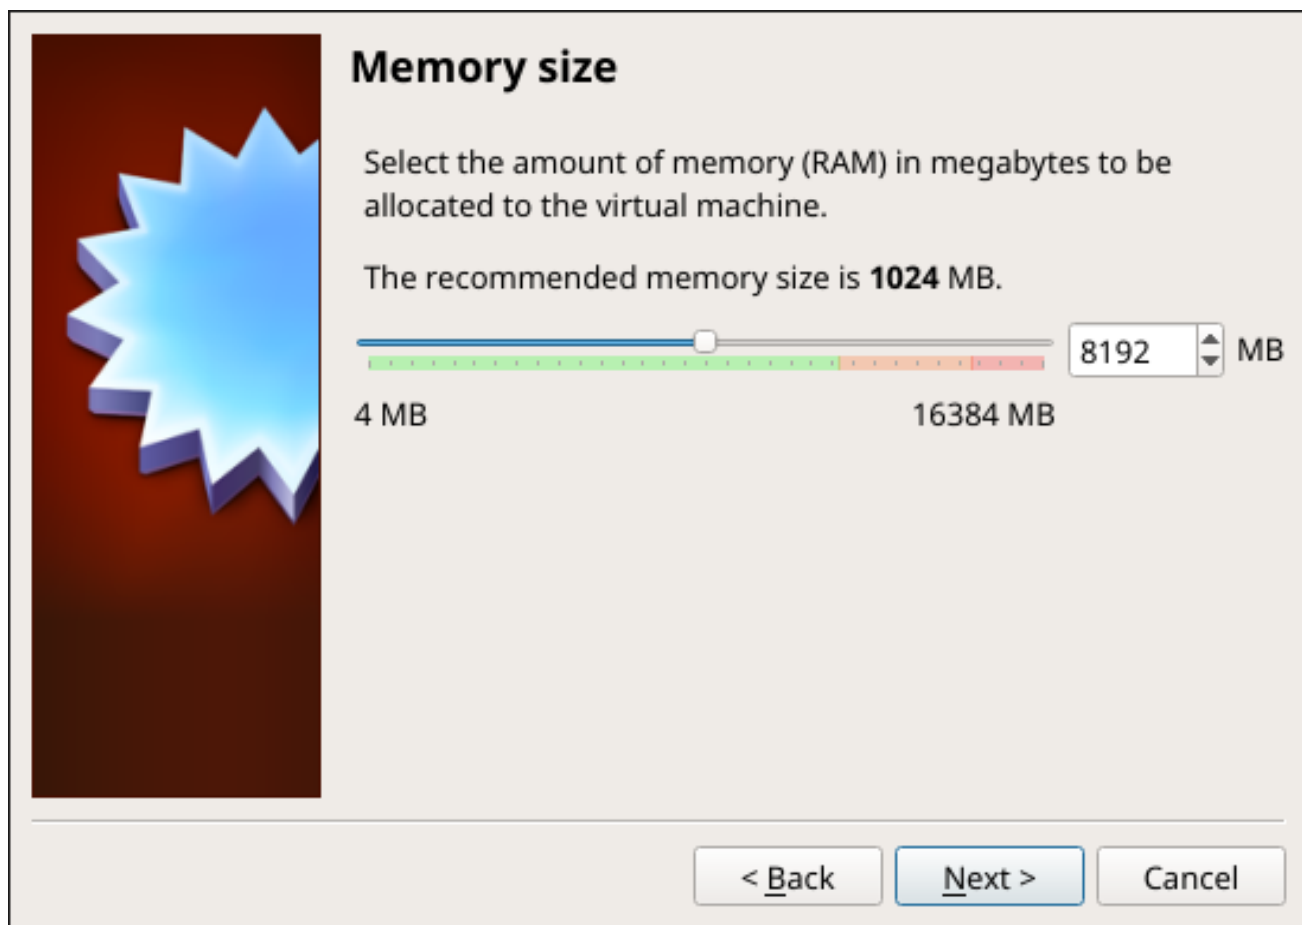


Fig. 2.17: Select the Amount of Memory Reserved for the Virtual Machine



Fig. 2.18: Select Existing or Create a New Virtual Hard Drive

Click *Create* to launch the *Create Virtual Hard Drive Wizard* shown in [Figure 2.19](#).



Fig. 2.19: Create New Virtual Hard Drive Wizard

Select *VDI* and click the *Next* button to see the screen in [Figure 2.20](#).




Fig. 2.20: 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 2.21](#) is shown.



File location and size

Please type the name of the new virtual hard disk file into the box below or click on the folder icon to select a different folder to create the file in.

test 

Select the size of the virtual hard disk in megabytes. This size is the limit on the amount of file data that a virtual machine will be able to store on the hard disk.

4.00 MB  2.00 TB 8.00 GB

< Back Create Cancel

Fig. 2.21: Select File Name and Size of Virtual Disk

This screen is used to set the size (or upper limit) of the virtual disk. **Set the default size to a minimum of 8 GiB.** 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 GiB and at least one data storage disk of at least 4 GiB.

Use the *Back* button to return to a previous screen if any values need to be modified. After making a selection and pressing *Create*, the new VM is created. The new virtual machine is listed in the left frame, as shown in the example in [Figure 2.22](#). Open the *Machine Tools* drop-down menu and select *Details* to see extra information about the VM.

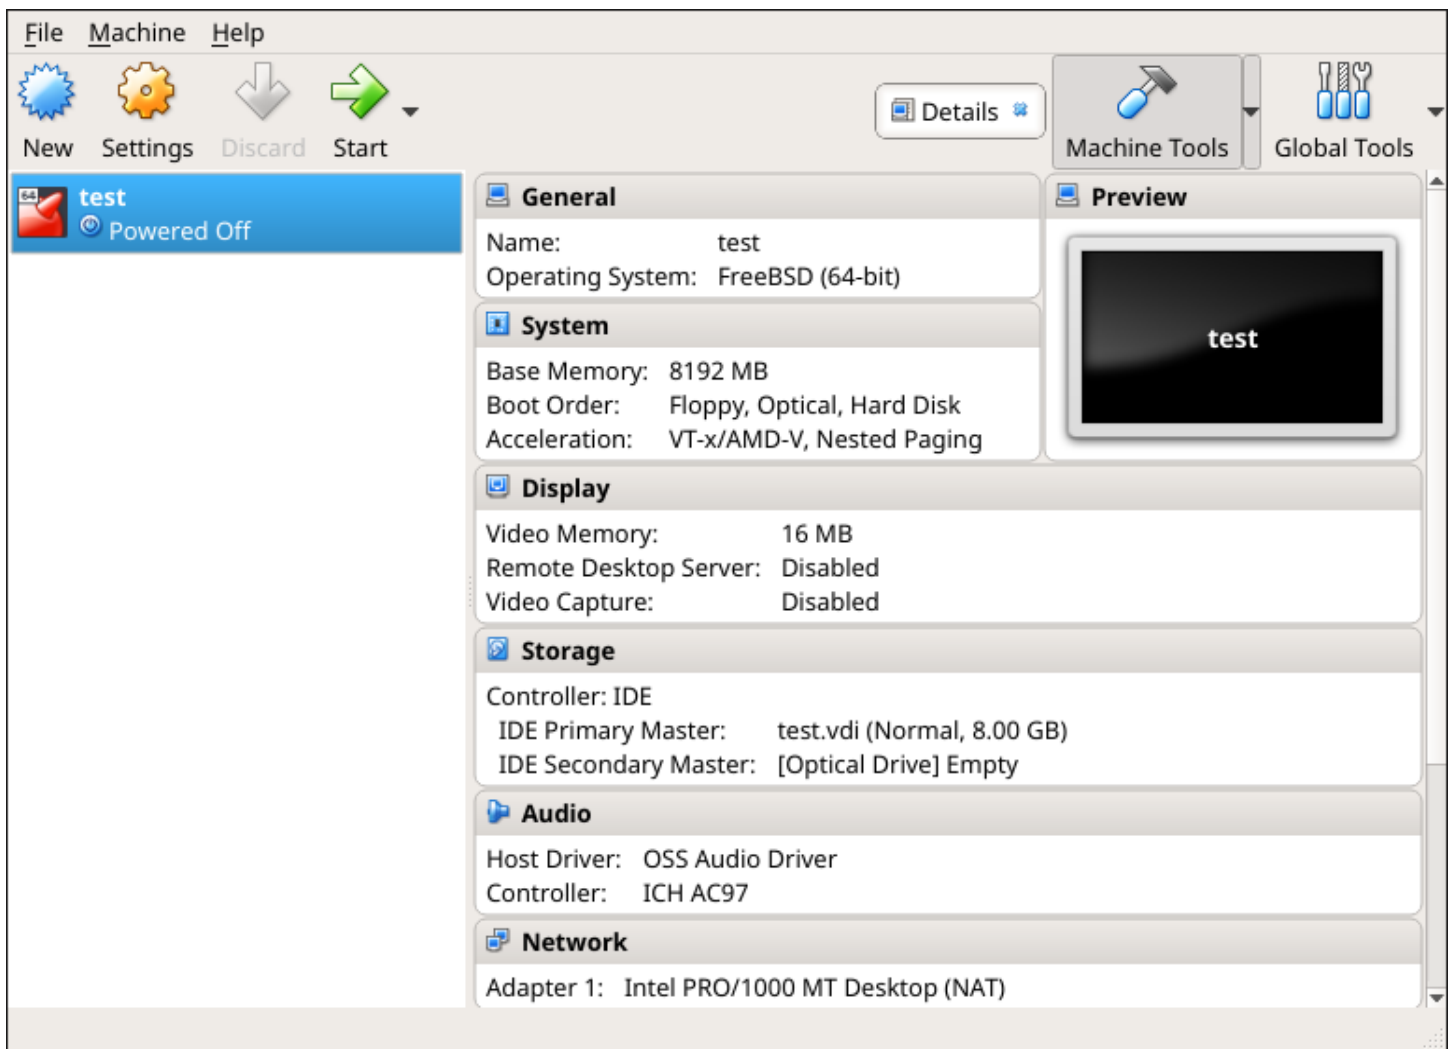


Fig. 2.22: The New Virtual Machine

Create the virtual disks to be used for storage. Highlight the VM and click *Settings* to open the menu. Click the *Storage* option in the left frame to access the storage screen seen in [Figure 2.23](#).

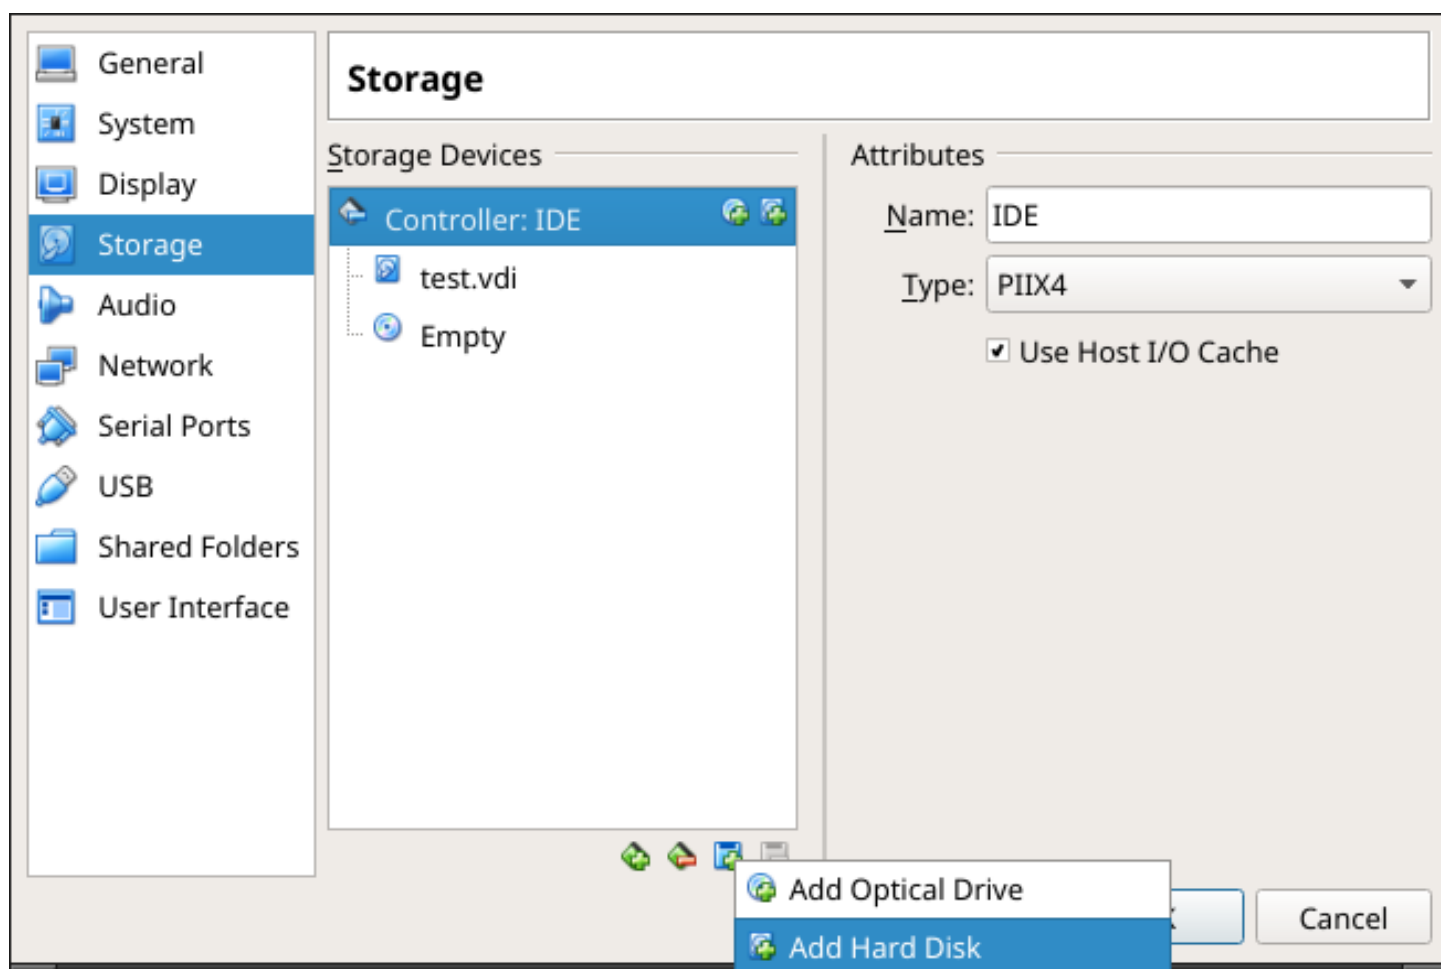


Fig. 2.23: Storage Settings of the Virtual Machine

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 Virtual Hard Disk* wizard seen in [Figure 2.19](#) and [2.20](#).

Create a disk large enough to hold the desired data. The minimum size is **4 GiB**. 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 2.24](#).




Fig. 2.24: Configuring ISO Installation Media

Click *Choose Virtual Optical 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 an `.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 there is 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* → *Adapter 1*. 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 2.25](#), the Intel Pro/1000 Ethernet card is attached to the network and has a device name of `em0`.

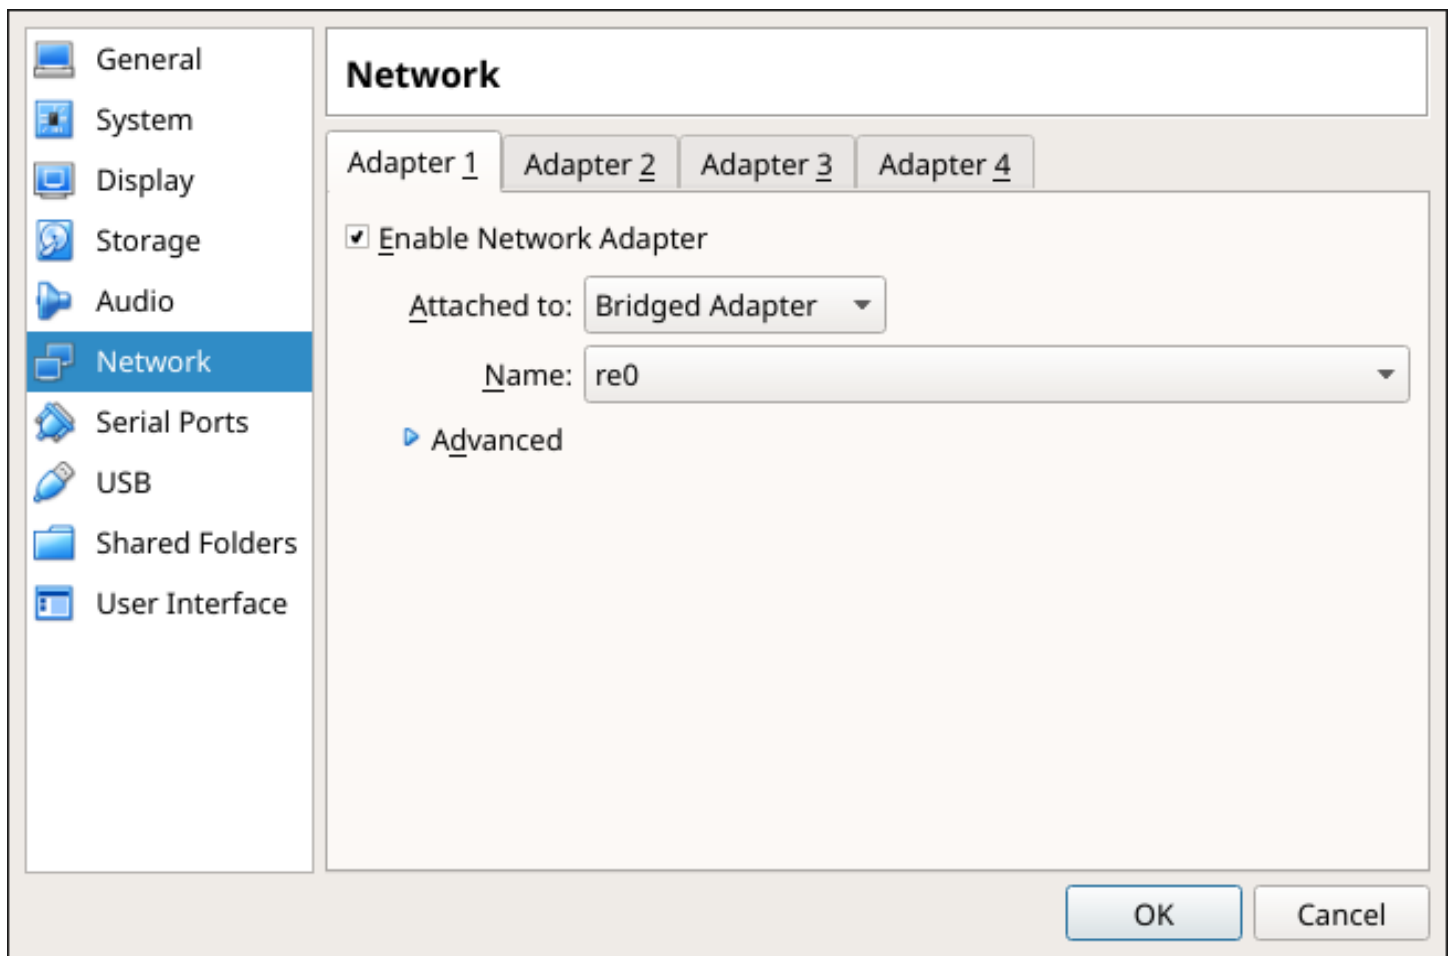


Fig. 2.25: Configuring a Bridged Adapter in VirtualBox

After configuration is complete, click the *Start* arrow and install FreeNAS® as described in [Performing the Installation](#) (page 21). After 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 *Optical* device in *Storage* or by unchecking *Optical* 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](https://www.vmware.com/products/esxi-and-esx.html?ClickID=bldvs6vzfq1svmyezgusneeezqfl6qddgne) (https://www.vmware.com/products/esxi-and-esx.html?ClickID=bldvs6vzfq1svmyezgusneeezqfl6qddgne). After the operating system is installed on the 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 2.26](#).

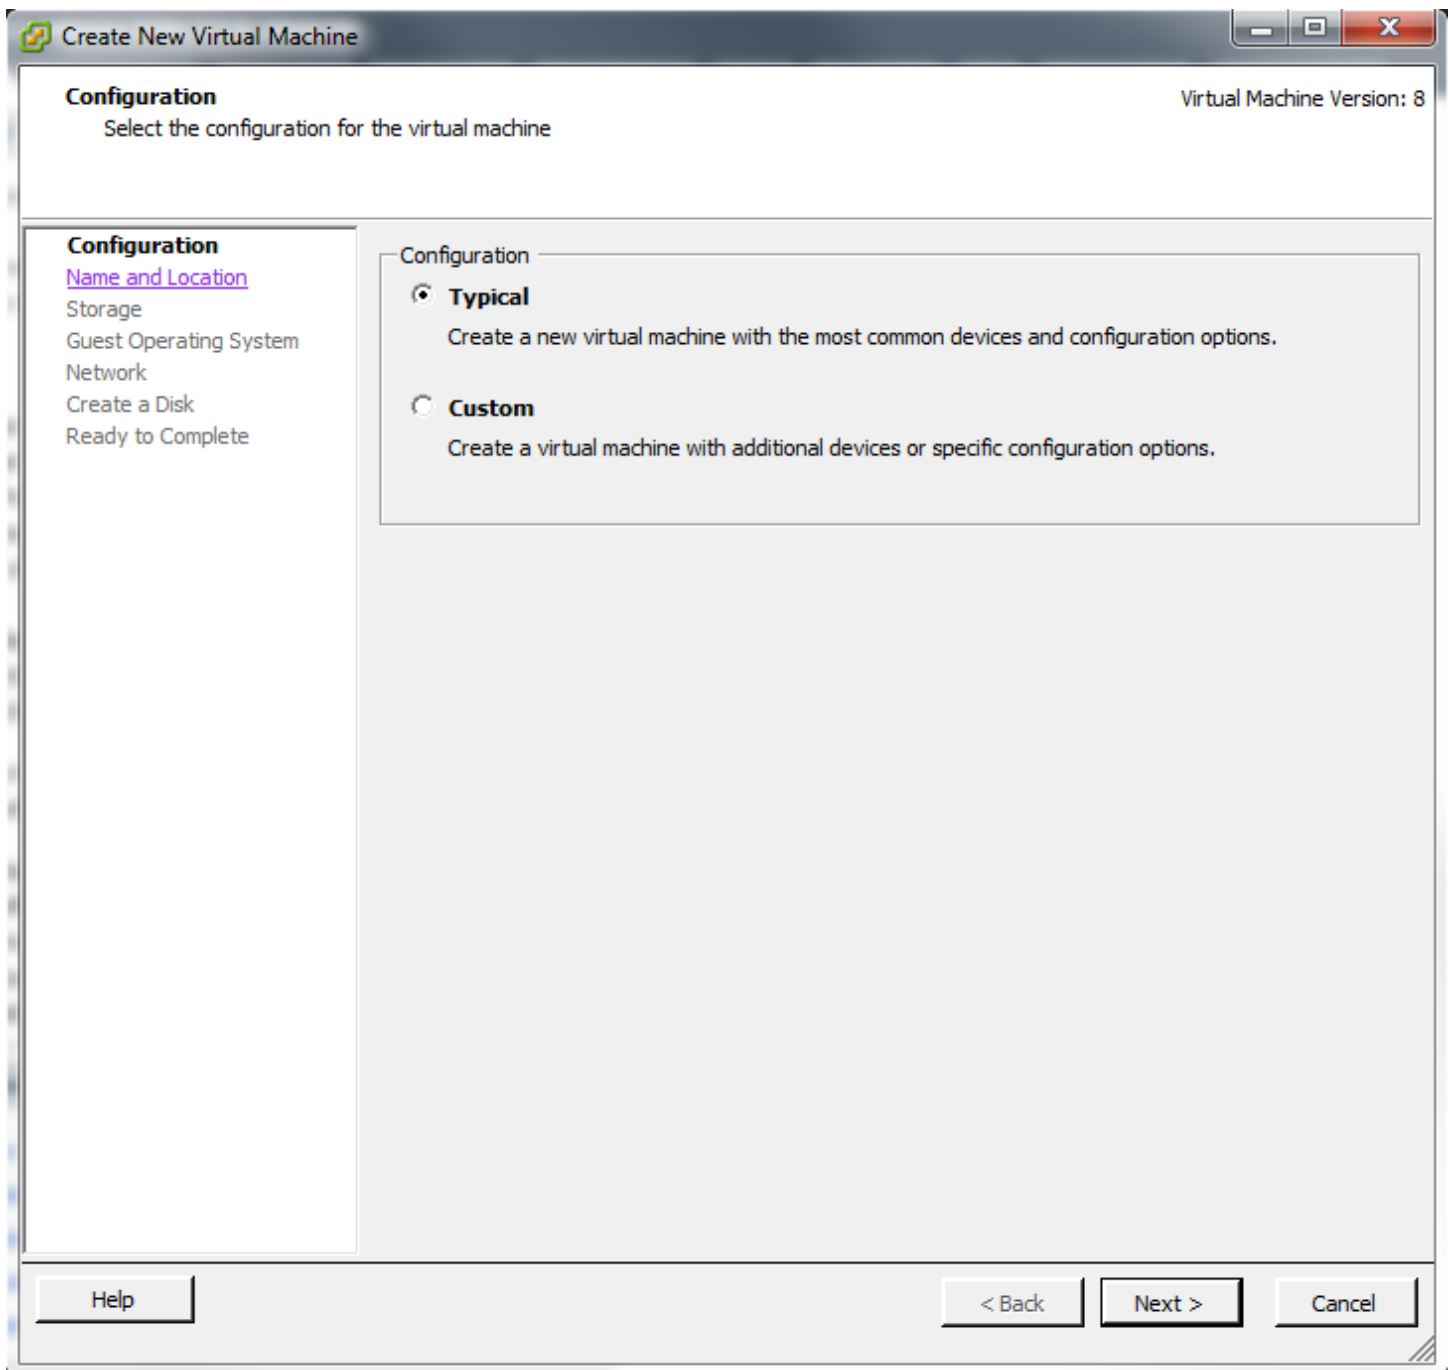


Fig. 2.26: New Virtual Machine Wizard

Click *Next* and enter a name for the virtual machine. Click *Next* and highlight a datastore. An example is shown in [Figure 2.27](#). Click *Next*. In the screen shown in [Figure 2.28](#), click *Other*, then select a FreeBSD 64-bit architecture.




Fig. 2.27: Select Datastore

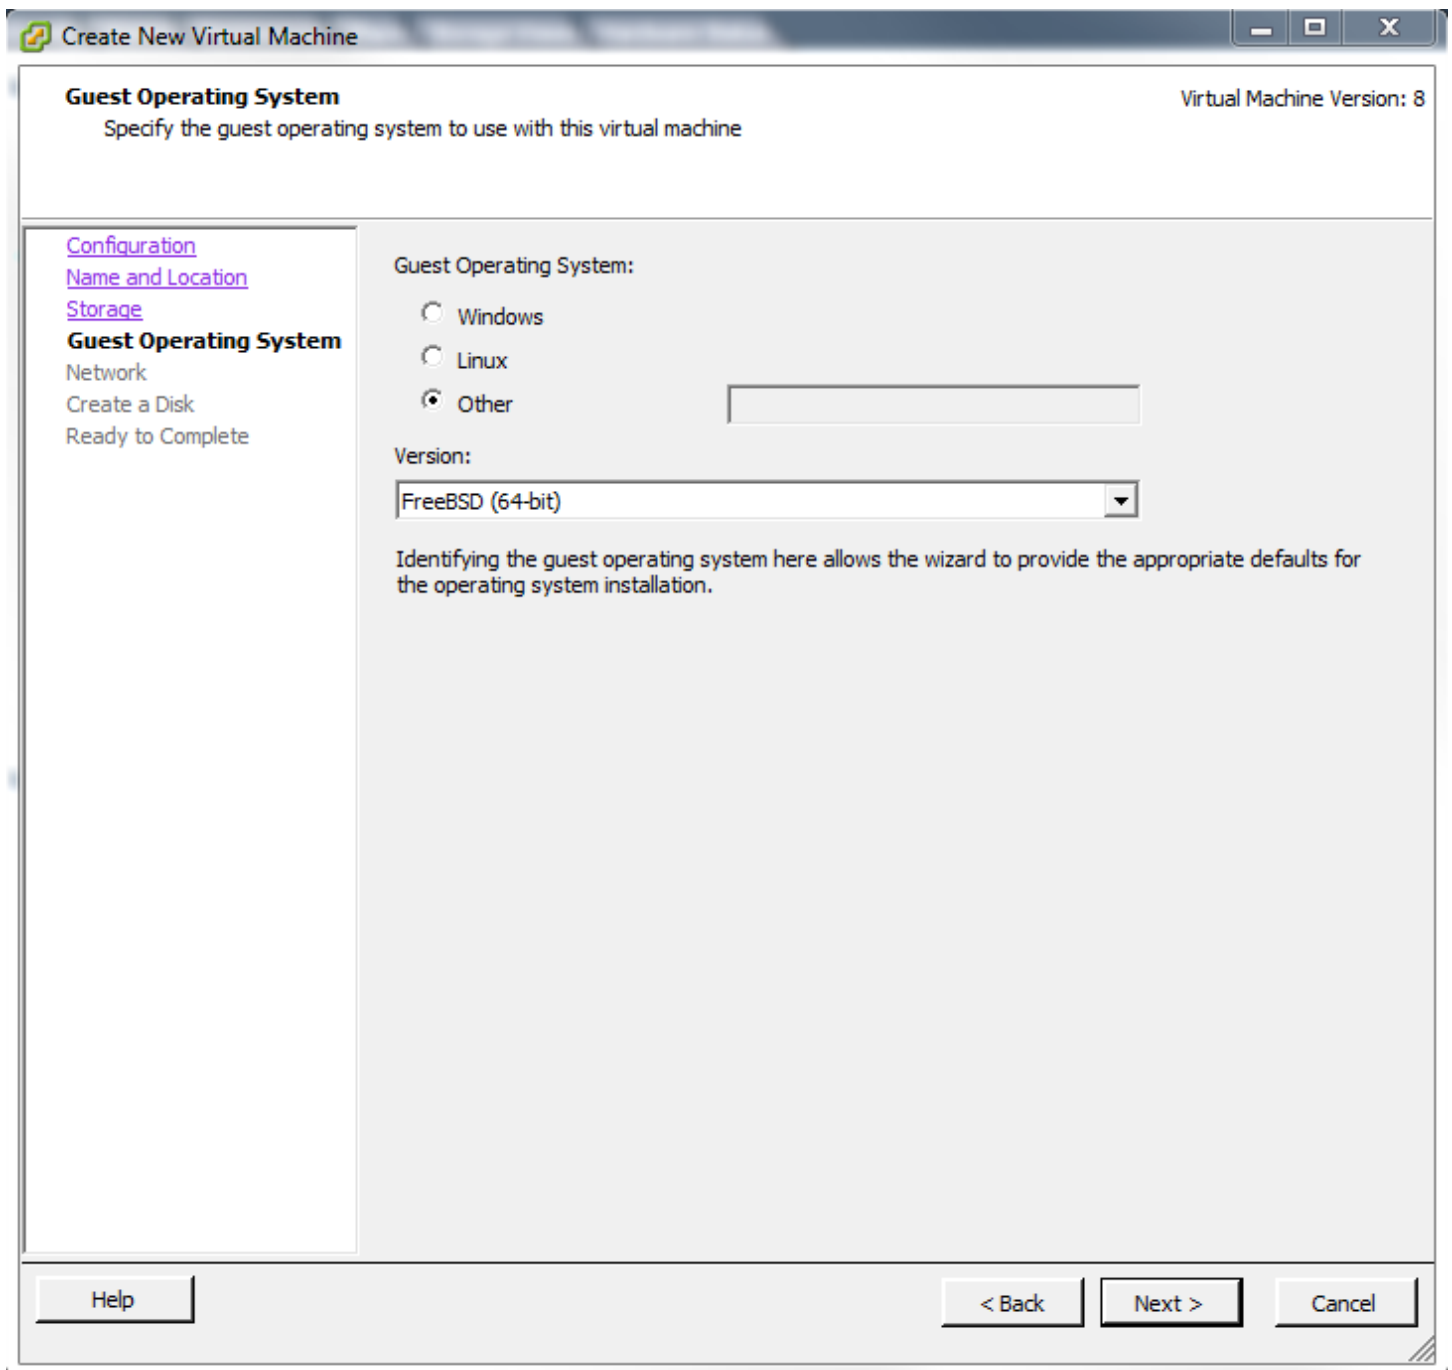


Fig. 2.28: Select Operating System

Click *Next* and create a virtual disk file of **8 GiB** to hold the FreeNAS® operating system, as shown in [Figure 2.29](#).

Create New Virtual Machine

Create a Disk
Specify the virtual disk size and provisioning policy

Virtual Machine Version: 8

[Configuration](#)
[Name and Location](#)
[Storage](#)
[Guest Operating System](#)
[Network](#)
Create a Disk
Ready to Complete

Datastore:

Available space (GB):

Virtual disk size:

☒ Thick Provision Lazy Zeroed
☐ Thick Provision Eager Zeroed
☐ Thin Provision

Fig. 2.29: 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 2.30](#).

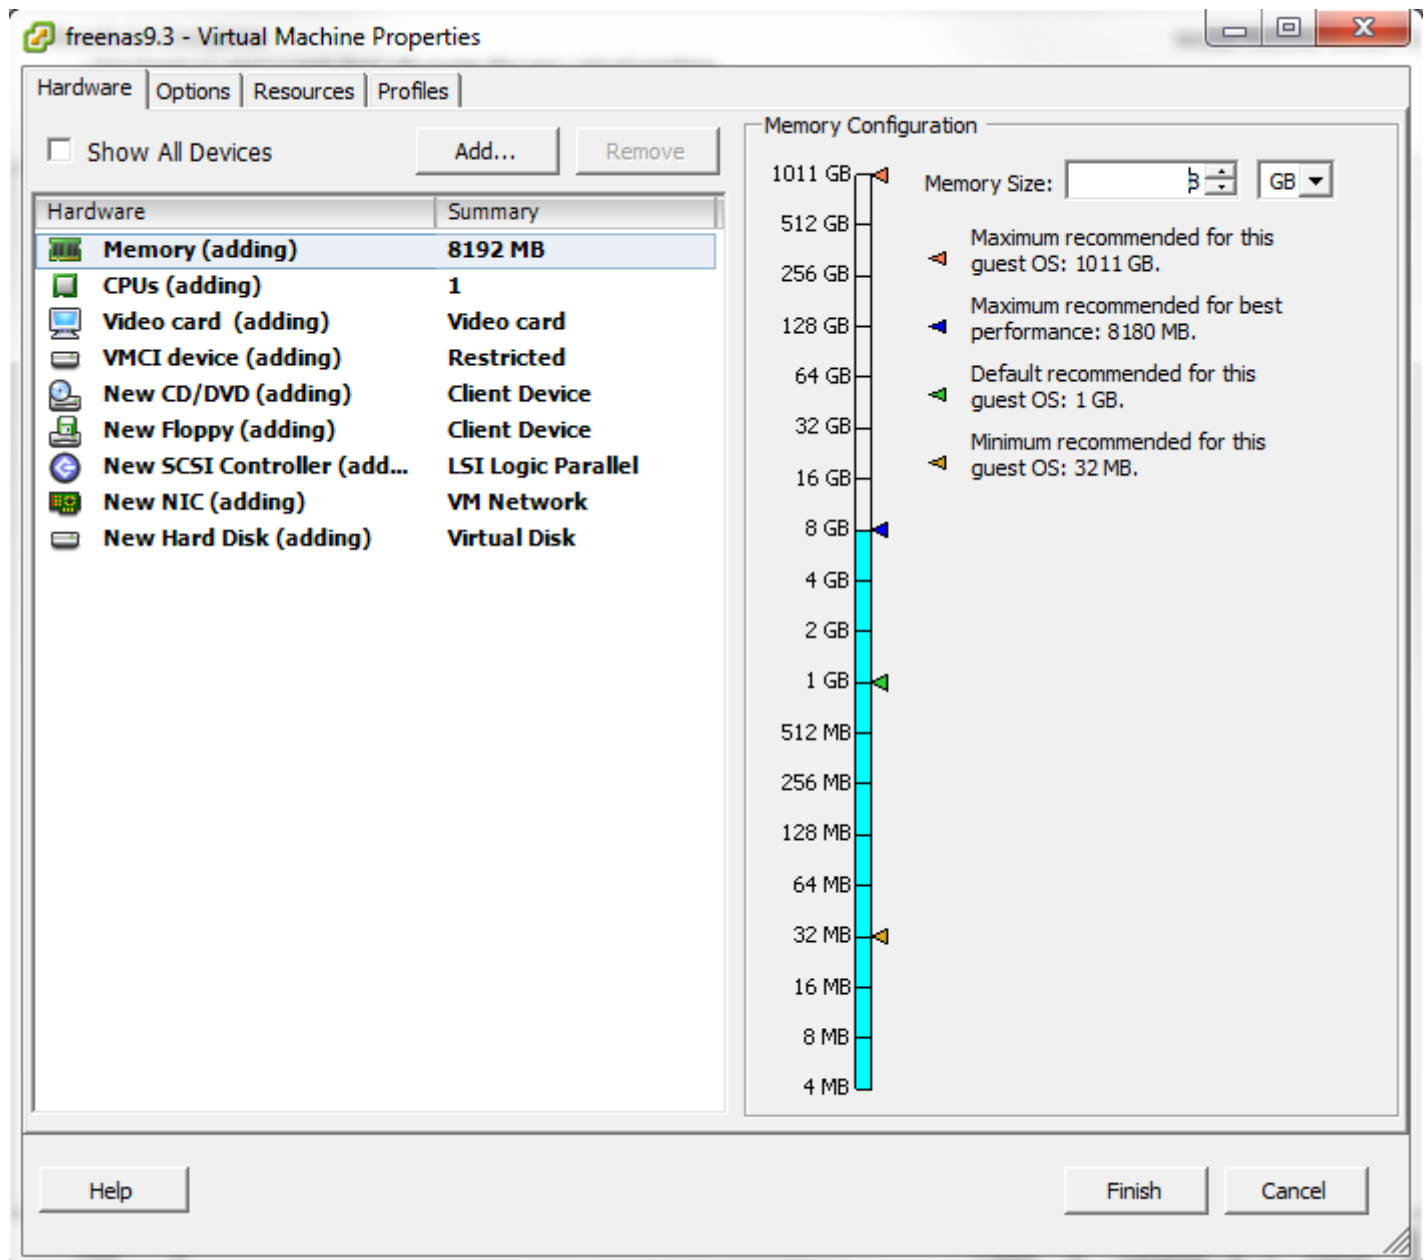


Fig. 2.30: Virtual Machine Settings

Increase the *Memory Configuration* to **at least 8192 MiB**.

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 2.31, 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.

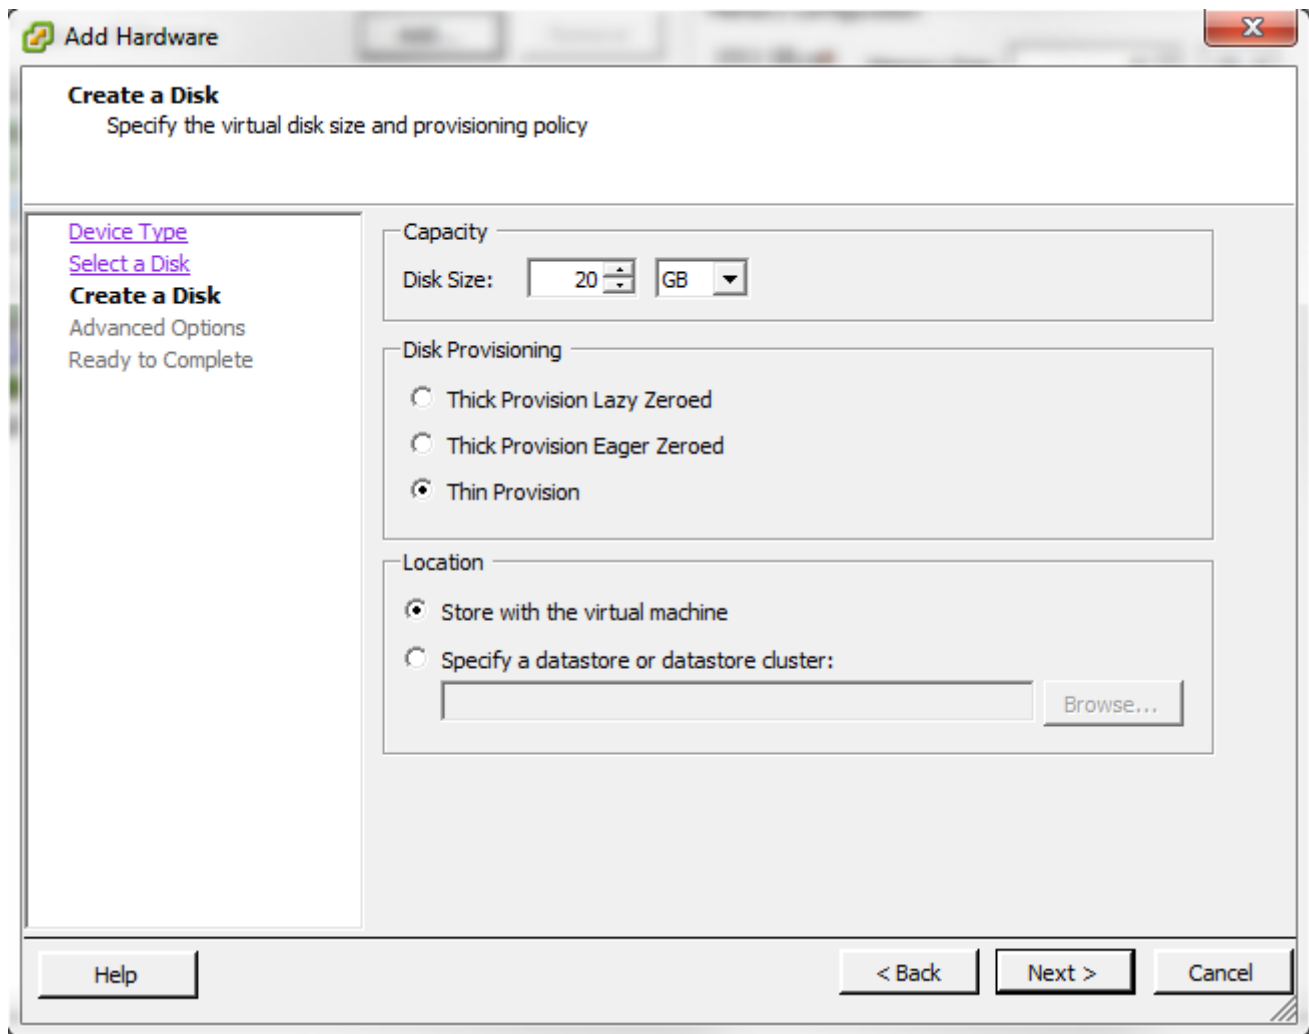


Fig. 2.31: Creating a Storage Disk

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 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 325), 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* → *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 → Pools → Import Pool* can be used to re-import pools.

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. After FreeNAS® successfully receives an IP address, the address is displayed so it can be used to access the web interface. The example in [Figure 3.1](#) shows a FreeNAS® system that is accessible at `http://10.0.0.102`.

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 1-11: 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](#).

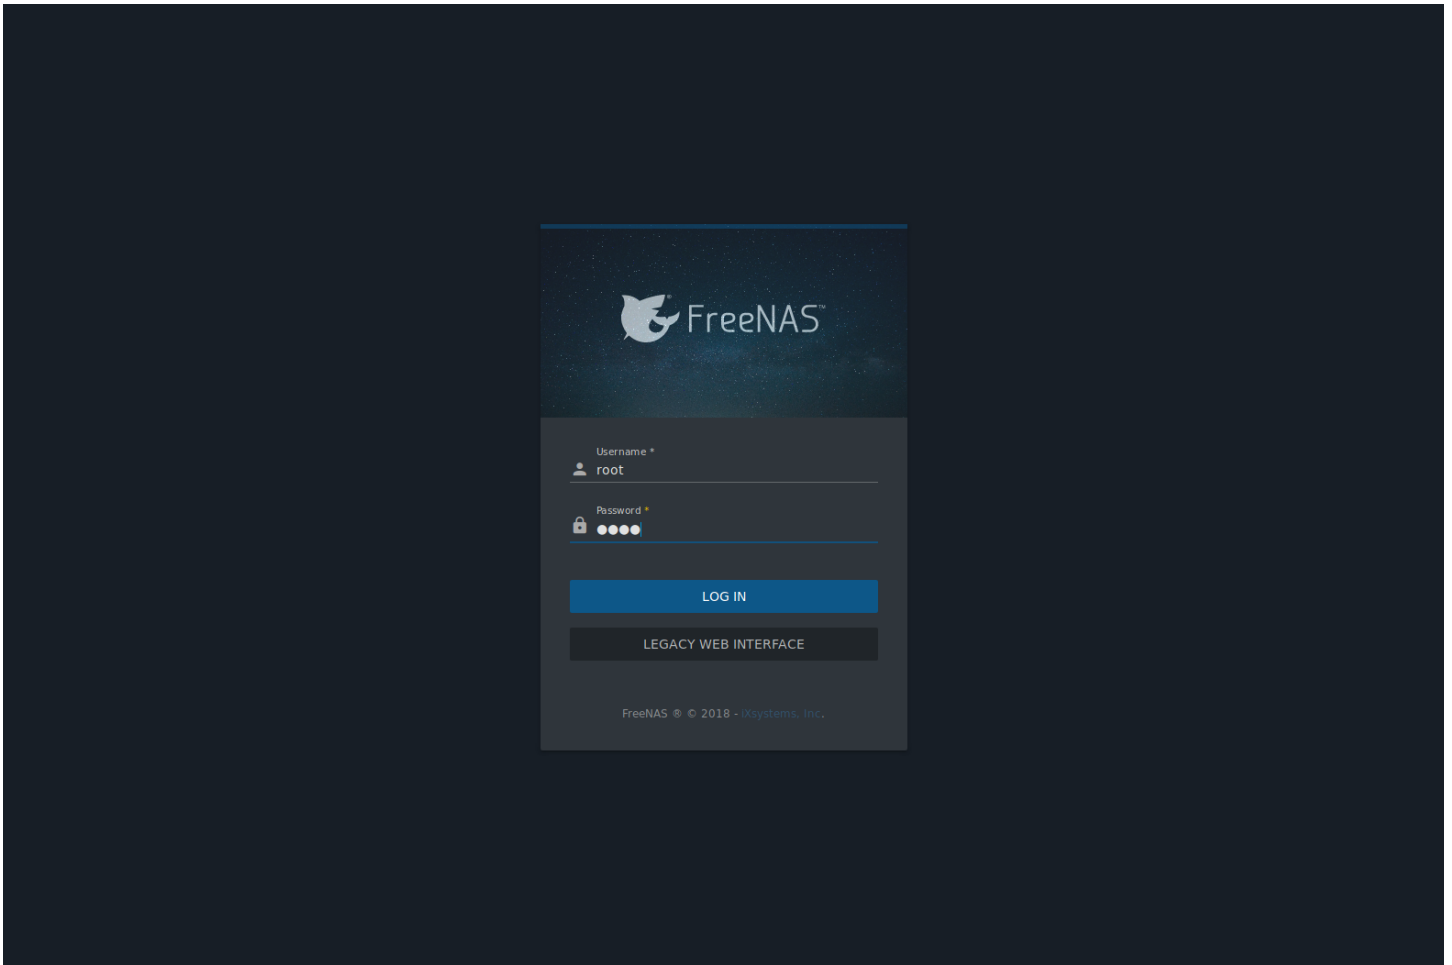


Fig. 3.2: Enter the Root Password

Note: The FreeNAS® web interface now uses Angular and a new, asynchronous middleware. To use the legacy Django web interface that was used before version 11.2, click *LEGACY WEB INTERFACE*. This User Guide only demonstrates the new Angular web interface.

Enter the password chosen during the installation. There is a prompt to set a root password if this was not set during the installation. The web interface is displayed as shown in [Figure 3.3](#).

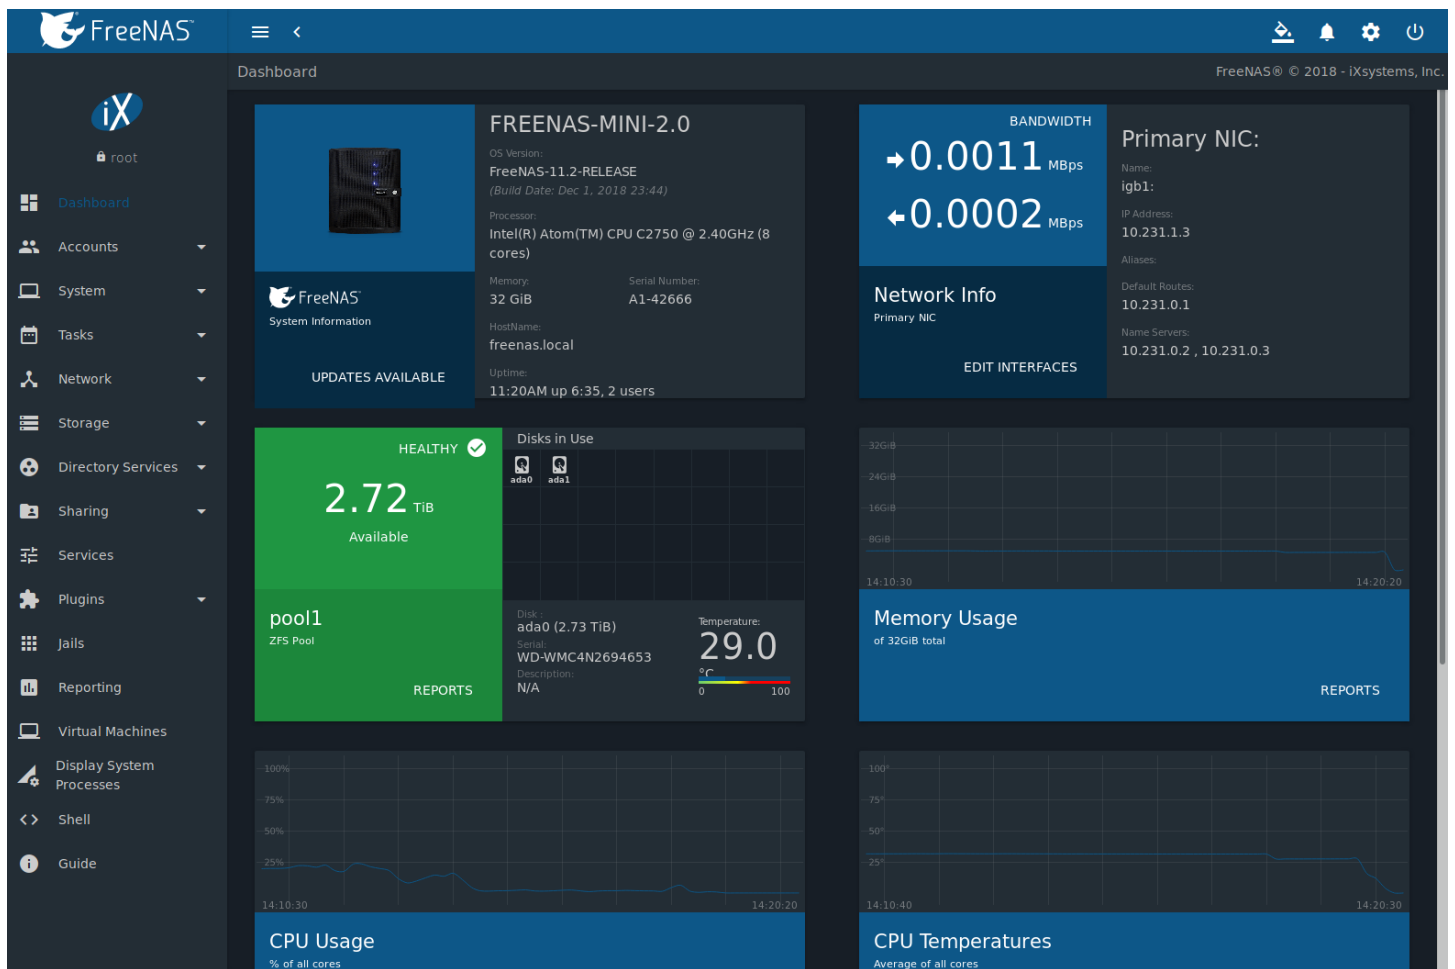


Fig. 3.3: FreeNAS® Graphical Configuration Menu

The rest of this User Guide describes the FreeNAS® web interface in more detail. The layout of this User Guide follows the order of the menu items in the tree located in the left frame of the web interface.

Note: It is important to use the web interface 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 the 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.

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

- Check for enabled proxy settings in the browser configuration, disable them, and try connecting again.
- ping the FreeNAS® system IP address from another computer on the same network.
- Try a different web browser if the user interface loads but is unresponsive or seems to be missing menu items. [Firefox](https://www.mozilla.org/en-US/firefox/all/) (<https://www.mozilla.org/en-US/firefox/all/>) is recommended.
- Make sure that the browser is set to allow cookies from the FreeNAS® system.

SETTINGS

The **Settings** menu has shortcuts to edit the `root` account settings and password, set interface preferences, view system information, and switch to the *Legacy Web Interface*.

4.1 Edit root Account

Click **Settings** and *Account* to begin editing the `root` account settings. This is the primary account used to log in and interact with the FreeNAS[®] system. See the *User Account Configuration table* (page 69) for details about each account option.

4.2 Change Password

Click **Settings** and *Change Password* to see a simplified *Change Password* form. This is used to quickly change the account password for the `root` and any other user account that is not built-in to FreeNAS[®].

Enter the *Username* and *Current Password* for the user account, then create and confirm a *New Password*. Click *SAVE* to update the account password.

4.3 Preferences

The FreeNAS[®] User Interface can be adjusted to match the user preferences. Go to the *Web Interface Preferences* page by clicking the **Settings** menu in the upper-right and clicking *Preferences*.

4.3.1 Web Interface Preferences

This page has options to adjust global settings in the web interface, manage custom themes, and create new themes. [Figure 4.1](#) shows the different options:

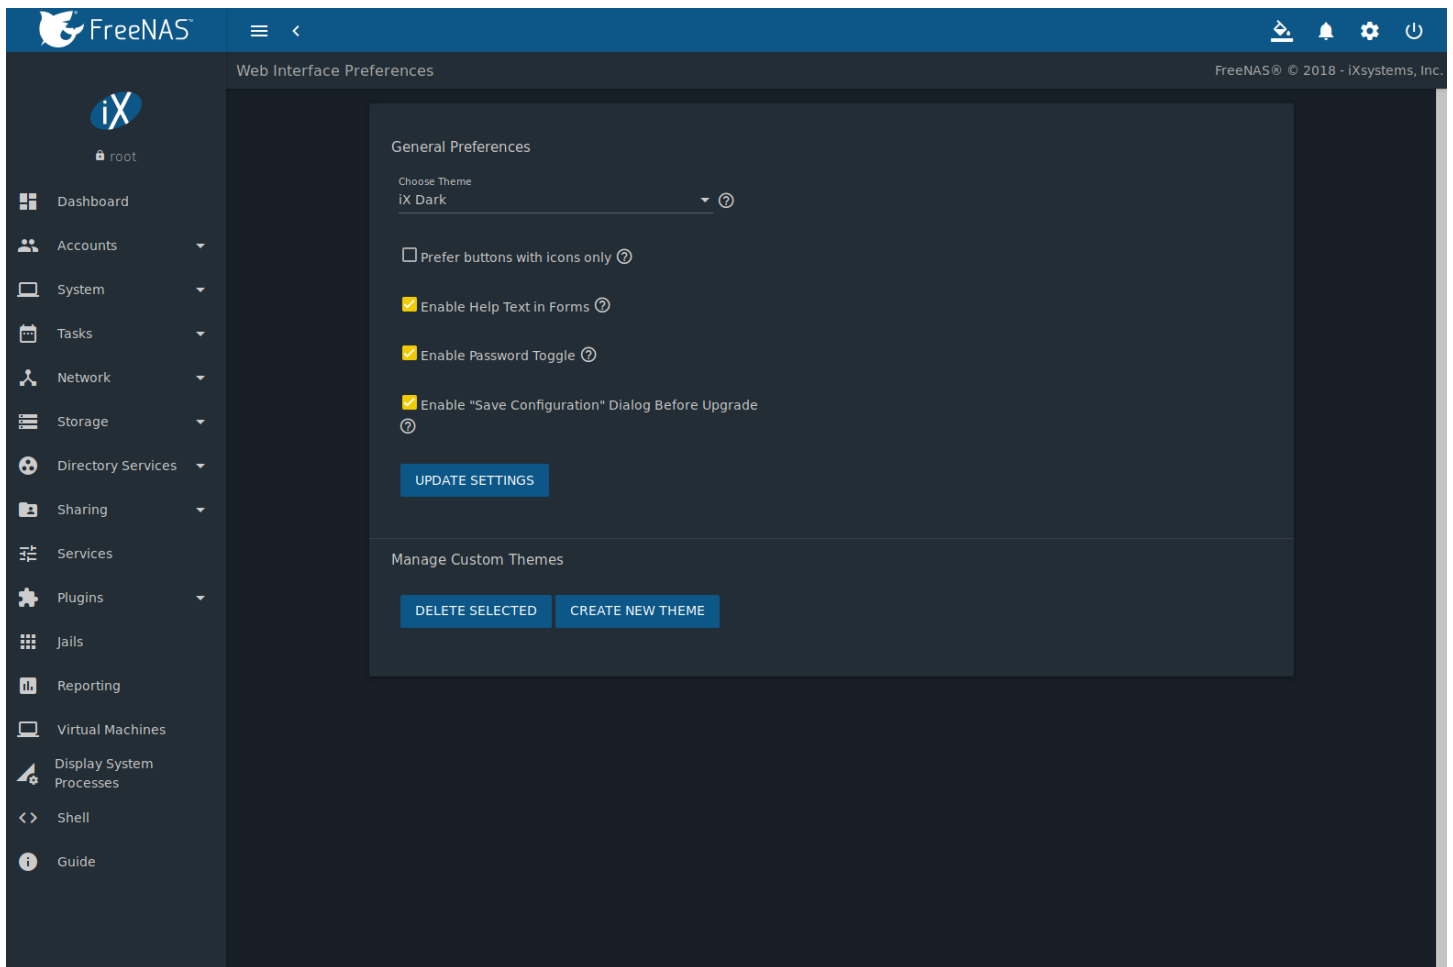


Fig. 4.1: Web Interface Preferences

These options are applied to the entire web interface:

- *Choose Theme*: Change the active theme. Custom themes are added to this list.
- *Enable Help Text in Forms*: Set to add viewable help text to each form in the web interface. Unset to hide all help text icons.
- *Enable Password Toggle*: Set to add the option to toggle between hidden or visible text for passwords in forms.
- *Enable "Save Configuration" Dialog Before Upgrade*: Shows a popup window to save the system configuration file on system upgrade.

Make any changes and click *UPDATE SETTINGS* to save the new selections.

4.3.2 Themes

The FreeNAS® web interface supports dynamically changing the active theme and creating new, fully customizable themes.

4.3.2.1 Theme Selector

Quickly change the active theme by using the theme selector. Look for the paint bucket icon in the upper-right corner of the web interface. Click the icon to see a list of different default and favorite themes. [Figure 4.2](#) shows an example:

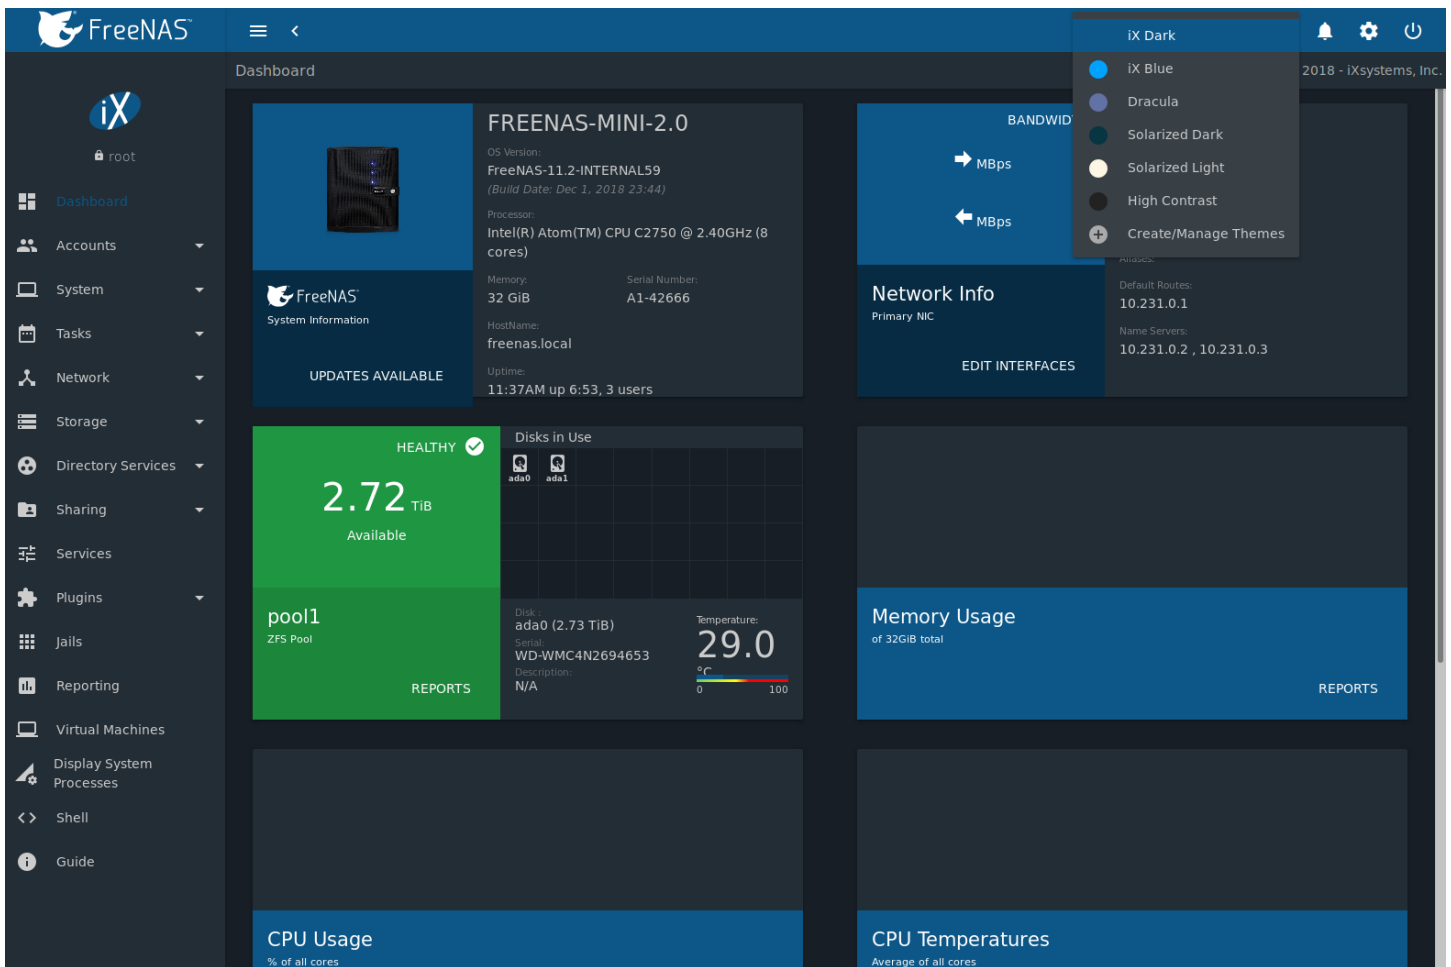


Fig. 4.2: Changing the FreeNAS® web interface theme

Click a theme to activate it.

Select *Manage Themes* to open the *Web Interface Preferences* page. The *Manage Custom Themes* column displays any created custom themes. Delete these themes by setting the options and clicking *DELETE SELECTED*.

Click *CREATE NEW THEME* to go to the *Create Custom Theme* page.

4.3.2.2 Create New Themes

This page is used to create and preview custom FreeNAS® themes. Figure 4.3 shows many of the theming and preview options:

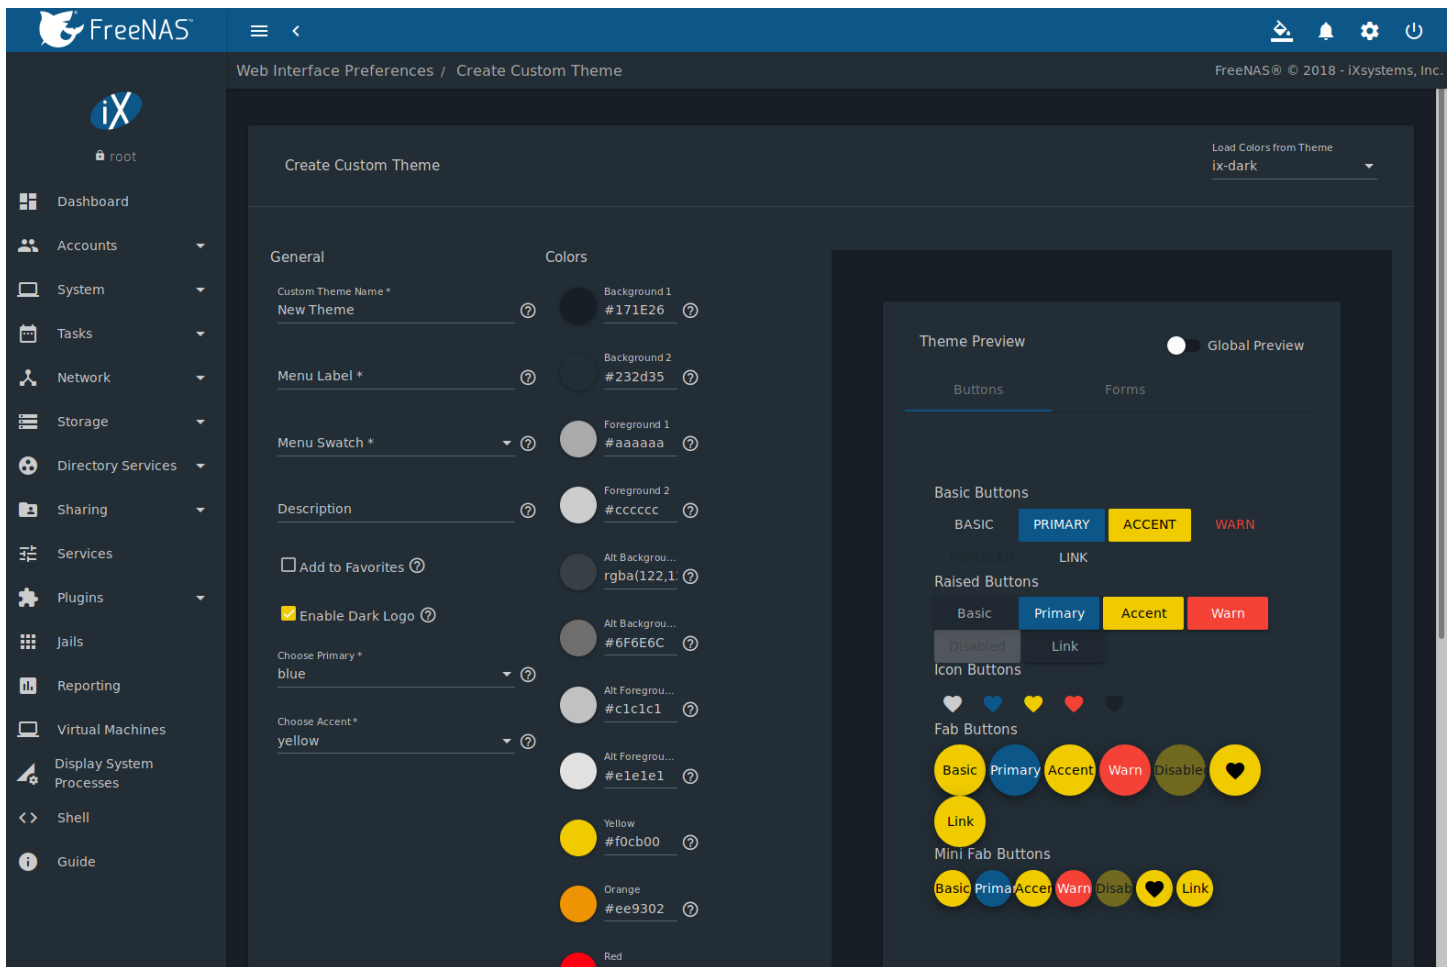


Fig. 4.3: Create and Preview a Custom Theme

Select an existing theme from the *Load Colors from Theme* drop-down menu in the upper-right to use the colors from that theme as the starting values for the new custom theme. [Table 4.1](#) describes each option:

Table 4.1: General Options for a New Theme

Setting	Value	Description
Custom Theme Name	string	Enter a name to identify the new theme.
Menu Label	string	Enter a short name to use for the FreeNAS® menus.
Menu Swatch	drop-down menu	Choose a color from the theme to display next to the menu entry of the custom theme.
Description	string	Enter a short description of the new theme.
Add to Favorites	checkbox	Set to add this theme to the Theme Selector (page 60).
Enable Dark Logo	checkbox	Set this to give the FreeNAS Logo a dark fill color.
Choose Primary	drop-down menu	Choose from either a generic color or import a specific color setting to use as the primary theme color. The primary color changes the top bar of the web interface and the color of many of the buttons.
Choose Accent	drop-down menu	Choose from either a generic color or import a specific color setting to use as the accent color for the theme. This color is used for many of the buttons and smaller elements in the web interface.

Choose the different *Colors* for this new theme after setting these general options. Click the color swatch to open a small

popup with sliders to adjust the color. Color values can also be entered as a hexadecimal value.


Changing any color value automatically updates the *Theme Preview* column. This section is completely interactive and shows how the custom theme is applied to all the different elements in the web interface.

Click *SAVE CUSTOM THEME* when finished with all the *General* and *Colors* options. The new theme will be immediately added to the list of available themes in *Web Interface Preferences*.


Click *Global Preview* to apply the unsaved custom theme to the current session of the FreeNAS® web interface. Activating *Global Preview* allows going to other pages in the web interface and live testing the new custom theme.

Note: Setting a custom theme as a *Global Preview* does **not** save that theme! Be sure to go back to *Preferences* → *Create Custom Theme*, complete any remaining options, and click *SAVE CUSTOM THEME* to save the current settings as a new theme.

4.4 About

Click  (Settings) and *About* to view a popup window with basic system information. This includes system *Version*, *Hostname*, *Uptime*, *IP* address, *Physical Memory*, *CPU Model*, and *Average Load*.

4.5 Legacy Web Interface

Click  (Settings) and *Legacy Web Interface* to switch to the previous FreeNAS® web interface. A popup window asks to confirm the choice. Click *CONTINUE* to log out and go to the log in screen for the Legacy web interface.

ACCOUNTS

Accounts is used to manage users and groups. This section contains these entries:

- [Groups](#) (page 64): used to manage UNIX-style groups on the FreeNAS® system.
- [Users](#) (page 67): used to manage UNIX-style accounts on the FreeNAS® system.

Each entry is described in more detail in this section.

5.1 Groups

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

Note: It is unnecessary to recreate the network users or groups when a directory service is running on the same network. Instead, import the existing account information into FreeNAS®. Refer to [Directory Services](#) (page 183) for details.

This section describes how to create a group and assign user accounts to it. The next section, [Users](#) (page 67), describes creating user accounts.

Click *Accounts* → *Groups* to see a screen like [Figure 5.1](#).

Accounts / Groups

FreeNAS® © 2018 - iXsystems, Inc.

Groups

Filter Groups

COLUMNS

ADD

Group	GID	Builtin	Permit Sudo	
data1	1000	no	no	⋮
user1	1001	no	no	⋮
wheel	0	yes	no	⋮
daemon	1	yes	no	⋮
kmem	2	yes	no	⋮
sys	3	yes	no	⋮
tty	4	yes	no	⋮
operator	5	yes	no	⋮
mail	6	yes	no	⋮
bin	7	yes	no	⋮
news	8	yes	no	⋮
man	9	yes	no	⋮
games	13	yes	no	⋮

1 - 13 of 42

1 2 3 4

Fig. 5.1: Group Management

The *Groups* page lists all groups, including those built in and used by the operating system. The table displays group names, group IDs (GID), built-in groups, and whether `sudo` is permitted. Clicking the ⋮ (Options) icon on a user-created group entry displays *Members*, *Edit*, and *Delete* options. Click *Members* to view and modify the group membership. Built-in groups are required by the FreeNAS® system, so they do not have *Edit* or *Delete* buttons.

The *ADD* button opens the screen shown in [Figure 5.2](#). [Table 5.1](#) summarizes the available options when creating a group.

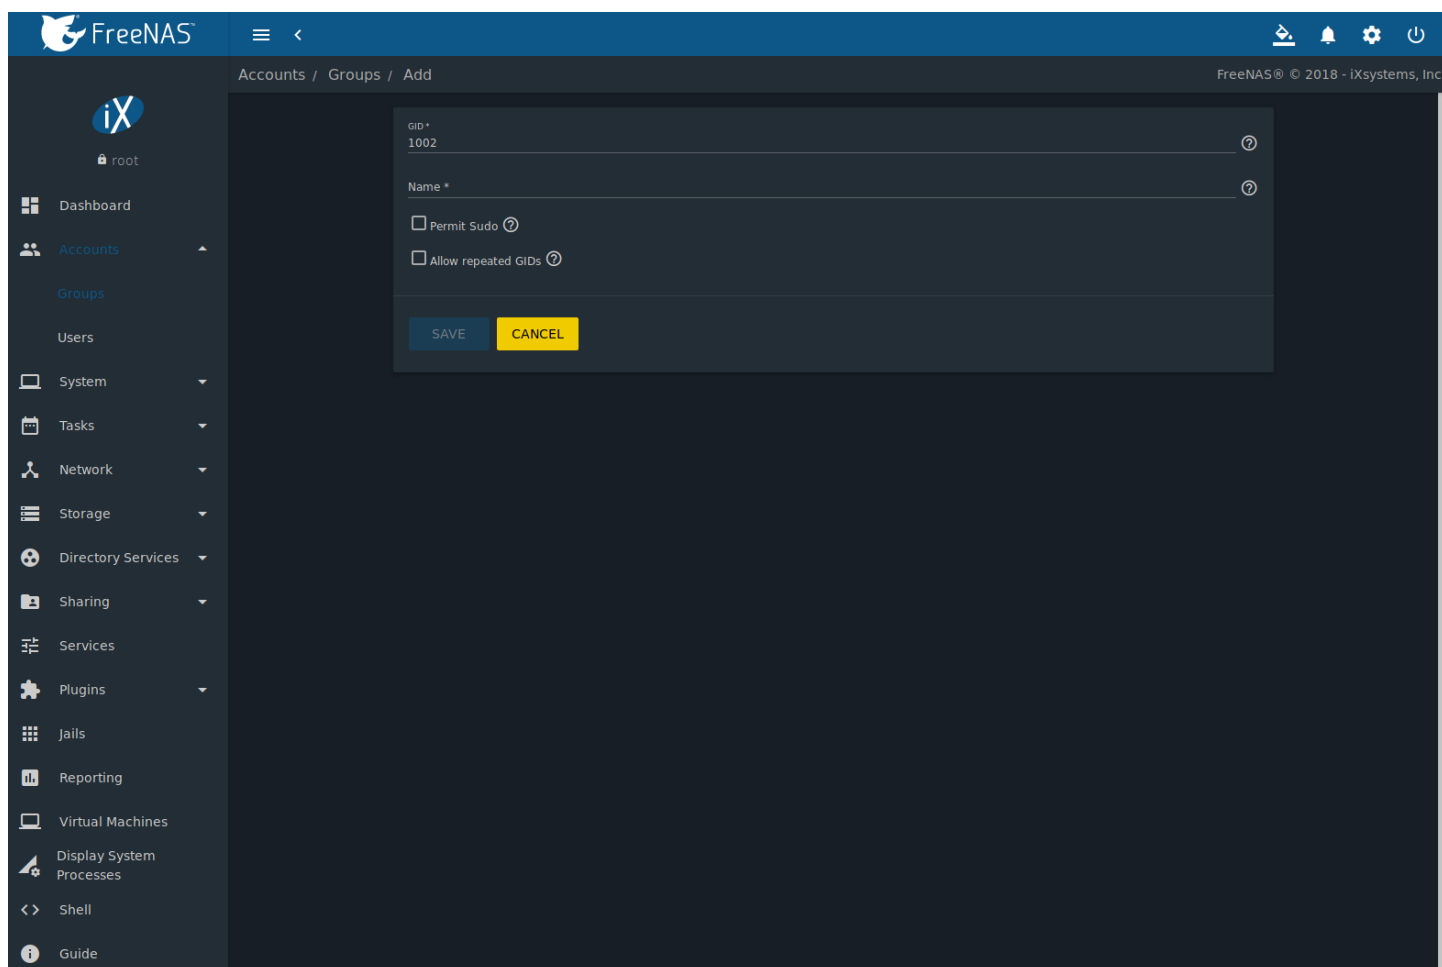


Fig. 5.2: Creating a New Group

Table 5.1: Group Creation Options

Setting	Value	Description
GID	string	The next available group ID is suggested. 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. Example: the <code>sshd</code> group has an ID of 22.
Name	string	Enter a name for the new group.
Permit Sudo	checkbox	Set to allow group members to use <code>sudo</code> (https://www.sudo.ws/). When using <code>sudo</code> , a user is prompted for their own password.
Allow repeated GIDs	checkbox	Set to allow 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 added to a group. Click ⓘ (Options) on the desired group then *Members*. Select the users in the *Members* list. This list shows all user accounts on the system. Next, click -> to move the users to the right frame. Press *SAVE* to add the users on the right frame to the group.

Figure 5.3, shows *user1* added as a member of group *data1*.

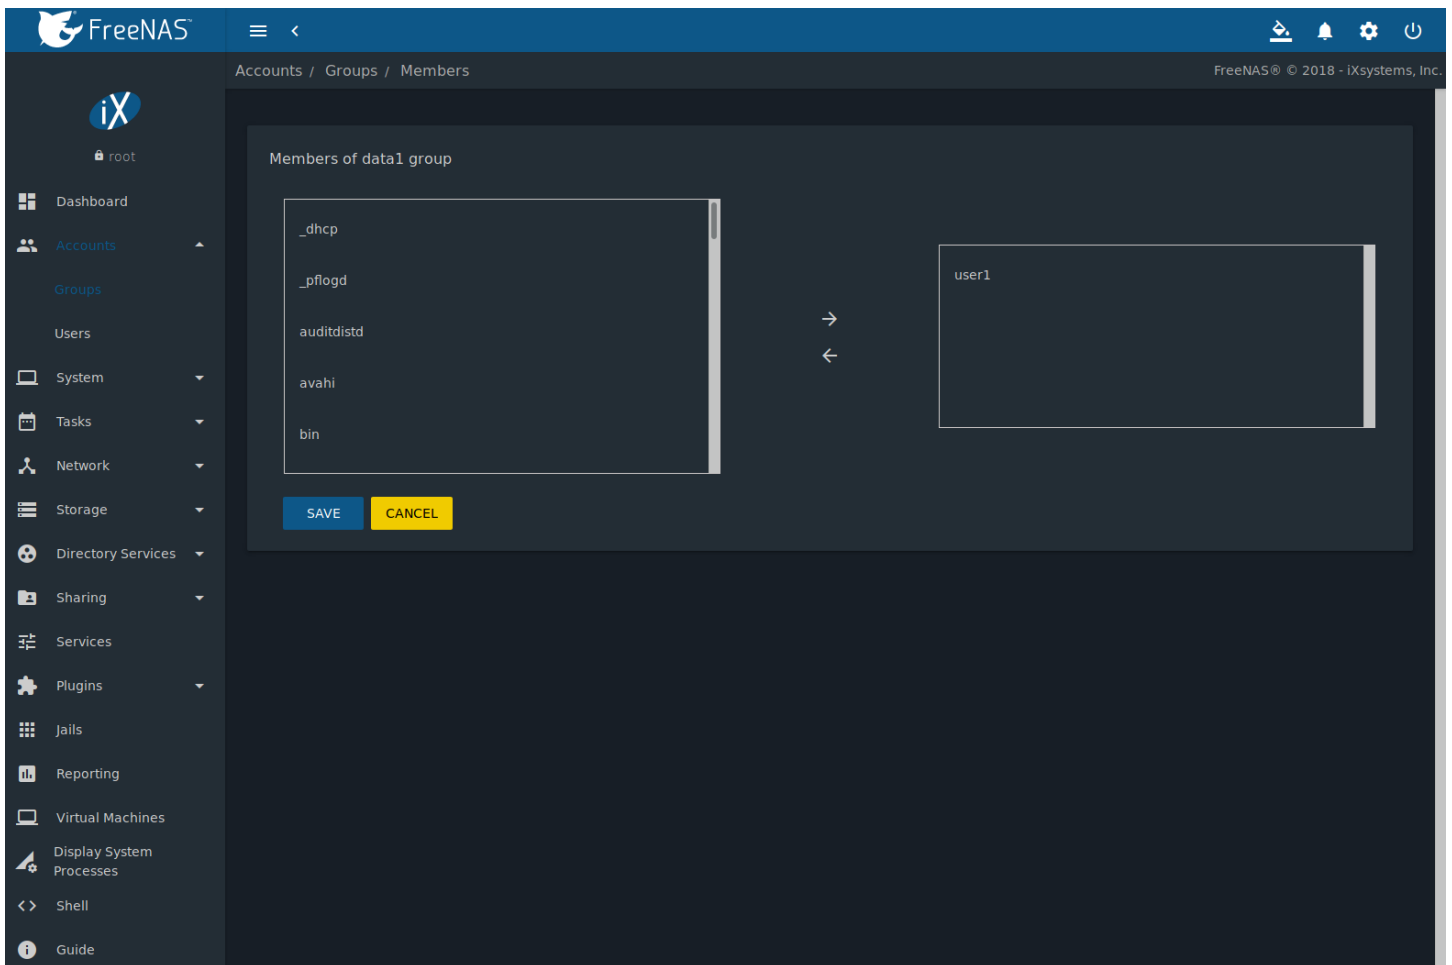


Fig. 5.3: Assigning a User to a Group

The *Delete* button deletes a group. The pop-up message asks if all users with this primary group should also be deleted, and to confirm the action. Note built-in groups do not have a *Delete* button.

5.2 Users

FreeNAS® supports users, groups, and permissions, allowing flexibility in configuring which users have access to the data stored on FreeNAS®. To assign permissions to shares, select one of these options:

1. Create a guest account for all users, or create a user account for every user in the network where the name of each account is the same as a login 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 the network uses a directory service, import the existing account information using the instructions in [Directory Services](#) (page 183).

Accounts → *Users* lists all system accounts installed with the FreeNAS® operating system, as shown in [Figure 5.4](#).

The screenshot shows the FreeNAS web interface with the 'Users' management page. The page has a dark theme. On the left is a sidebar with navigation icons and labels. The main content area shows a table of users. The table has columns: Username, Home directory, Shell, Full Name, and Lock User. The table lists several system users. At the top right of the table area, there is a search bar labeled 'Filter Users', a 'COLUMNS' dropdown, and an 'ADD' button. At the bottom right of the table, there is a pagination control showing '1 - 11 of 34' and a set of navigation arrows.

Username	Home directory	Shell	Full Name	Lock User
user1	/mnt/pool1/smb_user1	/bin/csh	user1	no
root	/root	/usr/local/bin/zsh	root	no
daemon	/root	/usr/sbin/nologin	Owner of many system f	no
operator	/	/usr/sbin/nologin	System &	no
bin	/	/usr/sbin/nologin	Binaries Commands anc	no
tty	/	/usr/sbin/nologin	Tty Sandbox	no
kmem	/	/usr/sbin/nologin	KMem Sandbox	no
games	/	/usr/sbin/nologin	Games pseudo-user	no
news	/	/usr/sbin/nologin	News Subsystem	no
man	/usr/share/man	/usr/sbin/nologin	Mister Man Pages	no
ftp	/nonexistent	/bin/csh		no

Fig. 5.4: Managing User Accounts

By default, each user entry displays the username, home directory, default shell, the user full name, and if the user is locked. This table is adjustable by setting the different column checkboxes above it. Set *Toggle* to display all options in the table.

Clicking a column name sorts the list by that value. An arrow indicates which column controls the view sort order. Click the arrow to reverse the sort order.

Click  (Options) on the user created account to display the *Edit* and *Delete* buttons. Note built-in users do not have a *Delete* button.

Note: Setting the email address for the built-in *root* user account is recommended 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)` (<https://www.freebsd.org/cgi/man.cgi?query=nologin>). For security reasons and to prevent breakage of system services, modifying the system accounts is discouraged.

The *ADD* button opens the screen shown in Figure 5.5. Table 5.2 summarizes the options that are available when user accounts are created or modified.

Warning: When using *Active Directory* (page 183), Windows user passwords must be set from within Windows.

FreeNAS® 11.2-RELEASE User Guide, Release 11.2

Accounts / Users / Add

FreeNAS® © 2018 - iXsystems, Inc.

Name & Contact

Full Name *

Username *

Email *

Password *

Confirm Password *

ID & Groups

User ID *

1001

☒ New Primary Group

Directories & Permissions

Home Directory

/nonexistent

Home Directory Permissions

	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"/>

Authentication

SSH Public Key

Enable password login

Yes

Shell

csh

☐ Lock User

☐ Permit Sudo

☐ Microsoft Account

SAVE CANCEL

Fig. 5.5: Adding or Editing a User Account

Table 5.2: User Account Configuration

Setting	Value	Description
Username	string	Enter an alphanumeric username of eight to sixteen characters. Keeping usernames to eight characters or less is recommended for compatibility with legacy clients. Note that \$ can only be used as the last character. Usernames cannot begin with a hyphen – or contain a space, tab, or these characters: , : + & # % ^ & () ! @ ~ * ? < > =
Full Name	string	This field is mandatory and may contain spaces.
Email	string	The email address associated with the account.
Password	string	Mandatory unless <i>Enable password login</i> is <i>No</i> . Cannot contain a ?. Click ⓘ (Show) to view or obscure the password characters.
Confirm Password	string	Required to match the value of <i>Password</i> .
User ID	integer	Grayed out if the user already exists. When creating an account, the next numeric ID is 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.
New Primary Group	checkbox	Set by default to create a new a primary group with the same name as the user. Unset to select a different primary group name.

Continued on next page

Table 5.2 – continued from previous page

Setting	Value	Description
Primary Group	drop-down menu	Unset <i>New Primary Group</i> to access this menu. For security reasons, FreeBSD will not give a user <code>su</code> permissions if <i>wheel</i> is their primary group. To give a user <code>su</code> access, add them to the <i>wheel</i> group in <i>Auxiliary groups</i> .
Auxiliary groups	drop-down menu	Select which groups the user will be added to.
Home Directory	browse button	Browse to the name of an existing pool or dataset where the user will be assigned permission to access.
Home Directory Permissions	checkboxes	Sets default Unix permissions of user's home directory. This is read-only for built-in users.
SSH Public Key	string	Paste the user's public SSH key to be used for key-based authentication. Do not paste the private key!
Enable password login	checkbox	Set to disable password logins and authentication to SMB shares. To undo this setting, set a password for the user with the <i>Edit</i> button for the user in <i>Users</i> . Setting this option grays out <i>Lock user</i> and <i>Permit Sudo</i> , which are mutually exclusive.
Shell	drop-down menu	Select the shell to use for local and SSH logins. See Table 5.3 for an overview of available shells.
Lock User	checkbox	Set to prevent the user from logging in until the account is unlocked. Setting this option grays out <i>Disable password login</i> , which is mutually exclusive.
Permit Sudo	checkbox	Set to allow members of the group to use <code>sudo</code> (https://www.sudo.ws/). When using <code>sudo</code> , a user is prompted for their own password.
Microsoft Account	checkbox	Set if the user is connecting from a Windows 8 or higher system.

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

Table 5.3: Available Shells

Shell	Description
<code>csh</code>	<i>C shell</i> (https://en.wikipedia.org/wiki/C_shell)
<code>sh</code>	<i>Bourne shell</i> (https://en.wikipedia.org/wiki/Bourne_shell)
<code>tcsh</code>	<i>Enhanced C shell</i> (https://en.wikipedia.org/wiki/Tcsh)
<code>bash</code>	<i>Bourne Again shell</i> (https://en.wikipedia.org/wiki/Bash_%28Unix_shell%29)
<code>ksh93</code>	<i>Korn shell</i> (http://www.kornshell.com/)
<code>mksh</code>	<i>mirBSD Korn shell</i> (https://www.mirbsd.org/mksh.htm)
<code>rbash</code>	<i>Restricted bash</i> (http://www.gnu.org/software/bash/manual/html_node/The-Restricted-Shell.html)
<code>rzsh</code>	<i>Restricted zsh</i> (http://www.csse.uwa.edu.au/programming/linux/zsh-doc/zsh_14.html)
<code>scponly</code>	Select <i>scponly</i> (https://github.com/scponly/scponly/wiki) to restrict the user's SSH usage to only the <code>scp</code> and <code>sftp</code> commands.
<code>zsh</code>	<i>Z shell</i> (http://www.zsh.org/)
<code>git-shell</code>	<i>restricted git shell</i> (https://git-scm.com/docs/git-shell)
<code>nologin</code>	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> .

Built-in user accounts needed by the system cannot be removed. A *Delete* button appears for custom users that were added by the system administrator. Clicking *Delete* opens a popup window to confirm the action and offer an option to keep the user primary group when the user is deleted.

SYSTEM

The System section of the administrative web interface contains these entries:

- [General](#) (page 71) configures general settings such as HTTPS access, the language, and the timezone
- [NTP Servers](#) (page 73) adds, edits, and deletes Network Time Protocol servers
- [Boot Environments](#) (page 75) creates, renames, and deletes boot environments. It also shows the condition of the Boot Pool.
- [Advanced](#) (page 79) configures advanced settings such as the serial console, swap space, and console messages
- [Email](#) (page 83) configures the email address to receive notifications
- [System Dataset](#) (page 84) configures the location where logs and reporting graphs are stored
- [Alert Services](#) (page 85) configures services used to notify the administrator about system events.
- [Alert Settings](#) (page 87) lists the available [Alert](#) (page 330) conditions and provides configuration of the notification frequency for each alert.
- [Cloud Credentials](#) (page 88) is used to enter connection credentials for remote cloud service providers
- [Tunables](#) (page 91) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- [Update](#) (page 94) performs upgrades and checks for system updates
- [CAs](#) (page 97): import or create internal or intermediate CAs (Certificate Authorities)
- [Certificates](#) (page 101): import existing certificates or create self-signed certificates
- [Support](#) (page 106): report a bug or request a new feature.

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

6.1 General

System → *General* is shown in [Figure 6.1](#).

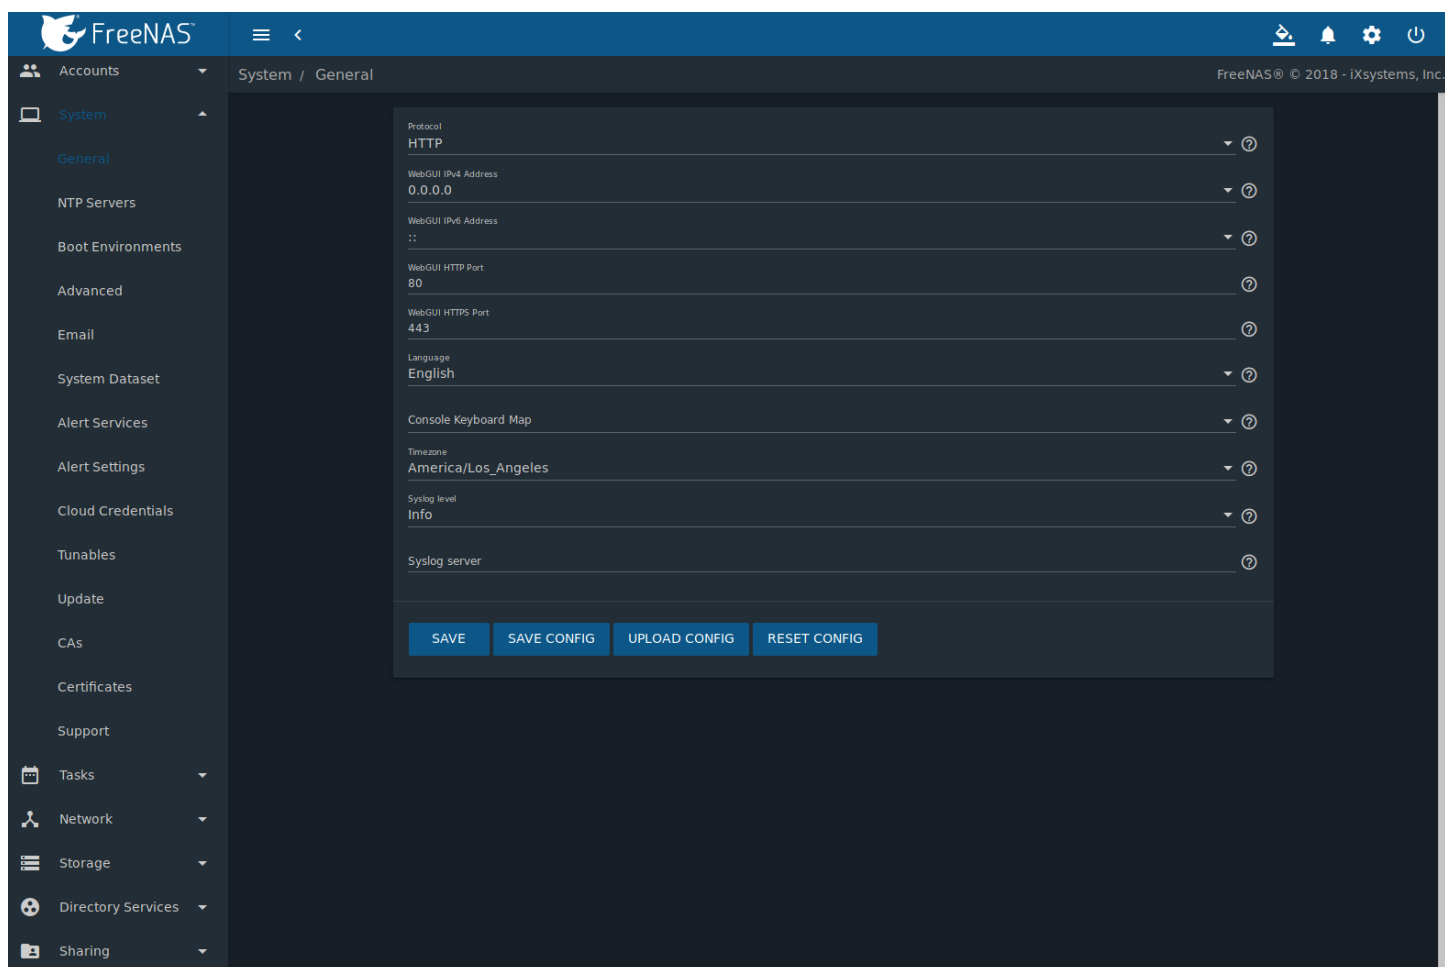


Fig. 6.1: General Screen

Table 6.1 summarizes the configurable settings in the General tab:

Table 6.1: General Configuration Settings

Setting	Value	Description
Protocol	drop-down menu	Set the web protocol to use when connecting to the web interface from a browser. To change the default <i>HTTP</i> to <i>HTTPS</i> or to <i>HTTP+HTTPS</i> , select a certificate in <i>GUI SSL Certificate</i> . If there are no certificates, create a CA (page 97) then a certificate (page 101).
WebGUI IPv4 Address	drop-down menu	Choose a recent IP addresses to limit the usage when accessing the web interface. The built-in HTTP server binds to the wildcard address of <i>0.0.0.0</i> (any address) and issues an alert if the specified address becomes unavailable.
WebGUI IPv6 Address	drop-down menu	Choose a recent IPv6 addresses to limit the usage when accessing the web interface. The built-in HTTP server binds to any address and issues an alert if the specified address becomes unavailable.
WebGUI HTTP Port	integer	Allow configuring a non-standard port for accessing the web interface over HTTP. Changing this setting can also require changing a Firefox configuration setting (https://www.redbrick.dcu.ie/~d_fens/articles/Firefox:_This_Address_is_Restricted).
WebGUI HTTPS Port	integer	Allow configuring a non-standard port for accessing the web interface over HTTPS.

Continued on next page

Table 6.1 – continued from previous page

Setting	Value	Description
GUI SSL Certificate	drop-down menu	Required for <i>HTTPS</i> . Browse to the location of the certificate to use for encrypted connections.
WebGUI HTTP -> HTTPS Redirect	checkbox	Set to redirect <i>HTTP</i> connections to <i>HTTPS</i> . <i>HTTPS</i> must be selected in <i>Protocol</i> .
Language	drop-down menu	Select a language. View the status of a language in the webui GitHub repository (https://github.com/freenas/webui/tree/master/src/assets/i18n) Refer to Contributing to FreeNAS® (page 334) for more information about supported languages.
Console Keyboard Map	drop-down menu	Select a keyboard layout.
Timezone	drop-down menu	Select a timezone.
Syslog level	drop-down menu	When <i>Syslog server</i> is defined, only logs matching this level are sent.
Syslog server	string	Select an <i>IP address_or_hostname:optional_port_number</i> to send logs to. Set to write log entries to both the console and the remote server.

After making any changes, click the *SAVE* button.

This screen also contains these buttons: **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 263) 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, Active Directory bind credentials, and cloud 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 new screen gives the option to *Include 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 *Include 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.

Reset Config: reset the configuration database to the default base version. This does not delete user *SSH* keys or any other data stored in a user 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.

6.2 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 the 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 physical location of the FreeNAS® system.

Click *System* → *NTP Servers* and *ADD* to add an NTP server. Figure 6.2 shows the configuration options. Table 6.2 summarizes the options available when adding or editing an NTP server. [ntp.conf\(5\)](https://www.freebsd.org/cgi/man.cgi?query=ntp.conf(5)) ([https://www.freebsd.org/cgi/man.cgi?query=ntp.conf\(5\)](https://www.freebsd.org/cgi/man.cgi?query=ntp.conf(5))) explains these options in more detail.

The screenshot shows the FreeNAS web interface with the 'System / NTP Servers / Add' page. The left sidebar contains a navigation menu with categories like Accounts, System, Tasks, Network, Storage, Directory Services, and Sharing. The main content area displays a form for adding a new NTP server. The form fields are: Address (text input), Burst (checkbox), IBurst (checkbox), Prefer (checkbox), Min. Poll (text input with value 6), Max. Poll (text input with value 10), and Force (checkbox). There are 'SAVE' and 'CANCEL' buttons at the bottom of the form. The top of the interface shows the FreeNAS logo and a copyright notice for 2018 - iXsystems, Inc.

Fig. 6.2: Add an NTP Server

Table 6.2: NTP Servers Configuration Options

Setting	Value	Description
Address	string	Enter the hostname or IP address of the NTP server.
Burst	checkbox	Recommended when <i>Max. Poll</i> is greater than 10. Only use on personal servers. Do not use with a public NTP server.
IBurst	checkbox	Speed up the initial synchronization, taking seconds rather than minutes.
Prefer	checkbox	This option is only recommended for highly accurate NTP servers, such as those with time monitoring hardware.
Min. Poll	integer	Minimum polling time in seconds. Must be a power of 2, and cannot be lower than 4 or higher than <i>Max. Poll</i> .
Max. Poll	integer	Maximum polling time in seconds. Must be a power of 2, and cannot be higher than 17 or lower than <i>Min. Poll</i> .

Continued on next page

Table 6.2 – continued from previous page

Setting	Value	Description
Force	checkbox	Force the addition of the NTP server, even if it is currently unreachable.

6.3 Boot Environments

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 an update fails, reboot the system and select the previous boot environment, using the instructions in *If Something Goes Wrong* (page 33), 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 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 6.3](#), FreeNAS® displays the condition and statistics of the *Boot Pool*. It also shows the two boot environments that 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 *Initial-Install* boot environment can be booted into if the system needs to be returned to a non-configured version of the installation.

The screenshot shows the FreeNAS web interface. The top navigation bar includes the FreeNAS logo and icons for search, notifications, settings, and power. The left sidebar contains a menu with categories like Accounts, System, Tasks, Network, Storage, Directory Services, and Sharing. The 'System' category is expanded, showing sub-items like General, NTP Servers, Boot Environments (selected), Advanced, Email, System Dataset, Alert Services, Alert Settings, Cloud Credentials, Tunables, Update, CAs, Certificates, and Support.

The main content area is titled 'System / Boot Environments'. It displays the following information:

- Boot Pool Condition : ONLINE
- Size: 29.5G
- Used: 764M(2%)
- Last Scrub Run on: Never
- Automatic scrub interval (in days): 7

Below this information are two buttons: 'SCRUB BOOT POOL' and 'BOOT POOL STATUS'.

The 'Boot Environments' section features a search bar, a 'Filter Boot Environments' button, and a 'COLUMNS' dropdown. Below this is a table with the following data:

Name	Active	Created	Space	Keep
default	Now/Reboot	12/3/2018 11:34 PM	762.2M	⋮
Initial-Install	-	12/3/2018 11:45 PM	2.2M	⋮

At the bottom of the table, it indicates '1 - 2 of 2'.

Fig. 6.3: 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.
- **Space:** displays the size of the boot environment.
- **Keep:** indicates whether or not this boot environment can be pruned if an update does not have enough space to proceed. Click ⋮ (Options) and *Keep* for an entry if that boot environment should not be automatically pruned.

Click ⋮ (Options) on an entry to see these configuration buttons:

- **Delete:** used to delete the highlighted entry, which also removes that entry from the boot menu. Since an activated entry cannot be deleted, this button does not appear for the active boot environment. 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 does not appear for the *default* boot environment as this entry is needed to return the system to the original installation state.
- **Clone:** makes a new boot environment from the selected boot environment.
- **Rename:** used to change the name of the boot environment.
- **Activate:** only appears on entries which are not currently set to *Active*. Changes the selected entry to the default boot entry on next boot. The status changes to *Reboot* and the current *Active* entry changes from *Now/Reboot* to *Now*, indicating that it was used on the last boot but will not be used on the next boot.

- **Keep:** 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.

There are also other options available.

- **Create:** makes a new boot environment from the active environment. The active boot environment contains the text *Now/Reboot* in the *Active* column. Only alphanumeric characters, underscores, and dashes are allowed in the name.
- **Scrub:** *Scrub Boot Pool* is 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 of the *Boot Environments* screen. 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 *Boot Pool Status* to see the status of the boot devices. [Figure 6.4](#), shows only one boot device, which is *ONLINE*.

Note: Using *Clone* to clone the active boot environment functions the same as using *Create*.

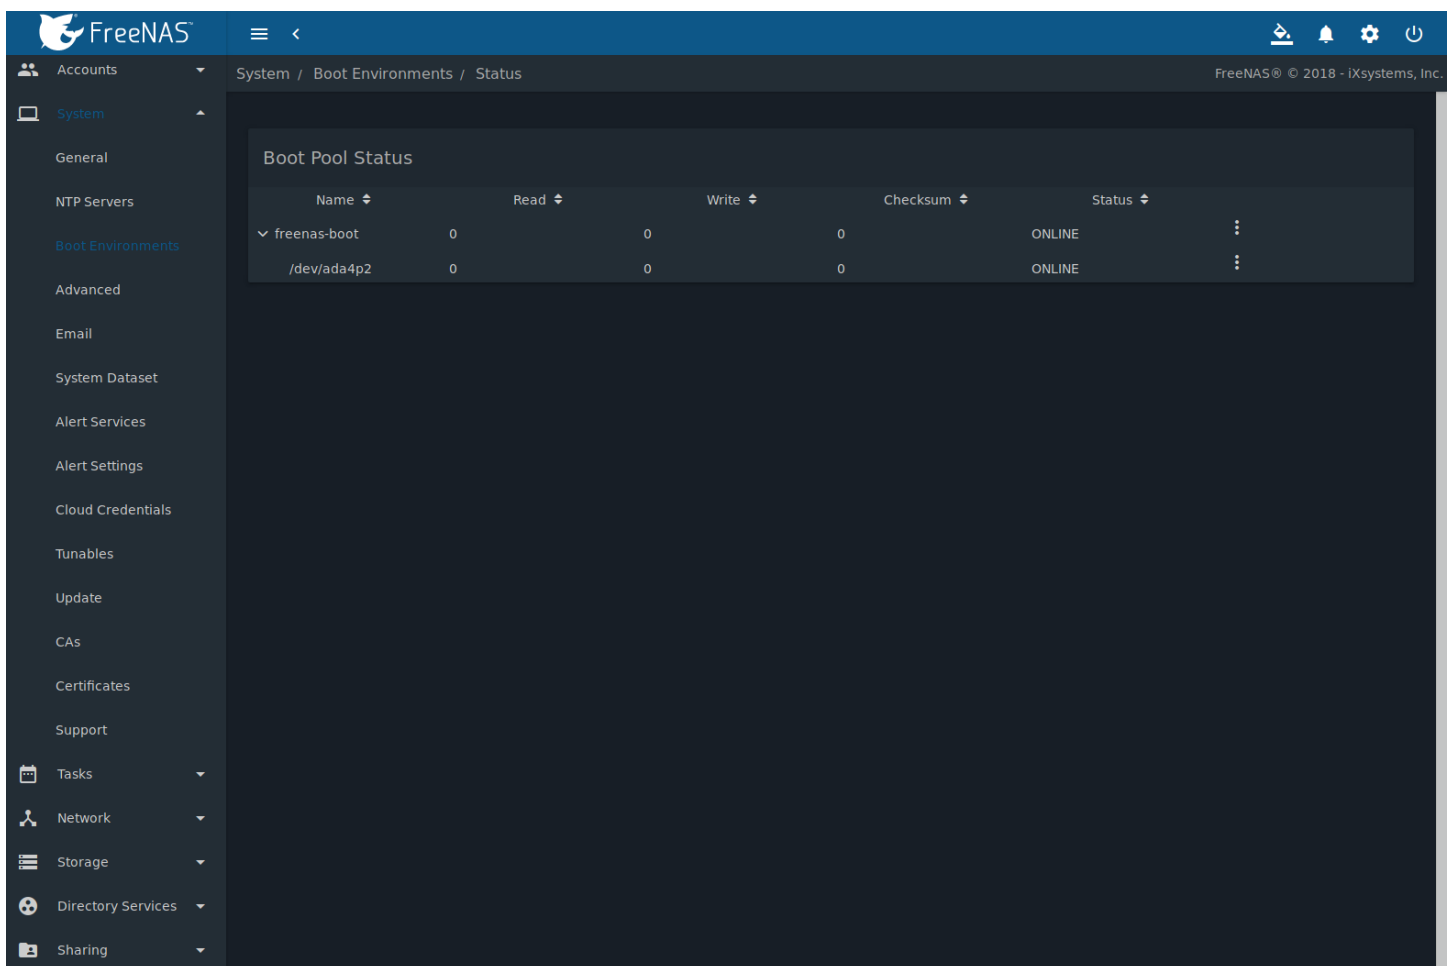



Fig. 6.4: Viewing the Status of the Boot Device

If the system has a mirrored boot pool, there will be a *Detach* option in addition to the *Replace* option. To remove a device from the boot pool, click  (Options) for the device and click *Detach*. 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 **the boot device cannot be replaced if it is the only boot device** because it contains the operating system itself.

6.3.1 Mirroring the Boot Device

If the system is currently booting from a device, another device can be added to create a mirrored boot device. If one device in a mirror fails, the remaining device can still be used to boot the system.

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 6.5](#), the user has gone to *System* → *Boot Environments*, and clicked the *BOOT POOL STATUS* button to display the current status of the boot device. As shown in [Figure 6.4](#), the *freenas-boot* pool is made of a single device, *ada0p2*. There is only one disk, indicated by the word *stripe*. To create a mirrored boot device, click ⓘ (Options) then *attach*. If another device is available, it appears in the *Member disk* drop-down menu. Select the desired device.

The *Use all disk space* option gives control of how much of the new device is made available to ZFS. The new device is partitioned to the same size as the existing device by default. Select *Use all disk space* to use all available space on the new 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 enabled, 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 *SAVE* to attach the new disk to the mirror.

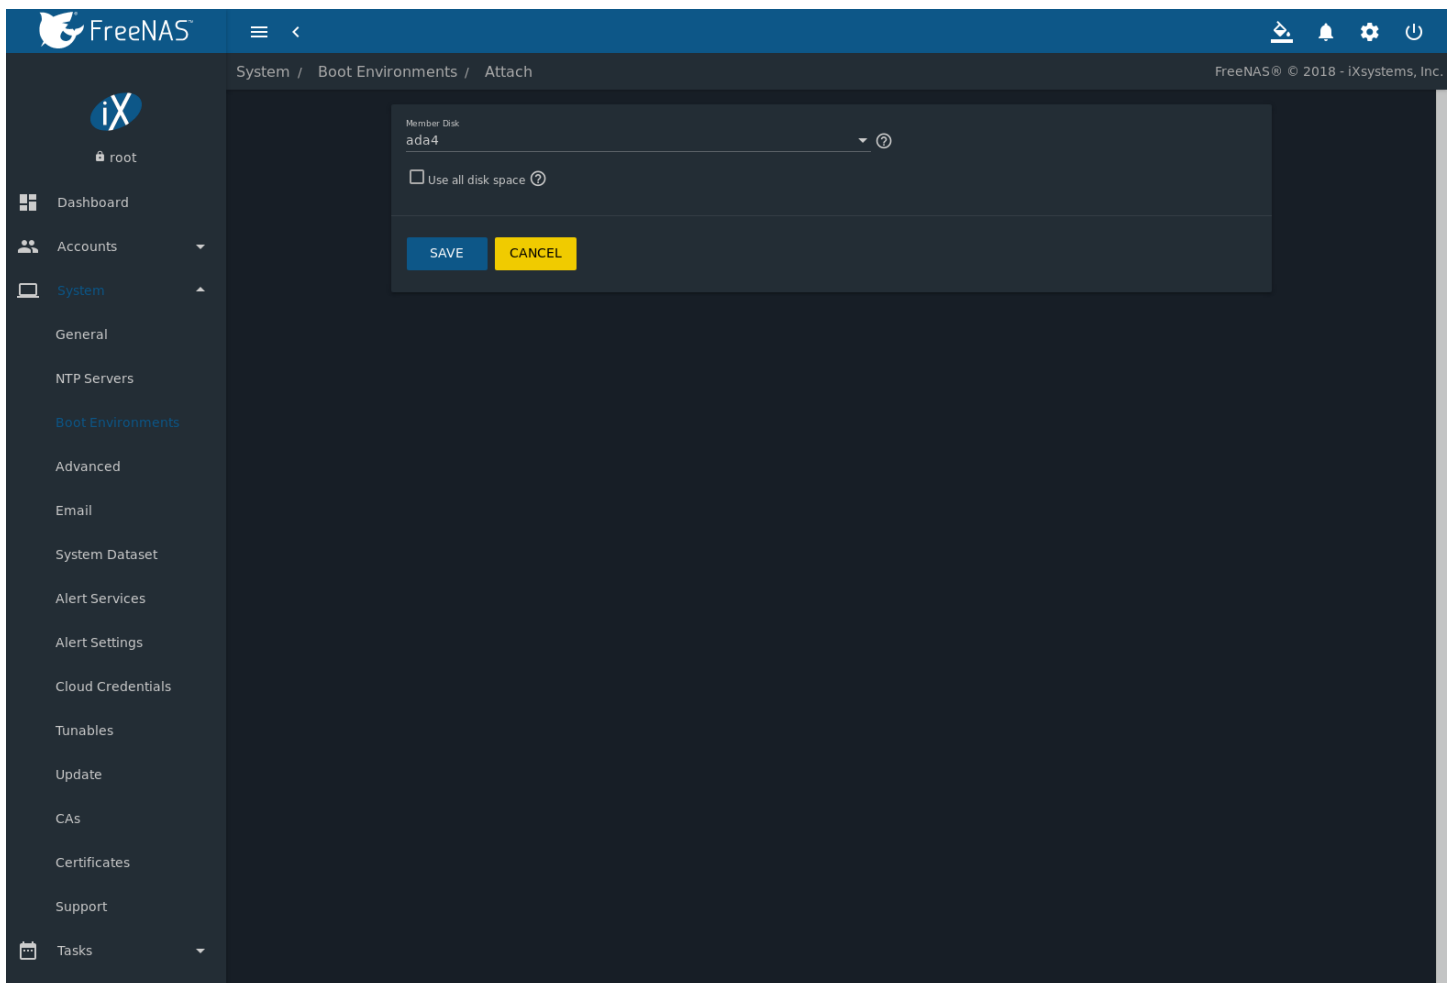


Fig. 6.5: Mirroring a Boot Device

After the mirror is created, the *Boot Pool Status* screen indicates that it is now a *mirror*. The number of devices in the mirror are shown as in [Figure 6.6](#).

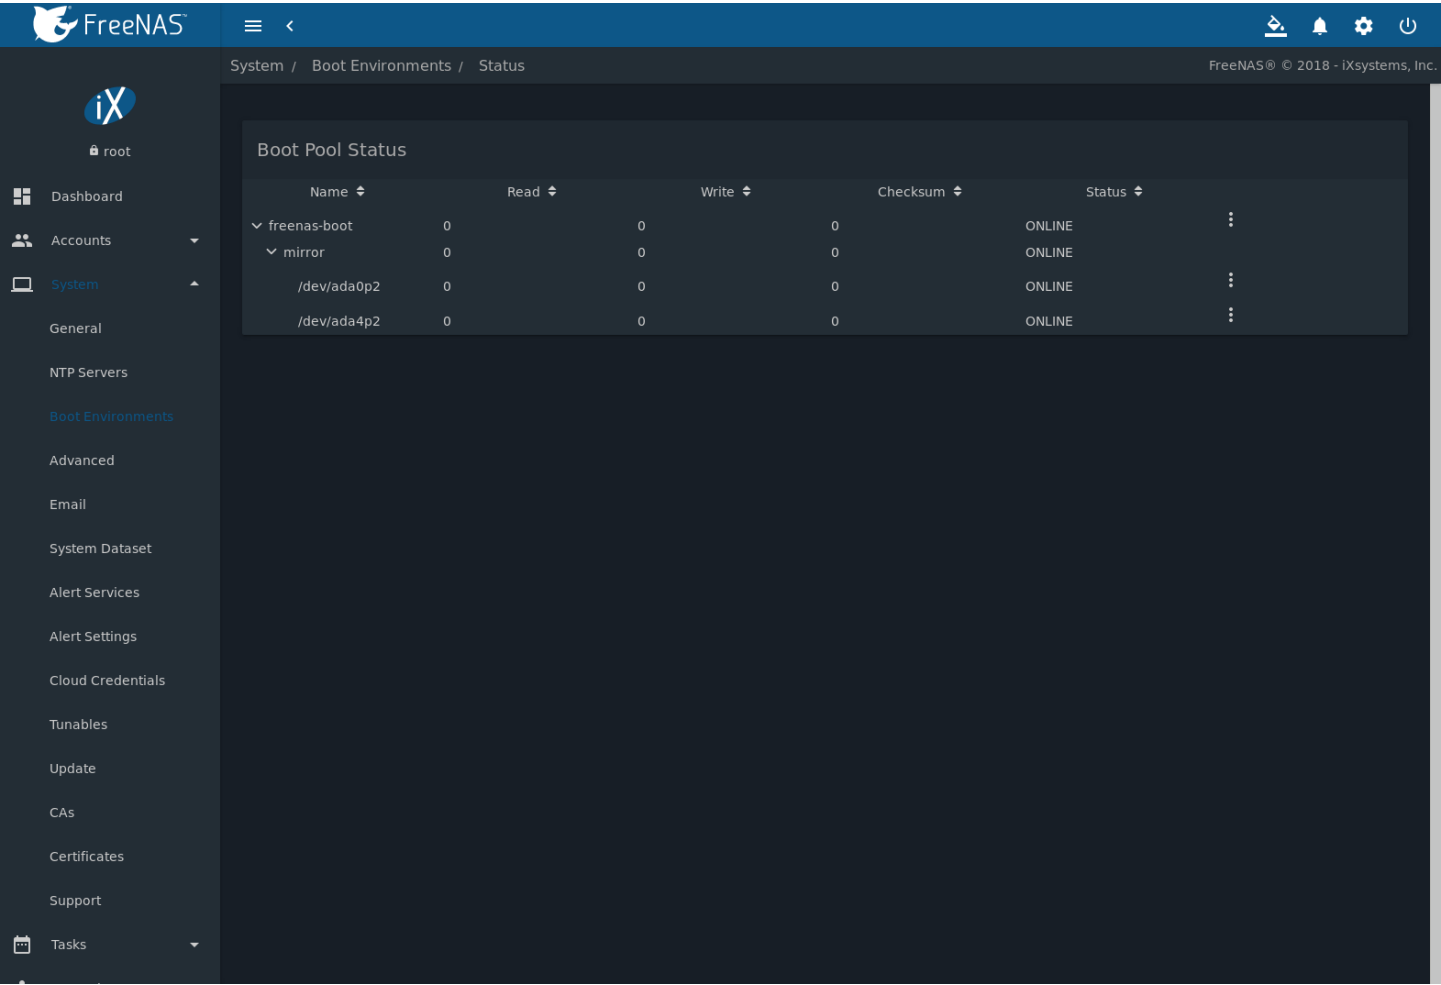


Fig. 6.6: Viewing the Status of a Mirrored Boot Device

6.4 Advanced

System → *Advanced* is shown in [Figure 6.7](#). The configurable settings are summarized in [Table 6.3](#).

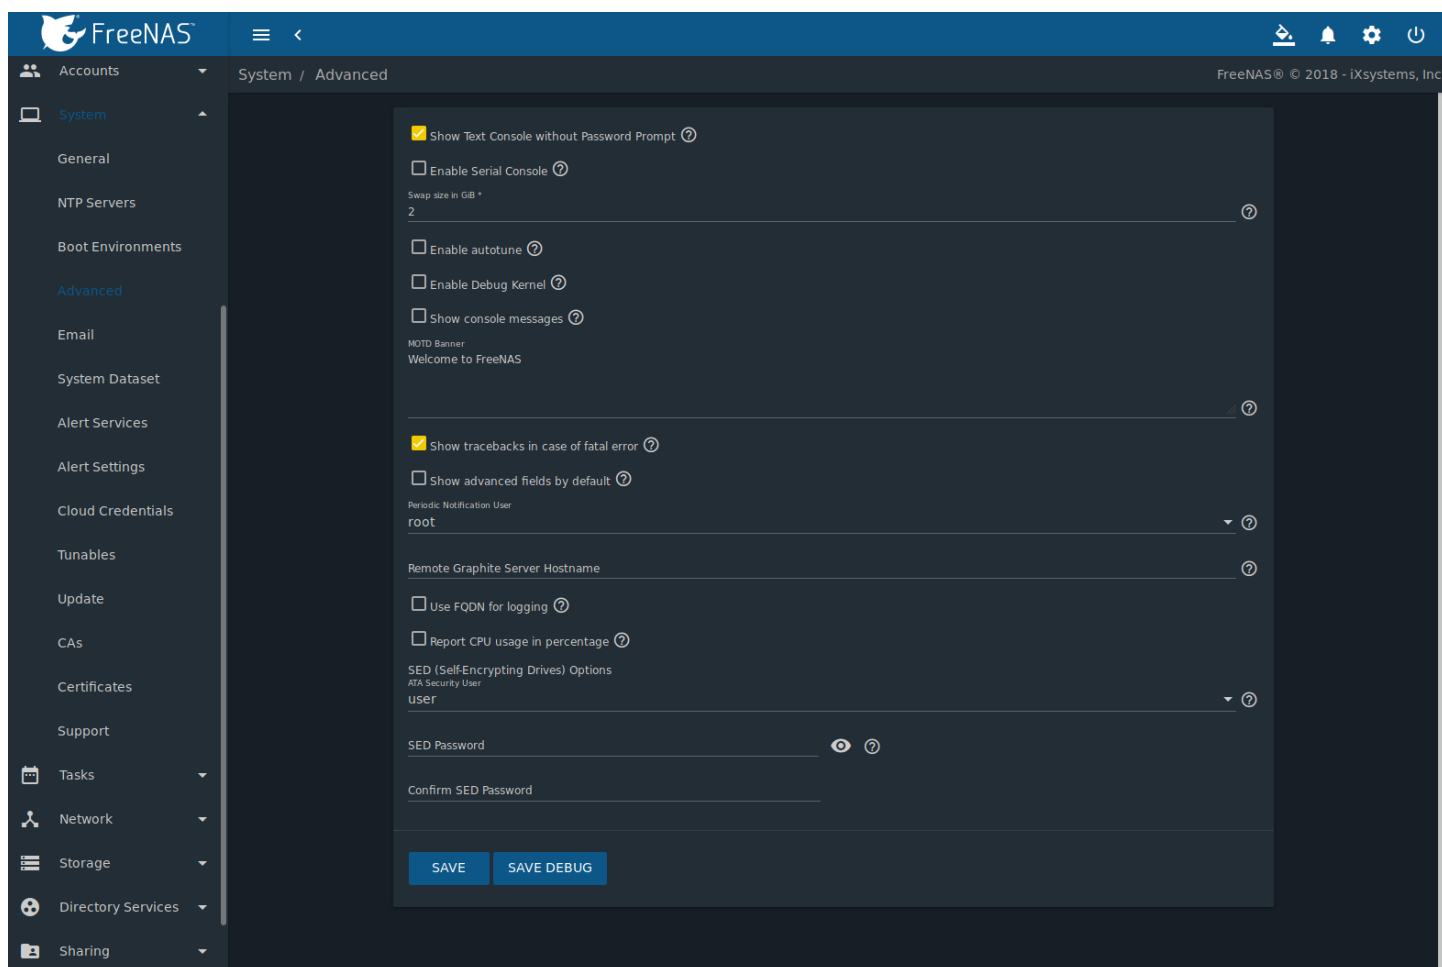


Fig. 6.7: Advanced Screen

Table 6.3: Advanced Configuration Settings

Setting	Value	Description
Show Text Console without Password Prompt	checkbox	Set for the system to immediately display the text console after booting. Unset to require logging into the system before the console menu is shown.
Enable Serial Console	checkbox	Do not enable this option if the serial port is disabled. Adds the <i>Serial Port</i> and <i>Serial Speed</i> fields.
Serial Port	string	Select the serial port address in hex.
Serial Speed	drop-down menu	Select the speed in bps used by the serial port.
Swap size in GiB	non-zero number	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. Setting to 0 disables swap creation completely. This is <i>strongly</i> discouraged.
Enable autotune	checkbox	Enable the Autotune (page 81) script which attempts to optimize the system based on the installed hardware. <i>Warning:</i> Autotuning is only used as a temporary measure and is not a permanent fix for system hardware issues.
Enable Debug Kernel	checkbox	Use a debug version of the kernel on the next boot.

Continued on next page

Table 6.3 – continued from previous page

Setting	Value	Description
Show console messages	checkbox	Set to display console messages in real time at bottom of browser. Click the console to bring up a scrollable screen. Enable the <i>Stop refresh</i> option in the scrollable screen to pause updating and deselect the option to continue to watch the messages as they occur.
MOTD banner	string	This message is shown when a user logs in with SSH.
Show tracebacks in case of fatal error	checkbox	Open a pop-up window of diagnostic information if a fatal error occurs.
Show advanced fields by default	checkbox	Show <i>Advanced Mode</i> fields by default.
Periodic Notification User	drop-down menu	Choose a user to receive security output emails. This output runs nightly, but only sends email when the system reboots or encounters an error.
Remote Graphite Server Hostname	string	IP address or hostname of a remote server running Graphite . (http://graphiteapp.org/)
Use FQDN for logging	checkbox	Include the Fully-Qualified Domain Name (FQDN) in logs to precisely identify systems with similar hostnames.
Report CPU usage in percentage	checkbox	Display CPU usage as percentages in Reporting (page 305).
ATA Security User	drop-down menu	User passed to <code>camcontrol security -u</code> for unlocking SEDs. Values are <i>User</i> or <i>Master</i> .
SED Password	string	Global password used to unlock Self-Encrypting Drives (page 81).
Reset SED Password	checkbox	Select to clear the <i>Password for SED</i> column of <i>Storage</i> → <i>Disks</i> .

Click the *SAVE* button after making any changes.

This tab also contains this button:

SAVE DEBUG: used to generate text files that contain diagnostic information. After the debug data is collected, the system prompts for a location to save the compressed .tgz file.

6.4.1 Autotune

FreeNAS® provides an autotune script which optimizes the system depending on the installed hardware. For example, if a pool exists on a system with limited RAM, the autotune script automatically adjusts some ZFS sysctl values in an attempt to minimize 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* option in *System* → *Advanced* is off by default. Enable this option to run the autotuner at boot. To run the script immediately, reboot the system.

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`.

6.4.2 Self-Encrypting Drives

FreeNAS® version 11.1-U5 introduced Self-Encrypting Drive (SED) support.

Three types of SED devices are supported:

- Legacy interface for older ATA devices. **Not recommended for security-critical environments**
- TCG OPAL 2 standard for newer consumer-grade devices (HDD or SSD over PCIe or SATA)
- TCG Enterprise standard for newer enterprise-grade SAS devices

The FreeNAS® middleware implements the security capabilities of [camcontrol](https://www.freebsd.org/cgi/man.cgi?query=camcontrol) (https://www.freebsd.org/cgi/man.cgi?query=camcontrol) (for legacy devices) and [sedutil-cli](https://www.mankier.com/8/sedutil-cli) (https://www.mankier.com/8/sedutil-cli) (for TCG devices). When managing SED devices from the command line, it is important to use `sedutil-cli` rather than `camcontrol` to access the full capabilities of the device. FreeNAS® provides the `sedhelper` wrapper script to ease SED device administration from the command line.

By default, SED devices are not locked until the administrator explicitly configures a global or per-device password and initializes the devices.

Once configured, the system automatically unlocks all SEDs during the boot process, without requiring manual intervention. This allows a pool to contain a mix of SED and non-SED devices.

A password-protected SED device protects the data stored on the device when the device is physically removed from the FreeNAS® system. This allows secure disposal of the device without having to first wipe its contents. If the device is instead removed to be repurposed on another system, it can only be unlocked if the password is known.

Warning: It is important to remember the password! Without it, the device is unlockable and its data remains unavailable. While it is possible to specify the PSID number on the label of the device with the `sedutil-cli` command, doing so will erase the contents of the device rather than unlock it. Always record SED passwords whenever they are configured or modified and store them in a safe place!

When SED devices are detected during system boot, the middleware checks for global and device-specific passwords. Devices with their own password are unlocked with their password and any remaining devices, without a device-specific password, are unlocked using the global password.


To configure a global password, go to *System* → *Advanced* → *SED Password* and enter the password. Recording the password and storing it in a safe place is recommended.

To determine which devices support SED and their device names:

```
sedutil-cli --scan
```

In the results:

- **no** indicates a non-SED device
- **1** indicates a legacy TCG OPAL 1 device
- **2** indicates a modern TCG OPAL 2 device
- **E** indicates a TCG Enterprise device

To specify a password for a device, go to *Storage* → *Disks*. Click  (Options) on the confirmed SED device, then *Edit*. Enter and confirm the password in the *SED Password* and *Confirm SED Password* fields. Disks that have a configured password show bullets in their row of the *SED Password* column of *Storage* → *Disks*. Conversely, the rows in that column will be empty for disks that do not support SED or which are unlocked using the global password.

Next, remember to initialize the devices:

```
sedhelper setup password
```

This command ensures that all detected SED disks are properly setup using the specified password.

Note: Rerun `sedhelper setup password` every time a new SED disk is placed in the system.

This command is used to unlock all available SED disks:

```
sedhelper unlock
```

6.5 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 330) events are also emailed to the *root* user account. Problems with *Scrub Tasks* (page 132) are reported separately in an email sent at 03:00AM.

Note: *S.M.A.R.T.* (page 256) 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. Go to *Accounts* → *Users*, click ⓘ (Options) and *Edit* for the *root* user. In the *Email* field, enter the email address on the remote system where email is to be sent, like *admin@example.com*. Click *SAVE* to save the settings.

Additional configuration is performed with *System* → *Email*, shown in Figure 6.8.

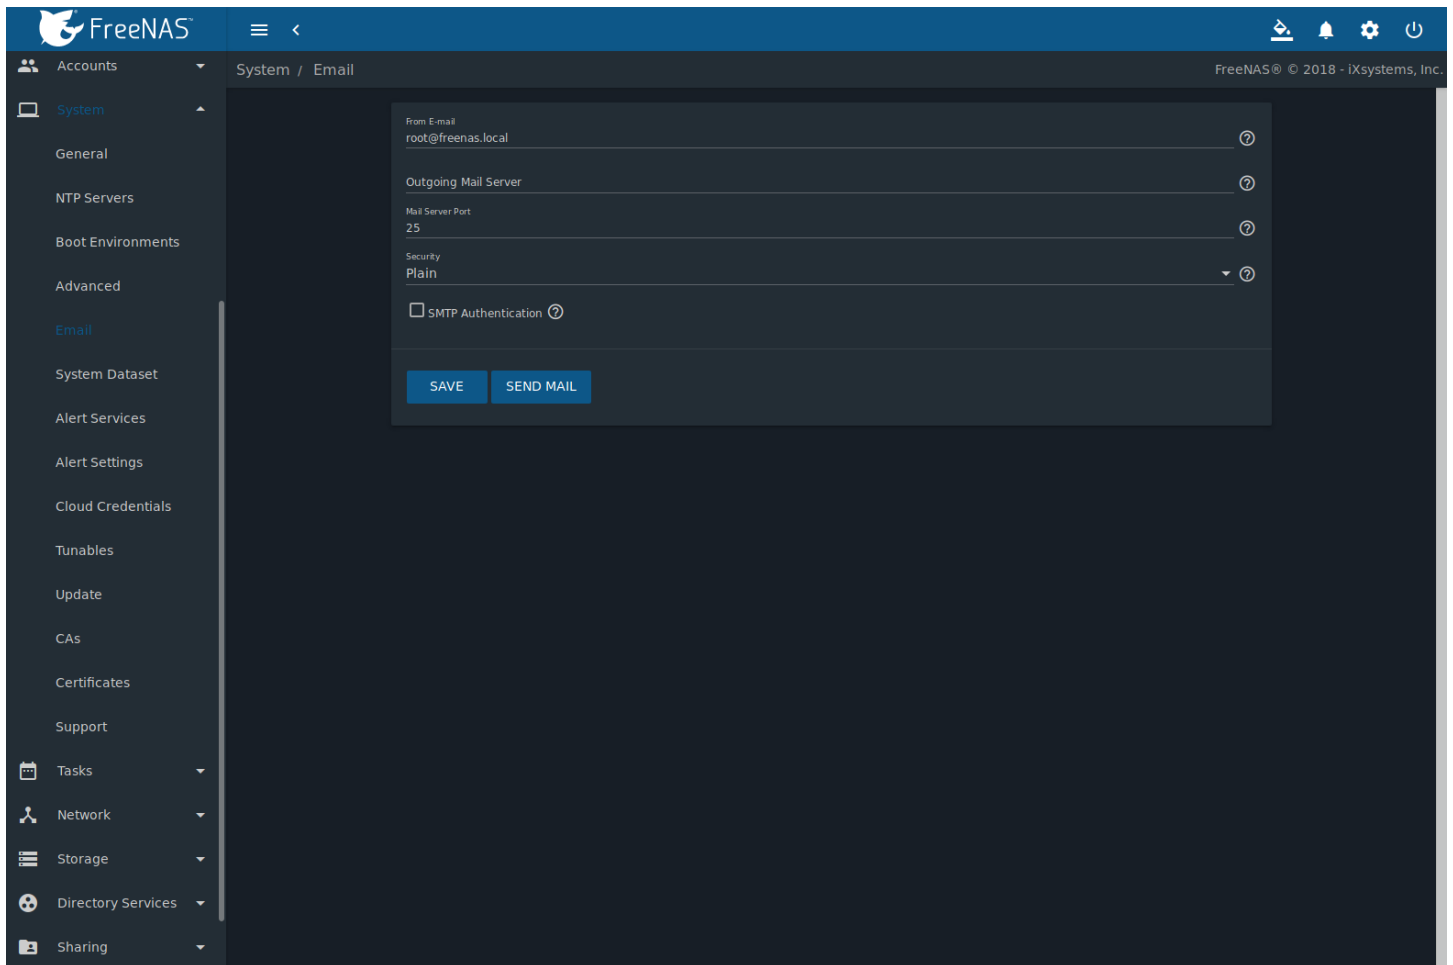


Fig. 6.8: Email Screen

Table 6.4: Email Configuration Settings

Setting	Value	Description
From E-mail	string	Setting a known <i>From</i> address can be helpful in filtering mail on the receiving system.
Outgoing Mail Server	string or IP address	Hostname or IP address of SMTP server used for sending this email.
Mail Server Port	integer	SMTP port number. Typically 25, 465 (secure SMTP), or 587 (submission).
Security	drop-down menu	Choose an encryption type. Choices are <i>Plain</i> , <i>SSL</i> , or <i>TLS</i> .
SMTP Authentication	checkbox	Enable or disable SMTP AUTH (https://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN SASL. If enabled, enter the required <i>Username</i> and <i>Password</i> .
Username	string	Enter the SMTP username if the SMTP server requires authentication.
Password	string	Enter the SMTP password if the SMTP server requires authentication.
Confirm Password	string	Confirm the SMTP previous password.

Click the *SEND MAIL* button to verify that the configured email settings are working. If the test email fails, double-check that the *Email* field of the *root* user is correctly configured by clicking the *Edit* button for the *root* account in *Accounts* → *Users*.

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/) (<https://forums.freenas.org/index.php?threads/are-you-having-trouble-getting-freenas-to-email-you-in-gmail.22517/>).

Note: The FreeNAS® user who receives periodic email is set in the *Periodic Notification User* field in *System* → *Advanced*.

6.6 System Dataset

System → *System Dataset*, shown in [Figure 6.9](#), is used to select the pool which contains 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 temporary or the information in the old dataset might be useful for later recovery.

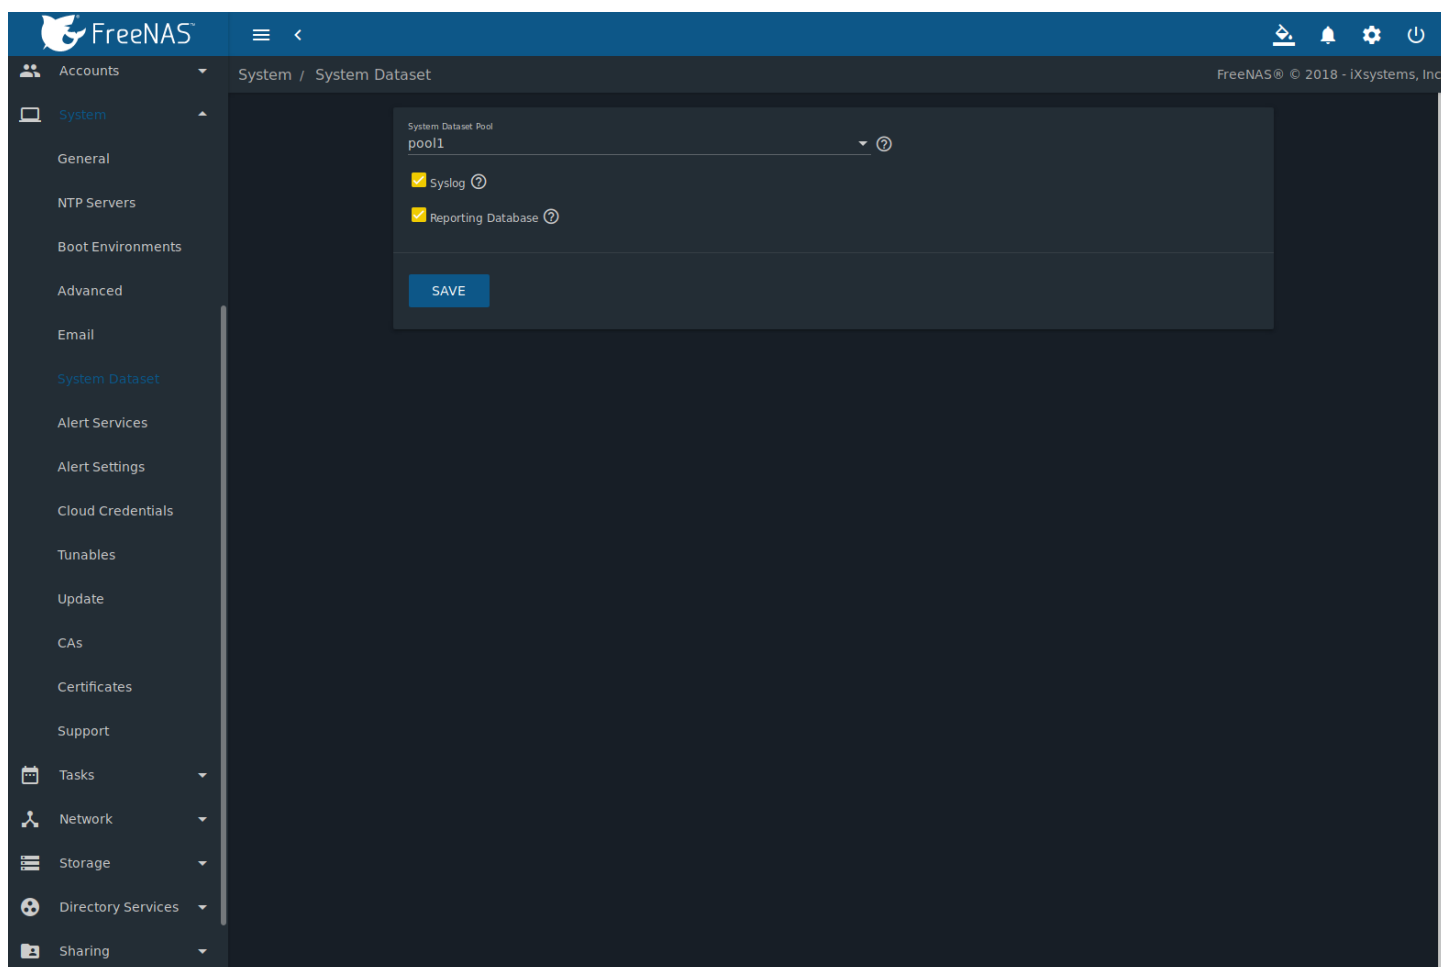


Fig. 6.9: System Dataset Screen

The system dataset can optionally be configured to also store the system log and [Reporting](#) (page 305) 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 pool to contain the system dataset.

To store the system log on the system dataset, enable the *Syslog* option.

To store the reporting information on the system dataset, enable the *Reporting Database* option. When this option is not enabled, 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.

6.7 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 330).

Available alert services:

- [AWS-SNS](https://aws.amazon.com/sns/) (<https://aws.amazon.com/sns/>)
- E-mail
- [Hipchat](https://www.stride.com) (<https://www.stride.com>)
- [InfluxDB](https://www.influxdata.com/) (<https://www.influxdata.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/>)
- [Slack](https://slack.com/) (<https://slack.com/>)
- [SNMP Trap](http://www.dpstele.com/snmp/trap-basics.php) (<http://www.dpstele.com/snmp/trap-basics.php>)
- [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 show the Alert Services screen, [Figure 6.10](#).

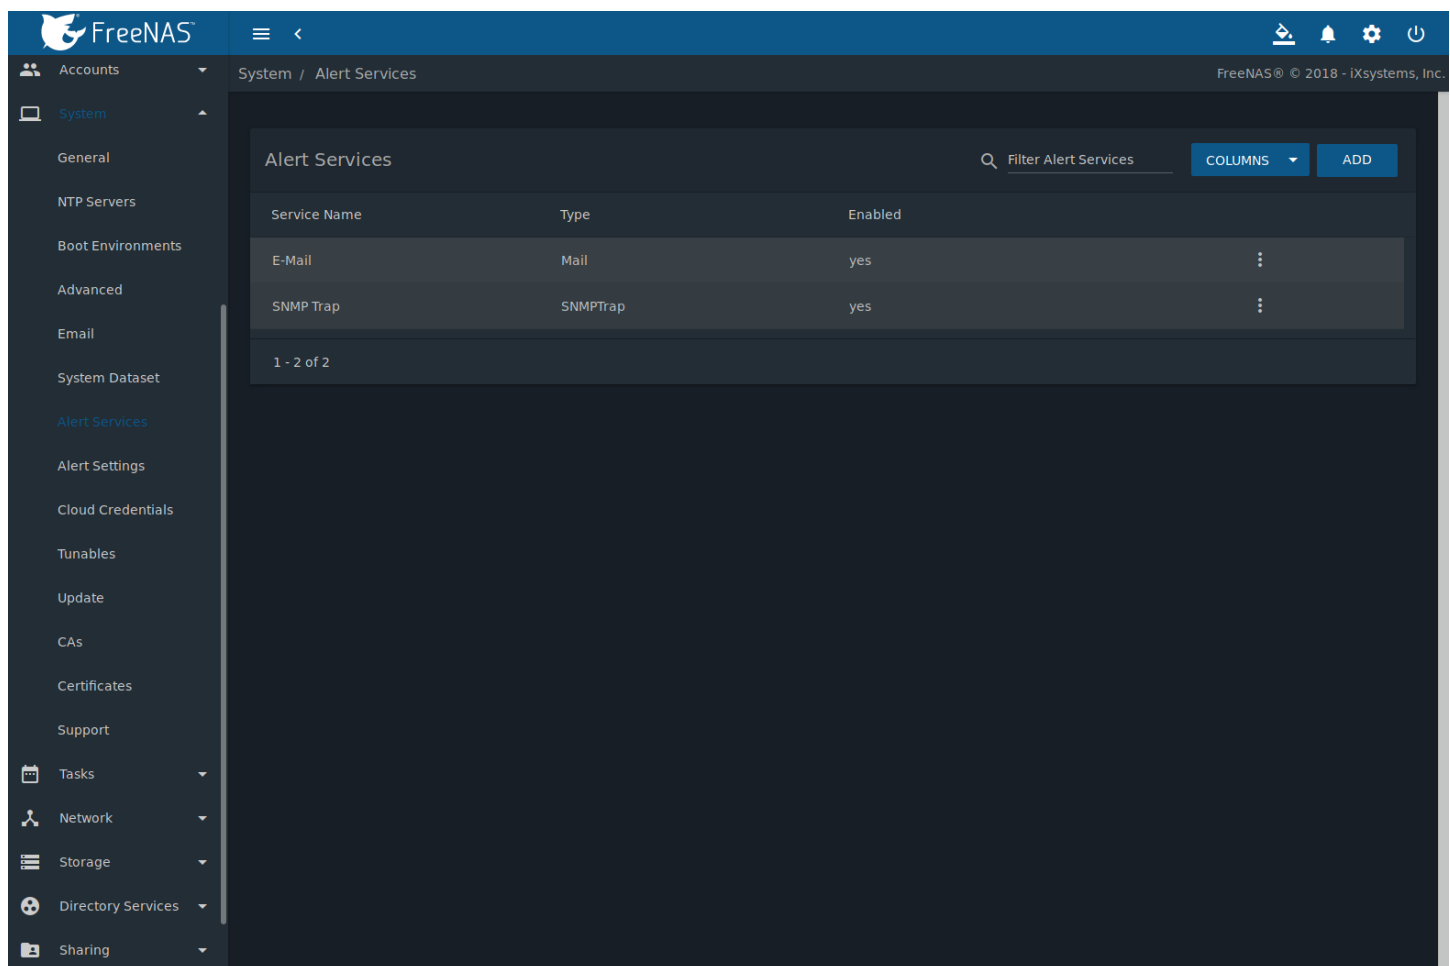


Fig. 6.10: Alert Services

Click *ADD* to display the *Add Alert Service* form, [Figure 6.11](#).

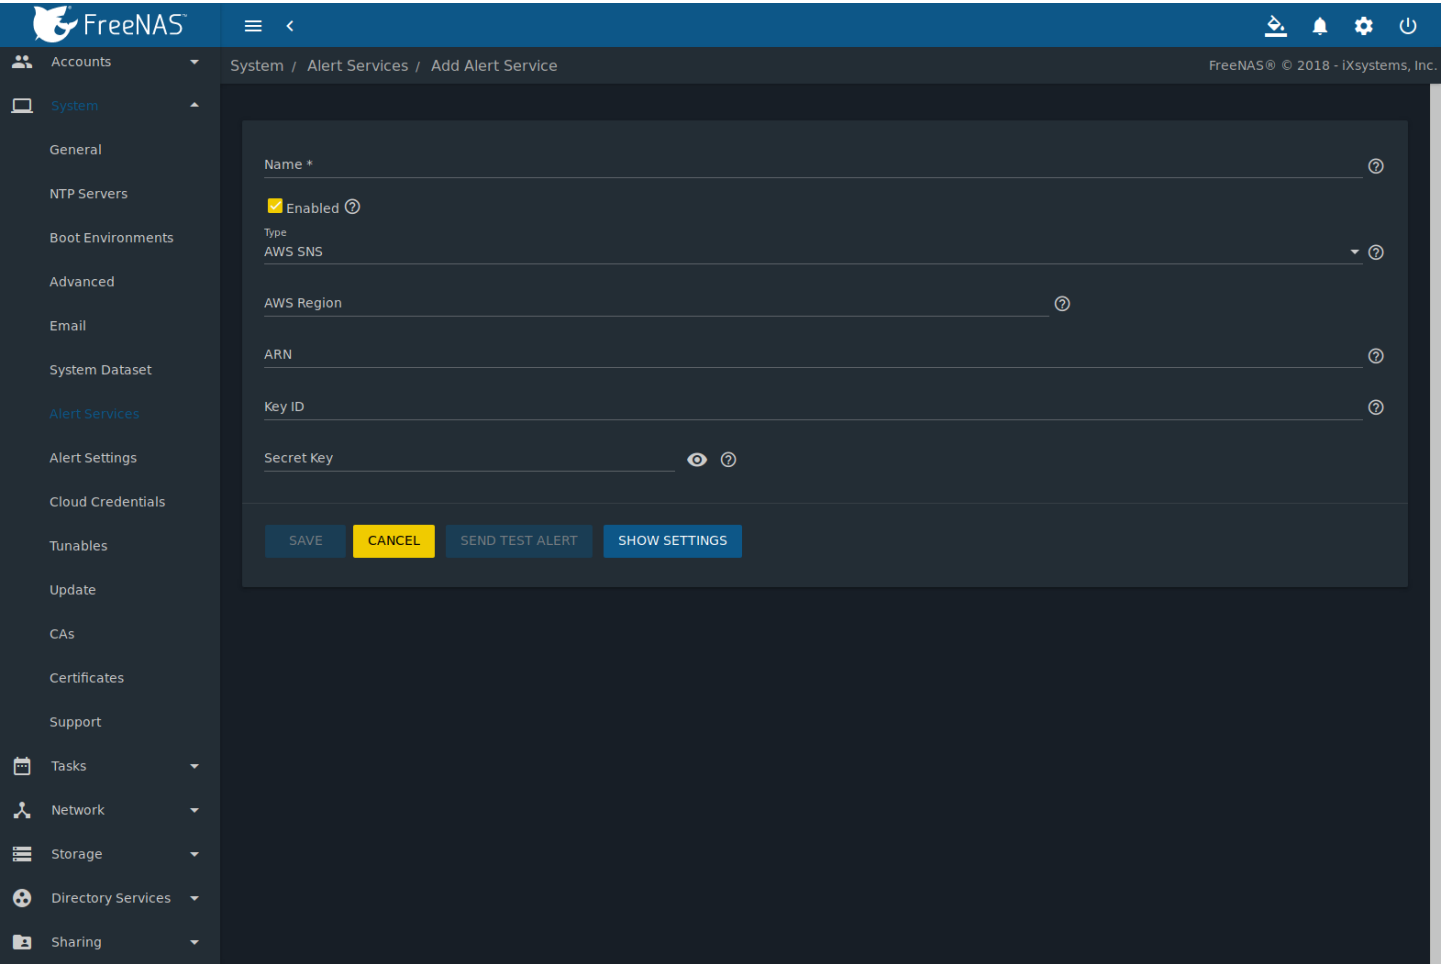


Fig. 6.11: Add Alert Service

Select the *Type* to choose an alert service to configure. The configurable fields and required information differ for each alert service. Set *Enabled* to activate the service. Enter any other required information and click *SAVE*.

Configure which alerts are sent to the alert service by clicking *SHOW SETTINGS*.

Click *SENDS TEST ALERT* to test the configured service.

All saved alert services are displayed in *System* → *Alert Services*. To delete an alert service, click ⓘ (Options) and *Delete*. To disable an alert service temporarily, click ⓘ (Options) and *Edit*, then unset the *Enabled* option.

6.8 Alert Settings

System → *Alert Settings* displays the notification frequency for each type of *Alert* (page 330). An example is shown in [Figure 6.12](#).

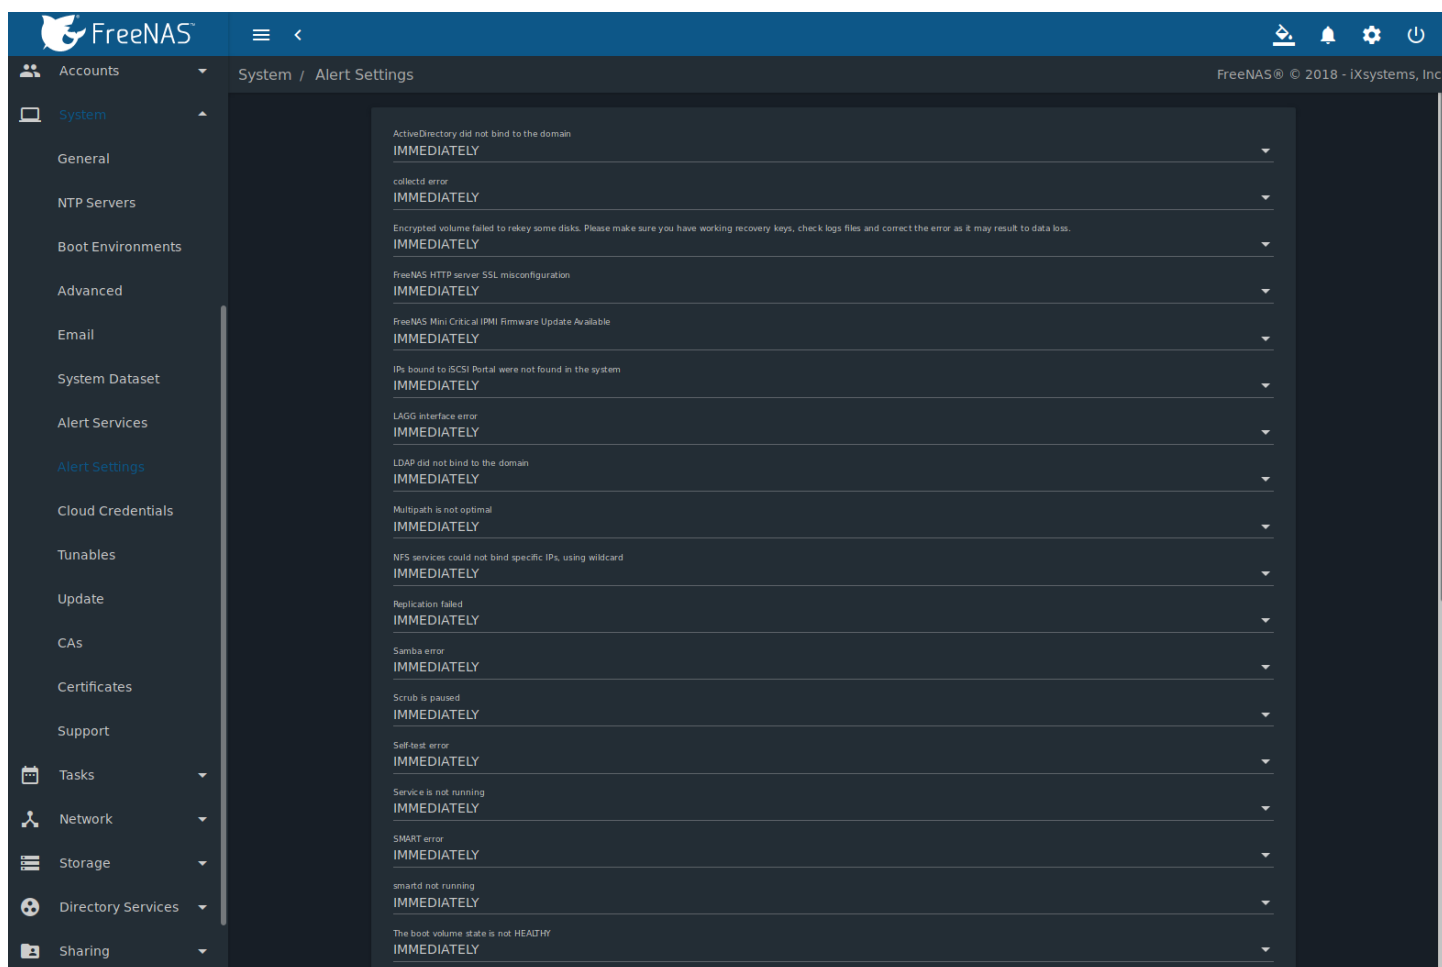


Fig. 6.12: Configure Alert Notification Frequency

To change the notification frequency of an alert, click its drop-down menu and select *IMMEDIATELY*, *HOURLY*, *DAILY*, or *NEVER*.

Note: To configure where alerts are sent, use [Alert Services](#) (page 85).

6.9 Cloud Credentials

FreeNAS® can use cloud services for features like [Cloud Sync Tasks](#) (page 134). The credentials to provide secure connections with cloud services are entered here. Amazon Cloud Drive, Amazon S3, Backblaze B2, Box, Dropbox, FTP, Google Cloud Storage, Google Drive, HTTP, Hubic, Mega, Microsoft Azure Blob Storage, Microsoft OneDrive, pCloud, SFTP, WebDAV, and Yandex are supported.

Click *System* → *Cloud Credentials* to see the screen shown in [Figure 6.13](#).

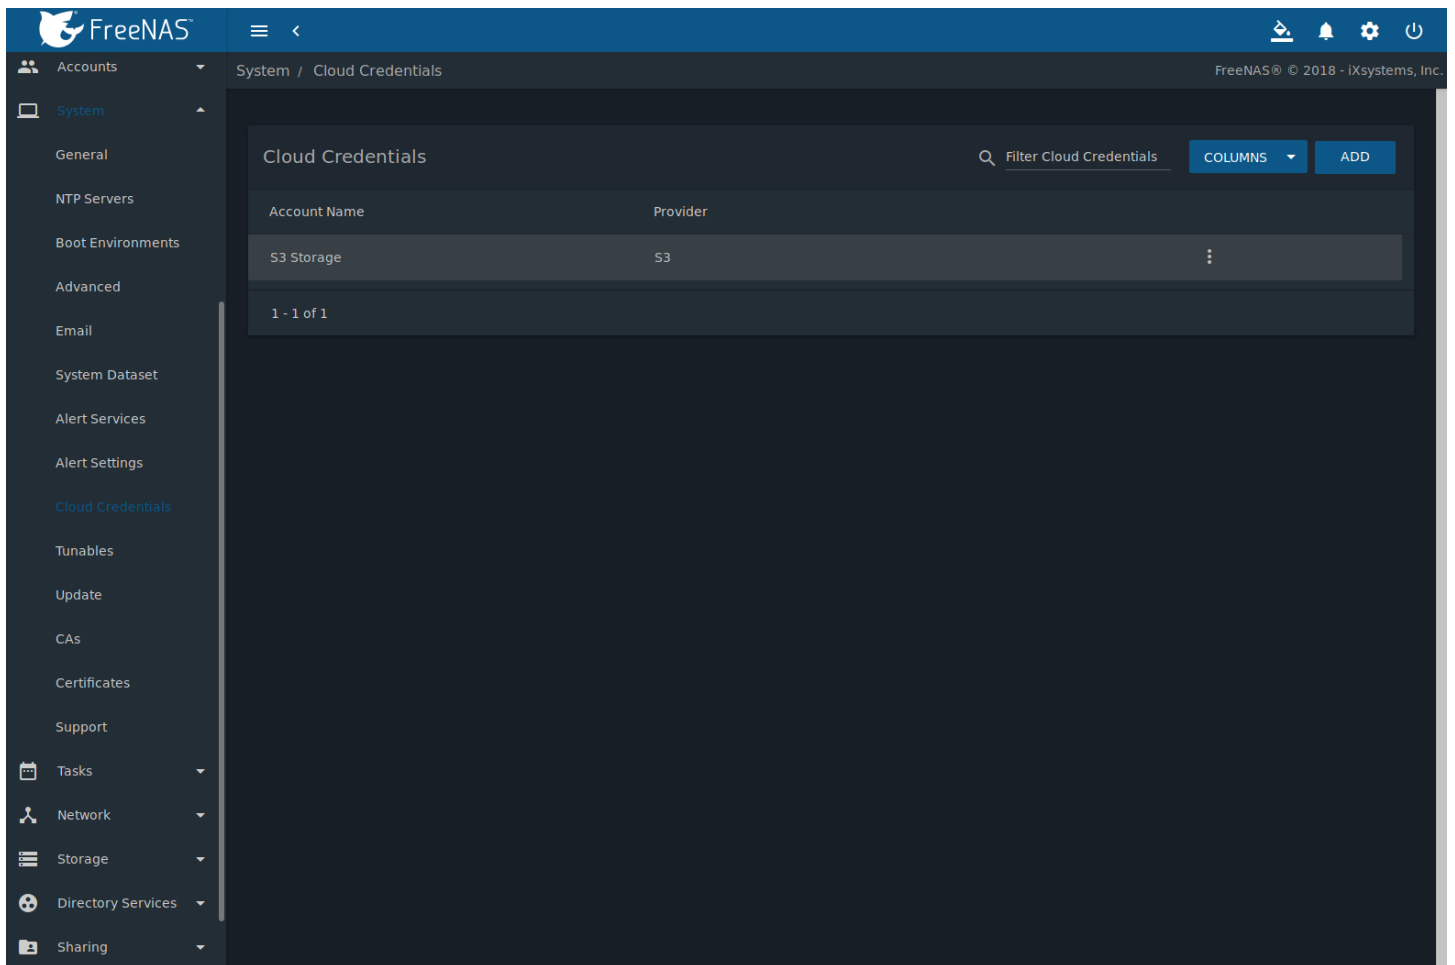



Fig. 6.13: Cloud Credentials List

The list shows the *Account Name* and *Provider* for each credential. There are options to *Edit* and *Delete* a credential after clicking  (Options) for a credential.

Click **ADD** to add a new cloud credential. Choose a *Provider* to display any specific options for that provider. [Figure 6.14](#) shows the form for an *Amazon Cloud Drive* provider:

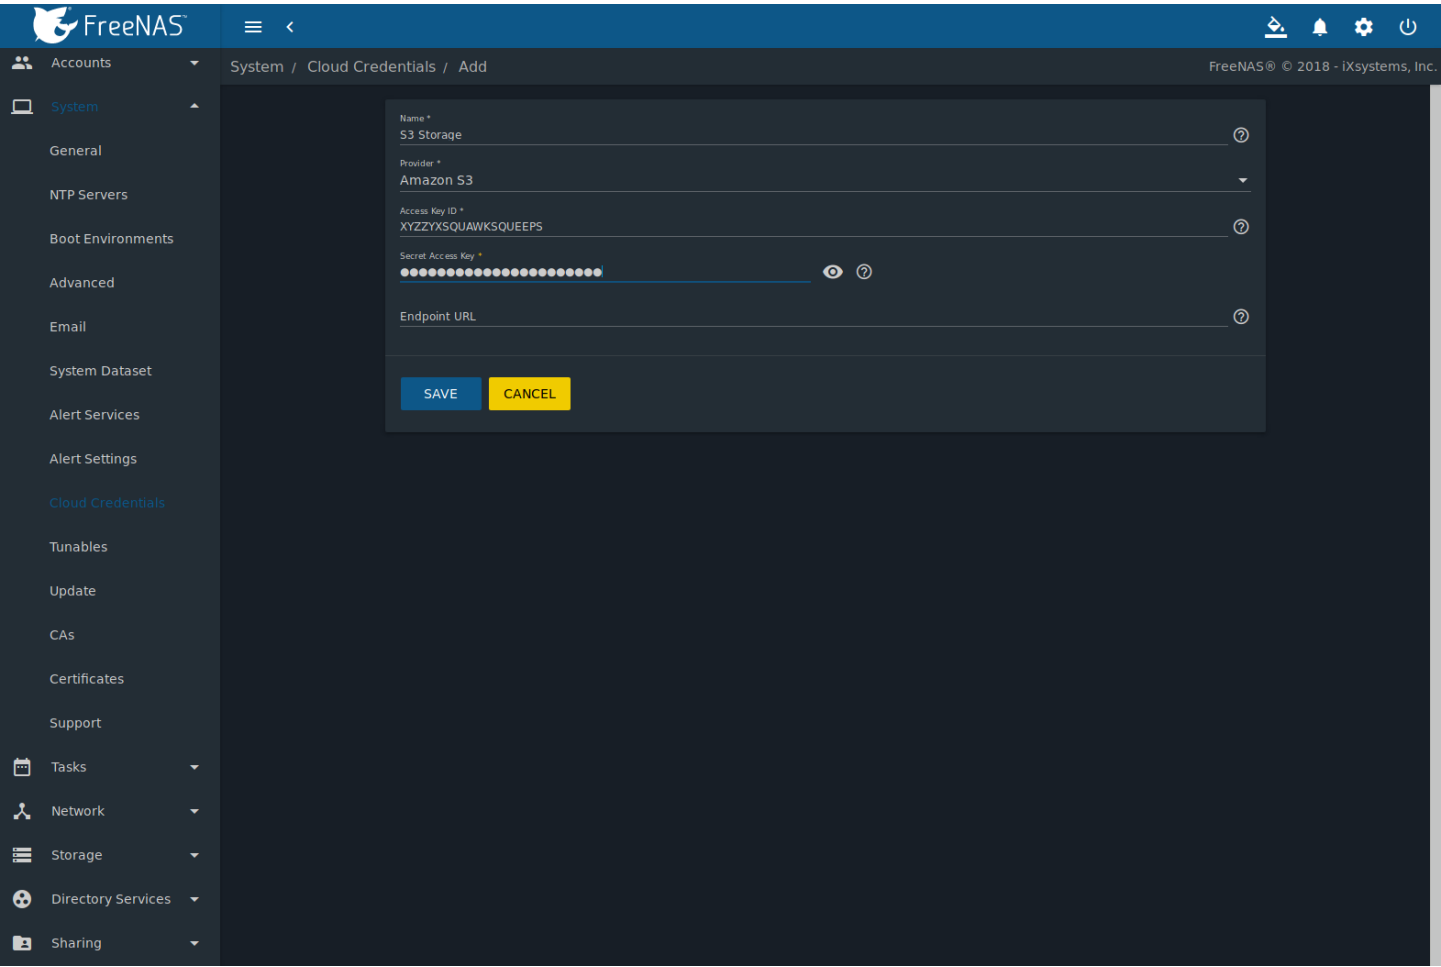


Fig. 6.14: Add Amazon Cloud Drive Credential

Enter a descriptive and unique name for the cloud credential in the *Name* field. The remaining options vary by *Provider*, and are shown in [Table 6.5](#).

Table 6.5: Cloud Credential Options

Provider	Setting	Description
Amazon Cloud Drive	Application Client ID, Application Key	Enter the Amazon application client ID and application key.
Amazon S3	Access Key ID, Secret Key	Enter the Amazon account access key and secret key.
Amazon S3	Endpoint URL	End point URL (https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteEndpoint) for a bucket configured for website hosting. Leave blank when using AWS. Enter endpoint URL if using custom S3 API. URL general for-format: <i>bucket-name.s3-website-region.amazonaws.com</i> . Refer to the AWS Documentation for a list of Simple Storage Service Websites Endpoints (https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_website_region_endp)

Continued on next page

Table 6.5 – continued from previous page

Provider	Setting	Description
Backblaze B2	Account ID or Application Key ID, Master Application Key or Application Key	Enter the Account ID and Master Application Key (https://help.backblaze.com/hc/en-us/articles/224991568-Where-can-I-find-my-Account-ID-and-Application-Key-) for the Backblaze B2 account. These are visible after logging into the account, clicking <i>Buckets</i> , and clicking <i>Show Account ID and Application Key</i> . An <i>Application Key</i> with limited permissions can be used in place of the <i>Account ID</i> and <i>Master Application Key</i> . Create a new <i>Application Key</i> and enter the key string in place of the <i>Master Application Key</i> and replace the <i>Account ID</i> with the <i>keyID</i> .
Box	Access Token	Enter the Box access token.
Dropbox	Access Token	Enter the Dropbox access token. The token is located on the App Console (https://www.dropbox.com/developers/apps). After creating an app, go to <i>Settings</i> and click the <i>Generate</i> button under the Generated access token field.
FTP	Host, Port	Enter the FTP host and port.
FTP	Username, Password	Enter the FTP username and password.
Google Cloud Storage	JSON Service Account Key	<i>Browse</i> to the location of the saved Google Cloud Storage key and select it.
Google Drive	Access Token, Team Drive ID	Enter the Google Drive Access Token. <i>Team Drive ID</i> is only used when connecting to a Team Drive (https://developers.google.com/drive/api/v3/reference/teamdrives). The ID is also the ID of the top level folder of the Team Drive.
HTTP	URL	Enter the URL.
Hubic	Access Token	Enter the access token.
Mega	Username, Password	Enter the Mega (https://mega.nz/) username and password.
Microsoft Azure Blob Storage	Account Name, Account Key	Enter the Azure Blob Storage account name and key.
Microsoft OneDrive	Access Token, Drive Account Type, Drive ID	Enter the access token. Choose the account type: <i>PERSONAL</i> , <i>BUSINESS</i> , or SharePoint (https://products.office.com/en-us/sharepoint/collaboration) <i>DOCUMENT_LIBRARY</i> . Enter the unique drive identifier. Open the Shell (page 325), enter <code>rcclone config</code> , and follow the prompts to find these values. The rcclone OneDrive documentation (https://rcclone.org/onedrive/) guides through the configuration process.
pCloud	Access Token	Enter the access token.
SFTP	Host, Port	Enter the SFTP host and port.
SFTP	Username, Password, key file path	Enter the SFTP username, password, and PEM-encoded private key file path.
WebDAV	URL, WebDAV service	Enter the URL and use the dropdown to select the WebDAV service.
WebDAV	Username, Password	Enter the username and password.
Yandex	Access Token	Enter the access token.

For Amazon S3, *Access Key* and *Secret Key* values are 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 unknown, a new key pair can be created on the same Amazon screen. The Google Cloud Storage *JSON Service Account Key* is found on the [Google Cloud Platform Console](#) (<https://console.cloud.google.com/apis/credentials>).

More details about individual *Provider* settings are available in the [rcclone documentation](#) (<https://rcclone.org/about/>).

6.10 Tunables

System → *Tunables* can be used to manage:

1. **FreeBSD sysctls:** a `sysctl(8)` (<https://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>) 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 before testing 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](https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/) (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/).

To add a loader, `sysctl`, or `rc.conf` option, go to *System* → *Tunables* and click *ADD* to access the screen shown in [Figure 6.15](#).

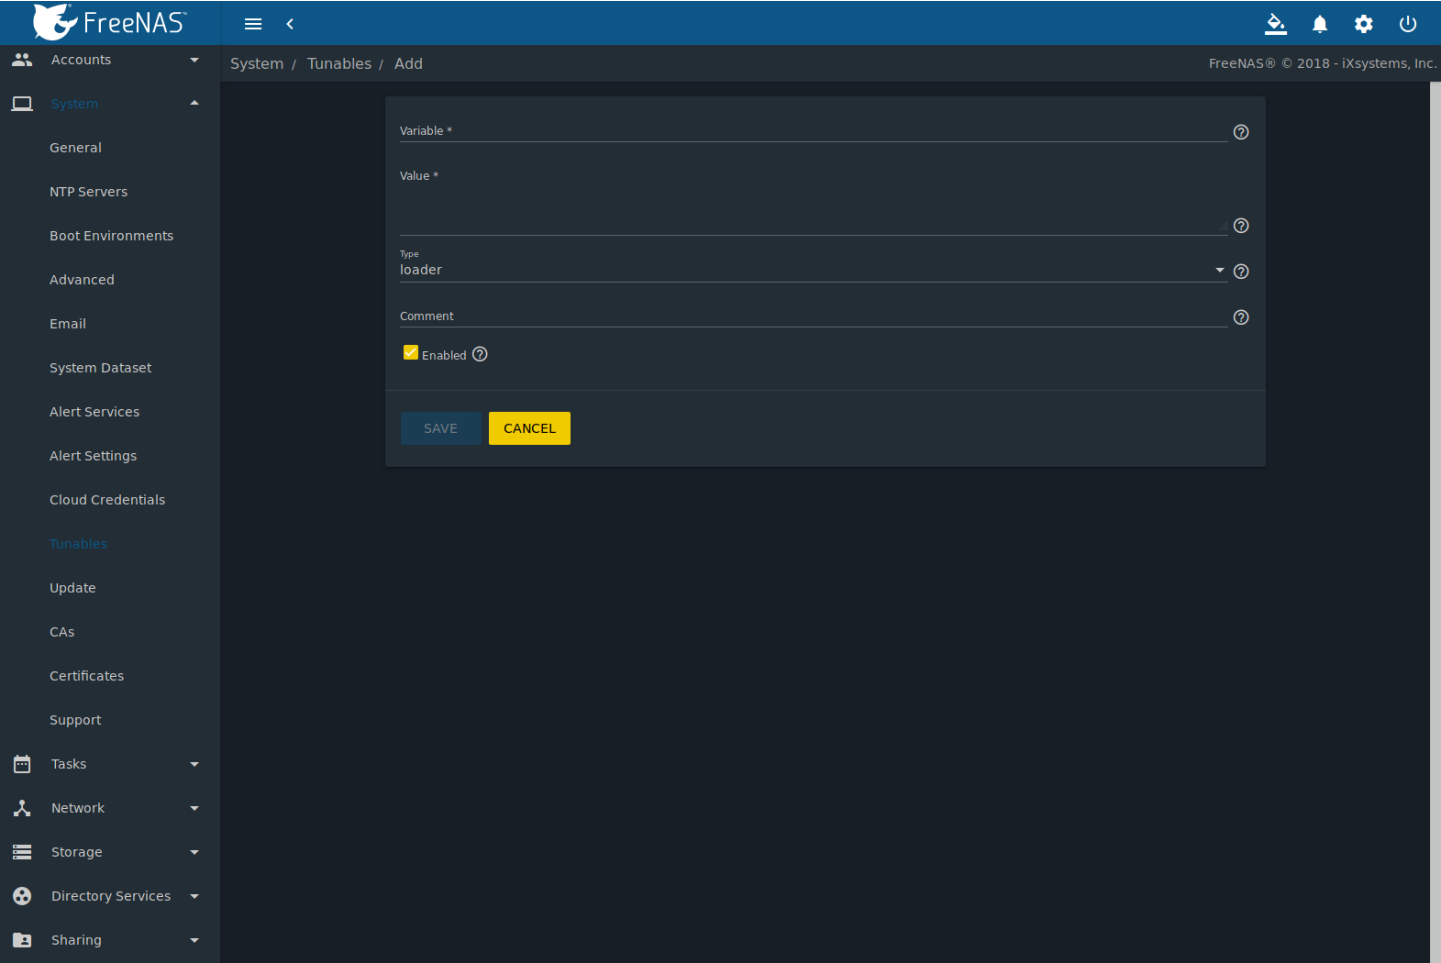




Fig. 6.15: Adding a Tunable

Table 6.6 summarizes the options when adding a tunable.

Table 6.6: Adding a Tunable

Setting	Value	Description
Variable	string	The name of the sysctl or driver to load.
Value	integer or string	Set a value for the <i>Variable</i> . Refer to the man page for the specific driver or the FreeBSD Handbook (https://www.freebsd.org/doc/en_US.ISO08859-1/books/handbook/) for suggested values.
Type	drop-down menu	Choices are <i>Loader</i> , <i>rc.conf</i> , and <i>Sysctl</i> .
Comment	string	Optional. Enter a description of this tunable.
Enabled	checkbox	Deselect this option 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 the *Enabled* option is deselected.

Existing tunables are listed in *System* → *Tunables*. To change the value of an existing tunable, click  (Options) and *Edit*. To remove a tunable, click  (Options) and *Delete*.

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 web interface does not display the sysctls that are pre-set when FreeNAS® is installed. FreeNAS® 11.2 ships with the 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 web interface does not display the loaders that are pre-set when FreeNAS® is installed. FreeNAS® 11.2 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. Changing the default tunables can make the system unusable.

The ZFS version used in 11.2 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.

6.11 Update

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

6.11.1 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. Most updates require a system reboot. Plan updates around 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 Environments](#) (page 75) to remove unneeded boot environments.

6.11.2 Updates and Trains

FreeNAS® uses 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 maintaining the ability to revert to a previous version of the operating system, using the instructions in [If Something Goes Wrong](#) (page 33). Digitally signed update files eliminate the need to manually download both an upgrade file and the associated checksum to verify file integrity.

FreeNAS® defines software branches, known as *trains*. There are several trains available for updates, but the web interface only displays trains that can be selected as an upgrade.

Update trains are labeled with a numeric version followed by a short description. The current version receives regular bug fixes and new features. Supported older versions of FreeNAS® only receive maintenance updates. Several specific words are used to describe the type of train:

- **STABLE:** Bug fixes and new features are available from this train. Upgrades available from a *STABLE* train are tested and ready to apply to a production environment.

- **Nightlies:** Experimental train used for testing future versions of FreeNAS®.
- **SDK:** Software Developer Kit train. This has additional tools for testing and debugging FreeNAS®.

Warning: The UI will warn if the currently selected train is not suited for production use. 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>.

6.11.3 Checking for Updates

Figure 6.16 shows an example of the *System* → *Update* screen.

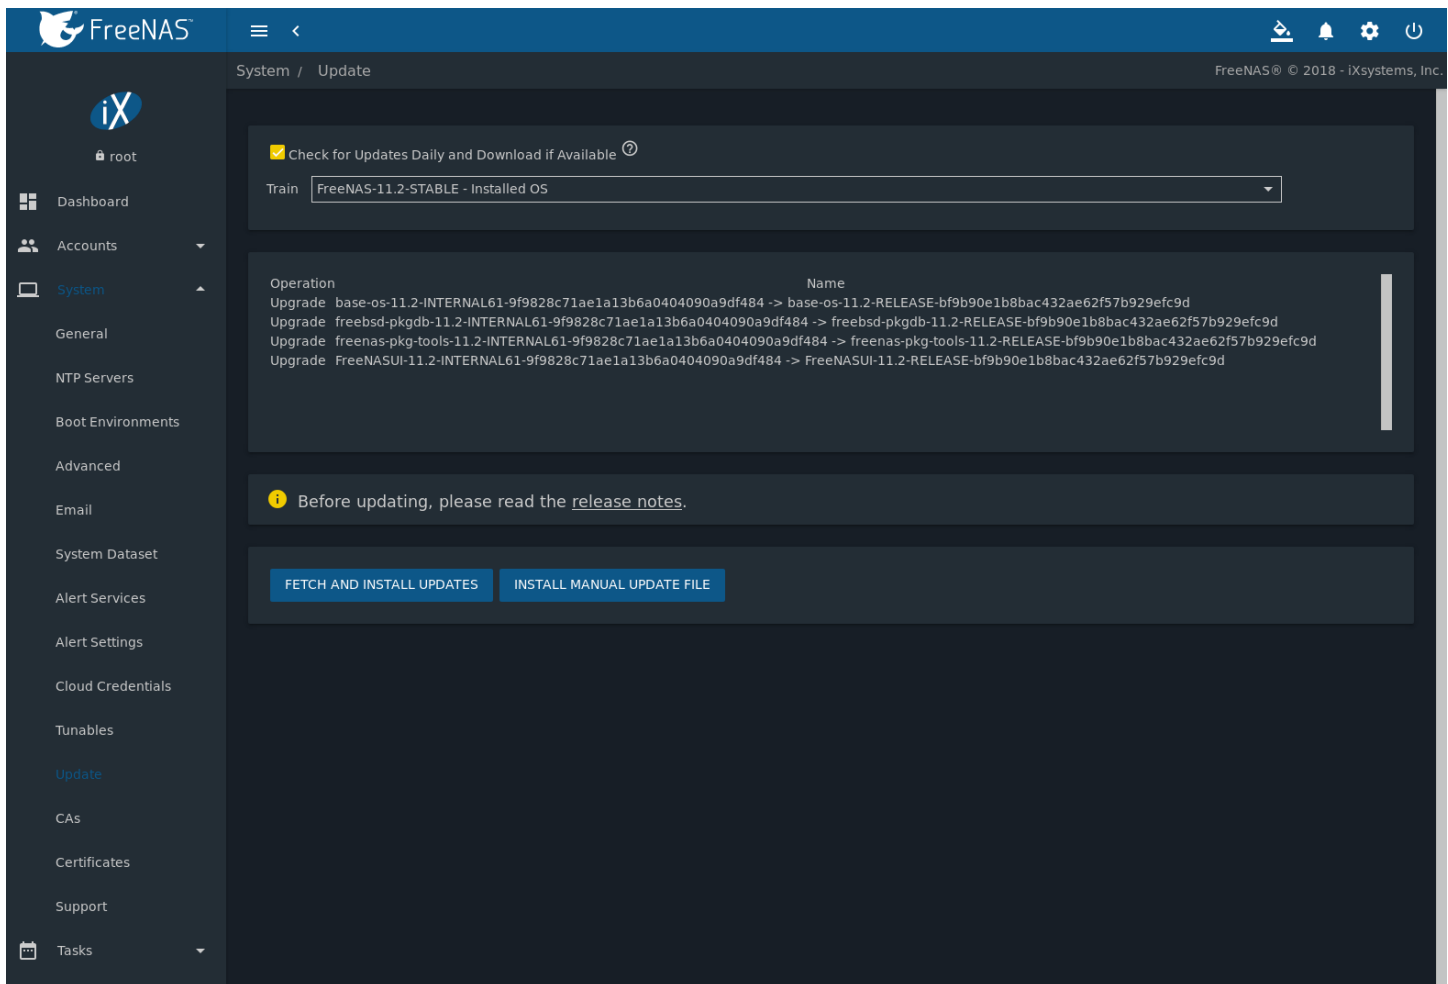


Fig. 6.16: Update Options

The system checks daily for updates and downloads an update if one is available. An alert is issued when a new update becomes available. The automatic check and download of updates is disabled by unsetting *Check for Updates Daily and Download if Available*.

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.

In the example shown in [Figure 6.17](#), information about the update is displayed along with a link to the *release notes*. It is important to read the release notes before updating to determine if any of the changes in that release impact the use of the system.

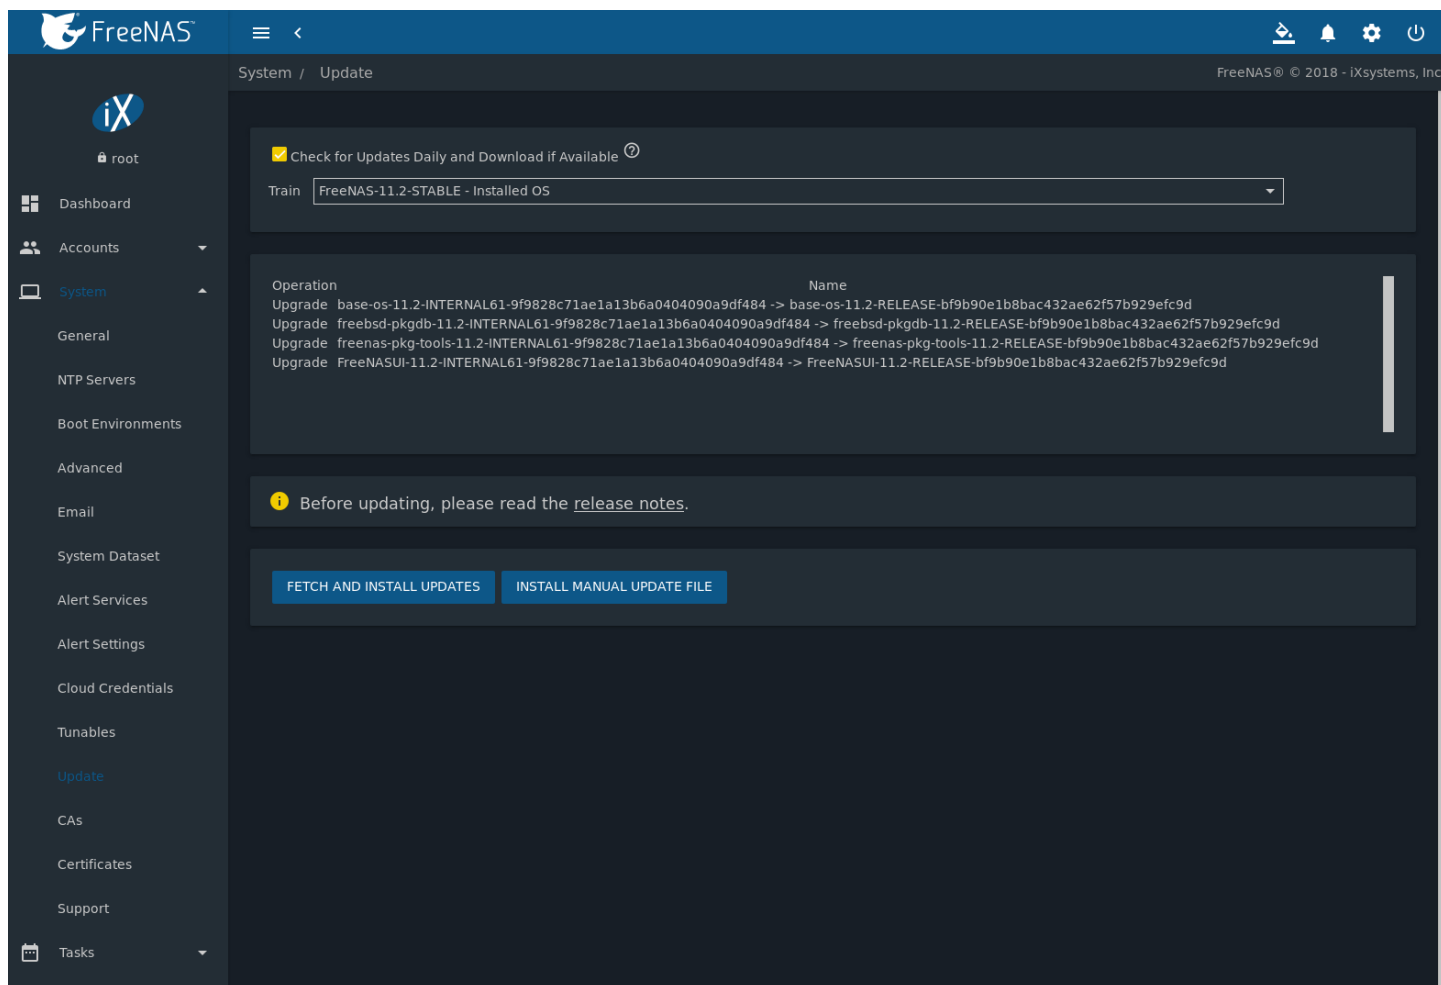


Fig. 6.17: Reviewing Updates

6.11.4 Saving the Configuration File

A dialog to save the system *configuration file* (page 73) appears before installing updates.

Note: The “Save Configuration” dialog can be disabled in ⓘ (Settings) *Preferences*, although this is *not* recommended. Saving backups of configuration files allows recovery of the system after a boot device failure.

Warning: Keep the system configuration file secure after saving it. The security information in the configuration file could be used for unauthorized access to the FreeNAS® system.

6.11.5 Applying Updates

Make sure the system is in a low-usage state as described above in *Preparing for Updates* (page 94).

Click *FETCH AND INSTALL UPDATES* to immediately download and install an update.

The “*Save Configuration*” (page 96) dialog appears so the current configuration can be saved to external media.

A confirmation window appears before the update is installed. When *Apply updates and reboot system after downloading* is set and, clicking *CONTINUE* downloads, applies the updates, and then automatically reboots the system. The update can be downloaded for a later manual installation by unsetting the *Apply updates and reboot system after downloading* option.

APPLY PENDING UPDATE is visible when an update is downloaded and ready to install. Click the button to see a confirmation window. Setting *Confirm* and clicking *CONTINUE* installs the update and reboots the system.

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 Environments* (page 75) 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 Environments*. Review the boot environments and remove the *Keep* attribute or delete any boot environments that are no longer needed.

During the update process a progress dialog appears. **Do not** interrupt the update until it completes.

6.11.6 Manual Updates

Updates can also be manually downloaded and applied using the *INSTALL MANUAL UPDATE FILE* button.

The “*Save Configuration*” (page 96) dialog appears so the current configuration can be saved to external media.

After clicking *INSTALL MANUAL UPDATE FILE*, choose a location to temporarily store the update file on the FreeNAS® system. Use *Browse* to navigate to the downloaded manual update file (`-manual-update-unsigned.tar`). Finally, click *APPLY UPDATE* to begin the update. A progress dialog is displayed during the update. **Do not** interrupt the update.

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

6.12 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 6.18 shows the screen after clicking *System → CAs*.

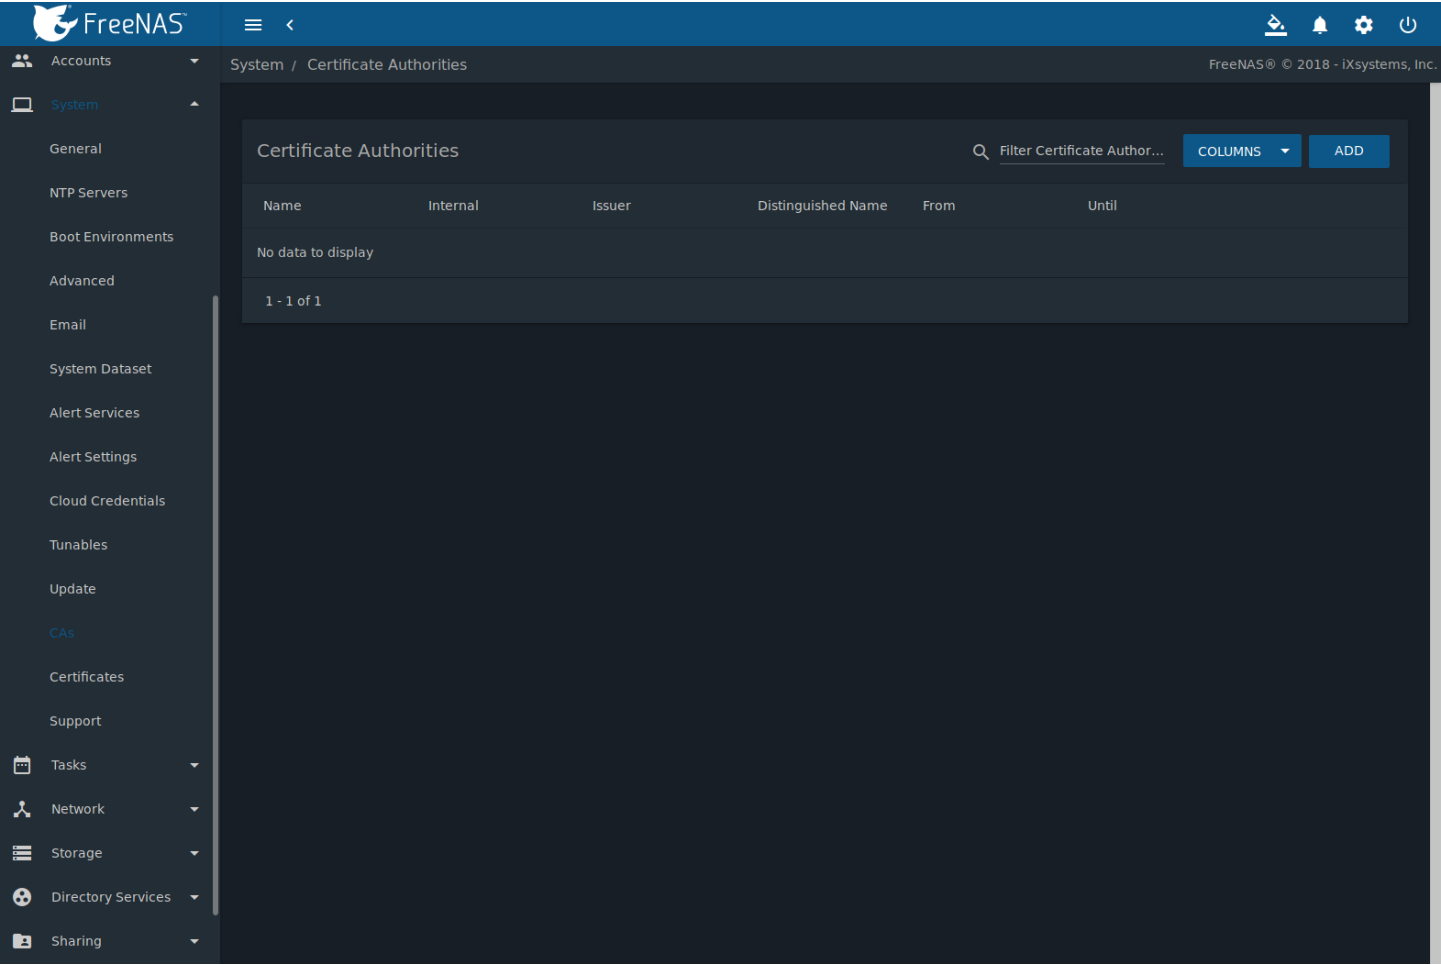


Fig. 6.18: Initial CA Screen

If the organization already has a CA, the CA certificate and key can be imported. Click *ADD* and set the *Type* to *Import CA* to see the configuration options shown in [Figure 6.19](#). The configurable options are summarized in [Table 6.7](#).

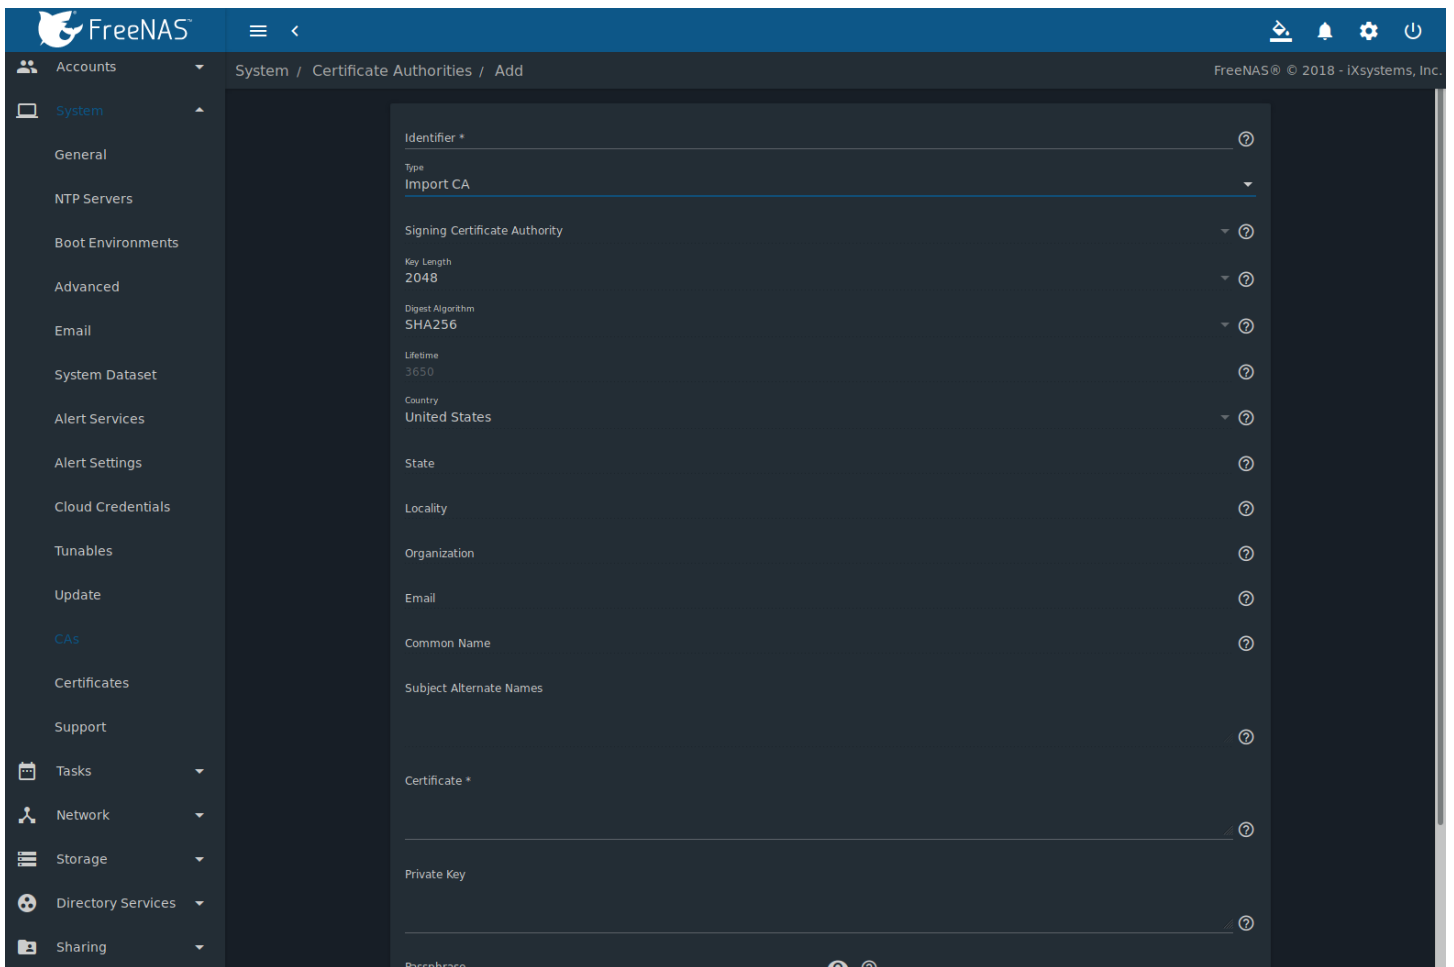


Fig. 6.19: Importing a CA

Table 6.7: Importing a CA Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore (<code>_</code>), and dash (<code>-</code>) characters.
Type	drop-down menu	Choose the type of CA. Choices are <i>Internal CA</i> , <i>Intermediate CA</i> , and <i>Import CA</i> .
Certificate	string	Mandatory. Paste in the certificate for the CA.
Private Key	string	If there is a private key associated with the <i>Certificate</i> , paste it here.
Passphrase	string	If the <i>Private Key</i> is protected by a passphrase, enter it here and repeat it in the “Confirm Passphrase” field.

To 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 *ADD* and set the *Type* to *Internal CA*. [Figure 6.20](#) shows the available options.

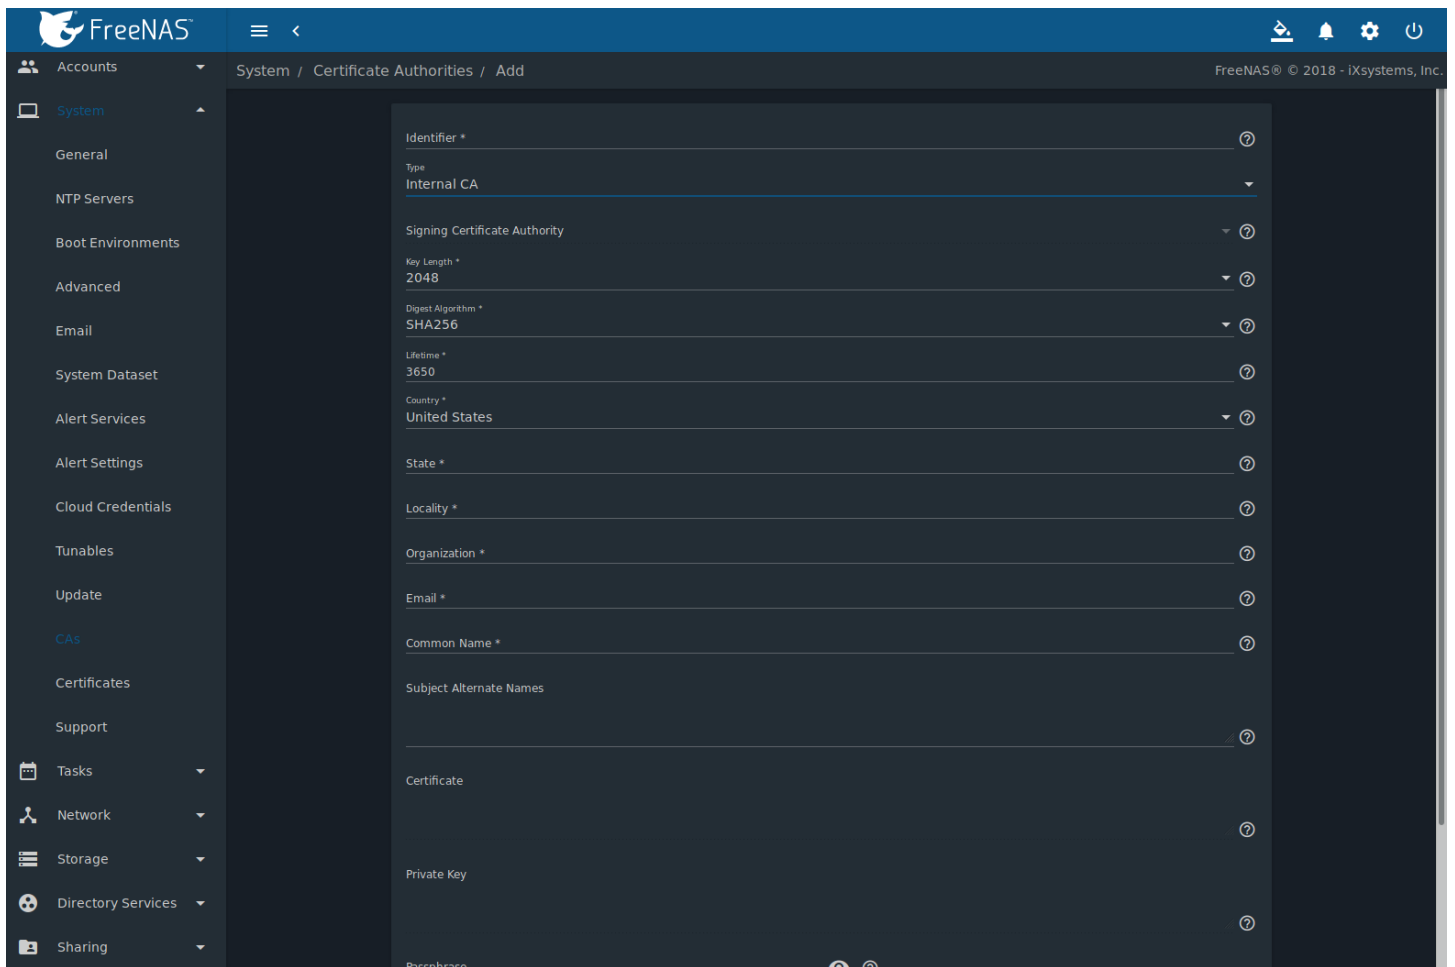


Fig. 6.20: Creating an Internal CA

The configurable options are described in [Table 6.8](#). When completing the fields for the certificate authority, supply the information for the organization.

Table 6.8: Internal CA Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore (<code>_</code>), and dash (<code>-</code>) characters.
Type	drop-down menu	Choose the type of CA. Choices are <i>Internal CA</i> , <i>Intermediate CA</i> , and <i>Import CA</i> .
Key Length	drop-down menu	For security reasons, a minimum of <i>2048</i> is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algorithm.
Lifetime	integer	The lifetime of a CA is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	Enter the state or province of the organization.
Locality	string	Enter the location of the organization.
Organization	string	Enter the name of the company or organization.
Email	string	Enter the email address for the person responsible for the CA.
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The <i>Common Name</i> must be unique within a certificate chain.

Continued on next page


Table 6.8 – continued from previous page

Setting	Value	Description
Subject Alternate Names	string	Multi-domain support. Enter additional space separated domain names.

To create an intermediate CA which is part of a certificate chain, set the *Type* to *Intermediate CA*. This screen adds one more option to the screen shown in [Figure 6.20](#):

- **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.

Imported or created CAs are 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 CA lifetime (in days), the common name of the CA, the date and time the CA was created, and the date and time the CA expires.

Click  (Options) on an existing CA to access these configuration buttons:

- **View:** use this option to view the contents of an existing *Certificate*, *Private Key*, or to edit the *Identifier*.
- **Sign CSR:** used to sign internal Certificate Signing Requests created using *System* → *Certificates* → *Create CSR*.
- **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.

6.13 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 97).

[Figure 6.21](#) shows the initial screen after clicking *System* → *Certificates*.

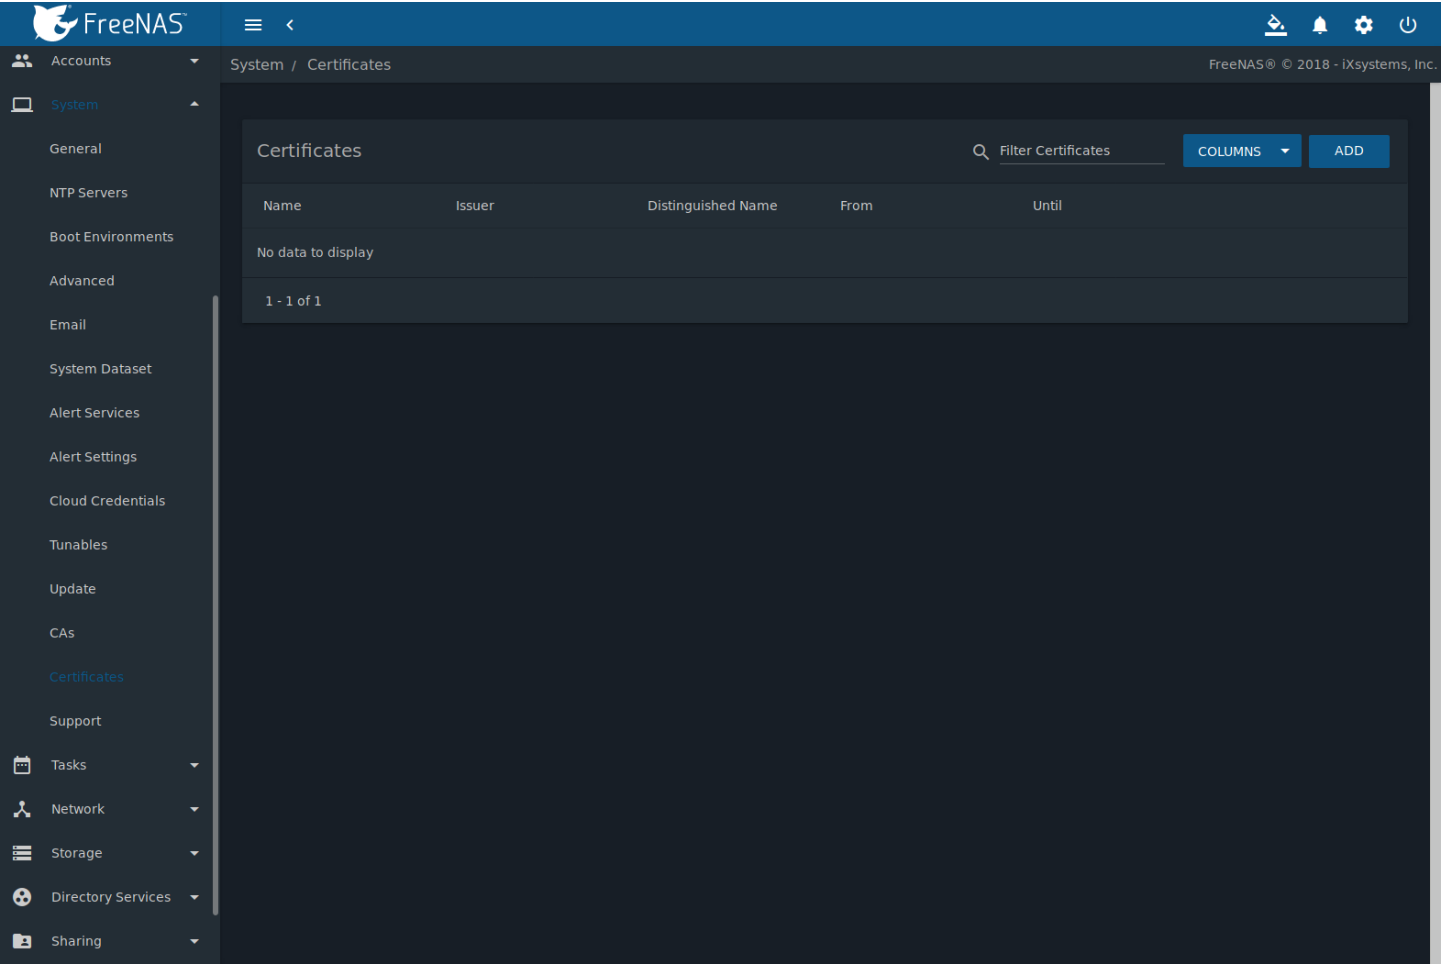


Fig. 6.21: Initial Certificates Screen

To import an existing certificate, click *ADD* and set the *Type* to *Import Certificate*. [Figure 6.22](#) shows the options. 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 6.9](#).

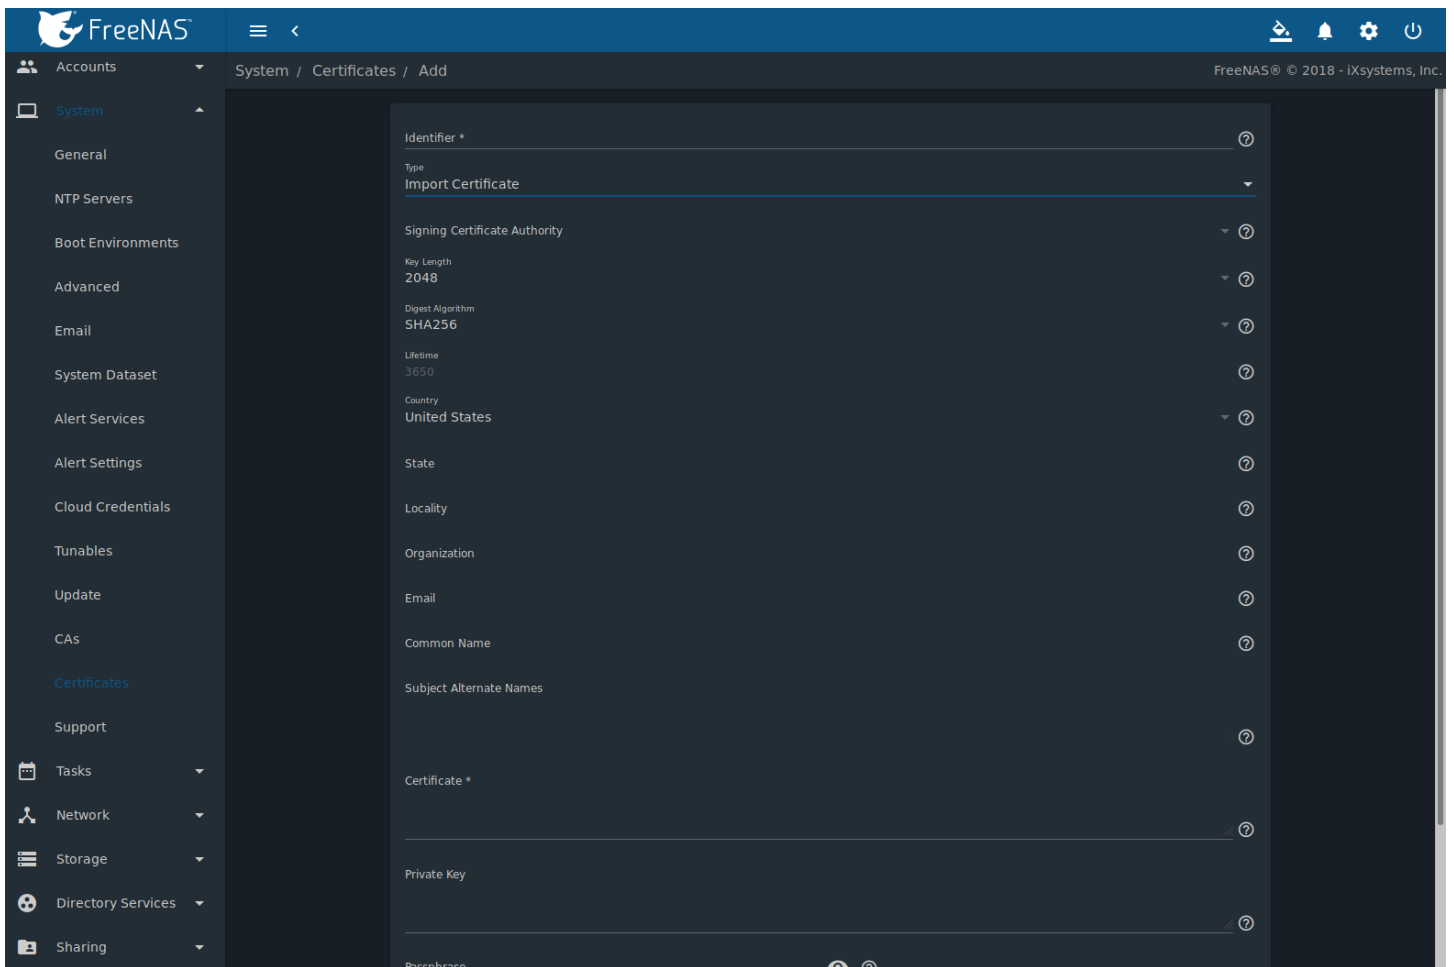


Fig. 6.22: Importing a Certificate

Table 6.9: Certificate Import Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric, underscore (_), and dash (-) characters.
Type	drop-down menu	Choose the type of certificate. Choices are <i>Internal Certificate</i> , <i>Certificate Signing Request</i> , and <i>Import Certificate</i> .
Certificate	string	Paste the contents of the certificate.
Private Key	string	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 <i>Confirm Passphrase</i> field.

To create a new self-signed certificate, set the *Type* to *Internal Certificate* to see the options shown in [Figure 6.23](#). The configurable options are summarized in [Table 6.10](#). When completing the fields for the certificate authority, use the information for the organization. Since this is a self-signed certificate, use the CA that was imported or created with [CAs](#) (page 97) as the signing authority.

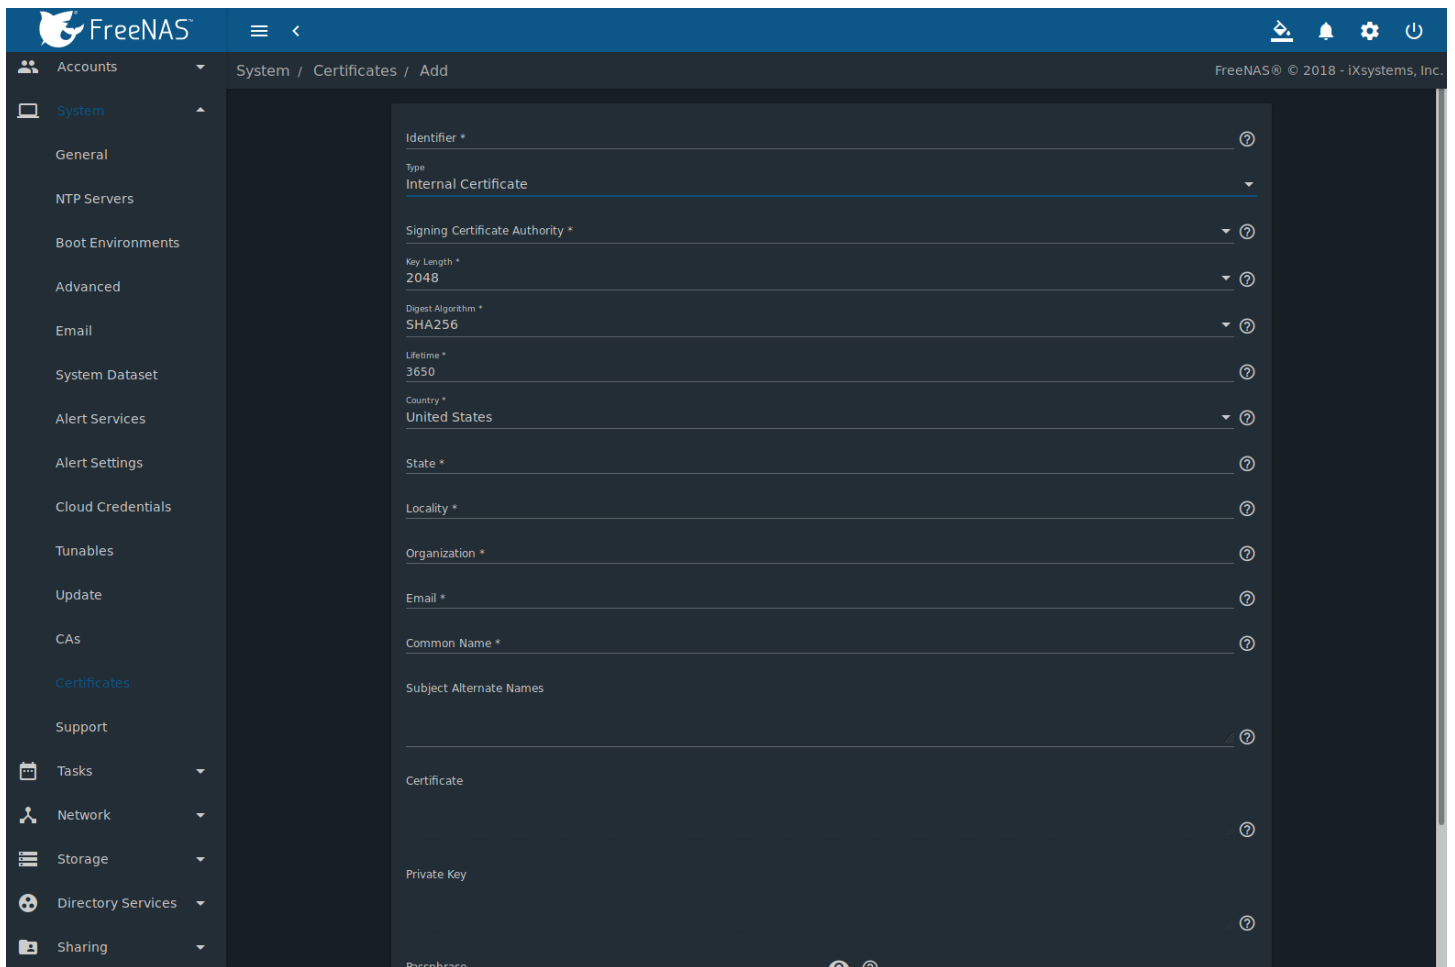


Fig. 6.23: Creating a New Certificate

Table 6.10: Certificate Creation Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric, underscore (_), and dash (-) characters.
Type	drop-down menu	Choose the type of certificate. Choices are <i>Internal Certificate</i> , <i>Certificate Signing Request</i> , and <i>Import Certificate</i> .
Signing Certificate Authority	drop-down menu	Select the CA which was previously imported or created using CAs (page 97).
Key Length	drop-down menu	For security reasons, a minimum of <i>2048</i> is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algorithm.
Lifetime	integer	The lifetime of the certificate is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	State or province of the organization.
Locality	string	Location of the organization.
Organization	string	Name of the company or organization.
Email	string	Enter the email address for the person responsible for the CA.
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The <i>Common Name</i> must be unique within a certificate chain.
Subject Alternate Names	string	Multi-domain support. Enter additional domain names and separate them with a space.

If the certificate is signed by an external CA, such as Verisign, instead create a certificate signing request. To do so, set the *Type* to *Certificate Signing Request*. The options from [Figure 6.23](#) display, 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 6.24](#), 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 with a new certificate set to *Import Certificate*. This makes the certificate available as a configurable option for encrypting connections.

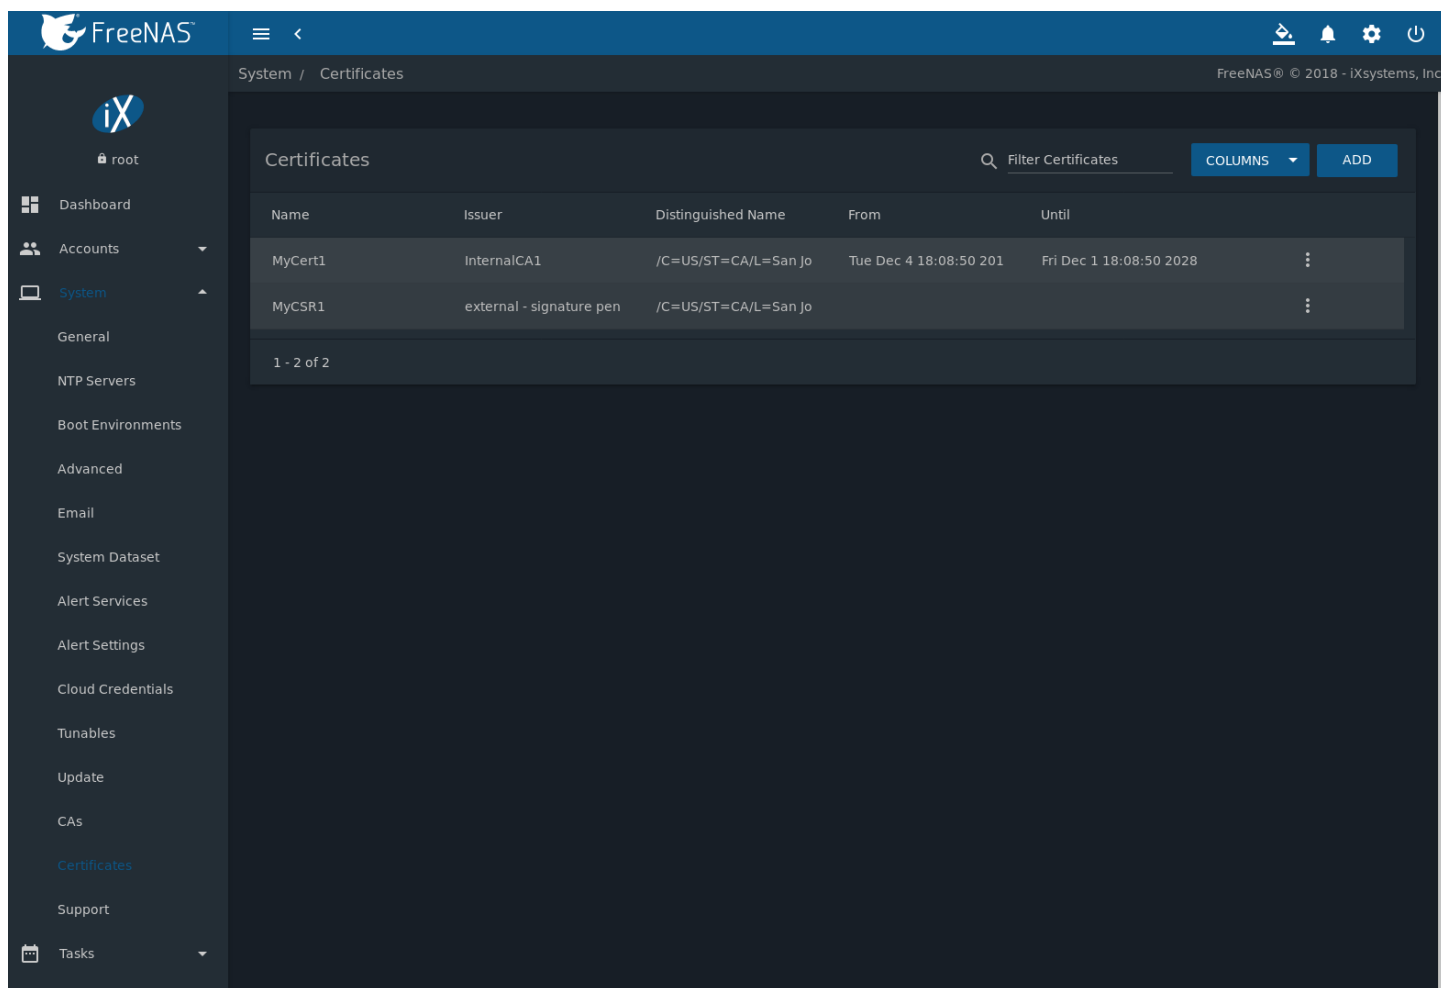



Fig. 6.24: Managing Certificates

Clicking  (Options) for an entry shows these configuration buttons:

- **View:** use this option to view the contents of an existing *Certificate*, *Private Key*, or to edit the *Identifier*.
- **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.

6.14 Support

The FreeNAS® *Support* option, shown in Figure 6.25, provides a built-in ticketing system for generating bug reports and feature requests.

The screenshot shows the FreeNAS web interface with the 'Support' menu item selected in the left sidebar. The main content area displays a registration form for the FreeNAS issue tracker. The form includes a header with instructions to search the issue tracker first, followed by fields for Username, Password, Type (a dropdown menu), an 'Attach Debug' checkbox, Subject, and Description. A 'SUBMIT' button is at the bottom of the form. The top of the interface shows the FreeNAS logo and navigation icons, and the bottom of the sidebar lists various system categories like Accounts, System, Tasks, Network, Storage, etc.

Fig. 6.25: Support Menu

This screen provides a built-in interface to the FreeNAS® issue tracker located at <https://redmine.ixsystems.com/projects/freenas/issues>. When using FreeNAS® bug tracker for the first time, 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 a similar issue is already present and has not been marked *Closed* or *Resolved*, comment on that issue, adding new information to help solve it. If similar issues have already been *Closed* or *Resolved*, create a new issue and refer to the previous issue.

Note: Update the system to the latest version of STABLE and retest before reporting an issue. Newer versions of the software might have already fixed the problem.

To generate a report using the built-in *Support* screen, complete these 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 until a registered *Username* and *Password* are entered. The field remains empty 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:** enabling this option is recommended so an overview of the system hardware, build string, and configuration is automatically generated and included with the ticket. Generating and attaching a debug to the ticket can take some time. An error will occur if the debug is more than the file size limit of 20 Mib.
- **Subject:** enter a descriptive title for the ticket. A good *Subject* makes it easy 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.

After 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 URL so to view status or add additional information to the report.

TASKS

The Tasks section of the web interface is used to configure repetitive tasks:

- [Cron Jobs](#) (page 108) schedules a command or script to automatically execute at a specified time
- [Init/Shutdown Scripts](#) (page 110) configures a command or script to automatically execute during system startup or shutdown
- [Rsync Tasks](#) (page 111) schedules data synchronization to another system
- [S.M.A.R.T. Tests](#) (page 116) schedules disk tests
- [Periodic Snapshot Tasks](#) (page 118) schedules automatic creation of filesystem snapshots
- [Replication Tasks](#) (page 120) automate the replication of snapshots to a remote system
- [Resilver Priority](#) (page 131) controls the priority of resilvers
- [Scrub Tasks](#) (page 132) schedules scrubs as part of ongoing disk maintenance
- [Cloud Sync Tasks](#) (page 134) schedules data synchronization to cloud providers

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

Note: By default, [Scrub Tasks](#) (page 132) are run once a month by an automatically-created task. [S.M.A.R.T. Tests](#) (page 116) and [Periodic Snapshot Tasks](#) (page 118) must be set up manually.

7.1 Cron Jobs

[cron\(8\)](https://www.freebsd.org/cgi/man.cgi?query=cron) (<https://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.

Navigate to *Tasks* → *Cron Jobs* and click *ADD* to create a cron job. [Figure 7.1](#) shows the configuration screen that appears.

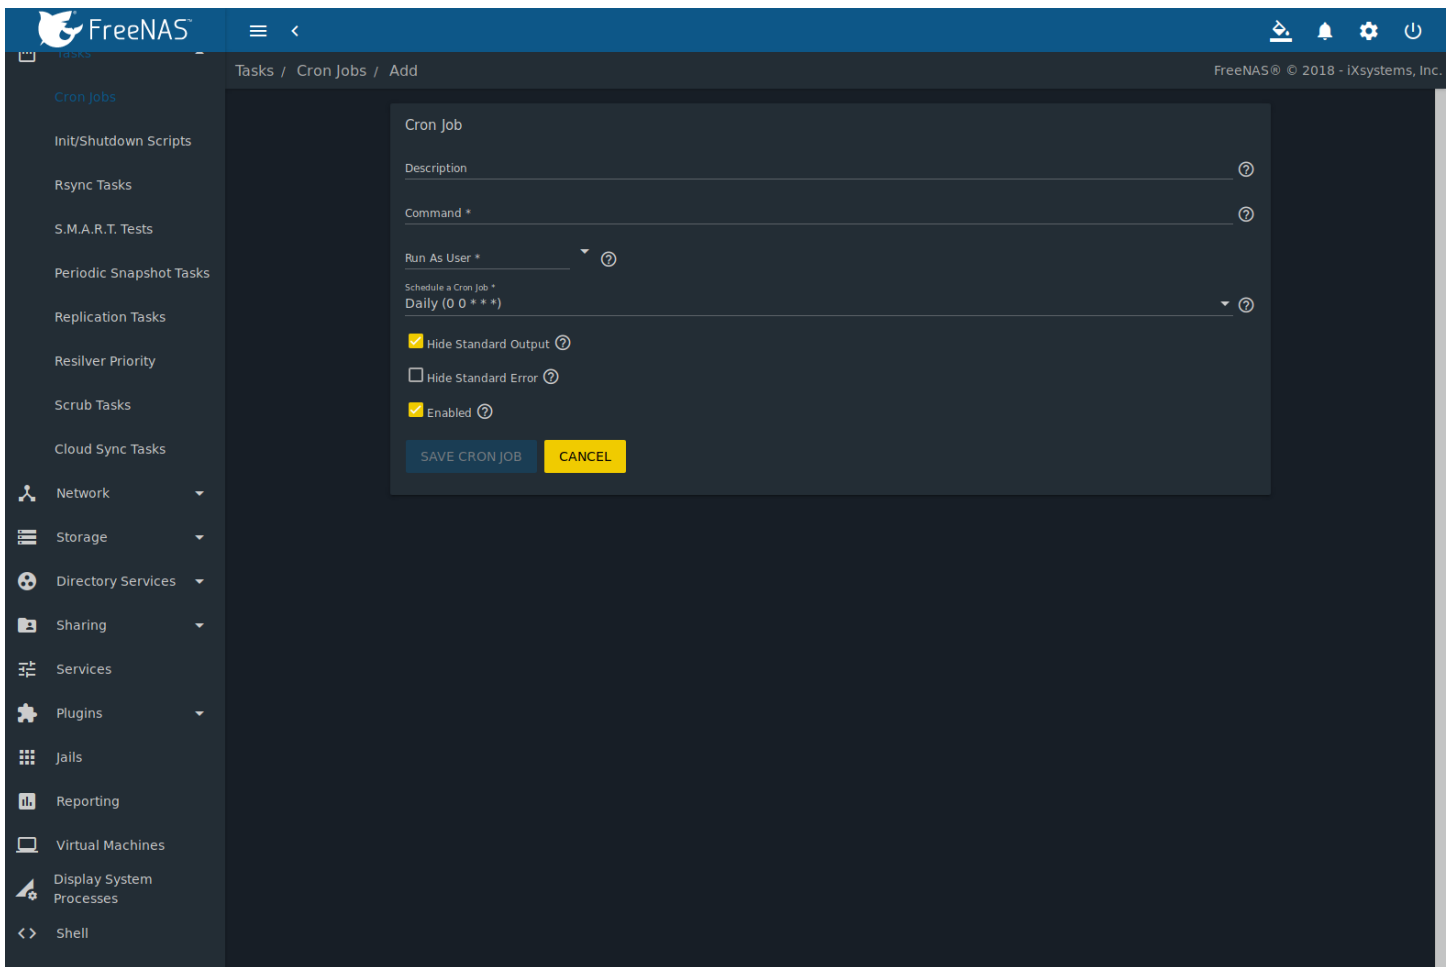


Fig. 7.1: Creating a Cron Job

Table 7.1 lists the configurable options for a cron job.

Table 7.1: Cron Job Options

Setting	Value	Description
Description	string	Enter a description of the cron job.
Command	drop-down menu	Enter the full path to the command or script to be run. If it is a script, testing it at the command line is recommended to ensure it works.
Run As User	string	Select a user account to run the command. The user must have permissions allowing them to run the command or script.
Schedule a Cron Job	drop-down menu	Select how often to run the cron job. Choices are <i>Hourly</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Select <i>Custom</i> to open the advanced scheduler.
Hide Standard Output	checkbox	Hide standard output (stdout) from the command. When unset, any standard output is mailed to the user account cron used to run the command.
Hide Standard Error	checkbox	Hide error output (stderr) from the command. When unset, any error output is mailed to the user account cron used to run the command.
Enable	checkbox	Enable this cron job. When unset, disable the cron job without deleting it.

Cron jobs are shown in *Tasks* → *Cron Jobs*. This table displays the user, command, description, schedule, and whether the job is enabled. This table is adjustable by setting the different column checkboxes above it. Set *Toggle* to display all options in the table. Click ⓘ (Options) for to show the *Run Now*, *Edit*, and *Delete* options.

Note: % symbols are automatically escaped and do not need to be prefixed with backslashes. For example, use `date`

'+%Y-%m-%d' in a cron job to generate a filename based on the date.

7.2 Init/Shutdown Scripts

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

Figure 7.2 shows the screen that opens after going to *Tasks* → *Init/Shutdown Scripts* and clicking *ADD*. Table 7.2 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.

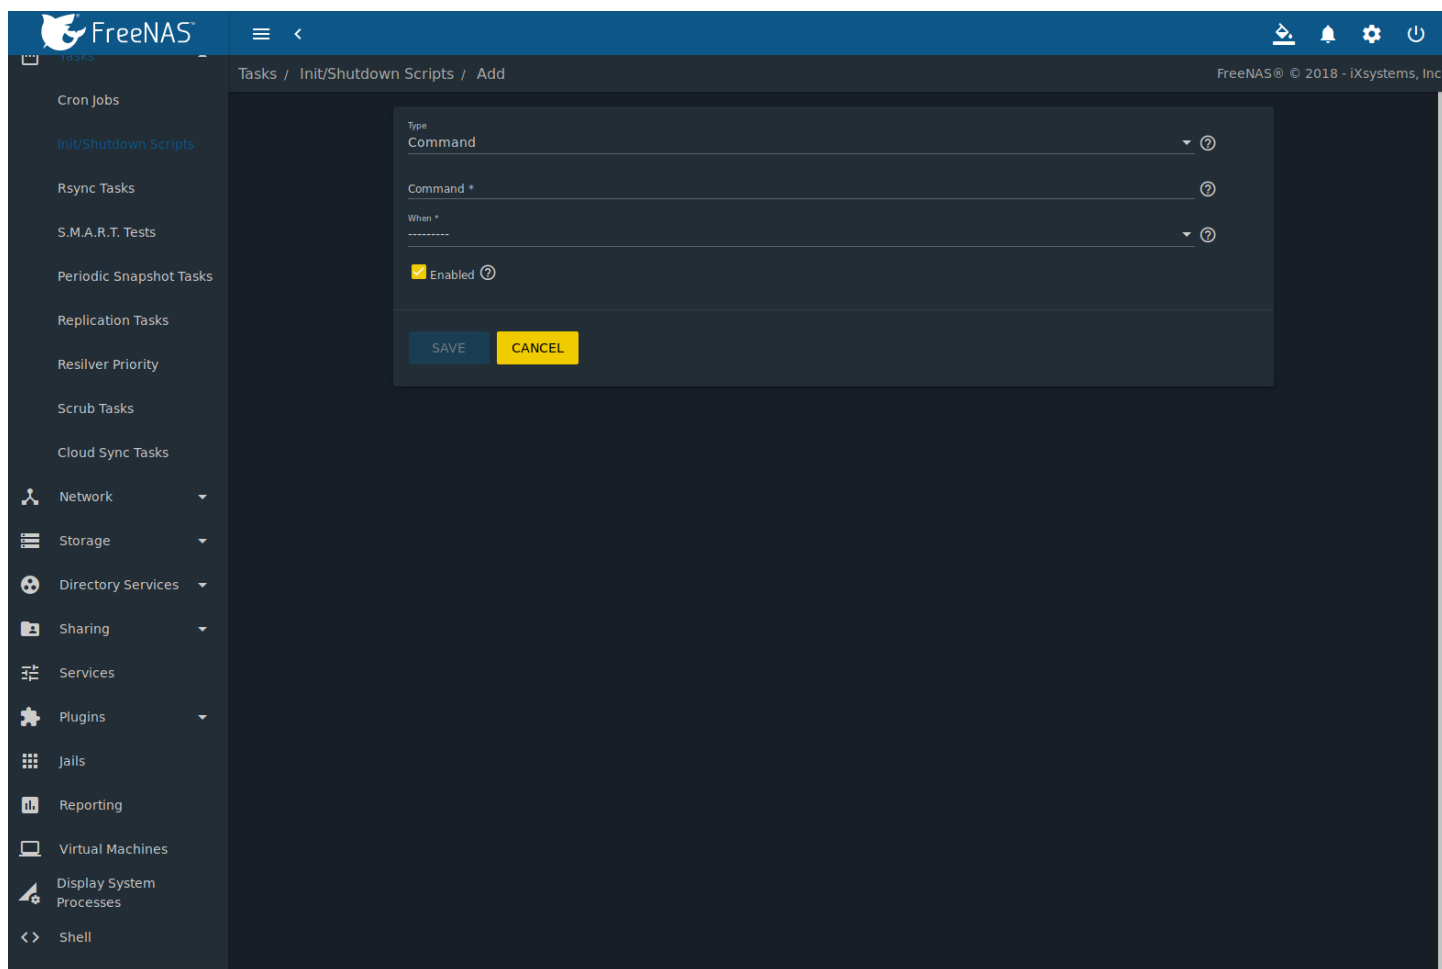


Fig. 7.2: Add an Init/Shutdown Script

Table 7.2: Options When Adding an Init/Shutdown Script

Setting	Value	Description
Type	drop-down menu	Select <i>Command</i> for an executable or <i>Script</i> for an executable script.
Command or Script	string	If <i>Command</i> is selected, enter the command plus any desired options. If <i>Script</i> is selected, browse to the location of the script.

Continued on next page

Table 7.2 – continued from previous page

Setting	Value	Description
When	drop-down menu	Select when the command or script runs. <i>Pre Init</i> is very early in boot process before mounting filesystems, <i>Post Init</i> is towards the end of boot process before FreeNAS services start, or at <i>Shutdown</i> .
Enabled	checkbox	Enable this task. Unset to disable the task without deleting it.

Init/Shutdown tasks are shown in *Tasks* → *Init/Shutdown Scripts*. Click  (Options) to see the *Edit* and *Delete* buttons.

7.3 Rsync Tasks

Rsync (<https://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 is 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 120) are often the faster and better solution.

Rsync is single-threaded and gains little from multiple processor cores. To see whether rsync is currently running, use `pgrep rsync` from the [Shell](#) (page 325).

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 the configured settings of the tree 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® web interface under *Services* → *Rsync Configure* → *Rsync Module*. In other operating systems, the module is defined in `rsyncd.conf(5)` (<https://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 7.3 shows the screen that appears after navigating to *Tasks* → *Rsync Tasks* and clicking *ADD*. Table 7.3 summarizes the configuration options available when creating an rsync task.

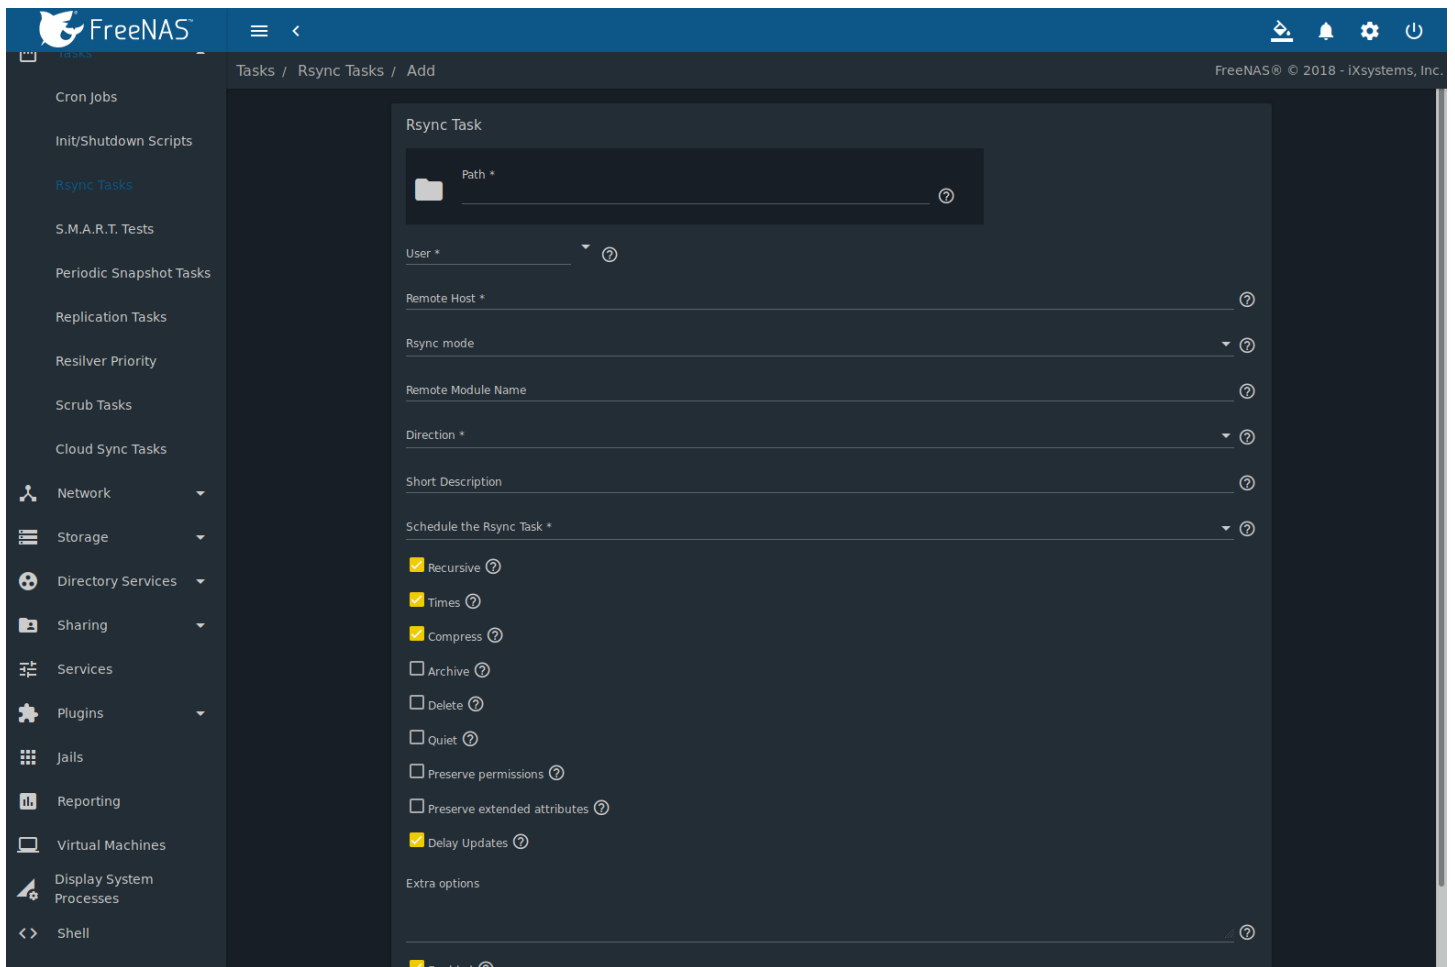


Fig. 7.3: Adding an Rsync Task

Table 7.3: Rsync Configuration Options

Setting	Value	Description
Path	browse button	<i>Browse</i> to the path to be copied. Path lengths cannot be greater than 255 characters.
User	drop-down menu	Select the user to run the rsync task. The user selected must have permissions to write to the specified directory on the remote host.
Remote Host	string	Enter the IP address or hostname of the remote system that will store the copy. Use the format <i>username@remote_host</i> if the username differs on the remote host.
Remote SSH Port	integer	Only available in <i>Rsync over SSH</i> mode. Allows specifying an SSH port other than the default of 22.
Rsync mode	drop-down menu	The choices are <i>Rsync Module</i> mode or <i>Rsync over SSH</i> mode
Remote Module Name	string	At least one module must be defined in rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html) of the rsync server or in the <i>Rsync Modules</i> of another system.
Remote Path	string	Only appears when using <i>Rsync over SSH</i> mode. Enter the existing path on the remote host to sync with, for example, <i>/mnt/pool</i> . Note that the path length cannot be greater than 255 characters.
Validate Remote Path	checkbox	Verifies the existence of the <i>Remote Path</i> .

Continued on next page

Table 7.3 – continued from previous page

Setting	Value	Description
Direction	drop-down menu	Direct the flow of the data to the remote host. Choices are <i>Push Pull</i> . Default is to push to a remote host.
Short Description	string	Enter a description of the rsync task.
Schedule the Rsync Task	drop-down menu	Choose how often to run the task. Choices are <i>Hourly, Daily, Weekly, Monthly, or Custom</i> . Select <i>Custom</i> to open the advanced scheduler.
Recursive	checkbox	Set to include all subdirectories of the specified directory. When unset, only the specified directory is included.
Times	checkbox	Set to preserve the modification times of files.
Compress	checkbox	Set to reduce the size of the data to transmit. Recommended for slow connections.
Archive	checkbox	When set, rsync is run recursively, preserving symlinks, permissions, modification times, group, and special files. When run as root, owner, device files, and special files are also preserved. Equivalent to <code>rsync -rlptgoD</code> .
Delete	checkbox	Set to delete files in the destination directory that do not exist in the source directory.
Quiet	checkbox	Set to suppress informational messages from the remote server.
Preserve permissions	checkbox	Set to preserve original file permissions. This is useful when the user is set to <i>root</i> .
Preserve extended attributes	checkbox	Extended attributes (https://en.wikipedia.org/wiki/Extended_file_attributes) are preserved, but must be supported by both systems.
Delay Updates	checkbox	Set to save the temporary file from each updated file to a holding directory until the end of the transfer when all transferred files are renamed into place.
Extra options	string	Additional <code>rsync(1)</code> (http://rsync.samba.org/ftp/rsync/rsync.html) options to include. Note: The * character must be escaped with a backslash (<code>*</code> .txt) or used inside single quotes. (<code>'*.txt'</code>)
Enabled	checkbox	Enable this rsync task. Unset to disable this rsync task without deleting it.

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

Created rsync tasks are listed in *Rsync Tasks*. Click  (Options) for an entry to display buttons for *Edit*, *Delete*, or *Run Now*.

7.3.1 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 pool 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*. 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* option is enabled so that the original permissions are not overwritten by the *root* user

On *PULL*, an rsync module is defined in *Services* → *Rsync Configure* → *Rsync Module, ADD*. 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](#) (page 253).

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

7.3.2 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 325) 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. Refer to [Key-based Authentication](https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/openssh.html#security-ssh-keygen) (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/openssh.html#security-ssh-keygen) for more information.

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 AAAAB3NzaC1yc2EAAAADAQABAAQCA1lBEXRgw1W8y8k+lXP1VR3xsmVSjtsoyIzV/PlQPoSrWotUQzqILq0SmUpViAAv4Ik3T8NtXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4hdcD7Y5mvU3MAEeDClt02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmastio0qmDDcp/kOxT+S6DFNDBY6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+19RaEXMRuTyQgqJB/rsRcmJX5fApdDmNfwrRSxLjDvUzfywnjFHlKk/+TQIT1gg1QQaj21PJD9pnDVF0AiJrWyWnR root@freenas.local
```

Go to *PULL* and paste (or append) the copied key into the *SSH Public Key* field of *Accounts* → *Users* → *root* → ⓘ (Options) → *Edit*, or the username of the specified rsync user account. The paste for the above example is shown in [Figure 7.4](#). When pasting the key, ensure that it is pasted as one long line and, if necessary, remove any extra spaces representing line breaks.

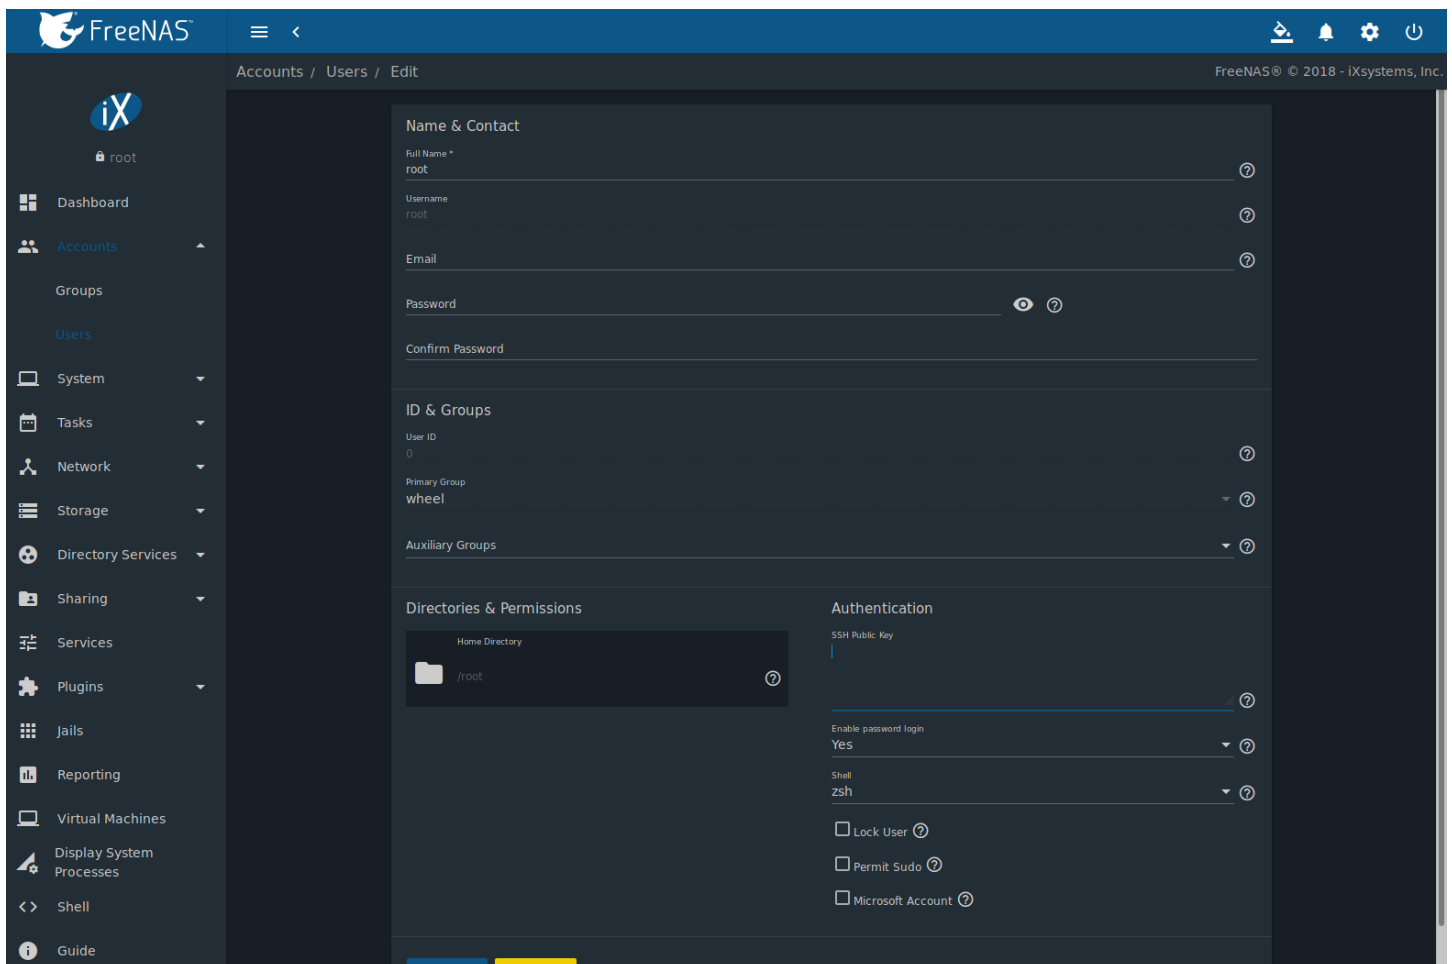


Fig. 7.4: Pasting the User SSH Public Key

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

Next, copy the host key of *PULL* using Shell on *PUSH*. The 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:

- 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* option is enabled 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.

7.4 S.M.A.R.T. Tests

S.M.A.R.T. (<https://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. Replace the drive when a failure is anticipated by S.M.A.R.T. Most modern ATA, IDE, and SCSI-3 hard drives support S.M.A.R.T. – refer to the drive documentation for confirmation.

Click *Tasks* → *S.M.A.R.T. Tests* and *ADD* to add a new scheduled S.M.A.R.T. test. [Figure 7.5](#) shows the configuration screen that appears. Tests are listed under *S.M.A.R.T. Tests*. After creating tests, check the configuration in *Services* → *S.M.A.R.T.*, then click the power button for the S.M.A.R.T. service in *Services* to activate the service. The S.M.A.R.T. service will not start if there are no pools.

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.

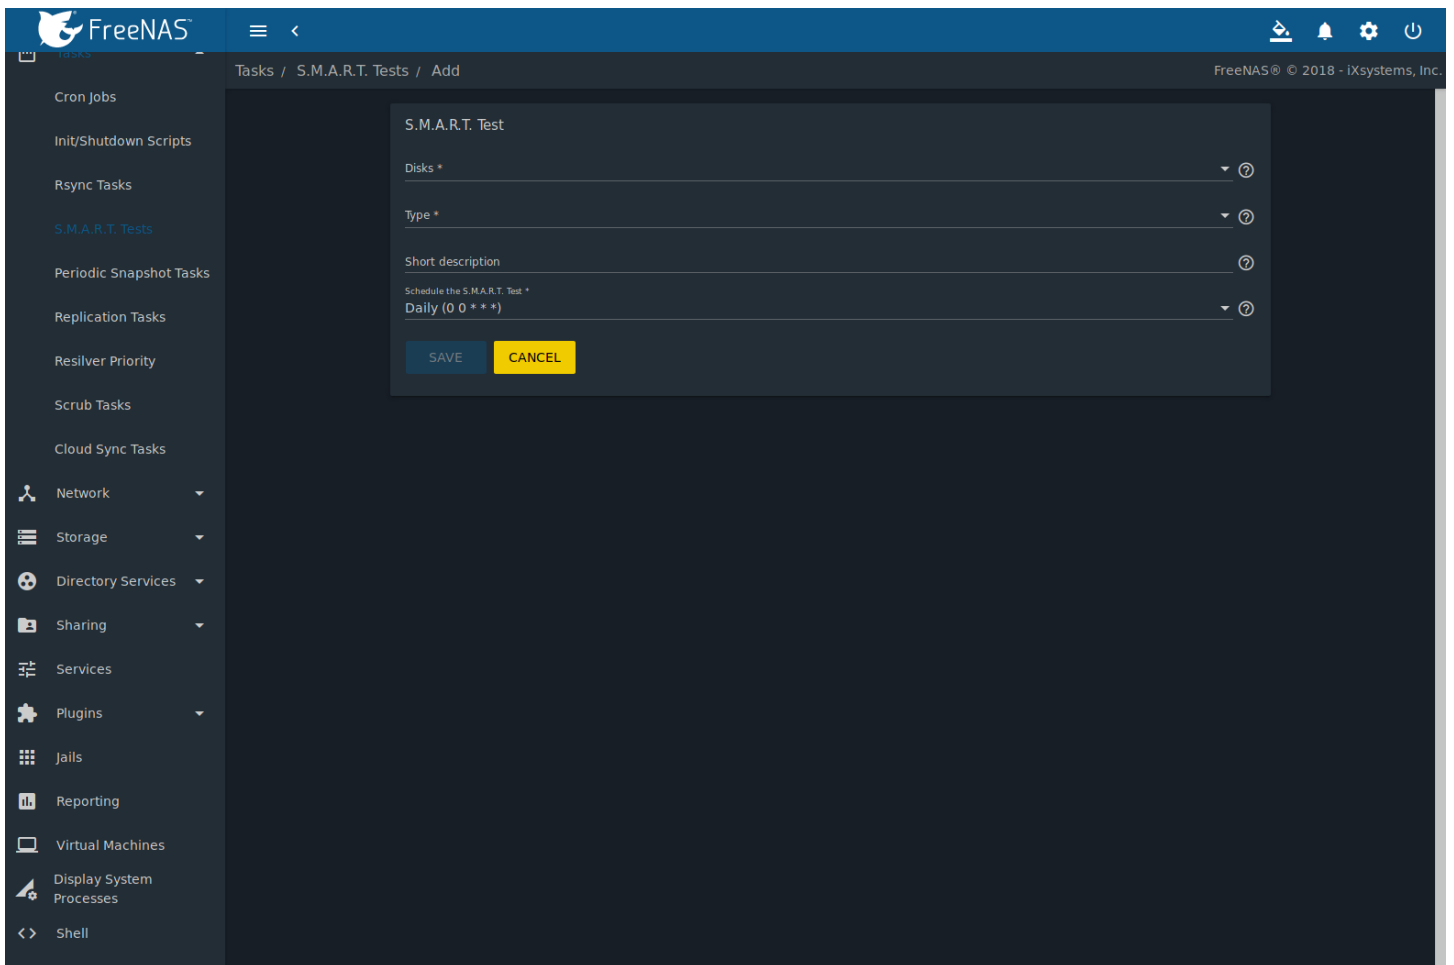


Fig. 7.5: Adding a S.M.A.R.T. Test

Table 7.4 summarizes the configurable options when creating a S.M.A.R.T. test.

Table 7.4: S.M.A.R.T. Test Options

Setting	Value	Description
Disks	drop-down menu	Select the disks to monitor.
Type	drop-down menu	Choose the test type. See smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) for descriptions of each type. Some test types will degrade performance or take disks offline. Avoid scheduling S.M.A.R.T. tests simultaneously with scrub or resilver operations.
Short description	string	Optional. Enter a description of the S.M.A.R.T. test.
Schedule the S.M.A.R.T. Test	drop-down menu	Choose how often to run the task. Choices are <i>Hourly</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Select <i>Custom</i> to open a visual scheduler for selecting minutes, hours, days, month, and days of week.

An example configuration is to schedule a *Short Self-Test* once a week and a *Long Self-Test* once a month. These tests do 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*, consider 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 325).

The results of a test can be checked from [Shell](#) (page 325) by specifying the name of the drive. For example, to see the results for disk `ada0`, type:

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

When an email address is entered in the *Email* field of *Services* → *S.M.A.R.T.* → *Configure*, the system sends an email to that address when a test fails. Logging information for S.M.A.R.T. tests can be found in `/var/log/daemon.log`.

7.5 Periodic Snapshot Tasks

A periodic snapshot task allows scheduling the creation of read-only versions of pools 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 keep 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, store them for a period of time, and store them on another system, typically using [Replication Tasks](#) (page 120). 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.

A pool must exist before a snapshot can be created. Creating a pool is described in [Pools](#) (page 154).

To create a periodic snapshot task, navigate to *Tasks* → *Periodic Snapshot Tasks* and click *ADD*. This opens the screen shown in [Figure 7.6](#). [Table 7.5](#) describes the fields in this screen.

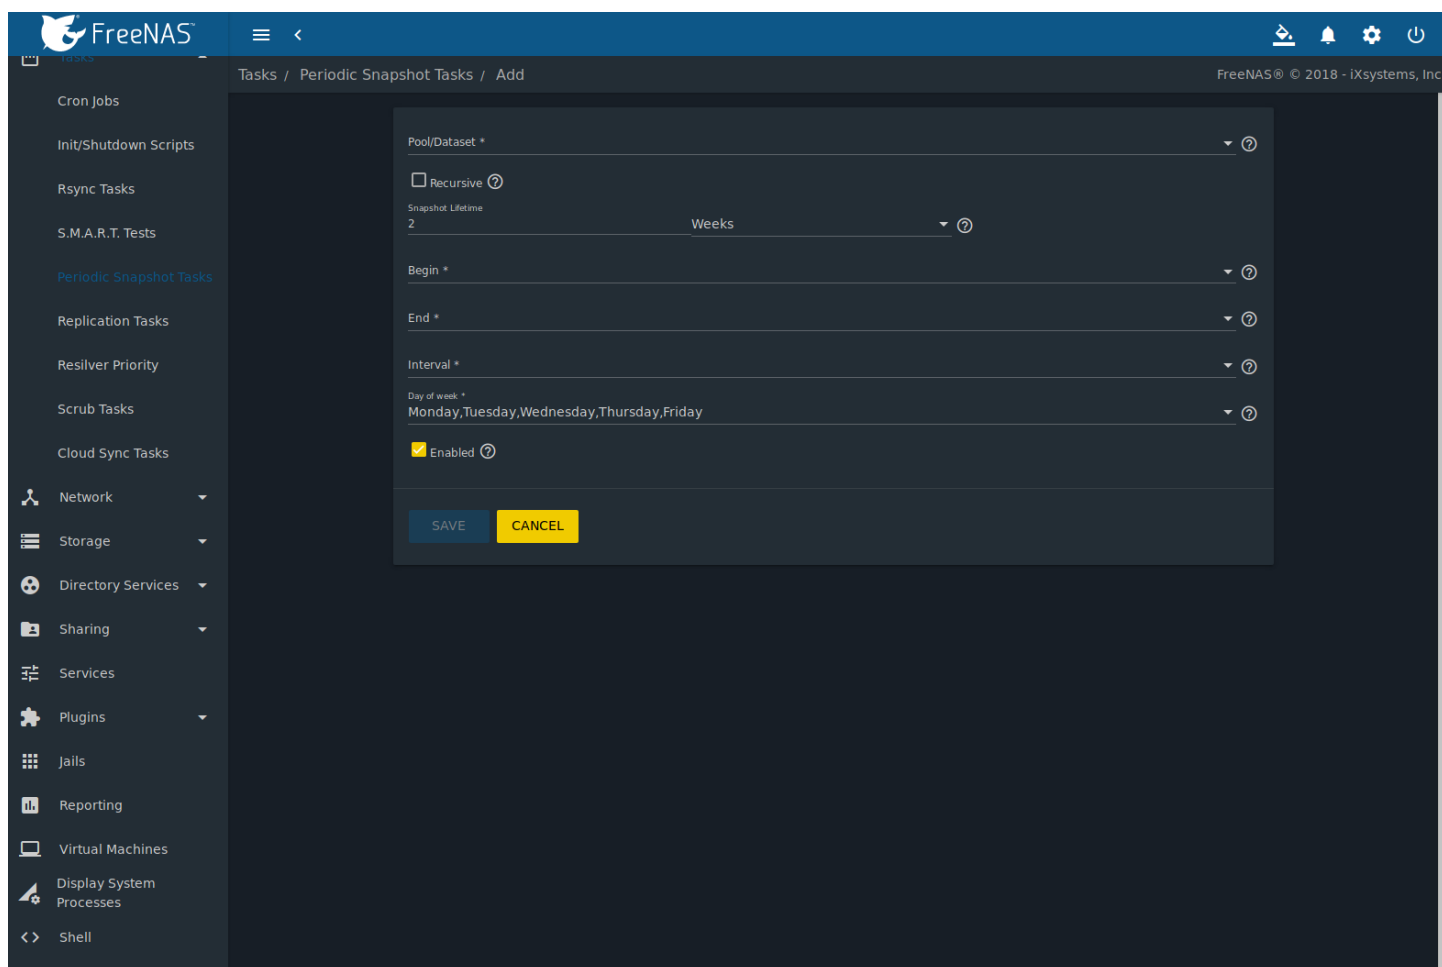



Fig. 7.6: Creating a Periodic Snapshot

Table 7.5: Options When Creating a Periodic Snapshot

Setting	Value	Description
Pool/Dataset	drop-down menu	Select an existing pool, dataset, or zvol.
Recursive	checkbox	Set this option to take separate snapshots of the pool or dataset and each of the child datasets. Deselect to take a single snapshot of the specified pool or dataset with no child datasets.
Snapshot Life-time	integer and drop-down menu	Define a length of time to retain the snapshot on this system. After the time expires, the snapshot is removed. Snapshots replicated to other systems are not affected.
Begin	drop-down menu	Choose the hour and minute when the system can begin taking snapshots.
End	drop-down menu	Choose the hour and minute when the system must stop taking snapshots.
Interval	drop-down menu	Define how often the system takes snapshots between <i>Begin</i> and <i>End</i> times.
Day of week	checkboxes	Choose the days of the week to take the snapshots.
Enabled	checkbox	Unset to disable the task without deleting it.

If the *Recursive* option is enabled, 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 *SAVE* when finished customizing the task. Defined tasks are listed alphabetically in *Periodic Snapshot Tasks*. Click  (Options) for an entry to display the *Edit* and *Delete* buttons.

7.6 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 118) 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 172) for more information on creating clones.

7.6.1 Examples: Common Configuration

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

7.6.1.1 Alpha (Source)

Alpha is the source computer with the data to be replicated. It is at IP address *10.0.0.102*. A *pool* (page 154) named *alphapool* has already been created, and a *dataset* (page 166) named *alphadata* has been created on that pool. 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.

Click *Tasks* → *Periodic Snapshot Tasks* and *ADD* to create a periodic snapshot of the source dataset. Add the *alphapool/alphadata* dataset to the *Pool/Dataset* field. [Figure 7.7](#) shows the configured periodic snapshot.




Fig. 7.7: Create a Periodic Snapshot for Replication

This example creates a snapshot of the *alphapool/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.

7.6.1.2 Beta (Destination)

Beta is the destination computer where the replicated data will be copied. It is at IP address *10.0.0.118*. A *pool* (page 154) named *betapool* has already been created.

Snapshots are transferred with *SSH* (page 263). 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*.

7.6.2 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 then *ADD*.

Select *alphapool/alphadata* as the dataset to replicate. *betapool* is the destination pool where *alphadata* snapshots are replicated. The *Setup mode* dropdown is set to *Semi-Automatic* as shown in Figure 7.8. 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* is enabled in *System* → *General* on the destination computer, set *Remote HTTP/HTTPS*

Port to the HTTPS port and ensure *Remote HTTPS* is enabled when creating the replication on the source computer.

The screenshot displays the 'Add Replication' configuration window in FreeNAS. The left sidebar contains a navigation menu with icons and labels for various system functions. The main content area is titled 'Tasks / Replication Tasks / Add Replication' and includes the following fields and options:

- Pool/Dataset ***: alphapool/alphadata
- Remote ZFS Pool/Dataset ***: betapool
- ☐ Recursively Replicate Child Dataset Snapshots
- ☐ Delete Stale Snapshots on Remote System
- Replication Stream Compression**: lz4 (fastest)
- Limit (kpbs)**: 0
- Begin Time**: 00:00:00
- End Time**: 23:59:00
- ☒ Enabled
- Setup Mode**: Semi-Automatic
- Remote Hostname ***: 10.231.1.3
- Remote HTTP/HTTPS Port**: 80
- ☒ Remote HTTPS
- Remote Auth Token**: (empty field)
- Encryption Cipher**: standard
- ☐ Dedicated User Enabled
- Dedicated User**: (empty field)
- Remote Hostkey**: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDDG0SiC78AFFQ+QNCN5qY6TvdT3Hi6u4/3PksGh63ofv4IvOGWYslwrrw3NNdICFvH3wHm91J4WKrohZw
- SCAN SSH KEY** button

Fig. 7.8: Add Replication Dialog, Semi-Automatic

The *Remote Auth Token* field expects a special token from the *Beta* computer. On *Beta*, navigate to *Tasks* → *Replication Tasks*, and click *REPLICATION TOKEN*. A dialog showing the temporary authorization token is shown as in [Figure 7.9](#).

Highlight the temporary authorization token string with the mouse and copy it.

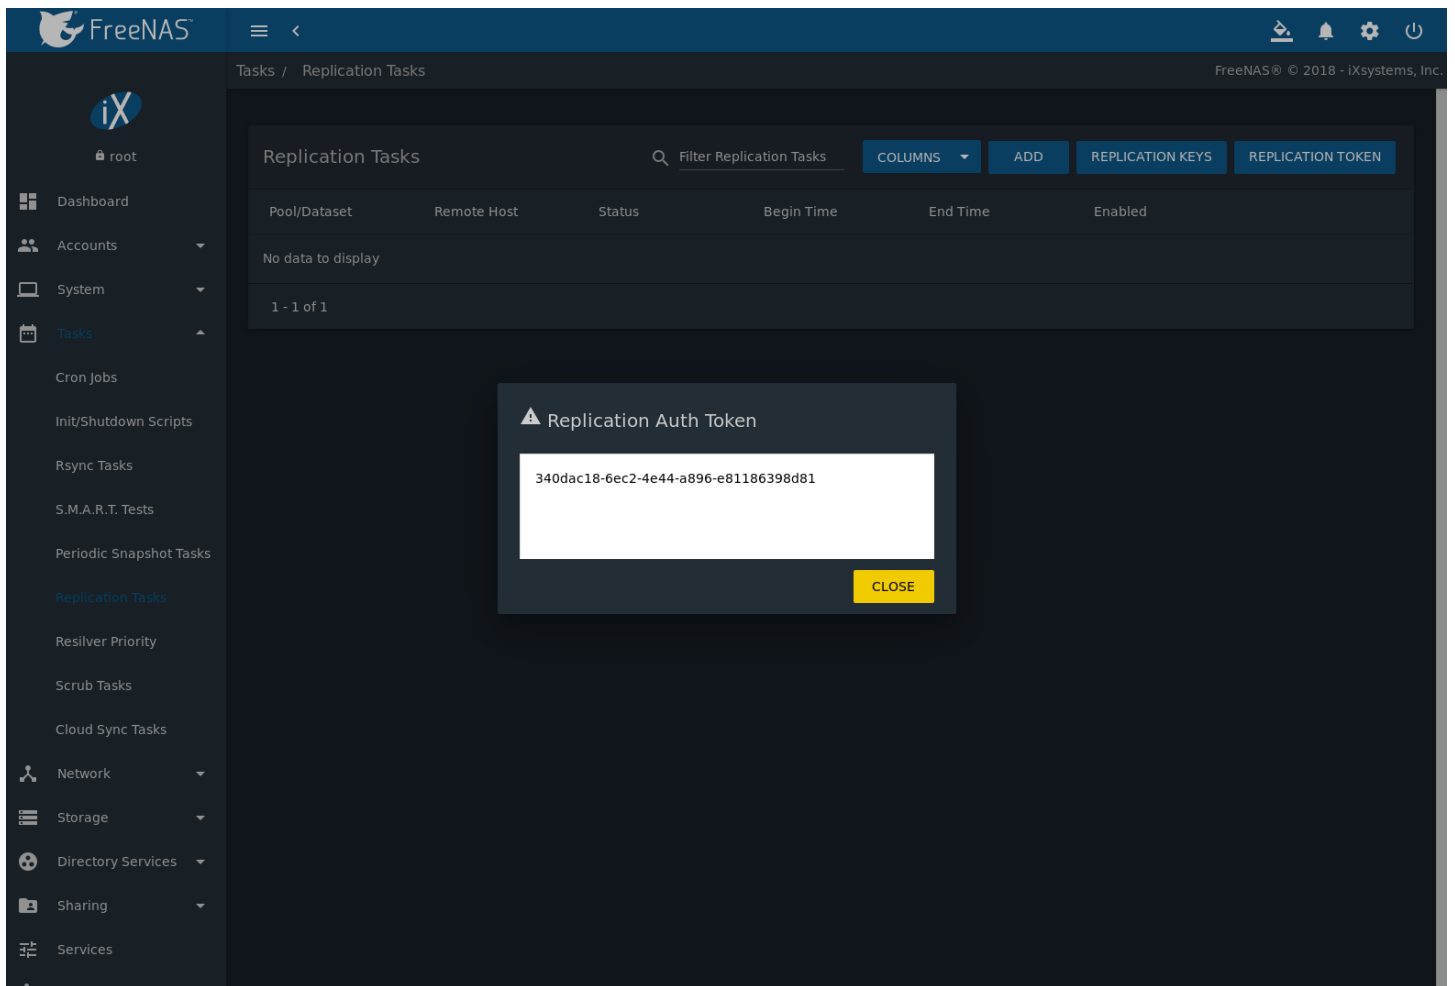


Fig. 7.9: Temporary Authentication Token on Destination

On the *Alpha* system, paste the copied temporary authorization token string into the *Remote Auth Token* field as shown in Figure 7.10.

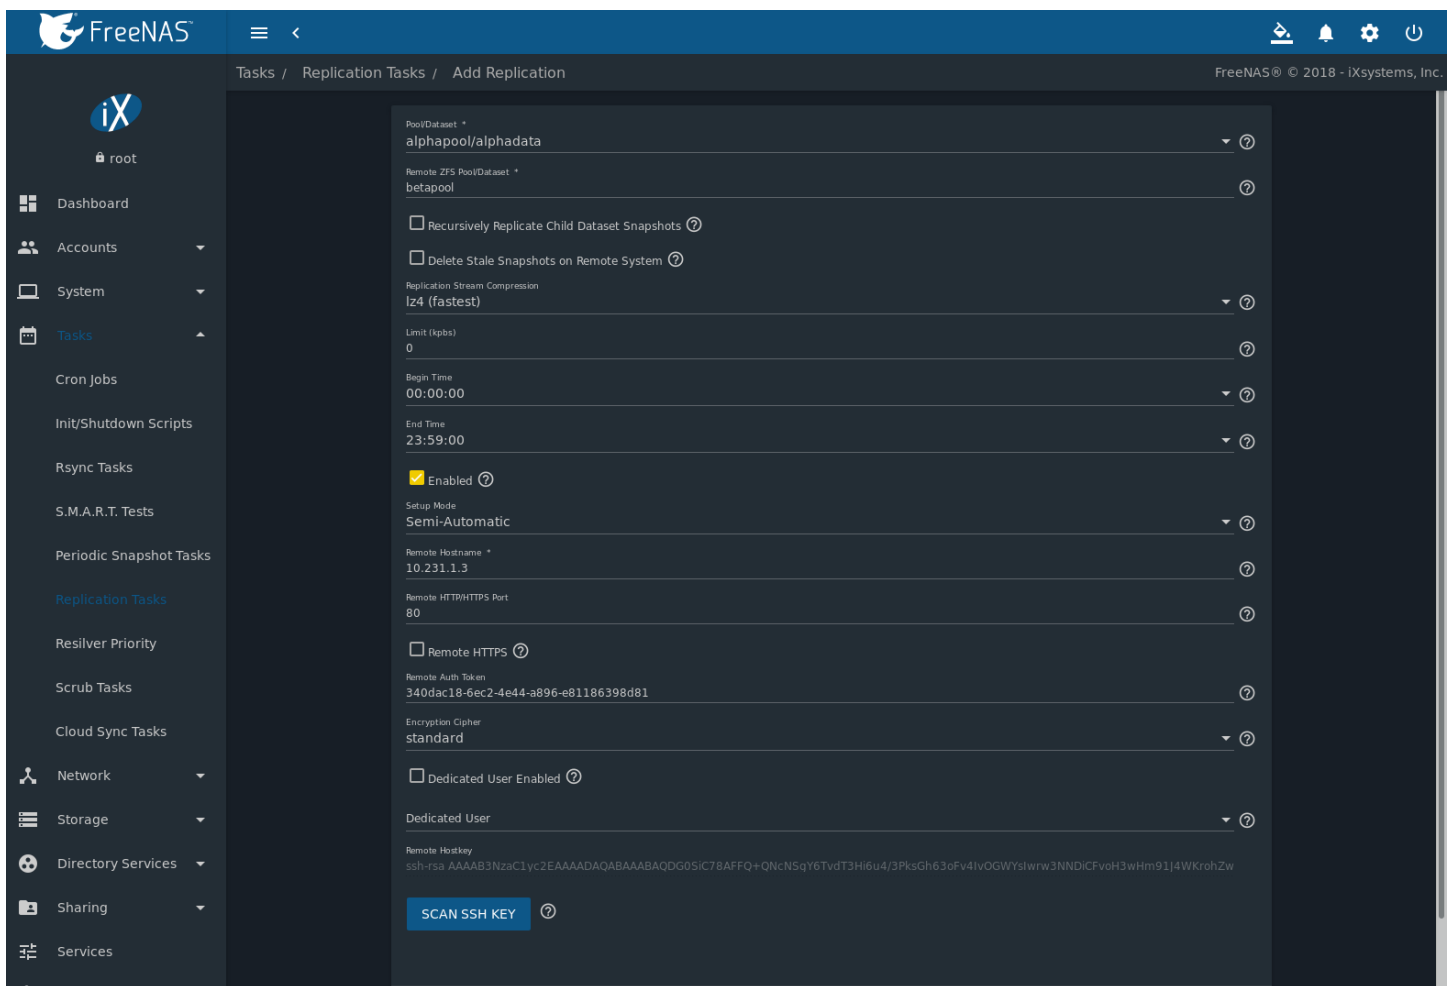


Fig. 7.10: Temporary Authentication Token Pasted to Source

Finally, click **SAVE** 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 130) 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.

7.6.3 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 120) has already been created, go to **Tasks** → **Periodic Snapshot Tasks**, click ⓘ (Options) for the task and **Delete** to remove it before continuing.

On *Alpha*, click **Accounts** → **Users** then **ADD**. Enter *repluser* for *Username*, enter */mnt/alphapool/repluser* in the *Home Directory* field, enter *Replication Dedicated User* for the *Full Name*, and set *Enable password login* to *No*. Leave the other fields at their default values, but note the *User ID* number. Click **SAVE** to create the user.

On *Beta*, the same dedicated user must be created as was created on the sending computer. Click **Accounts** → **Users** then

ADD. Enter the *User ID* number from *Alpha*, *repluser* for *Username*, enter `/mnt/betapool/repluser` in the *Home Directory* field, enter *Replication Dedicated User* for the *Full Name*, and set *Enable password login* to *No*. Leave the other fields at their default values. Click **SAVE** to create the user.

A dataset with the same name as the original must be created on the destination computer, *Beta*. Navigate to *Storage* → *Pools*, click *betapool*, then *Options* and *Add Dataset*. Enter *alphadata* as the *Name*, then click **SAVE**.

The replication user must be given permissions to the destination dataset. On *Beta*, open a *Shell* (page 325) and enter this command:

```
zfs allow -ldu repluser create,destroy,diff,mount,readonly,receive,release,send,userprop betapool/
↪alphadata
```

The destination dataset must also be set to read-only. Enter this command in the *Shell* (page 325):

```
zfs set readonly=on betapool/alphadata
```

The replication user must also be able to mount datasets. On *Beta*, go to *System* → *Tunables* and click **ADD**. Enter *vfs.usermount* for the *Variable*, *1* for the *Value*, and choose *Sysctl* from the *Type* drop-down. Click **SAVE**.

Back on *Alpha*, create a *periodic snapshot* (page 118) of the source dataset. [Figure 7.7](#) shows the configuration.

On *Alpha*, create the replication task by clicking *Replication Tasks* and click **ADD**. *alphapool/alphadata* is selected as the dataset to replicate. *betapool/alphadata* is the destination pool and dataset where *alphadata* snapshots are replicated.

The *Setup mode* dropdown is set to *Semi-Automatic* as shown in [Figure 7.8](#). 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* is enabled in *System* → *General* on the destination computer, set the *Remote HTTP/HTTPS Port* to the HTTPS port and enable the *Remote HTTPS* when creating the replication on the source computer.

The *Remote Auth Token* field expects a special token from the *Beta* computer. On *Beta*, click *Tasks* → *Replication Tasks*, then **REPLICATION TOKEN**. A dialog showing the temporary authorization token is shown as in [Figure 7.9](#).

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 7.10](#).

Set the *Dedicated User Enabled* option. Choose *repluser* in the *Dedicated User* drop-down.

Click **SAVE** 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 325) must be set on each new destination dataset.

7.6.4 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 pools and datasets, enabling SSH, and copying encryption keys will vary when the destination computer is not a FreeNAS® system.

7.6.4.1 Encryption Keys

A public encryption key must be copied from *Alpha* to *Beta* to allow a secure connection without a password prompt. On *Alpha*, navigate to *Tasks* → *Replication Tasks* and click *REPLICATION KEYS*. This produces the window shown in [Figure 7.11](#). Use the mouse to highlight the key data shown in the window, then copy it.

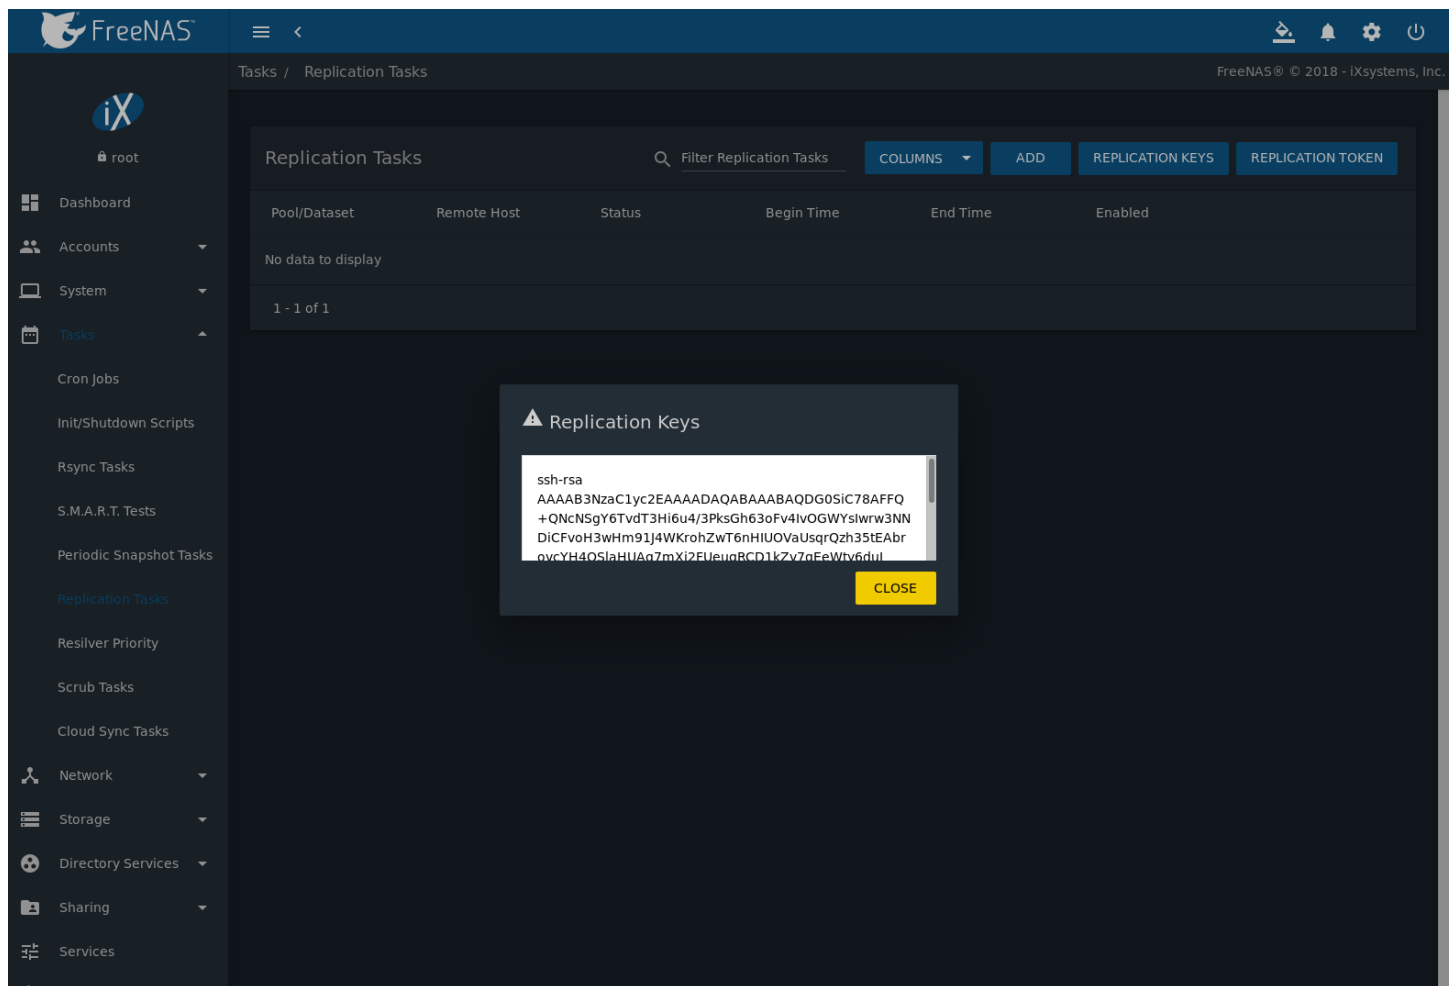



Fig. 7.11: Copy the Replication Key

On *Beta*, go to *Accounts* → *Users*. Click  (Options) for the *root* account, then *Edit*. Paste the copied key into the *SSH Public Key* field and click *SAVE* as shown in [Figure 7.12](#).

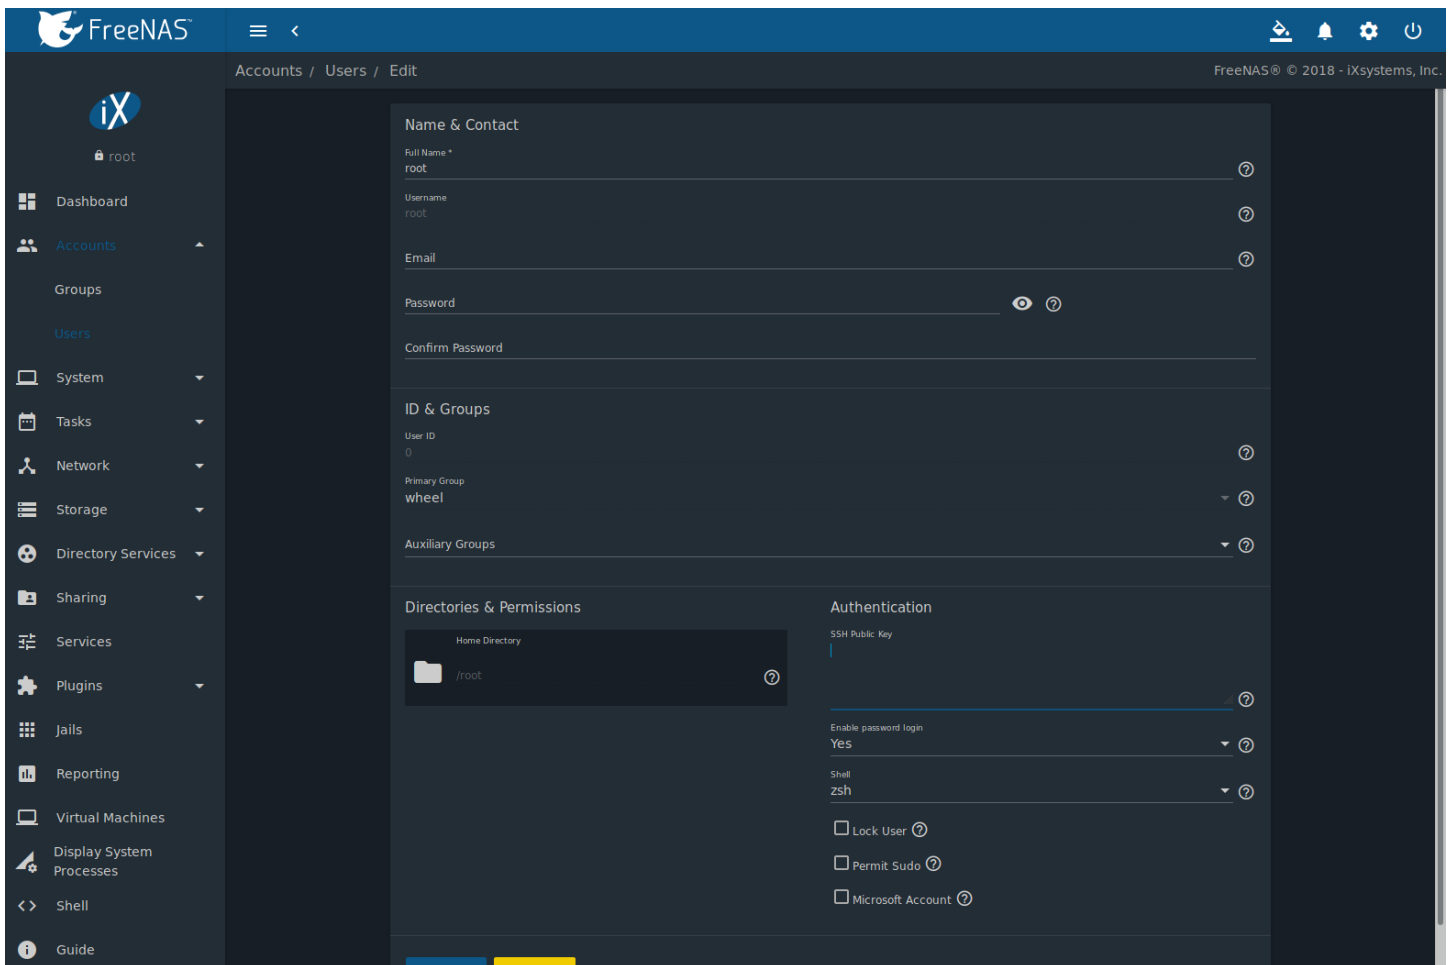


Fig. 7.12: Paste the Replication Key

Back on *Alpha*, create the replication task by clicking *Replication Tasks* and *ADD*. *alphapool/alphadata* is selected as the dataset to replicate. The destination pool is *betapool*. The *alphadata* dataset and snapshots are replicated there. The IP address of *Beta* is entered in the *Remote Hostname* field as shown in [Figure 7.13](#). A hostname can be entered here if local DNS resolves for that hostname.

Click the *SCAN SSH KEY* button to retrieve the SSH host keys from *Beta* and fill the *Remote Hostkey* field. Finally, click *SAVE* 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 130) for information about restricting when replication is allowed to run.

The screenshot shows the 'Add Replication' dialog in the FreeNAS web interface. The left sidebar contains navigation links like Dashboard, Accounts, System, Tasks, Cron Jobs, etc. The main area is titled 'Tasks / Replication Tasks / Add Replication'. The dialog fields are as follows:

- Pool/Dataset: alphapool/alphadata
- Remote ZFS Pool/Dataset: betapool
- ☐ Recursively Replicate Child Dataset Snapshots
- ☐ Delete Stale Snapshots on Remote System
- Replication Stream Compression: lz4 (fastest)
- Limit (kbps): 0
- Begin Time: 00:00:00
- End Time: 23:59:00
- ☒ Enabled
- Setup Mode: Manual
- Remote Hostname: 10.231.1.3
- Remote Port: 22
- Encryption Cipher: standard
- ☐ Dedicated User Enabled
- Dedicated User: (empty)
- Remote Hostkey: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDG0SIC78AFFQ+QNcNSgY6TvdT3Hi6u4/3PkGh63ofv4IvOGWYslwrw3NNDICFvH3wHm91J4WKrohZw

Buttons at the bottom: SCAN SSH KEY, SAVE, CANCEL.

Fig. 7.13: Add Replication Dialog

7.6.5 Replication Options

Table 7.6 describes the options in the replication task dialog.

Table 7.6: Replication Task Options

Setting	Value	Description
Pool/Dataset	drop-down menu	On the source computer with snapshots to replicate, choose an existing pool or dataset with an active periodic snapshot task.
Remote ZFS Pool/Dataset	string	Enter the pool or dataset on the remote or destination computer that will store snapshots. Example: poolname/datasetname, not the mountpoint or filesystem path.
Recursively Replicate Child Dataset Snapshots	checkbox	Set to include snapshots of child datasets from the primary dataset.
Delete Stale Snapshots on Remote System	checkbox	Set to delete snapshots from the remote system which are also no longer present on the source computer.
Replication Stream Compression	drop-down menu	Select a compression algorithm to reduce the size of the data being replicated. Choices are <i>lz4 (fastest)</i> , <i>pigz (all rounder)</i> , <i>plzip (best compression)</i> , or <i>Off (no compression)</i> .
Limit (kbps)	integer	Limit replication speed to the specified value in kbps. Default of 0 is unlimited.

Continued on next page

Table 7.6 – continued from previous page

Setting	Value	Description
Begin Time	drop-down menu	Set the time to start the replication task.
End Time	drop-down menu	Define the time the replication must start. A started replication task continues until it is finished.
Enabled	checkbox	Unset to disable the scheduled replication task without deleting it.
Setup Mode	drop-down menu	Choose the configuration mode for the remote system. Choices are <i>Manual</i> or <i>Semi-Automatic</i> . Note <i>Semi-Automatic</i> only works with remote version 9.10.2 or later.
Remote Hostname	string	Enter the IP address or DNS name of the remote system to receive the replication data.
Remote Port	string	Enter the port used by the SSH server on the remote system.
Encryption Cipher	drop-down menu	<i>Standard</i> provides the best security. <i>Fast</i> is less secure, but has better transfer rates for devices with limited cryptographic speed. <i>Disabled</i> is for networks where the entire path between sources and destinations is trusted.
Dedicated User Enabled	checkbox	Set to allow a user account other than root to be used for replication.
Dedicated User	drop-down menu	Select the user account to use for replication. Only available if <i>Dedicated User Enabled</i> is enabled.
Remote Hostkey	string	Paste the host key of the destination NAS configured for the Replication Task. Use the <i>SCAN SSH KEY</i> button to automatically retrieve the public host key of the remote system.

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.

Navigating to *Tasks* → *Replication Tasks* displays [Figure 7.14](#), the list of replication tasks. *Status* shows the current status of each replication task. The display is updated periodically, always showing the latest status.

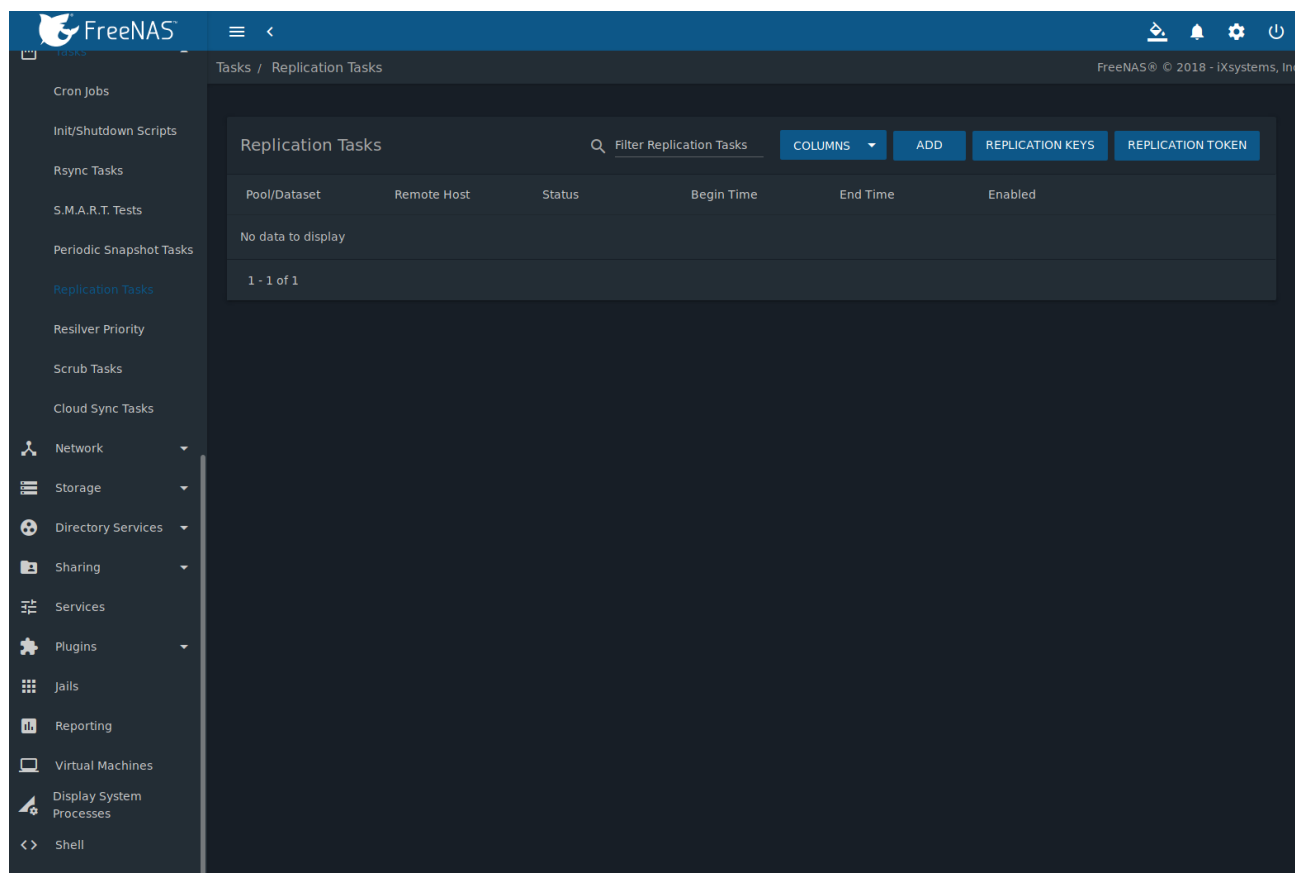


Fig. 7.14: Replication Task List

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.

7.6.6 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.

7.6.7 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 *Scrub Tasks* (page 132). 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.

7.6.8 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.

7.6.8.1 SSH

SSH (page 263) must be able to connect from the source system to the destination system with an encryption key. This is tested from *Shell* (page 325) by making an *SSH* (page 263) 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 325) 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 263) shell connection is open to the destination system, *Beta*.

If a password is requested, SSH authentication is not working. See [Figure 7.11](#) 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.

7.6.8.2 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.

7.6.8.3 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 325) 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 *alphapool/alphadata* dataset. A `@` symbol separates the name of the dataset from the name of the snapshot in the command.

```
zfs send alphapool/alphadata@auto-20161206.1110-2w | ssh -i /data/ssh/replication 10.0.0.118 zfs recv ↵
↵betapool
```

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 325) on *Beta* and running this command:

```
zfs destroy -R betapool/alphadata@auto-20161206.1110-2w
```

Then send the snapshot manually again. Snapshots on the destination system, *Beta*, are listed from the *Shell* (page 325) with `zfs list -t snapshot` or from *Storage* → *Snapshots*.

Error messages here can indicate any remaining problems.

7.7 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* menu 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 *Tasks* → *Resilver Priority* to display the screen shown in [Figure 7.15](#). [Table 7.7](#) describes the fields on this screen.

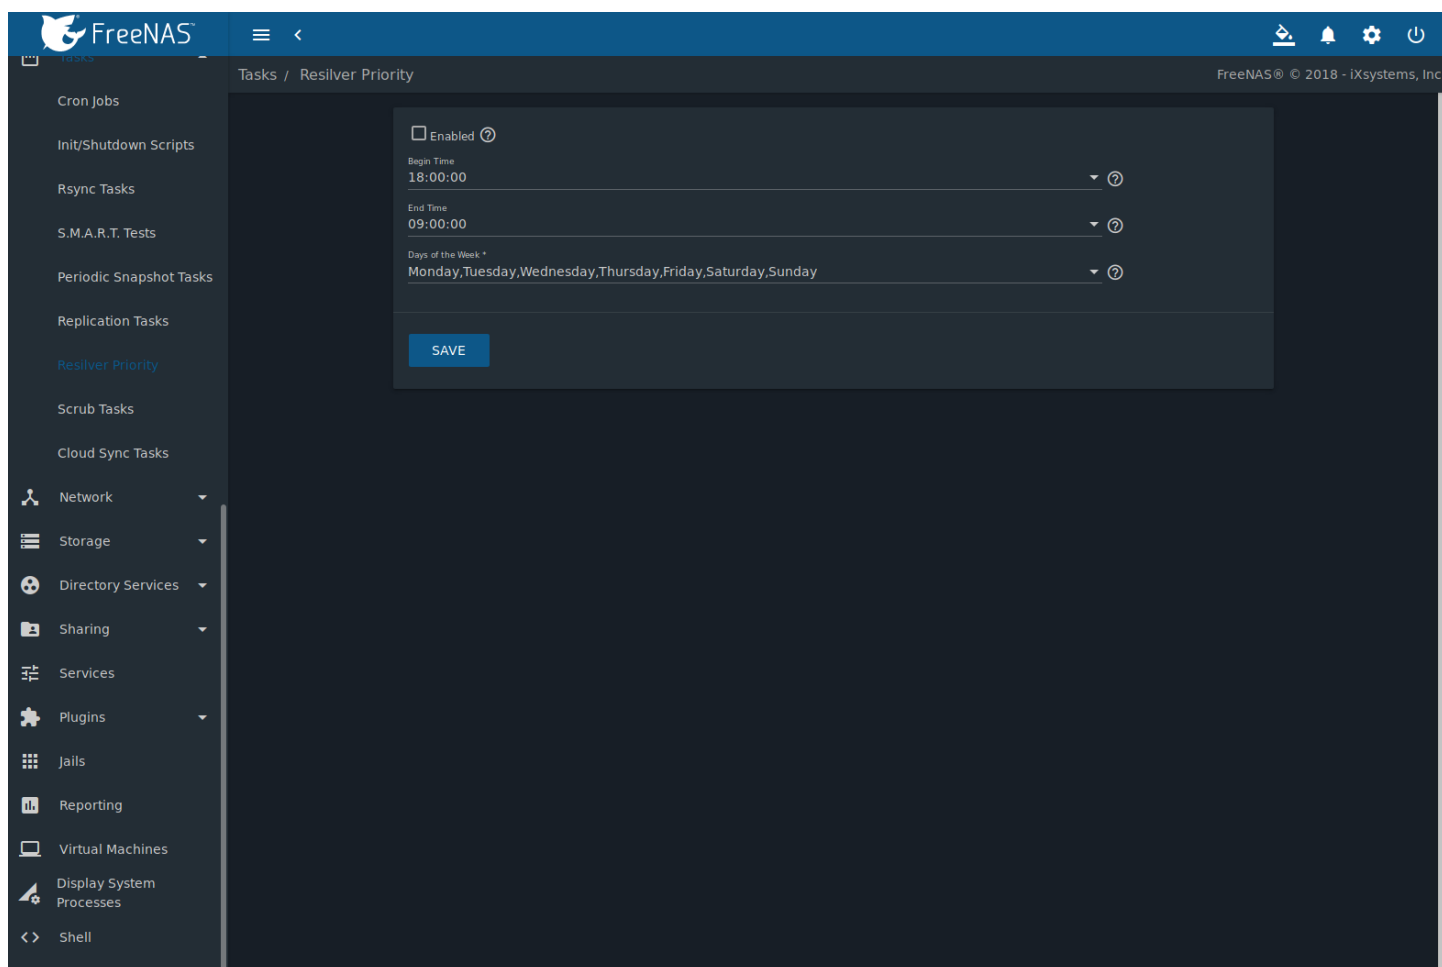


Fig. 7.15: Resilver Priority

Table 7.7: Resilver Priority Options

Setting	Value	Description
Enabled	checkbox	Set to run resilver tasks between the configured times.
Begin Time	drop-down	Choose the hour and minute when resilver tasks can be started.
End Time	drop-down	Choose the hour and minute when new resilver tasks can no longer be started. This does not affect active resilver tasks.
Days of the Week	checkboxes	Select the days to run resilver tasks.

7.8 Scrub Tasks

A scrub is the process of ZFS scanning through the data on a pool. 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.

It is recommended that each pool is 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 pool. 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 min-

imize impact to users. Make certain that scrubs and other disk-intensive activity like *S.M.A.R.T. Tests* (page 116) 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 116) of *Type Long Self-Test* to run once or twice a month.

Scrubs are scheduled and managed with *Tasks* → *Scrub Tasks*.

When a pool is created, a scrub is automatically scheduled. An entry with the same pool name is added to *Tasks* → *Scrub Tasks*. A summary of this entry can be viewed with *Tasks* → *Scrub Tasks*. [Figure 7.16](#) displays the default settings for the pool named `pool1`. In this example, *Options* and *Edit* for a pool is clicked to display the *Edit* screen. [Table 7.8](#) summarizes the options in this screen.

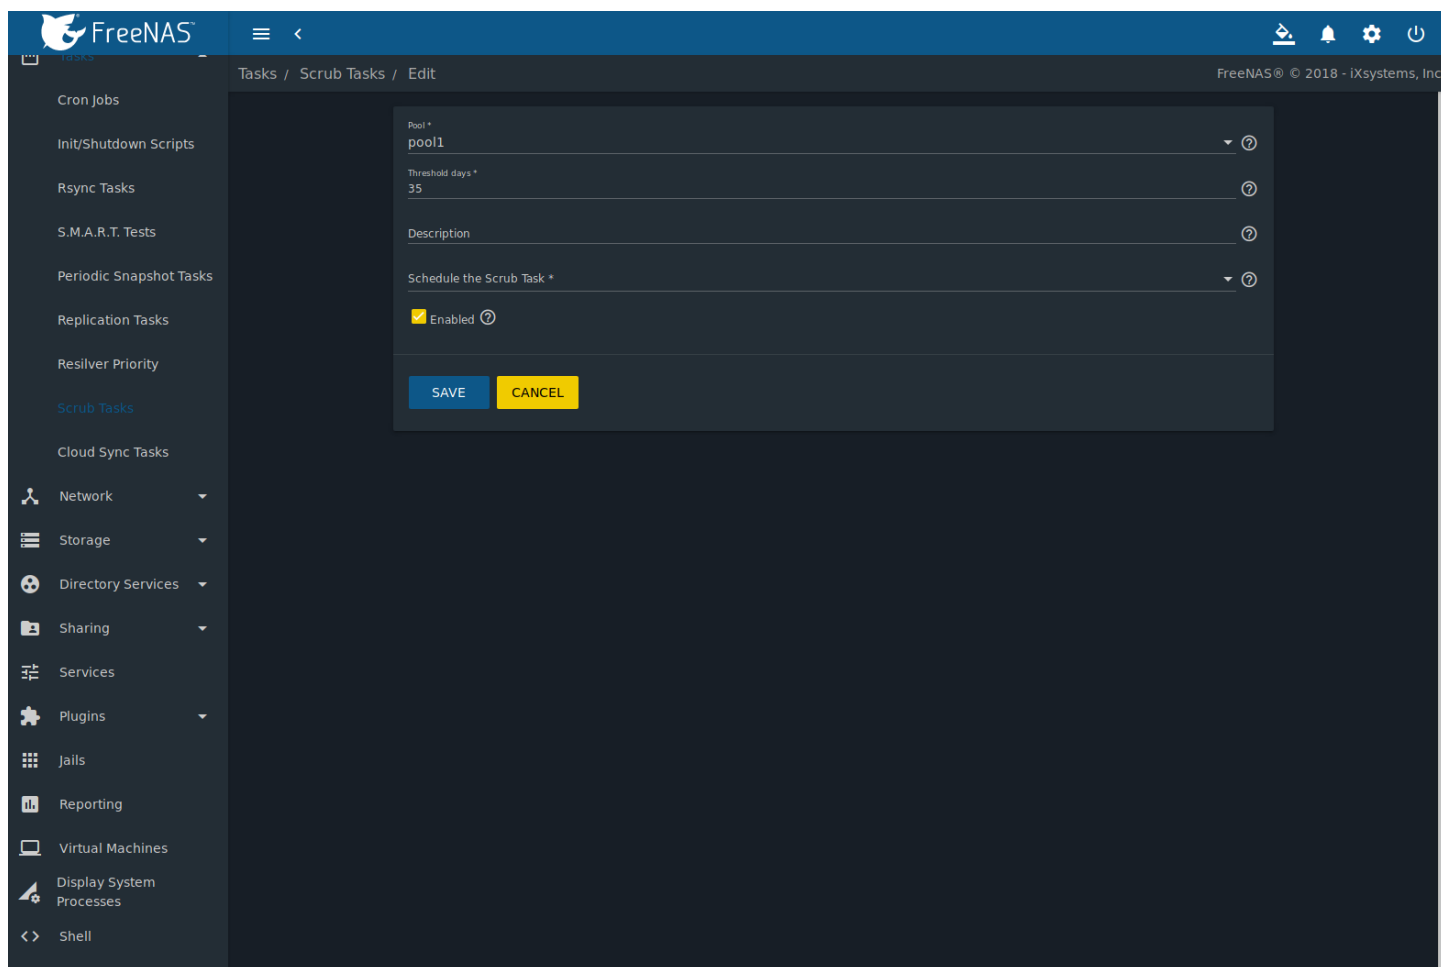


Fig. 7.16: Viewing Pool Default Scrub Settings

Table 7.8: ZFS Scrub Options

Setting	Value	Description
Pool	drop-down menu	Choose a pool to scrub.
Threshold days	string	Define the number of days to prevent a scrub from running after the last has completed. This ignores any other calendar schedule. The default is a multiple of 7 to ensure the scrub always occurs on the same weekday.
Description	string	Describe the scrub task.

Continued on next page

Table 7.8 – continued from previous page

Setting	Value	Description
Schedule the Scrub Task	drop-down menu	Choose how often to run the scrub task. Choices are <i>Hourly</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Select <i>Custom</i> to open a visual scheduler for selecting minutes, hours, days, month, and days of week.
Enabled	checkbox	Unset to disable the scheduled scrub without deleting it.

Note: Scrub tasks are run if and only if the threshold is met or exceeded *and* the task is scheduled to run on the date marked.

Review the default selections and, if necessary, modify them to meet the needs of the environment. Note that the *Threshold days* field is used to prevent scrubs from running too often, and overrides the schedule chosen in the other fields. Also, if a pool is locked or unmounted when a scrub is scheduled to occur, it will not be scrubbed.

Scheduled scrubs 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 deselecting the *Enabled* button for the scrub until the hardware can be upgraded.

7.9 Cloud Sync Tasks

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

Cloud Credentials (page 88) must be pre-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 cloud credentials have been configured, *Tasks* → *Cloud Sync Tasks* is used to define the schedule for running a cloud sync task. An example is shown in [Figure 7.17](#).

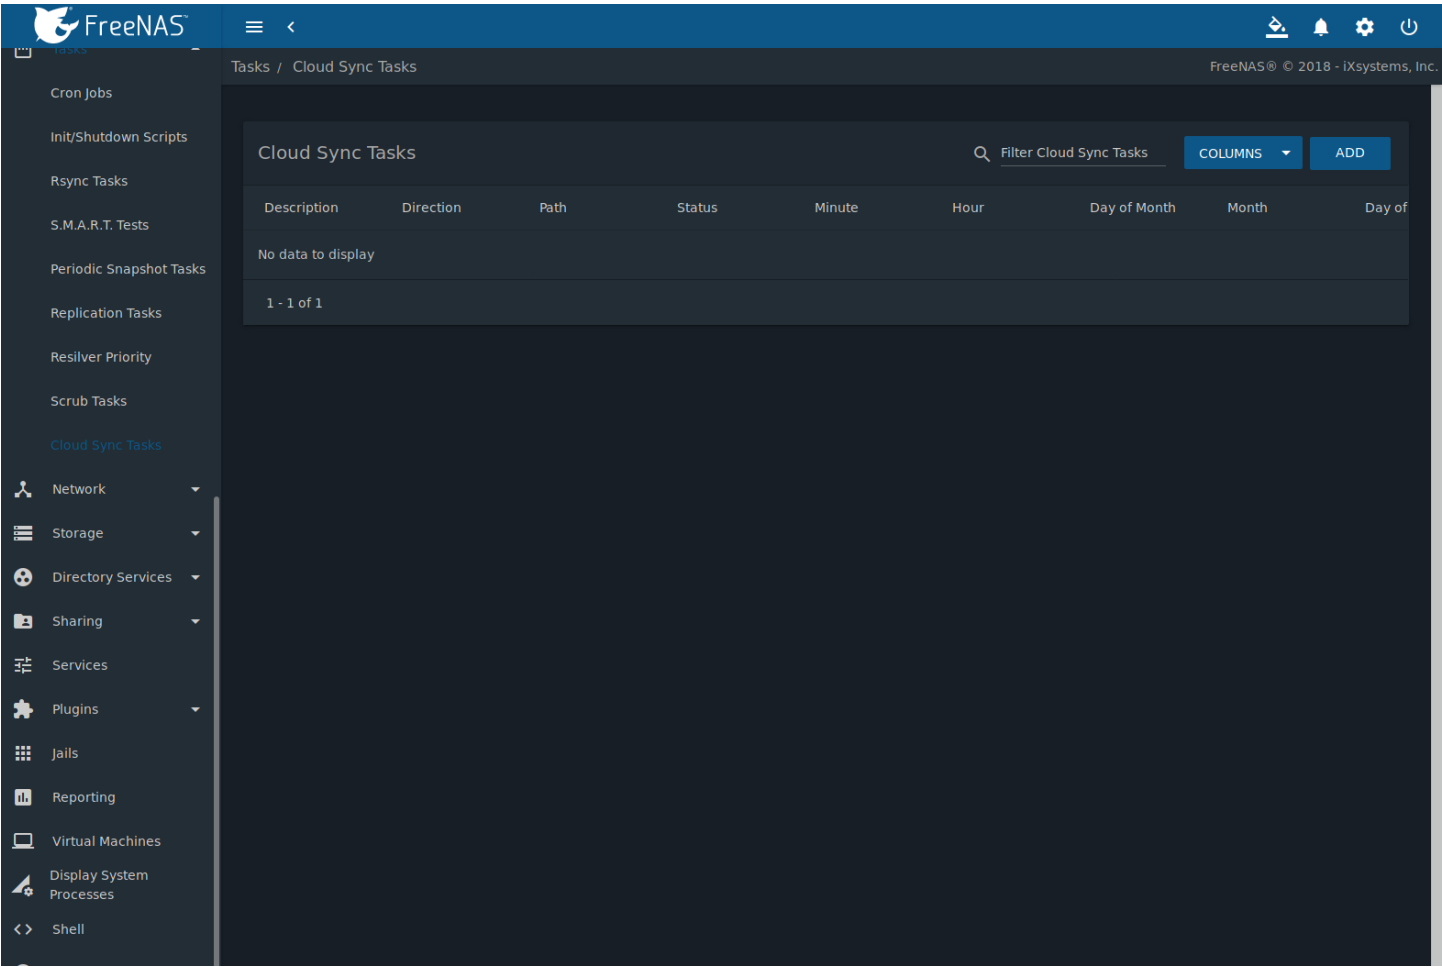


Fig. 7.17: Cloud Sync Status

When an existing task has run, a ✓ or ✕ is shown to reflect the success or failure of the task. Click either symbol to open the *Logs* window. This window displays logs related to the task that ran. Click *DOWNLOAD LOGS* to open a popup window to download the `.log` file.

Click *ADD* to display the *Add Cloud Sync* menu shown in [Figure 7.18](#).

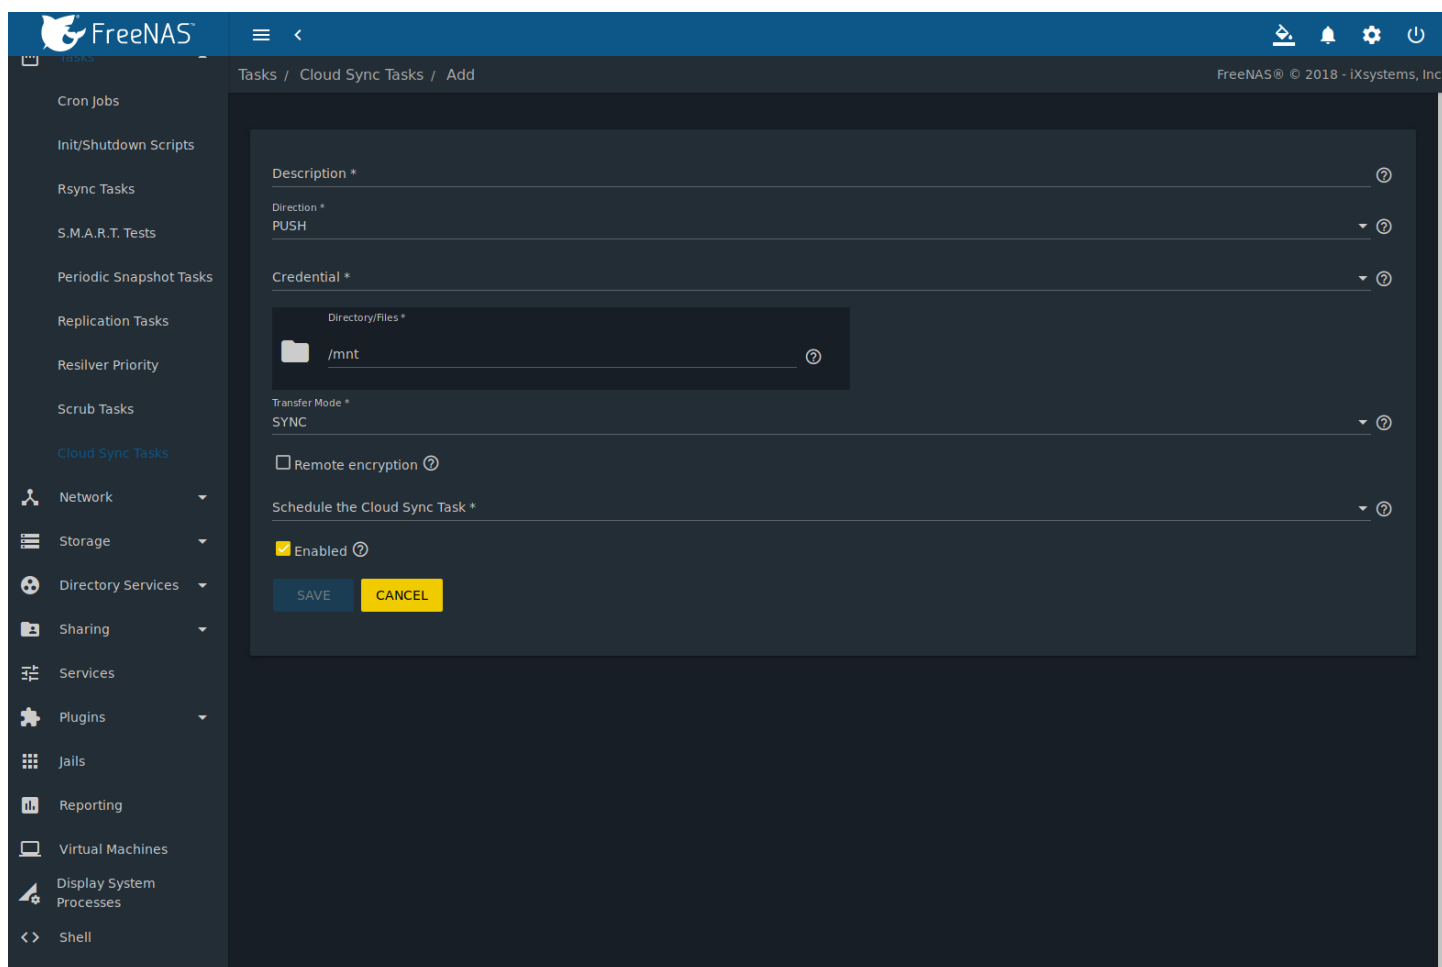


Fig. 7.18: Adding a Cloud Sync

Table 7.9 shows the configuration options for Cloud Syncs.

Table 7.9: Cloud Sync Options

Setting	Value Type	Description
Description	string	Enter a description of the Cloud Sync Task.
Direction	drop-down menu	<i>Push</i> sends data to cloud storage. <i>Pull</i> receives data from cloud storage.
Credential	drop-down menu	Select the cloud storage provider credentials from the list of available Cloud Credentials (page 88). The credential is tested and an error is displayed if a connection cannot be made. <i>SAVE</i> is disabled until a valid credential is entered.
Bucket/Container	drop-down menu	<i>Bucket</i> : Only appears when an S3 credential is the <i>Provider</i> . Select the pre-defined S3 bucket to use. <i>Container</i> : Only appears when a <code>AZUREBLOB</code> credential is selected for the <i>Credential</i> . Enter the name of the pre-configured Microsoft Azure Blob container.
Folder	string	The name of the pre-defined folder within the selected bucket or container.
Encryption	drop-down menu	Only appears when an S3 credential is the <i>Provider</i> . Choices are <i>None</i> (no encryption) or <i>AES-256</i> (encrypted).
Directory/Files	browse button	Select the directories or files to be sent to the cloud for <i>Push</i> syncs, or the destination to be written for <i>Pull</i> syncs. Be cautious about the destination of <i>Pull</i> jobs to avoid overwriting existing files.

Continued on next page

Table 7.9 – continued from previous page

Setting	Value Type	Description
Transfer Mode	drop-down menu	<i>Sync</i> makes files on the destination system identical to those on the source. Files that are removed from the source are also removed from the destination, similar to <code>rsync --delete</code> . <i>Copy</i> copies files from the source to the destination, skipping files that are identical, similar to <code>rsync</code> . <i>Move</i> copies files from the source to the destination, deleting files from the source after the copy, similar to <code>mv</code> .
Remote encryption	checkbox	Set to encrypt files before transfer and store the encrypted files on the remote system. rclone Crypt (https://rclone.org/crypt/) is used.
Filename encryption	checkbox	Only appears when <i>Remote encryption</i> is enabled. Set to encrypt the shared file names.
Encryption password	string	Only appears when <i>Remote encryption</i> is enabled. Enter the password to encrypt and decrypt remote data. <i>Warning:</i> Always save and back up this password. Losing the encryption password can result in data loss.
Encryption salt	string	Only appears when <i>Remote encryption</i> is enabled. Enter a long string of random characters for use as salt (https://searchsecurity.techtarget.com/definition/salt) for the encryption password. <i>Warning:</i> Save and back up the encryption salt value. Losing the salt value can result in data loss.
Schedule the Cloud Sync Task	drop-down menu	Choose how often to run the task. Choices are <i>Hourly</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Select <i>Custom</i> to open the advanced scheduler.
Enabled	checkbox	Enable this Cloud Sync Task. Unset to disable this Cloud Sync Task without deleting it.

Note: If the selected credential is incorrect it prompts for a correction. Click the *Fix Credential* button to return to the *System* → *Cloud Credentials* → *Edit* page for the selected credential.

To modify an existing cloud sync, click  (Options) to access the *Run Now*, *Edit*, and *Delete* options.

7.9.1 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.

Click *System* → *Cloud Credentials* and *ADD* to enter the credentials for storage on an Amazon AWS account. The credential is given the name *S3 Storage*, as shown in [Figure 7.19](#):

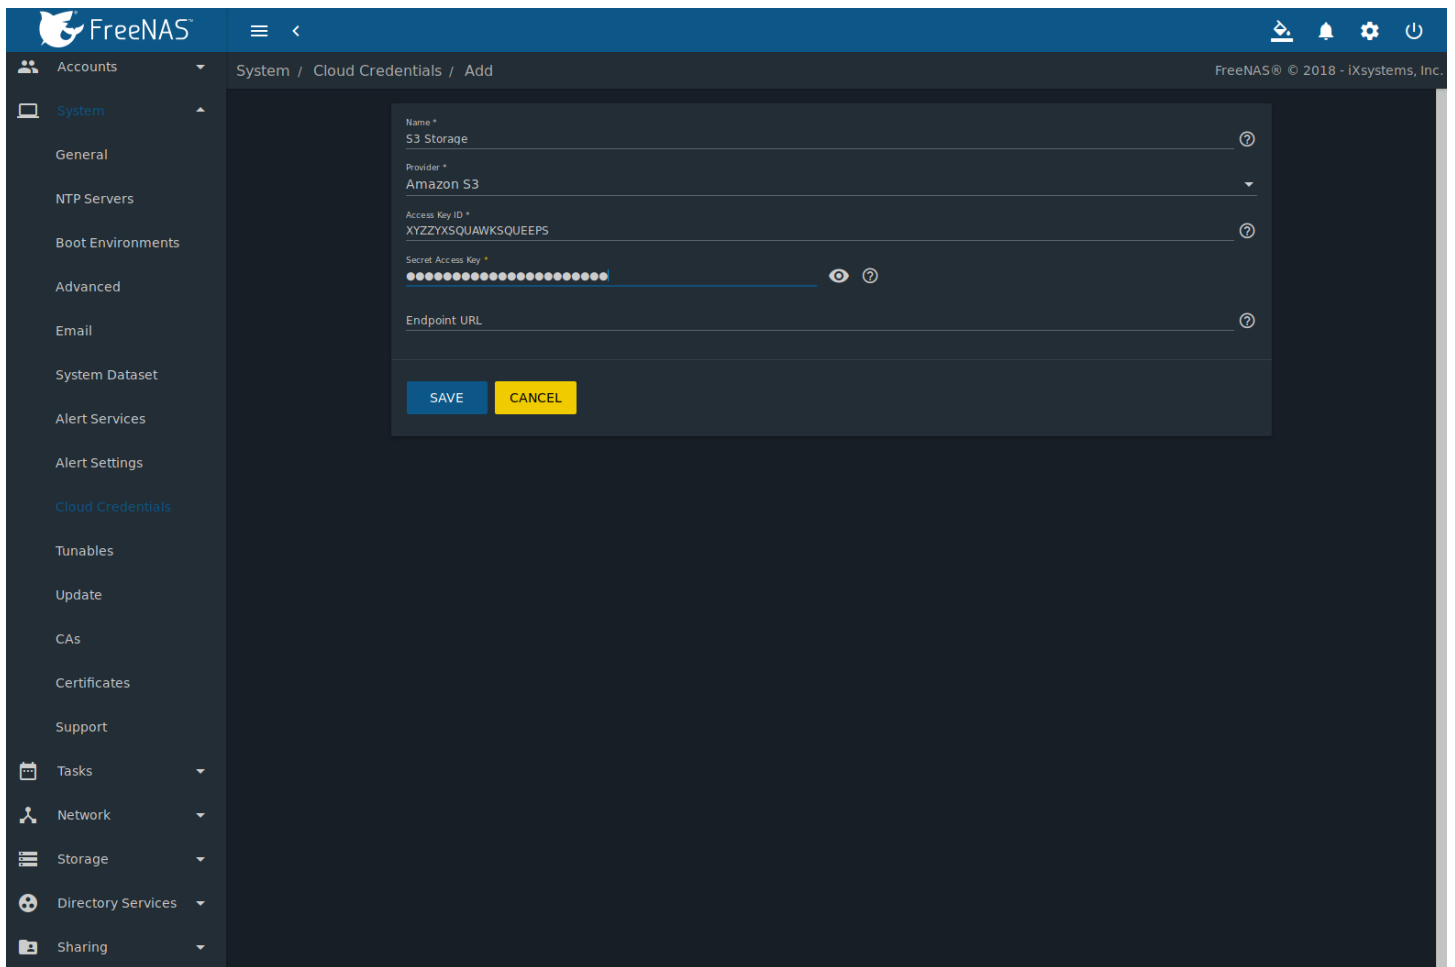


Fig. 7.19: Example: Adding Cloud Credentials

The local data to be sent to the cloud is a single file called `accounting-backup.bin` on the `smb-storage` dataset.

Click *Tasks* → *Cloud Sync* and *ADD* to create a cloud sync job. 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 *Directory/Files* is adjusted to the data file.

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 enabled by default, so this cloud sync will run at the next scheduled time.

The completed dialog is shown in [Figure 7.20](#):

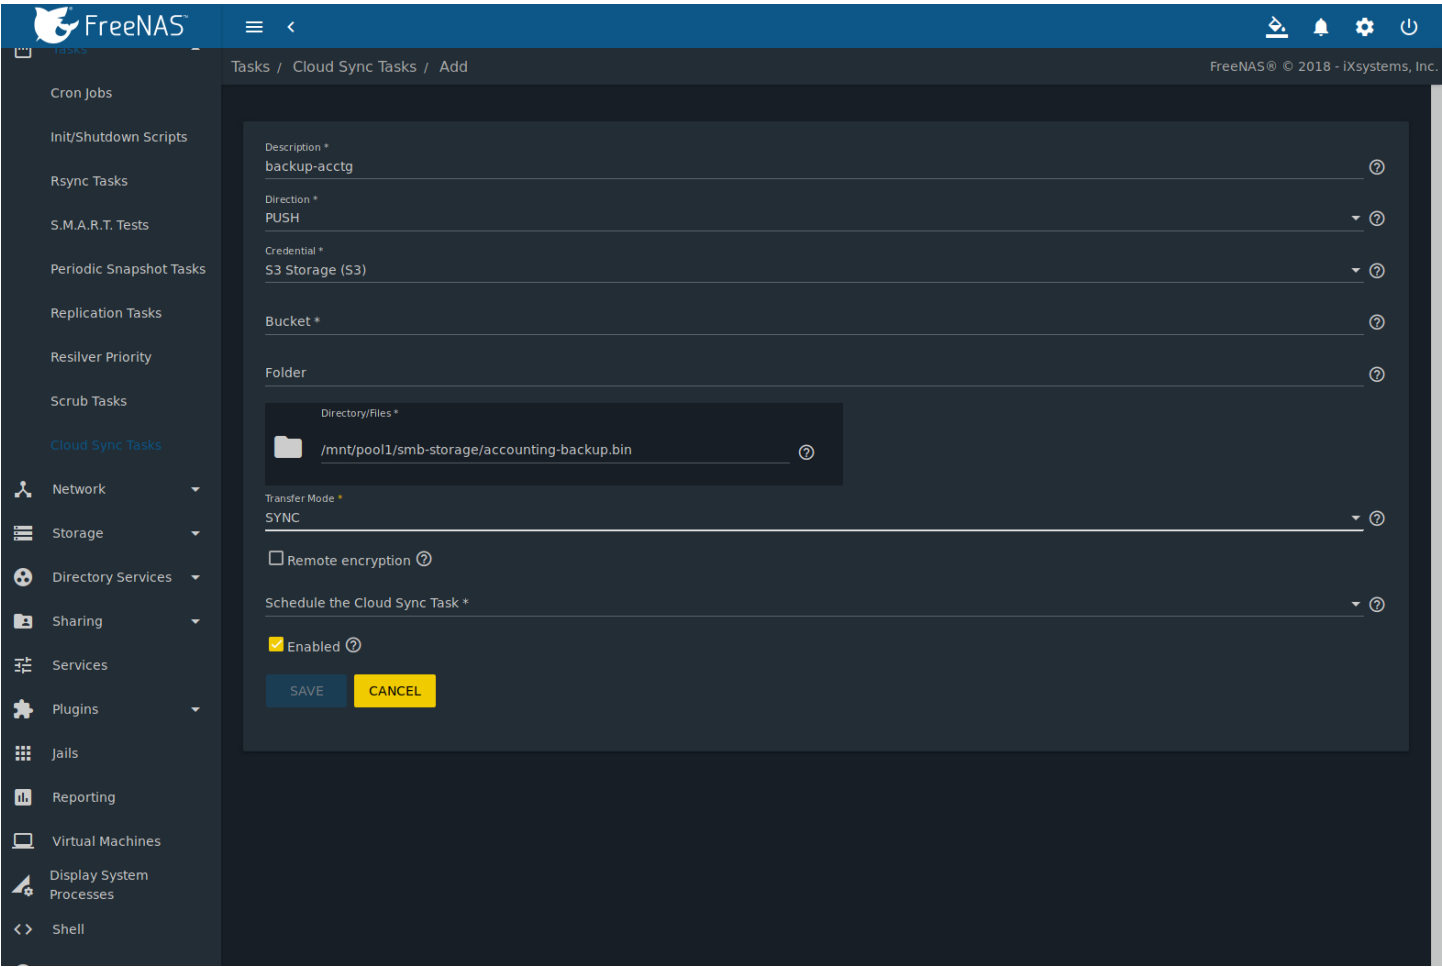


Fig. 7.20: Example: Adding a Cloud Sync

NETWORK

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

- [Global Configuration](#) (page 140): general network settings.
- [Interfaces](#) (page 142): settings for each network interface.
- [IPMI](#) (page 144): settings controlling connection to the appliance through the hardware side-band management interface if the user interface becomes unavailable.
- [Link Aggregations](#) (page 146): settings for network link aggregation and link failover.
- [Static Routes](#) (page 151): add static routes.
- [VLANs](#) (page 152): configure IEEE 802.1q tagging for virtual LANs.

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

Warning: Making changes to the network interface the web interface uses can result in losing connection to the FreeNAS® system! Misconfiguring network settings might require command line knowledge or physical access to the FreeNAS® system to fix. Be very careful when configuring [Interfaces](#) (page 142) and [Link Aggregations](#) (page 146).

8.1 Global Configuration

Network → *Global Configuration*, shown in [Figure 8.1](#), is for general network settings that are not unique to any particular network interface.

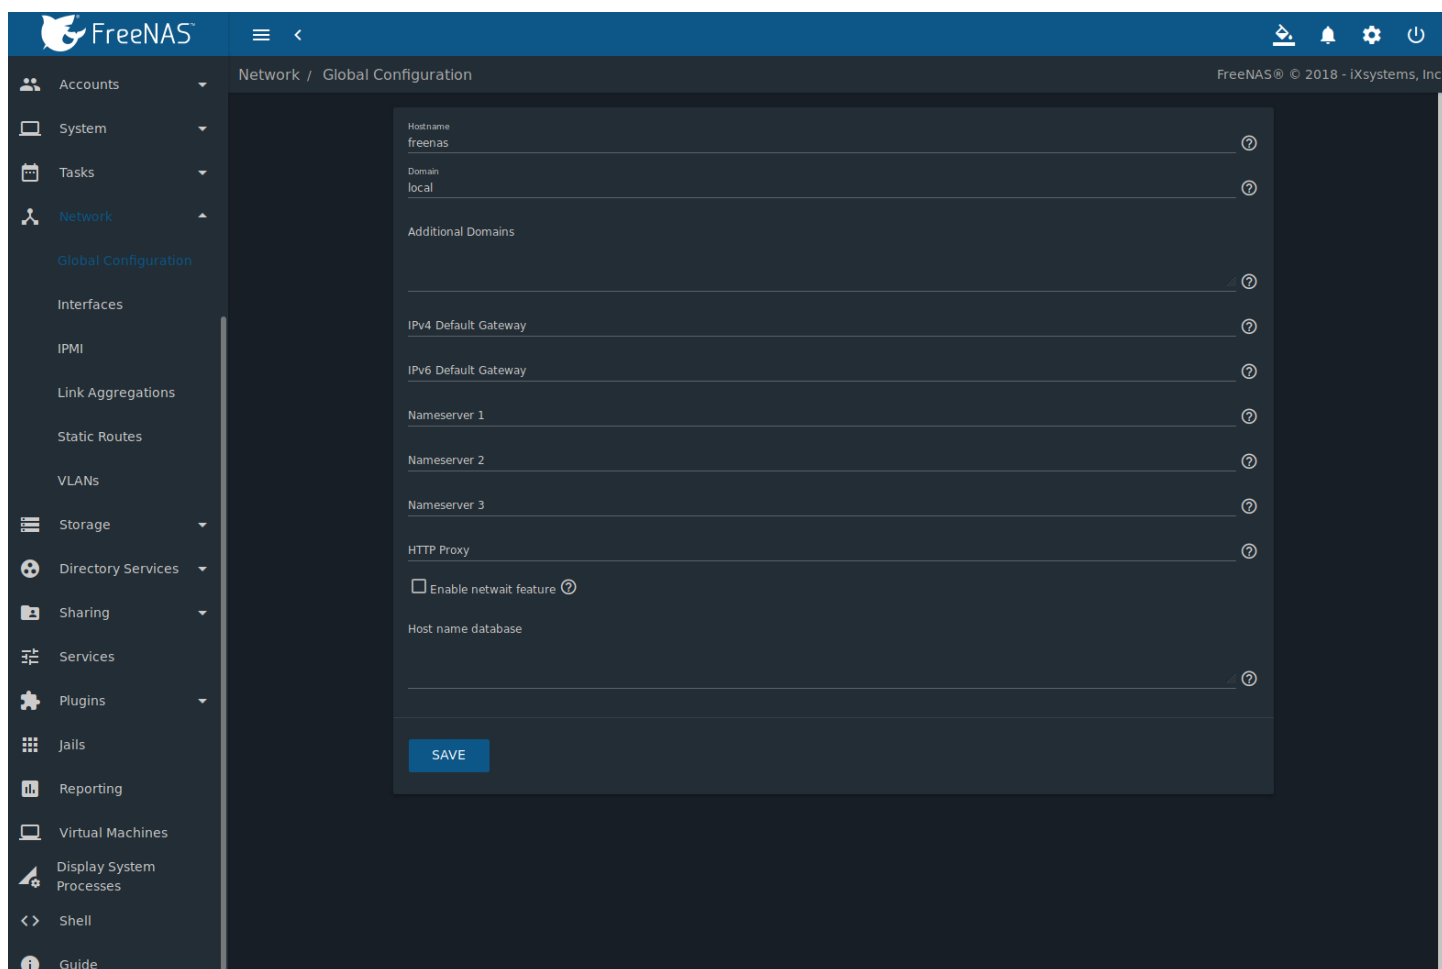


Fig. 8.1: Global Network Configuration

Table 8.1 summarizes the settings on the Global Configuration tab. *Hostname* and *Domain* fields are pre-filled as shown in Figure 8.1, but can be changed to meet requirements of the local network.

Table 8.1: Global Configuration Settings

Setting	Value	Description
Hostname	string	System host name. Cannot contain the underscore character.
Domain	string	System domain name.
Additional Do- mains	string	Additional space-delimited domains to search. Adding search domains can cause slow DNS lookups.
IPv4 Default Gateway	IP address	Typically not set. See this note about Gateways (page 142). If set, used instead of the default gateway provided by DHCP.
IPv6 Default Gateway	IP address	Typically not set. See this note about Gateways (page 142).
Nameserver 1	IP address	Primary DNS server.
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 do not start at boot until the interface is able to ping the addresses listed in the <i>Netwait IP list</i> .

Continued on next page

Table 8.1 – continued from previous page

Setting	Value	Description
Netwait IP list	string	Only appears when <i>Enable netwait feature</i> is set. Enter a space-delimited list of IP addresses to ping(8). Each address is tried until one is successful or the list is exhausted. Leave empty to use 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 <i>IP_address space hostname</i> where multiple hostnames can be used if separated by a space.

When using Active Directory, set the IP address of the realm 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 151) 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 the FreeNAS® system is protected by a properly configured firewall.

8.2 Interfaces

Network → *Interfaces* shows which interfaces are manually configured and allows adding or editing a manually configured interface.

See this [warning](#) (page 140) about changing the interface that the web interface uses.

[Figure 8.2](#) shows the screen that appears after clicking *ADD* from the *Interfaces* page. [Table 8.2](#) summarizes the configuration options shown when adding an interface or editing an existing interface.

Note: An interface can only be added when there is a NIC that has not already been configured. Clicking *ADD* when there are no NICs available will display a message across the bottom of the screen that `All interfaces are already in use..`

The screenshot shows the FreeNAS web interface with the 'Network / Interfaces / Add' page. The left sidebar contains navigation links for Accounts, System, Tasks, Network, Global Configuration, Interfaces, IPMI, Link Aggregations, Static Routes, VLANs, Storage, Directory Services, Sharing, Services, Plugins, Jails, Reporting, Virtual Machines, Display System Processes, Shell, and Guide. The main content area displays a form for adding a new interface. The form fields are: NIC (igb1), Interface Name (igb1), DHCP (unchecked), IPv4 Address, IPv4 Netmask, Auto configure IPv6 (unchecked), IPv6 Address, IPv6 Prefix Length, and an Options text area. At the bottom, there are buttons for SAVE, CANCEL, ADD ADDITIONAL IPV4 ALIAS, and ADD ADDITIONAL IPV6 ALIAS.

Fig. 8.2: Adding or Editing an Interface

Table 8.2: Interface Configuration Settings

Setting	Value	Description
NIC	drop-down menu	The FreeBSD device name of the interface. This is read-only when editing an interface.
Interface Name	string	Description of interface.
DHCP	checkbox	Requires static IPv4 or IPv6 configuration if unselected. Only one interface can be configured for DHCP.
IPv4 Address	IP address	Enter a static IP address if <i>DHCP</i> is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if <i>DHCP</i> is unset.
Auto configure IPv6	checkbox	Only one interface can be configured for this option. If unset, manual configuration is required to use IPv6.
IPv6 Address	IPv6 address	Must be unique on the network.
IPv6 Prefix Length	drop-down menu	Match the prefix used on the network.
Options	string	Additional parameters from ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig). Separate multiple parameters with a space. For example: <i>mtu 9000</i> increases the MTU for interfaces which support jumbo frames. See this note (page 151) about MTU and <i>lagg</i> interfaces.

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 informa-

tion. Check the subnet mask if an error is shown when setting the IP addresses on multiple interfaces. Set only the IPv4 **or** IPv6 address for the new interface.

8.3 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 is also used to give 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](https://www.ixsystems.com/blog/how-to-fix-the-ipmi-remote-management-vulnerability/) (<https://www.ixsystems.com/blog/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-asrocks-ipmi-virtual-console.53911/>) for more information.

IPMI is configured from *Network* → *IPMI*. The IPMI configuration screen, shown in [Figure 8.3](#), provides a shortcut to the most basic IPMI configuration. Those already familiar with IPMI management tools can use them instead. [Table 8.3](#) summarizes the options available when configuring IPMI with the FreeNAS® web interface.

Fig. 8.3: IPMI Configuration

Table 8.3: 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. The maximum length is 20 characters.
DHCP	checkbox	If left unset, <i>IPv4 Address</i> , <i>IPv4 Netmask</i> , and <i>IPv4 Default Gateway</i> must be set.
IPv4 Address	string	IP address used to connect to the IPMI web interface.
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.

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 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-) (<https://forums.freenas.org/index.php?resources/how-to-change-ipmi-sensor-thresholds->

using `ipmitool.35/`) for some examples.

8.4 Link Aggregations

FreeNAS® uses the FreeBSD `lagg(4)` (<https://www.freebsd.org/cgi/man.cgi?query=lagg>) interface to provide link aggregation and link failover support. A `lagg` interface allows combining multiple network interfaces into a single virtual interface. This provides fault-tolerance and high-speed multi-link throughput. The aggregation protocols supported by `lagg` both determine the ports to use for outgoing traffic and if 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 later 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 going to *System* → *Tunables* and clicking *ADD* to add a tunable. Set the *Variable* to `net.link.lagg.failover_rx_all`, the *Value* to a non-zero integer, and the *Type* to *Sysctl*.

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 situations there will be a single LAG which contains all ports. In the event of changes in physical connectivity, link aggregation quickly converges to a new configuration. LACP must be configured on the network 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.

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

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.

8.4.1 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.

8.4.2 Creating a Link Aggregation

Before creating a link aggregation, make sure that all interfaces to use in the lagg are not manually configured in *Network* → *Interfaces*. **Lagg creation fails if any of the included interfaces are manually configured.** See this [warning](#) (page 140) about changing the interface that the web interface uses.

To create a link aggregation, go to *Network* → *Link Aggregations* and click *ADD*. [Figure 8.4](#) shows the configuration options.

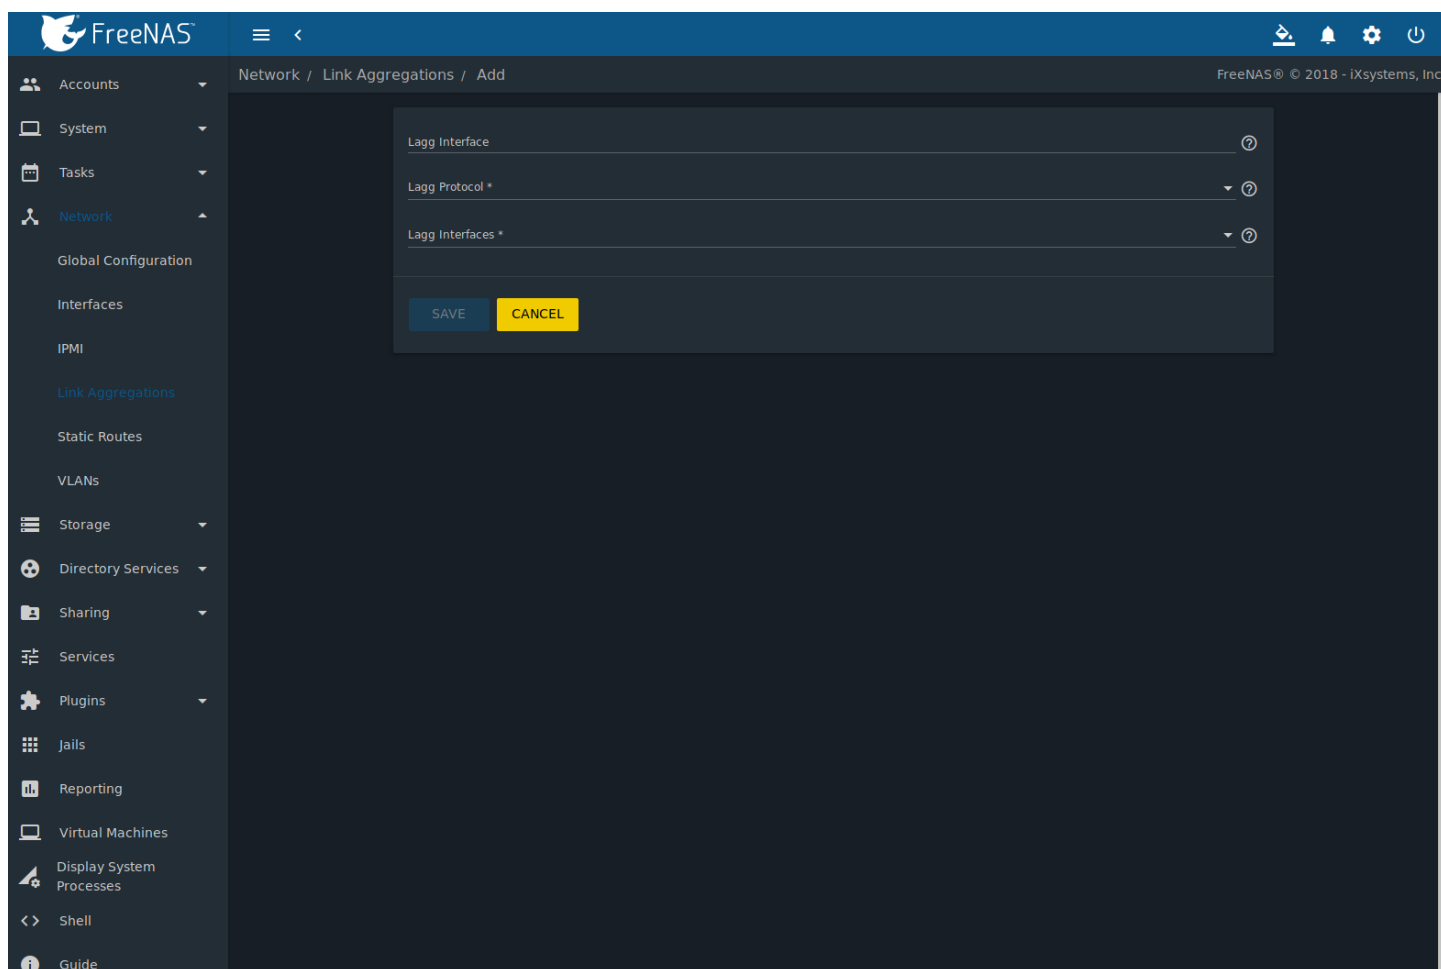


Fig. 8.4: Creating a Link Aggregation

Enter a descriptive name for the *Lagg Interface*. Next, select the desired *Lagg Protocol*. *LACP* is preferred. Choose *Failover*

when the network switch does not support LACP. Choose interfaces from the *Lagg Interfaces* drop-down menu to associate NICs with the lagg device and then click the **SAVE** button to save the new aggregation.

Note: If interfaces are installed but do not appear in the *Lagg Interfaces* list, check for a [FreeBSD driver](https://www.freebsd.org/releases/11.2R/hardware.html#ethernet) (https://www.freebsd.org/releases/11.2R/hardware.html#ethernet) for the interface.

After creating the link aggregation, go to *Network* → *Link Aggregations* and click ⓘ (Options) for the new lagg to view options to *Edit Interface*, *Edit Members*, and *Delete*.

Clicking *Edit Interface* for a lagg opens the configuration screen shown in Figure 8.5. Table 8.4 describes the options in this screen.

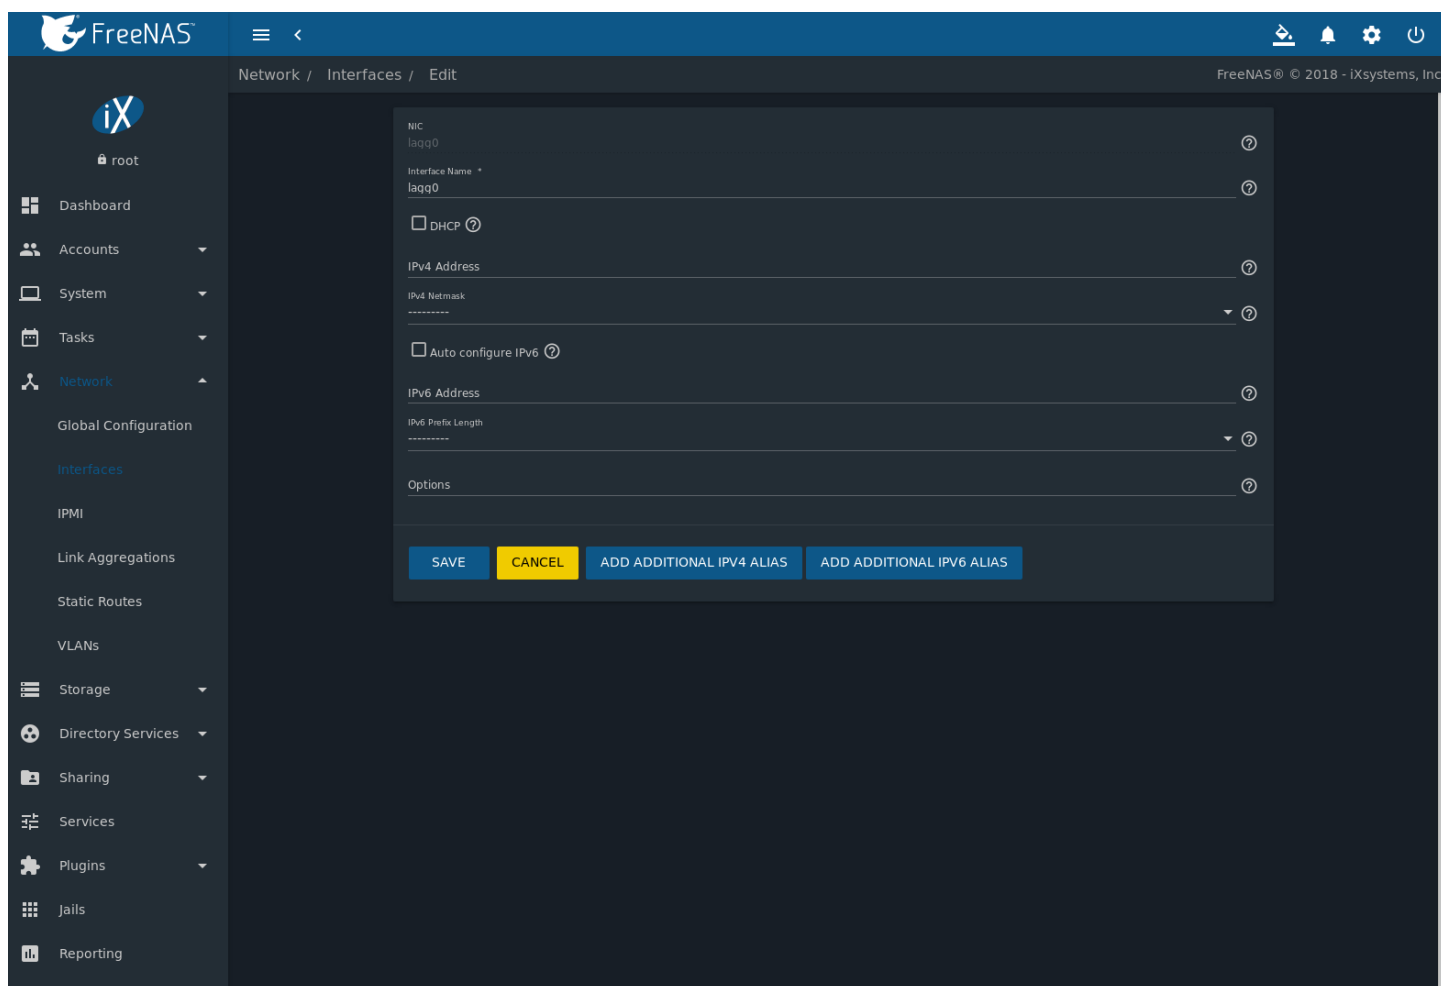


Fig. 8.5: Editing a lagg

Table 8.4: Configurable Options for a lagg


Setting	Value	Description
NIC	string	Read-only. Automatically assigned the next available numeric ID.
Interface Name	string	By default, this is the same as <i>NIC</i> . This can be changed to a more descriptive value.
DHCP	checkbox	Enable if the lagg device will get IP address info from DHCP server. The IP address of the new lagg can be set to DHCP only if no other interface uses DHCP.

Continued on next page

Table 8.4 – continued from previous page

Setting	Value	Description
IPv4 Address	string	Enter a static IP address if <i>DHCP</i> is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if <i>DHCP</i> is left unset.
Auto configure IPv6	checkbox	Set only if a DHCP server is available to provide IPv6 address information.
IPv6 Address	string	Optional.
IPv6 Prefix Length	drop-down menu	Required if an IPv6 address is entered.
Options	string	Additional ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig) options.

There are also buttons to add and remove extra IPv4 or IPv6 aliases.

In *Network* → *Link Aggregations*, click  (Options) and *Edit Members* for a lagg to see the *Members* screen, shown in [Figure 8.6](#).

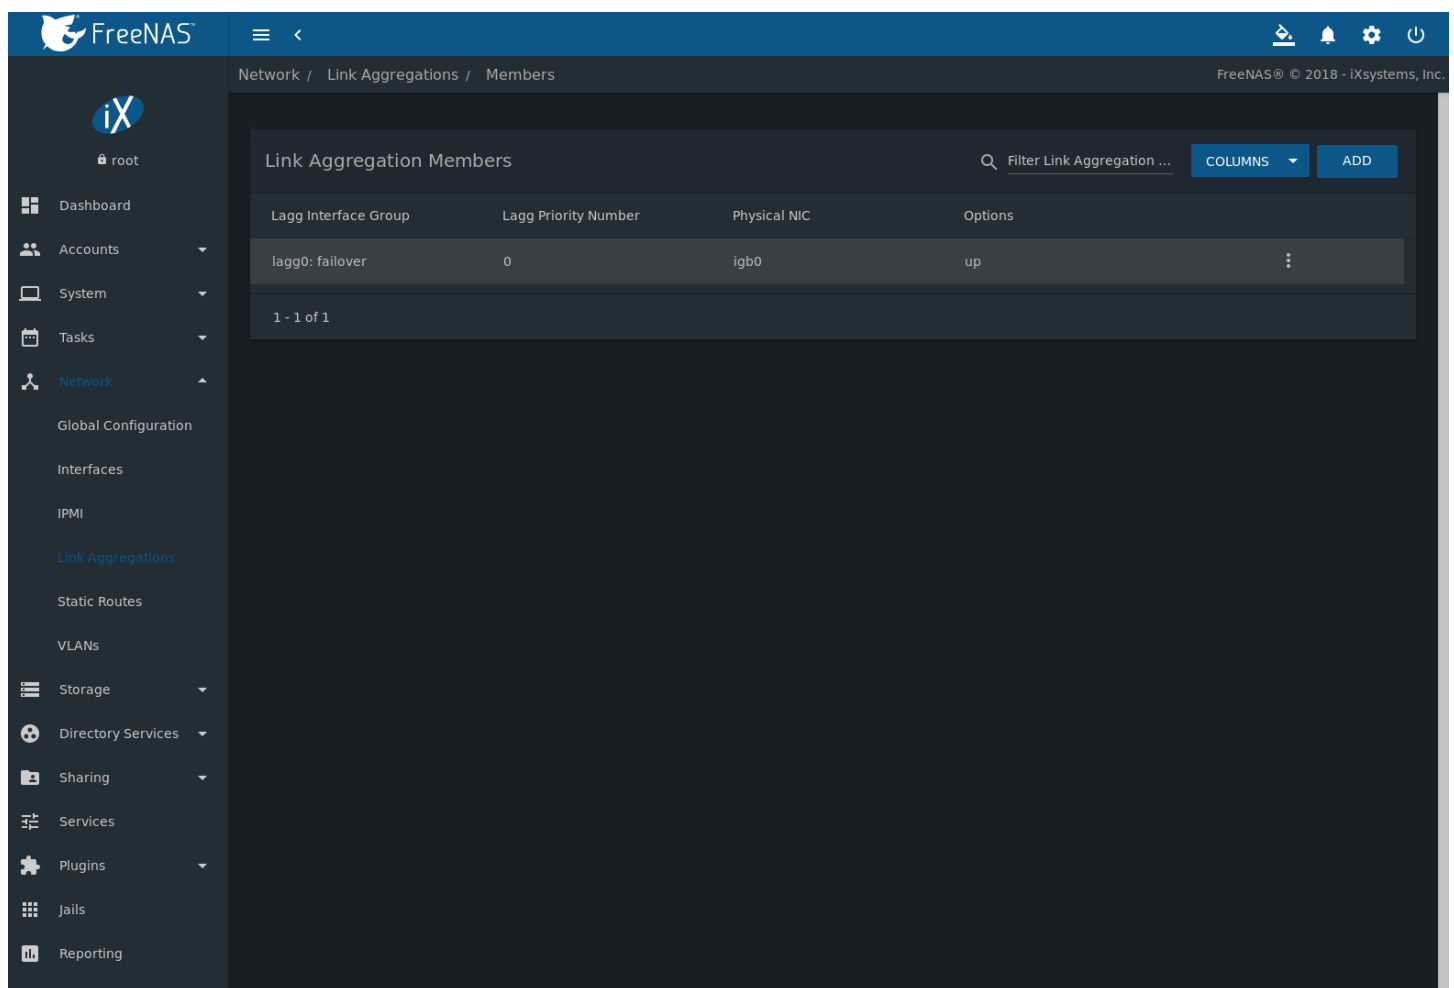


Fig. 8.6: Link Aggregation Members


Click  (Options) for an existing lagg member to see options to *Edit* and *Delete* it. Choose *Edit* to adjust an existing member. The configurable options are summarized in [Table 8.5](#).

Table 8.5: 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. This field only appears when a NIC is available.
Options	string	Additional parameters from <code>ifconfig(8)</code> (https://www.freebsd.org/cgi/man.cgi?query=ifconfig).

Click **ADD** to open the screen shown in [Figure 8.7](#).

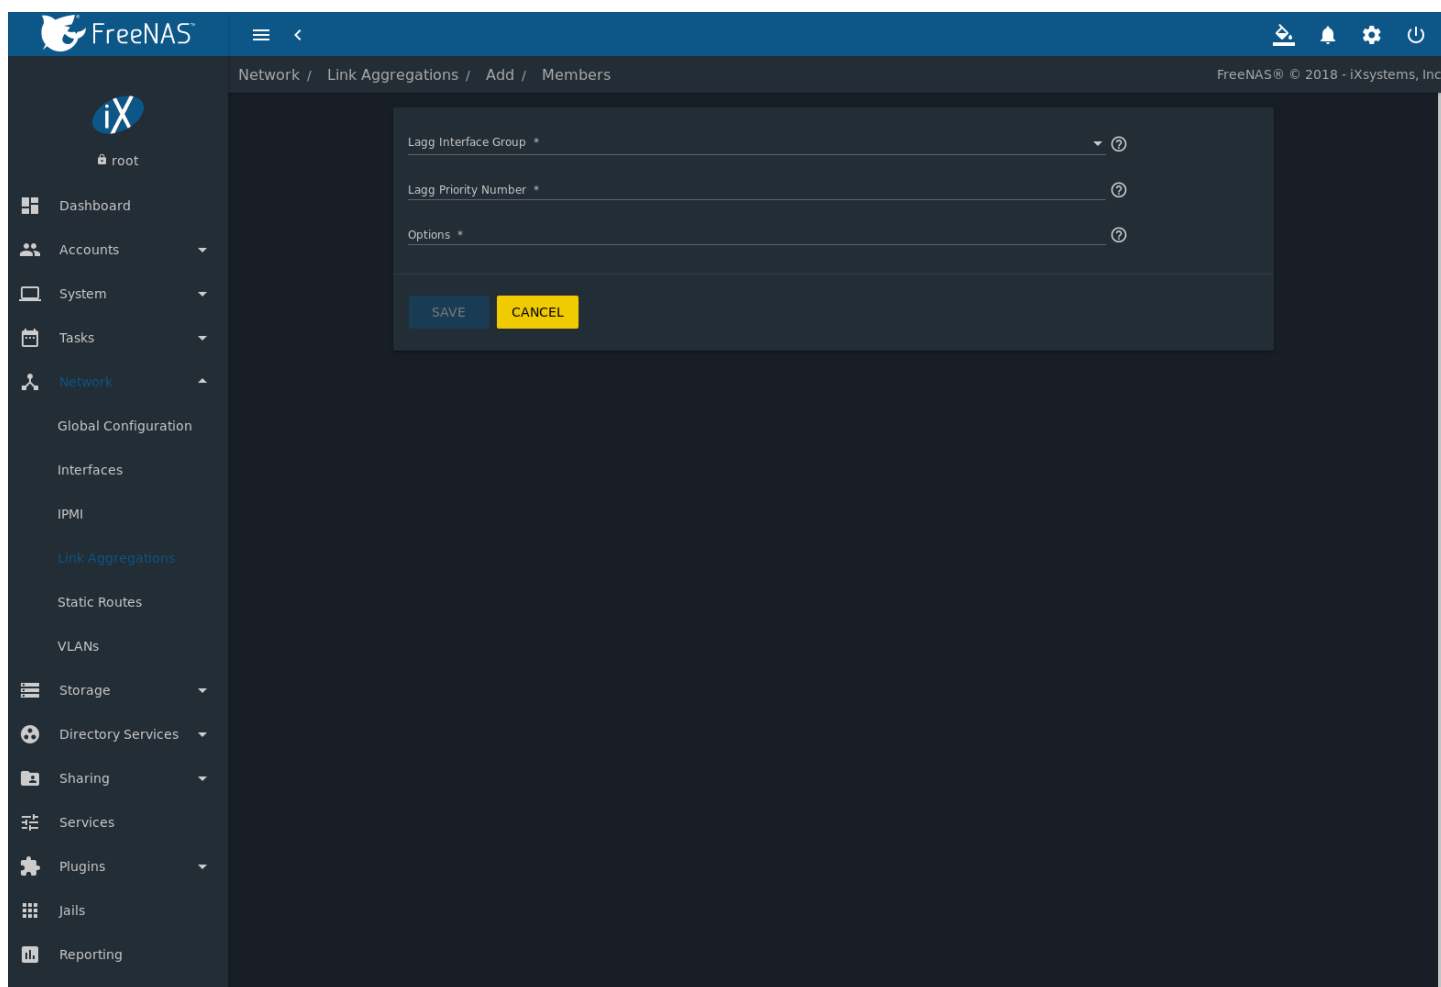


Fig. 8.7: Add Link Aggregation Member

The options are identical to the [Configuring a Member Interface](#) (page 150) table. Click **SAVE** to add the member to the list in *Network* → *Link Aggregations* → *Members*.

8.4.3 Link Aggregation Options

Options are set at the lagg level from the *Network* → *Link Aggregations* page. Click ⓘ (Options) and *Edit Members* for an existing lagg interface. Click ⓘ (Options) and *Edit* for the existing member. Scroll to the *Options* field.

To set options at the individual parent interface level, go to *Network* → *Interfaces*, and click ⓘ (Options) on the desired interface. Select *Edit*, and scroll to the *Options* field. Changes are typically made at the *lagg* level as each interface member inherits settings from the *lagg*. Configuring at the interface level requires repeating the configuration 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*, go to *Network* → *Interfaces*, click ⓘ (Options), and then *Edit* 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*.

Link aggregation load balancing can be tested with:

```
systat -ifstat
```

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

8.5 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 by going to *Network* → *Static Routes*, and clicking *ADD*. This is shown in [Figure 8.8](#).

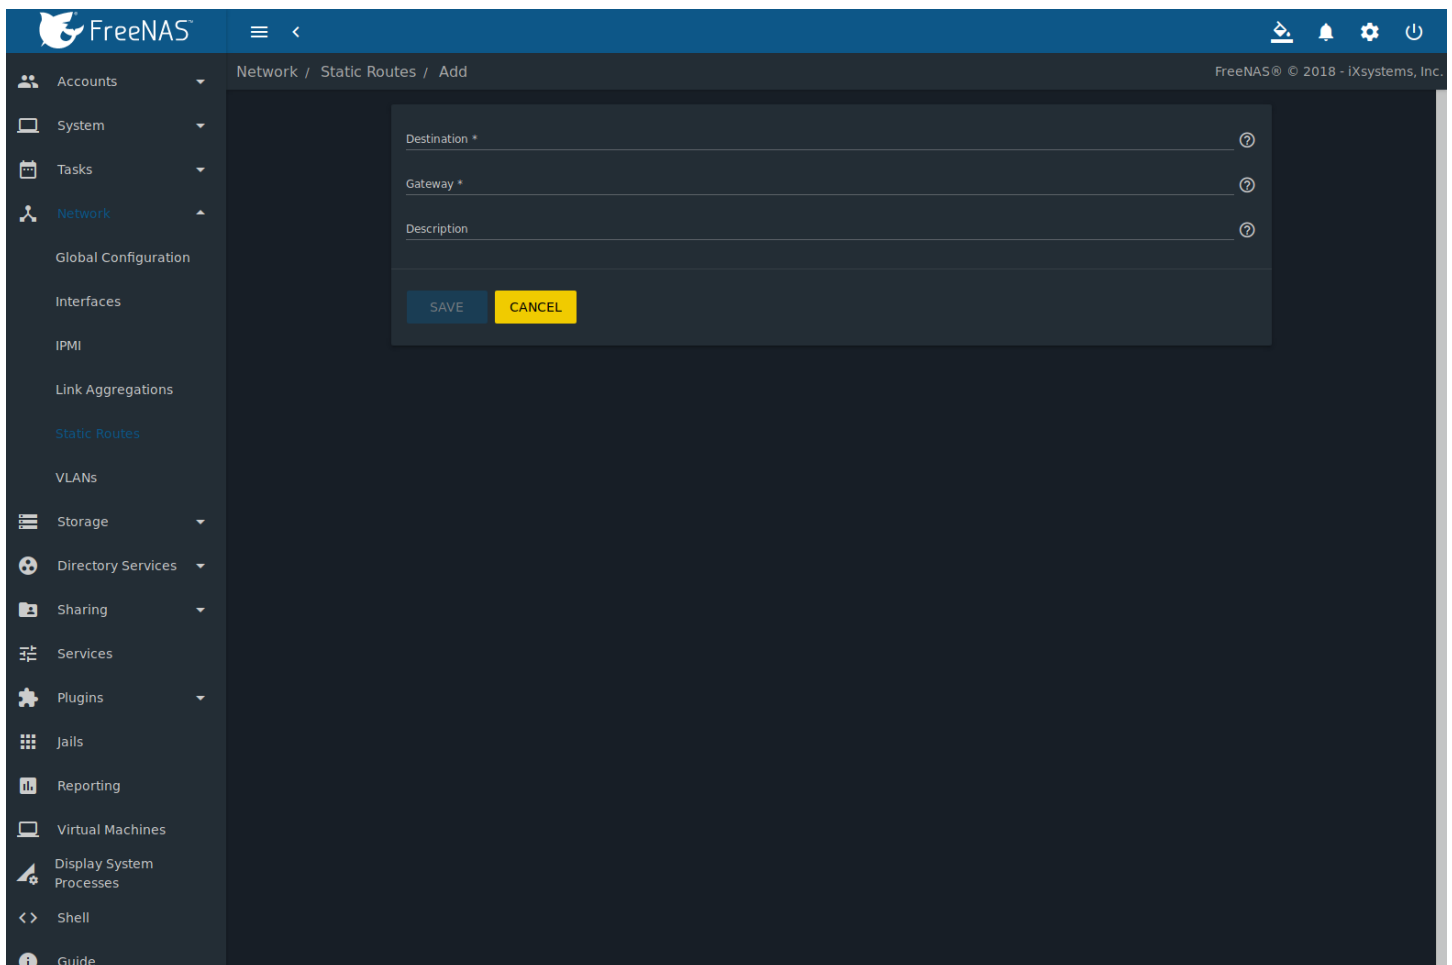


Fig. 8.8: Adding a Static Route

The available options are summarized in [Table 8.6](#).

Table 8.6: Static Route Options

Setting	Value	Description
Destination	integer	Use the format <i>A.B.C.D/E</i> where <i>E</i> is the CIDR mask.
Gateway	integer	Enter the IP address of the gateway.
Description	string	Optional. Add any notes about the route.

Added static routes are shown in *Network* → *Static Routes*. Click ⓘ (Options) on a route entry to access the *Edit* and *Delete* buttons.

8.6 VLANs

FreeNAS® uses FreeBSD's `vlan(4)` (<https://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)` (<https://www.freebsd.org/cgi/man.cgi?query=vlan>) for details.

Go to *Network* → *VLANs* and click *ADD* to see the screen shown in Figure 8.9.

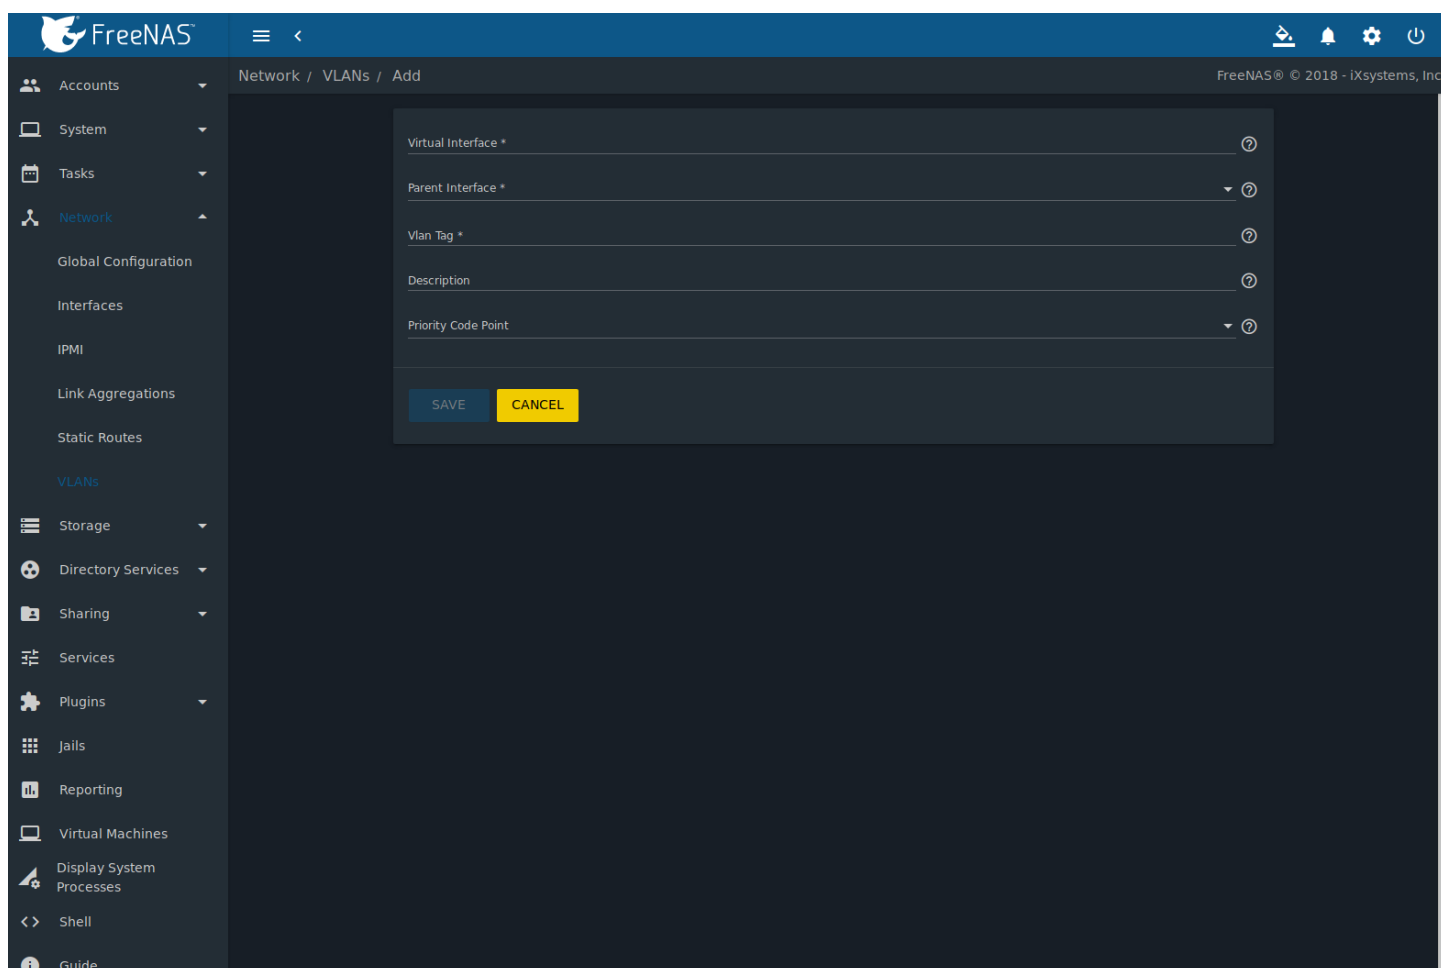


Fig. 8.9: Adding a VLAN

Table 8.7 summarizes the configurable fields.

Table 8.7: Adding a VLAN

Setting	Value	Description
Virtual Interface	string	Use the format <i>vlanX</i> where <i>X</i> 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. Newly created Link Aggregations (page 146) do not appear in the drop-down until the system is rebooted.
Vlan Tag	integer	Enter a number between 1 and 4095 which matches a numeric tag set up in the switched network.
Description	string	Optional. Enter any notes about this VLAN.
Priority Code Point	drop-down menu	Available 802.1p Class of Service ranges from <i>Best Effort (default)</i> to <i>Network Control (highest)</i> .

The parent interface of a VLAN must be up, but it can either have an IP address or be unconfigured, depending upon the requirements of the VLAN configuration. This makes it difficult for the web interface to do the right thing without trampling the configuration. To remedy this, add the VLAN, then select *Network* → *Interfaces*, and click *ADD*. Choose the parent interface from the *NIC* drop-down menu and in the *Options* field, type `up`. This brings up the parent interface. If an IP address is required, configure it using the rest of the options in the *ADD* screen.

Warning: Creating a VLAN causes an interruption to network connectivity. The web interface provides a warning about this interruption.

STORAGE

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

- [Pools](#) (page 154): create and manage storage pools.
- [Snapshots](#) (page 172): manage local snapshots.
- [VMware-Snapshots](#) (page 175): coordinate OpenZFS snapshots with a VMware datastore.
- [Disks](#) (page 176): view and manage disk options.
- [Importing a Disk](#) (page 181): import a **single** disk that is formatted with the UFS, NTFS, MSDOS, or EXT2 filesystem.

9.1 Pools

Storage → *Pools* is used to create and manage ZFS pools, datasets, and zvols.

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.**

9.1.1 Creating Pools

Before creating a pool, determine the level of required redundancy, how many disks will be added, and if any data exists on those disks. Creating a pool overwrites disk data, so save any required data to different media before adding disks to a pool.

Navigate to *Storage* → *Pools* and click *ADD* to open the screen shown in [Figure 9.1](#).

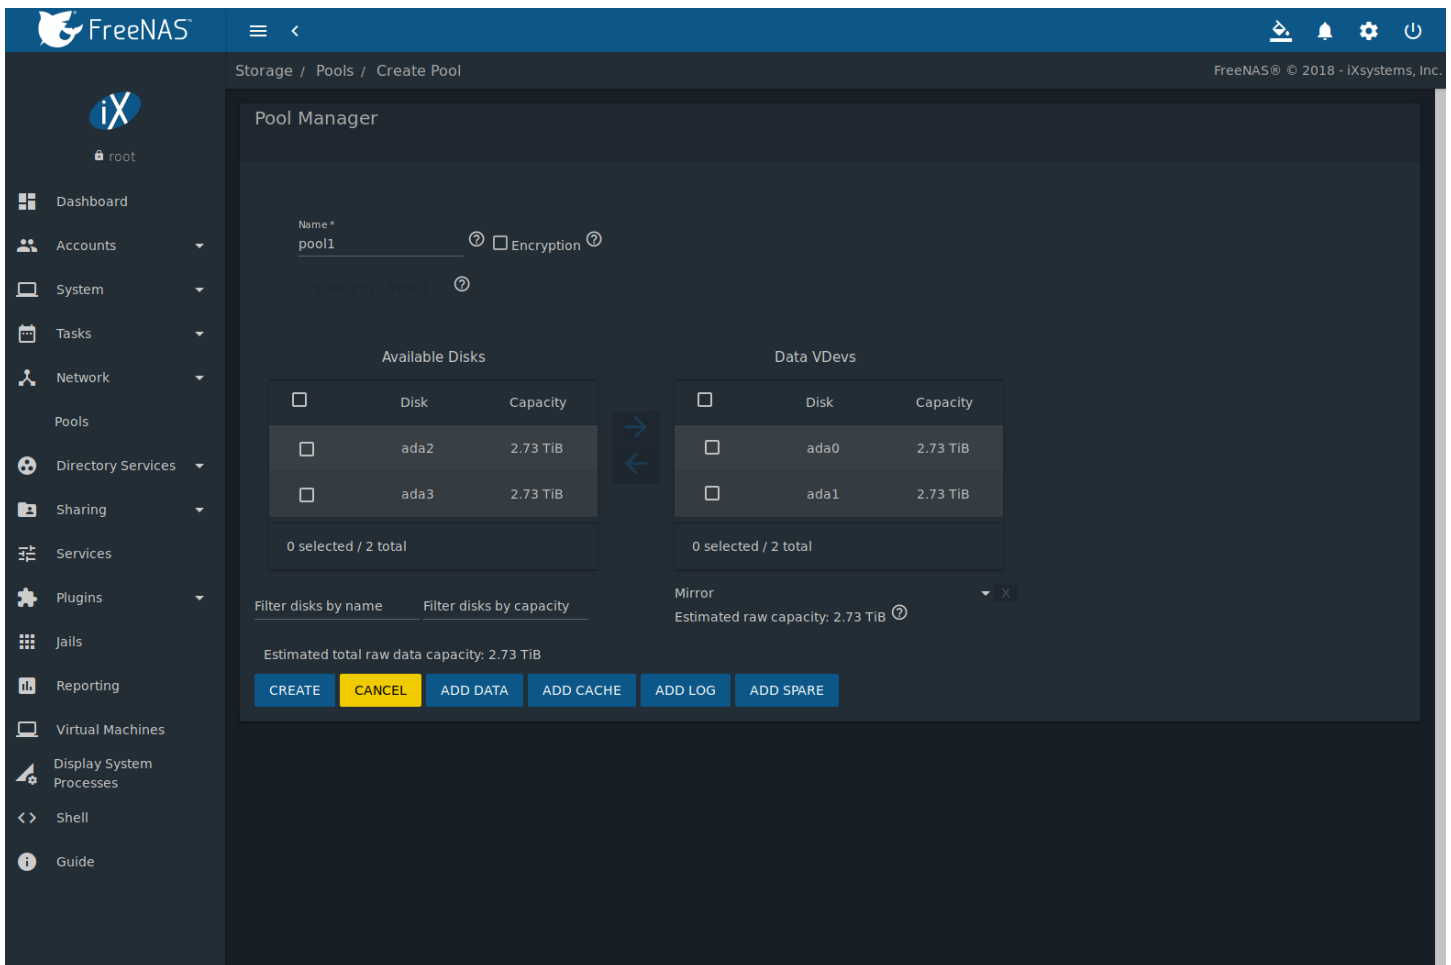


Fig. 9.1: Creating a Pool

Select *Create new pool* and click *CREATE POOL* to continue.

Enter a name for the pool in the *Name* field. Ensure that the chosen name conforms to these [naming conventions](https://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html) (https://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html). Choosing a name that will stick out in the logs is recommended, rather than generic names like “data” or “freenas”.

To encrypt data on the underlying disks as a protection against physical theft, set the *Encryption* option. A pop-up message shows a reminder to *Always back up the key!*. Without the key, the data on the disks is inaccessible. Select *Confirm* then click *I UNDERSTAND*.

Warning: Refer to the warnings in [Managing Encrypted Pools](#) (page 157) before enabling encryption!

In the *Available Disks* section, select the disks that will be added to the pool. Enter a value to *Filter disks by name* or *Filter disks by capacity* to change the order of displayed disks.

Note: The usable space of each disk in a pool is limited to the size of the smallest disk in the pool. Because of this, creating pools with disks of the same size is recommended.

After selecting the desired disks, click the right arrow to add them to the *Data VDevs* section. Any disks that appear in *Data VDevs* are used to create the pool. To remove a disk from that section, select the disk and use the left arrow to return it to the *Available Disks* section.

To add multiple *Data VDevs*, click *Add Data* for each required additional vdev.

Alternately, click the *SUGGEST LAYOUT* button which will add all of the disks and suggest an optimal layout for both redundancy and capacity.

The pool layout is dependent upon the number of disks added to *Data VDevs* and the number of available layouts increases as disks are added. To view the available layouts, ensure that at least one disk appears in *Data VDevs* and select the drop-down menu under this section. The web interface will automatically update the *Estimated total raw data capacity* when a layout is selected. 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

Warning: Refer to the [ZFS Primer](#) (page 354) for more information on redundancy and disk layouts. When more than five disks are used, consideration must be given to the optimal layout for the best performance and scalability. It is important to realize that different layouts of virtual devices (*vdevs*) affect which operations can be performed on that pool later. For example, drives can be added to a mirror to increase redundancy, but that is not possible with RAIDZ arrays.

After the desired layout is configured, click *CREATE*. A pop-up warning serves as a reminder that all disk contents will be erased. Click *Confirm*, then *CREATE POOL* to create the pool.

Note: To instead preserve existing data, click the *CANCEL* button and refer to [Importing a Disk](#) (page 181) and [Importing a Pool](#) (page 163) 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, create the pool, then restore the data to the new pool.

Depending on the size and number of disks, the type of controller, and whether encryption is selected, creating the pool may take some time. If the *Encryption* option was selected, a popup message provides a link to *Download Recovery Key*. Click the link and save the key to a safe location. When finished, click *DONE*.

The screen refreshes after the pool is created, showing the new pool in *Storage* → *Pools*.

In the example shown in [Figure 9.2](#), the created pool is named *pool1*. The description indicates that it is healthy. The *Used* and *Free* entries reflect the total size of the pool including disk parity.

Click the down arrow to see more details about the pool. This second entry has the same name and represents the implicit or root dataset. The *Used* and *Available* entries indicate the amount of disk space available for storage, after disk parity is subtracted. It also indicates the type of compression, the compression ratio, whether it is mounted as read-only, whether deduplication has been enabled, the mountpoint path, and any comments entered for the pool.

There is an option to *Upgrade Pool*. This upgrades the pool to the latest ZFS features, as described in [Upgrading a ZFS Pool](#) (page 35). This button does not appear if the pool is running the latest version of feature flags.

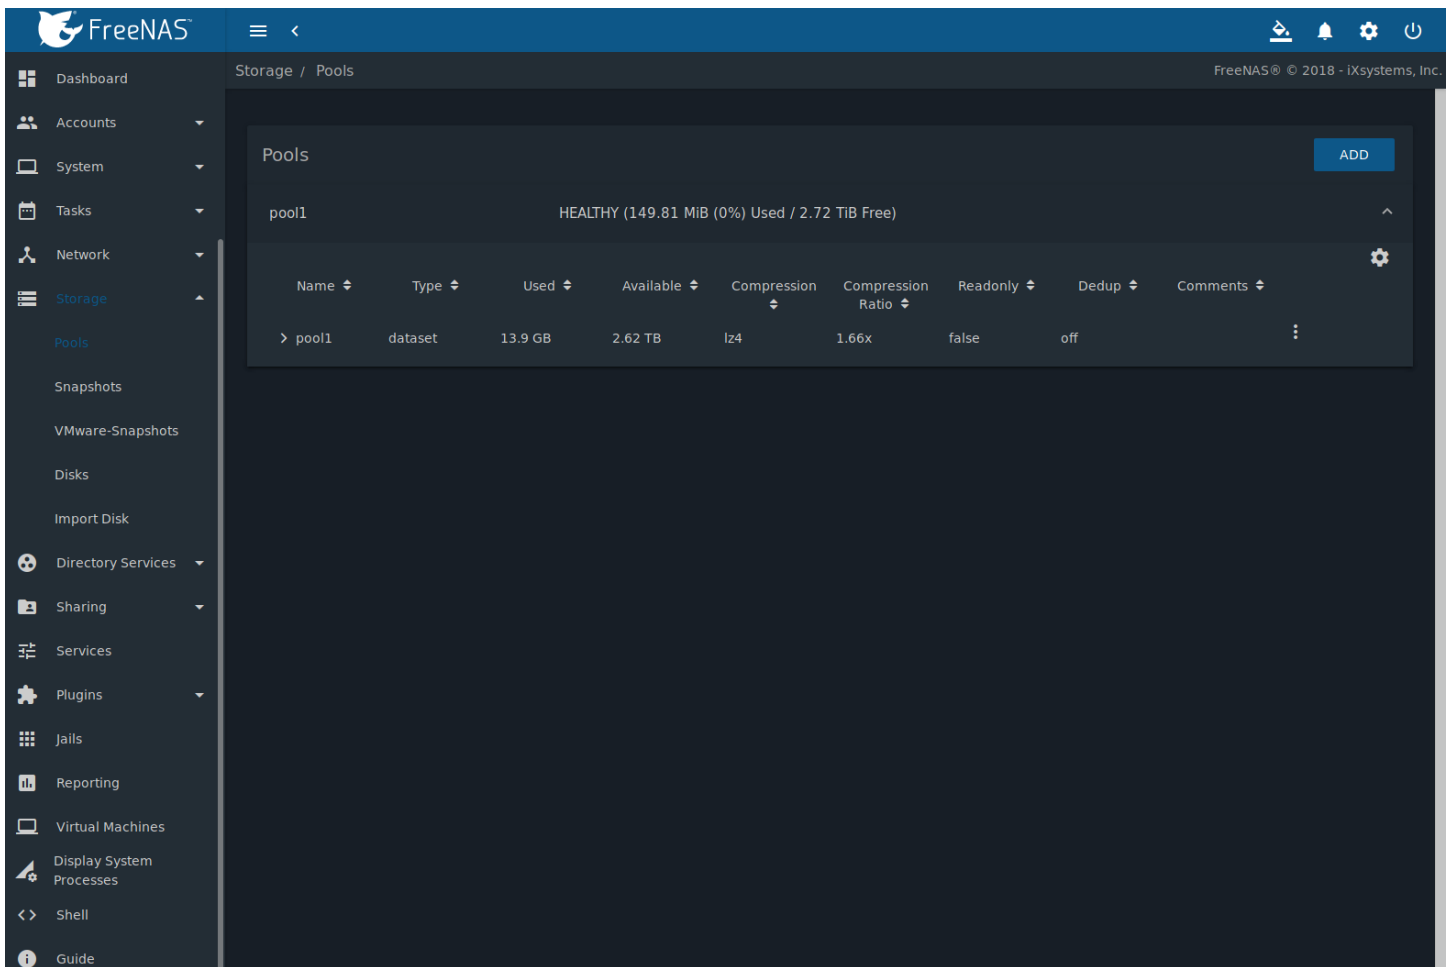


Fig. 9.2: Viewing Pools

9.1.2 Managing Encrypted Pools

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 web interface and that proper permissions are set on shares if sensitive data is stored on the system.

FreeNAS® supports [GELI](https://www.freebsd.org/cgi/man.cgi?query=geli) (<https://www.freebsd.org/cgi/man.cgi?query=geli>) full disk encryption for ZFS pools. It is important to understand the details before creating a pool with encryption:

- 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 pool is created on top of the encrypted devices.
- This type of encryption is primarily useful for users wanting the ability to remove disks from the pool without having to first wipe the disks of any sensitive data.
- 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 pool and 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.
- All drives in an encrypted pool are encrypted, including L2ARC (read cache) and SLOG (write cache). Drives added to an existing encrypted pool are encrypted with the same method specified when the pool was created.
- At present, there is no automated method to encrypt an existing, unencrypted pool. Instead, the data must be backed up, the existing pool destroyed, a new encrypted pool created, and the backup restored to the new pool.
- Hybrid pools are not supported. Added vdevs must match the existing encryption scheme. [Extending a Pool](#) (page 161) automatically encrypts a new vdev being added to an existing encrypted pool.

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

Note: Processors with support for the [AES-NI](https://en.wikipedia.org/wiki/AES_instruction_set) (https://en.wikipedia.org/wiki/AES_instruction_set) 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.

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

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 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 pool is decrypted and accessible when the system running. Protects “data at rest” only.
- *Key stored locally, with passphrase:* the encrypted pool is not accessible until the passphrase is entered by the FreeNAS® administrator.
- *Key not stored locally:* the encrypted pool 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 pool 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 pool is locked 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 pool, not individual users. When a pool is unlocked, data is accessible to all users with permissions to access it.

Note: [GELI](https://www.freebsd.org/cgi/man.cgi?query=geli) (<https://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 “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.

To manage the passphrase and keys on an encrypted pool, select the pool name in *Storage* → *Pools*, click the *Encryption Operations* (lock) icon, and select one of these operations:

Lock: Only appears after a passphrase has been created. When a pool is locked, its data is not accessible until the pool is unlocked by supplying the passphrase. For this reason, selecting this action will prompt to confirm. Once the pool is locked, its status will change to *LOCKED (Locked Used / Locked Free)*, *Pool Operations* are limited to *Detach* and *Extend*, and the *Encryption Operations* icon changes to the *Unlock* button.

Unlock the pool by clicking the *Unlock* icon and entering the passphrase *or* use the *Browse* button to load the recovery key. Only the passphrase is used when both a passphrase and a recovery key are entered. The services listed in *Restart Services* will restart when the pool is unlocked. This allows them to see the new pool and share or access data on it. Individual services can be prevented from restarting by clicking the *Restart Services* drop-down and unselecting them. However, a service that is not restarted might not be able to access the unlocked pool.

Create Passphrase: set and confirm a passphrase associated with the GELI encryption 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 you know should not be able to guess the passphrase). **Remember this passphrase. An encrypted pool cannot be reimported without it.** In other words, if the passphrase is forgotten, the data on the pool can become inaccessible if it becomes necessary to reimport the pool. Protect this passphrase, as anyone who knows it could reimport the encrypted pool, thwarting the reason for encrypting the disks in the first place.

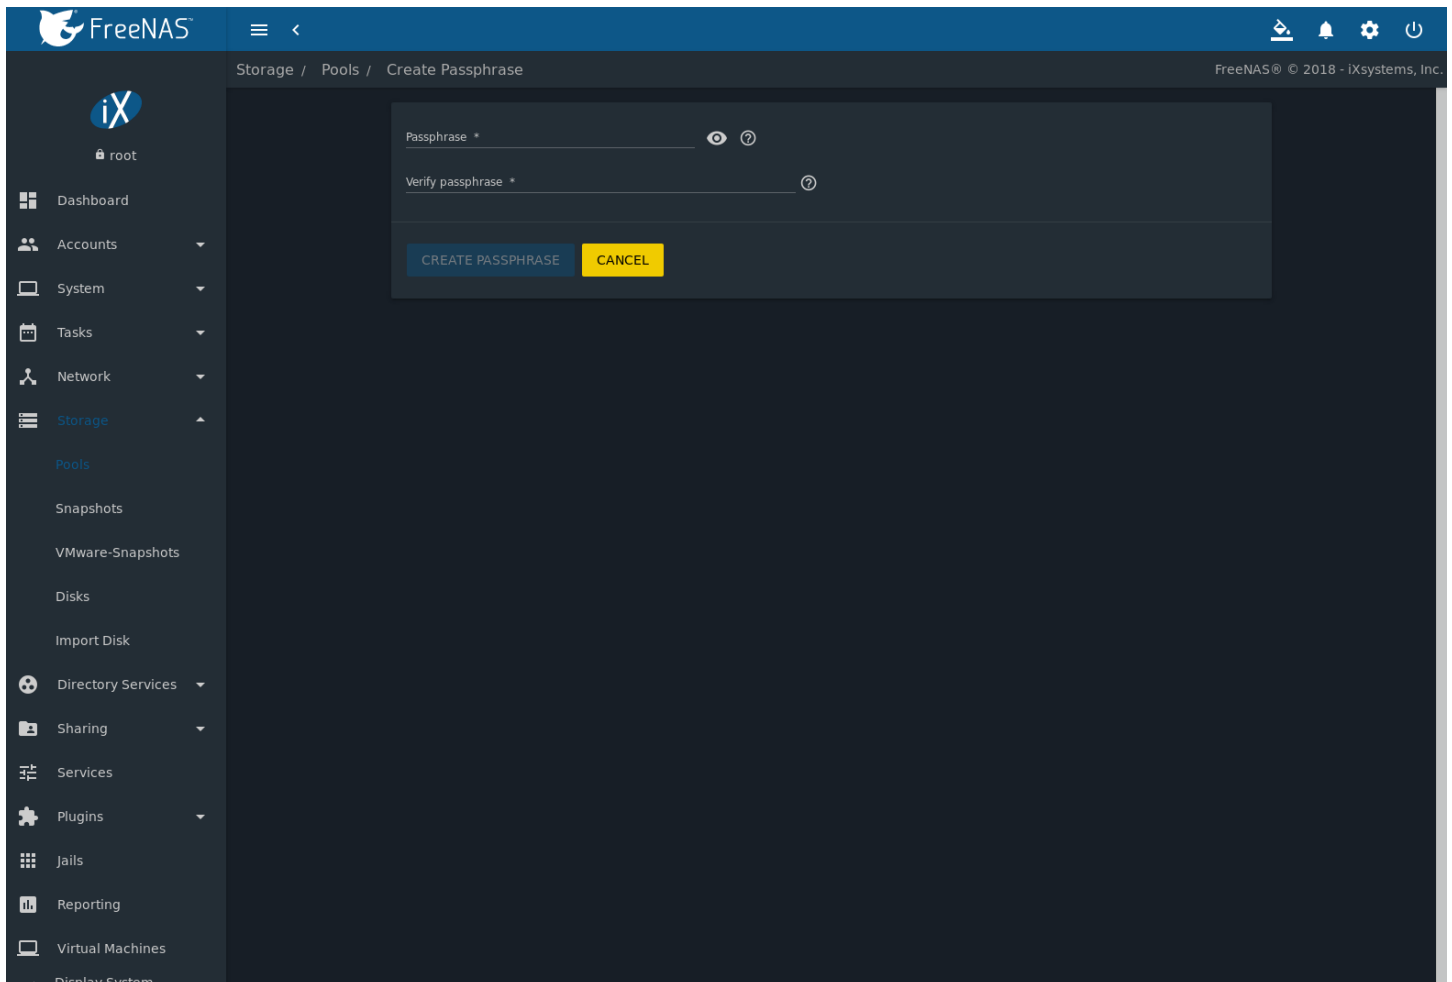


Fig. 9.3: Add a Passphrase to an Encrypted Pool

After the passphrase is set, the name of this button changes to *Change Passphrase* and the *Root Password* is also required to change the 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.

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.**

Delete 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: Protect the passphrase, recovery key, and encryption key. 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.


Encryption Rekey: 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.

Download Encrypt 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.**

9.1.3 Adding Cache or Log Devices



[Pools](#) (page 154) can be used either during or after pool creation to add an SSD as a cache or log device to improve performance of the pool under specific use cases. Before adding a cache or log device, refer to the [ZFS Primer](#) (page 354) to determine if the system will benefit or suffer from the addition of the device.

To add a Cache or Log device during pool creation, click the *Add Cache* or *Add Log* button. Select the disk from *Available Disks* and use the *right arrow* next to *Cache VDev* or *Log VDev* to add it to that section.

To add a device to an existing pool in *Storage* → *Pools*, click the pool name,  (Settings), then *Extend*. Click *Confirm* and *CONTINUE* to bypass the warning message. This will reopen the pool creation screen described in the previous paragraph, but with the pool name displayed as read-only.

..`index::` Remove cache or log device .. `_`Removing Cache or Log Devices:

9.1.4 Removing Cache or Log Devices

Cache or log devices can be removed by going to *Storage* → *Pools*. Choose the desired pool and click  (Settings) → *Status*. Choose the log or cache device to remove, then click  (Options) → *Remove*.

9.1.5 Adding Spare Devices

ZFS provides the ability to have “hot” *spares*. These are drives that are connected to a pool, but not in use. If the pool 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 pool, the spare is promoted to a full member of the pool.

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

To add a spare during pool creation, click the *Add Spare*. button. Select the disk from *Available Disks* and use the *right arrow* next to *Spare VDev* to add it to the section.

To add a device to an existing pool, click the pool name, ⓘ (Settings) icon, then *Extend*. Click *Confirm* and *CONTINUE* to bypass the warning message. This will reopen the pool creation screen described in the previous paragraph, but with the pool name displayed as read-only.

9.1.6 Extending a Pool

To increase the capacity of an existing pool, click the pool name, ⓘ (Settings), then *Extend*. A popup warning displays a reminder to stripe vdevs of the same size and type. Click *Confirm* and *CONTINUE* to continue.

Note: If the existing pool is encrypted, an additional warning message shows a reminder that **extending a pool resets the passphrase and recovery key**. After extending the pool, another popup message will provide a link to *Download Recovery Key*. Click the link and save the key to a safe location. When finished, click *DONE*.

When adding disks to increase the capacity of a pool, 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 pool. Extending a pool 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.

Warning: Make sure to select the same number of disks and disk layout when extending the pool!

9.1.7 Detaching a Pool

To export or destroy an existing pool, click the pool name, ⓘ (Settings), then *Detach*. The choice to retain or delete the contents of the pool depends upon the selections made in the screen shown in [Figure 9.4](#).

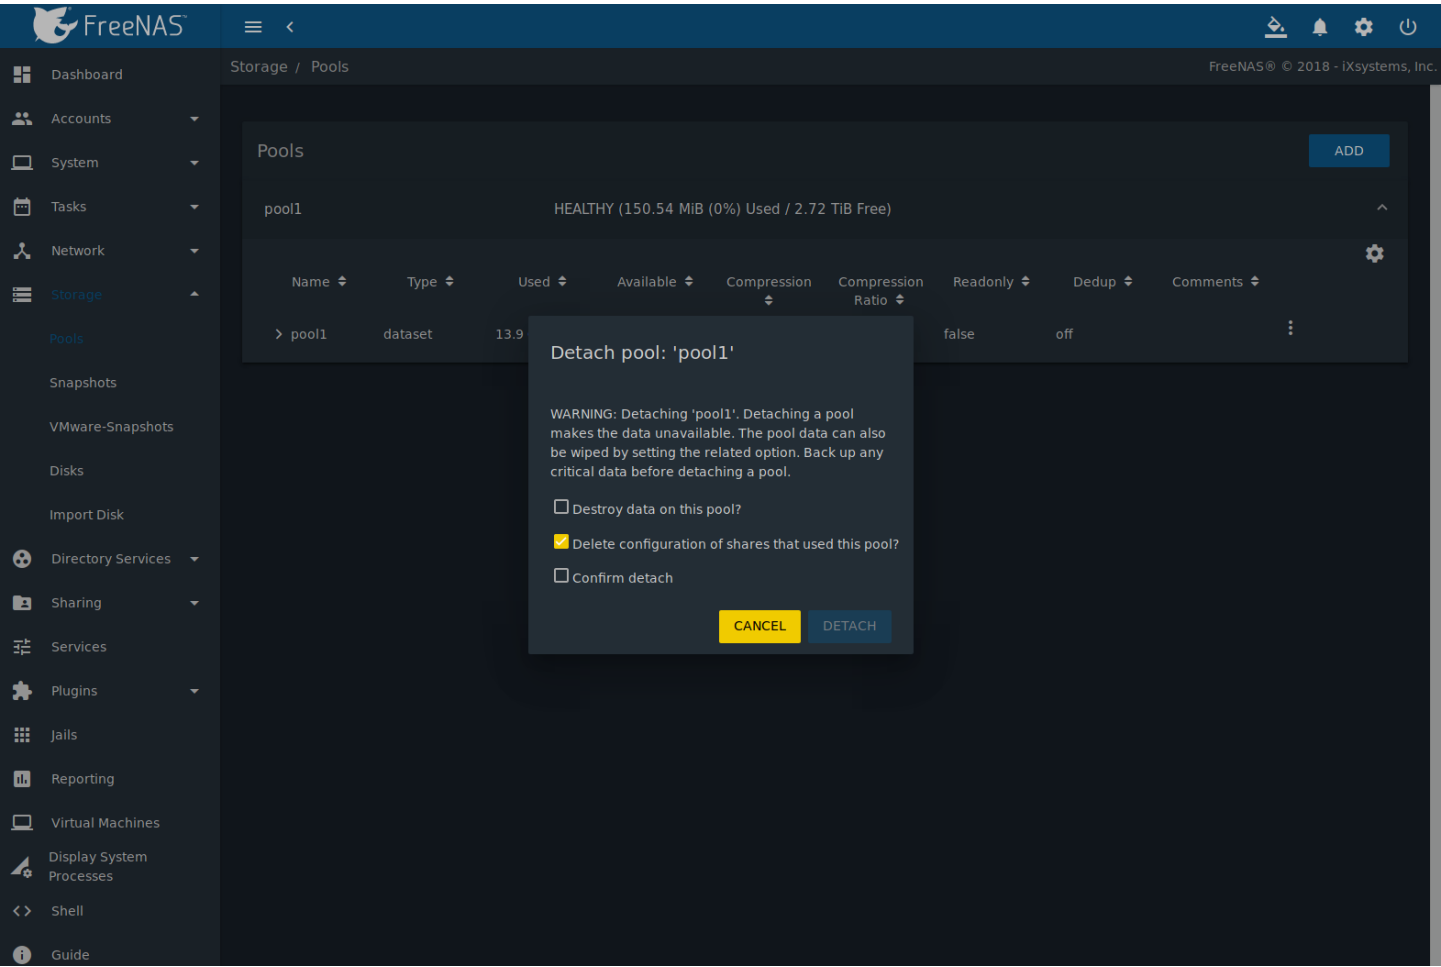


Fig. 9.4: Detach or Delete a Pool

Warning: Do not detach an encrypted pool if the passphrase has not been set! **An encrypted pool cannot be re-imported without a passphrase!** When in doubt, use the instructions in [Managing Encrypted Pools](#) (page 157) to set a passphrase.

The *Detach Pool* screen provides the options *Destroy data on this pool?*, *Confirm detach*, and *Delete configuration of shares that used this pool?*. An encrypted pool also displays a button to *DOWNLOAD KEY* for that pool.

Table 9.1: Detach Pool Options

Setting	Description
Destroy data on this pool?	Leave unset to keep existing data stored on the pool.
Delete configuration of shares that used this pool?	Leave unset to save the settings of the shares on the pool.
Confirm Detach	Confirm the detach process.

To detach the pool while retaining the data and share configurations, set the *Confirm detach* option and click *DETACH*. This operation allows the pool to be re-imported at a later time. For example, when moving a pool from one system to another, perform this detach action first to flush any unwritten data to disk, write data to the disk indicating that the export was done, and remove all knowledge of the pool from this system.

To instead destroy the data and share configurations on the pool, set all three options. This instructs the system to destroy the data on the pool, datasets, zvols, and the configuration of each share. It returns the individual disk to a raw state.

Danger: Before destroying a pool, ensure that any needed data has been backed up to a different pool or system.

9.1.8 Importing a Pool

This action can be used to reimport a detached pool, import a pool that was created on another system, or to import a pool after reinstalling an existing FreeNAS® system.

When physically installing ZFS pool disks from another system, use the `zpool export poolname` command or a web interface equivalent to export the pool on that system. Then shut it down and connect the drives to the FreeNAS® system. This prevents an “in use by another machine” error during the import to FreeNAS®.

Existing ZFS pools can be imported by clicking *Storage* → *Pools* and *ADD*. Select *Import an existing pool*, then click *NEXT* as shown in Figure 9.5.

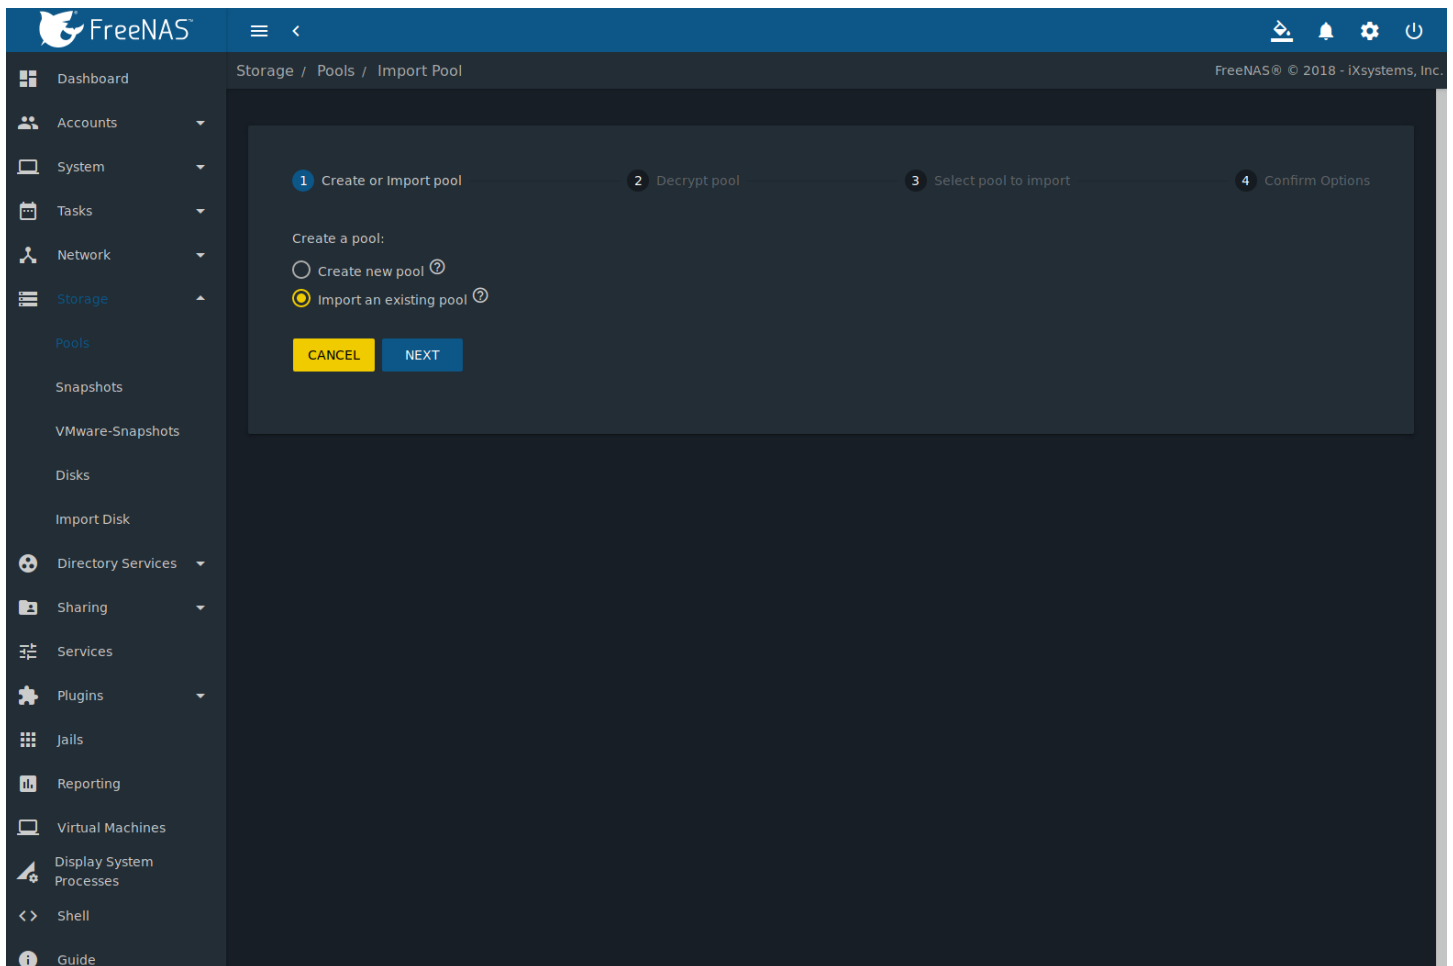


Fig. 9.5: Pool Import

To import a pool, click *No, continue with import* then *NEXT* as shown in Figure 9.6.

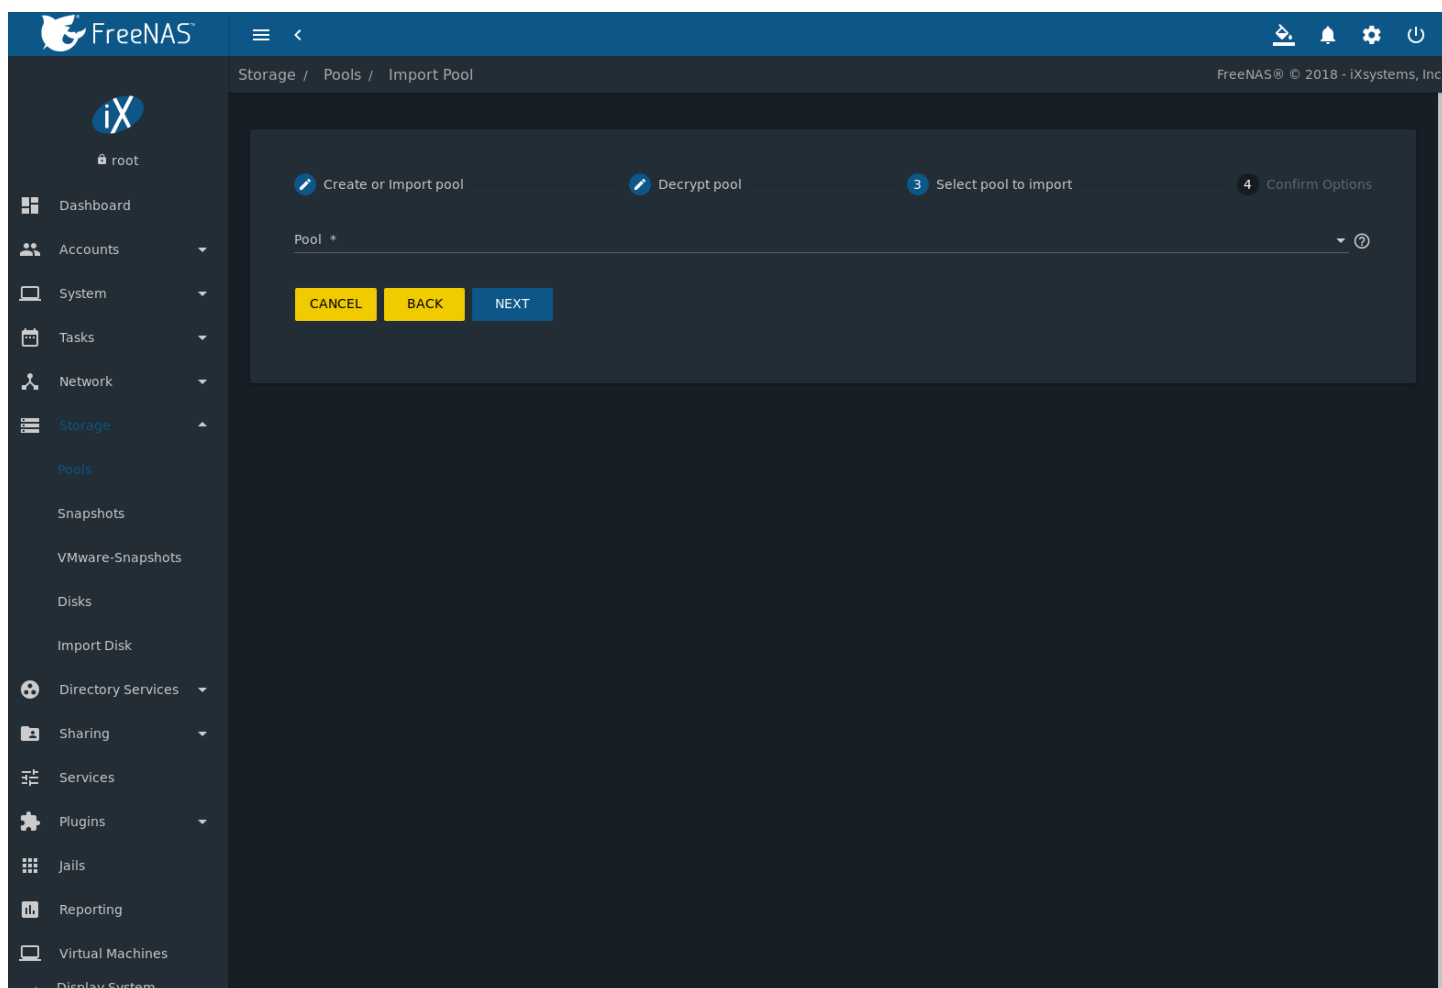


Fig. 9.6: Importing a Pool

Select the pool from the *Pool ** drop-down menu and click *NEXT* to confirm the options and *IMPORT* it.

If hardware is not being detected, run `camcontrol devlist` from [Shell](#) (page 325). 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 91).

Before importing a GELI-encrypted pool, disks must first be decrypted. Click *Yes, decrypt the disks*. This is shown in [Figure 9.7](#).

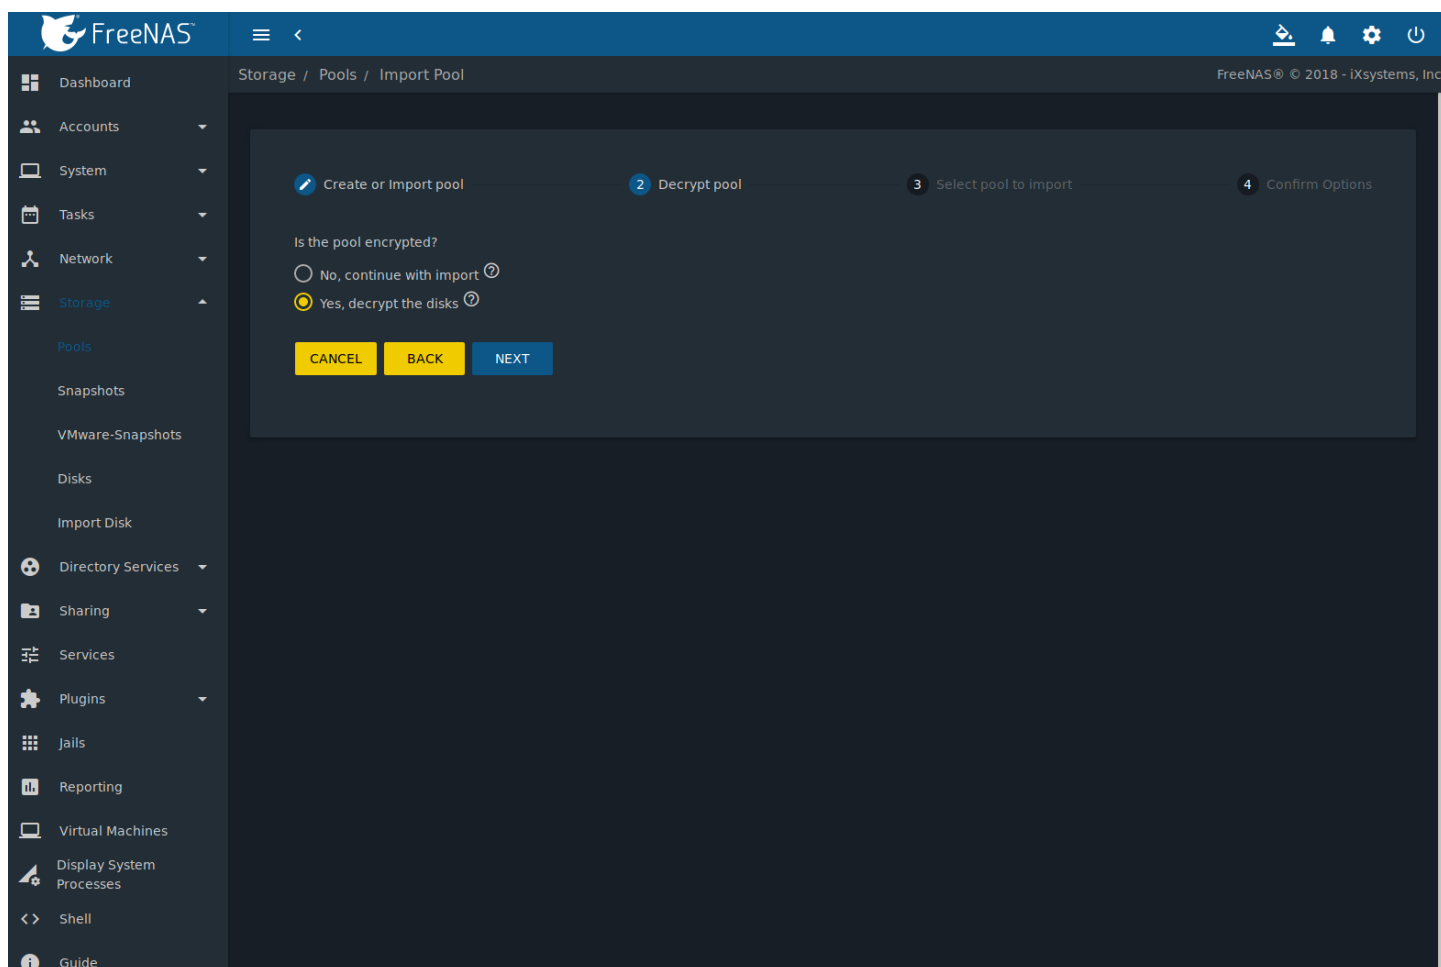


Fig. 9.7: Decrypting Disks Before Importing a Pool

Use the *Disks* dropdown menu to select the disks to decrypt. Click *Browse* to select an encryption key to upload. Enter the *Passphrase* associated with the key, then click *NEXT* to continue importing the pool.

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 Pools](#) (page 157) for instructions on managing keys.

Select the pool to import and confirm the settings. Click *IMPORT* to finish the process.

Note: For security reasons, GELI keys for encrypted pools 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 pool with Detach Pool, making sure that the *Destroy data on this pool?* is **not** selected. Then import the pool again. During the import, the GELI keys can be entered as described above.

9.1.9 Viewing Pool Scrub Status

Scrubs and how to set their schedule are described in more detail in [Scrub Tasks](#) (page 132).

To view the scrub status of a pool, click the pool name, (Settings), then *Status*. The resulting screen will display the status of a running scrub or the statistics from the last completed scrub.

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

9.1.10 Adding Datasets

An existing pool 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.

To create a dataset, select an existing pool in *Storage* → *Pools*, click ⓘ (Options), then select *Add Dataset*. This will display the screen shown in [Figure 9.8](#).

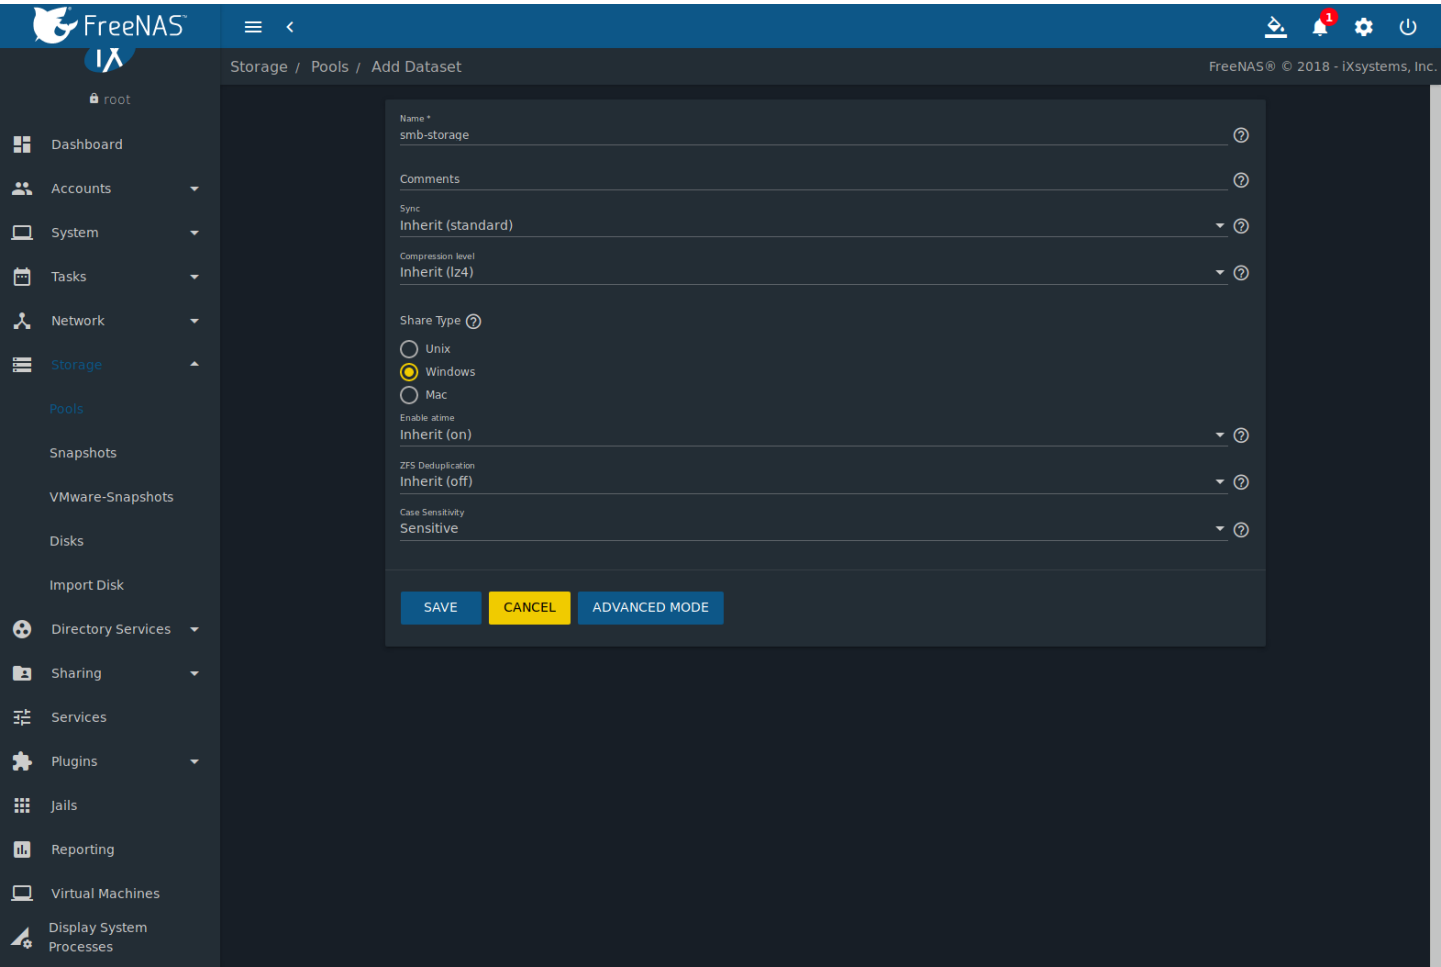



Fig. 9.8: Creating a ZFS Dataset

Table 9.2 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 enabling the *Show advanced fields by default* option in *System* → *Advanced*.

Table 9.2: Dataset Options

Setting	Value	Advanced Mode	Description
Name	string		This setting is mandatory. Enter a unique name for the dataset.
Comments	string		Enter any additional comments or user notes about this dataset.
Sync	drop-down menu		Sets the data write synchronization. <i>Inherit</i> inherits the sync settings from the parent dataset, <i>Standard</i> uses the sync settings that have been requested by the client software, <i>Always</i> waits for data writes to complete, and <i>Disabled</i> never waits for writes to complete.
Compression Level	drop-down menu		Refer to the section on Compression (page 168) 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 <i>UNIX</i> for an NFS share, <i>Windows</i> for a SMB share, or <i>Mac</i> for an AFP share.
Enable atime	Inherit, On, or Off		Choose <i>On</i> to update the access time for files when they are read. Choose <i>Off</i> to prevent producing log traffic when reading files. This can result in significant performance gains.
Quota for this dataset	integer	✓	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	✓	A specified value applies to both this dataset and any child datasets.
Reserved space for this dataset	integer	✓	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 otherwise take up all available free space.
Reserved space for this dataset and all children	integer	✓	A specified value applies to both this dataset and any child datasets.
ZFS Deduplication	drop-down menu		Read the section on Deduplication (page 168) before making a change to this setting.
Exec	drop-down menu	✓	Choices are <i>Inherit (on)</i> , <i>On</i> , or <i>Off</i> . Setting to <i>Off</i> will prevent the installation of Plugins (page 271) or Jails (page 281).
Read-only	drop-down menu	✓	Choices are <i>Inherit (off)</i> , <i>On</i> , or <i>Off</i> .
Snapshot directory	drop-down menu	✓	Choose if the <code>.zfs</code> snapshot directory is Visible or Invisible on this dataset.
Copies	drop-down menu	✓	Set the number of data copies on this dataset.
Record Size	drop-down menu	✓	While ZFS automatically adapts the record size dynamically to adapt to data, if the data has a fixed size (such as database records), matching its size might result in better performance. Warning: choosing a smaller record size than the suggested value can reduce disk performance and space efficiency.
Case Sensitivity	drop-down menu		Choices are <i>sensitive</i> (default, assumes filenames are case sensitive), <i>insensitive</i> (assumes filenames are not case sensitive), or <i>mixed</i> (understands both types of filenames).

After a dataset is created it appears in *Storage* → *Pools*. Click  (Options) on an existing dataset to configure these options:
Add Dataset: create a nested dataset, or a dataset within a dataset.

Add Zvol: add a zvol to the dataset. Refer to [Adding Zvols](#) (page 169) for more information about zvols.

Edit Options: edit the pool properties described in [Table 9.8](#). Note that the *Dataset Name*, and *Case Sensitivity* are read-only as they cannot be edited after dataset creation.

Edit Permissions: refer to [Setting Permissions](#) (page 170) for more information about permissions.

Delete Dataset: clicking this option will popup a warning as a reminder that this irreversible action will also delete all snapshots for the dataset. Set the *Confirm* option then click *DELETE DATASET* to destroy the dataset and all of its contents.

Promote Dataset: only appears on 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 deleted while the origin filesystem exists.

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use [Periodic Snapshot Tasks](#) (page 118).

9.1.10.1 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 GiB 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 GiB 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](https://constantin.glez.de/2011/07/27/zfs-to-dedupe-or-not-dedupe/) (<https://constantin.glez.de/2011/07/27/zfs-to-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 325). 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 at *Off*, 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 a snapshot of 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.

9.1.10.2 Compression

When selecting a compression type, balancing performance with the amount of disk space saved by compression is recommended. 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:** default and recommended compression method as it allows compressed datasets to operate at near real-time speed. This algorithm only compresses files that will benefit from compression.
- **GZIP:** levels 1, 6, and 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 which eliminates runs of zeroes.

If *OFF* is selected as the *Compression level* when creating a dataset or zvol, compression will not be used on that dataset/zvol. This is not recommended as using *LZ4* has a negligible performance impact and allows for more storage capacity.

9.1.11 Adding Zvols


A zvol is a feature of ZFS that creates a raw block device over ZFS. The zvol can be used as an *iSCSI* (page 248) device extent. To create a zvol, select an existing ZFS pool or dataset, click ⓘ (Options), then *Add Zvol* to open the screen shown in Figure 9.9.

Fig. 9.9: Adding a Zvol

The configuration options are described in Table 9.3.

Table 9.3: zvol Configuration Options

Setting	Value	Advanced Mode	Description
zvol name	string		Enter a short name for the zvol. Using a zvol name longer than 63-characters can prevent accessing zvols as devices. For example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent. This setting is mandatory.
Comments	string		Enter any notes about this zvol.
Size for this zvol	integer		Specify size and value such as <i>10 Gib</i> . If the size is more than 80% of the available capacity, the creation will fail with an “out of space” error unless <i>Force size</i> is also enabled.
Force size	checkbox		By default, the system will not create a zvol if that operation will bring the pool to over 80% capacity. While NOT recommended , enabling this option will force the creation of the zvol.
Sync	drop-down menu		Sets the data write synchronization. <i>Inherit</i> inherits the sync settings from the parent dataset, <i>Standard</i> uses the sync settings that have been requested by the client software, <i>Always</i> waits for data writes to complete, and <i>Disabled</i> never waits for writes to complete.
Compression level	drop-down menu		Compress data to save space. Refer to Compression (page 168) for a description of the available algorithms.
ZFS Deduplication	drop-down menu		ZFS feature to transparently reuse a single copy of duplicated data to save space. Warning: this option is RAM intensive. Read the section on Deduplication (page 168) before making a change to this setting.
Sparse	checkbox		Used to provide thin provisioning. Use with caution as writes will fail when the pool is low on space.
Block size	drop-down menu	✓	The default is based on the number of disks in the pool. This can be set to match the block size of the filesystem which will be formatted onto the iSCSI target. Warning: Choosing a smaller record size than the suggested value can reduce disk performance and space efficiency.

Click  (Options) next to the desired zvol in *Storage* → *Pools* to access the *Delete zvol* and *Edit zvol* options.


Similar to datasets, a zvol name cannot be changed, and destroying a zvol requires confirmation.

9.1.12 Setting Permissions

Setting permissions is an important aspect of managing data access. The web interface is meant to set the **initial** permissions for a pool or dataset to make it available as a share. Once a share is available, the client operating system is used to fine-tune the permissions of the files and directories that are created by the client.

[Sharing](#) (page 195) contains configuration examples for several types of permission scenarios. This section provides an overview of the options available for configuring the initial set of permissions.

Note: For users and groups to be available, they must either be first created using the instructions in [Accounts](#) (page 64) or imported from a directory service using the instructions in [Directory Services](#) (page 183). 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.

To set the permissions on a pool or dataset, select its entry in *Storage* → *Pools*, click  (Options), then *Edit Permissions*. This displays the screen shown in [Figure 9.10](#). [Table 9.4](#) lists the options in this screen.

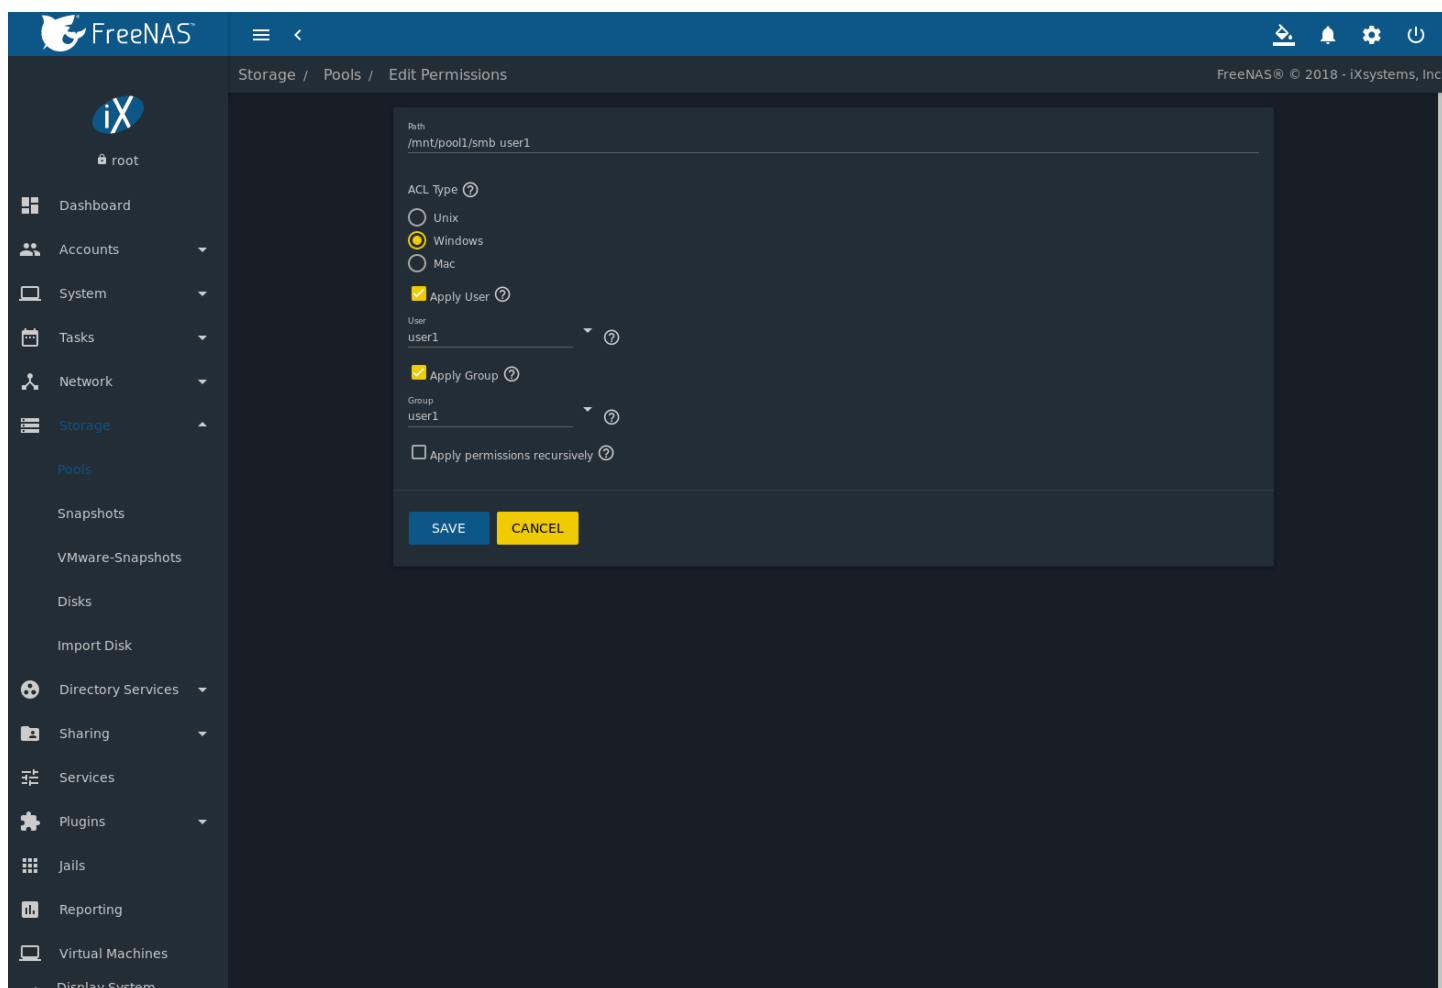


Fig. 9.10: Changing Permissions on a Dataset

Table 9.4: Permission Options

Setting	Value	Description
Path	string	Displays the path to the dataset or zvol directory.
ACL Type	bullet selection	Select the type that matches the type of client accessing. Choices are <i>Unix</i> , <i>Windows</i> or <i>Mac</i> . See description below this table.
Apply User	checkbox	Deselect to prevent new permission change from being applied to <i>User</i> , as described in the Note below this table.
User	drop-down menu	Select the user to control the permissions. Users manually created or imported from a directory service will appear in the drop-down menu.
Apply Group	checkbox	Deselect to prevent new permission change from being applied to <i>Group</i> , as described in the Note below this table.
Group	drop-down menu	Select the group to own the pool or dataset. Groups manually created or imported from a directory service will appear in the drop-down menu.
Apply Mode	checkbox	Unset to prevent new permission change from being applied to <i>Mode</i> , as described in the Note below this table.
Mode	checkboxes	Only applies to the <i>Unix</i> or <i>Mac</i> ACL Type so does not appear if <i>Windows</i> is selected. Sets the Unix-style permissions for owner, group, and other.
Apply permissions recursively	checkbox	If set, permissions will also apply to subdirectories. If data is already present on the pool or dataset, changing the permissions on the client side is recommended to prevent a performance lag.

Note: The *Apply User*, *Apply Group*, and *Apply Mode* options allow fine-tuning of the change permissions behavior. By default, all three options are enabled and FreeNAS® resets the *User*, *Group*, and *Mode* when the *SAVE* button is clicked. These options allow choosing which settings to change. For example, to change just the *Group* setting, unset the options for *Apply User* and *Apply Mode*.

The *Windows ACL Type* is used for [Windows \(SMB\) Shares](#) (page 211) or when the FreeNAS® system is a member of an Active Directory domain. This type adds ACLs to traditional Unix permissions. When the *Windows ACL Type* is selected, ACLs are set to the Windows defaults for new files and directories. A Windows client can be used to further fine-tune permissions as needed.

Warning: Changing a pool or dataset with *Windows* permissions back to *Unix* permissions will overwrite and destroy some of the extended permissions provided by Windows ACLs.

The *Unix ACL Type* is usually used with [Unix \(NFS\) Shares](#) (page 203). Unix 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 [Windows \(SMB\) Shares](#) (page 211).

The *Mac ACL Type* can be used with [Apple \(AFP\) Shares](#) (page 196).

9.1.13 View Multipaths

FreeNAS® uses `gmultipath(8)` (<https://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* → *Pools* → *View Multipaths*. Note that this option is not displayed in the *Storage* → *Pools* tree on systems that do not contain multipath-capable hardware.

9.2 Snapshots

Snapshots are scheduled using *Tasks* → *Periodic Snapshot Tasks*. To view and manage the listing of created snapshots, use *Storage* → *Snapshots*. An example is shown in [Figure 9.11](#).

Note: If snapshots do not appear, check that the current time configured in [Periodic Snapshot Tasks](#) (page 118) 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 325).

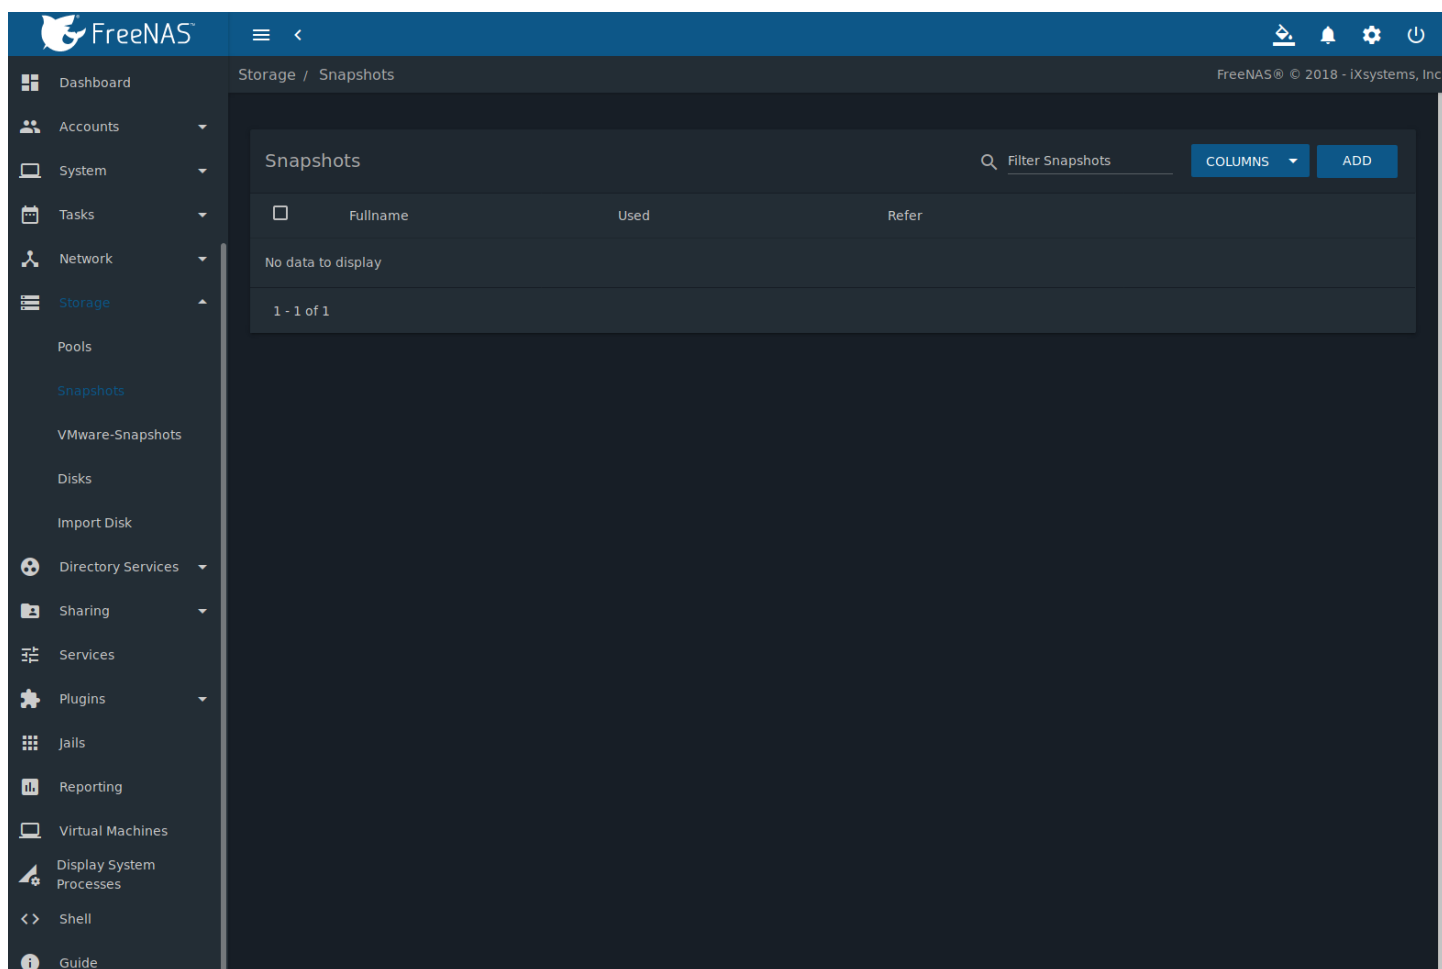



Fig. 9.11: Viewing Available Snapshots

Each entry in the listing includes the name of the snapshot, based on the volume/dataset name and time of the snapshot, and the amount of used and referenced data, where:

Used is the amount of space consumed by this dataset and all of its descendants. This value is checked against the dataset quota and reservation. The space used does not include the dataset 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 freed if this dataset is recursively deleted, 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 used space of the snapshot. Deleting a snapshot 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 325).

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 filesystem or snapshot it was created from, since its contents are identical.

To manage a snapshot, click  (Options) next to its entry. These actions are available from that menu:

Delete a pop-up message asks for confirmation. Child clones must be deleted before their parent snapshot can be deleted. 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.

Clone prompts for the name of the clone to create. A default name is provided that is based upon the name of the original snapshot but can be edited. Click the *SAVE* button to finish cloning the snapshot.

A clone is a writable copy of the snapshot. Since a clone is actually a dataset which can be mounted, it appears in the *Pools* screen rather than the *Snapshots* screen. By default, `-clone` is added to the name of a snapshot when a clone is created.

Rollback: only appears on the most recent snapshot. 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, delete the clone in the *Active Pools* tab.

This approach does not destroy any on-disk data and has no impact on replication.

A range of snapshots can be deleted. Set the left column checkboxes for each snapshot and click the *Delete* icon above the table. Be careful when deleting multiple snapshots.

Periodic snapshots can be configured to appear as shadow copies in newer versions of Windows Explorer, as described in [Configuring Shadow Copies](#) (page 221). Users can access the files in the shadow copy using Explorer without requiring any interaction with the FreeNAS® web interface.

To quickly search through the snapshots list by name, type a matching criteria into the *Filter Snapshots* text area. The listing will change to only display the snapshot names that match the filter text.

The *Items per page* drop-down menu is used to reduce or increase the amount of entries per page. Use the left or right arrows to scroll through a multi-page listing.

Warning: A snapshot and any files it contains will not be accessible or searchable if the mount path of the snapshot is longer than 88 characters. The data within the snapshot will be safe, and the snapshot will become accessible again when the mount path is shortened. For details of this limitation, and how to shorten a long mount path, see [Path and Name Lengths](#) (page 14).

9.2.1 Browsing a Snapshot Collection

All snapshots for a dataset are accessible as an ordinary hierarchical filesystem, which can be reached from a hidden `.zfs` file located at the root of every dataset. A user with permission to access that file can view and explore all snapshots for a dataset like any other files - from the `CLI` or via *File Sharing* services such as *Samba*, *NFS* and *FTP*. This is an advanced capability which requires some command line actions to achieve. In summary, the main changes to settings that are required are:

- Snapshot visibility must be manually enabled in the ZFS properties of the dataset.
- In Samba auxiliary settings, the `veto files` command must be modified to not hide the `.zfs` file, and the setting `zfsacl:expose_snapdir=true` must be added.

The effect will be that any user who can access the dataset contents will be able to view the list of snapshots by navigating to the `.zfs` directory of the dataset. They will also be able to browse and search any files they have permission to access throughout the entire snapshot collection of the dataset.

A user's ability to view files within a snapshot will be limited by any permissions or ACLs set on the files when the snapshot was taken. Snapshots are fixed as "read-only", so this access does not permit the user to change any files in the snapshots, or to modify or delete any snapshot, even if they had write permission at the time when the snapshot was taken.

Note: ZFS has a `zfs diff` command which can list the files that have changed between any two snapshot versions within a dataset, or between any snapshot and the current data.

9.3 VMware-Snapshots

Storage → *VMware-Snapshots* is used 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 172).

Figure 9.12 shows the menu for adding a VMware snapshot and Table 9.5 summarizes the available options.

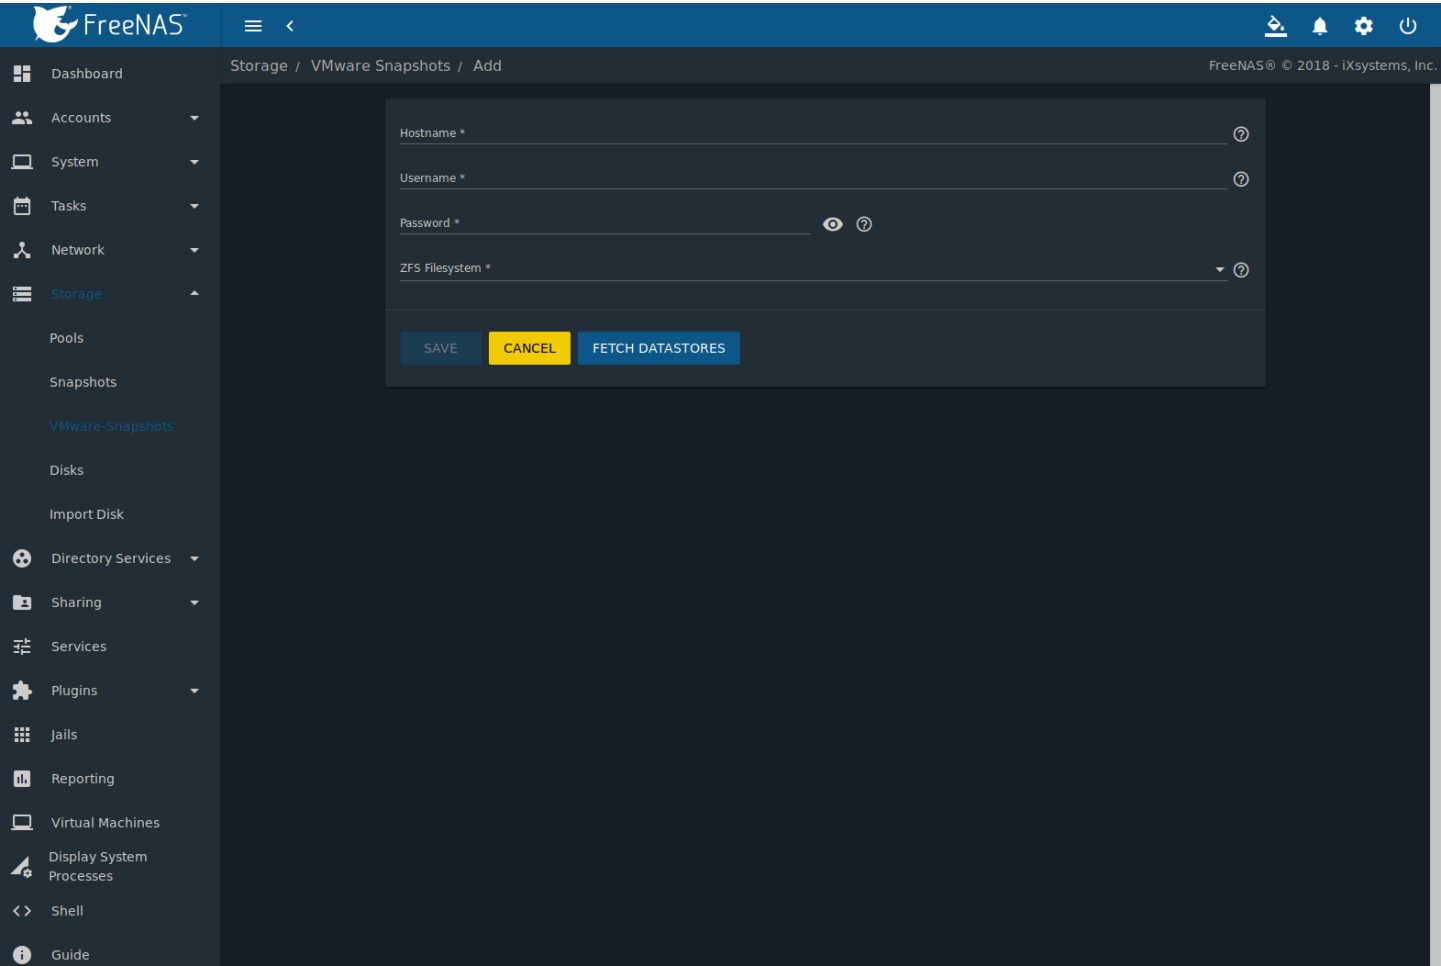


Fig. 9.12: Adding a VMware Snapshot

Table 9.5: VMware Snapshot Options

Setting	Value	Description
Hostname	string	Enter the IP address or hostname of the VMware host. When clustering, use the IP of the vCenter server for the cluster.

Continued on next page

Table 9.5 – continued from previous page

Setting	Value	Description
Username	string	Enter the username on the VMware host with permission to snapshot virtual machines.
Password	string	Enter the password associated with <i>Username</i> .
ZFS Filesystem	browse button	<i>Browse</i> to the filesystem to snapshot.
Datastore	drop-down menu	After entering the <i>Hostname</i> , <i>Username</i> , and <i>Password</i> , click <i>FETCH DATASTORES</i> to populate the menu, then select the datastore to be synchronized.

9.4 Disks

To view all of the disks recognized by the FreeNAS® system, use *Storage* → *Disks*. As seen in the example in [Figure 9.13](#), each disk entry displays its device name, pool membership, serial number, size, advanced power management settings, acoustic level settings, and if *S.M.A.R.T.* (page 256) testing is enabled. This table is adjustable by setting the different column checkboxes above it. Set *Toggle* to display all options in the table.

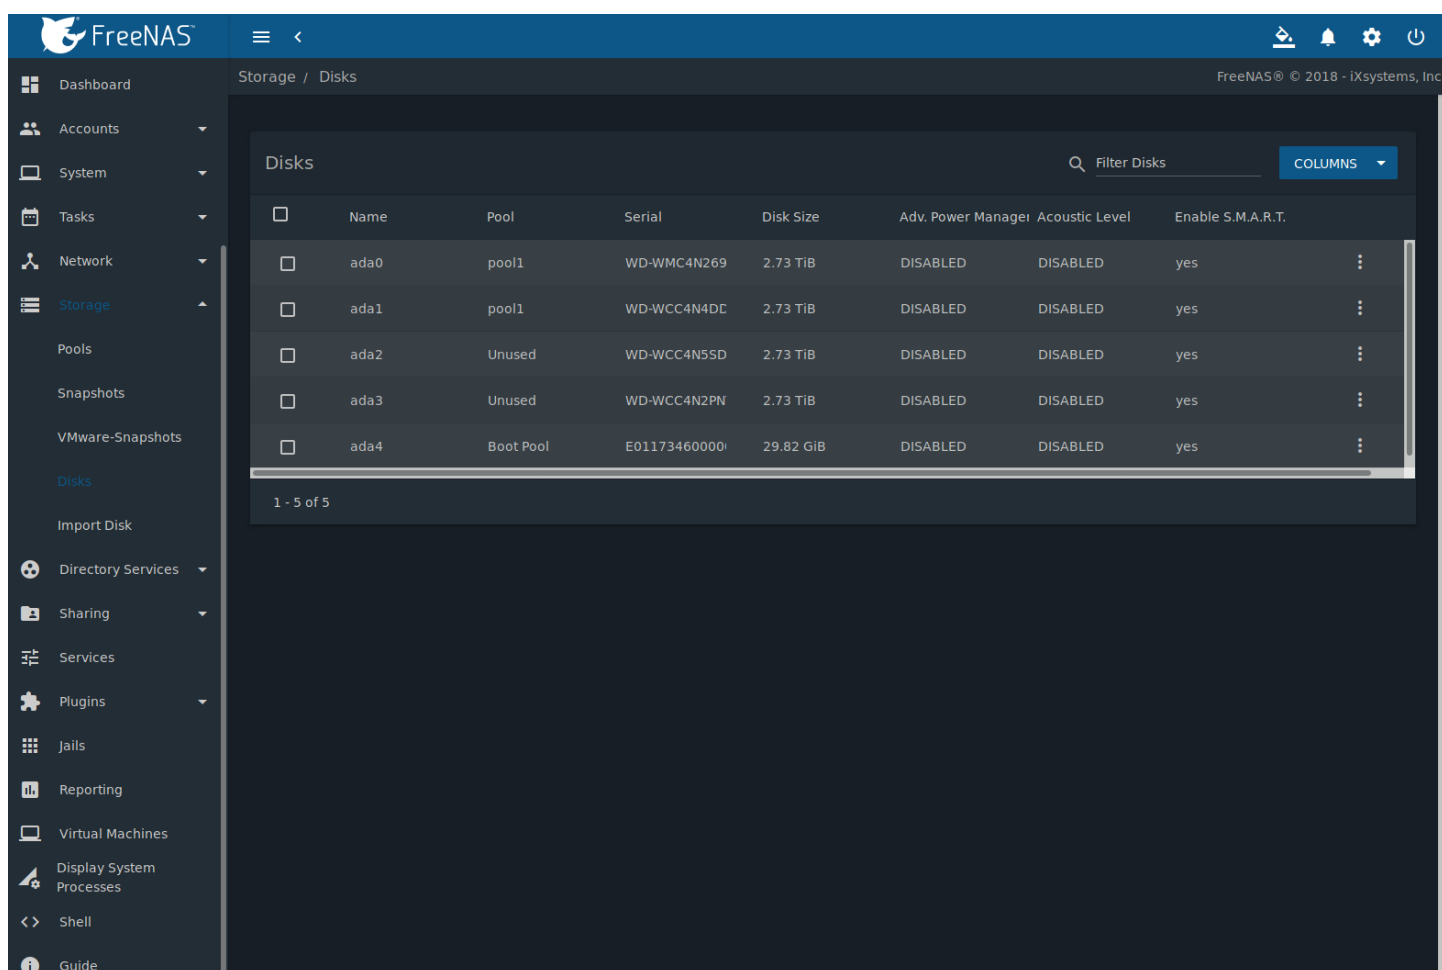



Fig. 9.13: Viewing Disks

To edit the options for a disk, click  (Options) on a disk, then *Edit* to open the screen shown in [Figure 9.14](#). [Table 9.6](#) lists the configurable options.

To bulk edit disks, set the checkbox for each disk in the table then click  (Edit Disks). The *Bulk Edit Disks* page displays

which disks are being edited and a short list of configurable options. The [Disk Options table](#) (page 177) indicates the options available when editing multiple disks.

To offline, online, or replace the device, see [Replacing a Failed Disk](#) (page 178).

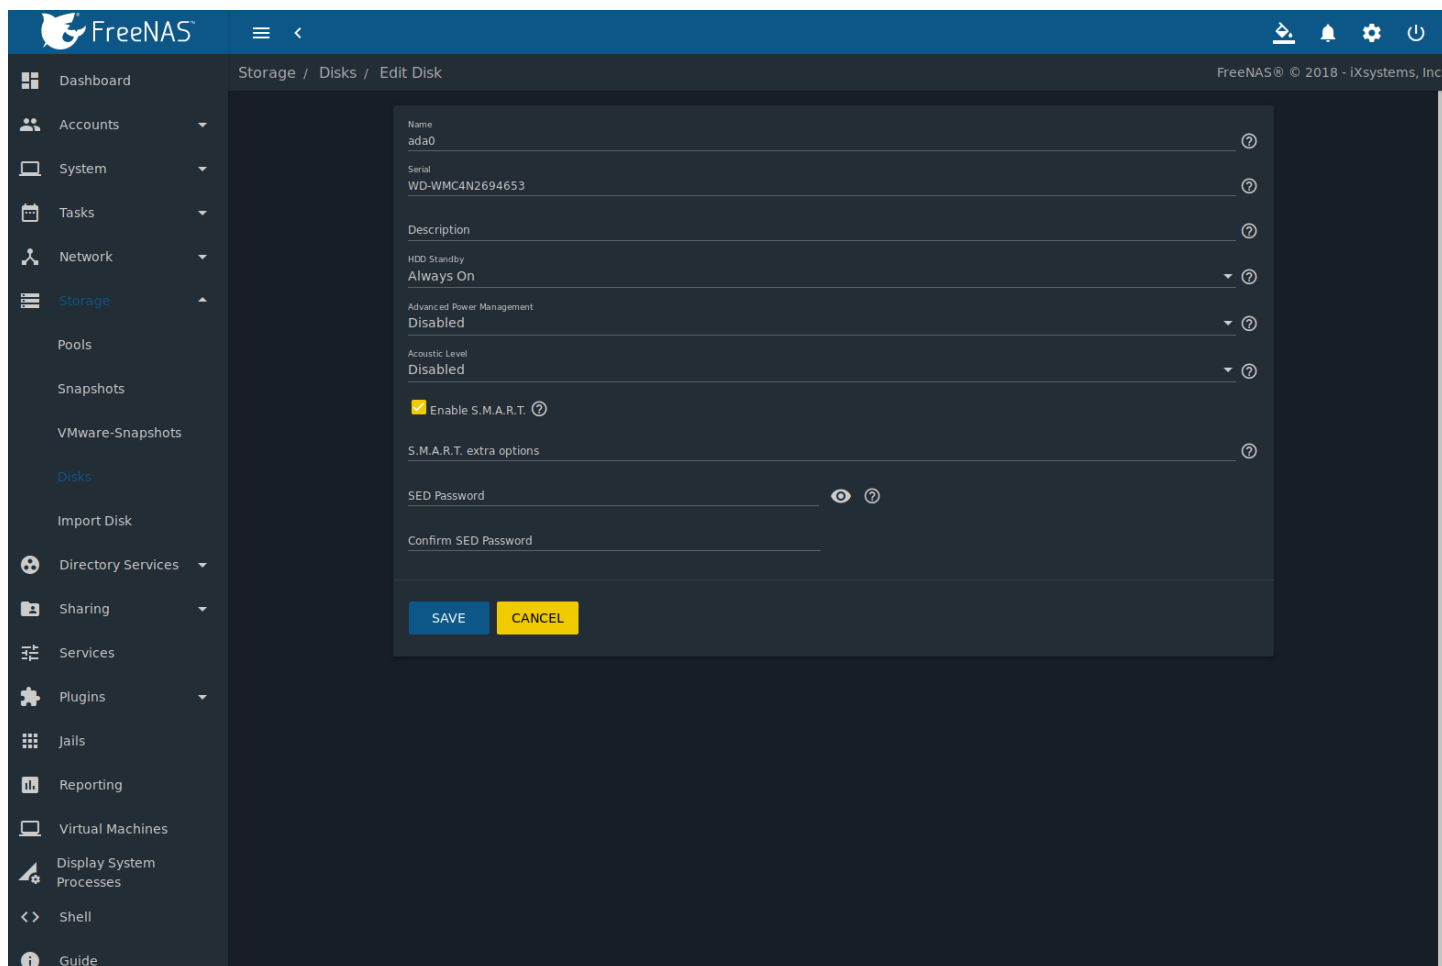


Fig. 9.14: Editing a Disk

Table 9.6: Disk Options

Setting	Value	Bulk Edit	Description
Name	string		This is the FreeBSD device name for the disk.
Serial	string		This is the serial number of the disk.
Description	string		Enter any notes about this disk.
HDD Standby	drop-down menu	✓	Indicates the time of inactivity in minutes before the drive enters standby mode 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	✓	Select a power management profile from the menu. The default value is <i>Disabled</i> .
Acoustic Level	drop-down menu	✓	Default is <i>Disabled</i> . Other values can be selected for disks that understand AAM (https://en.wikipedia.org/wiki/AAM).

Continued on next page

Table 9.6 – continued from previous page

Setting	Value	Bulk Edit	Description
Enable S.M.A.R.T.	checkbox	✓	Enabled by default if the disk supports S.M.A.R.T. Unsetting this option will disable any configured S.M.A.R.T. Tests (page 116) for the disk.
S.M.A.R.T. extra options	string	✓	Enter additional <code>smartctl(8)</code> (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) options.
SED Password	string		Enter and confirm the password which will be used for this device instead of the global SED password. Refer to Self-Encrypting Drives (page 81) for more information.

Tip: If the serial number for a disk is not displayed in this screen, use the `smartctl` command from [Shell](#) (page 325). For example, to determine the serial number of disk `ada0`, type `smartctl -a /dev/ada0 | grep Serial`.

The *Wipe* function is used to discard an unused disk.

Warning: Ensure all data is backed up and 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.


9.4.1 Replacing a Failed Disk

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 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 pool will be destroyed and must be recreated and the data restored from backup.


Note: If the pool is encrypted with GELI, refer to [Replacing an Encrypted Disk](#) (page 180) before proceeding.

Before physically removing the failed device, go to *Storage* → *Pools*. Select the pool name then click  (Settings). Select *Status* and locate the failed disk. Then perform these steps:

1. Click  (Options) on the disk entry, then *Offline* to change the disk status to OFFLINE. This step removes the device from the pool and prevents swap issues. If the hardware supports hot-pluggable disks, click the disk *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 status to OFFLINE fails with a “disk offline failed - no valid replicas” message, the pool must be scrubbed first with the *Scrub Pool* button in *Storage* → *Pools*. After the scrub completes, try *Offline* 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 web interface and locate the OFFLINE disk.

3. After the disk is replaced and is showing as OFFLINE, click  (Options) on the disk again and then *Replace*. Select the replacement disk from the drop-down menu and click the *REPLACE DISK* button. After clicking the *REPLACE DISK* button, the 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 116) screen.

In the example shown in [Figure 9.15](#), a failed disk is being replaced by disk *ada3* in the pool named *pool1*.

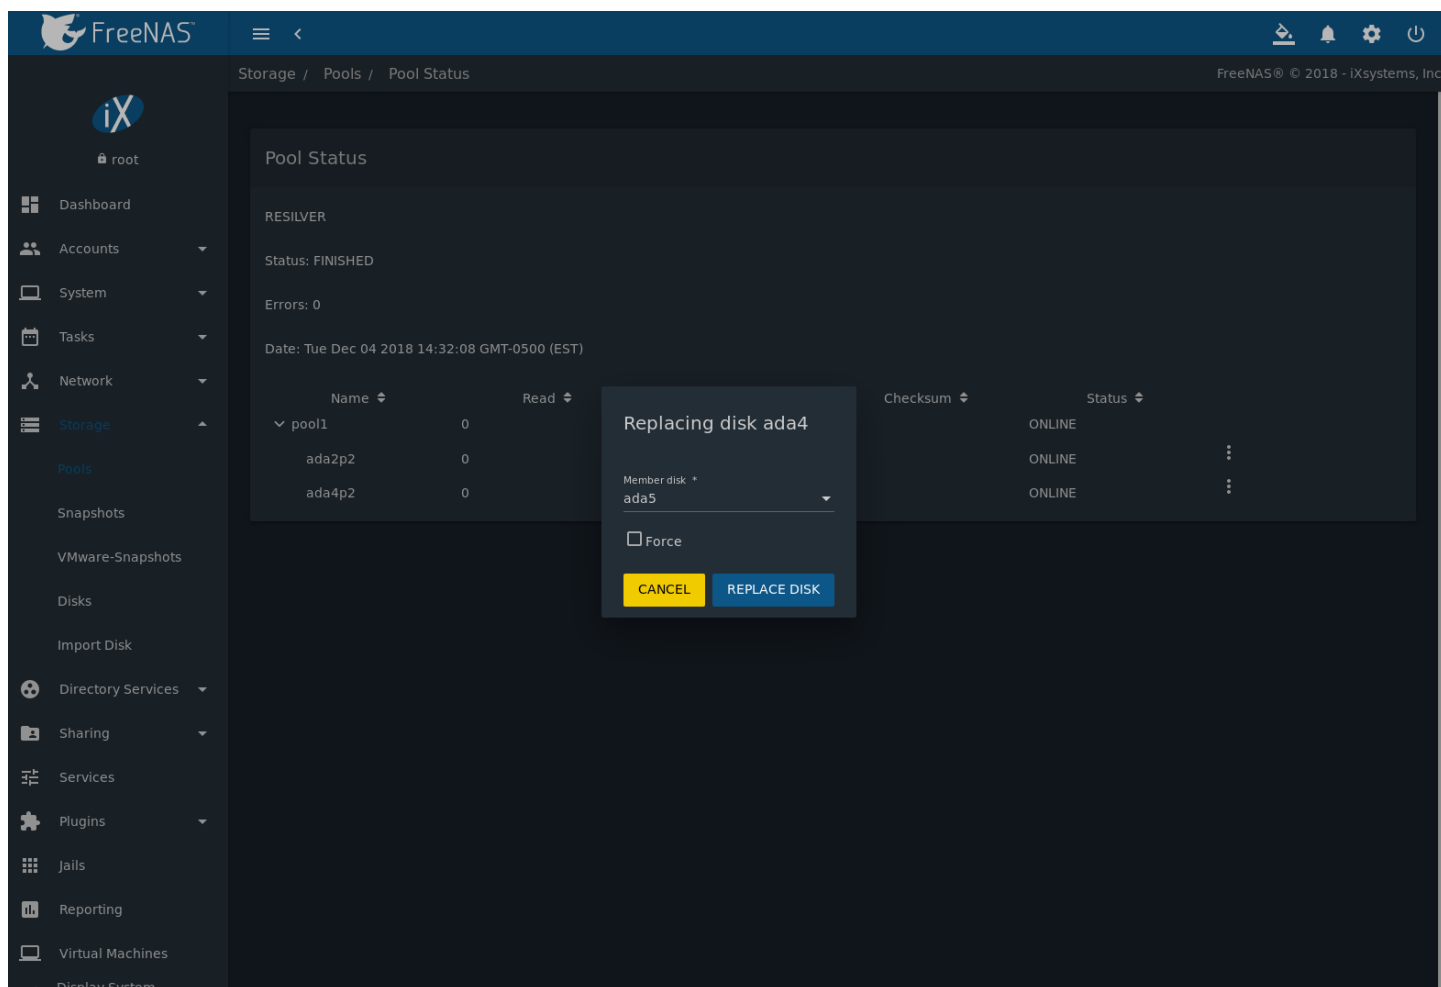


Fig. 9.15: Replacing a Failed Disk

After the resilver is complete, *Pools* shows a *Completed* resilver status and indicates any errors. [Figure 9.16](#) 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.

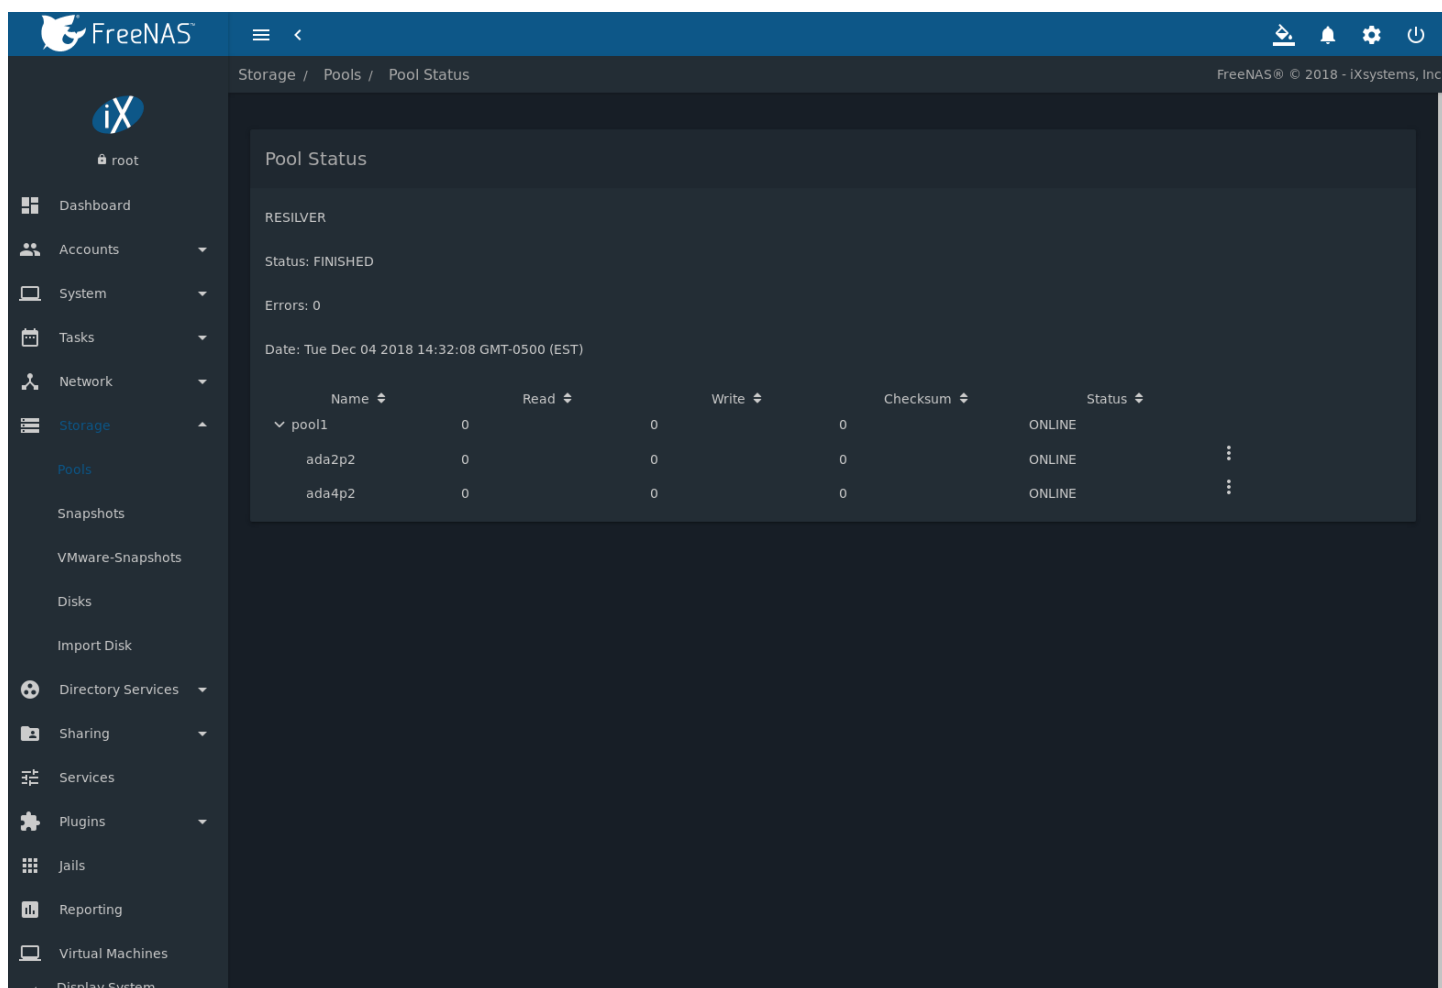


Fig. 9.16: Disk Replacement is Complete

9.4.1.1 Replacing an Encrypted Disk

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 [Managing Encrypted Pools](#) (page 157) **before** attempting to replace the failed drive. Then, follow steps 1 and 2 as described above. During step 3, there will be a prompt to enter and confirm the passphrase for the pool. Enter this information, then click **REPLACE DISK**. Wait until resilvering is complete.

Next, restore the encryption keys to the pool. **If this additional step is not performed before the next reboot, access to the pool might be permanently lost.**

1. Highlight the pool that contains the recently replaced disk and click *Add Recovery Key* to save the new recovery key. The old recovery key will no longer function, so it can be safely discarded.

9.4.1.2 Removing a Log or Cache Device

Added log or cache devices appear in *Storage → Pools → Pool Status*. Clicking the device enables the *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.

9.4.2 Replacing Disks to Grow a 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 [Pools](#) (page 154) 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* → *Pools*, select the pool to expand. Click  (Settings) then *Status*. Select a disk, click  (Options), then *Replace*. 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 Disk](#) (page 178). 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.

9.5 Importing a Disk

The *Pool* → *Import Disk* screen, shown in [Figure 9.17](#), is used to import a **single** disk that is formatted with a 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.

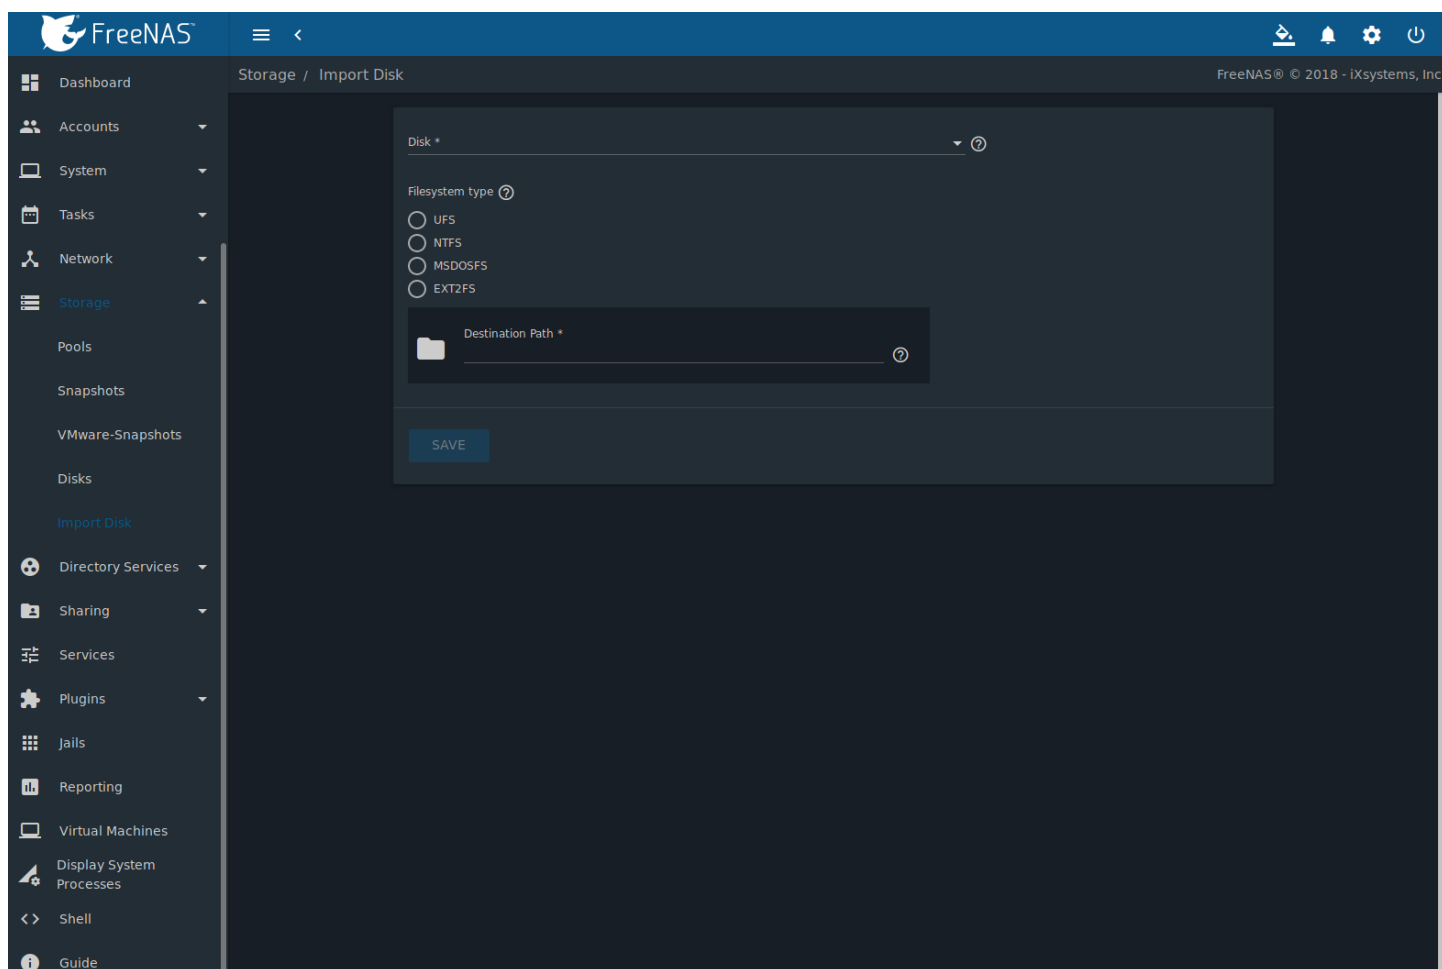


Fig. 9.17: Importing a Disk

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. If the *MSDOSFS* filesystem is selected, an additional *MSDOSFS locale* drop-down menu will display. Use this menu to select the locale if non-ASCII characters are present on the disk.

After clicking *SAVE*, the disk is mounted and its contents are copied to the specified dataset. The disk is unmounted after the copy operation completes.

DIRECTORY SERVICES

FreeNAS® supports integration with these directory services:

- *Active Directory* (page 183) (for Windows 2000 and higher networks)
- *LDAP* (page 188)
- *NIS* (page 191)

FreeNAS® also supports *Kerberos Realms* (page 192), *Kerberos Keytabs* (page 193), and the ability to add more parameters to *Kerberos Settings* (page 194).

This section summarizes each of these services and the available configuration options within the FreeNAS® web interface.

10.1 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/Setting_up_Samba_as_an_Active_Directory_Domain_Controller#Provisioning_a_Samba_Active_Di). Since AD provides authentication and authorization services for the users in a network, it is not necessary to recreate the same user accounts on the FreeNAS® system. Instead, configure the Active Directory service so 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® 11.2 before attempting Active Directory integration.

Ensure name resolution is properly configured before configuring the Active Directory service. `ping` the domain name of the Active Directory domain controller from *Shell* (page 325) on the FreeNAS® system. If the `ping` fails, check the DNS server and default gateway settings in *Network* → *Global Configuration* on the FreeNAS® system.

Add a DNS record for the FreeNAS® system on the Windows server and verify the hostname of the FreeNAS® system can be pinged from the domain controller.

Active Directory relies on Kerberos, 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.

To ensure both systems are set to the same time:

- use the same NTP server (set in *System* → *NTP Servers* on the FreeNAS® system)
- set the same timezone
- set either localtime or universal time at the BIOS level

Figure 10.1 shows *Directory Services* → *Active Directory* settings.

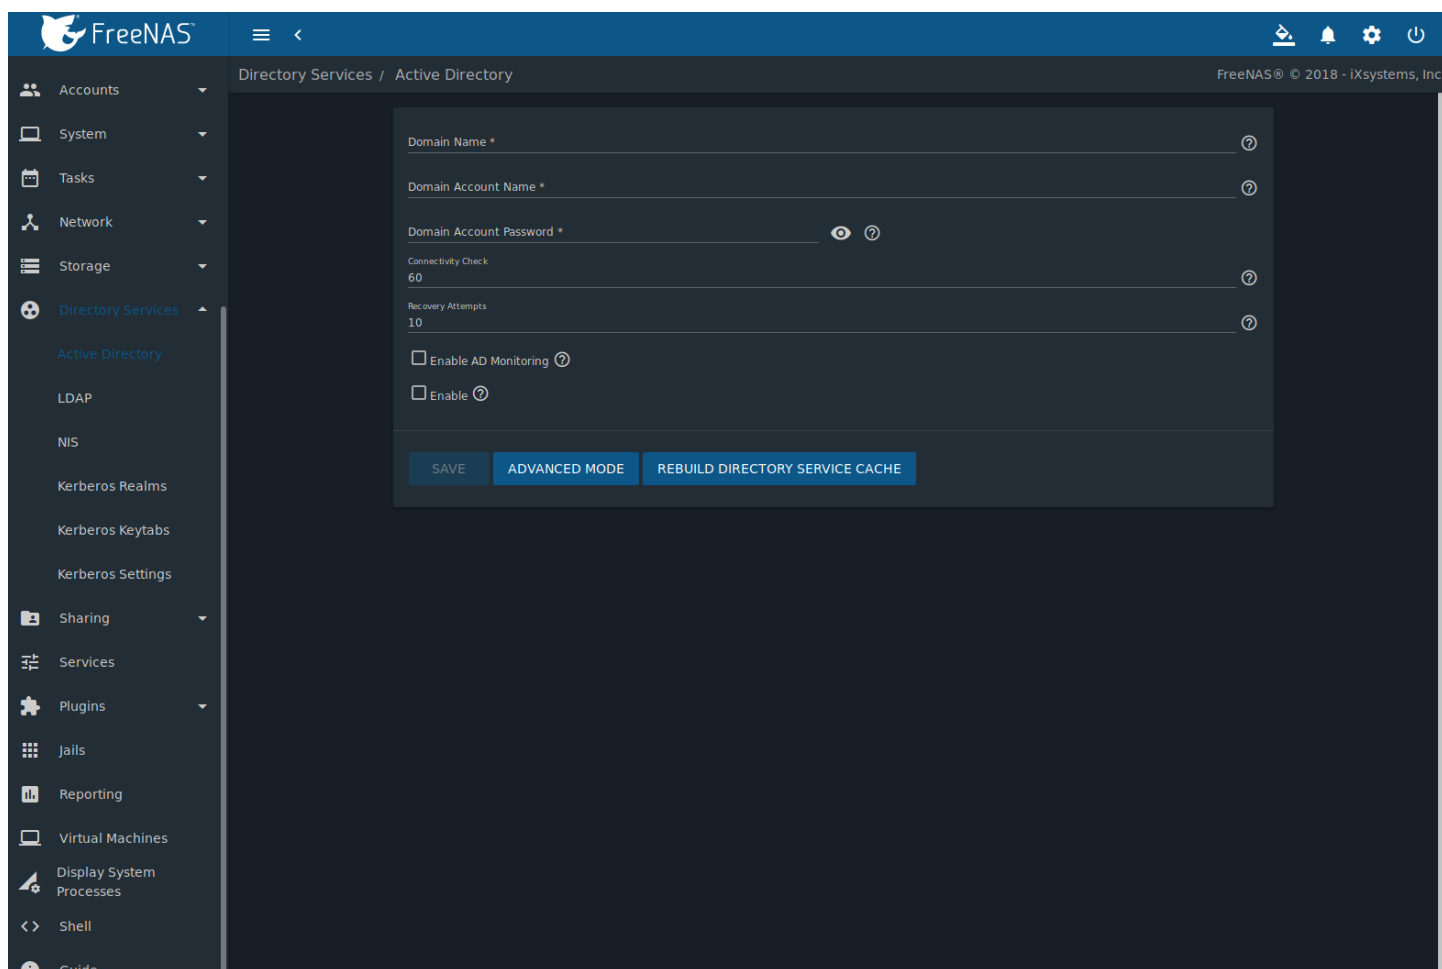


Fig. 10.1: Configuring Active Directory

Table 10.1 describes the configurable options. Some settings are only available in Advanced Mode. Click the *ADVANCED MODE* button to show the Advanced Mode settings. Go to *System* → *Advanced* and set the *Show advanced fields by default* option to always show advanced options.

Table 10.1: Active Directory Configuration Options

Setting	Value	Advanced Mode	Description
Domain Name	string		Name of the Active Directory domain (<i>example.com</i>) or child domain (<i>sales.example.com</i>). This field is mandatory. <i>Save</i> will be inactive until valid input is entered.
Domain Account Name	string		Name of the Active Directory administrator account. This field is mandatory. <i>Save</i> will be inactive until valid input is entered.
Domain Account Password	string		Password for the Active Directory administrator account. This field is mandatory. <i>Save</i> will be inactive until valid input is entered.
Connectivity Check	integer		How often for the system to verify Active Directory services are functioning. Enter a number of seconds.
Recovery Attempts	integer		Number of times to attempt reconnecting to the Active Directory server. Tries forever when set to 0.
Enable AD Monitoring	checkbox		Restart Active Directory automatically if the service disconnects.

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Encryption Mode	drop-down menu	✓	Choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i> .
Certificate	drop-down menu	✓	Select the certificate of the Active Directory server if SSL connections are used. If a certificate does not exist, create a Certificate Authority (page 97), then create a certificate on the Active Directory server. Import the certificate to the FreeNAS® system with the Certificates (page 101) menu.
Verbose logging	checkbox	✓	Set to log attempts to join the domain to <code>/var/log/messages</code> .
UNIX extensions	checkbox	✓	Only set if the AD server is explicitly configured to map permissions for UNIX users. Setting provides persistent UIDs and GIDs. Leave unset to map users and groups to the UID or GUID range configured in Samba.
Allow Trusted Domains	checkbox	✓	Only set when the network has active domain/forest trusts (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc757352(v=ws.10)) and managing file on multiple domains is required. Setting this option will generate more winbindd traffic and slow down filtering through user and group information.
Use Default Domain	checkbox	✓	Unset to prepend the domain name to the username. Unset to prevent name collisions when <i>Allow Trusted Domains</i> is set and multiple domains use the same username.
Allow DNS updates	checkbox	✓	Set to enable Samba to do DNS updates when joining a domain.
Disable FreeNAS Cache	checkbox	✓	Set to disable caching AD users and groups. This can help when unable to bind to a domain with a large number of users or groups.
Site Name	string	✓	The relative distinguished name of the site object in Active Directory.
Domain Controller	string	✓	Automatically be added to the SRV record for the domain. When multiple controllers are specified the FreeNAS® system selects the closest responding controller. Uses a short form of the FQDN. Example is <i>exampleserver</i> .
Global Catalog Server	string	✓	Ensure the hostname of the global catalog server to use is resolvable.
Kerberos Realm	drop-down menu	✓	Select the realm created using the instructions in Kerberos Realms (page 192).
Kerberos Principal	drop-down menu	✓	Browse to the location of the keytab created using the instructions in Kerberos Keytabs (page 193).
AD Timeout	integer	✓	Increase the number of seconds before timeout if the AD service does not immediately start after connecting to the domain.
DNS Timeout	integer	✓	Increase the number of seconds before a timeout occurs if AD DNS queries timeout.
Idmap backend	drop-down menu and Edit Idmap button	✓	Choose the backend to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. See Table 10.2 for a summary of the available backends. Click <i>Edit Idmap</i> to configure the selected backend.
Windbind NSS Info	drop-down menu	✓	Choose the schema to use when querying AD for user/group information. <i>rfc2307</i> uses the RFC2307 schema support included in Windows 2003 R2, <i>sfu</i> is for Services For Unix 3.0 or 3.5, and <i>sfu20</i> is for Services For Unix 2.0.

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Advanced Mode	Description
SASL wrapping	drop-down menu	✓	Choose how LDAP traffic is transmitted. Choices are <i>plain</i> (plain text), <i>sign</i> (signed only), or <i>seal</i> (signed and encrypted). Windows 2000 SP3 and newer can be configured to enforce signed LDAP connections.
Enable	checkbox		Set to enable the Active Directory service.
Netbios Name	string	✓	Limited to 15 characters. Automatically populated with the original hostname of the system. This must be different from the <i>Workgroup</i> name.
NetBIOS alias	string	✓	Limited to 15 characters.

Table 10.2 summarizes the backends which are available in the *ldmap 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/>) that 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 10.2: 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 <i>rid</i> , 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 or GIDs to use for user and group mappings and an optional size for the ranges.
fruit	Generate IDs as macOS does. The UID and GID can be identical on all FreeNAS® servers on the network. For use in <i>LDAP</i> (page 188) 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 is to 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, and limits the length of those names to 15 characters. If there are problems connecting to the realm, [verify](https://support.microsoft.com/en-us/help/909264/naming-conventions-in-active-directory-for-computers-domains-sites-and) (<https://support.microsoft.com/en-us/help/909264/naming-conventions-in-active-directory-for-computers-domains-sites-and>) the settings do not include any disallowed characters. The Administrator account password cannot contain the \$ character. If a \$ exists in the domain administrator 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

dataset.

The Active Directory users and groups that are imported to the FreeNAS® system are shown by typing commands in the FreeNAS® *Shell* (page 325):

- View users: `wbinfo -u`
- View groups: `wbinfo -g`

In addition, `wbinfo -t` tests the connection and, if successful, shows a message similar to:

```
checking the trust secret for domain YOURDOMAIN via RPC calls succeeded
```

To manually check that a specified user can authenticate, enter `net ads join -S dcname -U username`.

`getent passwd` and `getent group` can provide more troubleshooting information if no users or groups are listed in the output.

Tip: Sometimes network users do not appear in the drop-down menu of a *Permissions* screen but the `wbinfo` commands display these users. This is typically due to the FreeNAS® system taking longer than the default ten seconds to join Active Directory. Increase the value of *AD timeout* to 60 seconds.

To change a certificate, enable Advanced Mode, set the *Encryption Mode* to *Off*, then disable AD by unchecking *Enable*. Click *SAVE*. Select the new *Certificate*, set the *Encryption Mode* as desired, check *Enable* to re-enable AD, and click *SAVE* to restart AD.

10.1.1 Troubleshooting Tips

When running AD in a 2003/2008 mixed domain, [this forum post](https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) (<https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/>) has instructions 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 `host -t srv _ldap._tcp.domainname.com` to determine the SRV records of the network and 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 Tech-net article [How DNS Support for Active Directory Works](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759550(v=ws.10)) ([https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759550\(v=ws.10\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759550(v=ws.10))).

The realm used depends on the priority in the SRV DNS record. DNS can override the system Active Directory settings. When unable to connect to the correct realm, check the SRV records on the DNS server.

An expired password for the administrator account will cause `kinit` to fail. Ensure the password is still valid and double-check the password on the AD account being used does not include any spaces, 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 Organizational Unit (OU). When creating this entry, enter the FreeNAS® hostname in the *name* field. Make sure it is under 15 characters, the same name as the one set in the *Hostname* field in *Network* → *Global Configuration*, and the same *NetBIOS alias* in *Directory Service* → *Active Directory* → *Advanced* settings. Make sure the hostname of the domain controller is set in the *Domain Controller* field of *Directory Service* → *Active Directory* → *Advanced*.

10.1.2 If the System Does not Join the Domain

If the system will not join the Active Directory domain, run these commands in the order listed. `echo` commands will return a value of 0 and `klist` will show a Kerberos ticket:

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*.

Note: 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.

```
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 193):

```
service ix-sssd start
service sssd start
```

Finally, run these commands. `echo` returns a 0 unless something has gone wrong:

```
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 &
```

10.2 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 the 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](https://wiki.samba.org/index.php/4.1_smbldap-tools) (https://wiki.samba.org/index.php/4.1_smbldap-tools). The LDAP server must support SSL/TLS and the certificate for the LDAP server CA must be imported with *System* → *CAs* → *Import CA*. Non-CA certificates are not currently supported.

Tip: Apple's [Open Directory](https://manuals.info.apple.com/MANUALS/0/MA954/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) (https://manuals.info.apple.com/MANUALS/0/MA954/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) is an LDAP-compatible directory service into which FreeNAS® can be integrated. The forum post [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/>) has more information.

Figure 10.2 shows the LDAP Configuration section from *Directory Services* → *LDAP*.

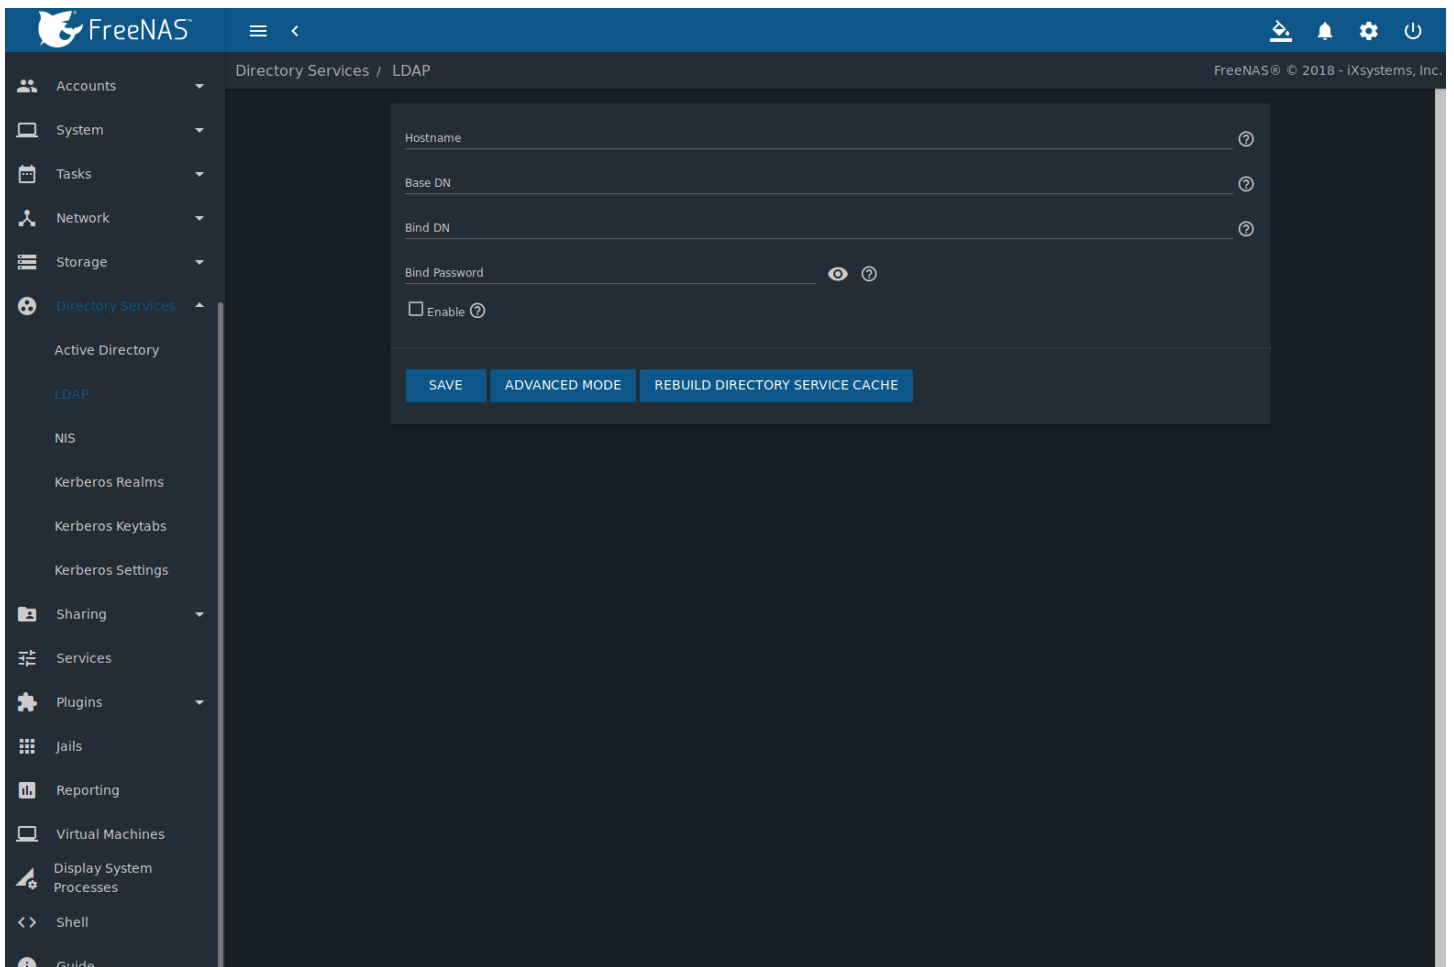


Fig. 10.2: Configuring LDAP

Table 10.3 summarizes the available configuration options. Some settings are only available in Advanced Mode. Click the *ADVANCED MODE* button to show the Advanced Mode settings. Go to *System* → *Advanced* and set the *Show advanced fields by default* option to always show advanced options.

Those new to LDAP terminology should read the [OpenLDAP Software 2.4 Administrator's Guide](http://www.openldap.org/doc/admin24/) (<http://www.openldap.org/doc/admin24/>).

Table 10.3: LDAP Configuration Options

Setting	Value	Advanced Mode	Description
Hostname	string		Hostname or IP address of the LDAP server.
Base DN	string		Top level of the LDAP directory tree to be used when searching for resources (Example: <i>dc=test,dc=org</i>).
Bind DN	string		Administrative account name on the LDAP server (Example: <i>cn=Manager,dc=test,dc=org</i>).
Bind Password	string		Password for the <i>Bind DN</i> . Click <i>SHOW/HIDE PASSWORDS</i> to view or obscure the password characters.
Allow Anonymous Binding	checkbox	✓	Instruct the LDAP server to disable authentication and allow read and write access to any client.
User Suffix	string	✓	Optional suffix to add to a name when the user account is added to the LDAP directory (Example: dept. company name).

Continued on next page

Table 10.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Group Suffix	string	✓	Optional suffix to add to a name when the group is added to the LDAP directory (Example: dept. or company name).
Password Suffix	string	✓	Optional suffix to add to the password when the password is added to the LDAP directory.
Machine Suffix	string	✓	Optional suffix to add to the name when the system is added to the LDAP directory (Example: server, accounting).
SUDO Suffix	string	✓	The suffix for LDAP-based users that need superuser access.
Kerberos Realm	drop-down menu	✓	The realm created using the instructions in Kerberos Realms (page 192).
Kerberos Principal	drop-down menu	✓	The location of the principal in the keytab created as described in Kerberos Keytabs (page 193).
Encryption Mode	drop-down menu	✓	Choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i> . Note: <i>SSL</i> or <i>TLS</i> and a <i>Certificate</i> must be selected for authentication to work.
Certificate	drop-down menu	✓	The LDAP CA certificate. The certificate for the LDAP server CA must first be imported using the <i>System</i> → <i>Certificates</i> menu. A certificate is required to use authentication
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	✓	The backend used to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. See Table 10.2 for a summary of the available backends. Click <i>EDIT IDMAP</i> to configure the selected backend.
Samba Schema	checkbox	✓	Set if LDAP authentication for SMB shares is required and the LDAP server is already configured with Samba attributes.
Auxiliary Parameters	string	✓	Additional options for sssd.conf(5) (https://jhrozek.fedorapeople.org/sssd/1.11.6/man/sssd.conf.5.html).
Schema	drop-down menu	✓	If <i>Samba Schema</i> is set, select the schema to use. Choices are <i>rfc2307</i> and <i>rfc2307bis</i> .
Enable	checkbox		Unset to disable the configuration without deleting it.
Netbios Name	string	✓	Limited to 15 characters. Automatically populated with the original hostname of the system. This must be different from the <i>Workgroup</i> name.
NetBIOS alias	string	✓	Limited to 15 characters.

Note: FreeNAS® automatically appends the root DN. This means the scope and root DN are not to 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 dataset after configuring the LDAP service. Type `getent passwd` in the FreeNAS® *Shell* (page 325) to verify the users have been imported. Type `getent group` to verify 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) (<http://www.openldap.org/doc/admin24/appendix-common-errors.html>) for common errors and how to fix them. When troubleshooting LDAP, open the FreeNAS® *Shell* (page 325) and look for error messages in `/var/log/auth.log`.

To clear LDAP users and groups from FreeNAS®, go to *Directory Services* → *LDAP*, clear the *Hostname* field, unset *Enable*, and click *SAVE*. Confirm LDAP users and groups are cleared by going to the *Shell* and viewing the output of the `getent passwd` and `getent group` commands.

10.3 NIS

The Network Information Service (NIS) maintains and distributes a central directory of Unix user and group information, host-names, email aliases, and other text-based tables of information. If an NIS server is running on the network, the FreeNAS® system can be configured to import the users and groups from the NIS directory.

Click the *Rebuild Directory Service Cache* button if a new NIS user needs immediate access to FreeNAS®. This occurs automatically once a day as a cron job.

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 10.3 shows the *Directory Services* → *NIS* section. Table 10.4 summarizes the configuration options.

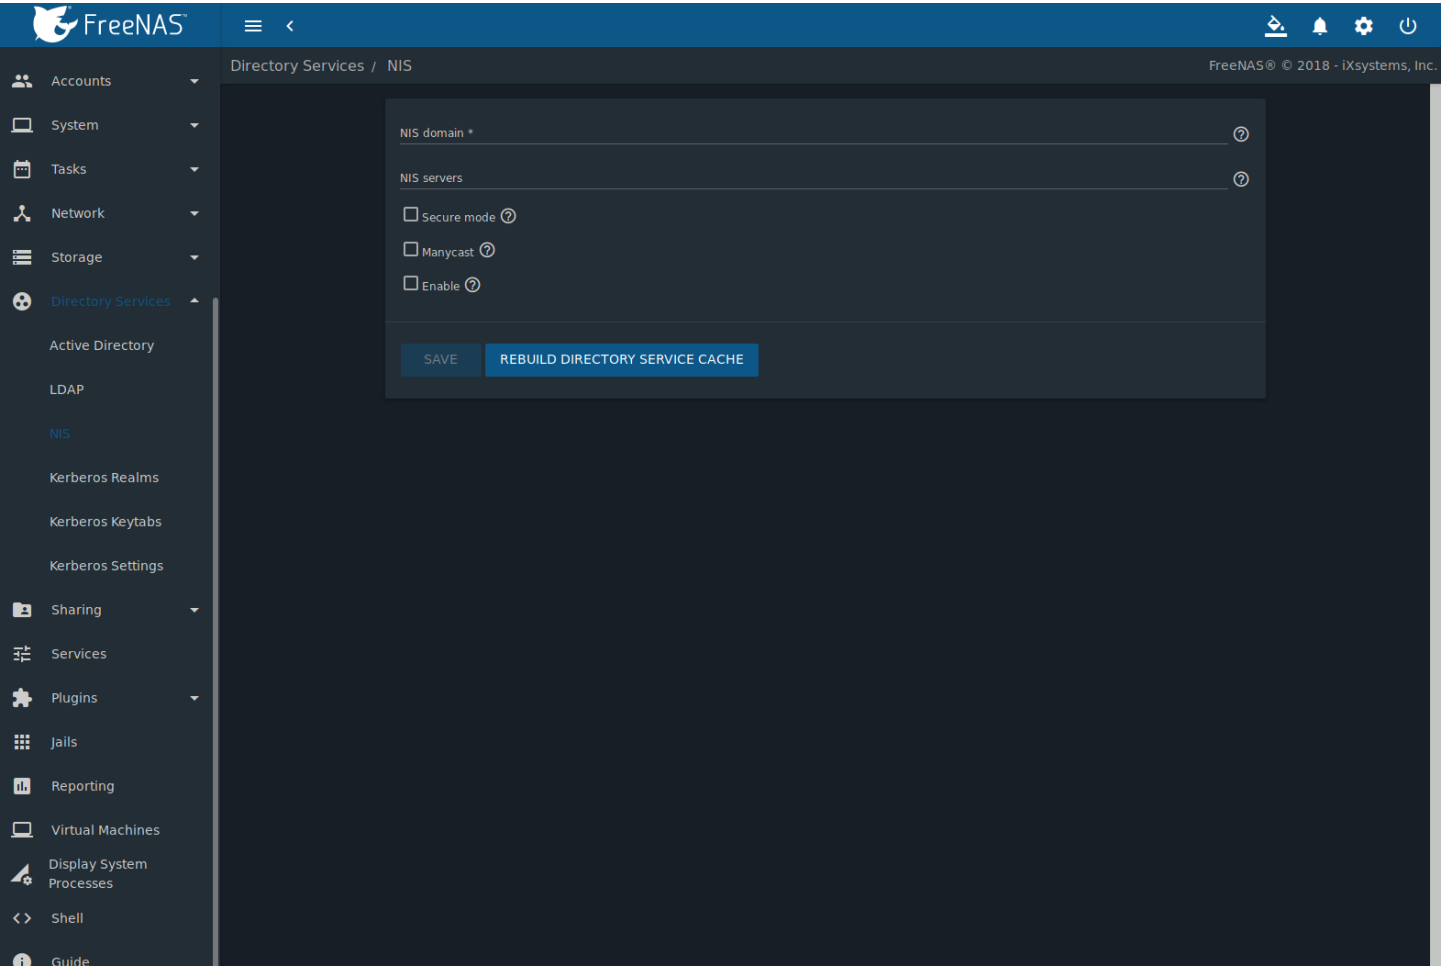


Fig. 10.3: NIS Configuration

Table 10.4: NIS Configuration Options

Setting	Value	Description
NIS domain	string	Name of NIS domain.
NIS servers	string	Comma-delimited list of hostnames or IP addresses.

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Description
Secure mode	checkbox	Set to have <code>ypbind(8)</code> (https://www.freebsd.org/cgi/man.cgi?query=ypbind) refuse to bind to any NIS server not running as root on a TCP port over 1024.
Manycast	checkbox	Set to have <code>ypbind</code> to bind to the server that responds the fastest. This is useful when no local NIS server is available on the same subnet.
Enable	checkbox	Unset to disable the configuration without deleting it.

10.4 Kerberos Realms

A default Kerberos realm is created for the local system in FreeNAS®. *Directory Services* → *Kerberos Realms* can be used to view and add Kerberos realms. If the network contains a Key Distribution Center (KDC), click *ADD* to add the realm. The configuration screen is shown in [Figure 10.4](#).

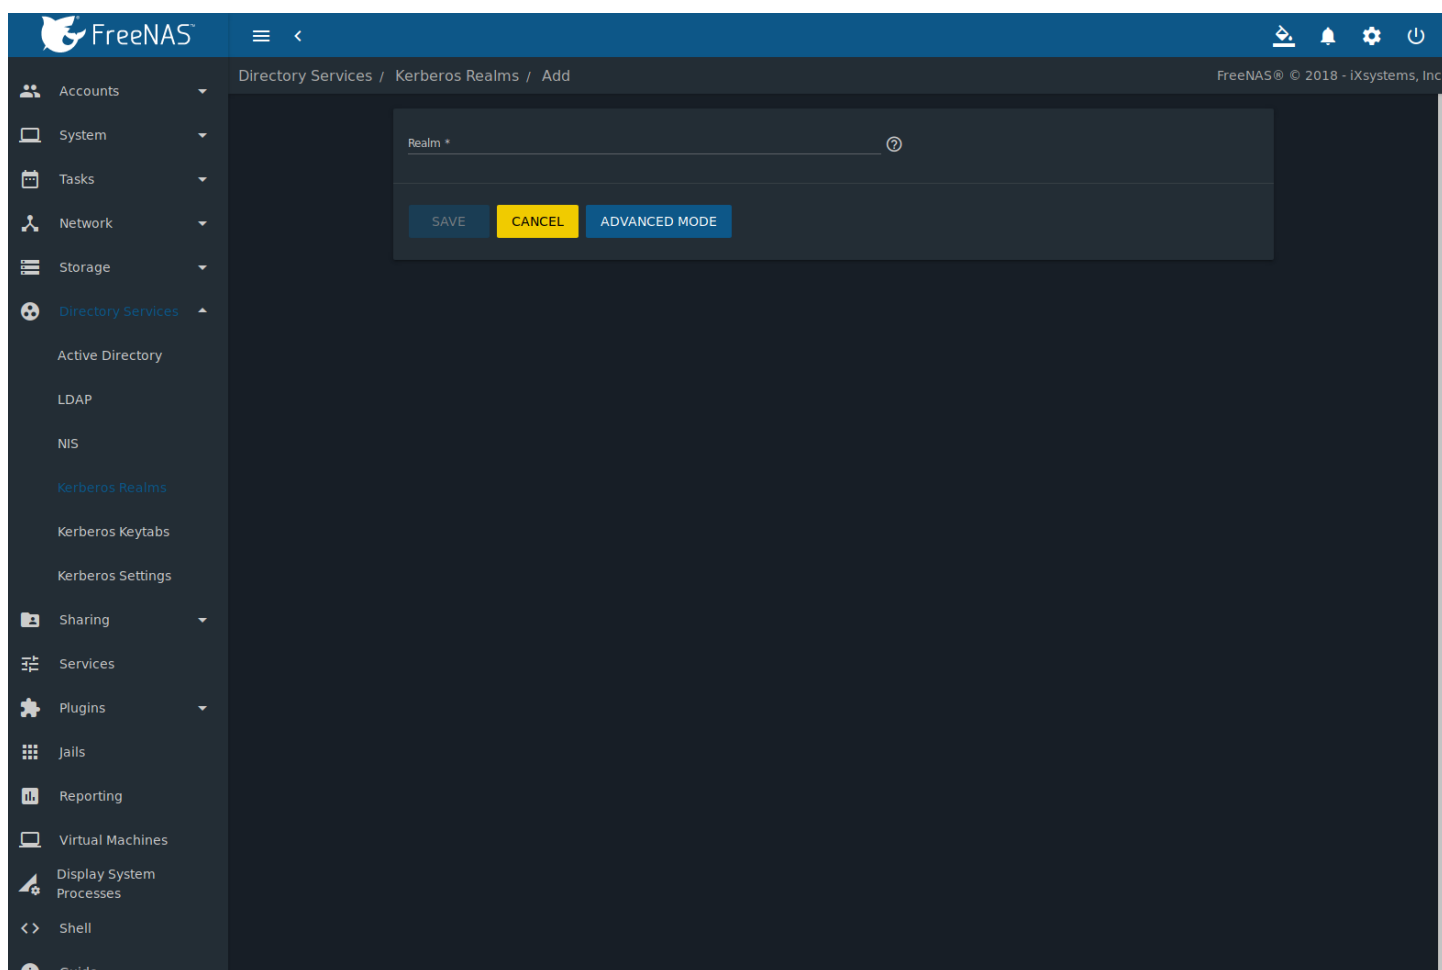


Fig. 10.4: Adding a Kerberos Realm

[Table 10.5](#) summarizes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click *ADVANCED MODE* or configure the system to always display these settings by setting *Show advanced fields by default* in *System* → *Advanced*.

Table 10.5: Kerberos Realm Options

Setting	Value	Advanced Mode	Description
Realm	string		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.

10.5 Kerberos Keytabs

Kerberos keytabs are used to do Active Directory or LDAP joins without a password. This means 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 the `ktpass` (<https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass>) command:

```
ktpass.exe /out freenas.keytab /princ http/useraccount@EXAMPLE.COM /mapuser useraccount /ptype KRB5_NT_
↪PRINCIPAL /crypto ALL /pass userpass
```

where:

- `freenas.keytab` is the file to upload to the FreeNAS® server.
- `useraccount` is the name of the user account for the FreeNAS® server generated in [Active Directory Users and Computers](https://technet.microsoft.com/en-us/library/aa998508(v=exchg.65).aspx) ([https://technet.microsoft.com/en-us/library/aa998508\(v=exchg.65\).aspx](https://technet.microsoft.com/en-us/library/aa998508(v=exchg.65).aspx)).
- `http/useraccount@EXAMPLE.COM` is the principal name written in the format *host/user.account@KERBEROS.REALM*. By convention, the kerberos realm is written in all caps, but make sure the case used for the *Kerberos Realm* (page 192) matches the realm name. See [this note](https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass#BKMK_remarks) (https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass#BKMK_remarks) about using `/princ` for more details.
- `userpass` is the password associated with `useraccount`.

Setting `/crypto` to `ALL` allows using all supported cryptographic types. These keys can be specified instead of `ALL`:

- `DES-CBC-CRC` is used for compatibility.
- `DES-CBC-MD5` adheres more closely to the MIT implementation and is used for compatibility.
- `RC4-HMAC-NT` uses 128-bit encryption.
- `AES256-SHA1` uses AES256-CTS-HMAC-SHA1-96 encryption.
- `AES128-SHA1` uses AES128-CTS-HMAC-SHA1-96 encryption.

This will create a keytab with sufficient privileges to grant tickets.

After the keytab is generated, add it to the FreeNAS® system using *Directory Services* → *Kerberos Keytabs* → *Add Kerberos Keytab*.

To instruct the Active Directory service to use the keytab, select the installed keytab using the drop-down *Kerberos Principal* menu in *Directory Services* → *Active Directory* Advanced Mode. When using a keytab with Active Directory, make sure that username and userpass in the keytab matches the Domain Account Name and Domain Account Password fields in *Directory Services* → *Active Directory*.

To instruct LDAP to use a principal from the keytab, select the principal from the drop-down *Kerberos Principal* menu in *Directory Services* → *LDAP* Advanced Mode.

10.6 Kerberos Settings

Configure additional Kerberos parameters in the *Directory Services* → *Kerberos Settings* section. Figure 10.5 shows the fields available:

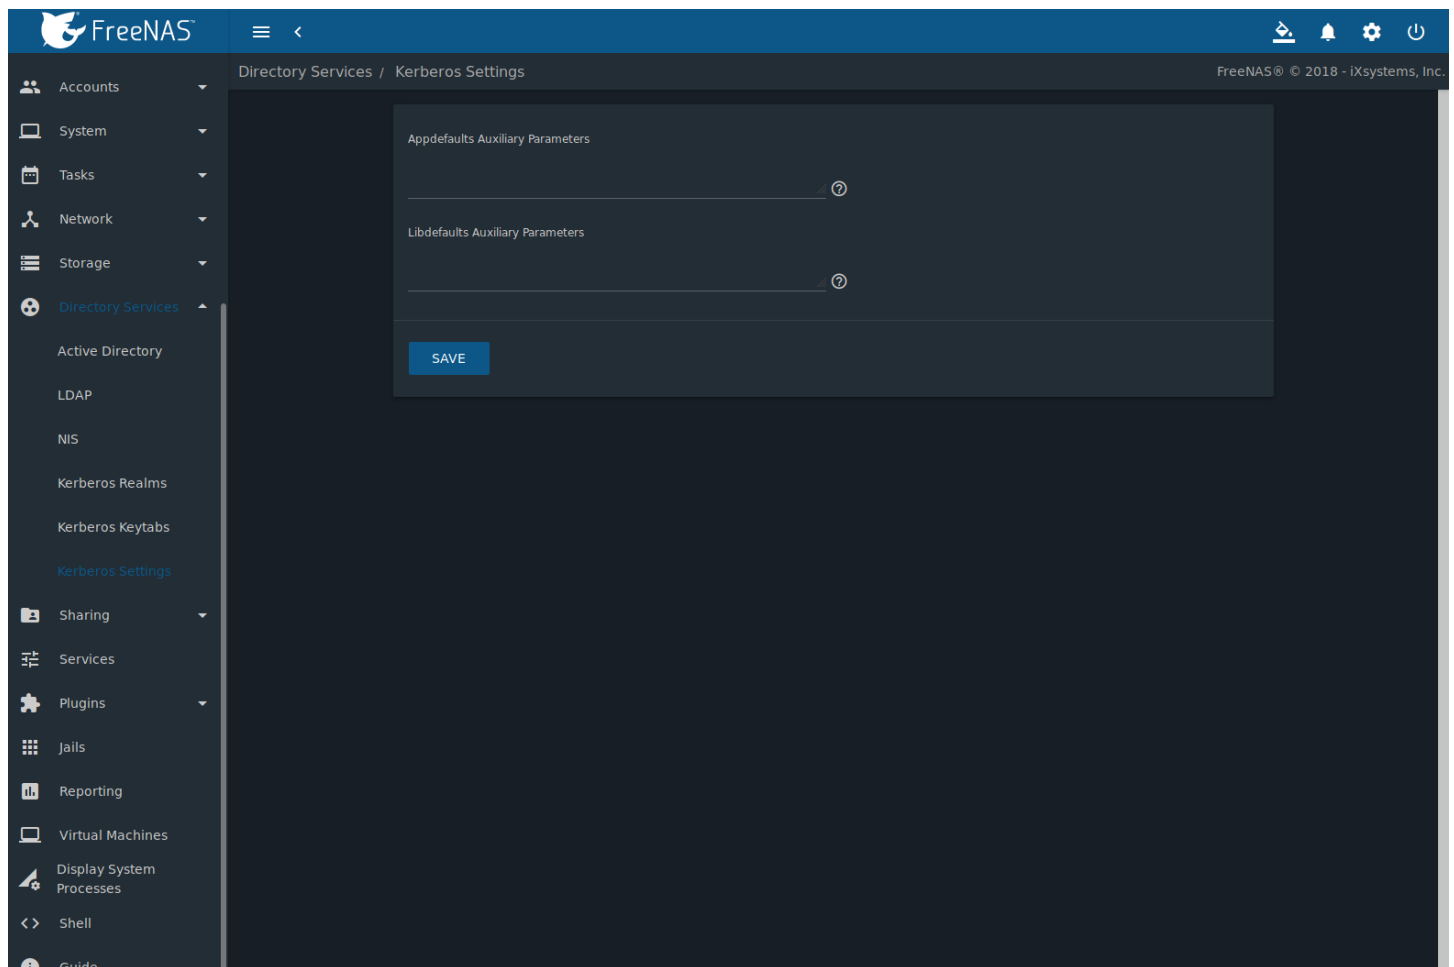


Fig. 10.5: Additional Kerberos Settings

- **Appdefaults Auxiliary Parameters:** Define any additional settings for use by some Kerberos applications. The available settings and syntax is 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:** Define any 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).

SHARING

Shares are created to make part or all of a pool 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.

Note: Shares are created to provide and control access to an area of storage. Before creating shares, making 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 is recommended. 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 196): Apple File Protocol shares are often used when the client computers all run macOS. Apple has slowly shifted to preferring [SMB](#) (page 211) for modern networks, although Time Machine still requires AFP.
- [Unix \(NFS\)](#) (page 203): Network File System shares are accessible from macOS, 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 210): WebDAV shares are accessible using an authenticated web browser (read-only) or [WebDAV client](#) (https://en.wikipedia.org/wiki/WebDAV#Client_support) running on any operating system.
- [SMB](#) (page 211): Server Message Block shares, also known as Common Internet File System (CIFS) shares, are accessible by Windows, macOS, 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 223): 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 243) 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](#) (<https://winscp.net/eng/index.php>), consider using the [SSH](#) (page 263) 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 pool 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 pool 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 pool 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 pool, and use that single type of share or service. To support multiple types of shares, divide the pool 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 237).

11.1 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. It then provides configuration examples for configuring Time Machine to back up to a dataset on the FreeNAS® system and for connecting to the share from a macOS client.

Create a share by clicking *Sharing* → *Apple (AFP)*, then *ADD*.

New AFP shares are visible in the *Sharing* → *Apple (AFP)* menu.

The configuration options shown in [Figure 11.1](#) appear after clicking ⓘ (Options) on an existing share, and selecting the *Edit* option. The values showing for these options will vary, depending upon the information given when the share was created.

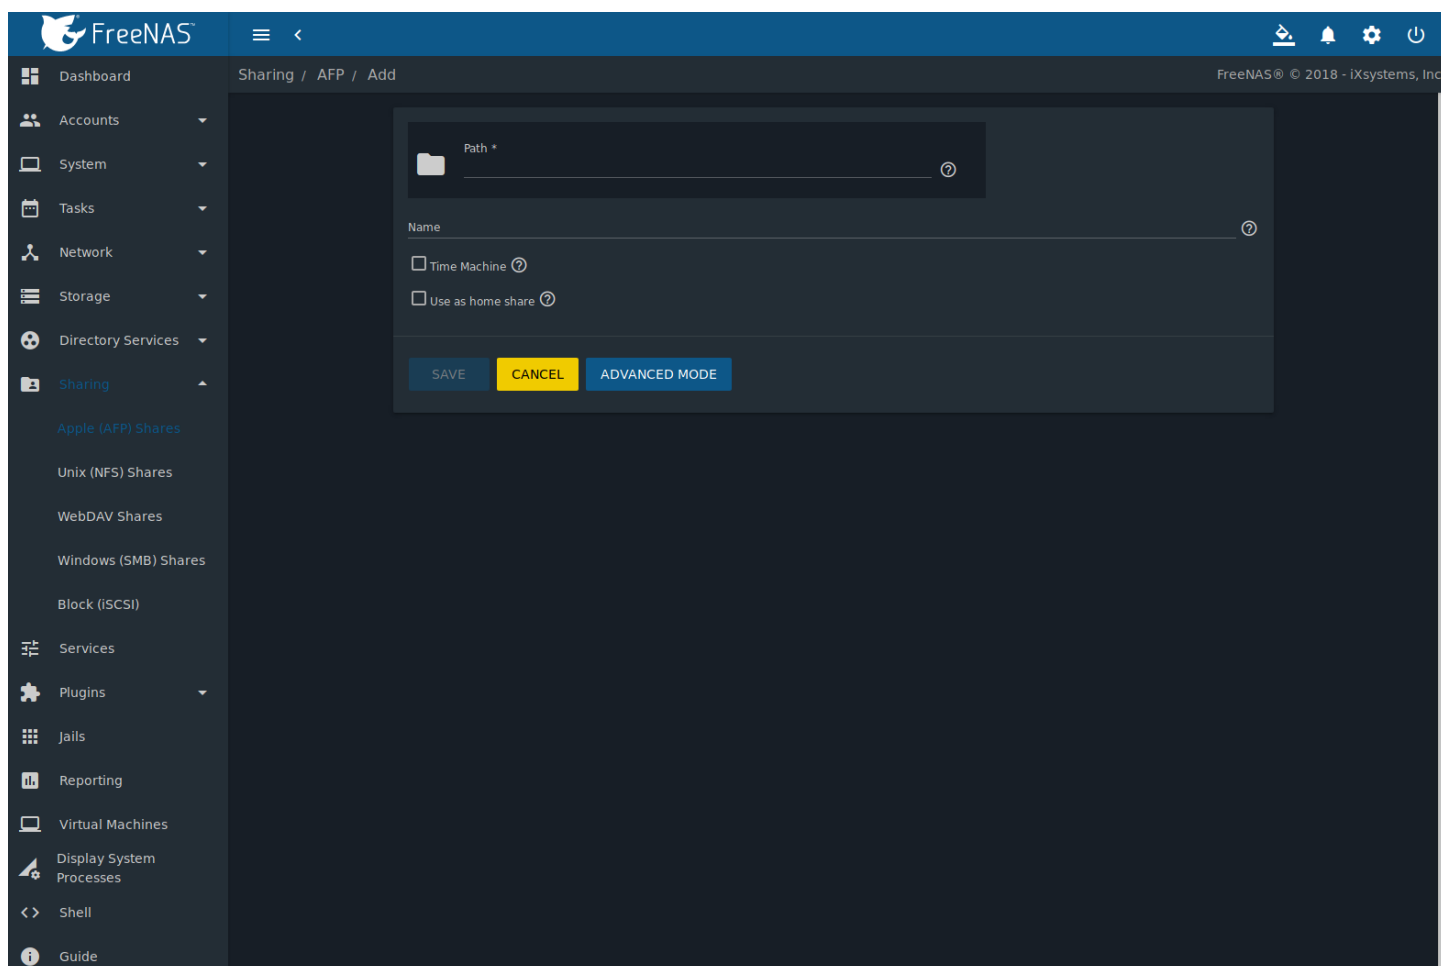


Fig. 11.1: Creating an AFP Share

Note: [Table 11.1](#) summarizes the options available to fine-tune an AFP share. Leaving these options at the default settings is recommended as changing them can 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 11.1: AFP Share Configuration Options

Setting	Value	Advanced Mode	Description
Path	browse button		Browse to the pool or dataset to share. Do not nest additional pools, datasets, or symbolic links beneath this path because Netatalk does not fully support that.
Name	string		Enter the pool name that appears in macOS after selecting <i>Go → Connect to server</i> in the Finder menu. Limited to 27 characters and cannot contain a period.
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		Set to advertise FreeNAS® as a Time Machine disk so it can be found by Macs. Setting multiple shares for Time Machine use is not recommended. When multiple Macs share the same pool, low disk space issues and intermittently failed backups can occur.
Use as home share	checkbox		Set to allow the share to host user home directories. Only one share can be used as the home share.
Zero Device Numbers	checkbox	✓	Enable when the device number is not constant across a reboot.
No Stat	checkbox	✓	If set, AFP does not stat the pool path when enumerating the pools list. Useful for automounting or pools created by a pre-exec script.
AFP3 UNIX Privs	checkbox	✓	Set to enable Unix privileges supported by OSX 10.5 and higher. Do not enable if the network contains Mac OS X 10.4 clients or lower as they do not support this feature.
Default file permissions	checkboxes	✓	Only works with Unix ACLs. New files created on the share are set with the selected permissions.
Default directory permissions	checkboxes	✓	Only works with Unix ACLs. New directories created on the share are set with the selected permissions.
Default umask	integer	✓	Umask is used for newly created files. Default is 000 (anyone can read, write, and execute).
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma, space, or tab.
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Separate entries with a comma, space, or tab.
Auxiliary Parameters	string		Enter any additional afp.conf (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered by other option fields.

11.1.1 Creating AFP Guest Shares

AFP supports guest logins, meaning that macOS 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 click the sliding button to turn on the service. Click ⓘ (Configure) to open the screen shown in [Figure 11.2](#). For *Guest Account*, use the drop-down to select *Nobody*, set *Guest Access*, and click *SAVE*.

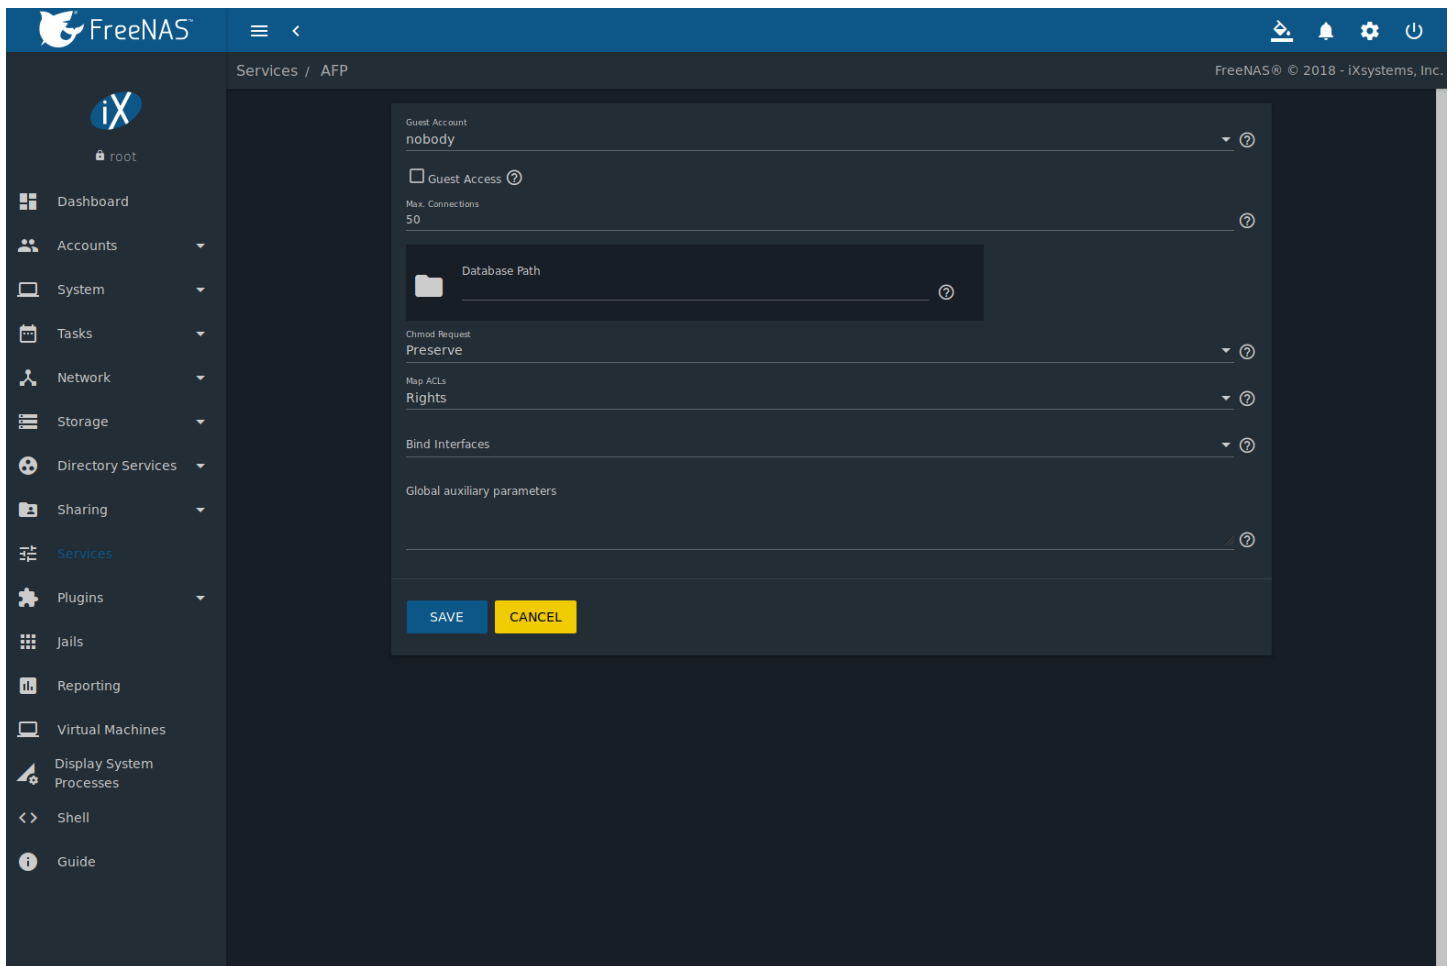


Fig. 11.2: Creating a Guest AFP Share

Next, create a dataset for the guest share. Refer to [Adding Datasets](#) (page 166) for more information about dataset creation. After creating the dataset for the guest share, go to *Storage* → *Pools*, click the ⓘ (Options) button for the dataset, then click *Edit Permissions*. Complete the fields shown in [Figure 11.3](#).

1. **ACL Type:** Select *Mac*.
2. **User:** Use the drop-down to select *Nobody*.
3. Click *SAVE*.

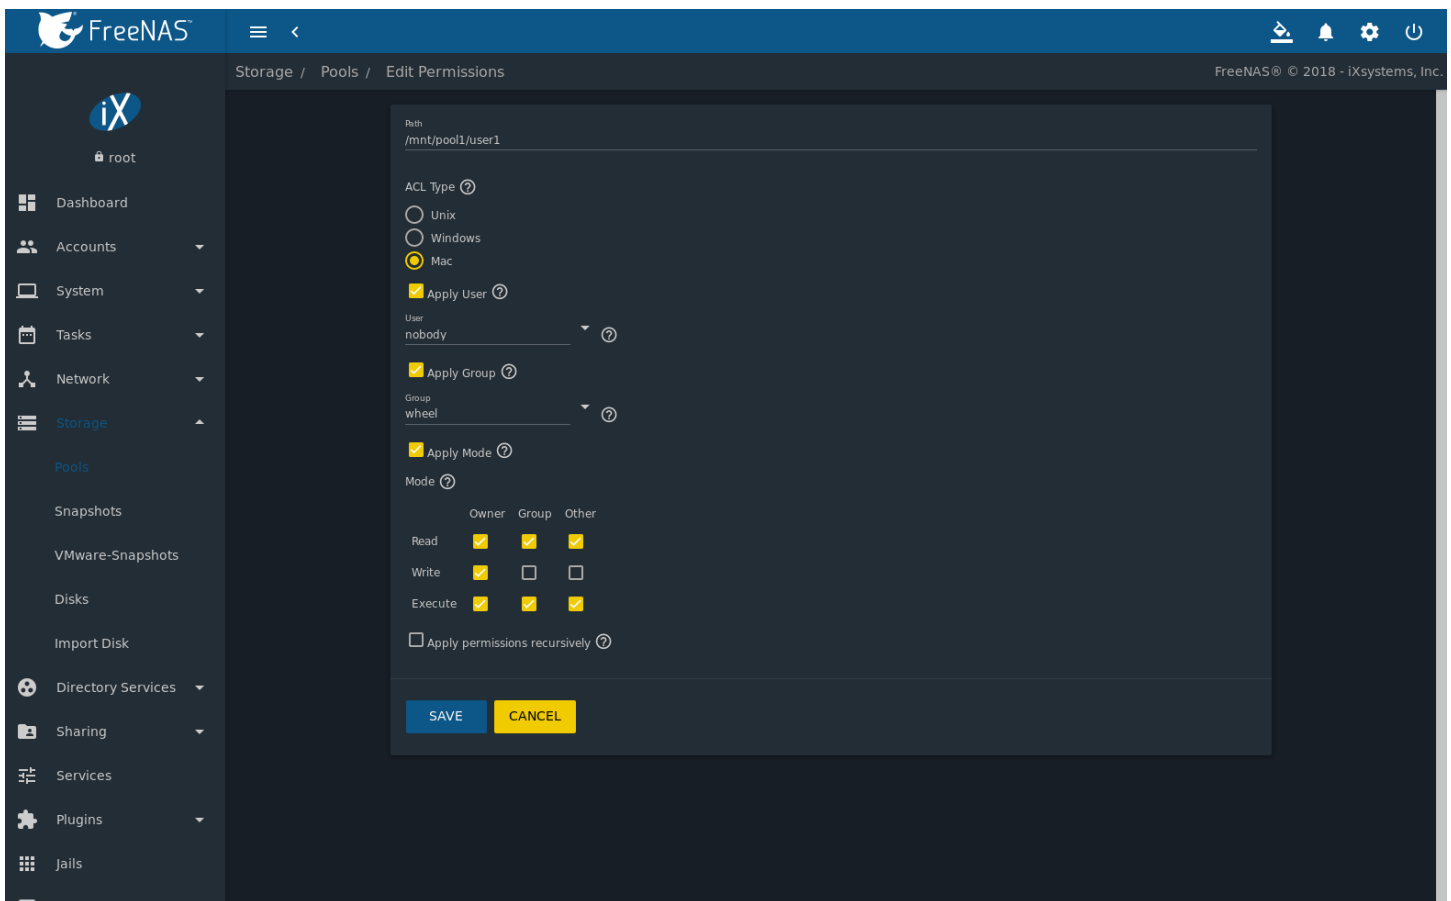


Fig. 11.3: Editing Dataset Permissions for Guest AFP Share

To create a guest AFP share:

1. Go to *Sharing* → *Apple (AFP) Shares* and click *ADD*.
2. *Browse* to the dataset created for the guest share.
3. Fill out the other required fields, then press *SAVE*.

macOS users can connect to the guest AFP share by clicking *Go* → *Connect to Server*. In the example shown in Figure 11.4, the user 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. 11.4: Connect to Server Dialog

To disconnect from the pool, click the *eject* button in the *Shared* sidebar.

11.1.2 Creating Authenticated and Time Machine Shares

macOS 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. Creating a separate Time Machine share for each user that will be using Time Machine to backup their macOS system to FreeNAS® is recommended. The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

Before creating a Time Machine or authenticated share, go to *Storage* → *Pools* to make a dataset for the share. For more information about dataset creation, refer to [Adding Datasets](#) (page 166).

After creating the dataset for the guest share, go to *Storage* → *Pools*, click the ⓘ (Options) button for the dataset, then *Edit Permissions*.

Enter the following information as shown in [Figure 11.5](#).

1. **ACL Type:** Select *Mac*.
2. **User:** Use the drop-down to select the desired user account. If the user does not yet exist on the FreeNAS® system, go to *Accounts* → *Users* to create one. Refer to [users](#) (page 67) in this guide for more information about creating a user.
3. **Group:** Use the drop-down to select the desired group name. If the group does not yet exist on the FreeNAS® system, go to *Accounts* → *Groups* to create one. Refer to [groups](#) (page 64) in this guide for more information about creating a group.
4. Click *SAVE*.

To create an authenticated or Time Machine share:

1. Go to *Sharing* → *AFP* and click *ADD*.
2. *Browse* to the dataset created for the share.
3. When creating a Time Machine share, enable the *Time Machine* option.
4. Fill out the other required fields.

5. Click **SAVE**.

To configure multiple authenticated or Time Machine shares, repeat this process for each user. The new shares appear in *Sharing* → *Apple (AFP)*.

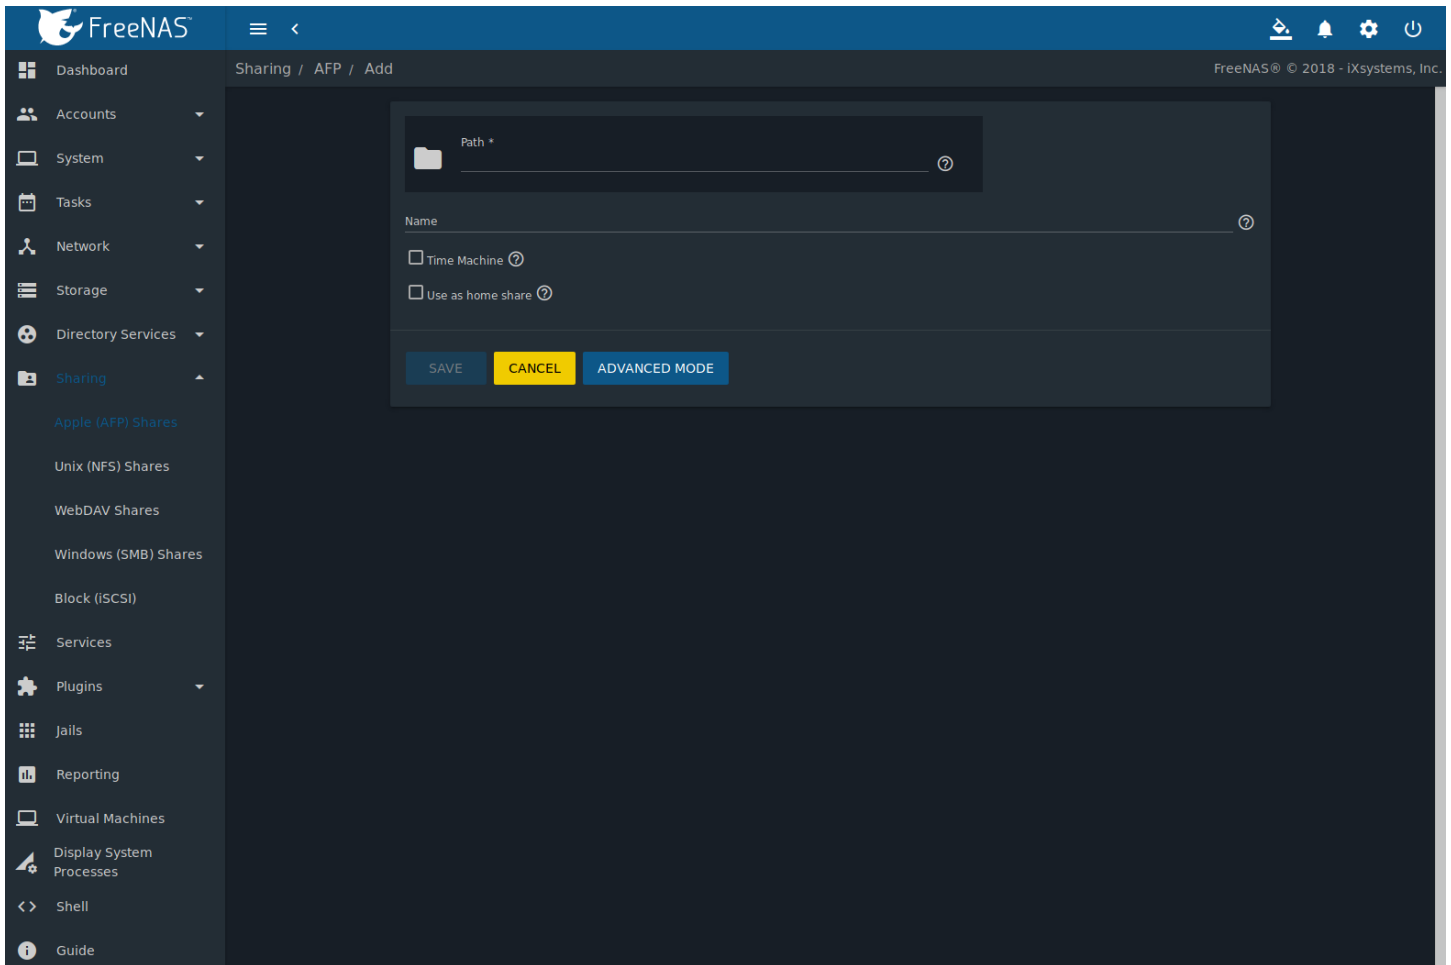


Fig. 11.5: Creating an Authenticated or Time Machine Share

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 macOS is ~21 GiB in size.

To configure a quota, go to *Sharing* → *Apple (AFP)*, click ⓘ (Options) on the existing Time Machine share, then *Edit*. In the example shown in Figure 11.6, the Time Machine share name is *backup_user1*. Enter a value in the *Time Machine Quota* field, then click **SAVE**. In this example, the Time Machine share is restricted to 200 GiB.

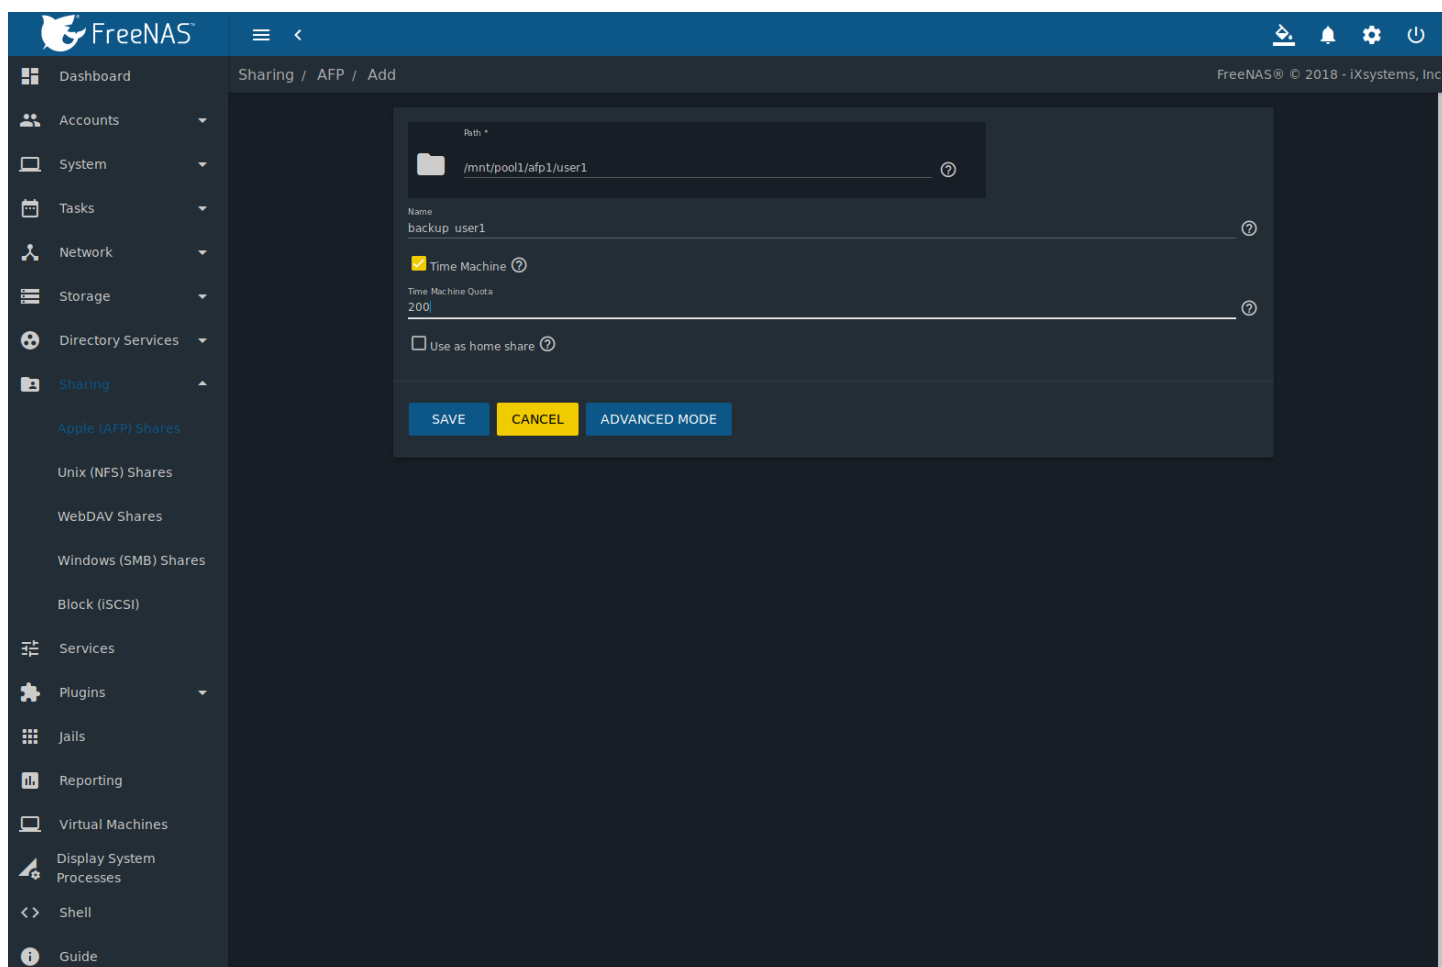


Fig. 11.6: Setting a Quota

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 macOS client, go to *System Preferences* → *Time Machine* which opens the screen shown in [Figure 11.7](#). Click *ON* and a pop-up menu shows the FreeNAS® system as a backup option. In this 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. 11.7: Configuring Time Machine on macOS

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](https://community.netgear.com/t5/Stora-Legacy/Solution-to-quot-Time-Machine-could-not-complete-the-backup/td-p/294697) (<https://community.netgear.com/t5/Stora-Legacy/Solution-to-quot-Time-Machine-could-not-complete-the-backup/td-p/294697>).

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.

11.2 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 presented in [Running ZFS over NFS as a VMware Store](https://tinyurl.com/archive-zfs-over-nfs-vmware) (<https://tinyurl.com/archive-zfs-over-nfs-vmware>).

Before creating an NFS share, go to *Services* → *NFS* and click the sliding button to turn on the service.

Next, go to *Storage* → *Pools* to create a dataset for the share. Refer to [Adding Datasets](#) (page 166) for more information about dataset creation.

Create an NFS share by going to *Sharing* → *Unix (NFS) Shares* and clicking *ADD*. [Figure 11.8](#) shows an example of creating an NFS share.

For the required *Path* field, browse to the dataset created for the share. In the *Comment* field enter a share name, note that if left empty, the share name is the list of selected Path entries. Click the *SAVE* button to create the share.

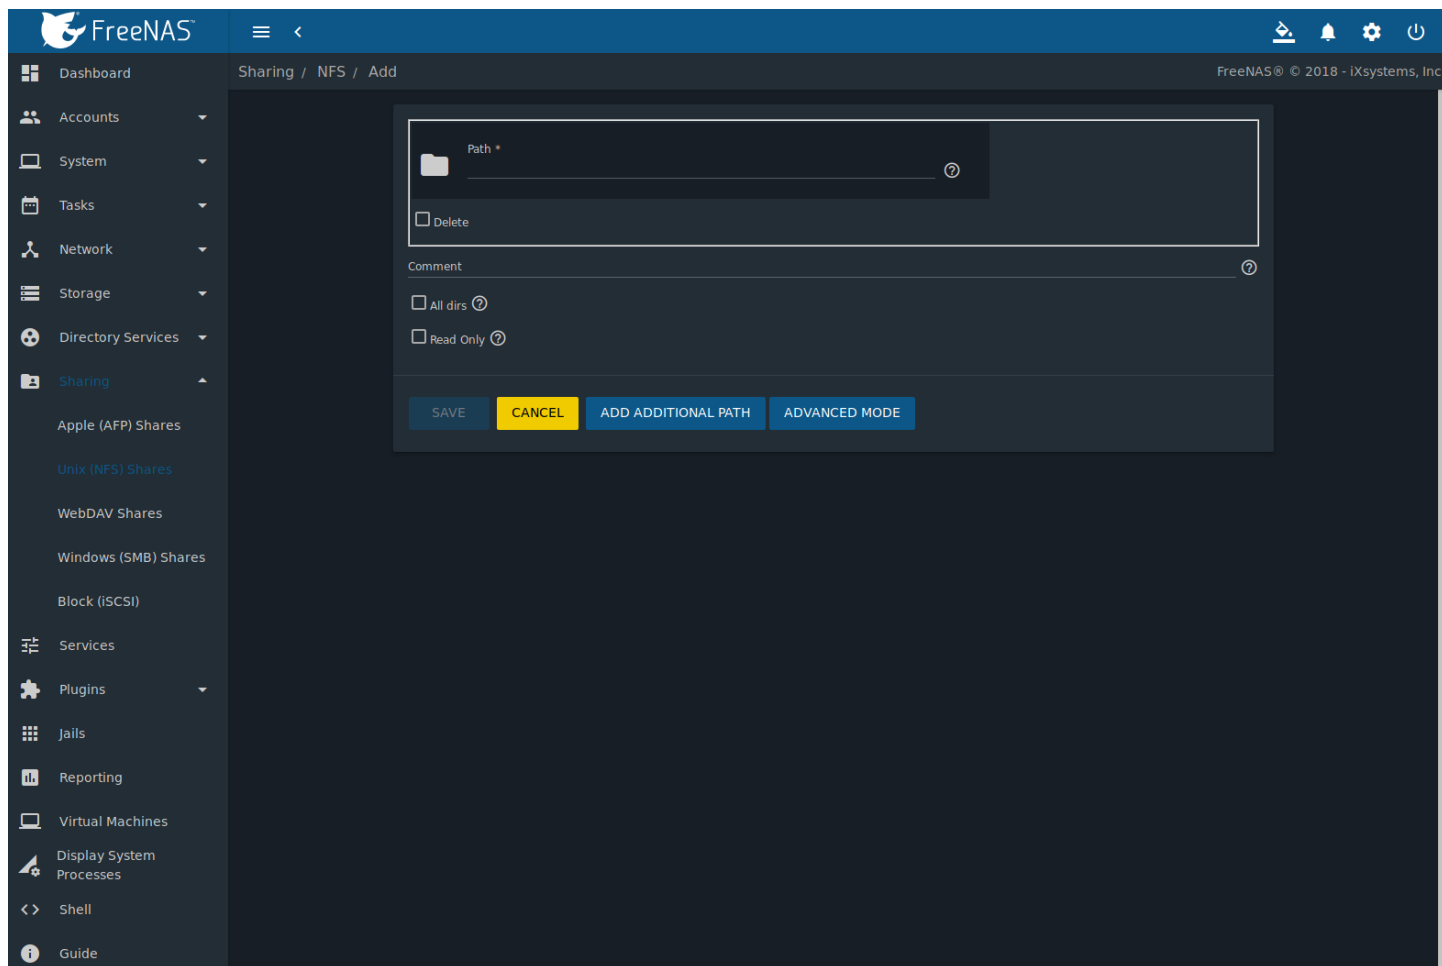



Fig. 11.8: NFS Share Creation

Edit an NFS share by going to *Sharing* → *Unix (NFS)*, clicking the  (Options) button for the desired share, then *Edit*. In the example shown in [Figure 11.9](#), the configuration screen is open for the *nfs_share1* share.

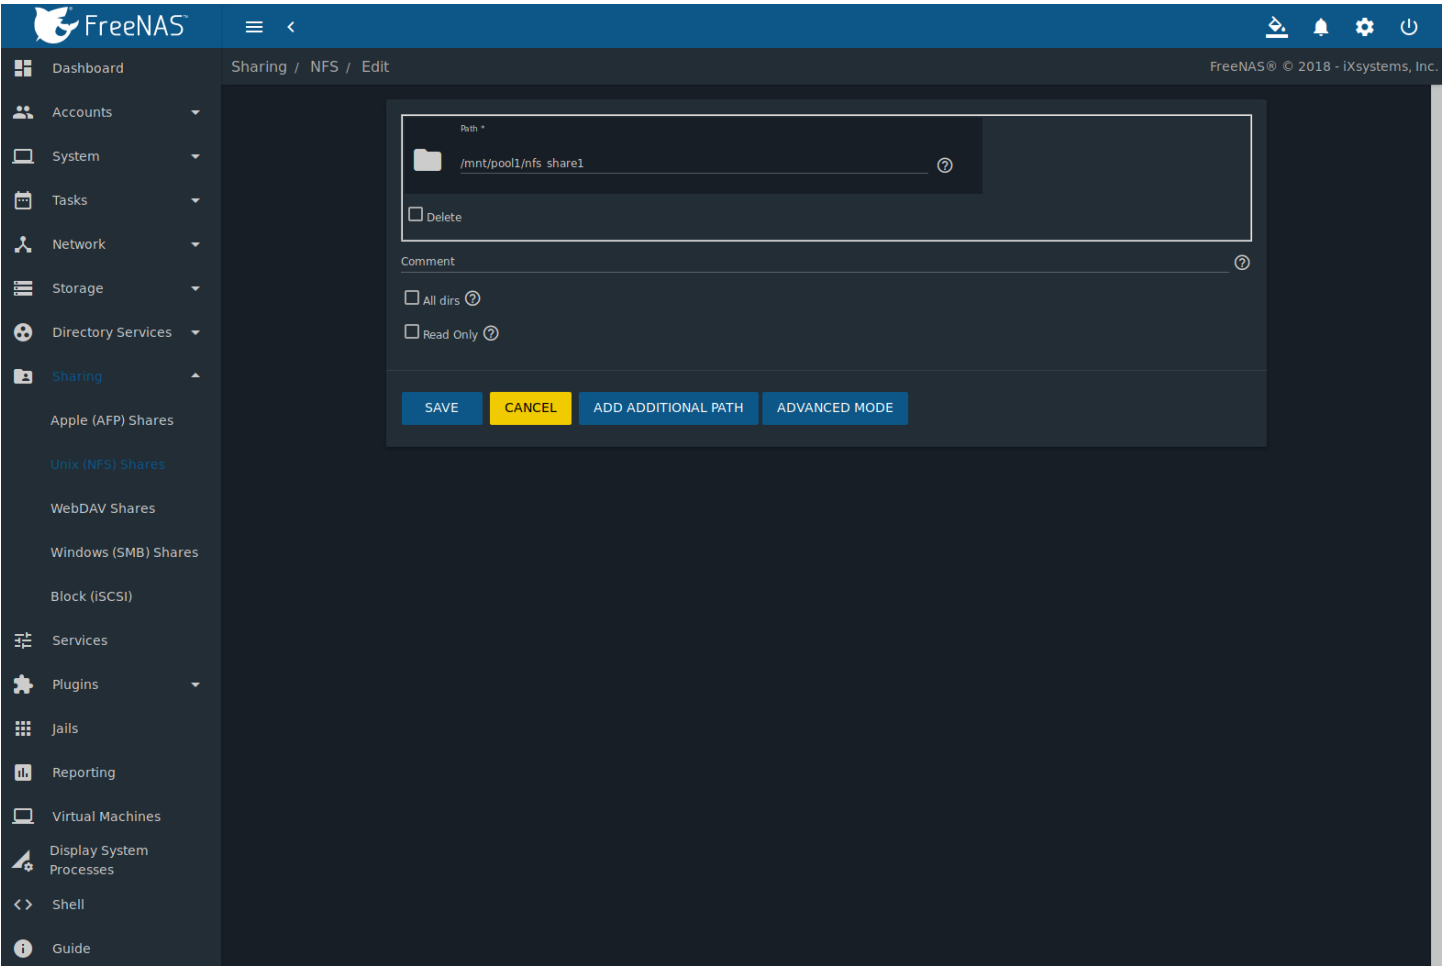


Fig. 11.9: NFS Share Settings

Table 11.2 summarizes the available configuration options in this screen. Some settings are only available by clicking the *ADVANCED MODE* button.

Table 11.2: NFS Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		The required full path to the pool or dataset to share. Click <i>ADD ADDITIONAL PATH</i> to configure multiple paths.
Comment	string		Set the share name. If left empty, share name is the list of selected <i>Path</i> entries.
All dirs	checkbox		Set to allow the client to mount any subdirectory within the <i>Path</i> .
Read Only	checkbox		Set to prohibit writing to the share.
Quiet	checkbox	✓	Set to inhibit some syslog diagnostics to avoid error messages. See exports(5) (https://www.freebsd.org/cgi/man.cgi?query=exports) for examples.
Authorized networks	string	✓	Space-delimited list of allowed networks in network/mask CIDR notation. Example: <i>1.2.3.0/24</i> . Leave empty to allow all.
Authorized Hosts and IP addresses	string	✓	Space-delimited list of allowed IP addresses or hostnames. Leave empty to allow all.

Continued on next page

Table 11.2 – continued from previous page

Setting	Value	Advanced Mode	Description
Maproot User	drop-down menu	✓	When a user is selected, the <i>root</i> user is limited to the permissions of that user.
Maproot Group	drop-down menu	✓	When a group is selected, the <i>root</i> user is also limited to the permissions of that group.
Mapall User	drop-down menu	✓	The specified permissions of that user are used by all clients.
Mapall Group	drop-down menu	✓	The specified permissions of that group are used by all clients.
Security	selection	✓	Only appears if <i>Enable NFSv4</i> is enabled in <i>Services</i> → <i>NFS</i> → <i>Configure</i> . Choices are <i>sys</i> or these Kerberos options: <i>krb5</i> (authentication only), <i>krb5i</i> (authentication and integrity), or <i>krb5p</i> (authentication and privacy).

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 as the web interface does not allow both. The *Mapall* options supersede the *Maproot* options. To restrict only the *root* user permissions, set the *Maproot* option. To restrict permissions of all users, set the *Mapall* options.
3. Each pool or dataset is considered to be its own filesystem and NFS is not able to cross filesystem boundaries.
4. The network and host must be unique per share and per filesystem or directory. Since `/etc/exports` does not act like an ACL, the rule to apply is undefined among overlapping networks or when using the same share with multiple hosts.
5. The *All dirs* 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 pool named `pool1` 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/pool1/dataset1` and `/mnt/pool1/dataset1/directory1`

Instead, set a *Path* of `/mnt/pool1/dataset1` and check the *All dirs* 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/pool1/dataset1`

Second NFS share:

- *Authorized Networks* set to `20.0.0.0/8`
- *Path* set to `/mnt/pool1/dataset1/directory1`

Note that this requires the creation of two shares. It cannot be done with only one share.

11.2.1 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 pool or dataset that is being shared, change the owner and group to *nobody* and set the permissions according to the desired 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 pool or dataset. For example, even if the *root* user is able to connect, it will not gain *root* access to the share.

11.2.2 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/pool1/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/pool1/nfs_share1*
 - *Authorized Networks*: *192.168.2.0/24*
 - *All dirs* option is enabled
 - *MapAll User* is set to *nobody*
 - *MapAll Group* is set to *nobody*

11.2.2.1 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/pool1/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/pool/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* (<https://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/pool1/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
```

11.2.2.2 From Microsoft

Windows NFS client support varies with versions and releases. For best results, use [Windows \(SMB\) Shares](#) (page 211).

11.2.2.3 From macOS

To mount the NFS share from a macOS client, go to *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 pool or dataset being shared by NFS. The example shown in [Figure 11.10](#) continues with the example of *192.168.2.2:/mnt/pool1/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 11.11](#), */mnt/data* has one folder named *images*. The user can now copy files to and from the share.

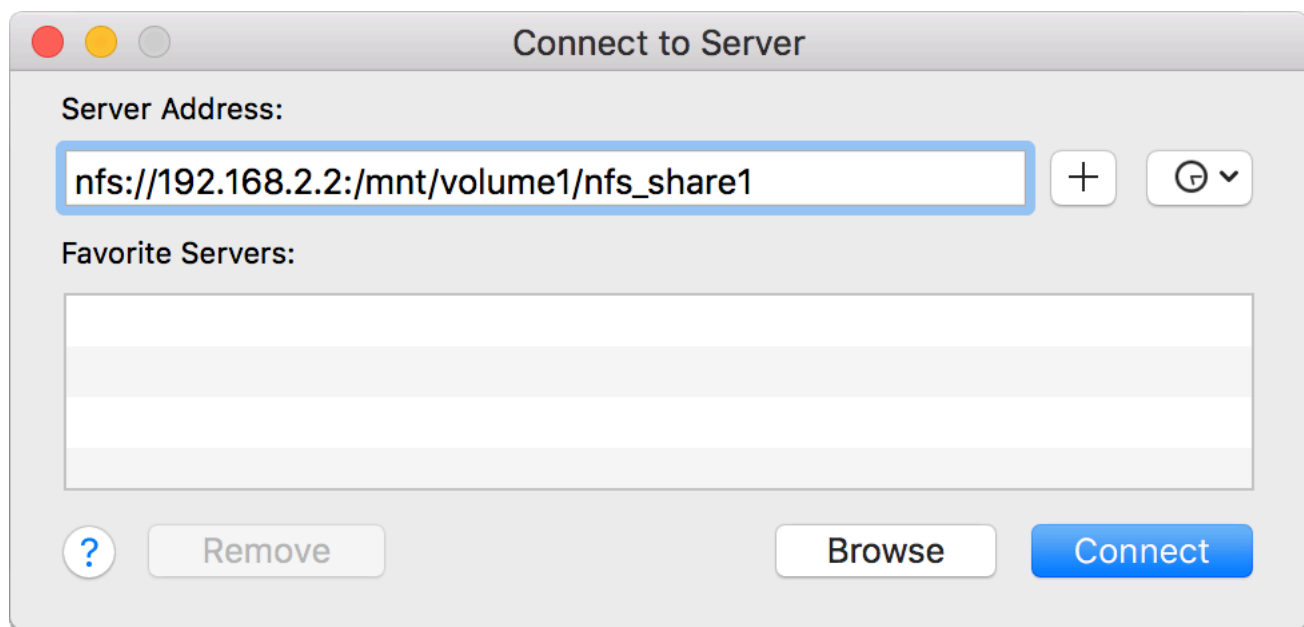


Fig. 11.10: Mounting the NFS Share from macOS




Fig. 11.11: Viewing the NFS Share in Finder

11.2.3 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 database* 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 → Configure* and enable the *Serve UDP NFS clients* option.

The `nfsstat -c` or `nfsstat -s` commands can be helpful to detect problems from the [Shell](#) (page 325). A high proportion of retries and timeouts compared to reads usually indicates network problems.

11.3 WebDAV Shares

In FreeNAS®, WebDAV shares can be created so that authenticated users can browse the contents of the specified pool, 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* → *CONFIGURE*.
- **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 firewall only allows access to authorized systems.
- **port_number:** is configured in *Services* → *WebDAV* → *CONFIGURE*. If the FreeNAS® system is to be accessed using a public IP address, consider changing the default port number and ensure that the network firewall only allows access to authorized systems.
- **share_name:** is configured by clicking *Sharing* → *WebDAV Shares*, then *ADD*.

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* → *CONFIGURE*.

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, go to *Sharing* → *WebDAV Shares* and click *ADD*, which will open the screen shown in [Figure 11.12](#).

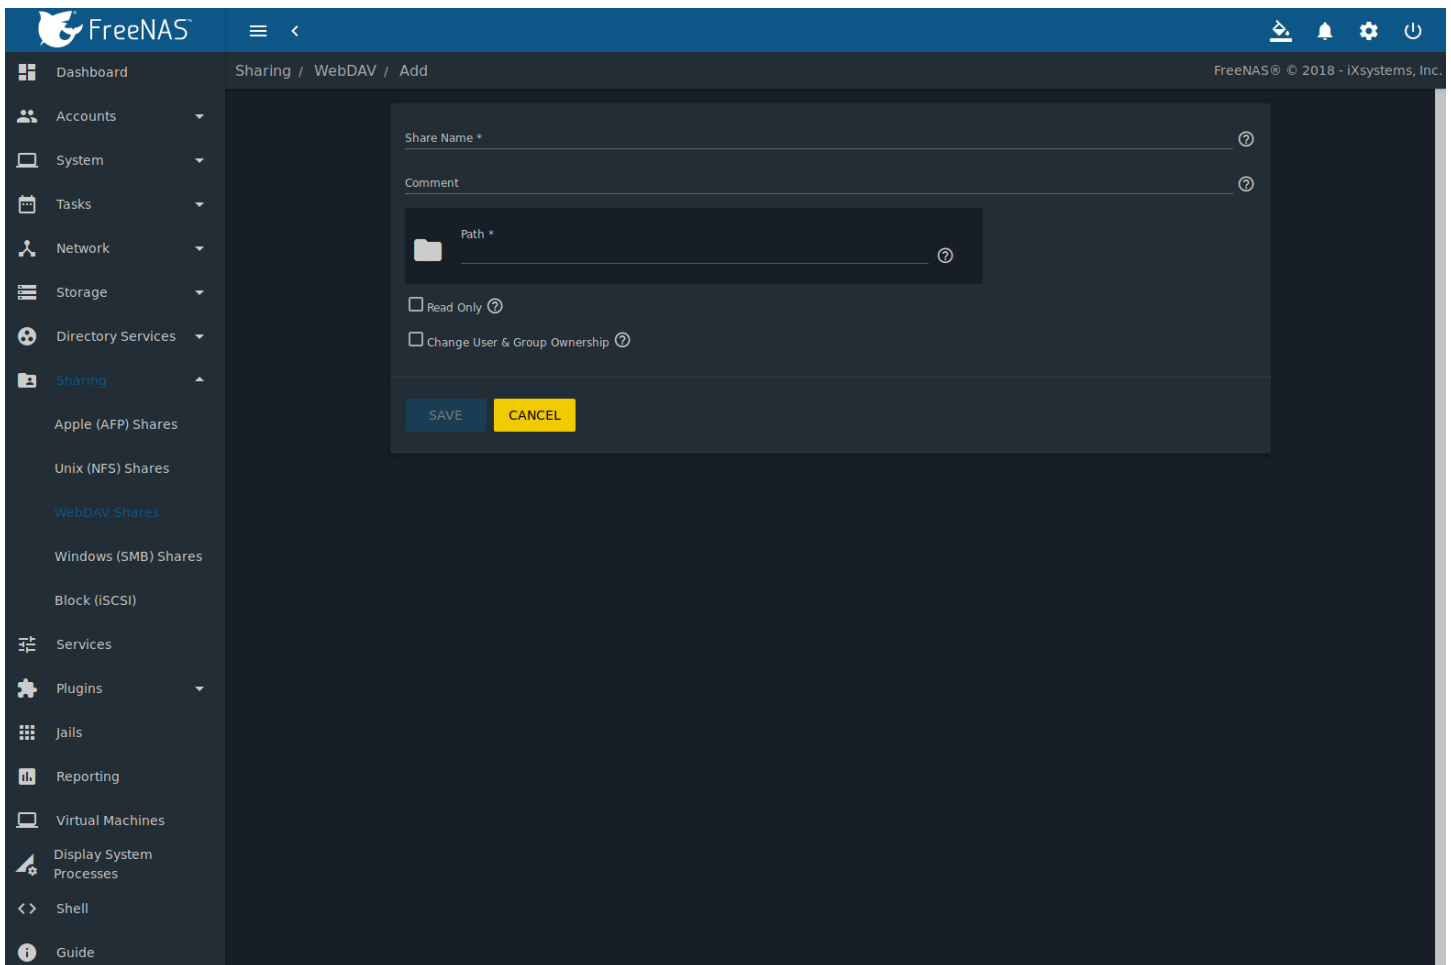


Fig. 11.12: Adding a WebDAV Share

Table 11.3 summarizes the available options.

Table 11.3: WebDAV Share Options

Setting	Value	Description
Share Name	string	Enter a name for the share.
Comment	string	Optional.
Path	browse button	Browse to the pool or dataset to share.
Read Only	checkbox	Set to prohibit users from writing to the share.
Change User & Group Ownership	checkbox	Enable to automatically set the share contents to the <i>webdav</i> user and group.

Click **SAVE** to create the share. Then, go to *Services* → *WebDAV* and click the  (Power) button to turn on the service.

After the service starts, review the settings in *Services* → *WebDAV* → *CONFIGURE* 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 269).

11.4 Windows (SMB) Shares

FreeNAS® uses [Samba](https://www.samba.org/) (<https://www.samba.org/>) to share pools using Microsoft's SMB protocol. SMB is built into the Windows and macOS operating systems and most Linux and BSD systems pre-install the Samba client in order to provide support

for SMB. If the distro did not, install the Samba client using the distro software repository.

The SMB protocol supports many different types of configuration scenarios, ranging from the simple to 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://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc733145(v=ws.11))) (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc733145(v=ws.11)).

This chapter starts by summarizing the available configuration options. It demonstrates some common configuration scenarios as well as offering some troubleshooting tips. Reading through this entire chapter before creating any SMB shares is recommended to gain a better understanding of the configuration scenario that meets the specific network requirements.

Warning: SMB1 is disabled by default for security. If legacy clients are unable to connect to the share, open [Shell](#) (page 325), type `sysctl freenas.services.smb.config.server_min_protocol=NT1`, then restart the [SMB](#) (page 258) service. If legacy clients can then connect to the share, the change can be made permanent by going to [Tunables](#) (page 91), creating a tunable with a *Variable* of `freenas.services.smb.config.server_min_protocol`, a *Value* of `NT1`, and a *Type* of `Sysctl`.

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 325) for a list of active connections and users.

Figure 11.13 shows the configuration screen that appears after clicking *Sharing* → *Windows (SMB Shares)*, then *ADD*.

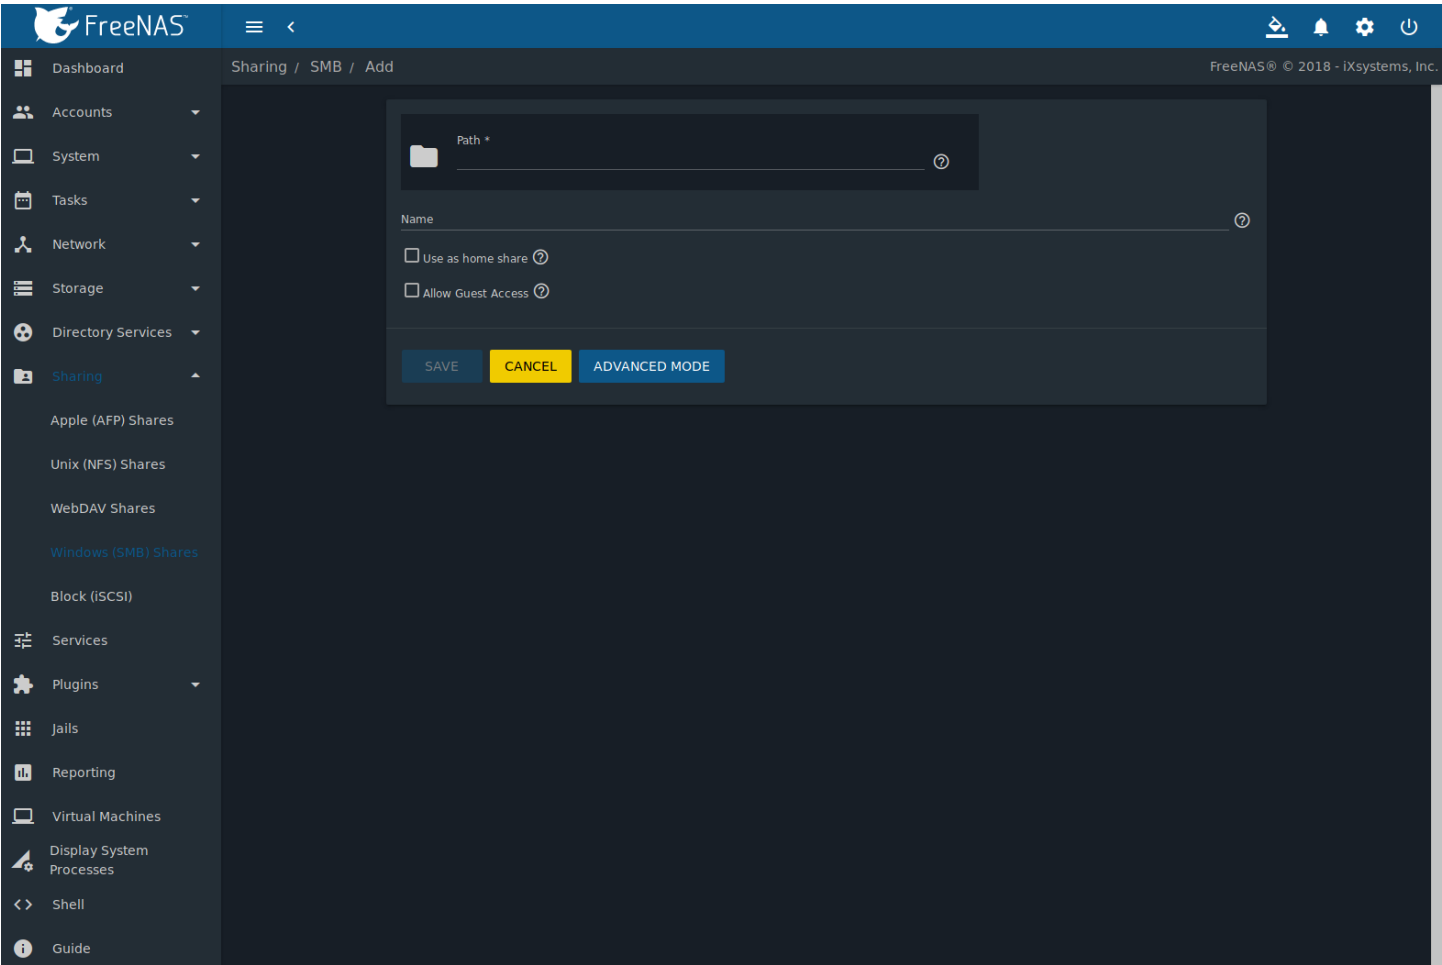


Fig. 11.13: Adding an SMB Share

Table 11.4 summarizes the options available when creating a SMB share. Some settings are only configurable 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(5)) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf>) provides more details for each configurable option.

Table 11.4: Options for a SMB Share

Setting	Value	Advanced Mode	Description
Path	browse button		Select pool or dataset/directory to share.
Name	string		Enter a mandatory name for the share.
Use as home share	checkbox		Set to allow this share to hold user home directories. Only one share can be the home share. Note that lower case names for user home directories are strongly recommended, as Samba maps usernames to all lower case. For example, the username John will be mapped to a home directory named john. If the <i>Path</i> to the home share includes an upper case username, delete the existing user and recreate it in <i>Accounts</i> → <i>Users</i> with an all lower case <i>Username</i> . Return to <i>Sharing</i> → <i>SMB</i> to create the home share, and select the <i>Path</i> that contains the new lower case username.

Continued on next page

Table 11.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Default Permissions	checkbox		When enabled, the ACLs grant read and write for owner or group and read-only for others. Only leave unset if creating a share on a system that already has custom ACLs configured.
Export Read Only	checkbox	✓	Set to prohibit write access to this share.
Browsable to Network Clients	checkbox	✓	When set, users see the contents of <i>/homes</i> , which includes the home directories of other users. When unset, users only see their own home directory.
Export Recycle Bin	checkbox	✓	When set, deleted files are moved to a hidden <i>.recycle</i> in the root folder of the share. The <i>.recycle</i> directory can be deleted to reclaim space and is automatically recreated when a file is deleted.
Show Hidden Files	checkbox	✓	Set to disable the Windows <i>hidden</i> attribute on a new Unix hidden file. Unix hidden filenames start with a dot: <i>.foo</i> . Existing files are not affected.
Allow Guest Access	checkbox		Set to allow access to this share without a password. See SMB (page 258) service for more information about guest user permissions.
Only Allow Guest Access	checkbox	✓	Requires <i>Allow guest access</i> to also be enabled. Forces guest access for all connections.
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma, space, or tab.
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Allowed hosts take Specify <i>ALL</i> and list any hosts from <i>Hosts Allow</i> to have those hosts take precedence.
VFS Objects	selection	✓	Adds virtual file system modules to enhance functionality. Table 11.5 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 <i>smb4.conf</i> parameters not covered by other option fields.

Here are some notes about *ADVANCED MODE* settings:

- Hostname lookups add some time to accessing the SMB share. If only using IP addresses, unset the *Hostnames Lookups* setting in *Services* → *SMB* → *i□□* (Configure).
- When the *Browsable to Network Clients* option is selected, the share is visible through Windows File Explorer or through *net view*. When the *Use as home share* option is selected, deselecting the *Browsable to Network Clients* option hides the share named *homes* so that only the dynamically generated share containing the authenticated user home directory will be visible. By default, the *homes* share and the user 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 pool 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>).

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 selecting the *NTLMv1 auth* option in *Services* → *SMB* → *i□□* (Configure).

Table 11.5 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/old/Samba3-HOWTO/VFS.html) (https://www.samba.org/samba/docs/old/Samba3-HOWTO/VFS.html) and the [vfs_* man pages](https://www.samba.org/samba/docs/current/man-html/) (https://www.samba.org/samba/docs/current/man-html/) for more details.

Table 11.5: 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 <i>audit</i> 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 macOS support by providing the SMB2 AAPL extension and Netatalk interoperability. Automatically loads <i>catia</i> and <i>streams_xattr</i> but read the caveat in NOTE below table.
full_audit	Record selected client operations to the system log.
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 <i>offline</i> attribute. This can prevent Windows Explorer from reading files just to make thumbnail images.

Continued on next page

Table 11.5 – continued from previous page

Value	Description
posix_eadb	Provides Extended Attributes (EAs) support so they can be used on filesystems which do not provide native support for EAs.
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 https://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*. macOS 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 macOS 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 drop-down menu:

- **recycle:** moves deleted files to the recycle directory instead of deleting them. Controlled by *Export Recycle Bin* in the *SMB share options* (page 213).
- **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.

11.4.1 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:

1. Go to *Sharing* → *Windows (SMB) Shares* and click *ADD*.
2. Fill out the the fields as shown in [Figure 11.14](#).
3. Enable the *Allow guest access* option.
4. Press *SAVE*.

Note: If a dataset for the share has not been created, refer to [Adding Datasets](#) (page 166) to find out more about dataset creation.

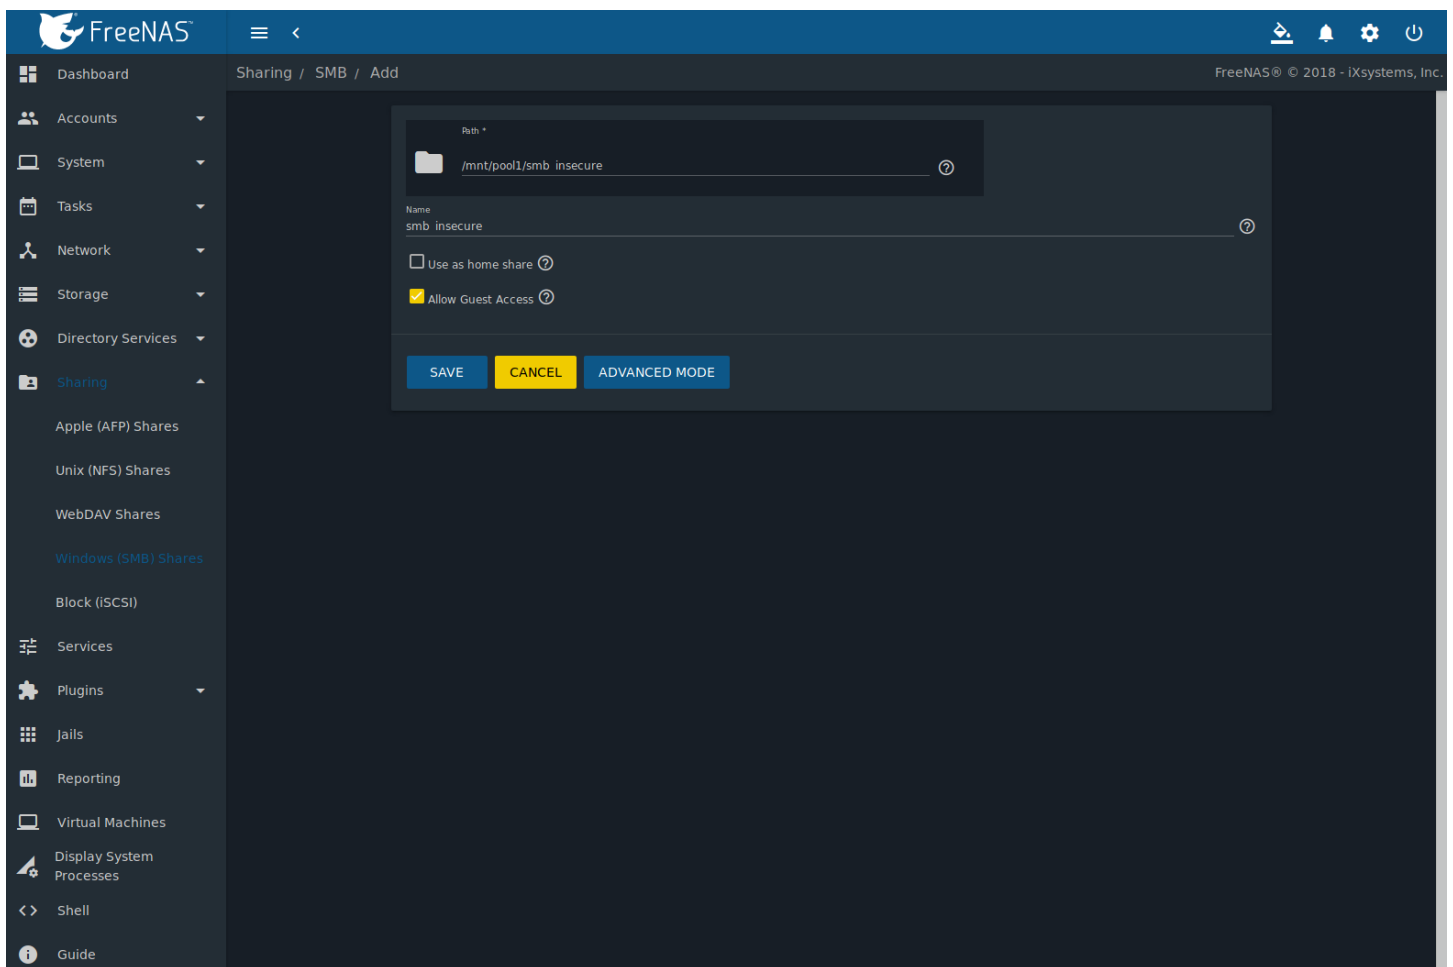


Fig. 11.14: Creating an Unauthenticated SMB Share

The new share appears in *Sharing* → *Windows (SMB) Shares*.

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.

11.4.2 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.

Before creating an authenticated SMB share, go to *Storage* → *Pools* to make a dataset for the share. For more information about dataset creation, refer to [Adding Datasets](#) (page 166).

After creating the dataset, go to *Storage* → *Pools* and click the ⓘ (Options) button for the desired dataset. Click *Edit Permissions* and fill out the information as shown in [Figure 11.15](#).

1. **ACL Type:** Select *Windows*.
2. **User:** If the user does not yet exist on the FreeNAS® system, go to *Accounts* → *Users* to create one. Refer to [Users](#) (page 67) for more information about creating a user. After the user has been created, use the drop-down to select the user account.
3. **Group:** Use the drop-down to select the desired group name. If the group does not yet exist on the FreeNAS® system, go to *Accounts* → *Groups* to create one. Refer to [Groups](#) (page 64) for more information about creating a group.
4. Click **SAVE**.

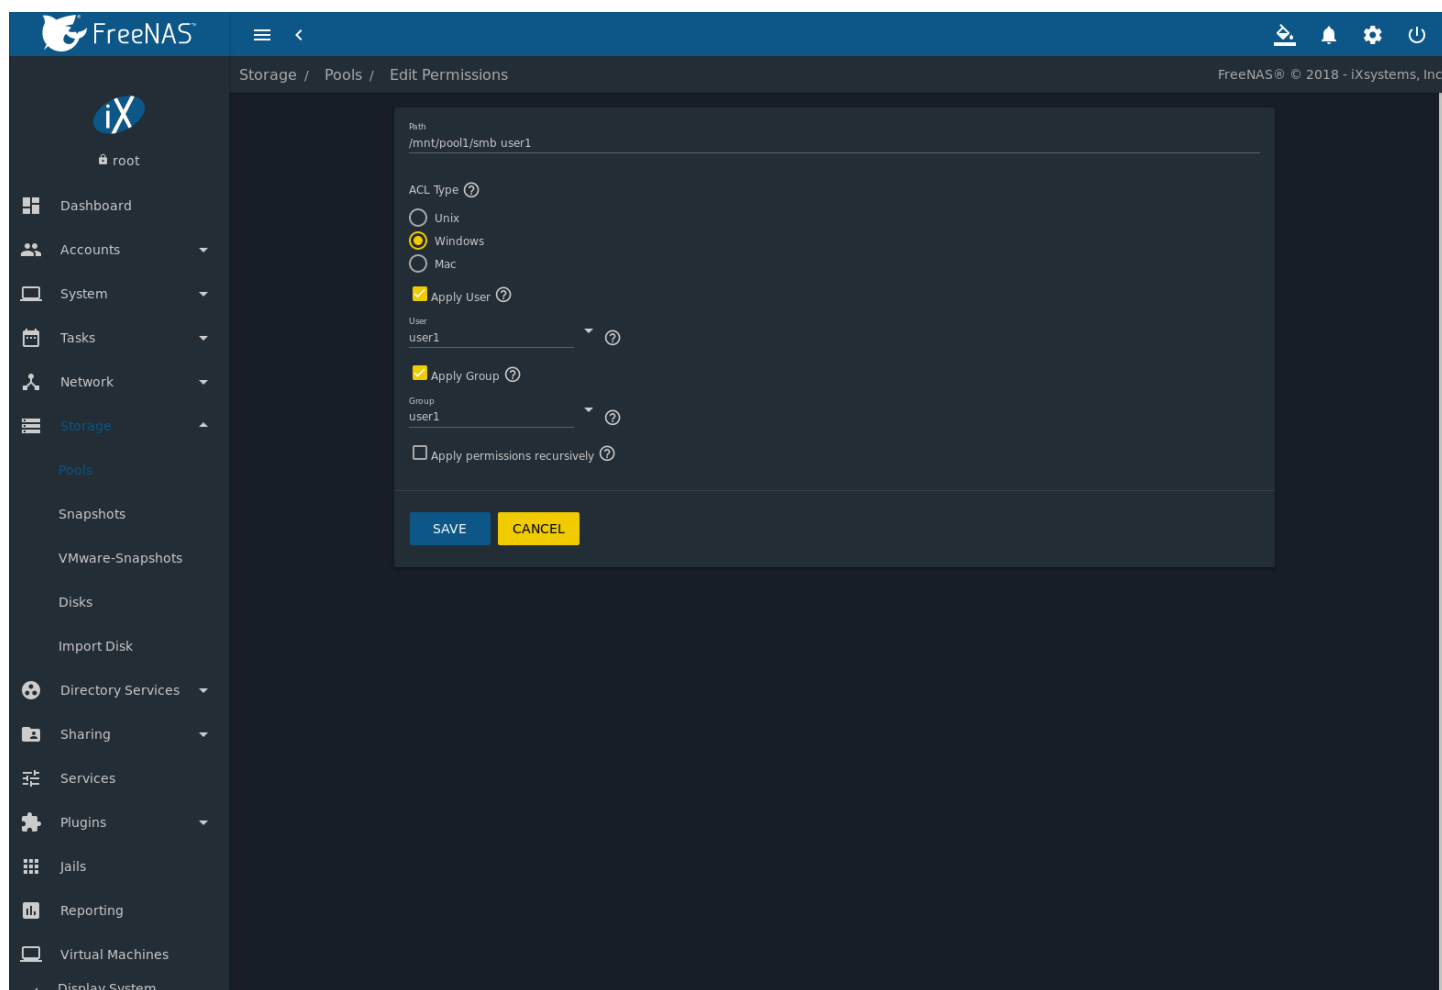


Fig. 11.15: Editing Dataset Permissions for Authenticated SMB Share

To create an authenticated SMB share, go to *Sharing* → *Windows (SMB) Shares* and click **ADD**, as shown in [Figure 11.16](#). Browse to the dataset created for the share and enter a name for the share. Press **SAVE** to create the share.

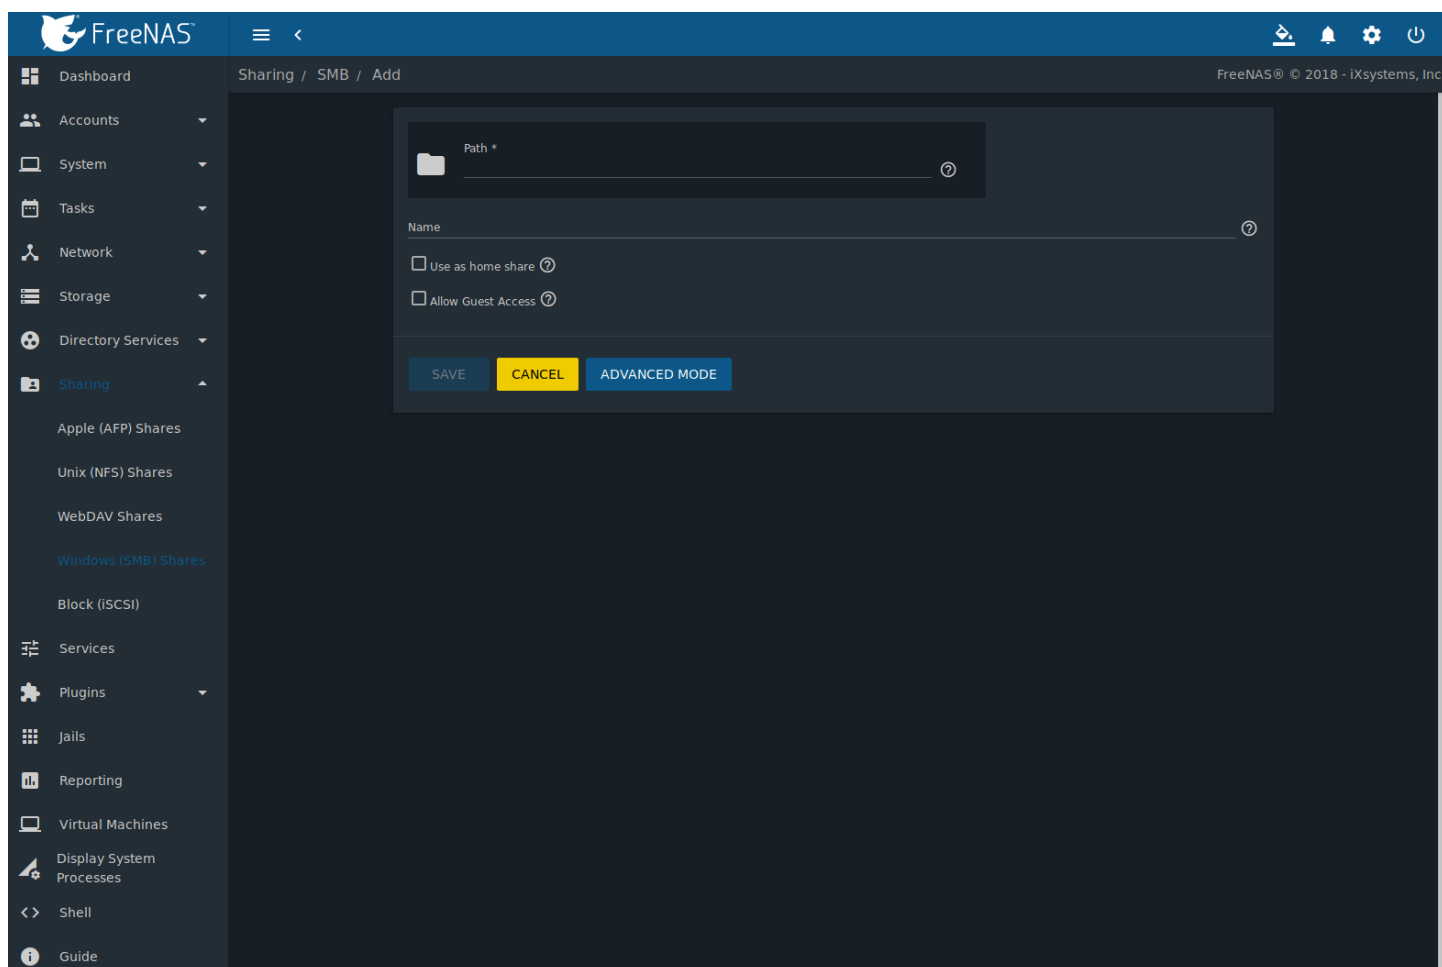


Fig. 11.16: Creating an Authenticated SMB Share

To configure multiple authenticated shares, repeat for each user. The new shares are also added to *Sharing* → *Windows (SMB) Shares*.

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*. After clicking *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. When authenticating to a share, if problems occur and the username and password are correct, type `cmd` in the *Search programs and files* box and use the following command to see if the share is already authenticated. 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, a prompt to authenticate appears.

11.4.3 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 can be used to 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 a periodic snapshot task is created on a ZFS pool 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 no previous versions of files to restore are visible, use Windows Update to ensure 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 pool or dataset, not on a directory.
- Datasets are filesystems and shadow copies cannot traverse filesystems. To see the shadow copies in the child datasets, create separate shares for them.
- Shadow copies will not work with a manual snapshot. Creating a periodic snapshot task for the pool or dataset being shared by SMB or a recursive task for a parent dataset is recommended.
- 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*.
- Appropriate permissions must be configured on the pool or 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® web interface. 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 218) 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. Go to *Tasks* → *Periodic Snapshot Tasks* and click **ADD** to create at least one periodic snapshot task. There are two options for snapshot tasks. One is to create a snapshot task for each user's dataset. In this example the datasets are `/mnt/volume1/user1` and `/mnt/volume1/user2`. Another option is to create one periodic snapshot task for the entire volume, `/mnt/volume1` in this case. **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 the 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* and click *Options* on an existing share. Click *Edit* then *ADVANCED MODE*. Use the *Periodic Snapshot Task* drop-down menu to 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/pool1/user1` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/pool1/user1` dataset and the share named `/mnt/pool1/user2` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/pool1/user2` dataset.
3. Verify that the SMB service is running in *Services*.

Figure 11.17 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. 11.17: Viewing Previous Versions within Explorer

11.5 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://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)) ([https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628\(v=ws.11\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11))), meaning that file copies happen locally, rather than over the network. It also supports the [VAAI](#) (page 358) (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 [Adding Zvols](#) (page 169) and use it to create a device extent, as described in [Extents](#) (page 231).

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*.

The rest of this section describes these steps in more detail.

11.5.1 Target Global Configuration

Sharing → *Block (iSCSI)* → *Target Global Configuration*, shown in [Figure 11.18](#), contains settings that apply to all iSCSI shares. [Table 11.6](#) summarizes the settings that are 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.

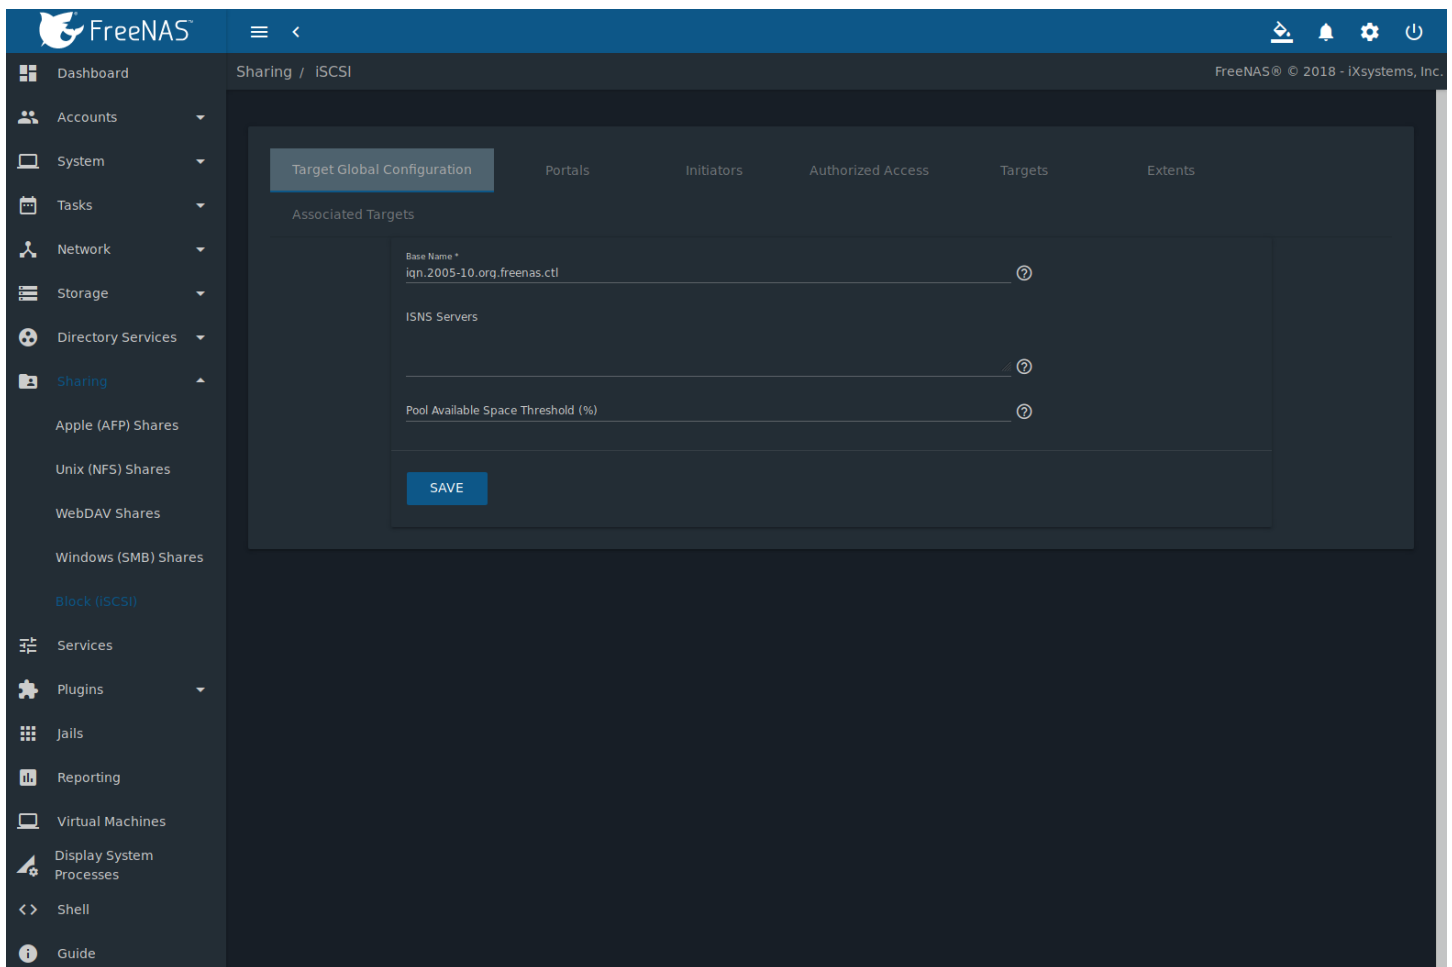


Fig. 11.18: iSCSI Target Global Configuration Variables

Table 11.6: 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.html) if unfamiliar with this format.
ISNS Servers	string	Enter the hostnames or IP addresses of ISNS servers to be registered with iSCSI targets and portals of the system. Separate each entry with a space.
Pool Available Space Threshold	integer	Enter the percentage of free space to in the pool. When this percentage is reached, the system issues an alert, but only if zvols are used. See VAAI (page 358) Threshold Warning for more information.

11.5.2 Portals

A portal specifies the IP address and port number to be used for iSCSI connections. Go to *Sharing* → *Block (iSCSI)* → *Portals* and click *ADD* to display the screen shown in [Figure 11.19](#).

[Table 11.19](#) summarizes the settings that can be configured when adding a portal. To assign additional IP addresses to the portal, click the link *Add extra Portal IP*.

Fig. 11.19: Adding an iSCSI Portal

Table 11.7: Portal Configuration Settings

Setting	Value	Description
Comment	string	Enter an optional description. Portals are automatically assigned a numeric group ID.
Discovery Auth Method	drop-down menu	<i>iSCSI</i> (page 248) supports multiple authentication methods that are used by the target to discover valid devices. <i>None</i> allows anonymous discovery while <i>CHAP</i> and <i>Mutual CHAP</i> both require authentication.
Discovery Auth Group	drop-down menu	Select a user created in <i>Authorized Access</i> if the <i>Discovery Auth Method</i> is set to <i>CHAP</i> or <i>Mutual CHAP</i> .
IP address	drop-down menu	Select the IPv4 or IPv6 address associated with an interface or the wild-card address of <i>0.0.0.0</i> (any interface).
Port	integer	TCP port used to access the iSCSI target. Default is <i>3260</i> .

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 these addresses:

192.168.1.1/24

192.168.2.1/24

192.168.3.1/24

192.168.4.1/24

A portal containing the first two IP addresses (group ID 1) and a portal containing the remaining two IP addresses (group ID 2) could be created. Then, a target named A with a Portal Group ID of 1 and a second target named B with a Portal Group ID of 2 could be created. 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.

11.5.3 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, go to *Sharing* → *Block (iSCSI)* → *Initiators* and click *ADD* as shown in Figure 11.20.

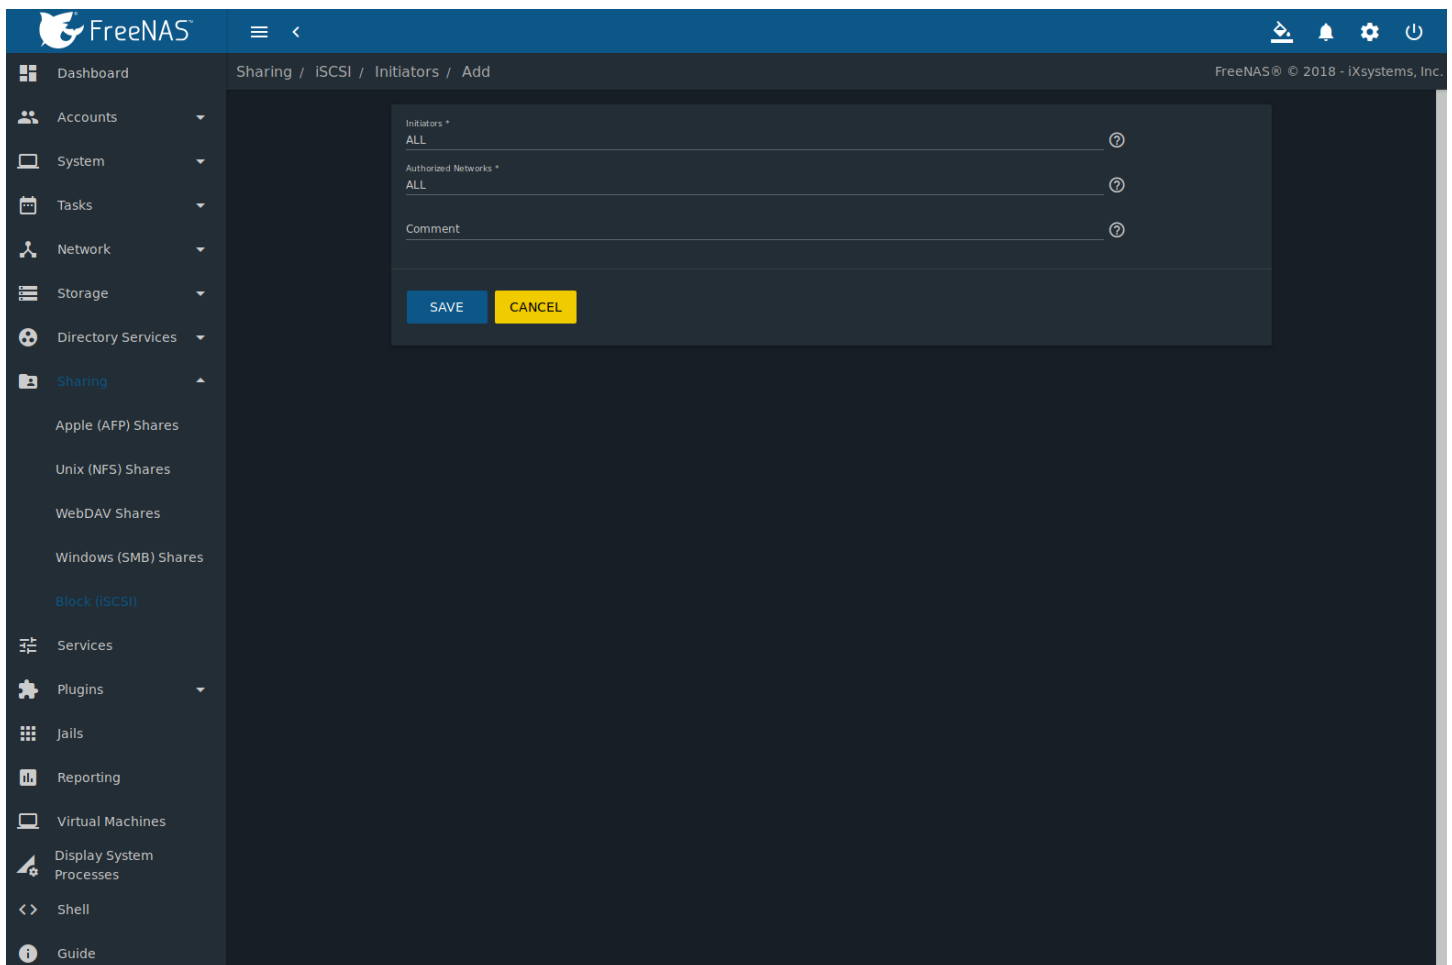


Fig. 11.20: Adding an iSCSI Initiator

Table 11.8 summarizes the settings that can be configured when adding an initiator.

Table 11.8: Initiator Configuration Settings

Setting	Value	Description
Initiators	string	Use <i>ALL</i> keyword or a list of initiator hostnames separated by spaces.
Authorized Networks	string	Network addresses that can use this initiator. Use <i>ALL</i> or list network addresses with a <i>CIDR</i> (https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing) mask. Separate multiple addresses with a space: <code>192.168.2.0/24 192.168.2.1/12</code> .
Comment	string	Notes or a description of the initiator.

In the example shown in Figure 11.21, two groups are 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 ⓘ (Options) on an initiator 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.

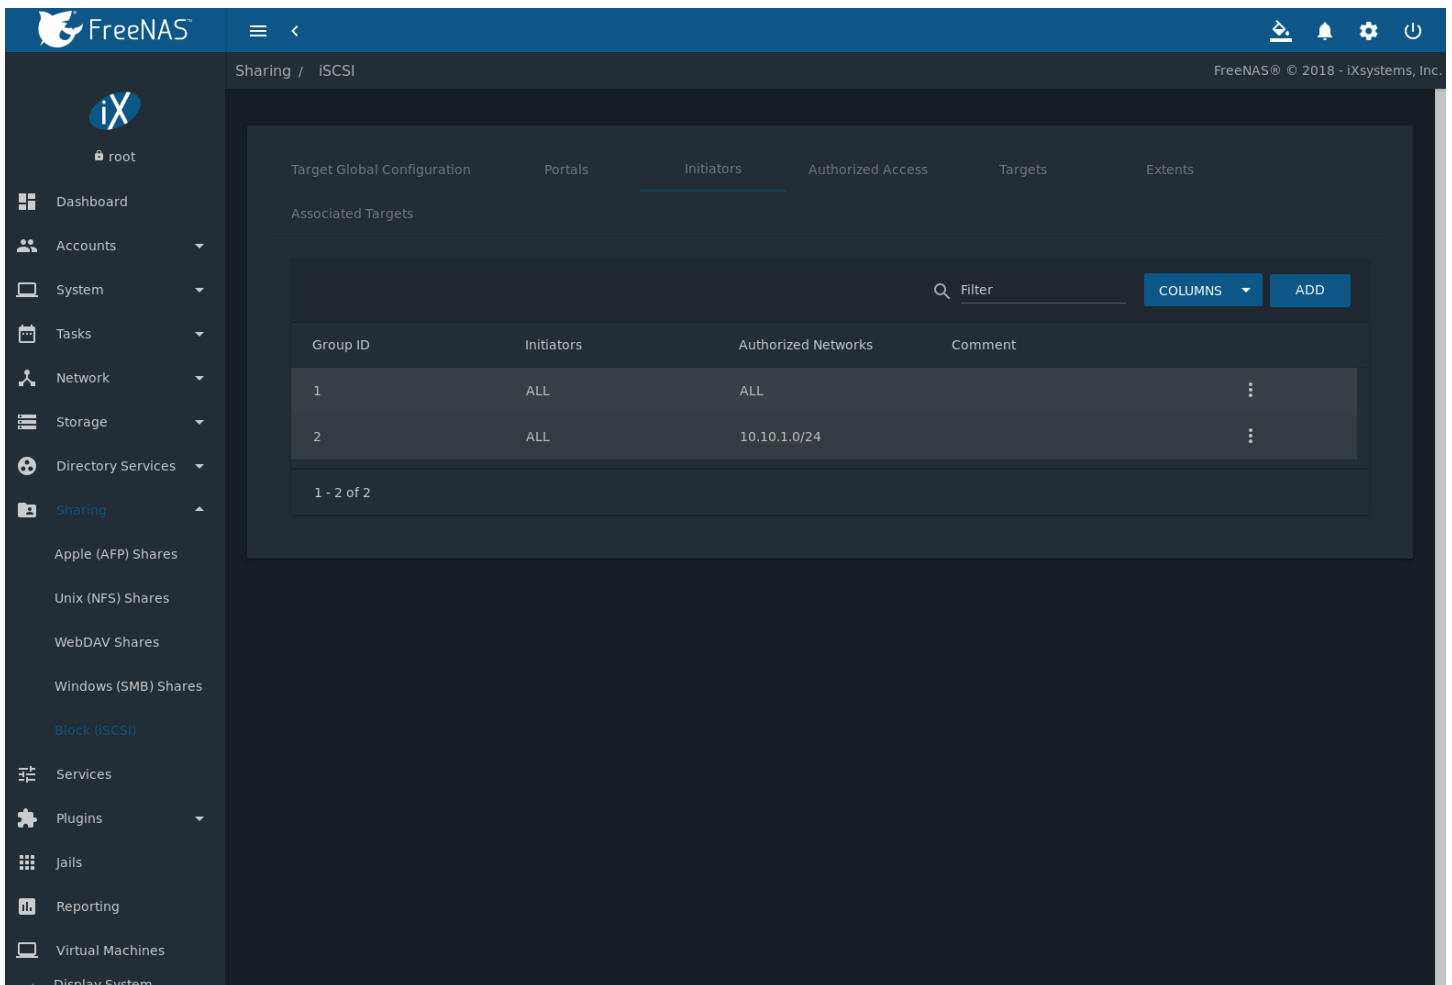


Fig. 11.21: Sample iSCSI Initiator Configuration

11.5.4 Authorized Accesses

When using CHAP or mutual CHAP to provide authentication, creating an authorized access is recommended. Do this by going to *Sharing* → *Block (iSCSI)* → *Authorized Access* and clicking *ADD*. The screen is shown in [Figure 11.22](#).

Note: This screen sets login authentication. This is different from discovery authentication which is set in *Global Configuration* (page 140).

Fig. 11.22: Adding an iSCSI Authorized Access

Table 11.9 summarizes the settings that can be configured when adding an authorized access:

Table 11.9: Authorized Access Configuration Settings

Setting	Value	Description
Group ID	integer	Allows different groups to be configured with different authentication profiles. Example: all users with a Group ID of 1 will inherit the authentication profile associated with Group 1
User	string	Enter 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	Enter and confirm a password for <i>User</i> . Must be between 12 and 16 characters.

Continued on next page

Table 11.9 – continued from previous page

Setting	Value	Description
Peer User	string	Only input when configuring mutual CHAP. In most cases it will need to be the same value as <i>User</i> .
Peer Secret	string	Enter and confirm the mutual secret password which must be different than the Secret . Required if <i>Peer User</i> is set.

Note: CHAP does not work with GlobalSAN initiators on macOS.

New authorized accesses are visible from the *Sharing* → *Block (iSCSI)* → *Authorized Access* menu. In the example shown in [Figure 11.23](#), 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.

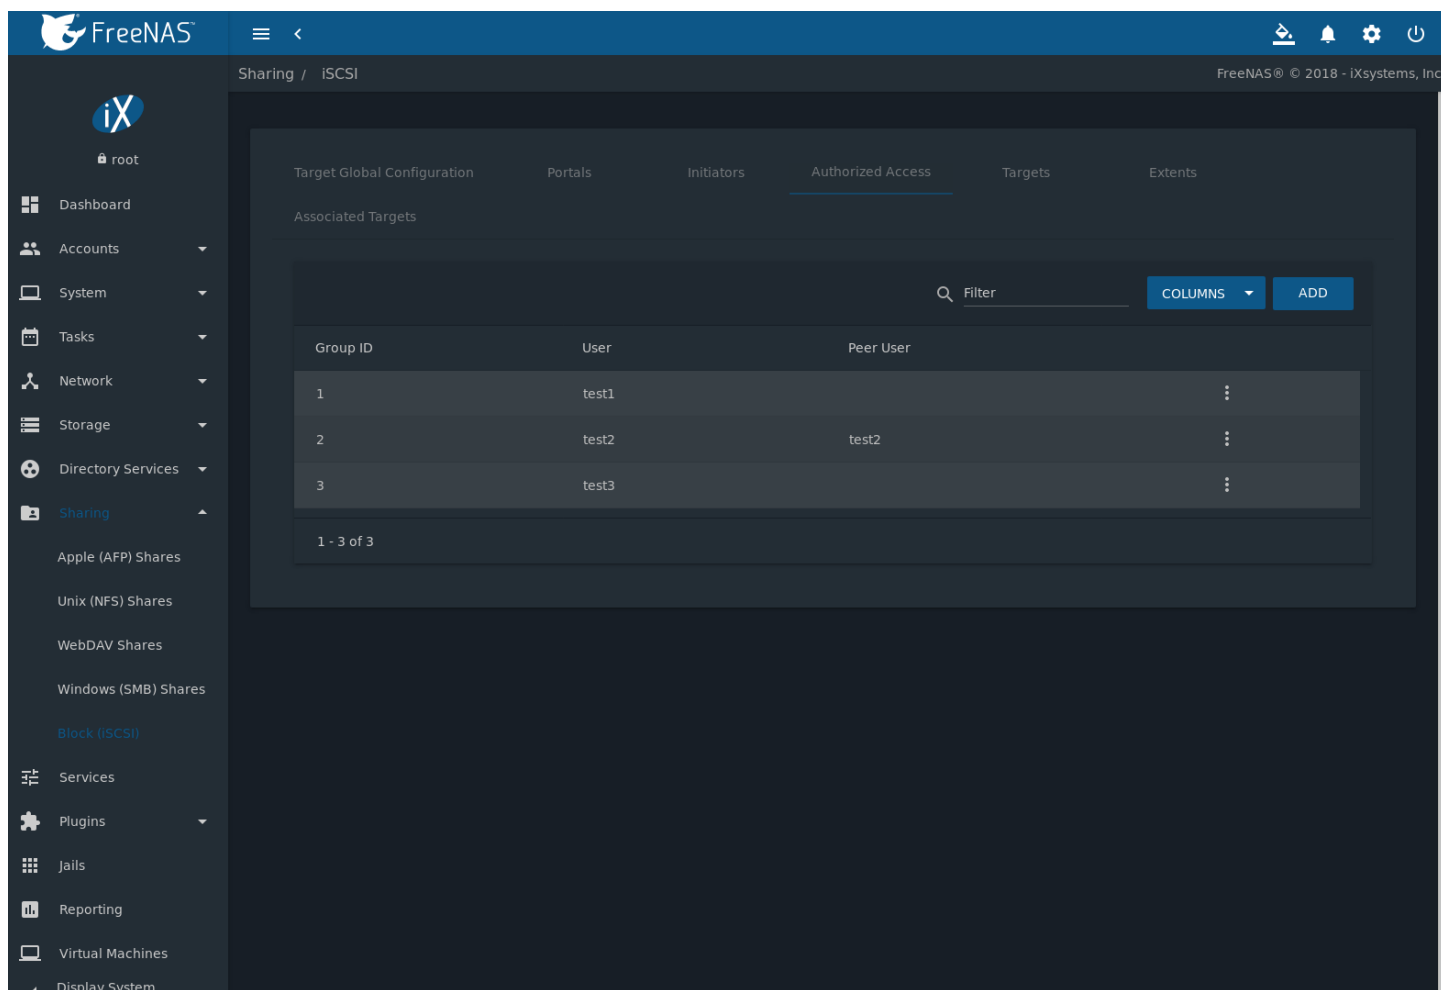


Fig. 11.23: Viewing Authorized Accesses

11.5.5 Targets

Next, create a Target by going to *Sharing* → *Block (iSCSI)* → *Targets* and clicking *ADD* as shown in [Figure 11.24](#). A target combines a portal ID, allowed initiator ID, and an authentication method. [Table 11.10](#) 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).

The screenshot shows the FreeNAS web interface with the 'Add iSCSI Target' form. The left sidebar contains navigation links: Dashboard, Accounts, System, Tasks, Network, Storage, Directory Services, Sharing (selected), Services, Plugins, Jails, Reporting, Virtual Machines, Display System Processes, Shell, and Guide. The 'Sharing' menu is expanded, showing options like Apple (AFP) Shares, Unix (NFS) Shares, WebDAV Shares, Windows (SMB) Shares, and Block (iSCSI). The main content area shows the 'Add iSCSI Target' form with the following fields:

- Target Name *
- Target Alias
- Portal Group ID * (set to None)
- Initiator Group ID
- Auth Method (set to None)
- Authentication Group number

At the bottom of the form are three buttons: SAVE, CANCEL, and ADD EXTRA ISCSI GROUP.

Fig. 11.24: Adding an iSCSI Target

Table 11.10: Target Settings

Setting	Value	Description
Target Name	string	Required value. Base name will be appended automatically if it does not start with <i>iqn</i> .
Target Alias	string	Enter an 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: <i>None</i> , <i>Auto</i> , <i>CHAP</i> , or <i>Mutual CHAP</i> .
Authentication Group number	drop-down menu	Select <i>None</i> or an integer. This number represents the number of existing authorized accesses.

11.5.6 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](https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html) (https://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 358) 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 pool 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 235).

To add an extent, go to *Sharing* → *Block (iSCSI)* → *Extents* and click *ADD*. In the example shown in [Figure 11.25](#), the device extent is using the `export` zvol that was previously created from the `/mnt/pool1` pool.

[Table 11.11](#) summarizes the settings that can be configured when creating an extent. Note that **file extent creation fails unless the name of the file to be created is appended to the pool or dataset name**.

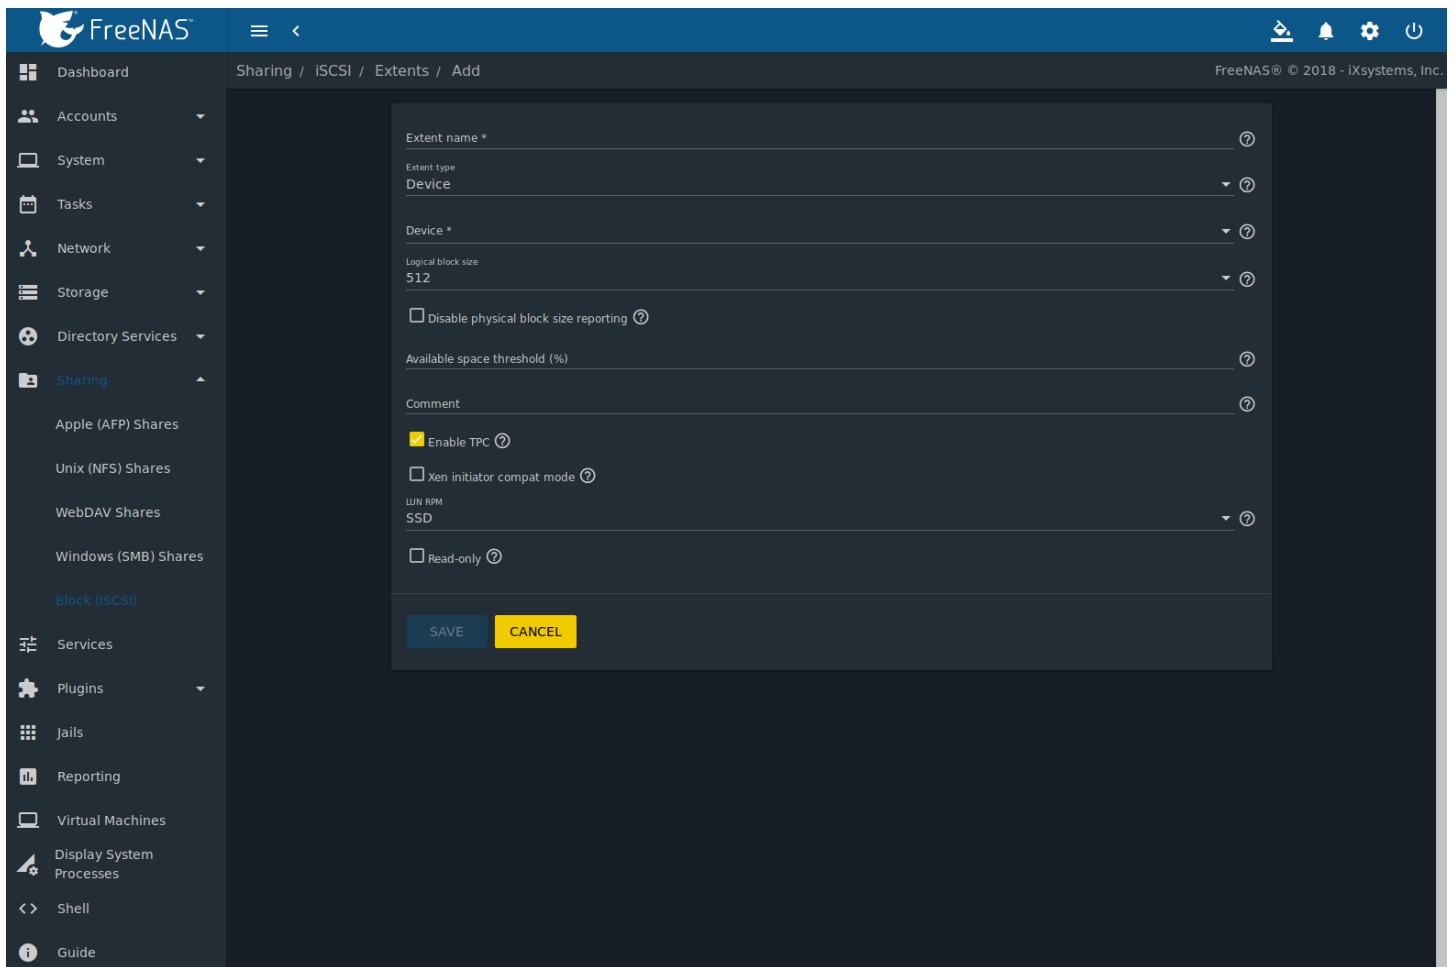


Fig. 11.25: Adding an iSCSI Extent

Table 11.11: Extent Configuration Settings

Setting	Value	Description
Extent name	string	Enter the extent name. If the <i>Extent size</i> is not 0, it cannot be an existing file within the pool or dataset.
Extent type	drop-down menu	Select from <i>File</i> or <i>Device</i> .
Path to the extent	browse button	Only appears if <i>File</i> is selected. Browse to an existing file and use 0 as the <i>Extent size</i> , or browse to the pool or dataset, click <i>Close</i> , append the <i>Extent Name</i> to the path, and specify a value in <i>Extent size</i> . Extents cannot be created inside the jail root directory.
Extent size	integer	Only appears if <i>File</i> 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.
Device	drop-down menu	Only appears if <i>Device</i> is selected. Select the unformatted disk, controller, zvol, zvol snapshot, or HAST device.
Logical block size	drop-down menu	Only override the default if the initiator requires a different block size.
Disable physical block size reporting	checkbox	Set if the initiator does not support physical block size values over 4K (MS SQL). Setting can also prevent constant block size warnings (https://www.virtten.net/2016/12/the-physical-block-size-reported-by-the-device-is-not-supported/) when using this share with ESXi.

Continued on next page

Table 11.11 – continued from previous page

Setting	Value	Description
Available space threshold	string	Only appears if <i>File</i> or a zvol is selected. When the specified percentage of free space is reached, the system issues an alert. See VAAI (page 358) Threshold Warning.
Comment	string	Enter an optional comment.
Enable TPC	checkbox	If enabled, 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	Set this option 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	Set this option to prevent the initiator from initializing this LUN.

11.5.7 Associated Targets

The last step is associating an extent to a target by going to *Sharing* → *Block (iSCSI)* → *Associated Targets* and clicking *ADD*. The screen is shown in [Figure 11.26](#). Use the drop-down menus to select the existing target and extent. Click *SAVE* to add an entry for the LUN.

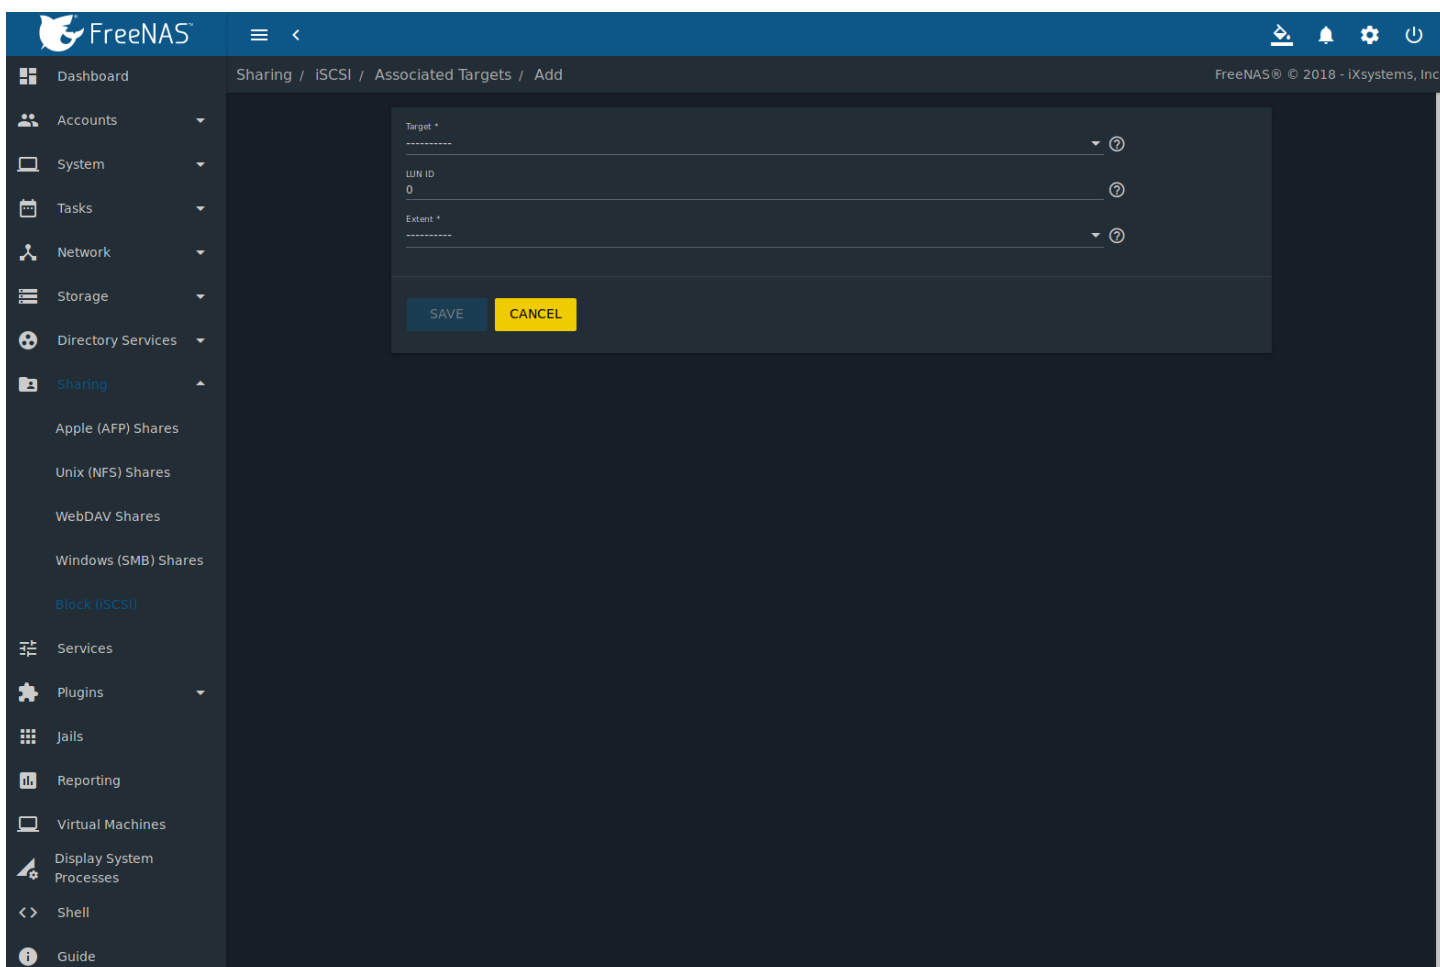


Fig. 11.26: Associating a Target With an Extent

Table 11.12 summarizes the settings that can be configured when associating targets and extents.

Table 11.12: Associated Target Configuration Settings

Setting	Value	Description
Target	drop-down menu	Select an existing target.
LUN ID	integer	Type a value between 0 and 1023. Note that some initiators expect a value below 256. Enter 0 to statically assign the next available ID.
Extent	drop-down menu	Select an existing extent.

Always associating extents to targets in a one-to-one manner is recommended, even though the web interface 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. Clearing the initiator connections to a LUN before deleting it is recommended.

After iSCSI has been configured, remember to start the service in *Services* → *iSCSI* by clicking the  (Power) button.

11.5.8 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://techgenix.com/Connecting-Windows-7-iSCSI-SAN/) (<http://techgenix.com/Connecting-Windows-7-iSCSI-SAN/>). 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](https://www.pluralsight.com/blog/software-development/freenas-8-iscsi-target-windows-7) (<https://www.pluralsight.com/blog/software-development/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\(8\)](https://www.freebsd.org/cgi/man.cgi?query=iscontrol) (<https://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) (<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) (<http://netbsd.gw.com/cgi-bin/man-cgi?iscsi-initiator++NetBSD-current>) comes with NetBSD, and [iscsid\(8\)](http://man.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man8/iscsid.8?query=iscsid) (<http://man.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 the distribution should the command not exist on the 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\(i\)](https://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/) (<https://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](https://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) (https://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.

11.5.9 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.

11.5.9.1 Zvol Based LUN

To grow a zvol-based LUN, go to *Storage* → *Pools*, click ⓘ (Options) on the zvol to be grown, then click *Edit zvol*. In the example shown in Figure 11.27, the current size of the zvol named *zvol1* is 4 GiB.

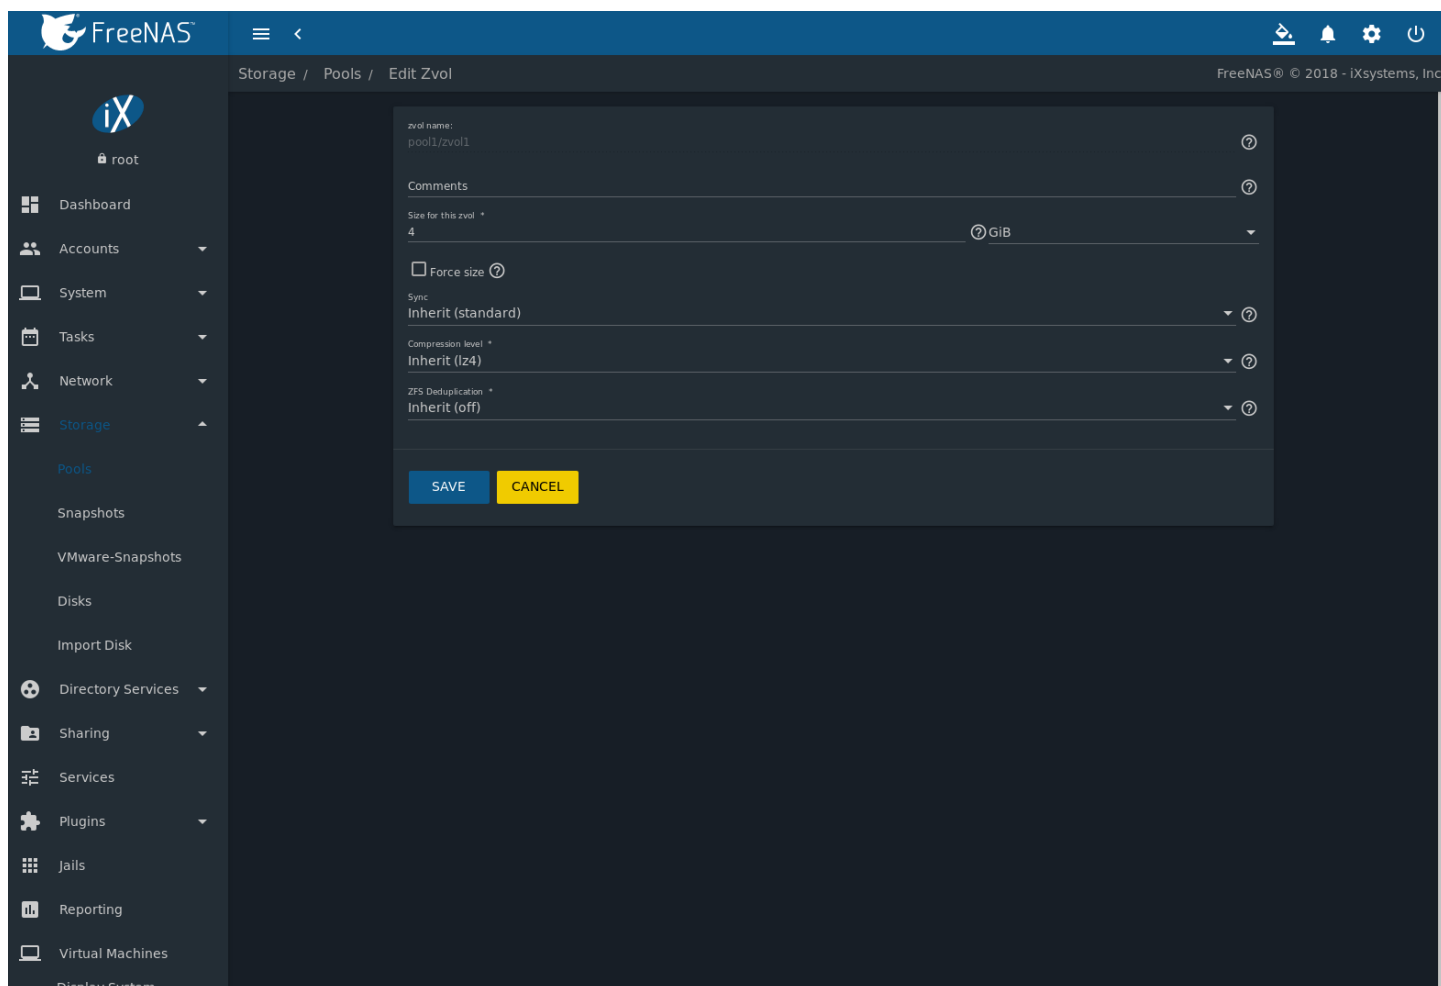



Fig. 11.27: Editing an Existing Zvol

Enter the new size for the zvol in the *Size for this zvol* field and click *SAVE*. The new size for the zvol is immediately shown in the *Used* column of the *Storage* → *Pools* table.


Note: The web interface does not allow reducing 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 pool size.

11.5.9.2 File Extent Based LUN

To grow a file extent-based LUN:

Go to *Services* → *iSCSI* → *CONFIGURE* → *Extents*. Click  (Options), then *Edit*. Ensure the *Extent Type* is set to file and enter the *Path to the extent*. Open the [Shell](#) (page 325) to grow the file extent. This example grows `/mnt/pool1/data` by 2 GiB:

```
truncate -s +2g /mnt/pool1/data
```

Return to *Services* → *iSCSI* → *CONFIGURE* → *Extents*, click  (Options) on the desired file extent, then click *Edit*. Set the size to `0` as this causes the iSCSI target to use the new size of the file.

SERVICES

Services that ship with FreeNAS® are configured, started, or stopped in *Services*. FreeNAS® includes these built-in services:

- *AFP* (page 238)
- *Domain Controller* (page 240)
- *Dynamic DNS* (page 242)
- *FTP* (page 243)
- *iSCSI* (page 248)
- *LLDP* (page 248)
- *Netdata* (page 249)
- *NFS* (page 250)
- *Rsync* (page 252)
- *S3* (page 255)
- *S.M.A.R.T.* (page 256)
- *SMB* (page 258)
- *SNMP* (page 261)
- *SSH* (page 263)
- *TFTP* (page 265)
- *UPS* (page 267)
- *WebDAV* (page 269)

This section demonstrates starting a FreeNAS® service and the available configuration options for each FreeNAS® service.

12.1 Configure Services

The *Services* page, shown in [Figure 12.1](#), lists all services. The list has options to activate the service, set a service to *Start Automatically* at system boot, and configure a service. 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* (<https://en.wikipedia.org/wiki/S.M.A.R.T.>). Other services default to *off* until started.

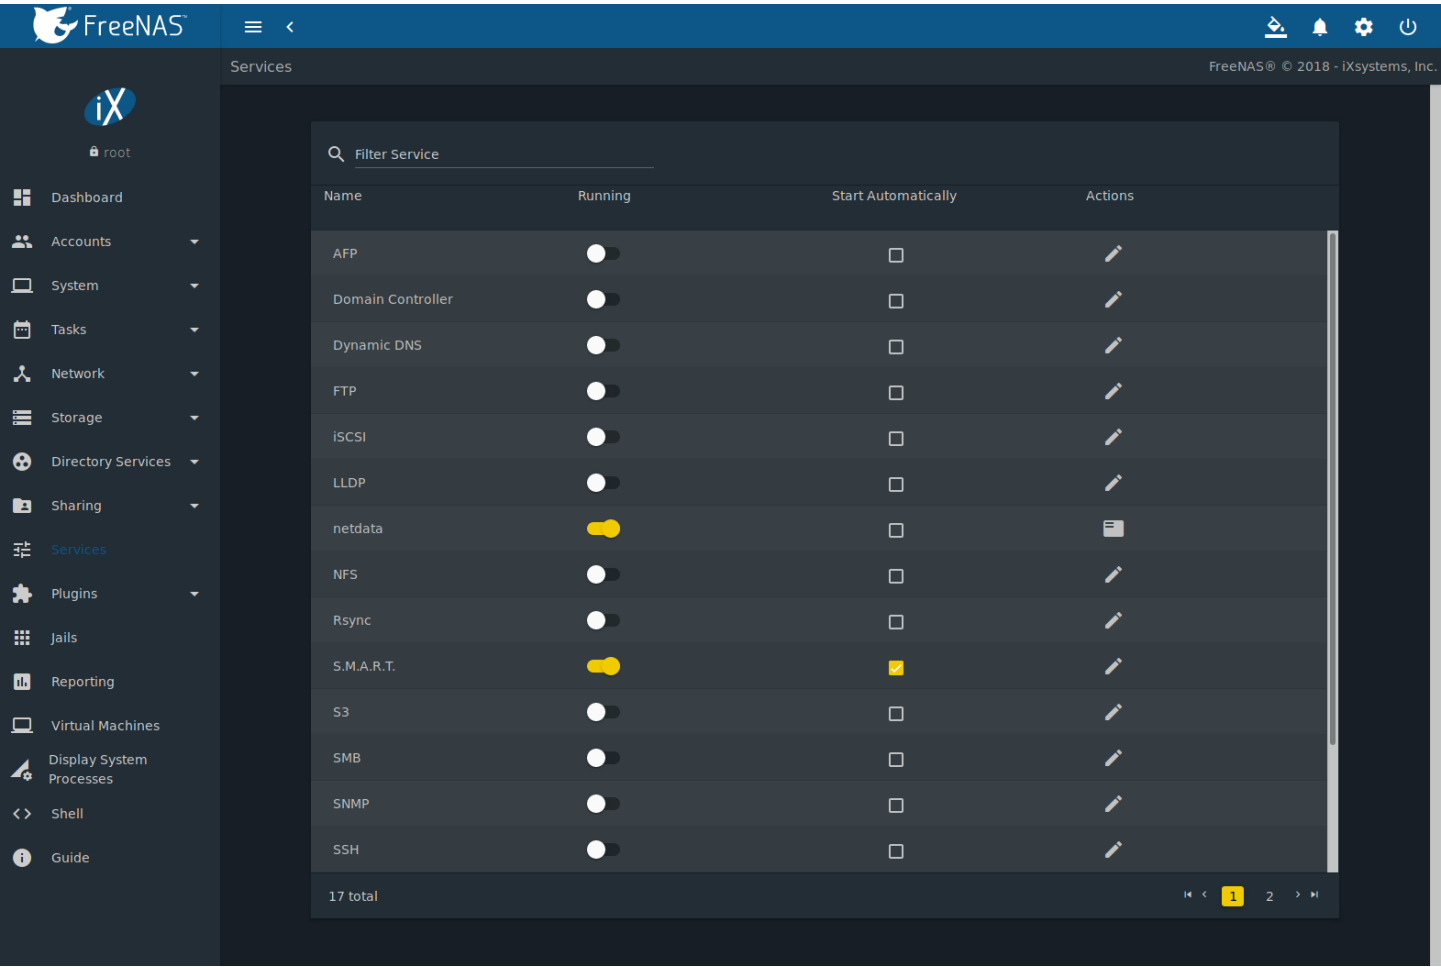


Fig. 12.1: Configure Services

Stopped services show the sliding button on the left. Active services show the sliding button on the right. Click the slider to start or stop a service. Stopping a service shows a confirmation dialog.

Tip: Using a proxy server can prevent the list of services from being displayed. If a proxy server is used, do not configure it to proxy local network 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  (Configure).

If a service does not start, go to *System* → *Advanced* and enable *Show console messages*. 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 325) and type `more /var/log/messages`.

12.2 AFP

The settings that are configured when creating AFP shares in are specific to each configured AFP share. An AFP share is created by navigating to *Sharing* → *Apple (AFP)*, and clicking *ADD*. In contrast, global settings which apply to all AFP shares are configured in *Services* → *AFP* → *Configure*.

Figure 12.2 shows the available global AFP configuration options which are described in Table 12.1.

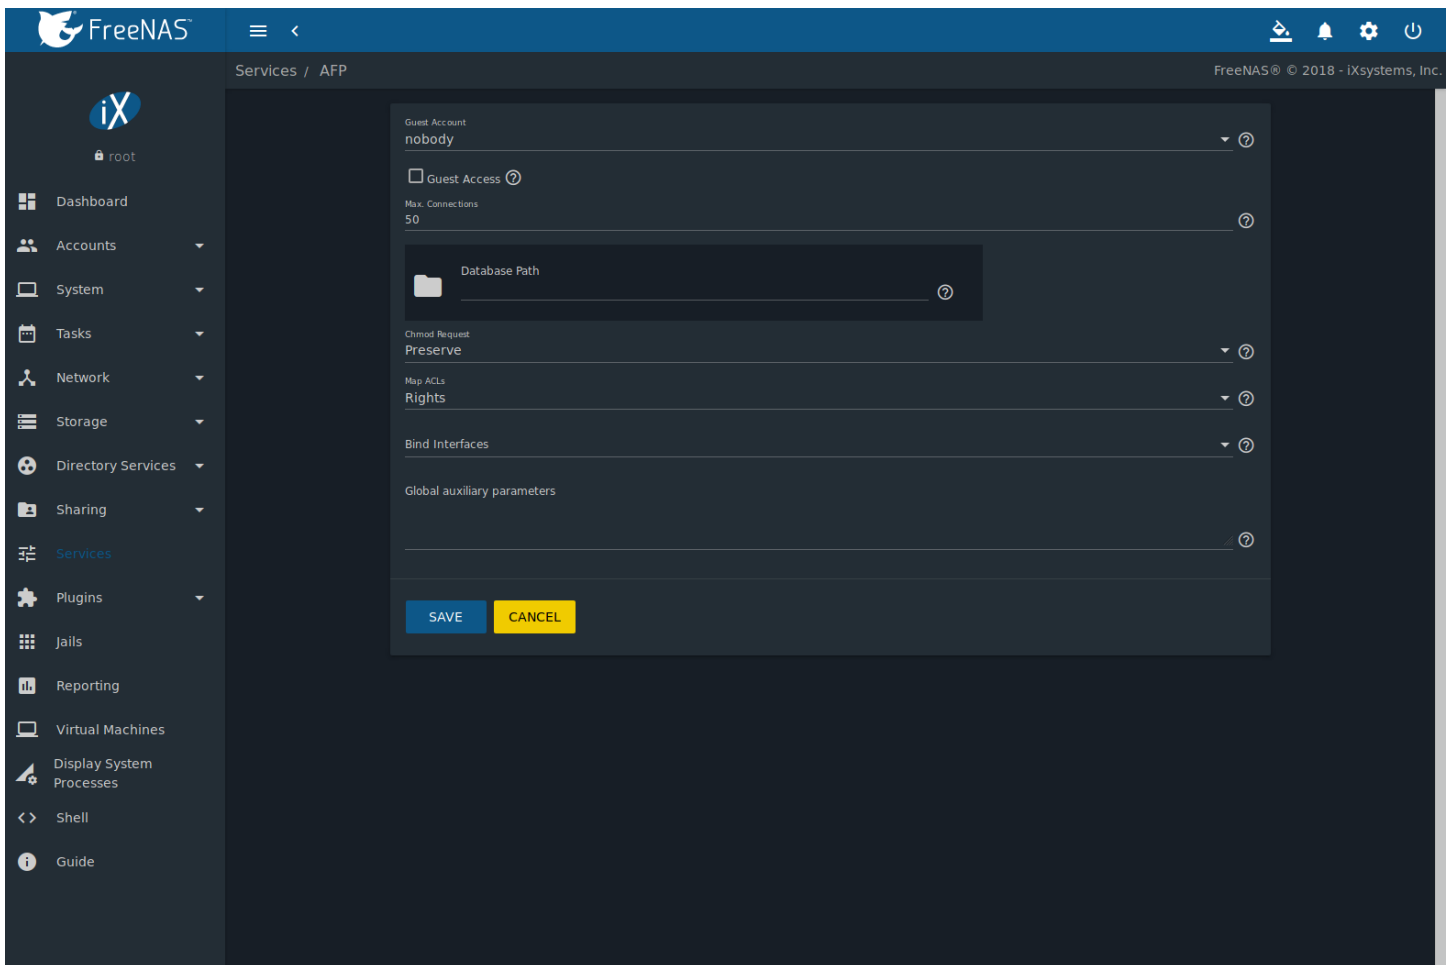


Fig. 12.2: Global AFP Configuration

Table 12.1: Global AFP Configuration Options

Setting	Value	Description
Guest Account	drop-down menu	Select an account to use for guest access. The account must have permissions to the pool or dataset being shared.
Guest Access	checkbox	If enabled, clients are not prompted to authenticate before accessing AFP shares.
Max. Connections	integer	Maximum number of simultaneous connections permitted via AFP. The default limit is 50.
Database Path	browse button	Sets the database information to be stored in the path. Default is the root of the pool. The path must be writable even if the pool is read only.
Chmod Request	drop-down menu	Set how ACLs are handled. Choices are: <i>Ignore</i> , <i>Preserve</i> , or <i>Simple</i> .
Map ACLs	drop-down menu	Choose mapping of effective permissions for authenticated users: <i>Rights</i> (default, Unix-style permissions), <i>Mode</i> (ACLs), or <i>None</i> .
Bind Interfaces	selection	Specify the IP addresses to listen for FTP connections. Select the desired IP addresses in the list to add them to the <i>Bind Interfaces</i> list.
Global auxiliary parameters	string	Additional afp.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered elsewhere in this screen.

12.2.1 Troubleshooting AFP

Check for error messages in `/var/log/afp.log`.

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 325), replacing the path to the problematic AFP share:

```
dbd -rf /path/to/share
```

This command can take some time, depending upon the size of the pool or dataset being shared. The CNID database is wiped and rebuilt from the CNIDs stored in the AppleDouble files.

12.3 Domain Controller

FreeNAS® can be configured to act either as the domain controller for a network or to join an existing *Active Directory* (page 183) 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 183) network to access its authentication and authorization services, configure *Active Directory* (page 183) instead.

Note that configuring a domain controller is a complex process that requires a good understanding of how *Active Directory* (page 183) works. While *Services* → *Domain Controller* → *Configure* makes it easy to enter the needed settings into the web 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). 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 12.3 shows the configuration screen for creating a domain controller and Table 12.2 summarizes the available options.

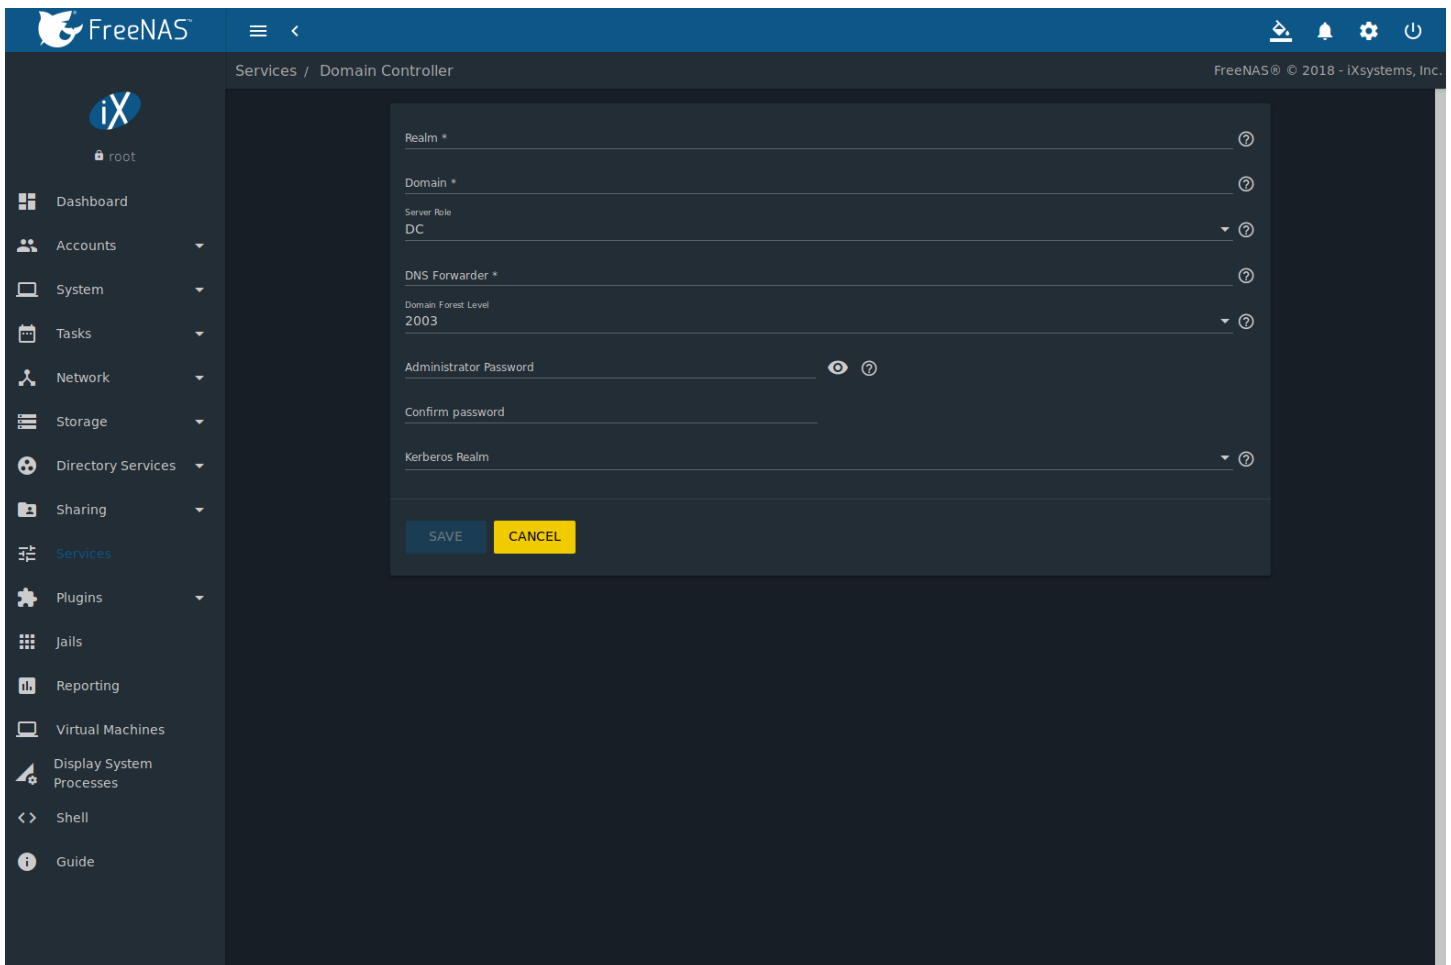


Fig. 12.3: Domain Controller Settings

Table 12.2: Domain Controller Configuration Options

Setting	Value	Description
Realm	string	Enter a capitalized DNS realm name.
Domain	string	Enter a 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	Enter the IP address of the DNS forwarder. Required for recursive queries when <i>SAMBA_INTERNAL</i> is selected.
Domain Forest Level	drop-down menu	Choices are <i>2000</i> , <i>2003</i> , <i>2008</i> , <i>2008_R2</i> , <i>2012</i> , or <i>2012_R2</i> . Refer to Understanding Active Directory Domain Services (AD DS) Functional Levels (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc754918(v=ws.10)) .
Administrator Password	string	Enter and confirm the password to be used for the <i>Active Directory</i> (page 183) administrator account.
Kerberos Realm	drop-down menu	Auto-populates with information from the <i>Realm</i> when the settings in this screen are saved.

12.3.1 Samba Domain Controller Backup

A `samba_backup` script is available to back up Samba4 domain controller settings is available. From the [Shell](#) (page 325), run `/usr/local/bin/samba_backup --usage` to show the input options.

12.4 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 access to the FreeNAS® system even if the IP address changes. DDNS requires registration with a DDNS service such as [DynDNS](https://dyn.com/dns/) (<https://dyn.com/dns/>).

Figure 12.4 shows the DDNS configuration screen and Table 12.3 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* → *Dynamic DNS*.

Fig. 12.4: Configuring DDNS

Table 12.3: DDNS Configuration Options

Setting	Value	Description
Provider	drop-down menu	Several providers are supported. If a specific provider is not listed, select <i>Custom Provider</i> and enter the information in the <i>Custom Server</i> and <i>Custom Path</i> fields.
CheckIP Server SSL	string	Set to use HTTPS for the connection to the <i>CheckIP Server</i> .
CheckIP Server	string	Enter the name and port of the server that reports the external IP address. Example: <i>server.name.org:port</i> .

Continued on next page

Table 12.3 – continued from previous page

Setting	Value	Description
CheckIP Path	string	Enter the path requested by the <i>CheckIP</i> Server to determine the user IP address.
Use SSL		Set to use HTTPS for the connection to the server that updates the DNS record.
Domain name	string	Enter a fully qualified domain name. Example: <i>yourname.dyndns.org</i> .
Username	string	Enter the username used to log in to the provider and update the record.
Password	string	Enter the password used to log in 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 (<https://he.net>) website for *Password*.

12.5 FTP

FreeNAS® uses the [proftpd](http://www.proftpd.org/) (<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 concerned about sensitive data, see [Encrypting FTP](#) (page 248).

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 12.5](#) shows the configuration screen for *Services → FTP → Configure*. 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 setting the *Show advanced fields by default* option in *System → Advanced*.

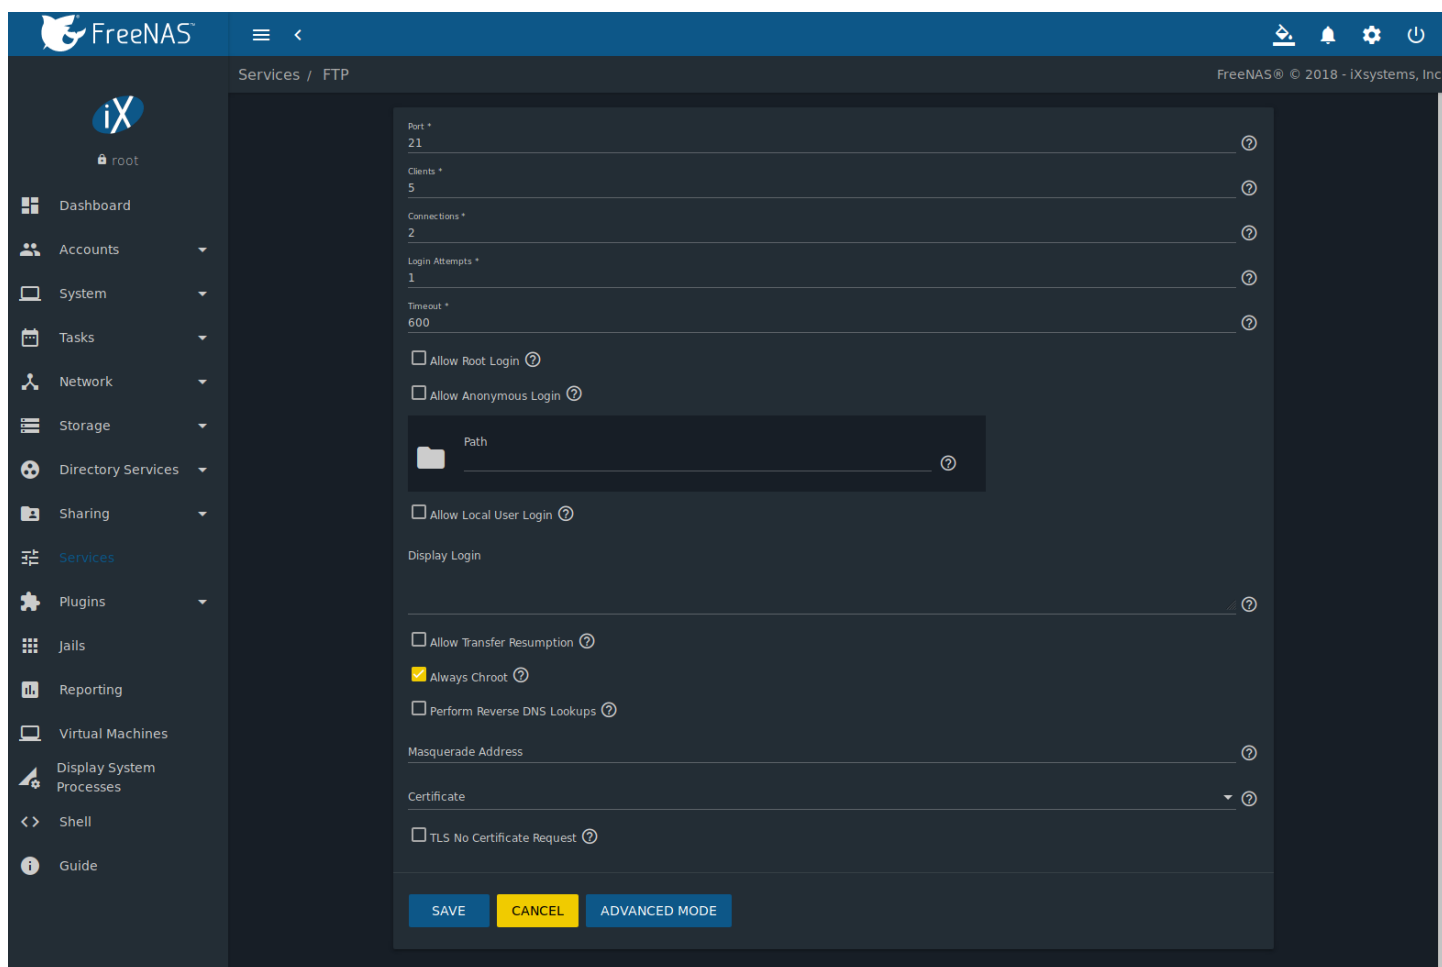


Fig. 12.5: Configuring FTP

Table 12.4 summarizes the available options when configuring the FTP server.

Table 12.4: FTP Configuration Options

Setting	Value	Advanced Mode	Description
Port	integer		Set the port the FTP service listens on.
Clients	integer		Maximum number of simultaneous clients.
Connections	integer		Set the maximum number of connections per IP address. <i>0</i> means unlimited.
Login Attempts	integer		Enter the maximum number of attempts before the 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		Setting this option is discouraged as it increases security risk.
Allow Anonymous Login	checkbox		Set to allow anonymous FTP logins with access to the directory specified in <i>Path</i> .
Path	browse button		Set the root directory for anonymous FTP connections.
Allow Local User Login	checkbox		Required if <i>Anonymous Login</i> is disabled.
Display Login	string		Specify the message displayed to local login users after authentication. Not displayed to anonymous login users.

Continued on next page

Table 12.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Allow Transfer Re-umption	checkbox		Set to allow FTP clients to resume interrupted transfers.
Always Chroot	checkbox		When set a local user is only allowed access to their home directory when they are a member of the <i>wheel</i> group.
Perform Reverse DNS Lookups	checkbox		Set to 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.
Certificate	drop-down menu		Select the SSL certificate to be used for TLS FTP connections. Go to <i>System</i> → <i>Certificates</i> to create a certificate.
TLS No Certificate Request	checkbox		Set if the client cannot connect, and it is suspected the client is not properly handling server certificate requests.
File Permission	checkboxes	✓	Sets default permissions for newly created files.
Directory Permission	checkboxes	✓	Sets default permissions for newly created directories.
Enable FXP (https://en.wikipedia.org/wiki/File_eXchange_Protocol)	checkbox	✓	Set to enable the File eXchange Protocol. This is discouraged as it makes the server vulnerable to FTP bounce attacks.
Require IDENT Authentication	checkbox	✓	Setting this option results in timeouts if <code>identd</code> is not running on the client.
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	✓	Defined in KiB/s, default of 0 means unlimited.
Local User Download Bandwidth	integer	✓	Defined in KiB/s, default of 0 means unlimited.
Anonymous User Upload Bandwidth	integer	✓	Defined in KiB/s, default of 0 means unlimited.
Anonymous User Download Bandwidth	integer	✓	Defined in KiB/s, default of 0 means unlimited.
Enable TLS	checkbox	✓	Set to enable encrypted connections. Requires a certificate to be created or imported using Certificates (page 101).
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	✓	Setting this option 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 set, the user 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 set, the user password may be sent unencrypted.
TLS Common Name Required	checkbox	✓	When set, the common name in the certificate must match the FQDN of the host.
TLS Enable Diagnostics	checkbox	✓	If set when troubleshooting a connection, logs more verbosely.
TLS Export Certificate Data	checkbox	✓	If set, exports the certificate environment variables.
TLS No Certificate Request	checkbox	✓	Set if the client cannot connect and it is suspected the client is poorly handling the server certificate request.

Continued on next page

Table 12.4 – continued from previous page

Setting	Value	Advanced Mode	Description
TLS No Empty Fragments	checkbox	✓	Setting this option is not recommended as it bypasses a security mechanism.
TLS No Session Reuse Required	checkbox	✓	Setting this option reduces the security of the connection. Only use if the client does not understand reused SSL sessions.
TLS Export Standard Vars	checkbox	✓	If enabled, sets several environment variables.
TLS DNS Name Required	checkbox	✓	If set, the client DNS name must resolve to its IP address and the cert must contain the same DNS name.
TLS IP Address Required	checkbox	✓	If set, the client certificate must contain the IP address that matches the IP address of the client.
Auxiliary Parameters	string	✓	Used to add proftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=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>
```

12.5.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the FreeNAS® system is not accessible from the Internet and everyone in the internal network needs easy access to the stored data. Anonymous FTP does not require 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 pool or dataset to be shared in *Storage* → *Pools* → *Edit Permissions*:
 - *User*: select the built-in *ftp* user from the drop-down menu
 - *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 Unix ACLs are used even if Windows clients are accessing FreeNAS® via FTP.

2. Configure anonymous FTP in *Services* → *FTP* → *Configure* by setting these attributes:
 - *Allow Anonymous Login*: set this option
 - *Path*: browse to the pool/dataset/directory to be shared
3. Start the FTP service in *Services*. Click the sliding button on the *FTP* row. The FTP service takes a second or so to start. The sliding button moves to the right when the service is running.
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 12.6](#), The user has entered this information into the Filezilla client:

- IP address of the FreeNAS® server: *192.168.1.113*
- *Username*: *anonymous*
- *Password*: the email address of the user

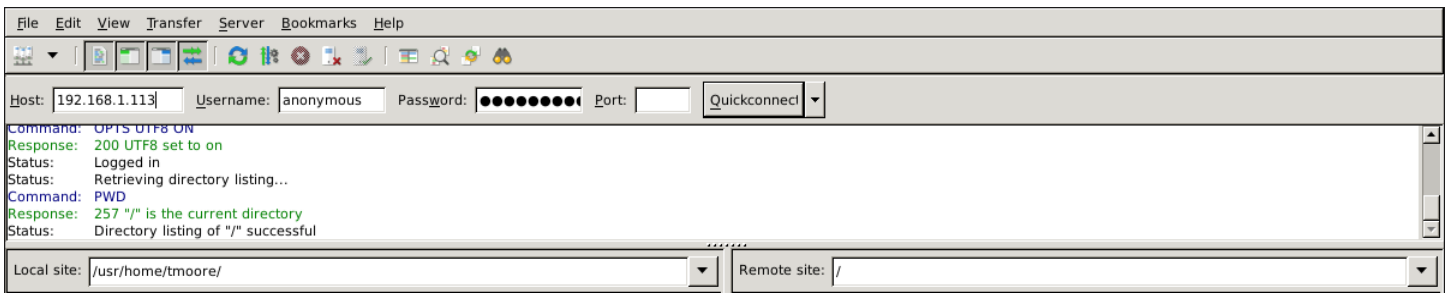


Fig. 12.6: Connecting Using Filezilla

The messages within the client indicate the FTP connection is successful. The user can now navigate the contents of the root folder on the remote site. This is the pool or dataset 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).

12.5.2 FTP in chroot

If users are required to authenticate before accessing the data on the FreeNAS® system, either create a user account for each user or import existing user accounts using [Active Directory](#) (page 183) or [LDAP](#) (page 188). Create a ZFS dataset for *each* user, then chroot each user so 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 a user home directory is limited to the size of the quota.

To configure this scenario:

1. Create a ZFS dataset for each user in *Storage* → *Pools*. Click the (Options) button, then *Add Dataset*. 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. When [Active Directory](#) (page 183) or [LDAP](#) (page 188) are not being used, create a user account for each user by navigating to *Accounts* → *Users*, and clicking *ADD*. 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 by navigating to *Storage* → *Pools*, and clicking the (Options) on the desired dataset. Click the *Edit Permissions* button, then assign a user account as *User* of that dataset. 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 Unix ACLs are always used, even if Windows clients will be accessing FreeNAS® via FTP.

4. Configure FTP in *Services* → *FTP* → *Configure* with these attributes:
 - *Path*: browse to the parent pool containing the datasets.
 - Make sure the options for *Allow Root Login* and *Allow Anonymous Login* are **unselected**.
 - Select the *Allow Local User Login* option to enable it.
 - Select the *Always Chroot* option to enable it.
5. Start the FTP service in *Services* → *FTP*. Click the sliding button on the *FTP* row. The FTP service takes a second or so to start. The sliding button moves to the right to show the service is running.
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 is associated with a dataset, and the *Password* for that user. The messages will indicate 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 pool but the dataset created for that user. The user can transfer files between the local site (their system) and the remote site (their dataset on the FreeNAS® system).

12.5.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

1. Import or create a certificate authority using the instructions in [CAs](#) (page 97). Then, import or create the certificate to use for encrypted connections using the instructions in [Certificates](#) (page 101).
2. In *Services* → *FTP* → *Configure*, click *ADVANCED*, choose the certificate in *Certificate*, and set the *Enable TLS* option.
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 *SAVE* to accept the certificate and negotiate an encrypted connection.
4. To force encrypted connections, select *On* for the *TLS Policy*.

12.5.4 Troubleshooting FTP

The FTP service will not start if it cannot resolve the system hostname to an IP address with DNS. To see if the FTP service is running, open [Shell](#) (page 325) 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*, enable *Show console messages*, and click *SAVE*. Go to *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 hostname and IP address, or add an entry for the IP address of the FreeNAS® system in the *Network* → *Global Configuration Host name database* field.

12.6 iSCSI

Refer to [Block \(iSCSI\)](#) (page 223) for instructions on configuring iSCSI. Start the iSCSI service in *Services* by clicking the sliding button in the *iSCSI* row.

Note: A warning message is shown the iSCSI service stops when initiators are connected. Open the [Shell](#) (page 325) and type `ctladm islist` to determine the names of the connected initiators.

12.7 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 the network contains managed switches, configuring and starting the LLDP service will tell the FreeNAS® system to advertise itself on the network.

[Figure 12.7](#) shows the LLDP configuration screen and [Table 12.5](#) summarizes the configuration options for the LLDP service.

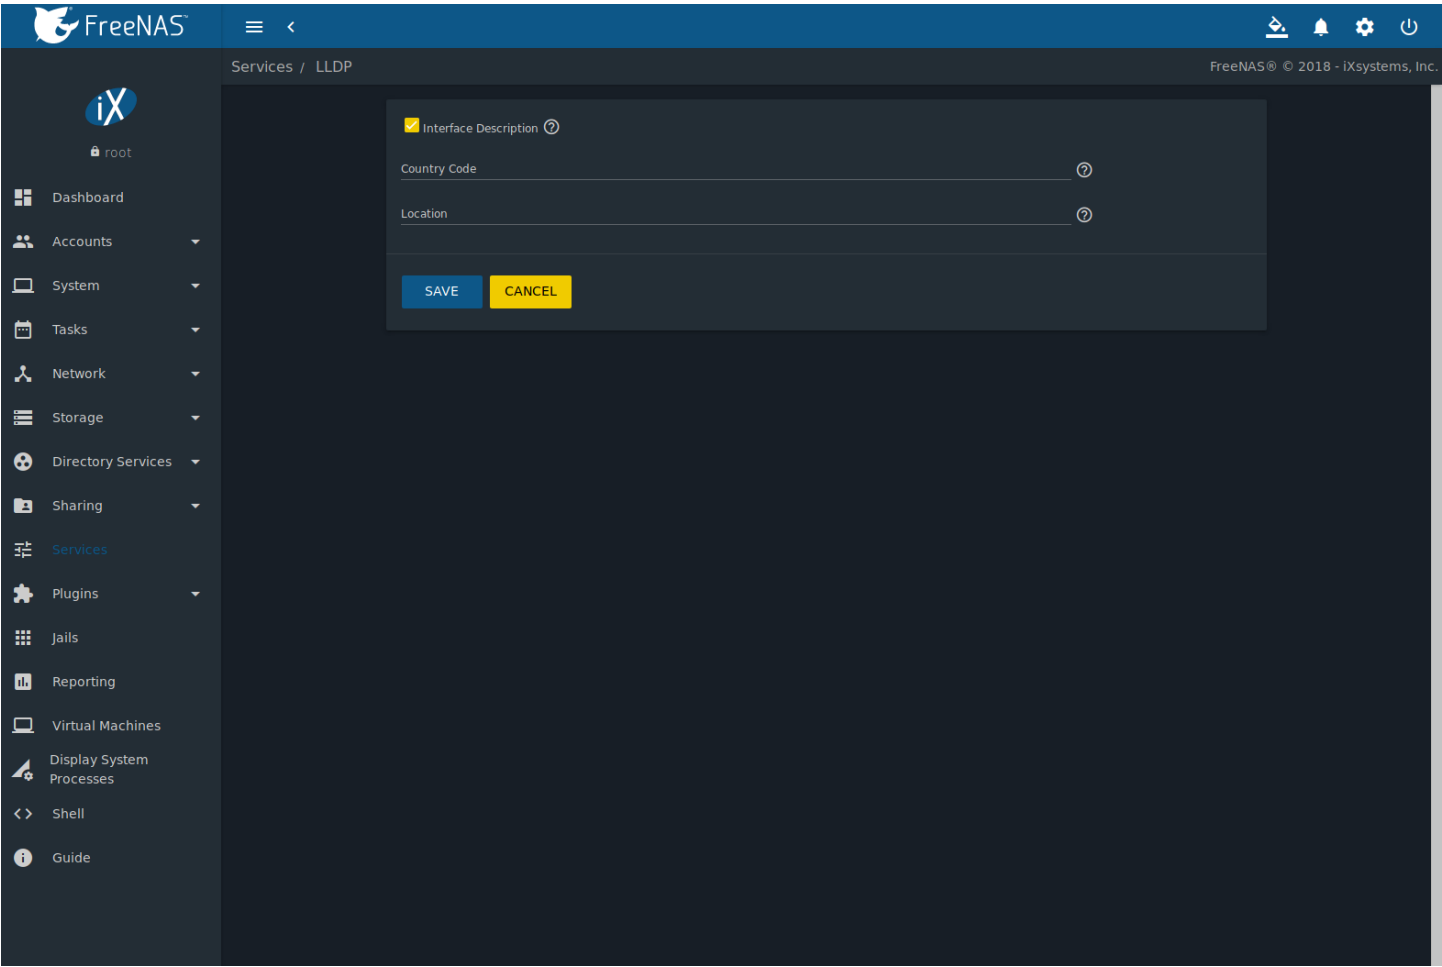


Fig. 12.7: Configuring LLDP

Table 12.5: LLDP Configuration Options

Setting	Value	Description
Interface De- scription	checkbox	Set to enable receive mode and to save and received peer information 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.

12.8 Netdata

Netdata is a real-time performance and monitoring system. It displays data as web dashboards.


Go to *Services* and click the sliding button in the *netdata* row to turn on the netdata service. Click  (Launch) to open the netdata web dashboard in a new browser tab. [Figure 12.8](#) shows an example:



Fig. 12.8: Netdata Web Dashboard

More information on configuring and using Netdata is available at the [Netdata website \(https://my-netdata.io/\)](https://my-netdata.io/).

12.9 NFS

The settings that are configured when creating NFS shares are specific to each configured NFS share. An NFS share is created by going to *Sharing* → *Unix (NFS) Shares* and clicking *ADD*. Global settings which apply to all NFS shares are configured in *Services* → *NFS* → *Configure*.

Figure 12.9 shows the configuration screen and Table 12.6 summarizes the configuration options for the NFS service.

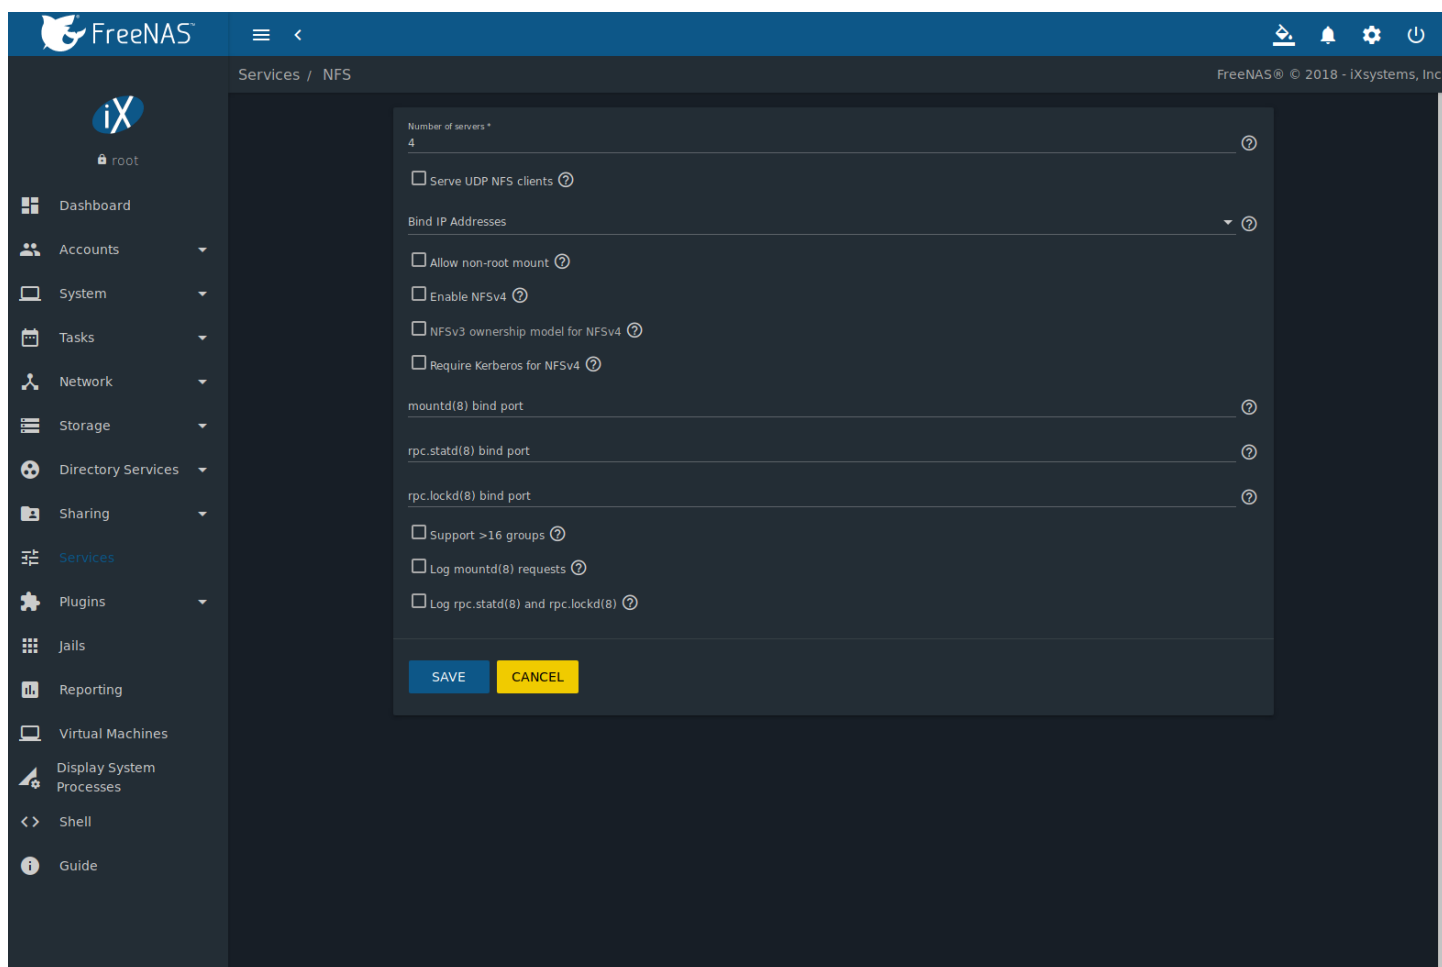


Fig. 12.9: Configuring NFS

Table 12.6: NFS Configuration Options

Setting	Value	Description
Number of servers	integer	Specify how many servers to create. Increase 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	Set if NFS clients need to use UDP.
Bind IP Addresses	drop-down	Select IP addresses to listen on for NFS requests. When all options are unset, NFS listens on all available addresses.
Allow non-root mount	checkbox	Set only if required by the NFS client.
Enable NFSv4	checkbox	Set to switch from NFSv3 to NFSv4. The default is NFSv3.
NFSv3 ownership model for NFSv4	checkbox	Grayed out unless <i>Enable NFSv4</i> is selected and, in turn, grays out <i>Support >16 groups</i> which is incompatible. Set this option if NFSv4 ACL support is needed without requiring the client and the server to sync users and groups.
Require Kerberos for NFSv4	checkbox	Set to force NFS shares to fail if the Kerberos ticket is unavailable.
mountd(8) bind port	integer	Optional. Specify the port that mountd(8) (https://www.freebsd.org/cgi/man.cgi?query=mountd) binds to.

Continued on next page

Table 12.6 – continued from previous page

Setting	Value	Description
rpc.statd(8) bind port	integer	Optional. Specify the port that rpc.statd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) binds to.
rpc.lockd(8) bind port	integer	Optional. Specify the port that rpc.lockd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) binds to.
Support >16 groups	checkbox	Set this option if any users are members of more than 16 groups (useful in AD environments). Note this assumes group membership is configured correctly on the NFS server.
Log mountd(8) requests	checkbox	Enable logging of mountd(8) (https://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) (https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) and rpc.lockd(8) (https://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.

12.10 Rsync

Services → *Rsync* is used to configure an rsync server when using rsync module mode. Refer to [Rsync Module Mode](#) (page 113) for a configuration example.

This section describes the configurable options for the `rsyncd` service and rsync modules.

12.10.1 Configure Rsyncd

Figure 12.10 shows the rsyncd configuration screen which is accessed from *Services* → *Rsync* → *Configure*.

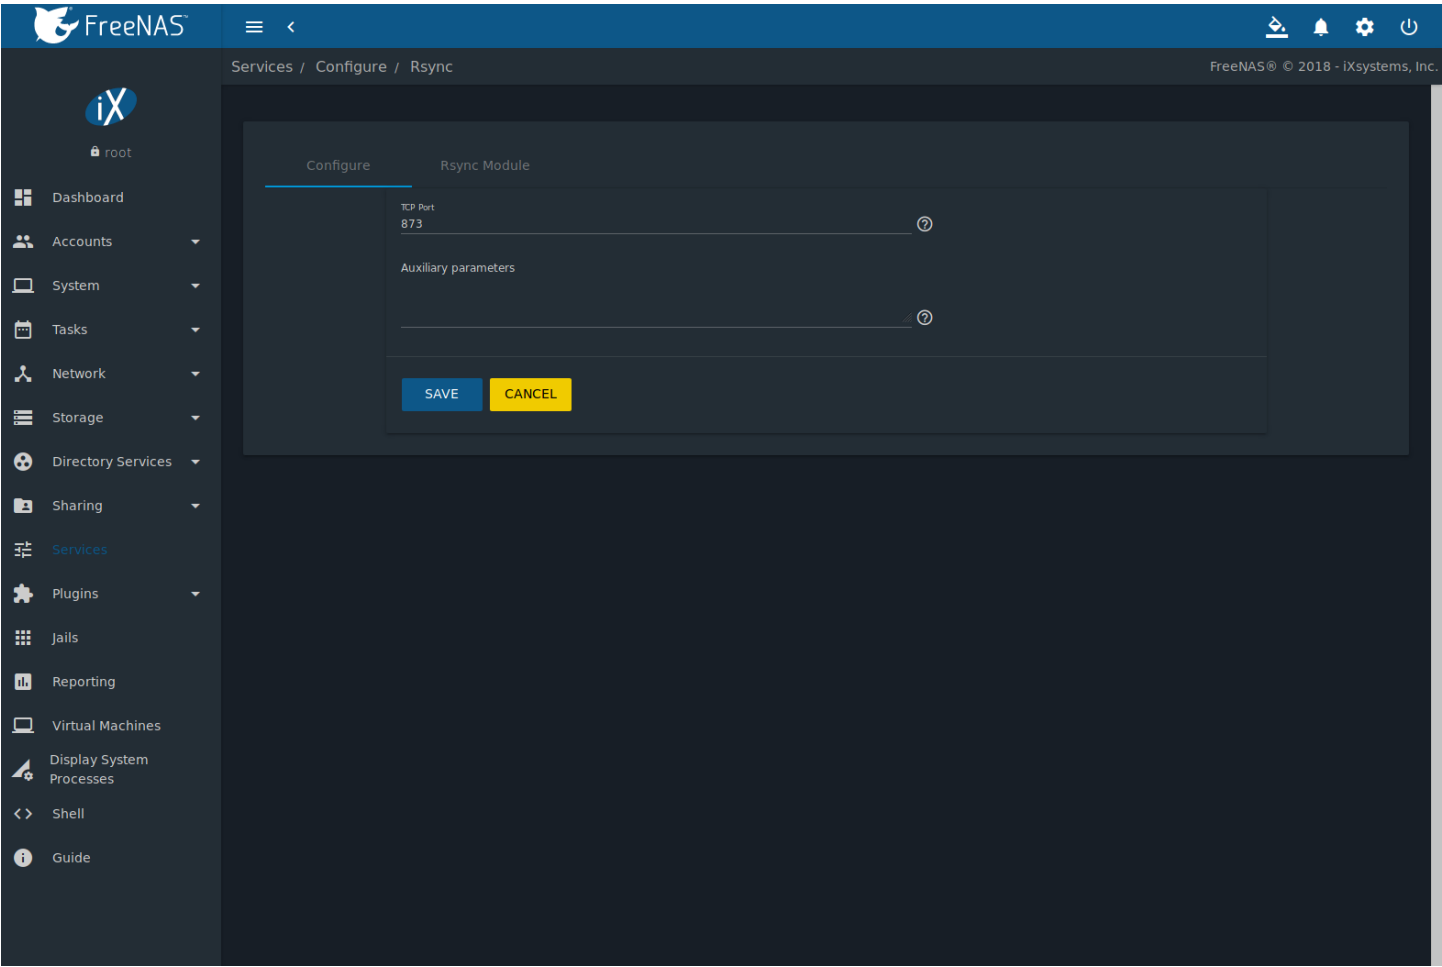


Fig. 12.10: Rsyncd Configuration

Table 12.7 summarizes the configuration options for the rsync daemon:

Table 12.7: Rsyncd Configuration Options

Setting	Value	Description
TCP Port	integer	Port for <code>rsyncd</code> to listen on, default is 873.
Auxiliary parameters	string	Enter any additional parameters from rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

12.10.2 Rsync Modules

Figure 12.11 shows the configuration screen that appears after navigating *Services* → *Rsync* → *Configure* → *Rsync Module*, and clicking *ADD*.

Table 12.8 summarizes the configuration options available when creating a rsync module.

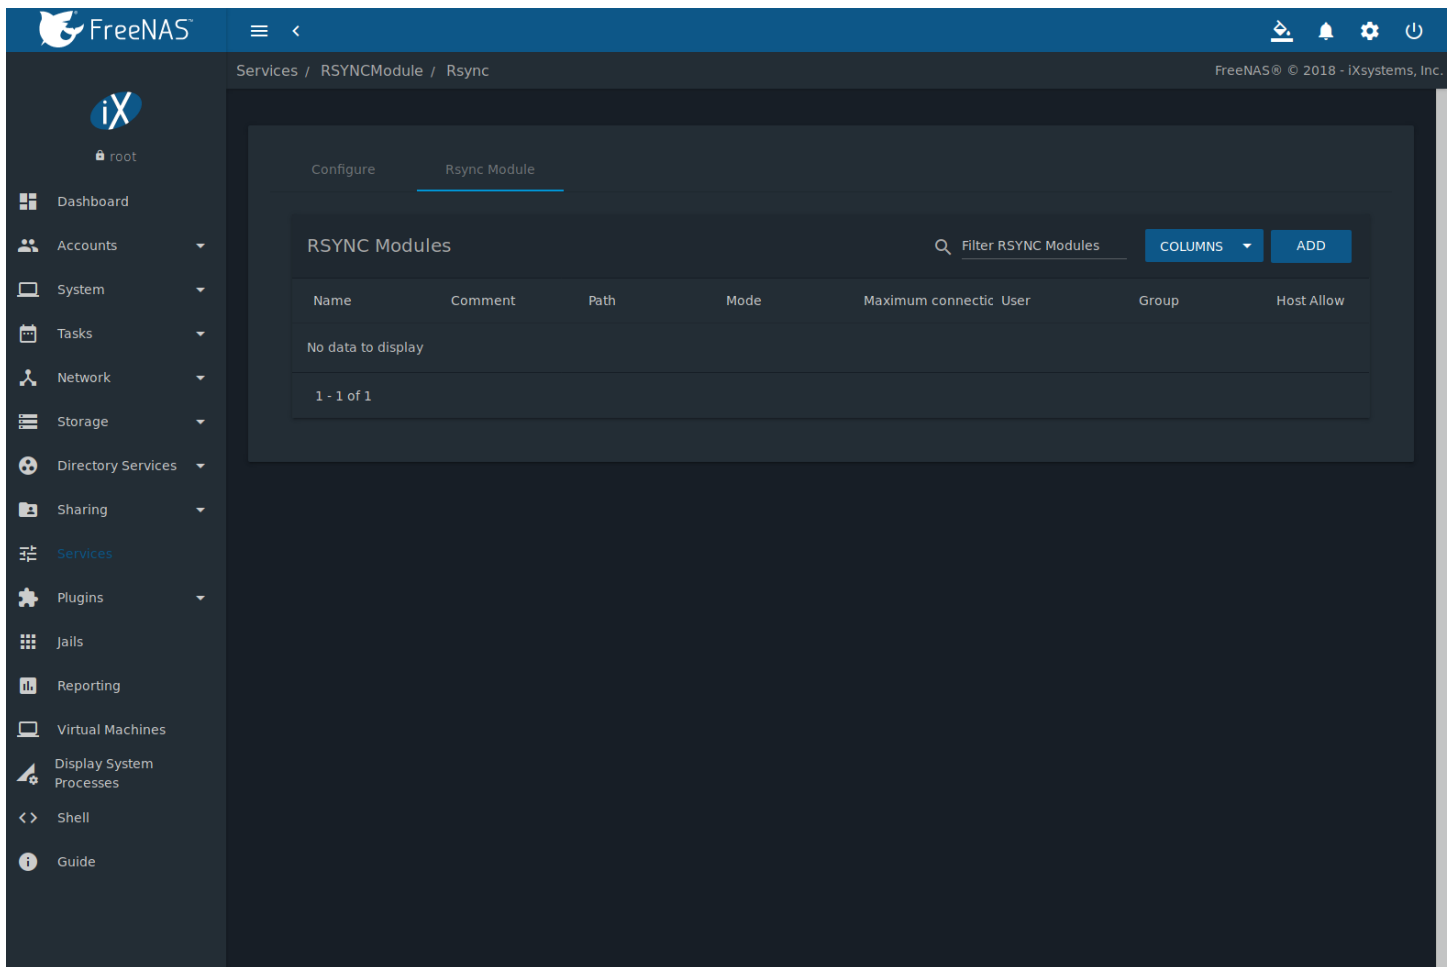


Fig. 12.11: Adding an Rsync Module

Table 12.8: Rsync Module Configuration Options

Setting	Value	Description
Name	string	This must match the setting on the rsync client.
Comment	string	Optional description.
Path	browse button	Browse to the pool or dataset to store the received data.
Access Mode	drop-down menu	Choices are <i>Read and Write</i> , <i>Read-only</i> , or <i>Write-only</i> .
Maximum connections	integer	0 is unlimited.
User	drop-down menu	Select the user to control file transfers to and from the module.
Group	drop-down menu	Select the group to control file transfers to and from the module.
Hosts Allow	string	From rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf). Enter a list of patterns to match with the hostname and IP address of a connecting client. Separate patterns with whitespace or comma.
Hosts Deny	string	See rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf) for allowed formats.
Auxiliary parameters	string	Enter any additional parameters from rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf)

12.11 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 12.12 shows the S3 service configuration screen and Table 12.9 summarizes the configuration options. After configuring the S3 service, start it in *Services*.

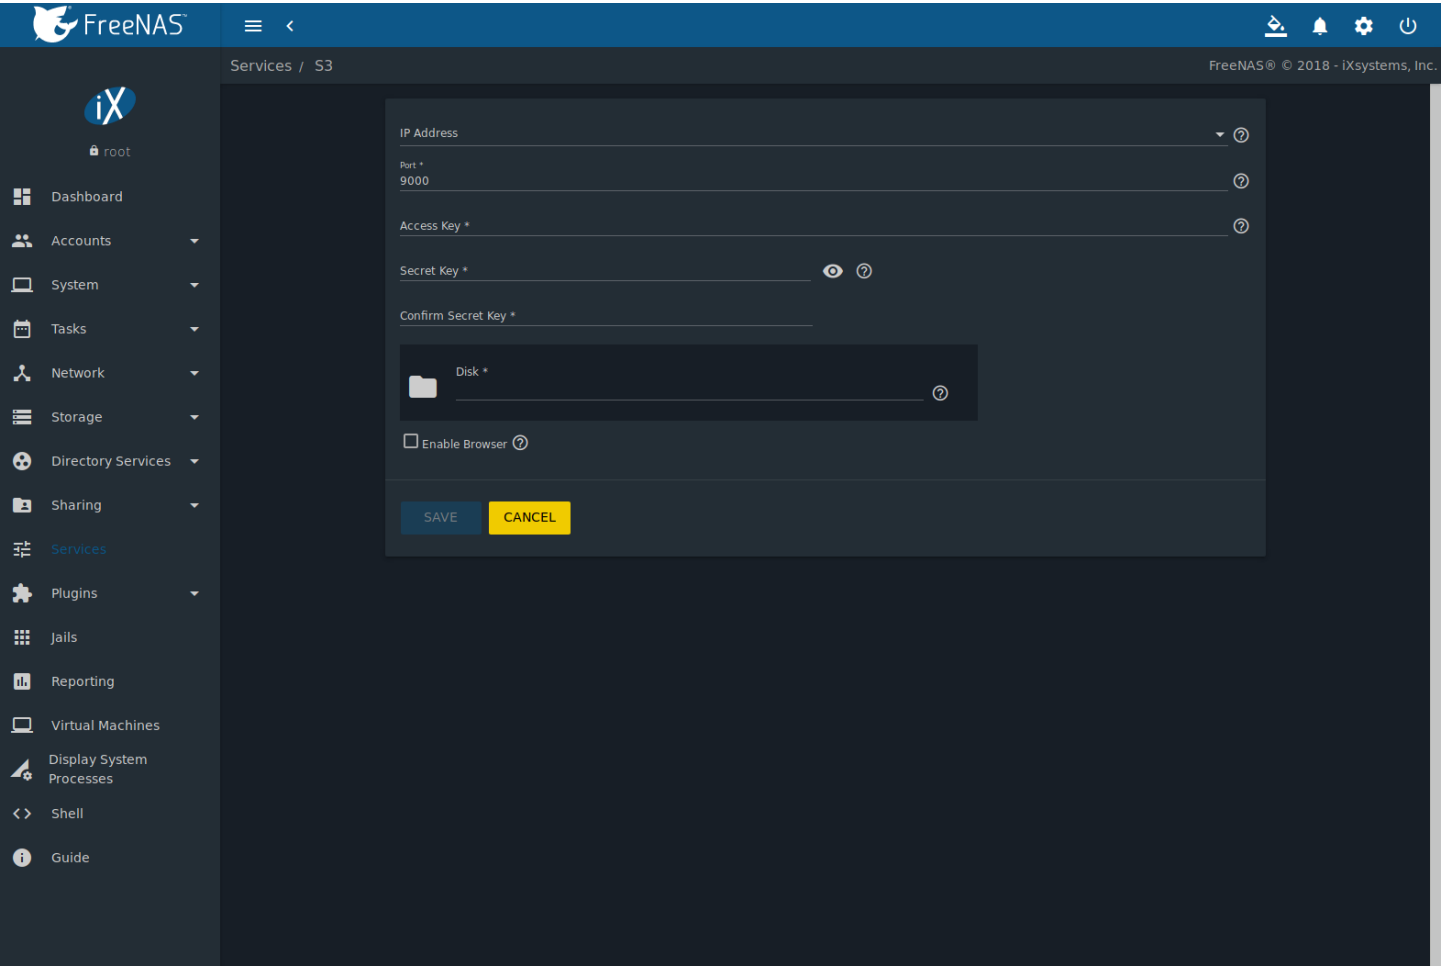


Fig. 12.12: Configuring S3

Table 12.9: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	Enter the IP address to run the S3 service. 0.0.0.0 sets the server to listen on all addresses.
Port	string	Enter the TCP port on which to provide the S3 service. Default is 9000.
Access Key	string	Enter the S3 user name. This username must contain only alphanumeric characters and be between 5 and 20 characters long.
Secret Key	string	Enter the password to be used by connecting S3 systems. The key must contain only alphanumeric characters and be at least 8 but no more than 40 characters long.
Confirm Secret Key	string	Re-enter the S3 password to confirm.

Continued on next page

Table 12.9 – continued from previous page

Setting	Value	Description
Disk	browse	Browse to the directory for the S3 filesystem. This overrides all permissions of that directory and all of its subdirectories to <i>minio:minio</i> . Create a separate dataset (page 166) for Minio to avoid any issues with directory permissions.
Enable Browser	checkbox	Set to enable the web user interface for the S3 service.
Certificate	drop-down menu	Add the SSL certificate (page 101) to be used for secure S3 connections.

12.12 S.M.A.R.T.

S.M.A.R.T., or [Self-Monitoring, Analysis, and Reporting Technology](https://en.wikipedia.org/wiki/S.M.A.R.T.) (<https://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 [Scrub Tasks](#) (page 132).

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\(8\)](https://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in) (<https://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. S.M.A.R.T. tests are created by navigating to *Tasks* → *S.M.A.R.T. Tests*, and clicking *ADD*.
2. Enabling or disabling S.M.A.R.T. for each disk member of a pool in *Storage* → *Pools*. 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 in *Services*.

Figure 12.13 shows the configuration screen that appears after clicking *Services* → *S.M.A.R.T* → *Configure*.

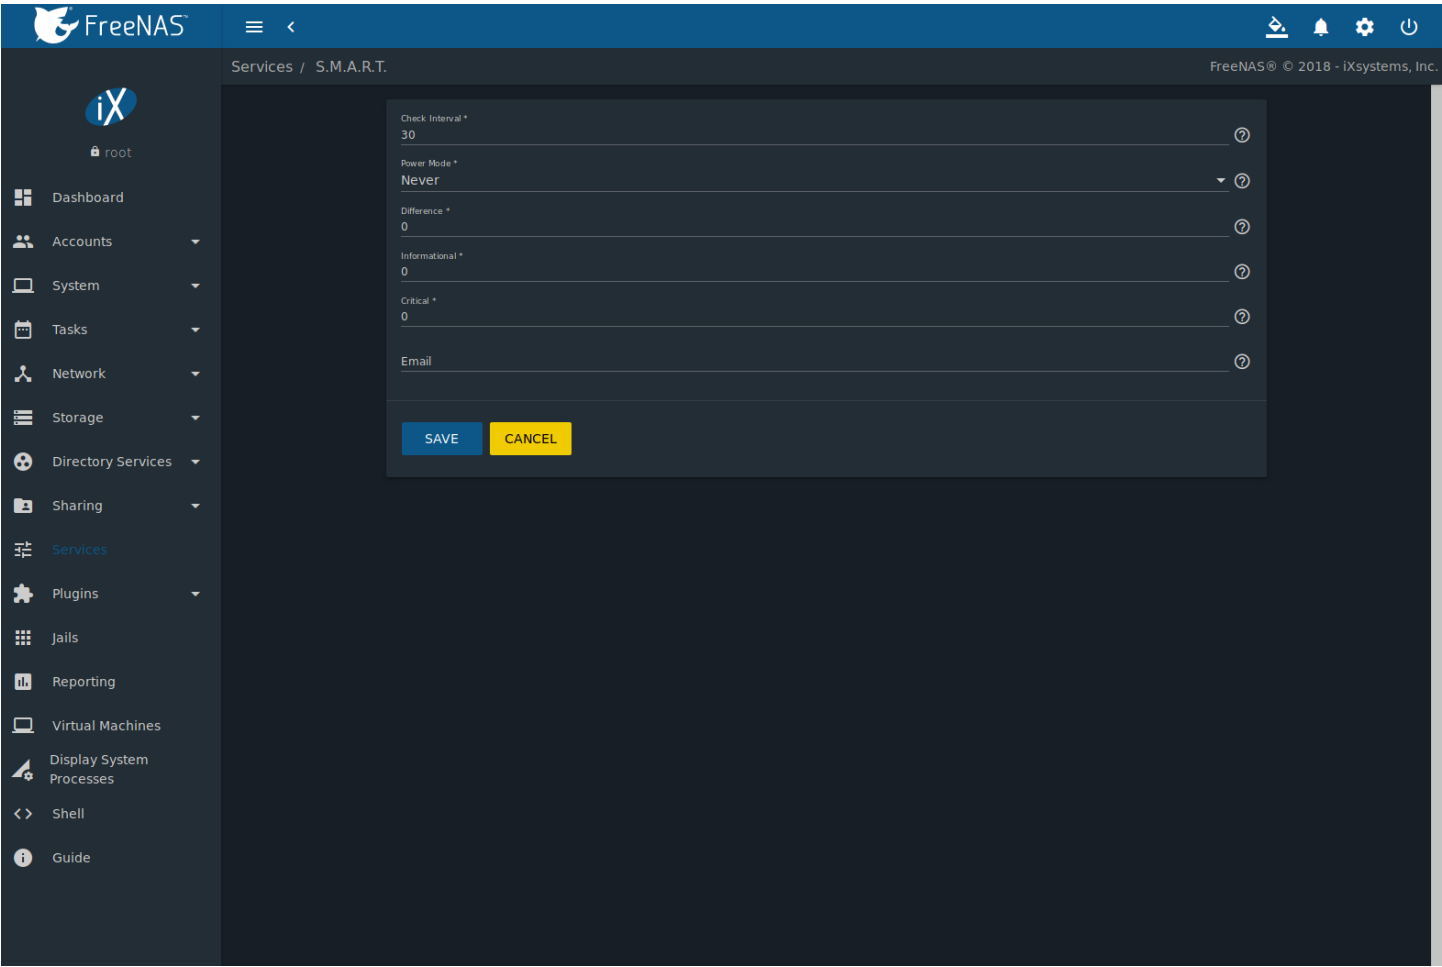


Fig. 12.13: S.M.A.R.T Configuration Options

Note: `smartd` wakes up at the configured *Check Interval*. It checks the times configured in *Tasks* → *S.M.A.R.T. Tests* to see if a test must begin. Since the smallest time increment for a test is an hour, 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 activates every two hours.

Table 12.10 summarizes the options in the S.M.A.R.T configuration screen.

Table 12.10: S.M.A.R.T Configuration Options

Setting	Value	Description
Check Interval	integer	Define in minutes how often <code>smartd</code> activates to check if any tests are configured to run.
Power Mode	drop-down menu	Tests are not performed if the system enters the specified power mode. Choices are: <i>Never</i> , <i>Sleep</i> , <i>Standby</i> , or <i>Idle</i> .
Difference	integer in degrees Celsius	Enter number of degrees in Celsius. S.M.A.R.T reports if the temperature of a drive has changed by N degrees Celsius since the last report. Default of 0 disables this option.
Informational	integer in degrees Celsius	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a log level of LOG_INFO if the temperature is higher than the threshold. Default of 0 disables this option.

Continued on next page

Table 12.10 – continued from previous page

Setting	Value	Description
Critical	integer in degrees Celsius	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a log level of LOG_CRIT and send an email if the temperature is higher than the threshold. Default of 0 disables this option.
Email	string	Enter email address to receive S.M.A.R.T. alerts. Use a space to separate multiple email addresses.

12.13 SMB

The settings configured when creating SMB shares are specific to each configured SMB share. An SMB share is created by navigating to *Sharing* → *Windows (SMB) Shares*, and clicking *ADD*. In contrast, global settings which apply to all SMB shares are configured in *Services* → *SMB* → *Configure*.

Note: After starting the SMB service, it can take several minutes for the *master browser election* (<https://www.samba.org/samba/docs/old/Samba3-HOWTO/NetworkBrowsing.html#id2581357>) to occur and for the FreeNAS® system to become available in Windows Explorer.

Figure 12.14 shows the global SMB configuration options which are described in Table 12.11. This configuration screen is really a front-end to `smb4.conf` (<https://www.freebsd.org/cgi/man.cgi?query=smb4.conf>).

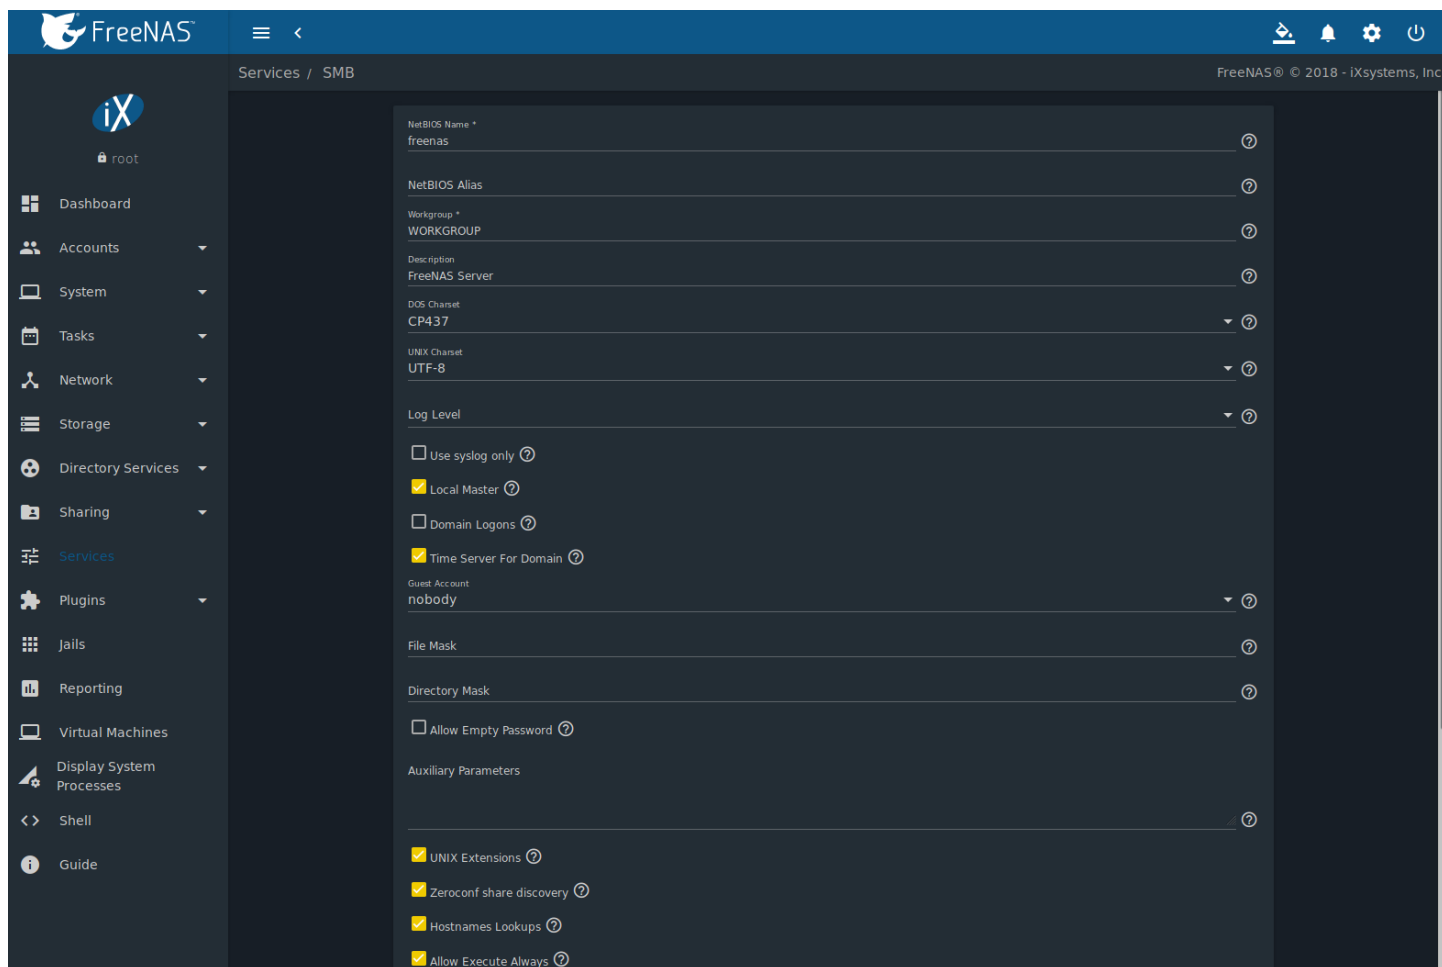


Fig. 12.14: Global SMB Configuration

Table 12.11: Global SMB Configuration Options

Setting	Value	Description
NetBIOS Name	string	Automatically populated with the original hostname of the system. Limited to 15 characters. It must be different from the <i>Workgroup</i> name.
NetBIOS Alias	string	Enter an alias. Limited to 15 characters.
Workgroup	string	Must match the Windows workgroup name. This setting is ignored if the Active Directory (page 183) or LDAP (page 188) service is running.
Description	string	Enter a server description. Optional.
DOS Charset	drop-down menu	The character set Samba uses when communicating with DOS and Windows 9x/ME clients. Default is <i>CP437</i> .
UNIX Charset	drop-down menu	Default is <i>UTF-8</i> which supports all characters in all languages.
Log Level	drop-down menu	Choices are <i>Minimum</i> , <i>Normal</i> , or <i>Debug</i> .
Use syslog only	checkbox	Set to log authentication failures in <code>/var/log/messages</code> instead of the default of <code>/var/log/samba4/log.smbd</code> .
Local Master	checkbox	Set to determine if the system participates in a browser election. Disable when network contains an AD or LDAP server or Vista or Windows 7 machines are present.
Domain Logons	checkbox	Set if it is necessary to provide netlogin service for older Windows clients.
Time Server for Domain	checkbox	Set to determine if the system advertises itself as a time server to Windows clients. Disable when network contains an AD or LDAP server.
Guest Account	drop-down menu	Select the account to be used for guest access. Default is <i>nobody</i> . Account must have permission to access the shared pool or dataset. If Guest Account user is deleted, resets to <i>nobody</i> .
File Mask	integer	Overrides default file creation mask of <i>0666</i> which creates files with read and write access for everybody.
Directory Mask	integer	Overrides default directory creation mask of <i>0777</i> which grants directory read, write and execute access for everybody.
Allow Empty Password	checkbox	Set to allow users to press <code>Enter</code> when prompted for a password. Requires the username/password be the same as the Windows user account.
Auxiliary Parameters	string	Add any <code>smb.conf</code> options not covered elsewhere in this screen. See the Samba Guide (https://www.oreilly.com/openbook/samba/book/appb_02.html) for additional settings.
UNIX Extensions	checkbox	Set to allow 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	Set to allow using hostnames rather than IP addresses in the <i>Hosts Allow</i> or <i>Hosts Deny</i> fields of a SMB share. Unset if IP addresses are used to avoid the delay of a host lookup.
Allow Execute Always	checkbox	When set, Samba allows the user to execute a file, even if that user's permissions are not set to execute.
Obey Pam Restrictions	checkbox	Unset to allow cross-domain authentication, and users and groups to be managed on another forest. Unsetting this option also allows permissions to be delegated from Active Directory (page 183) users and groups to domain admins on another forest.
NTLMv1 Auth	checkbox	Set to allow NTLMv1 authentication. Required by Windows XP clients and sometimes by clients in later versions of Windows.
Bind IP Addresses	checkboxes	Select the IP addresses SMB will listen for. Both IPv4 and IPv6 addresses are supported.
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.

Continued on next page

Table 12.11 – continued from previous page

Setting	Value	Description
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 258) cannot be disabled while *Active Directory* (page 183) is enabled.

12.13.1 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 pool or dataset being shared by SMB and the share becomes inaccessible, log out and back in to 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 they access the system, 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](https://www.oreilly.com/openbook/samba/book/ch05_04.html) (https://www.oreilly.com/openbook/samba/book/ch05_04.html) explains in more detail.

If a particular user cannot connect to a SMB share, ensure 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 macOS users, try disabling *UNIX Extensions* and restarting the SMB service.

If the SMB service will not start, run this command from *Shell* (page 325) 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* → *Configure* and verify that *Server maximum protocol* is set to *SMB2*.

Using a dataset for SMB sharing is recommended. 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 the ACLs are already destroyed by using `chmod`, `winacl` can be used to fix them. Type `winacl` from *Shell* (page 325) for usage instructions.

The [Common Errors](https://www.samba.org/samba/docs/old/Samba3-HOWTO/domain-member.html#id2573692) (https://www.samba.org/samba/docs/old/Samba3-HOWTO/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) (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 259).
- 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 [Auxiliary Parameters](#) (page 213) for the share.
- Disable as many *VFS Objects* as possible in the [share settings](#) (page 213). Many have performance overhead.

12.14 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/) (http://net-snmp.sourceforge.net/) to provide SNMP. When starting the SNMP service, this 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 12.15](#) shows the *Services* → *SNMP* → *Configure* screen. [Table 12.12](#) summarizes the configuration options.

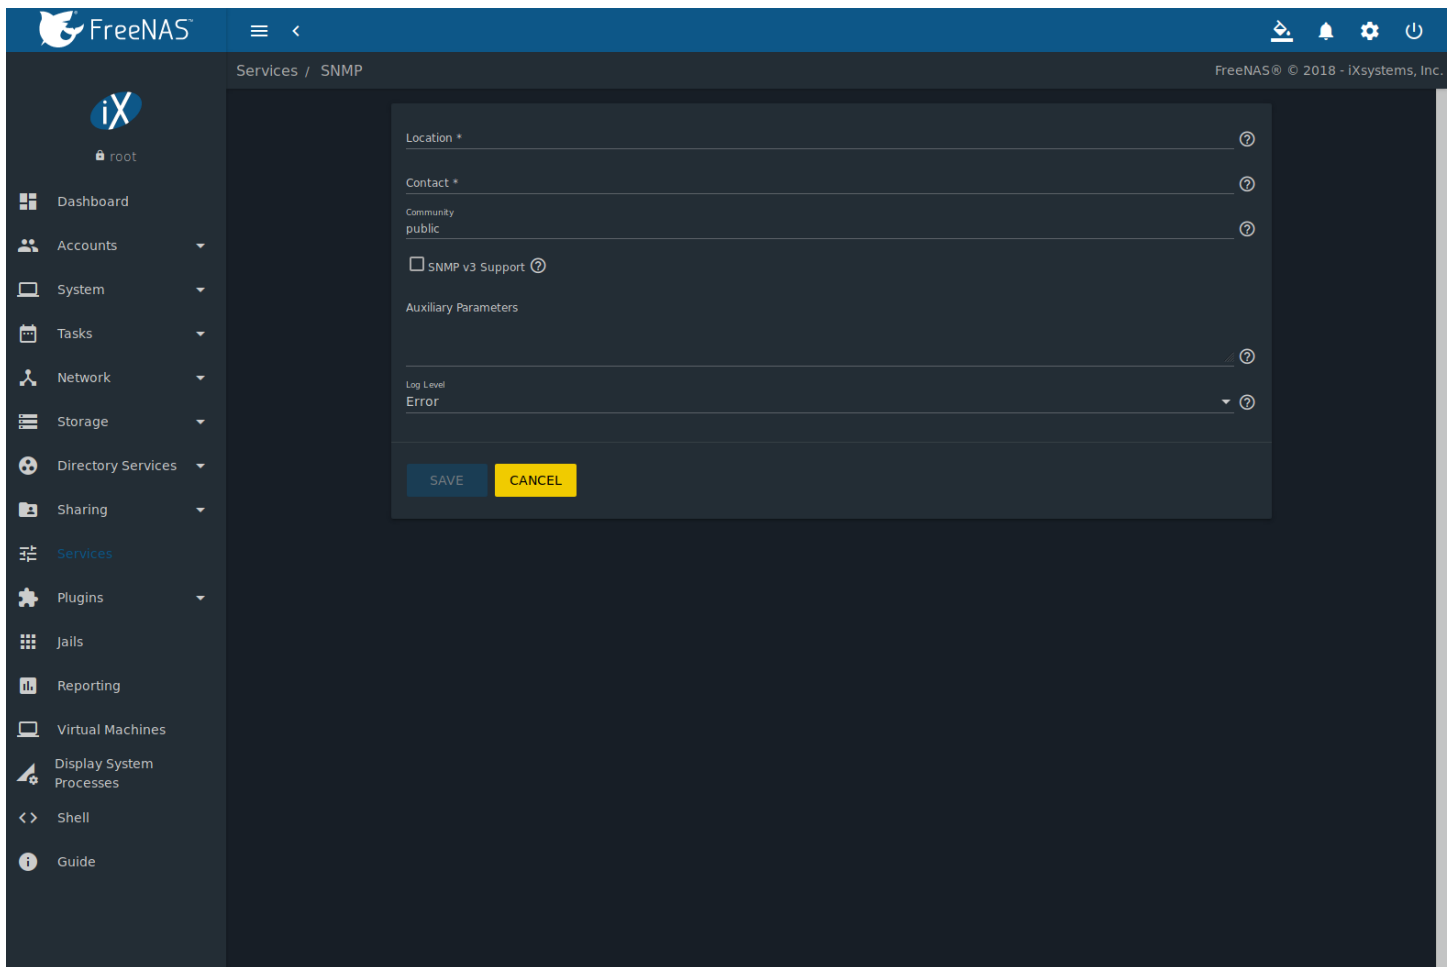


Fig. 12.15: Configuring SNMP

Table 12.12: SNMP Configuration Options

Setting	Value	Description
Location	string	Optional description of the system location.
Contact	string	Optional. Enter the administrator email address.
Community	string	Default is <i>public</i> . Change this for security reasons! The value can only contain alphanumeric characters, underscores, dashes, periods, and spaces. Leave empty for SNMPv3 networks.
SNMP v3 Support	checkbox	Set to enable support for SNMP version 3.
Username	string	Only applies if <i>SNMP v3 Support</i> is set. 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 about configuring this and the <i>Authentication Type</i> , <i>Password</i> , <i>Privacy Protocol</i> , and <i>Privacy Passphrase</i> fields.
Authentication Type	drop-down menu	Only applies if <i>SNMP v3 Support</i> is enabled. Choices are <i>MD5</i> or <i>SHA</i> .
Password	string	Only applies if <i>SNMP v3 Support</i> is enabled. Enter and confirm a password of at least eight characters.
Privacy Protocol	drop-down menu	Only applies if <i>SNMP v3 Support</i> is enabled. Choices are <i>AES</i> or <i>DES</i> .
Privacy Passphrase	string	If not specified, <i>Password</i> is used.

Continued on next page

Table 12.12 – continued from previous page

Setting	Value	Description
Auxiliary Parameters	string	Enter additional snmpd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=snmpd.conf) options. Add one option for each line.
Log Level	drop-down menu	Choices range from the least log entries (<i>Emergency</i>) to the most (<i>Debug</i>)

12.15 SSH

Secure Shell (SSH) is used to transfer files securely over an encrypted network. When a FreeNAS® system is used as an SSH server, the users in the network must 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 12.16 shows the *Services* → *SSH* → *Configure* screen.

Note: After configuring SSH, remember to start it in *Services* by clicking the sliding button in the *SSH* row. The sliding button moves to the right when the service is running.

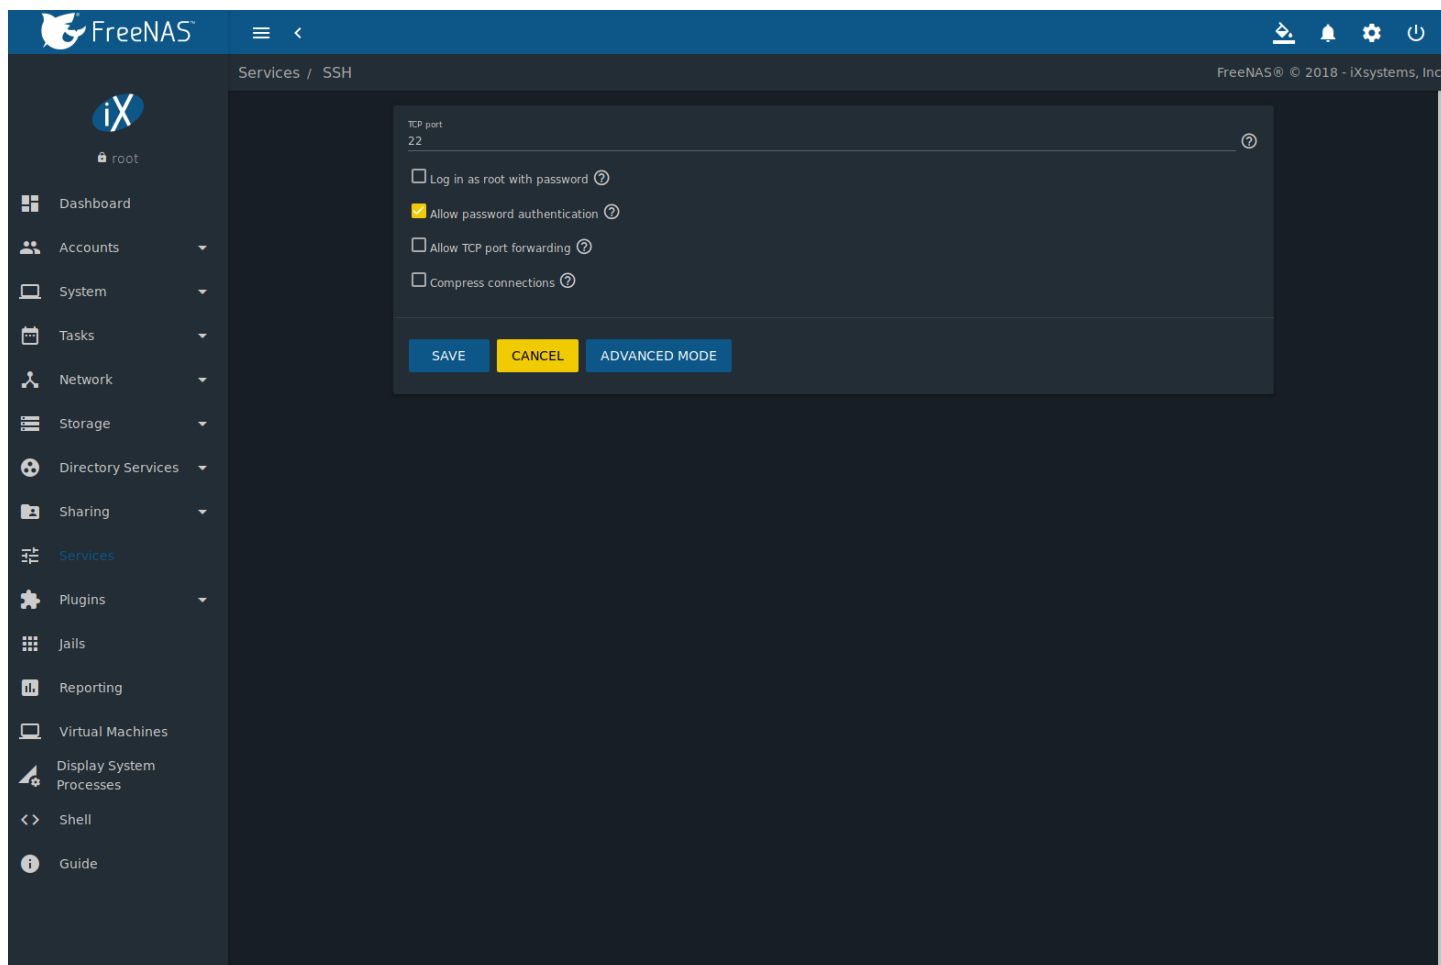


Fig. 12.16: SSH Configuration

Table 12.13 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 enabling the *Show advanced fields by default* option in *System* → *Advanced*.

Table 12.13: SSH Configuration Options


Setting	Value	Advanced Mode	Description
Bind interfaces	selection	✓	By default, SSH listens on all interfaces unless specific interfaces are selected in this drop-down menu.
TCP port	integer		Port to open for SSH connection requests. 22 by default.
Log in as root with password	checkbox		As a security precaution, root logins are discouraged and disabled by default. If enabled, password must be set for the <i>root</i> user in <i>Users</i> .
Allow password authentication	checkbox		Unset to require key-based authentication for all users. This requires additional setup (http://the.earth.li/~sgtatham/putty/0.55/htmldoc/Chapter8.html) on both the SSH client and server.
Allow kerberos authentication	checkbox	✓	Ensure Kerberos Realms (page 192) and Kerberos Keytabs (page 193) are configured and FreeNAS® can communicate with the Kerberos Domain Controller (KDC) before enabling this option.
Allow TCP port forwarding	checkbox		Set to allow users to bypass firewall restrictions using the SSH port forwarding feature (https://www.symantec.com/connect/articles/ssh-port-forwarding).
Compress connections	checkbox		Set to attempt to reduce latency over slow networks.
SFTP log level	drop-down menu	✓	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) level of the SFTP server.
SFTP log facility	drop-down menu	✓	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) facility of the SFTP server.
Extra options	string	✓	Add any additional sshd_config(5) (https://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 can prevent the SSH service from starting.

A few [sshd_config\(5\)](#) (https://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 more concurrent SSH connections are required.

12.15.1 SCP Only

When SSH is configured, authenticated users with a user account can use `ssh` to log into the FreeNAS® system over the network. User accounts are created by navigating to *Accounts* → *Users*, and clicking **ADD**. The user home directory is the pool or dataset specified in the *Home Directory* field of the FreeNAS® account for that user. While the SSH login defaults to the user home directory, users are able to navigate outside their home directory, which can pose a security risk.

It is possible to allow users to use `scp` and `sftp` 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 *Accounts* → *Users*, click  (Options) for the user, and then **Edit**. Change the *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. `sftp` and `scp` will work but `ssh` will fail.

Note: Some utilities like WinSCP and Filezilla can bypass the `scponly` shell. This section assumes users are accessing the system using the command line versions of `scp` and `sftp`.

12.15.2 Troubleshooting SSH

Keywords listed in `sshd_config(5)` (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) are case sensitive. This is important to remember when adding any *Extra options*. The configuration will not function as intended if the upper and lowercase letters of the keyword are not an exact match.

If 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 the configuration as an SSH user account to ensure the user is limited by the configuration and they have permission to transfer files within the intended directories. If the user account is experiencing problems, the SSH error messages are specific in describing the problem. Type this command within *Shell* (page 325) to read these messages as they occur:

```
tail -f /var/log/messages
```

Additional messages regarding authentication errors are found in `/var/log/auth.log`.

12.16 TFTP

Trivial File Transfer Protocol (TFTP) is a light-weight version of FTP typically 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 network devices, configure and start the TFTP service. Starting the TFTP service opens UDP port 69.

Figure 12.17 shows the TFTP configuration screen and Table 12.14 summarizes the available options.

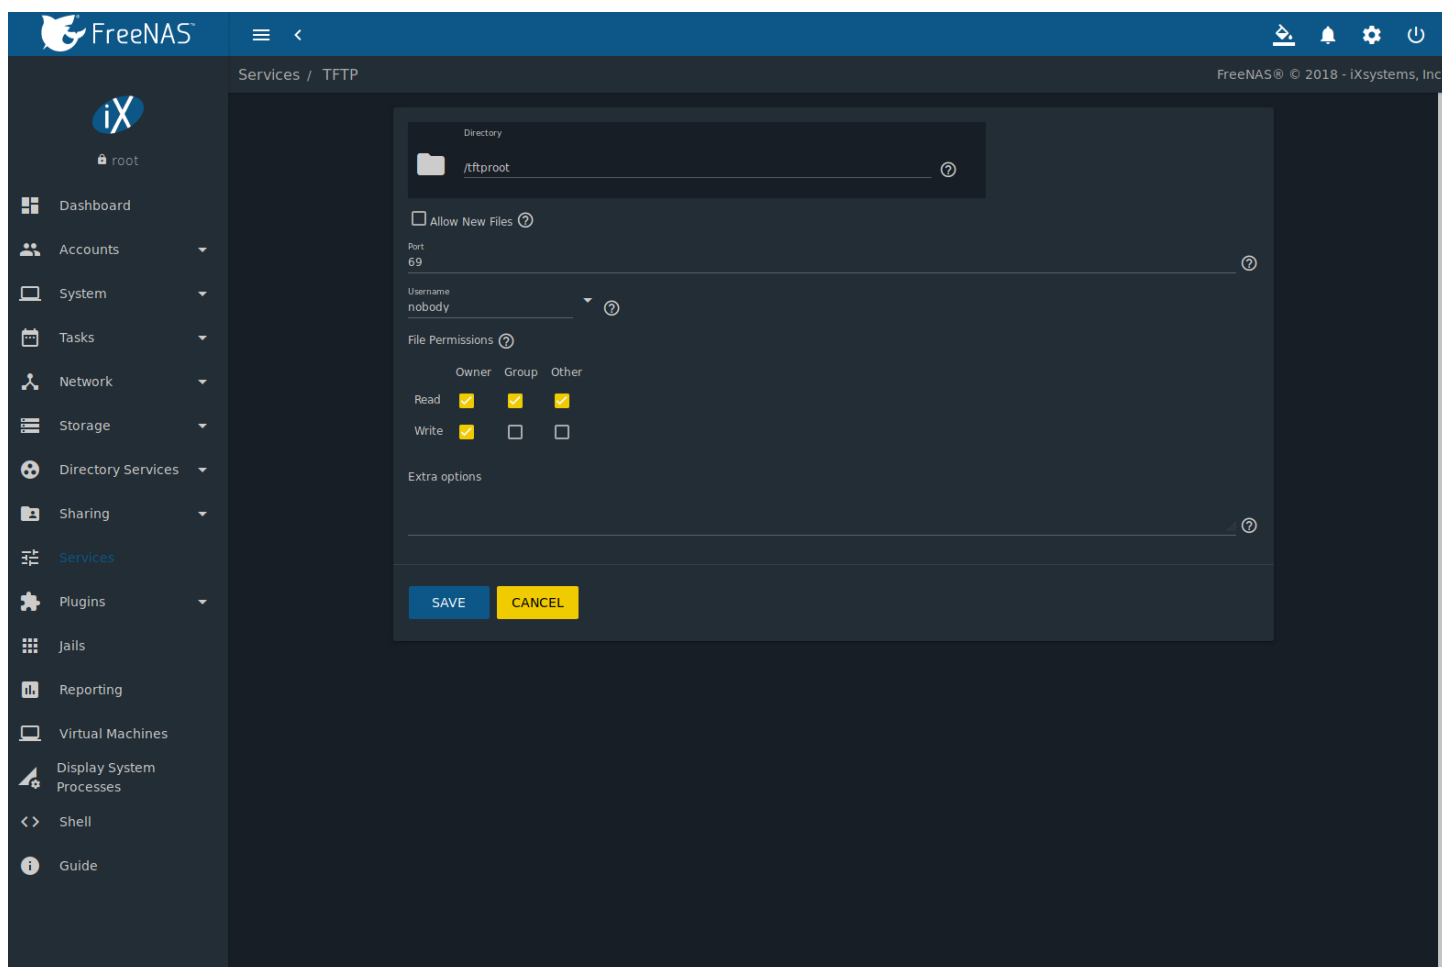


Fig. 12.17: TFTP Configuration

Table 12.14: 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 documentation for details.
Allow New Files	checkbox	Set when network devices need to send files to the system. For example, to back up their configuration.
Port	integer	Enter a UDP port to listen for TFTP requests, <i>69</i> by default.
Username	drop-down menu	Select the account to use for TFTP requests. This account must have permission to the <i>Directory</i> .
File Permissions	checkboxes	Set permissions for newly created files. The default is everyone can read and only the owner can write. Some devices require less strict permissions.
Extra options	string	Add more options from tftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=tftpd) Add one option on each line.

12.17 UPS

FreeNAS® uses [NUT](https://networkupstools.org/) (<https://networkupstools.org/>) (Network UPS Tools) to provide UPS support. If the FreeNAS® system is connected to a UPS device, configure the UPS service in *Services* → *UPS* → *Configure*.

Figure 12.18 shows the UPS configuration screen:

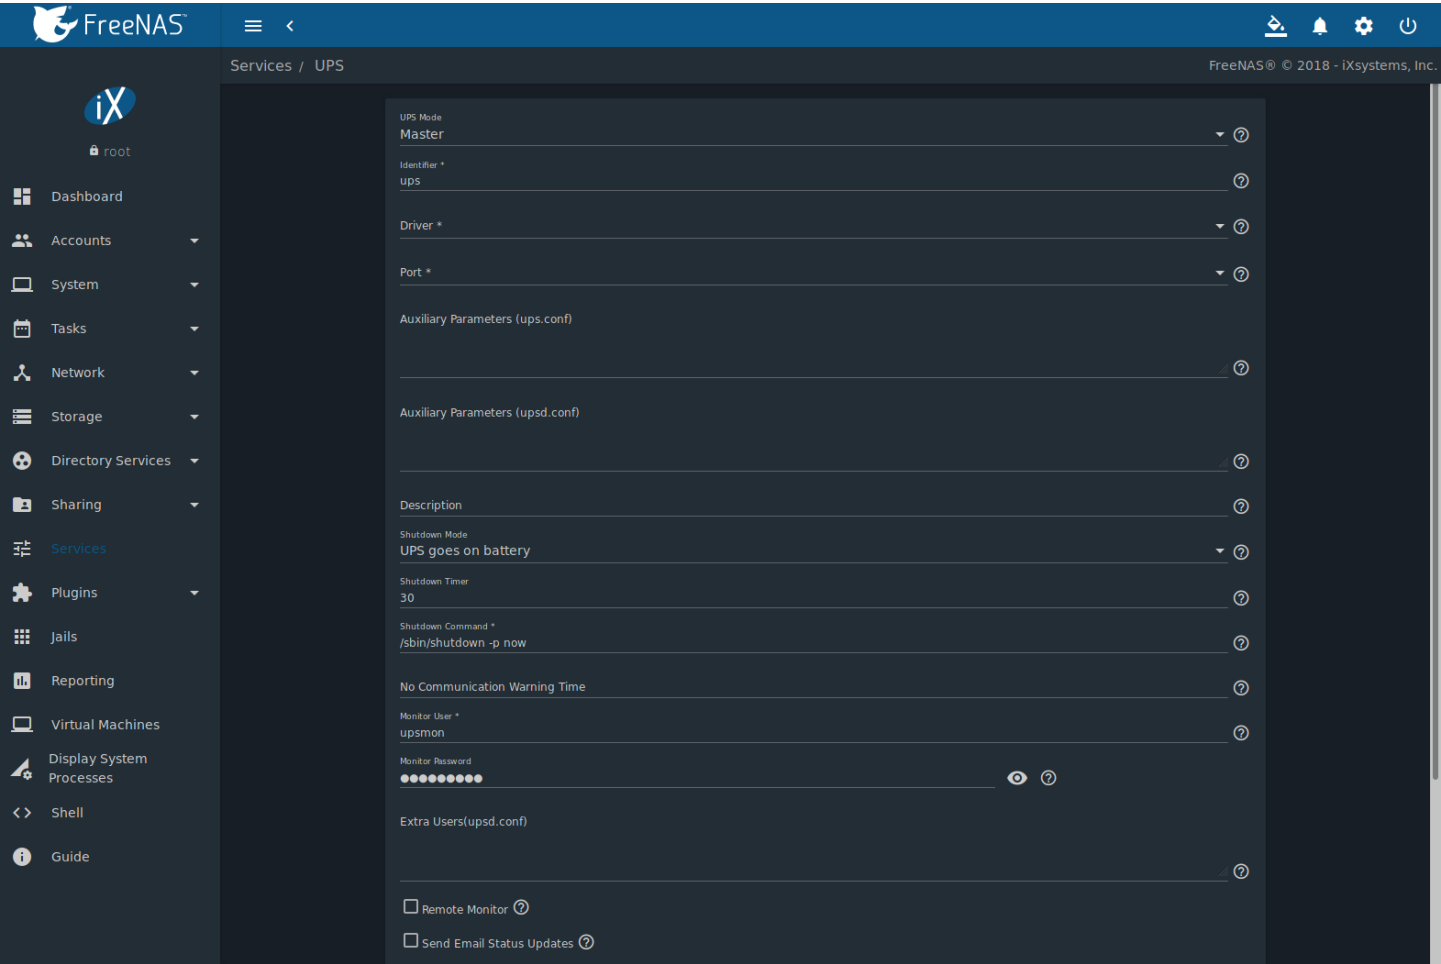


Fig. 12.18: UPS Configuration Screen

Table 12.15 summarizes the options in the UPS Configuration screen.

Table 12.15: UPS Configuration Options

Setting	Value	Description
UPS Mode	drop-down menu	Select <i>Master</i> if the UPS is plugged directly into the system serial port. The UPS will remain the last item to shut down. Select <i>Slave</i> to have the system shut down before <i>Master</i> .
Identifier	string	Describe the UPS device. Can contain alphanumeric, period, comma, hyphen, and underscore characters.

Continued on next page

Table 12.15 – continued from previous page

Setting	Value	Description
Driver / Remote Host	drop-down menu	For a list of supported devices, see the Network UPS Tools compatibility list (https://networkupstools.org/stable-hcl.html). <i>Driver</i> becomes <i>Remote Host</i> when <i>UPS Mode</i> is set to <i>Slave</i> . The IP address of the system configured as the <i>UPS Master</i> system. See this post (https://forums.freenas.org/index.php?resources/configuring-ups-support-for-single-or-multiple-freenas-servers.30/) for more details about configuring multiple systems with a single UPS.
Port / Hostname / Remote Port	drop-down menu	<i>Port</i> : Enter the serial or USB port the UPS is plugged into (see NOTE (page 268)). Choosing an <code>snmp</code> value in the <i>Driver</i> field changes the <i>Port</i> field to <i>Host-name</i> : Enter the IP address or hostname of the SNMP UPS device. <i>Port</i> becomes <i>Remote Port</i> when the <i>UPS Mode</i> is set to <i>Slave</i> . The open network port number of the <i>UPS Master</i> system. The default port is 3493.
Auxiliary Parameters (ups.conf)	string	Enter any additional options from ups.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=ups.conf).
Auxiliary Parameters (upsd.conf)	string	Enter any additional options from upsd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=upsd.conf).
Description	string	Optional. Describe the UPS service.
Shutdown Mode	drop-down menu	Choose when the UPS initiates shutdown. Choices are <i>UPS goes on battery</i> and <i>UPS reaches low battery</i> .
Shutdown Timer	integer	Select a value in seconds for the UPS to wait before initiating shutdown. Shutdown will not occur if the power is restored while the timer is counting down. This value only applies when <i>Shutdown Mode</i> is set to <i>UPS goes on battery</i> .
Shutdown Command	string	Enter the command to run to shut down the computer when battery power is low or shutdown timer runs out.
No Communication Warning Time	string	Enter a value in seconds to wait before alerting that the service cannot reach any UPS. Warnings continue until the situation is fixed.
Monitor User	string	Enter a user to associate with this service. The recommended default user is <i>upsmon</i> .
Monitor Password	string	Default is the known value <i>fixmepass</i> . Change this password to enhance system security. The new password cannot contain a space or #.
Extra Users(upsd.conf)	string	Enter the accounts with administrative access. See upsd.users(5) (https://www.freebsd.org/cgi/man.cgi?query=upsd.users) for examples.
Remote Monitor	checkbox	Set for the default configuration to listen on all interfaces using the known values of user: <i>upsmon</i> and password: <i>fixmepass</i> .
Send Email Status Updates	checkbox	Set to enables the FreeNAS® system to send email updates to the configured <i>Email</i> field.
Email	email address	Enter any email addresses to receive status updates. Separate multiple addresses with a semicolon (;).
Email Subject	string	Enter a subject line for email status updates.
Power Off UPS	checkbox	Set for the UPS to power off after shutting down the FreeNAS® system.

Note: For USB devices, the easiest way to determine the correct device name is to enable the *Show console messages* option in *System* → *Advanced*. Plug in the USB device and look for a `/dev/ugen` or `/dev/uhid` device name in the console messages.

Tip: Some UPS models might be unresponsive with the default polling frequency. This can show in FreeNAS® logs as a recurring error like: `libusb_get_interrupt: Unknown error`.

If this error occurs, increase the polling frequency by adding an entry to *Auxiliary Parameters (ups.conf)*: `pollinterval = 10`. The default polling frequency is two seconds.

`upsc(8)` (<https://www.freebsd.org/cgi/man.cgi?query=upsc>) 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* (page 325) using this syntax:

```
upsc ups@localhost
```

The `upsc(8)` (<https://www.freebsd.org/cgi/man.cgi?query=upsc>) man page gives some other usage examples.

`upscmd(8)` (<https://www.freebsd.org/cgi/man.cgi?query=upscmd>) can be used to send commands directly to the UPS, assuming 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.

12.17.1 Multiple Computers with One UPS

A UPS with adequate capacity can 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](https://networkupstools.org/docs/user-manual.chunked/index.html) (<https://networkupstools.org/docs/user-manual.chunked/index.html>) and [NUT User Manual Pages](https://networkupstools.org/docs/man/index.html#User_man) (https://networkupstools.org/docs/man/index.html#User_man).

12.18 WebDAV

The WebDAV service can be configured to provide a file browser over a web connection. Before starting this service, at least one WebDAV share must be created by navigating to *Sharing* → *WebDAV Shares*, and clicking *ADD*. Refer to [WebDAV Shares](#) (page 210) for instructions on how to create a share and connect to it after the service is configured and started.

The settings in the WebDAV service apply to all WebDAV shares. [Figure 12.19](#) shows the WebDAV configuration screen. [Table 12.16](#) summarizes the available options.

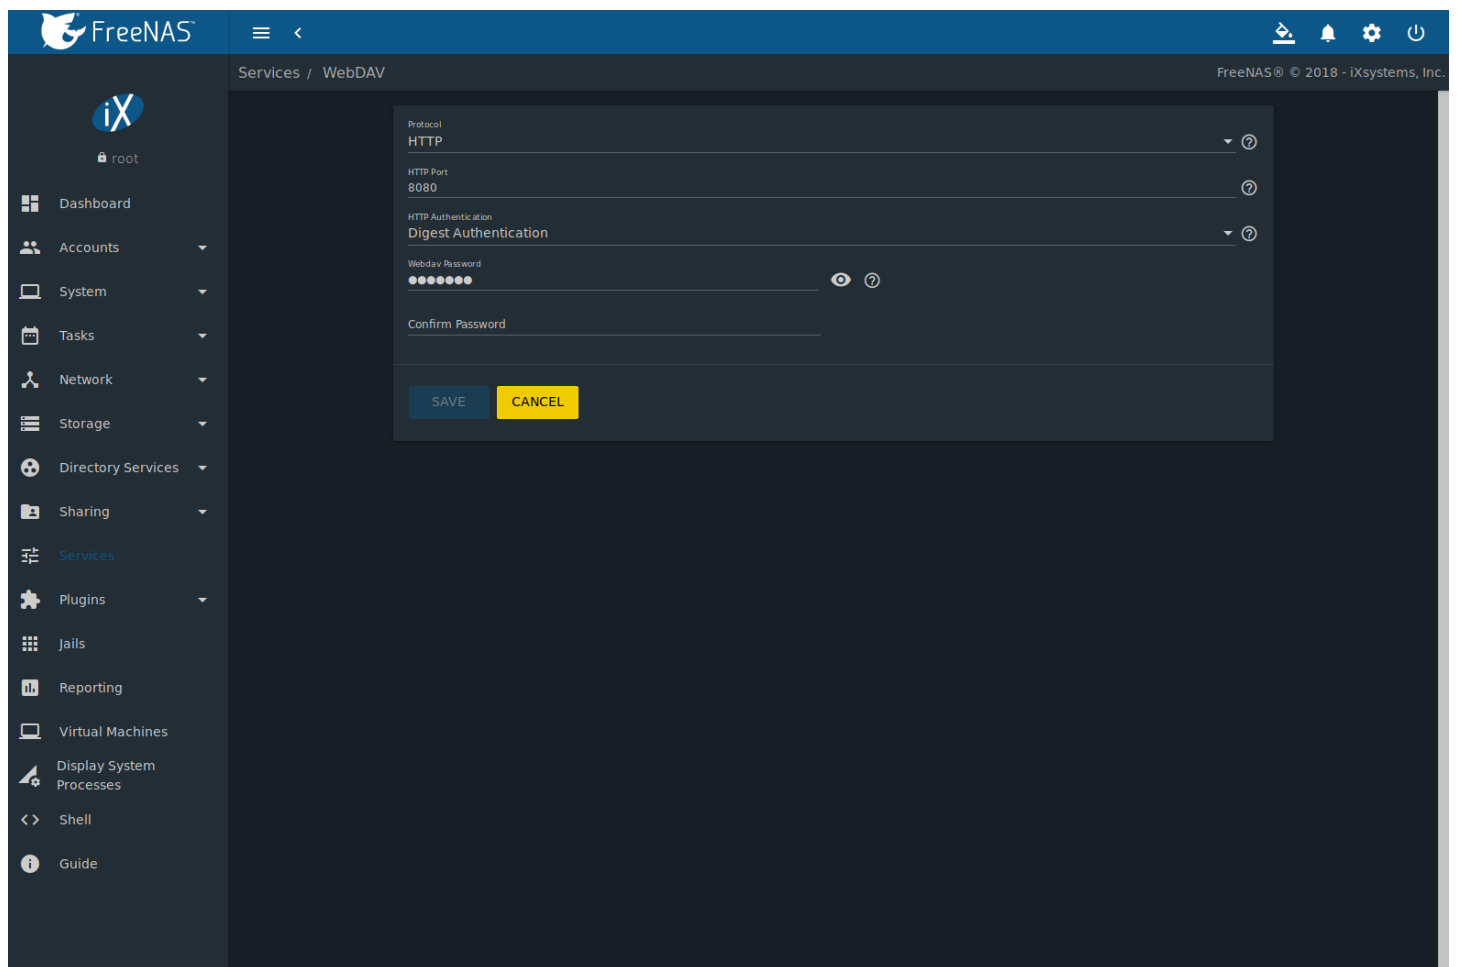


Fig. 12.19: WebDAV Configuration Screen

Table 12.16: WebDAV Configuration Options

Setting	Value	Description
Protocol	drop-down menu	<i>HTTP</i> keeps the connection unencrypted. <i>HTTPS</i> encrypts the connection. <i>HTTP+HTTPS</i> allows both types of connections.
HTTP Port	string	Specify a port for unencrypted connections. The default port <i>8080</i> is recommended. Do not use a port number already being used by another service.
HTTPS Port	string	Specify a port for encrypted connections. The default port <i>8081</i> is recommended. Do not use a port number already being used by another service.
Webdav SSL Certificate	drop-down menu	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 <i>No Authentication</i> , <i>Basic Authentication</i> (unencrypted) or <i>Digest Authentication</i> (encrypted).
Webdav Password	string	Default is <i>davtest</i> . Change this password as it is a known value.

PLUGINS

Warning: This section describes the plugin system implemented in the 11.2 release of FreeNAS®. Any plugins created or installed with a previous version of FreeNAS® must be managed with the *Legacy Web Interface* (page 63).

FreeNAS® provides the ability to extend the built-in NAS services by providing two methods for installing additional software. *Plugins* (page 271) allow the user to browse, install, and configure pre-packaged software from the web interface. This method is easy to use, but provides a limited amount of available software. Each plugin is automatically installed into its own limited *FreeBSD jail* (https://en.wikipedia.org/wiki/Freebsd_jail) that cannot install additional software.

Jails (page 281) provide more control over software installation, but requires working from the command line and a good understanding of networking basics and software installation on FreeBSD-based systems.

Look through the *Plugins* (page 271) and *Jails* (page 281) sections to become familiar with the features and limitations of each. Choose the method that best meets the needs of the application.

Note: *Jail Storage* (page 281) must be configured before plugins are available on FreeNAS®. This means having a suitable *pool* (page 154) created to store plugins.

13.1 Install

A plugin is a self-contained application installer designed to integrate into the FreeNAS® web interface. A plugin offers several advantages:

- the FreeNAS® web interface provides a browser for viewing the list of available plugins
- the FreeNAS® web interface provides buttons for installing, starting, managing, and deleting plugins
- if the plugin has configuration options, a screen will be added to the FreeNAS® web interface for these options to be configured

To install a plugin, click *Plugins* → *Available*. *Figure 13.1* shows some of the available plugins.

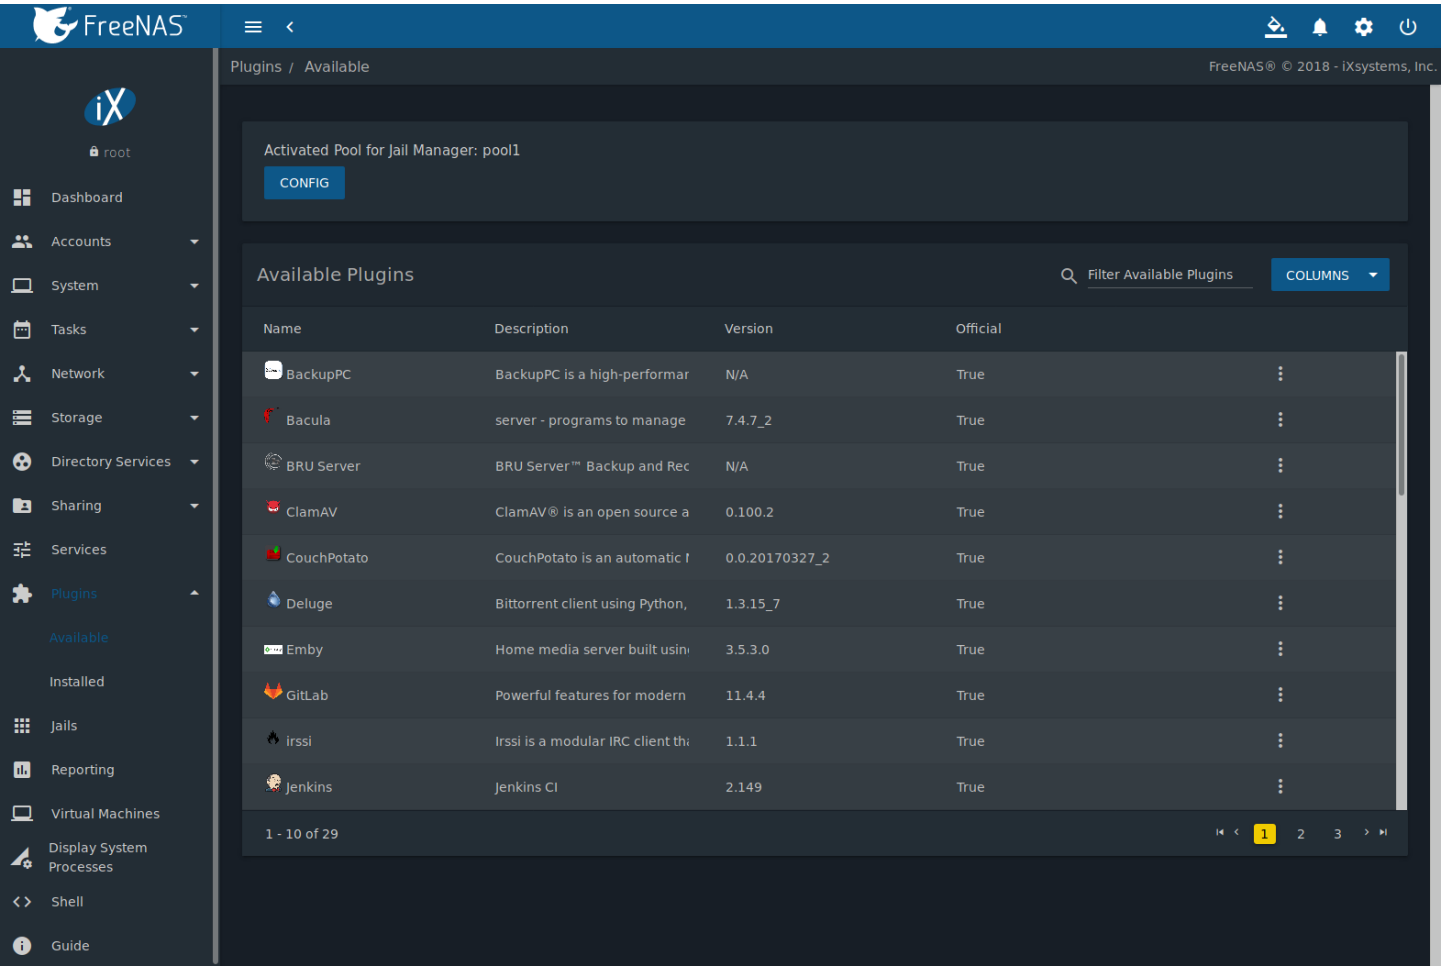



Fig. 13.1: Viewing the List of Available Plugins

The *Available Plugins* page lists the plugin name, description, current version, and whether the plugin is officially supported.

Note: If the list of available plugins is not displayed, open *Shell* (page 325) and verify that the FreeNAS® system can ping an address on the Internet. If it cannot, add a default gateway address and/or DNS server address in *Network* → *Global Configuration*.

Click  (Options) and *Install* for the desired plugin. Set *DHCP* to automatically configure IP settings, or manually enter an IPv4 or IPv6 address. Click *SAVE*. In the example shown in [Figure 13.2](#), Plex Media Server is selected for installation.

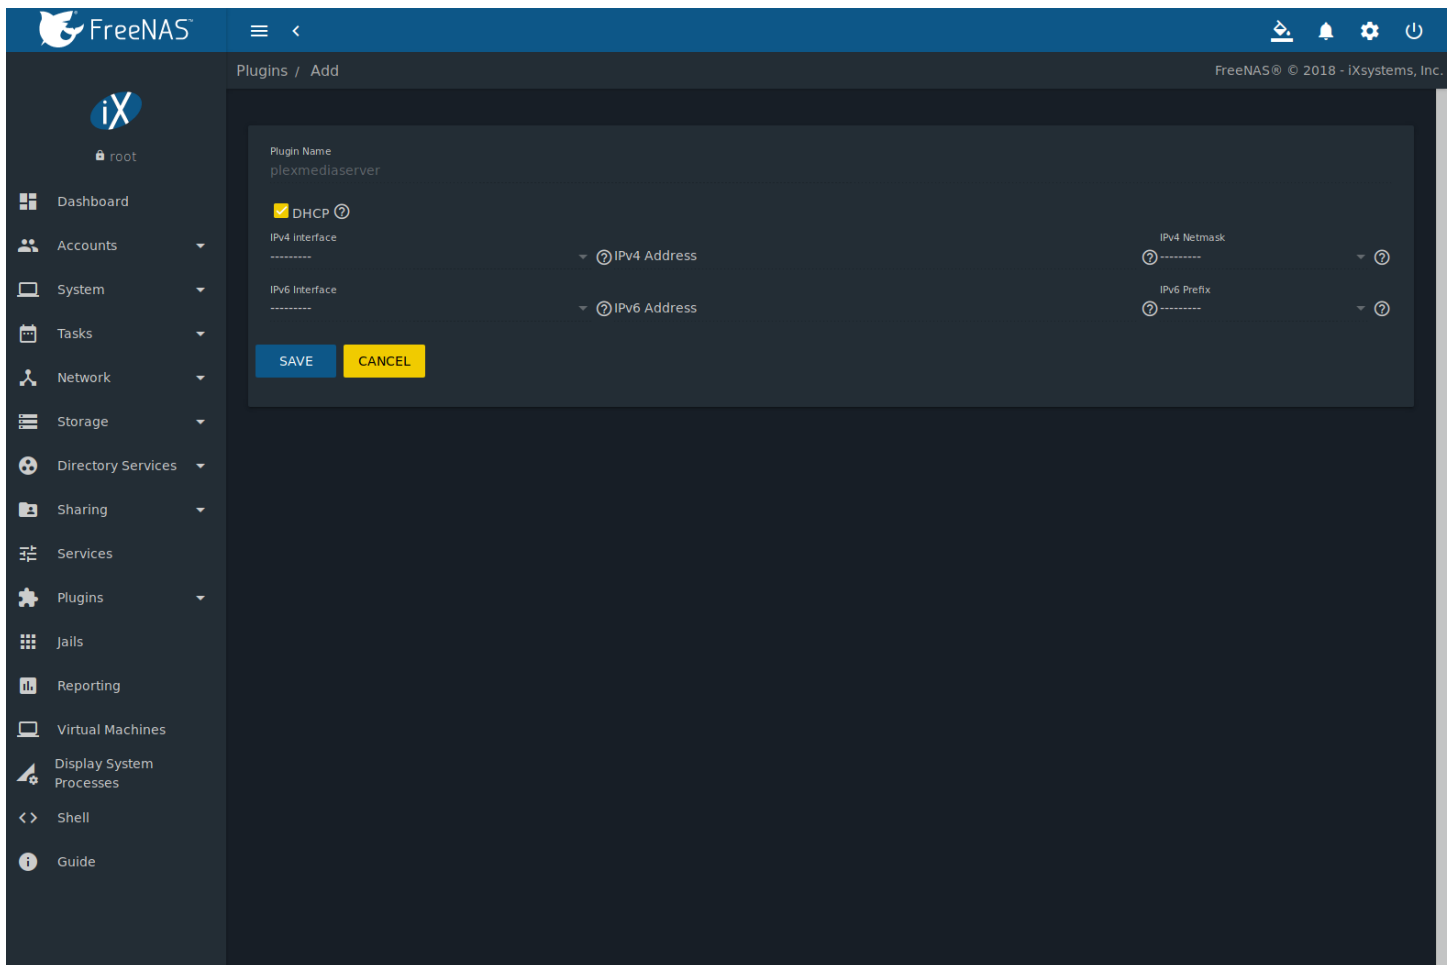


Fig. 13.2: Installing the Plex Plugin

The installation takes a few minutes because the system downloads and configures a jail to store the plugin application. A confirmation message displays at the bottom of the screen after successfully installing a plugin. When applicable, post-install notes are displayed after a successful install. Installed plugins appear in the *Plugins* → *Installed* page as shown in Figure 13.3.

Tip: Installed plugins are also added to the *Jails* page. This page is also used to manage installed software.

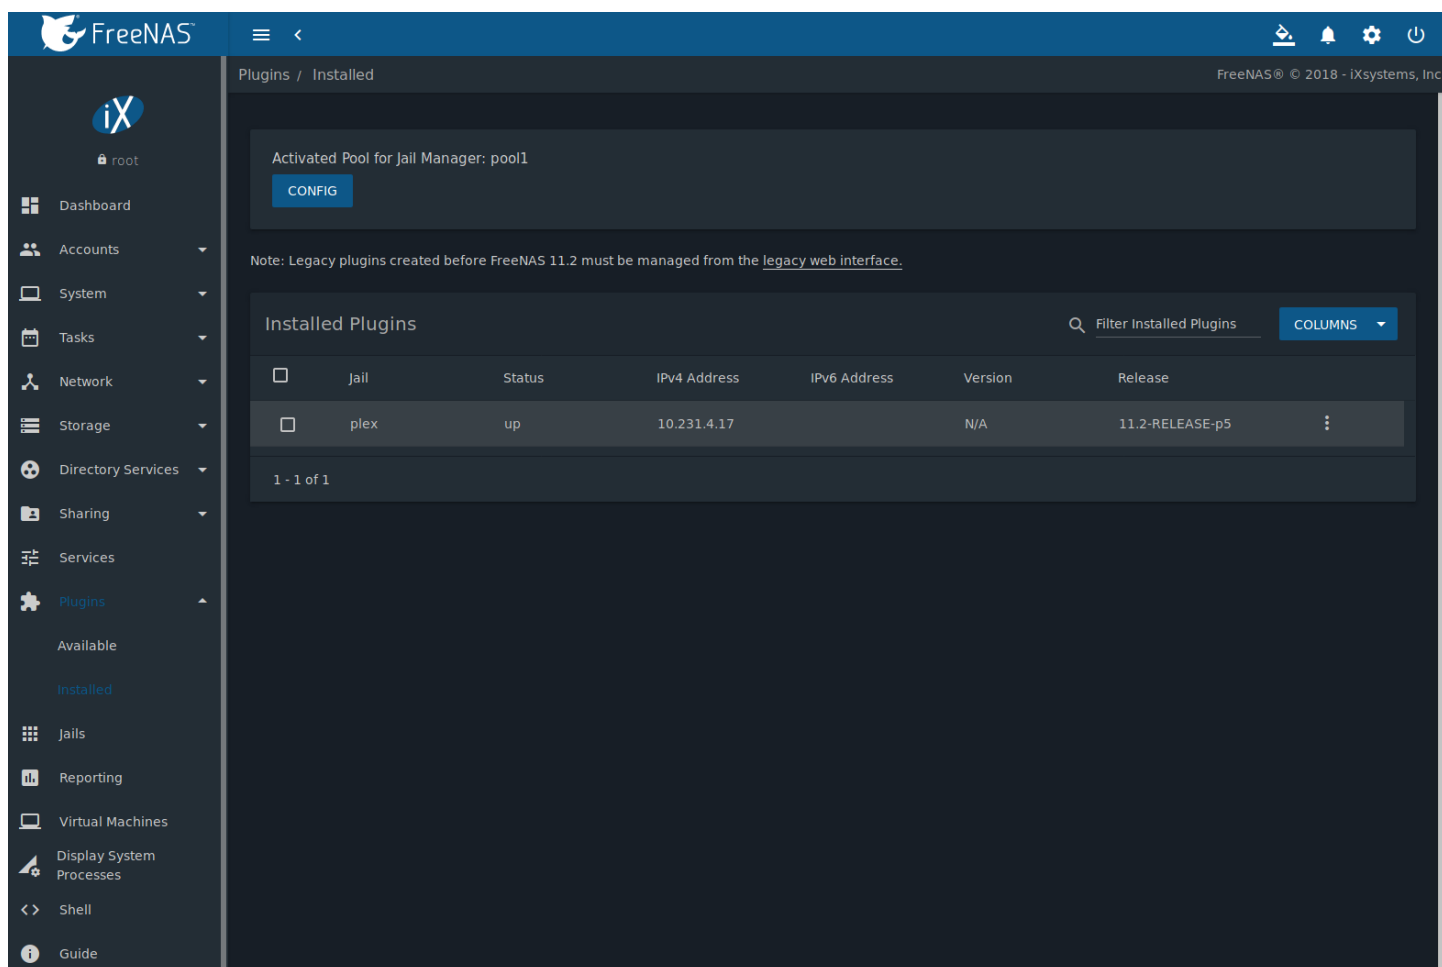




Fig. 13.3: Viewing Installed Plugins

The entry in the *Plugins* → *Installed* section displays the plugin jail name, status, IPv4 and IPv6 addresses, plugin application version, and FreeBSD release.

The plugin must be started before the installed application is available. Click  (Options) and *Start*. The plugin *Status* changes to *up* when it starts successfully.


Stop and immediately start an *up* plugin by clicking  (Options) and *Restart*.

Click  (Options) and *Management* to open a management or configuration screen for the application. For example, clicking *Management* for an installed Plex plugin opens the Plex web interface in a new browser tab.

Note: Not all plugins have a functional management option. See [Managing Jails](#) (page 289) for more instructions about interacting with a plugin jail with the shell.


Always review plugin configuration options before attempting to start it. Some plugins have options that need to be set before their service will successfully start. To help with installing a new application, check the website of the application to see what documentation is available.

If the application requires access to the data stored on the FreeNAS® system, click the entry for the associated jail in the *Jails* page and add a storage as described in [Additional Storage](#) (page 293).

Click  (Options) and *Shell* for the plugin jail in the *Jails* page. This will give access to the shell of the jail containing the application to complete or test the configuration.

If a plugin jail fails to start, open the plugin jail shell from the *Jail* page and type `tail /var/log/messages` to see if any errors were logged.

13.2 Update

When a newer version of a plugin becomes available in the official repository, update the plugin jail by navigating to the *Jails* page and clicking  (Options) and *Update*.

13.3 Delete

Installing a plugin creates an associated jail. Deleting a plugin deletes the associated jail because it is no longer required. **Before** deleting a plugin, make sure that there is no data or configuration in the jail that needs to be saved. Back up that data **first** if needed.

In the example shown in [Figure 13.4](#), Quasselcore has been installed and the *Delete* button has been clicked. A pop-up message asks for verification that the plugin is to be deleted. **This is the only warning.** The plugin and the associated jail are permanently deleted when *Confirm* is set and *DELETE* is clicked.

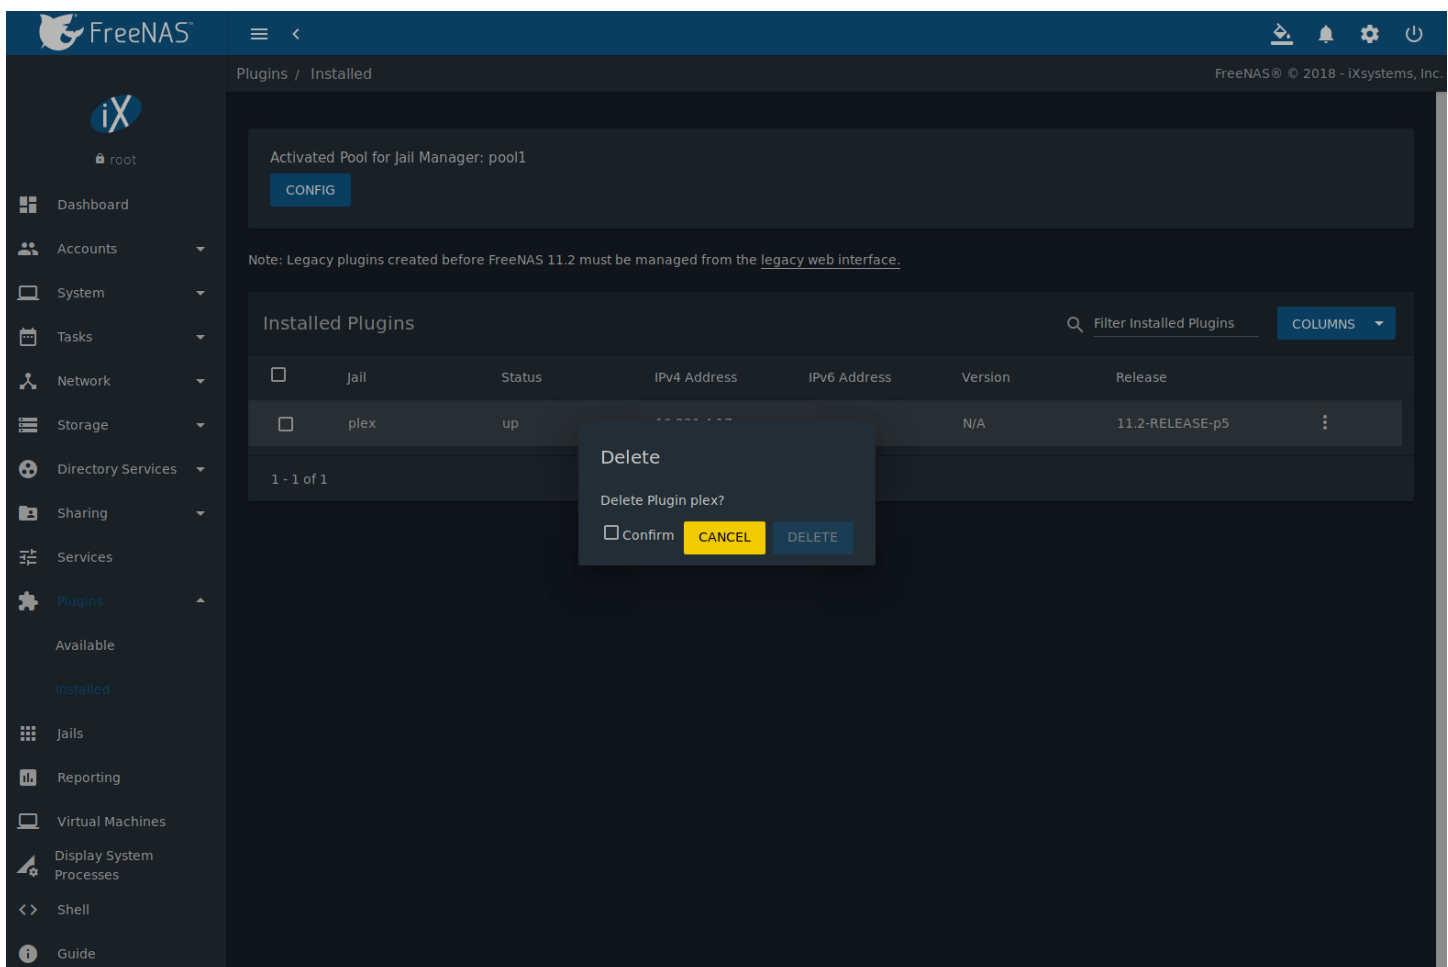


Fig. 13.4: Deleting an Installed Plugin

13.4 Create a Plugin

If an application is not available as a plugin, it is possible to create a new plugin for FreeNAS® in a few steps. This requires an existing [GitHub](https://github.com) (<https://github.com>) account.

Create a new artifact repository on [GitHub](https://github.com) (<https://github.com>).

Refer to [table 13.1](#) for the files to add to the artifact repository.

Table 13.1: FreeNAS® Plugin Artifact Files

Directory/File	Description
post_install.sh	This script is run <i>inside</i> the jail after it is created and any packages installed. Enable services in <code>/etc/rc.conf</code> that need to start with the jail and apply any configuration customizations with this script.
ui.json	JSON file that accepts the key or value options. For example: <code>adminportal: "http://%%IP%%/"</code> designates the web-interface of the plugin.
overlay/	Directory of files overlaid on the jail after install. For example, <code>usr/local/bin/myfile</code> is placed in the <code>/usr/local/bin/myfile</code> location of the jail. Can be used to supply custom files and configuration data, scripts, and any other type of customized files to the plugin jail.
settings.json	JSON file that manages the settings interface of the plugin. Required fields include: <ul style="list-style-type: none"> • <code>"servicerestart" : "service foo restart"</code> Command to run when restarting the plugin service after changing settings. • <code>"serviceget" : "/usr/local/bin/myget"</code> Command used to get values for plugin configuration. Provided by the plugin creator. The command accepts two arguments for key or value pair. • <code>"options" :</code> This subsection contains arrays of elements, starting with the “key” name and required arguments for that particular type of setting. See options subsection example (page 276) below.

This example `settings.json` file is used for the *Quasselcore* plugin. It is also available online in the [iocage-plugin-quassel artifact repository](https://github.com/freenas/iocage-plugin-quassel/blob/master/settings.json) (<https://github.com/freenas/iocage-plugin-quassel/blob/master/settings.json>).

```
{
  "servicerestart": "service quasselcore restart",
  "serviceget": "/usr/local/bin/quasselget",
  "serviceset": "/usr/local/bin/quasselset",
  "options": {
    "adduser": {
      "type": "add",
      "name": "Add User",
      "description": "Add new quasselcore user",
      "requiredargs": {
        "username": {
          "type": "string",
          "description": "Quassel Client Username"
        },
        "password": {
          "type": "password",
          "description": "Quassel Client Password"
        },
        "fullname": {
          "type": "string",
          "description": "Quassel Client Full Name"
        }
      },
      "optionalargs": {
        "adminuser": {
          "type": "bool",
          "description": "Can this user administrate quasselcore?"
        }
      }
    }
  },
}
```



```

    "port": {
      "type": "int",
      "name": "Quassel Core Port",
      "description": "Port for incoming quassel connections",
      "range": "1024-32000",
      "default": "4242",
      "requirerestart": true
    },
    "sslmode": {
      "type": "bool",
      "name": "SSL Only",
      "description": "Only accept SSL connections",
      "default": true,
      "requirerestart": true
    },
    "ssloption": {
      "type": "combo",
      "name": "SSL Options",
      "description": "SSL Connection Options",
      "requirerestart": true,
      "default": "tlsallow",
      "options": {
        "tlsrequire": "Require TLS",
        "tlsallow": "Allow TLS",
        "tlsdisable": "Disable TLS"
      }
    },
    "deluser": {
      "type": "delete",
      "name": "Delete User",
      "description": "Remove a quasselcore user"
    }
  }
}

```

Create and submit a new JSON file for the plugin:

Clone the `iocage-ix-plugins` (<https://github.com/freenas/iocage-ix-plugins>) GitHub repository.

Tip: Full tutorials and documentation for GitHub and `git` commands are available on [GitHub Guides](https://guides.github.com/) (<https://guides.github.com/>).

On the local copy of `iocage-ix-plugins`, create a new file for the plugin to be added to FreeNAS®. The naming convention is `pluginname.json`. For example, the *Transmission* plugin has a `.json` file named `transmission.json`.

Add fields to this `.json` file. [table 13.2](#) lists and describes each required entry.

Table 13.2: Plugin `.json` File Contents

Data Field	Description
"name":	Name of the plugin.
"release":	FreeBSD release to use for the plugin jail.
"artifact":	URL of the plugin artifact repo.
"pkgs":	Port of the plugin.
"packagesite":	CDN the plugin jail uses. Default for the TrueOS CDN is http://pkg.cdn.trueos.org/iocage .

Continued on next page

Table 13.2 – continued from previous page

Data Field	Description
"fingerprints":	"function": Default is <i>sha256</i> . "fingerprint": The pkg fingerprint for the artifact repo. Default is <i>226efd3a126fb86e71d60a37353d17f57af816d1c7ecad0623c21f0bf73eb0c7</i>
"official":	Defines if this an official iXsystems supported plugin. Enter <i>true</i> or <i>false</i> .

Here is `quasselcore.json` reproduced as an example:

```
{
  "name": "Quasselcore",
  "release": "11.1-RELEASE",
  "artifact": "https://github.com/freenas/iocage-plugin-quassel.git",
  "pkgs": [
    "irc/quassel-core"
  ],
  "packagesite": "http://pkg.cdn.trueos.org/iocage",
  "fingerprints": {
    "iocage-plugins": [
      {
        "function": "sha256",
        "fingerprint": "226efd3a126fb86e71d60a37353d17f57af816d1c7ecad0623c21f0bf73eb0c7"
      }
    ]
  },
  "official": true
}
```

The correct directory and package name of the plugin application must be used for the "pkgs": value. Find the package name and directory by searching [FreshPorts](https://www.freshports.org/) (<https://www.freshports.org/>) and checking the "To install the port:" line. For example, the *Quasselcore* plugin uses the directory and package name `/irc/quassel-core`.

Now edit `iocage-ix-plugins/INDEX`. Add an entry for the new plugin that includes these fields:

- "MANIFEST": Add the name of the newly created `plugin.json` file here.
- "name": Use the same name used within the `.json` file.
- "icon": Most plugins will have a specific icon. Search the web and save the icon to the `icons/` directory as a `.png`. The naming convention is `pluginname.png`. For example, the *Transmission* plugin has the icon file `transmission.png`.
- "description": Describe the plugin in a single sentence.
- "official": Specify if the plugin is supported by iXsystems. Enter *false*.

See the [INDEX](https://github.com/freenas/iocage-ix-plugins/blob/master/INDEX) (<https://github.com/freenas/iocage-ix-plugins/blob/master/INDEX>) for examples of `INDEX` entries.

Submit the plugin

Open a pull request for the [iocage-ix-plugins](https://github.com/freenas/iocage-ix-plugins) repo (<https://github.com/freenas/iocage-ix-plugins>). Make sure the pull request contains:

- the new `plugin.json` file.
- the plugin icon `.png` added to the `icons/` directory.
- an update to the `INDEX` file with an entry for the new plugin.
- a link to the artifact repository populated with all required plugin files.

13.4.1 Test a Plugin

Warning: Installing experimental plugins is not recommended for general use of FreeNAS®. This feature is meant to help plugin creators test their work before it becomes generally available on FreeNAS®.

Plugin pull requests are merged into the `master` branch of the [iocage-ix-plugins](https://github.com/freenas/iocage-ix-plugins) (<https://github.com/freenas/iocage-ix-plugins>) repository. These plugins are not available in the web interface until they are tested and added to a *RELEASE* branch of the repository. It is possible to test an in-development plugin by using this `iocage` command: `iocage fetch -P --name PLUGIN IPADDRESS_PROPS --branch 'master'`

This will install the plugin, configure it with any chosen properties, and specifically use the `master` branch of the repository to download the plugin.

Here is an example of downloading and configuring an experimental plugin with the FreeNAS® *Shell*:

```
[root@freenas ~]# iocage fetch -P --name mineos ip4_addr="em0|10.231.1.37/24" --branch 'master'
Plugin: mineos
  Official Plugin: False
  Using RELEASE: 11.2-RELEASE
  Using Branch: master
  Post-install Artifact: https://github.com/jsegaert/iocage-plugin-mineos.git
  These pkgs will be installed:
...

...
Running post_install.sh
Command output:
...

...
Admin Portal:
http://10.231.1.37:8443
[root@freenas ~]#
```

This plugin appears in the *Jails* and *Plugins* → *Installed* screens as `mineos` and can be tested with the FreeNAS® system.

13.5 Official Plugins

[table 13.3](#) lists and describes all plugins supported by iXsystems. Adding “unofficial” plugins to FreeNAS® is supported by following the process outlined in [Create a Plugin](#) (page 275).

Table 13.3: Official FreeNAS® plugins

Name	Description
BackupPC (http://backuppc.sourceforge.net/)	BackupPC is a high-performance, enterprise-grade system for backing up Linux, WinXX and MacOSX PCs and laptops to a server disk.
Bacula (https://www.baculasystems.com/)	Bacula is an open-source, enterprise-level computer backup system for heterogeneous networks.
BRU Server (http://www.tolisgroup.com/client-server-cross-platform-backup.html)	BRU Server™ Backup and Recovery Software by TOLIS Group, Inc.
ClamAV (https://www.clamav.net/)	ClamAV is an open source antivirus engine for detecting trojans, viruses, malware & other malicious threats.
CouchPotato (https://couchpota.to/)	CouchPotato is an automatic NZB and torrent downloader.
Deluge (https://deluge-torrent.org/)	Bittorrent client using Python, and libtorrent-rasterbar.
Emby (https://emby.media/)	Home media server built using mono and other open source technologies.

Continued on next page

Table 13.3 – continued from previous page

Name	Description
GitLab (https://about.gitlab.com/)	GitLab is a fully integrated software development platform.
irssi (https://irssi.org/)	Irssi is an IRC client.
Jenkins (https://jenkins.io/)	Jenkins is a self-contained, open source automation server which can be used to automate all sorts of tasks related to building, testing, and delivering or deploying software.
Jenkins (LTS) (https://jenkins.io/download/lts/)	Jenkins Long-Term Support releases.
Madsonic (http://beta.madsonic.org/pages/index.jsp)	Open-source web-based media streamer and jukebox.
MineOS (https://minecraft.codeemo.com/)	Self-contained Minecraft server.
Nextcloud (https://nextcloud.com/)	Access, share and protect files, calendars, contacts, communication and more at home and in the enterprise environment.
PlexMediaServer (https://www.plex.tv/)	The Plex media server system.
Plex Media Server (PlexPass) (https://www.plex.tv/plex-pass/)	Premium service for Plex media server system.
qBittorrent (http://qbittorrent.org/)	qBittorrent is a cross-platform client for the BitTorrent protocol that is released under the GNU GPL, version 2.
Quasselcore (https://quassel-irc.org/)	Quassel Core is a daemon/headless IRC client, part of Quassel, that supports 24/7 connectivity. Quassel Client can also be attached to it.
radarr (https://radarr.video/)	A fork of Sonarr to work with movies in the style of Couchpotato.
Redmine (http://www.redmine.org/)	Flexible project management web application.
Resilio Sync (https://www.resilio.com/)	Formerly known as BitTorrent Sync. Resilient, fast and scalable file sync software for enterprises and individuals.
Sonarr (https://sonarr.tv/)	PVR for Usenet and BitTorrent users.
Subsonic (http://www.subsonic.org/pages/index.jsp)	Open-source web-based media streamer and jukebox.
Syncthing (https://syncthing.net/)	Personal cloud sync.
Tarsnap (https://www.tarsnap.com/)	Online encrypted backup service (client).
Transmission (https://transmissionbt.com/)	Fast and lightweight daemon BitTorrent client.
WeeChat (https://weechat.org/)	WeeChat is a free and open-source Internet Relay Chat client, which is designed to be light and fast.
XMRig (https://github.com/xmrig/xmrig)	XMRig is a high performance Monero (XMR) CPU miner
ZoneMinder (https://zoneminder.com/)	A full-featured, open source, state-of-the-art video surveillance software system.

If there are any difficulties using a plugin, refer to the official documentation for that application.

JAILS

Warning: This section describes installing and using jails on FreeNAS® version 11.2 or later. Any jails created with a previous version of FreeNAS® must be managed with the [Legacy Web Interface](#) (page 63).

Jails are a lightweight, operating-system-level virtualization. One or multiple services can run in a jail, isolating those services from the host FreeNAS® system. FreeNAS® uses the [iocage](https://github.com/iocage/iocage) (<https://github.com/iocage/iocage>) utility for jail management. Jails are also used as the basis for FreeNAS® [Plugins](#) (page 271). The main differences between a user-created jail and a plugin are that plugins are preconfigured and usually provide only a single service.

By default, jails run the [FreeBSD](https://www.freebsd.org/) (<https://www.freebsd.org/>) operating system. These jails are independent instances of FreeBSD. The jail uses the host hardware and runs on the host kernel, avoiding most of the overhead usually associated with virtualization. The jail installs FreeBSD software management utilities so FreeBSD packages or ports can be installed from the jail command line. This allows for FreeBSD ports to be compiled and FreeBSD packages to be installed from the command line of the jail.

It is important to understand that users, groups, installed software, and configurations within a jail are isolated from both the FreeNAS® host operating system and any other jails running on that system.

During creation, set the *VNET* option to provide the jail with an independent networking stack. The jail is then able to broadcast an IP address, which is required by some applications.

The ability to create multiple jails offers flexibility regarding software management. For example, an administrator can choose to provide application separation by installing different applications in each jail, to create one jail for all installed applications, or to mix and match how software is installed into each jail.

14.1 Jail Storage

A *pool* (page 154) must be created before using jails or [Plugins](#) (page 271). Make sure the pool has enough storage for all the intended jails and plugins. The *Jails* screen displays a message and button to *CREATE POOL* if no pools exist on the FreeNAS® system.

Multiple pools can be activated to store iocage jails and plugins. After a pool is created, the *Jails* page displays an *Activated Pool* section. This shows which pool and iocage dataset is active with FreeNAS®. Click *CONFIG* to view the option to choose another pool or dataset to activate with iocage. *ACTIVATE* another pool to refresh the *Jails* list with any jails that exist on the chosen pool or dataset.

Jails and downloaded FreeBSD release files are stored in a dataset named `iocage/`.

Notes about the `iocage/` dataset:

- At least 10 GiB of free space is recommended.
- Cannot be located on a [Share](#) (page 195).
- [iocage](http://iocage.readthedocs.io/en/latest/index.html) (<http://iocage.readthedocs.io/en/latest/index.html>) automatically uses the first pool that is not a root pool for the FreeNAS® system.

- A `defaults.json` file contains default settings used when a new jail is created. The file is created automatically if not already present. If the file is present but corrupted, `iocage` shows a warning and uses default settings from memory.
- Each new jail installs into a new child dataset of `iocage/`. For example, with the `iocage/jails` dataset in `pool1`, a new jail called `jail1` installs into a new dataset named `pool1/iocage/jails/jail1`.
- FreeBSD releases are fetched as a child dataset into the `/iocage/download` dataset. This dataset is then extracted into the `/iocage/releases` dataset to be used in jail creation. The dataset in `/iocage/download` can then be removed without affecting the availability of fetched releases or an existing jail.
- `iocage/` datasets on activated pools are independent of each other and do **not** share any data.

14.2 Creating Jails

FreeNAS® has two options to create a jail. The *Jail Wizard* makes it easy to quickly create a jail. *ADVANCED JAIL CREATION* is an alternate method, where every possible jail option is configurable. There are numerous options spread across four different primary sections. This form is recommended for advanced users with very specific requirements for a jail.

14.2.1 Jail Wizard

New jails can be created quickly by going to *Jails* → *ADD*. This opens the wizard screen shown in [Figure 14.1](#).

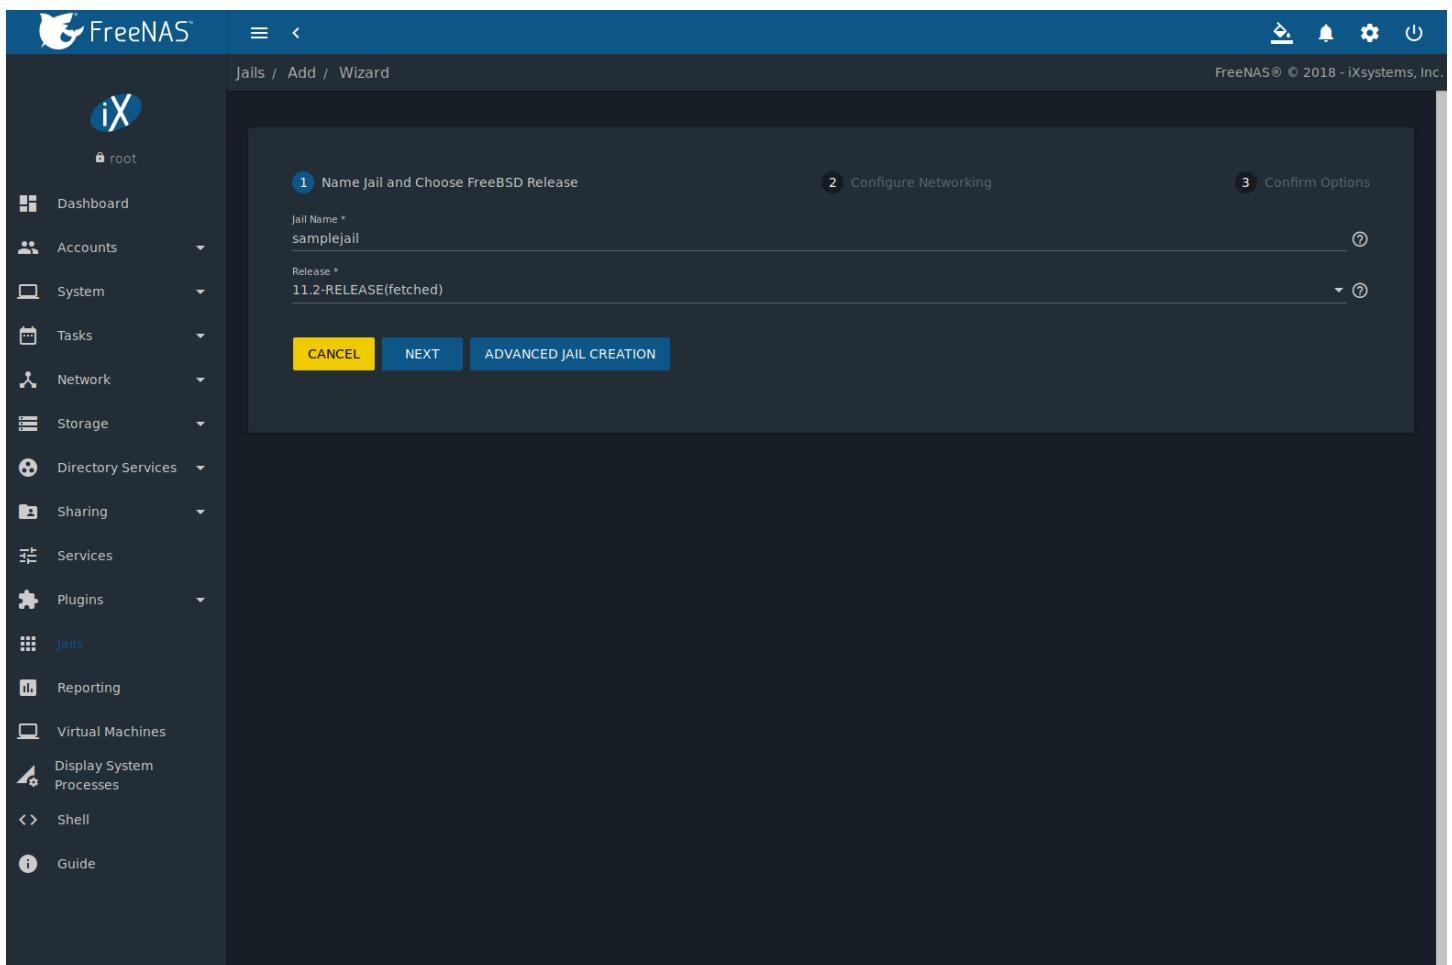


Fig. 14.1: Jail Creation Wizard

The wizard provides the simplest process to create and configure a new jail. Enter a *Jail Name*. Jail names can only contain alphanumeric characters (Aa-Zz, 123), dashes (-), and underscores (_). Choose the version of FreeBSD to install for this jail. Previously downloaded versions display (fetched) next to their entry in the list.

Click *NEXT* to see a simplified list of networking options. The jail can be set to automatically configure IPv4 with *DHCP* and *VNET* or IPv4 and IPv6 can be configured manually. Multiple interfaces are supported in the *IPv4 Address* and *IPv6 Address* fields by entering a comma delimited list of interfaces, addresses, and netmask in the format `interface|ipaddress/netmask`.

Click *NEXT* to view a summary screen of the chosen jail options. Click *SUBMIT* to create the new jail. After a few moments, the new jail is added to the primary jails list.

Tip: Versions of FreeBSD are downloaded the first time they are used in a jail. Additional jails created with the same version of FreeBSD are created faster because the download has already been completed.

14.2.2 Advanced Jail Creation

The advanced jail creation form is opened by clicking *Jails* → *ADD* then *Advanced Jail Creation*. The screen in Figure 14.2 is shown.

The screenshot shows the 'Advanced Jail Creation' form in the FreeNAS web interface. The form is titled 'Basic Properties' and contains the following fields and options:

- Name ***: A text input field with a help icon.
- Release ***: A dropdown menu with a help icon.
- DHCP Autoconfigure IPv4**: A checkbox with a help icon.
- VNET**: A checkbox with a help icon.
- Berkeley Packet Filter**: A checkbox with a help icon.
- IPv4 Interface**: A dropdown menu with a help icon.
- IPv4 Address**: A text input field with a help icon.
- IPv4 Netmask**: A dropdown menu with a help icon.
- IPv4 Default Router**: A text input field with a help icon.
- Autoconfigure IPv6**: A checkbox with a help icon.
- IPv6 Interface**: A dropdown menu with a help icon.
- IPv6 Address**: A text input field with a help icon.
- IPv6 Prefix**: A dropdown menu with a help icon.
- IPv6 Default Router**: A text input field with a help icon.
- Notes**: A text input field with a help icon.
- Auto-start**: A checkbox with a help icon.

A blue **NEXT** button is located at the bottom right of the form. Below the form, there are three expandable sections: **Jail Properties**, **Network Properties**, and **Custom Properties**.

Fig. 14.2: Creating a Jail

A usable jail can be quickly created by setting only the required values, the *Jail Name* and *Release*. Additional settings are in the *Jail Properties*, *Network Properties*, and *Custom Properties* sections. Table 14.1 shows the available options of the *Basic Properties* of a new jail.

Table 14.1: Basic Properties

Setting	Value	Description
Name	string	Required. Jail names can only contain alphanumeric characters (Aa-zz, 123), dashes (-), and underscores (_).
Release	drop-down menu	Required. Choose the version of FreeBSD to download and install for the jail. Previously downloaded versions of FreeBSD display (fetched) next to the entry in the list and do not need to be fetched again.
DHCP Autoconfigure IPv4	checkbox	Automatically configure IPv4 networking with an independent VNET stack. <i>VNET</i> and <i>Berkeley Packet Filter</i> must also be checked. If not set, ensure the defined address in <i>IPv4 Address</i> does not conflict with an existing address.
VNET	checkbox	Use VNET to emulate network devices for this jail and a create a fully virtualized per-jail network stack. See VNET(9) (https://www.freebsd.org/cgi/man.cgi?query=vnet) for more details.
Berkeley Packet Filter	checkbox	Use the Berkeley Packet Filter to data link layers in a protocol independent fashion. Unset by default to avoid security vulnerabilities. See BPF(4) (https://www.freebsd.org/cgi/man.cgi?query=bpf) for more details.
IPv4 Interface	drop-down menu	Choose a network interface to use for this IPv4 connection.
IPv4 Address	string	This and the other IPv4 settings are grayed out if <i>DHCP autoconfigure IPv4</i> is set. Configures the interface to use for network or internet access for the jail. Enter an IPv4 address for this IP jail. Example: <i>192.168.0.10</i> .
IPv4 Netmask	drop-down menu	Choose a subnet mask for this IPv4 Address.
IPv4 Default Router	string	Type <i>none</i> or a valid IP address. Setting this property to anything other than <i>none</i> configures a default route inside a VNET jail.
Auto Configure IPv6	checkbox	Set to use SLAAC (Stateless Address Auto Configuration) to autoconfigure IPv6 in the jail.
IPv6 Interface	drop-down menu	Choose a network interface to use for this IPv6 connection.
IPv6 Address	string	Configures network or internet access for the jail. Type the IPv6 address for VNET and shared IP jails. Example: <i>2001:0db8:85a3:0000:0000:8a2e:0370:7334</i> .
IPv6 Prefix	drop-down menu	Choose a prefix for this IPv6 Address.
IPv6 Default Router	string	Type <i>none</i> or a valid IP address. Setting this property to anything other than <i>none</i> configures a default route inside a VNET jail.
Notes	string	Enter any notes or comments about the jail.
Auto-start	checkbox	Start the jail at system startup.

Similar to the [Jail Wizard](#) (page 282), configuring the basic properties, then clicking *SAVE* is often all that is needed to quickly create a new jail. To continue configuring more settings, click *NEXT* to proceed to the *Jail Properties* section of the form. [Table 14.2](#) describes each of these options.

Table 14.2: Jail Properties

Setting	Value	Description
devfs_ruleset	integer	Number of the devfs(8) (https://www.freebsd.org/cgi/man.cgi?query=devfs) ruleset to enforce when mounting <i>devfs</i> in the jail. The default value of 0 means no ruleset is enforced. Mounting <i>devfs</i> inside a jail is only possible when the <i>allow_mount</i> and <i>allow_mount_devfs</i> permissions are enabled and <i>enforce_stafs</i> is set to a value lower than 2.
exec.start	string	Commands to run in the jail environment when a jail is created. Example: <code>sh /etc/rc</code> . See jail(8) (https://www.freebsd.org/cgi/man.cgi?query=jail) for more details.
exec.stop	string	Commands to run in the jail environment before a jail is removed and after any <i>exec_prestop</i> commands are complete. Example: <code>sh /etc/rc.shutdown</code> .
exec_prestart	string	Commands to run in the system environment before a jail is started.
exec_poststart	string	Commands to run in the system environment after a jail is started and after any <i>exec_start</i> commands are finished.
exec_prestop	string	Commands to run in the system environment before a jail is stopped.
exec_poststop	string	Commands to run in the system environment after a jail is started and after any <i>exec_start</i> commands are finished.
exec.clean	checkbox	Run commands in a clean environment. The current environment is discarded except for <code>\$HOME</code> , <code>\$SHELL</code> , <code>\$TERM</code> and <code>\$USER</code> . <code>\$HOME</code> and <code>\$SHELL</code> are set to the target login. <code>\$USER</code> is set to the target login. <code>\$TERM</code> is imported from the current environment. The environment variables from the login class capability database for the target login are also set.
exec_timeout	integer	The maximum amount of time in seconds to wait for a command to complete. If a command is still running after the allotted time, the jail is terminated.
stop_timeout	integer	The maximum amount of time in seconds to wait for the jail processes to exit after sending a SIGTERM signal. This happens after any <i>exec_stop</i> commands are complete. After the specified time, the jail is removed, killing any remaining processes. If set to 0, no SIGTERM is sent and the jail is immediately removed.
exec_jail_user	string	Enter either <code>root</code> or a valid <i>username</i> . Inside the jail, commands run as this user.
exec.system_jail_user	string	Set this boolean option to <i>True</i> to look for the <i>exec.jail_user</i> in the system passwd(5) (https://www.freebsd.org/cgi/man.cgi?query=passwd) file <i>instead</i> of the jail <i>passwd</i> .
exec.system_user	string	Run commands in the jail as this user. By default, commands are run as the current user.
mount.devfs	checkbox	Mount a devfs(5) (https://www.freebsd.org/cgi/man.cgi?query=devfs) filesystem on the chrooted <code>/dev</code> directory and apply the ruleset in the <i>devfs_ruleset</i> parameter to restrict the devices visible inside the jail.
mount.fdescfs	checkbox	Mount an fdescfs(5) (https://www.freebsd.org/cgi/man.cgi?query=fdescfs) filesystem in the jail <code>/dev/fd</code> directory.

Continued on next page

Table 14.2 – continued from previous page

Setting	Value	Description
enforce_statfs	drop-down	Determine which information processes in a jail are able to obtain about mount points. The behavior of multiple syscalls is affected: statfs(2) (https://www.freebsd.org/cgi/man.cgi?query=statfs), fstatfs(2) (https://www.freebsd.org/cgi/man.cgi?query=fstatfs), getfsstat(2) (https://www.freebsd.org/cgi/man.cgi?query=getfsstat), fhstatfs(2) (https://www.freebsd.org/cgi/man.cgi?query=fhstatfs), and other similar compatibility syscalls. All mount points are available without any restrictions if this is set to 0. Only mount points below the jail chroot directory are available if this is set to 1. Set to 2, the default option only mount points where the jail chroot directory is located are available.
children.max	integer	Number of child jails allowed to be created by the jail or other jails under this jail. A limit of 0 restricts the jail from creating child jails. <i>Hierarchical jails</i> in the jail(8) (https://www.freebsd.org/cgi/man.cgi?query=jail) man page explains the finer details.
login_flags	string	Flags to pass to login(1) (https://www.freebsd.org/cgi/man.cgi?query=login) when logging in to the jail using the console function.
securelevel	integer	Value of the jail securelevel (https://www.freebsd.org/doc/faq/security.html) sysctl. A jail never has a lower securelevel than the host system. Setting this parameter allows a higher securelevel. If the host system securelevel is changed, jail securelevel will be at least as secure. Securelevel options are: 3, 2, 1, 0, and -1.
sysvmsg	drop-down	Allow or deny access to SYSV IPC message primitives. Set to <i>Inherit</i> : All IPC objects on the system are visible to the jail. Set to <i>New</i> : Only objects the jail created using the private key namespace are visible. The system and parent jails have access to the jail objects but not private keys. Set to <i>Disable</i> : The jail cannot perform any sysvmsg related system calls.
sysvsem	drop-down	Allow or deny access to SYSV IPC semaphore primitives. Set to <i>Inherit</i> : All IPC objects on the system are visible to the jail. Set to <i>New</i> : Only objects the jail creates using the private key namespace are visible. The system and parent jails have access to the jail objects but not private keys. Set to <i>Disable</i> : The jail cannot perform any sysvsem related system calls.
sysvshm	drop-down	Allow or deny access to SYSV IPC shared memory primitives. Set to <i>Inherit</i> : All IPC objects on the system are visible to the jail. Set to <i>New</i> : Only objects the jail creates using the private key namespace are visible. The system and parent jails have access to the jail objects but not private keys. Set to <i>Disable</i> : The jail cannot perform any sysvshm related system calls.
allow.set_hostname	checkbox	Allow the jail hostname to be changed with hostname(1) (https://www.freebsd.org/cgi/man.cgi?query=hostname) or sethostname(3) (https://www.freebsd.org/cgi/man.cgi?query=sethostname).
allow.sysvipc	checkbox	Choose whether a process in the jail has access to System V IPC primitives. Equivalent to setting <i>sysvmsg</i> , <i>sysvsem</i> , and <i>sysvshm</i> to <i>Inherit</i> . <i>Deprecated in FreeBSD 11.0 and later!</i> Use <i>sysvmsg</i> , <i>sysvsem</i> , and <i>sysvshm</i> instead.

Continued on next page

Table 14.2 – continued from previous page

Setting	Value	Description
allow.raw_sockets	checkbox	Allow raw sockets. Utilities like ping(8) (https://www.freebsd.org/cgi/man.cgi?query=ping) and traceroute(8) (https://www.freebsd.org/cgi/man.cgi?query=traceroute) require raw sockets to operate inside a jail. When set, the source IP addresses are enforced to comply with the IP address bound to the jail, ignoring the IP_HDRINCL flag on the socket.
allow.chflags	checkbox	Treat jail users as privileged and allow the manipulation of system file flags. <i>securelevel</i> constraints are still enforced.
allow.mount	checkbox	Allow privileged users inside the jail to mount and unmount filesystem types marked as jail-friendly.
allow.mount.devfs	checkbox	Allow privileged users inside the jail to mount and unmount the devfs(5) device filesystem (https://www.freebsd.org/cgi/man.cgi?query=devfs). This permission is only effective when <i>allow_mount</i> is set and <i>enforce_statfs</i> is set to a value lower than 2.
allow.mount.nullfs	checkbox	Allow privileged users inside the jail to mount and unmount the nullfs(5) file system (https://www.freebsd.org/cgi/man.cgi?query=nullfs). This permission is only effective when <i>allow_mount</i> is set and <i>enforce_statfs</i> is set to a value lower than 2.
allow.mount.procfs	checkbox	Allow privileged users inside the jail to mount and unmount the procfs(5) file system (https://www.freebsd.org/cgi/man.cgi?query=procfs). This permission is only effective when <i>allow_mount</i> is set and <i>enforce_statfs</i> is set to a value lower than 2.
allow.mount.tmpfs	checkbox	Allow privileged users inside the jail to mount and unmount the tmpfs(5) file system (https://www.freebsd.org/cgi/man.cgi?query=tmpfs). This permission is only effective when <i>allow_mount</i> is set and <i>enforce_statfs</i> is set to a value lower than 2.
allow.mount.zfs	checkbox	Allow privileged users inside the jail to mount and unmount the ZFS file system. This permission is only effective when <i>allow_mount</i> is set and <i>enforce_statfs</i> is set to a value lower than 2. The ZFS(8) (https://www.freebsd.org/cgi/man.cgi?query=zfs) man page has information on how to configure the ZFS filesystem to operate from within a jail.
allow.quotas	checkbox	Allow the jail root to administer quotas on the jail filesystems. This includes filesystems the jail shares with other jails or with non-jailed parts of the system.
allow.socket_af	checkbox	Allow access to other protocol stacks beyond IPv4, IPv6, local (UNIX), and route. Warning: jail functionality does not exist for all protocol stacks.

Click *NEXT* to view all jail *Network Properties*. These are shown in [Table 14.3](#):

Table 14.3: Network Properties

Setting	Value	Description
interfaces	string	Enter up to four interface configurations in the format <i>interface:bridge</i> , separated by a comma (,). The left value is the virtual VNET interface name and the right value is the bridge name where the virtual interface is attached.

Continued on next page

Table 14.3 – continued from previous page

Setting	Value	Description
host_domainname	string	Enter an NIS Domain name (https://www.freebsd.org/doc/handbook/network-nis.html) for the jail.
host.hostname	string	Enter a hostname for the jail. By default, the system uses the jail UUID.
exec.fib	integer	Enter a number to define the routing table (FIB) to set when running commands inside the jail.
ip4.saddrsel	checkbox	Only available when the jail is not configured to use VNET. Disables IPv4 source address selection for the jail in favor of the primary IPv4 address of the jail.
ip4	drop-down	Control the availability of IPv4 addresses. Set to <i>Inherit</i> : allow unrestricted access to all system addresses. Set to <i>New</i> : restrict addresses with <i>ip4_addr</i> . Set to <i>Disable</i> : stop the jail from using IPv4 entirely.
ip6.saddrsel	string	Only available when the jail is not configured to use VNET. Disables IPv6 source address selection for the jail in favor of the primary IPv6 address of the jail.
ip6	drop-down	Control the availability of IPv6 addresses. Set to <i>Inherit</i> : allow unrestricted access to all system addresses. Set to <i>New</i> : restrict addresses with <i>ip6_addr</i> . Set to <i>Disable</i> : stop the jail from using IPv6 entirely.
resolver	string	Add lines to <code>resolv.conf</code> in file. Example: <i>nameserver IP;search domain.local</i> . Fields must be delimited with a semicolon (;), this is translated as new lines in <code>resolv.conf</code> . Enter <i>none</i> to inherit <code>resolv.conf</code> from the host.
mac_prefix	string	Optional. Enter a valid MAC address vendor prefix. Example: <i>E4F4C6</i>
vnet_default_interface	drop-down	Default network interface used for the VNET bridge interface in the jail. Only takes effect when <i>VNET</i> is set and bridge interfaces are not active.
vnet0_mac	string	Optional. Enter a valid MAC address for the VNET0 interface.
vnet1_mac	string	Optional. Enter a valid MAC address for the VNET1 interface.
vnet2_mac	string	Optional. Enter a valid MAC address for the VNET2 interface.
vnet3_mac	string	Optional. Enter a valid MAC address for the VNET3 interface.

The final set of jail properties are contained in the *Custom Properties* section. [Table 14.4](#) describes these options.

Table 14.4: Custom Properties

Setting	Value	Description
owner	string	The owner of the jail. Can be any string.
priority	integer	The numeric start priority for the jail at boot time. Smaller values mean a higher priority. At system shutdown, the priority is <i>reversed</i> . Example: 99
hostid	string	A new a jail hostid, if necessary. Example hostid: <i>1a2bc345-678d-90e1-23fa-4b56c78901de</i> .
comment	string	Comments about the jail.
depends	string	Specify any jails the jail depends on. Child jails must already exist before the parent jail can be created.
mount.procfs	checkbox	Allow mounting of a procfs(5) (https://www.freebsd.org/cgi/man.cgi?query=procfs) filesystems in the jail <code>/dev/proc</code> directory.
mount_linprocfs	checkbox	Allow mounting of a linprocfs(5) (https://www.freebsd.org/cgi/man.cgi?query=linprocfs) filesystem in the jail.
host_time	checkbox	Synchronize the time between jail and host.

Continued on next page

Table 14.4 – continued from previous page

Setting	Value	Description
jail_zfs	checkbox	Enable automatic ZFS jailing inside the jail. The assigned ZFS dataset is fully controlled by the jail. Note: <i>allow_mount</i> , <i>enforce_statfs</i> , and <i>allow_mount_zfs</i> must all be set for ZFS management inside the jail to work correctly.
jail_zfs_dataset	string	Define the dataset to be jailed and fully handed over to a jail. Enter a ZFS filesystem name without a pool name. <i>jail_zfs</i> must be set for this option to work.
jail_zfs_mountpoint	string	The mountpoint for the <i>jail_zfs_dataset</i> . Example: <i>/data/example-dataset-name</i>

Click *SAVE* when the desired jail properties have been set. New jails are added to the primary list in the *Jails* menu.

14.3 Managing Jails

Clicking *Jails* shows a list of installed jails. An example is shown in [Figure 14.3](#).

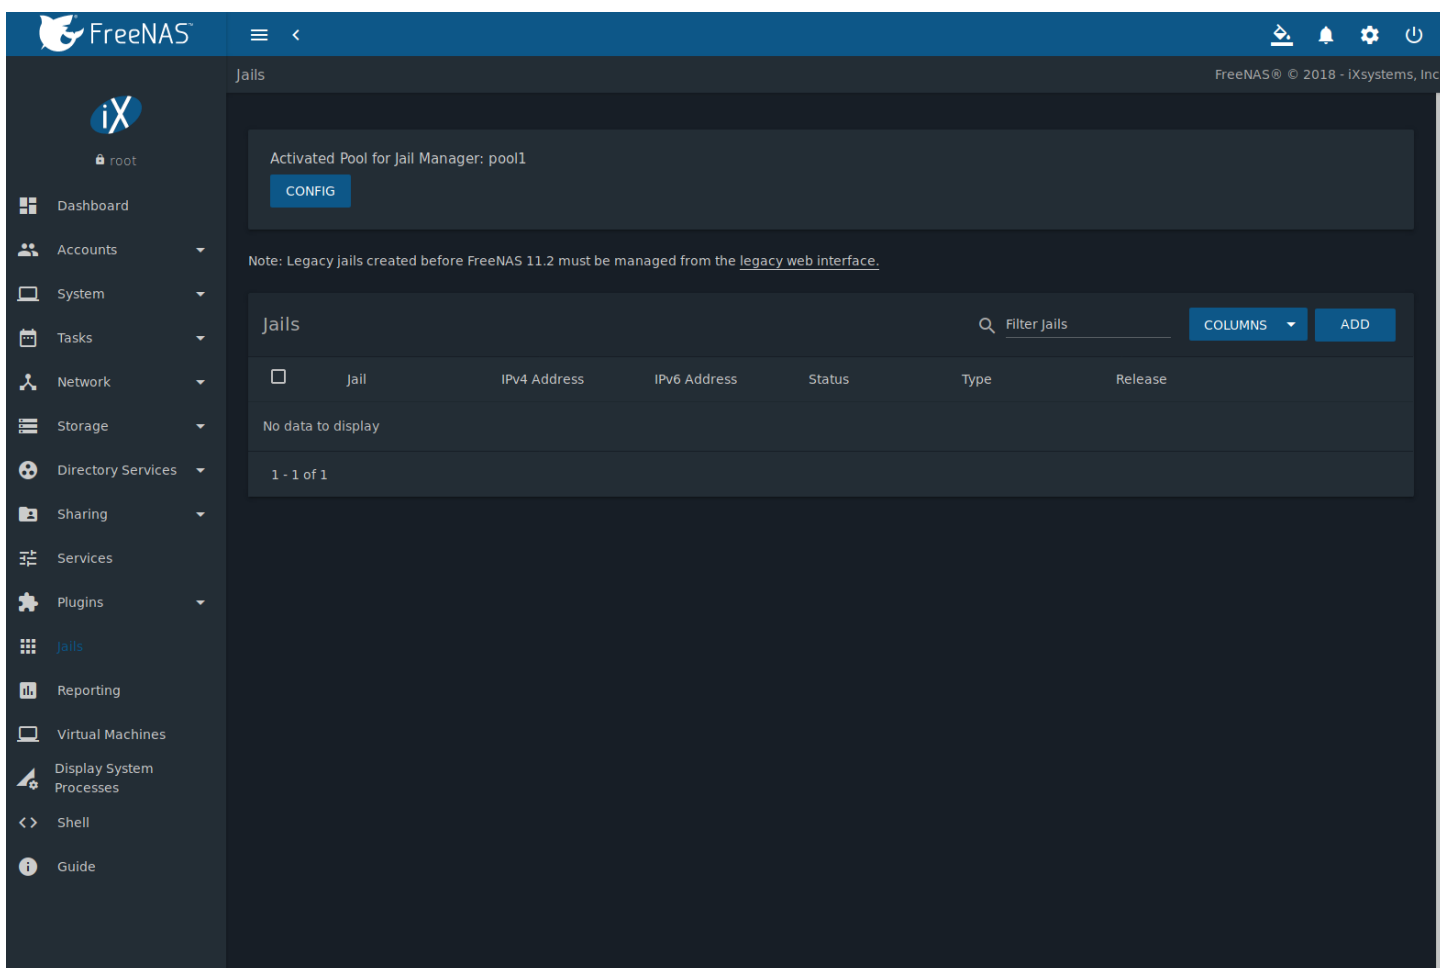


Fig. 14.3: Jail Overview Section

Table 14.5 describes each column.

Table 14.5: Jail Overview Information

Column Name	Description
Jail	The name of the jail.
IPv4 Address	Listing of configured IPv4 addresses.
IPv6 Address	Listing of configured IPv6 addresses.
Status	<i>up</i> indicates the jail is running and <i>down</i> indicates the jail is stopped.
Type	Indicates the installation method where <i>jail</i> was installed using <i>jails</i> (page 281) and <i>plug-inv2</i> was installed using <i>Plugins</i> (page 271).
Release	The FreeBSD version the jail is based on.
ⓘ (Options)	Click to display the options shown in Figure 14.4.

Operations can be applied to multiple jails by selecting those jails with the checkboxes on the left. After selecting one or more jails, icons appear which can be used to ⓘ (Start), ⓘ (Stop), ⓘ (Update), or ⓘ (Delete) those jails.

Click ⓘ (Options) for a jail to see all options for that jail. Figure 14.4 shows the menu that appears.

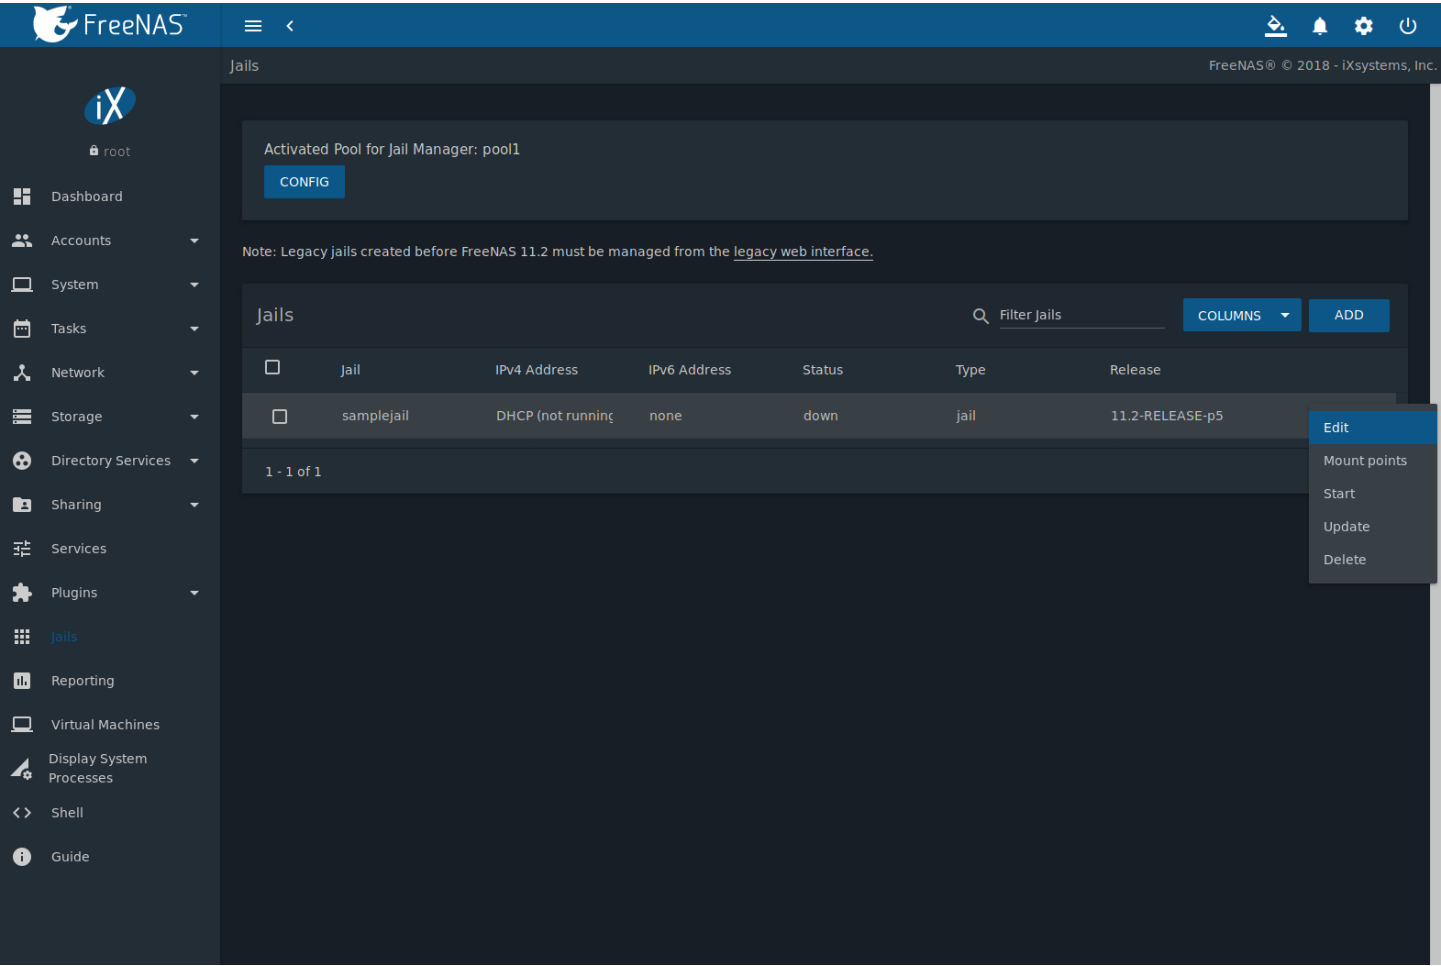


Fig. 14.4: Jail Options Menu

Table 14.6 describes each option available for a jail.


Warning: Modify the IP address information for a jail by using ⓘ (Options) *Edit* instead of issuing the networking commands directly from the command line of the jail. This ensures the changes are saved and will survive a jail or FreeNAS® reboot.


Table 14.6: Jail Option Menu Entry Descriptions

Option	Description
Edit	Used to modify the settings described in Table 14.5 . A popup error will display if the jail's <i>Status</i> is <i>up</i> . A jail cannot be edited while it is running.
Mount points	Open the <i>Mount Points</i> list. Select an existing mount point to <i>Edit</i> or click <i>ADD</i> to open the <i>Add Mount Point</i> screen. A mount point gives a jail access to storage located elsewhere on the system. A jail must be stopped before adding, editing, or deleting a <i>Mount Point</i> . See Additional Storage (page 293) for more details.
Restart	Stop and immediately start an <i>up</i> jail.
Start	Start a jail that has a current <i>Status</i> of <i>down</i> .
Stop	Stop a jail that has a current <i>Status</i> of <i>up</i> .
Update	Runs <code>freebsd-update</code> (https://www.freebsd.org/cgi/man.cgi?query=freebsd-update) to update the jail to the latest patch level of the installed FreeBSD release.
Shell	Access a <i>root</i> command prompt to interact with a jail directly from the command line. Type <code>exit</code> to leave the command prompt.
Delete	Delete the jail, all of the jail's contents, and all associated Snapshots (page 172). Back up the jail's data, configuration, and programs first. There is no way to recover the contents of a jail after deletion!

Note: Menu entries change depending on the jail state. For example, a stopped jail does not have a *Stop* or *Shell* option.

14.3.1 Jail Updates and Upgrades



Click  (Options) → *Update* to update a jail to the most current patch level of the installed FreeBSD release. This does **not** change the release.

To *upgrade* a jail to newer release of FreeBSD, stop the jail and click  (Options) → *Edit* for the jail. Open the *Release* drop-down menu, choose a newer RELEASE of FreeBSD, and click *SAVE*. Upgrading a jail can take an extended amount of time, depending on connection speed and if the chosen RELEASE is already fetched on the system.

Tip: It is possible to [manually remove](#) (page 167) unused releases from the `/iocage/releases/` dataset after upgrading a jail. The release **must** not be in use by any jail on the system!

14.3.2 Accessing a Jail Using SSH

The ssh daemon `sshd(8)` (<https://www.freebsd.org/cgi/man.cgi?query=sshd>) must be enabled in a jail to allow SSH access to that jail from another system.

The jail *Status* must be up before the *Shell* option is available. If the jail is not up, start it by clicking *Jails* →  (Options) → *Start* for the desired jail. Click  (Options) → *Shell* to start a shell on the jail. A jail root shell is shown in this example:

```
Last login: Fri Apr 6 07:57:04 on pts/12
FreeBSD 11.1-STABLE (FreeNAS.amd64) #0 0ale9f753(freenas/11-stable): Fri Apr 6 04:46:31 UTC 2018

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@jailexamp:~ #
```

Tip: A root shell can also be opened for a jail using the FreeNAS® UI *Shell*. Open the *Shell*, then type `iocage console jailname`.

Enable sshd:

```
sysrc sshd_enable="YES"
sshd_enable: NO -> YES
```

Tip: Using `sysrc` to enable sshd verifies that sshd is enabled.

Start the SSH daemon: `service sshd start`

The first time the service runs, the jail RSA key pair is generated and the key fingerprint is displayed.

Add a user account with `adduser`. Follow the prompts, Enter will accept the default value offered. Users that require *root* access must also be a member of the *wheel* group. Enter *wheel* when prompted to *invite user into other groups?* []:

```
root@jailexamp:~ # adduser
Username: jailuser
Full name: Jail User
Uid (Leave empty for default):
Login group [jailuser]:
Login group is jailuser. Invite jailuser into other groups? []: wheel
Login class [default]:
Shell (sh csh tcsh git-shell zsh rzsh nologin) [sh]: csh
Home directory [/home/jailuser]:
Home directory permissions (Leave empty for default):
Use password-based authentication? [yes]:
Use an empty password? (yes/no) [no]:
Use a random password? (yes/no) [no]:
Enter password:
Enter password again:
Lock out the account after creation? [no]:
Username   : jailuser
Password   : *****
Full Name  : Jail User
Uid        : 1002
Class      :
Groups     : jailuser wheel
Home       : /home/jailuser
Home Mode  :
Shell      : /bin/csh
Locked     : no
OK? (yes/no): yes
adduser: INFO: Successfully added (jailuser) to the user database.
Add another user? (yes/no): no
```



```
Goodbye!
root@jailexamp:~
```

After creating the user, set the jail *root* password to allow users to use *su* to gain superuser privileges. To set the jail *root* password, use *passwd*. Nothing is echoed back when using *passwd*

```
root@jailexamp:~ # passwd
Changing local password for root
New Password:
Retype New Password:
root@jailexamp:~ #
```

Finally, test that the user can successfully *ssh* into the jail from another system and gain superuser privileges. In the example, a user named *jailuser* uses *ssh* to access the jail at 192.168.2.3. The host RSA key fingerprint must be verified the first time a user logs in.


```
ssh jailuser@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:
```


Note: Every jail has its own user accounts and service configuration. These steps must be repeated for each jail that requires SSH access.

14.3.3 Additional Storage

Jails can be given access to an area of storage outside of the jail that is configured on the FreeNAS® system. It is possible to give a FreeBSD jail access to an area of storage on the FreeNAS® system. This is useful for applications or plugins that store large amounts of data or if an application in a jail needs access to data stored on the FreeNAS® system. For example, Transmission is a plugin that stores data using BitTorrent. The %brand\$ external storage is added using the [mount_nullfs\(8\)](https://www.freebsd.org/cgi/man.cgi?query=mount_nullfs(8)) ([https://www.freebsd.org/cgi/man.cgi?query=mount_nullfs\(8\)](https://www.freebsd.org/cgi/man.cgi?query=mount_nullfs(8))) mechanism, which links data that resides outside of the jail as a storage area within a jail.

The *Mount points* section of a jail shows any added storage and allows adding more storage.

Note: A jail must have a *Status* of *down* before adding a new mount point. Click  (Options) and *Stop* for a jail to change the jail *Status* to *down*.

Storage can be added by clicking *Jails* →  (Options) → *Mount points* for the desired jail. The *Mount points* section is a list of all of the currently defined mount points.

Go to *Mount points* → *ADD* to add storage to a jail. This opens the screen shown in [Figure 14.5](#).

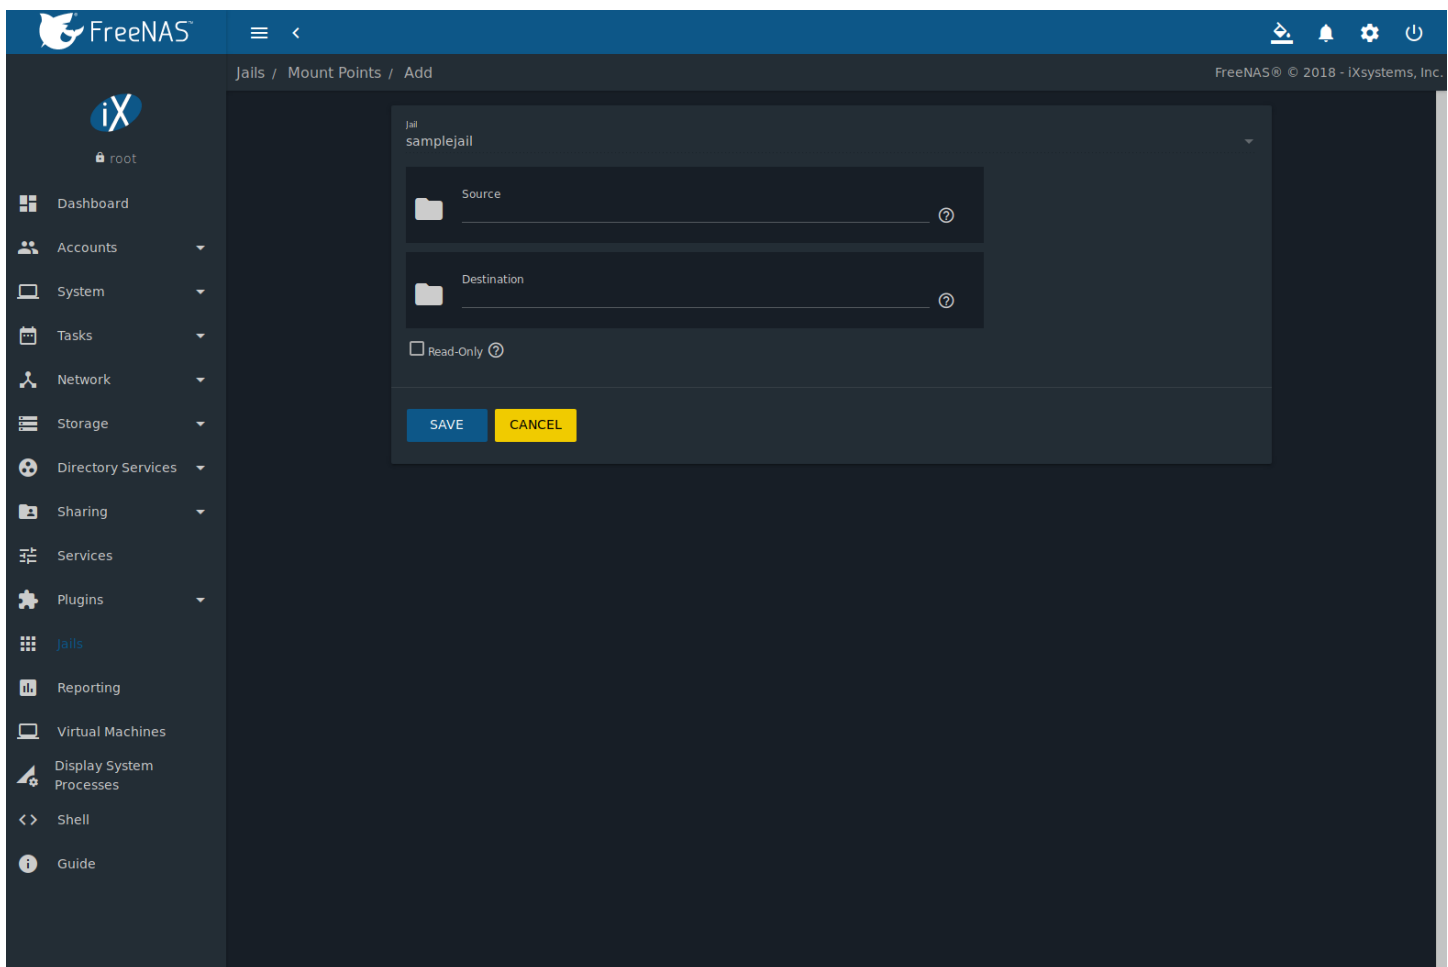


Fig. 14.5: Adding Storage to a Jail

Browse to the *Source* and *Destination*, where:

- *Source*: is the directory or dataset on the FreeNAS® system which will be accessed by the jail. FreeNAS® creates the directory if it does not exist. This directory must reside outside of the pool 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*: Browse to an existing and **empty** directory within the jail to link to the *Source* storage area. It is also possible to add / and a name to the end of the path and FreeNAS® automatically creates a new directory. New directories created must be **within** the jail directory structure. Example: `/mnt/iocage/jails/samplejail/root/new-destination-directory`.

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 will be given access to existing data or a new storage area will be allocated.
4. If the jail accesses existing data, edit the permissions of the pool 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 jail *Mount points* → *ADD* to select the the *Source* of the data and the *Destination* where it will be mounted in the jail.

To prevent writes to the storage, click *Read-Only*.

After storage has been added or created, it appears in the *Mount points* for that jail. In the example shown in [Figure 14.6](#), a dataset named `pool1/smb-storage` has been chosen as the *Source* as it contains the files stored on the FreeNAS® system. The user entered `/mnt/iocage/jails/samplejail/root/mounted` as the directory to be mounted in the *Destination* field. To users inside the jail, this data will appear to be in the `/root/mounted` directory.

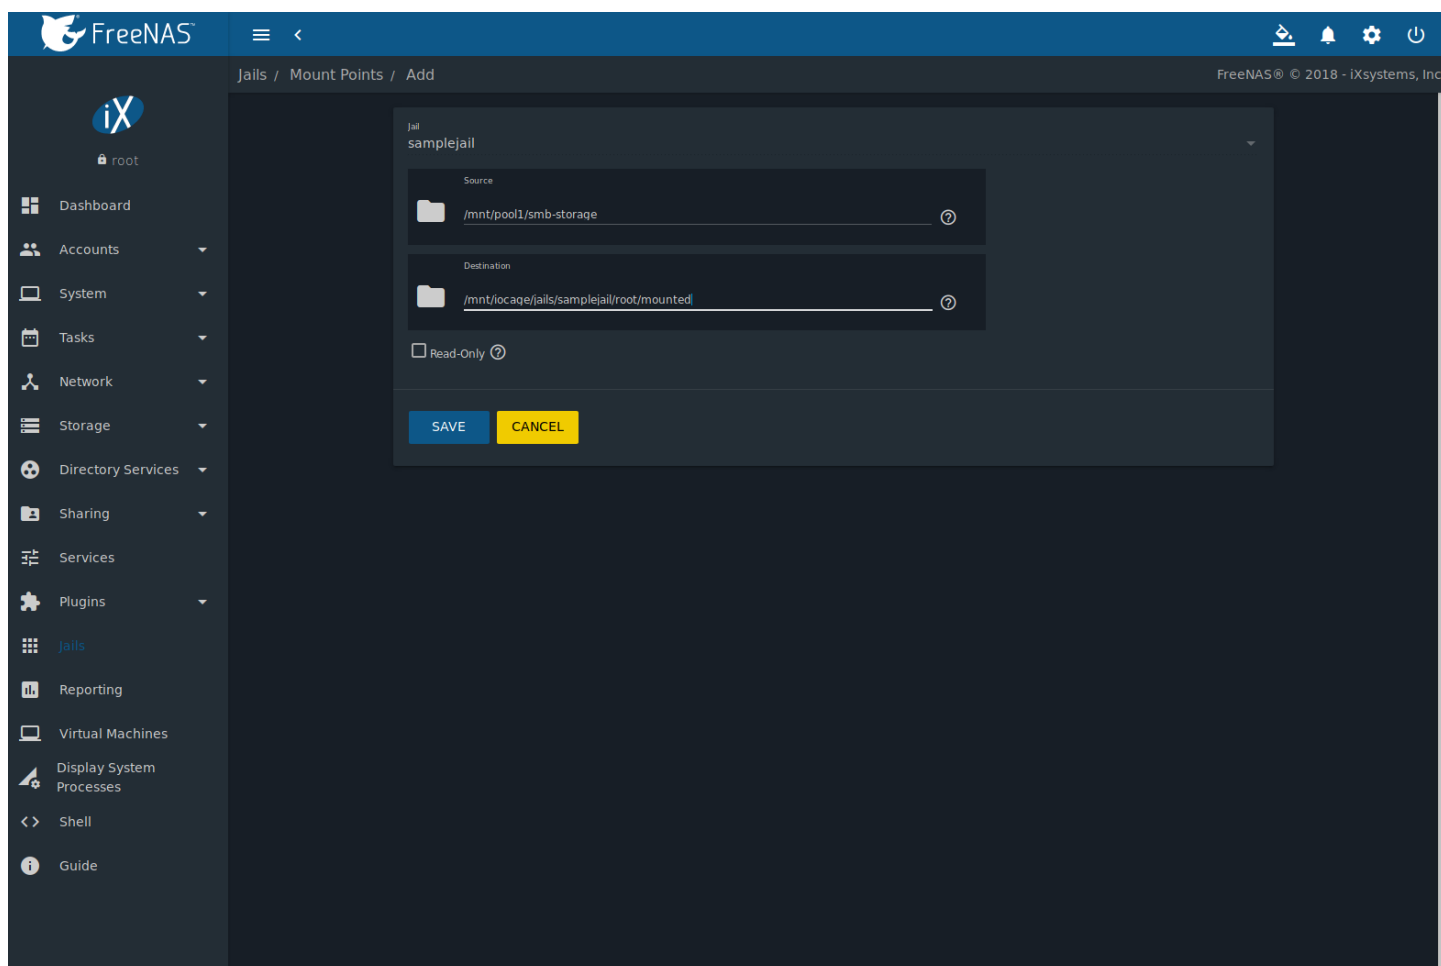


Fig. 14.6: Example Storage

Storage is automatically mounted as it is created.

Note: Mounting a dataset does not automatically mount any child datasets inside it. Each dataset is a separate filesystem,

so child datasets must each have separate mount points.

Click  (Options) → *Delete* to delete the storage.

Warning: Remember that added storage is 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. This leaves the data intact but not accessible from the jail.

14.4 Jail Software

A jail is created with no software aside from the core packages installed as part of the selected version of FreeBSD. Software in a jail is managed by going to the *Shell* and logging into the jail with `iocage console {jailname}`. In this example, the user has logged into *testjail01*:

```
[root@freenas ~]# iocage console testjail01
FreeBSD 11.1-STABLE (FreeNAS.amd64) #0 35e0ef284(freenas/11-stable): Mon Apr  9 17:44:36 UTC 2018

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@testjail01:~ #
```

Tip: See *Using iocage* (page 302) for more details about different `iocage` commands.

14.4.1 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 and 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. Most of that software is available as packages. One way to find FreeBSD software is to use the search bar at [FreshPorts.org](https://www.freshports.org/) (<https://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 the command `pkg install audiotag`

When prompted, press `y` to complete the installation. Messages will show the download and installation status.

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`.

14.4.2 Compiling FreeBSD Ports

Compiling a port is another option. Compiling ports offer 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 resources to use the `pkg install` command instead.

The [FreshPorts.org](https://www.freshports.org/) (<https://www.freshports.org/>) listing shows whether a port has any configurable compile options. [Figure 14.7](#) shows the *Configuration Options* for audiotag.



If you buy from Amazon USA, please support us by using [this link](#).

Follow us

[Blog](#)
[Twitter](#)
[Status page](#)

Port details			Login
audiotag Command-line tool for mass tagging/renaming of audio files 0.19_1 audio Z=1 🔍 🔥 🔧			User Login Create account
<p>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 🔍</p> <p>Port Added: 15 Apr 2008 13:43:37 Also Listed In: multimedia License: GPLv2+</p> <p>Audiotag is a command-line tool for mass tagging/renaming of audio files it supports the vorbis comment, id3 tags, and MP4 tags.</p> <p>WWW: https://github.com/Daenyth/audiotag</p> <p>SVNWeb : Homepage : PortsMon Pseudo-pkg-plist information, but much better, from make generate-plist ▼ Expand this list (4 items)</p>			Servers and bandwidth provided by New York Internet , SuperNews , and RootBSD
<hr/>			This site
<p>What is FreshPorts? About the authors Issues FAQ How big is it? The latest upgrade! Privacy Forums Blog Contact</p>			
<hr/>			Search
<p>Enter Keywords:</p> <div><input type="text"/></div> <div><input type="button" value="go"/></div> <p>more...</p>			
<hr/>			Latest Vulnerabilities
			moodle31 Dec 03
			moodle33 Dec 03
			rubygem-activejob4 Dec 02
			rubygem-activejob5 Dec 02
			rubygem-activejob50 Dec 02
			uriparser Dec 02
			gitlab-ce Nov 28
			messagelib Nov 28
			payara Nov 28
			phpmailer* Nov 23
			phpmailer6* Nov 23
			php56-imag* Nov 22
			gitlab-ce Nov 20
			linux-flashplayer Nov 20
			powerdns Nov 19
			12 vulnerabilities affecting 51 ports have been reported in the past 14 days
			* - modified, not new All vulnerabilities
			Last updated: 2018-12-03 02:25:50
			Ports
			Home Categories Deleted ports Sanity Test Failures Newsfeeds
			Statistics
			Graphs NEW Graphs (javascript) Traffic
			Calculated hourly:
			Port count 36155
			Broken 122
			Deprecated 1190
			Ignore 383
			Forbidden 3
			Restricted 162
			No CDROM 74
			Vulnerable 39
			Expired 4
			Set to expire 1163
			Interactive 0
			new 24 hours 3
			new 48 hours 8
			new 7 days 24
			new fortnight 58
			new month 123
<hr/>			
Commit History - (may be incomplete: see SVNWeb link above for full details)			
Date	By	Description	
02 Dec 2016 09:21:59 🔍 🔥 0.19_1	mat 🔍	http://github.com redirects to https://github.com , spare everyone a redirect. Sponsored by: Absolight	
25 May 2016 15:43:34 🔍 🔥 0.19_1	mat 🔍	Remove useless WRKSRC definitions. While there, correct DEV_WARNINGS when they occur. Sponsored by: Absolight	
01 Apr 2016 13:29:17 🔍 🔥 0.19_1	mat 🔍	Remove \${PORTSDIR}/ from dependencies, Mk and categories a, b, and c. With hat: portmgr Sponsored by: Absolight	
12 Jan 2016 16:20:32 🔍 🔥 0.19_1	amdmi3 🔍	Convert LICENSE= "GPLxx # or later" to "GPLxx+" Approved by: portmgr blanket	
18 Nov 2015 10:14:05 🔍 🔥 0.19_1	amdmi3 🔍	- Clarify LICENSE - Add LICENSE_FILE - Add NO_ARCH - Switch to options helpers	
01 Jun 2014 13:03:14 🔍 🔥 0.19_1	ohauer 🔍	- USE_(BZIP2 XZ) -> USES= tar:(bzip2 xz)	
24 Nov 2013 18:38:39 🔍 🔥 0.19_1	bapt 🔍	Remove cruft	
24 Nov 2013 18:36:37 🔍 🔥 0.19_1	bapt 🔍	Support staging Use options helpers	
20 Sep 2013 14:36:37 🔍 🔥 0.19_1	bapt 🔍	Add NO_STAGE all over the place in preparation for the staging support (cat: audio)	
03 Aug 2013 13:44:01 🔍 🔥 0.19_1	mat 🔍	- Convert to new perl framework - Remove MAKE_JOBS_SAFE=yes, it's the default.	
06 May 2013 22:26:27 🔍 🔥 0.19_1	bapt 🔍	Use shebangfix	
23 Mar 2013 19:36:24 🔍 🔥 0.19_1	bapt 🔍	Fix USE_GITHUB in combinaison with MASTER_SITE= GHC which breaks WRKSRC Drop maintainership	
31 Dec 2012 11:31:44 🔍 🔥 0.19_1	bapt 🔍	- Trim headers - Switch to USE_GITHUB - Remove useless LICENSE_FILE license being plain GPLv2 - Various cleanup	

299

This port has five configurable options: *DOCS*, *FLAC*, *ID3*, *MP4*, and *VORBIS*. Stars (*) show which options are enabled.

Packages use default options. Ports let the user select options.

The Ports Collection must be installed in the jail before ports can be compiled. Inside the jail, use the `portsnap` utility. This command downloads the ports collection and extracts it to the `/usr/ports/` directory of the jail:

```
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
```

The first time this command is run, the configure screen shown in Figure 14.8 is displayed:

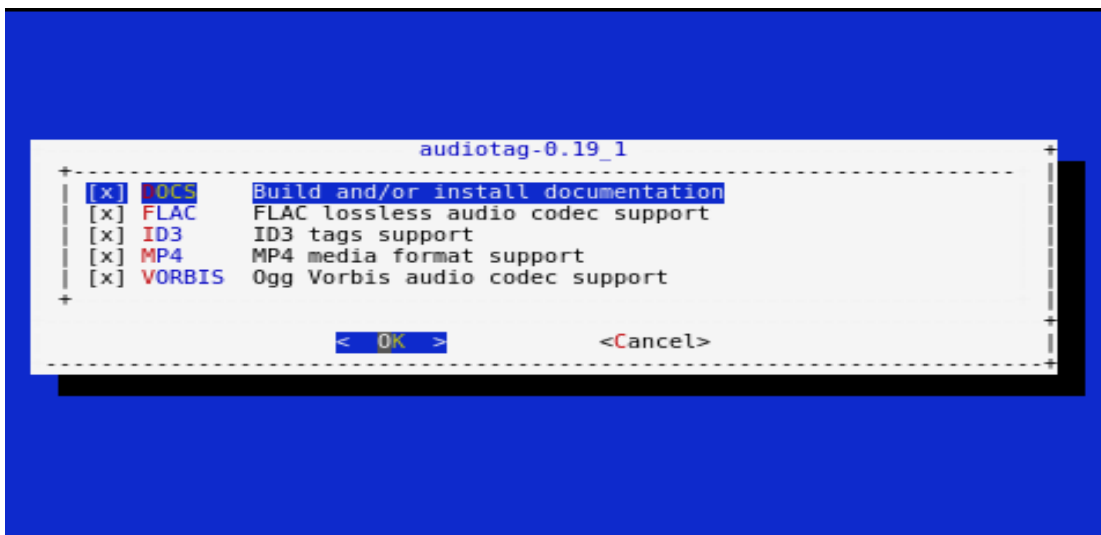


Fig. 14.8: Configuration Options for Audiotag Port

Use the arrow keys to select an option and press `spacebar` to toggle the value. Press `Enter` when satisfied with the jail options. The port begins to compile and install.

Note: After options have been set, the configuration screen is normally not shown again. Use `make config` to display the screen and change options before rebuilding the port with `make clean install clean`.

Many ports depend on other ports. Those other ports also have configuration screens that are shown before compiling begins. It is a good idea to watch the compile until it finishes and the command prompt returns.

Installed ports are registered in the same package database that manages packages. The `pkg info` can be used to determine which ports were installed.

14.4.3 Starting Installed Software

After packages or ports are installed, they must be configured and started. Configuration files are usually in `/usr/local/etc` or a subdirectory of it. Many FreeBSD packages contain a sample configuration file as a reference. Take some time to

read the software documentation 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, test starting the service by running the script with the `onestart` option. For example, with `openvpn` installed in the jail, these commands are run to 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 any error messages if an issue is found. Most startup failures are related to a misconfiguration 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
```

14.5 Using iocage

Beginning with FreeNAS® 11.0, 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**.

iocage has several options to help users:

- There is built-in help displayed by entering `iocage --help | less`. Each subcommand also has help. Display help by adding the `--help` flag after the subcommand name. For example, `iocage activate --help` shows help for the `activate` subcommand.
- The **iocage** manual page is accessed by typing `man iocage | less`.
- 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>).

14.5.1 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* has been 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 are 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 destroy examplejail

This will destroy jail examplejail

Are you sure? [y/N]: y
Destroying examplejail
```

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 the properties for **examplejail**. 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.
```

REPORTING

Reporting displays several graphs, as seen in Figure 15.1. Choose a category from the drop-down menu to view those graphs. There are also options to change the graph view and number of graphs on each page.

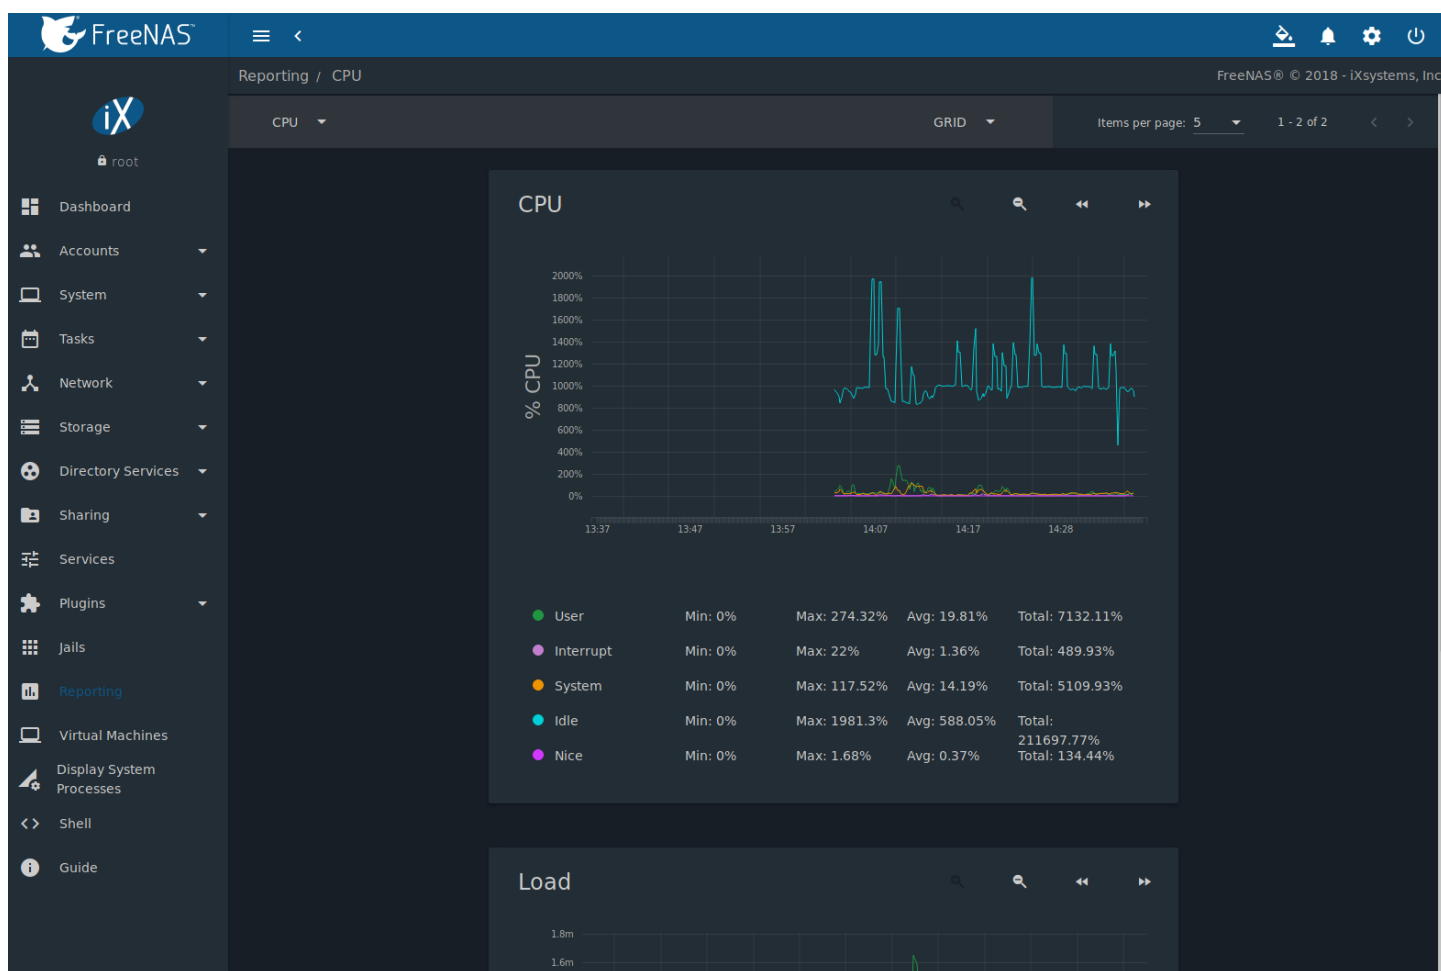


Fig. 15.1: Reporting Graphs

FreeNAS® uses [collectd](https://collectd.org/) (<https://collectd.org/>) to provide reporting statistics. For a clearer picture, hover over a point in the graph to show exact numbers for that point in time. 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 scroll through the output.

Graphs are grouped by category on the Reporting page:

- CPU

- [CPU](https://collectd.org/wiki/index.php/Plugin: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. It also shows the load on the CPU for a short term, mid term, and long term.
- *Disk*
 - [Disk](https://collectd.org/wiki/index.php/Plugin:Disk) (<https://collectd.org/wiki/index.php/Plugin:Disk>) shows read and write statistics on I/O, percent busy, latency, operations per second, pending I/O requests, and disk temperature. Choose the *DEVICES* and *METRICS* to view the selected metrics for the chosen devices.
- *Memory*
 - [Memory](https://collectd.org/wiki/index.php/Plugin:Memory) (<https://collectd.org/wiki/index.php/Plugin:Memory>) displays memory usage.
 - [Swap](https://collectd.org/wiki/index.php/Plugin: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, used, and reserved space for each pool and dataset. However, the disk space used by an individual zvol is not displayed as it is a block device.
- *System*
 - [Processes](https://collectd.org/wiki/index.php/Plugin:Processes) (<https://collectd.org/wiki/index.php/Plugin:Processes>) displays the number of processes. It is grouped by state.
- *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, demand data, demand meta-data, and prefetch data.

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* option. When deselected, data files are recorded in a temporary filesystem and copied hourly to on-disk files.

When *System* → *System Dataset Reporting database* is enabled, data files are written directly to the [System Dataset](#) (page 84).

Warning: Reporting data is frequently written and should not be stored on the boot pool or boot device.

[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://graphiteapp.org/) (<http://graphiteapp.org/>) server.

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 281), although there is additional overhead. A portion of system RAM is assigned to each VM, and each VM uses a *zvol* (page 169) 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>) 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 325) 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 325) 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>) `-H` option 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.

16.1 Creating VMs

Click on *Virtual Machines* and *ADD* to open the wizard shown in [Figure 16.1](#):

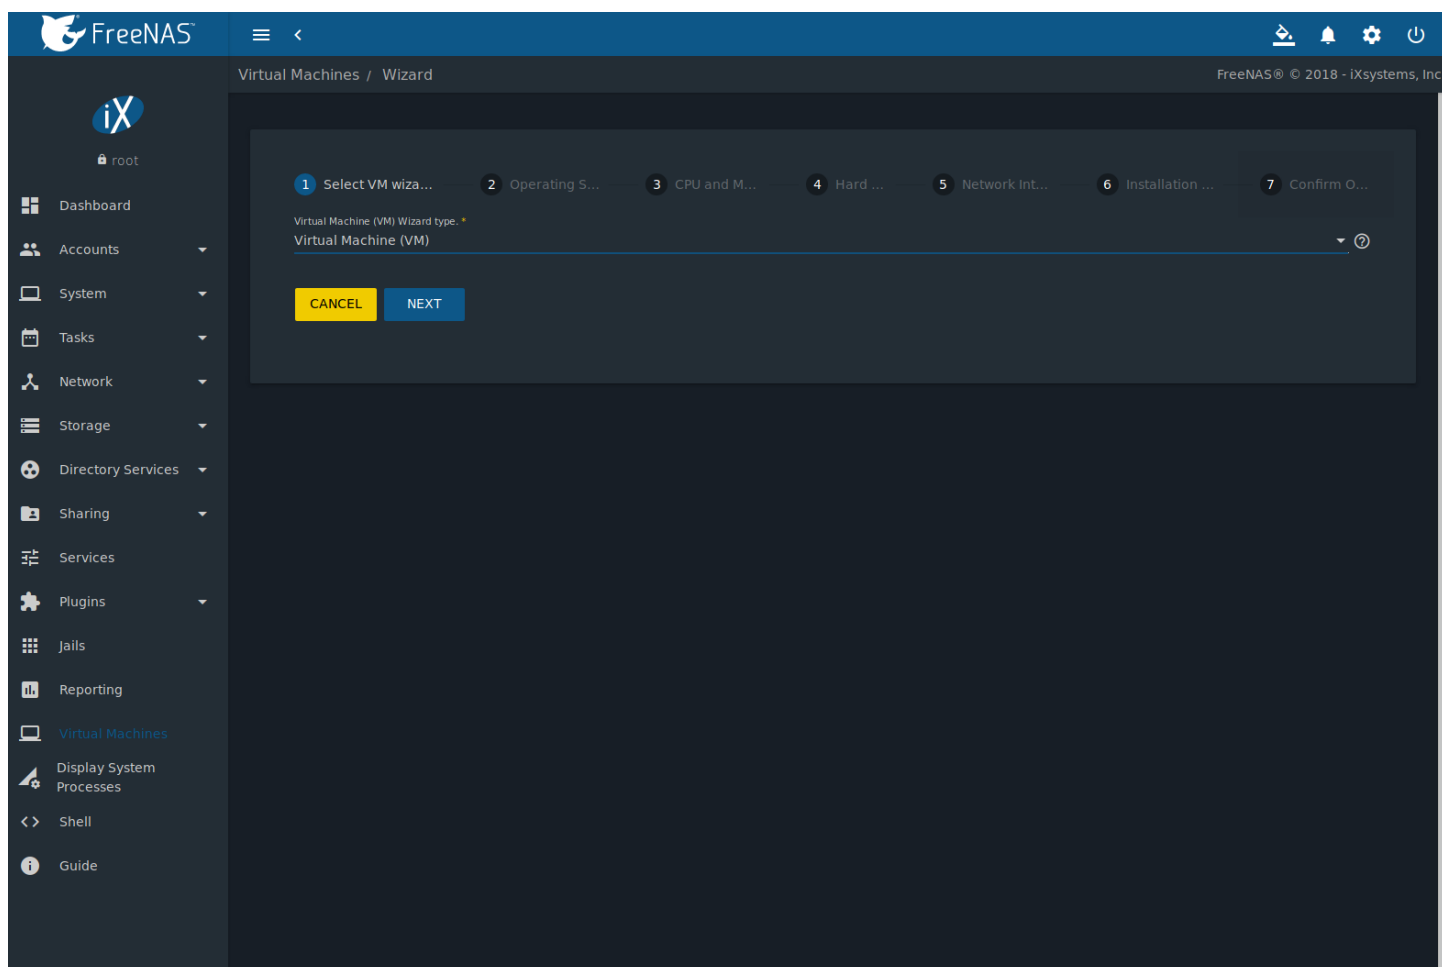


Fig. 16.1: Add VM

Virtual machine configuration options are described in [Table 16.1](#).

Table 16.1: VM Wizard Options

Screen #	Setting	Value	Description
1	Virtual Machine (VM) Wizard type	drop-down menu	Choose to create either a standard VM or a Docker Host.
2	Guest Operating System	drop-down menu	Choose the VM operating system type. Choices are: <i>Windows</i> , <i>Linux</i> , or <i>FreeBSD</i> . See this guide (https://github.com/FreeBSD-UPB/freebsd/wiki/How-to-launch-different-guest-OS) for detailed instructions about using a different guest OS.
2	VM Name	string	Enter an alphanumeric name to identify the VM. An error occurs if another VM is already using the same name.
2	Boot Method	drop-down menu	Select <i>UEFI</i> for newer operating systems, or <i>UEFI-CSM</i> for (Compatibility Support Mode) older operating systems that only understand BIOS booting.
2	Start on Boot	checkbox	Set to start the VM when the system boots.
2	Enable VNC	checkbox	Set to to activate a Virtual Network Computing (VNC) remote connection. Requires <i>UEFI</i> booting.

Continued on next page

Table 16.1 – continued from previous page

Screen #	Setting	Value	Description
3	Virtual CPUs	integer	Select the number of virtual CPUs to allocate to the VM. The maximum is 16 unless the host CPU limits the maximum. The VM operating system might also have operational or licensing restrictions on the number of CPUs.
3	Memory Size (MiB)	integer	Allocate the amount of RAM in mebibytes (https://simple.wikipedia.org/wiki/Mebibyte) for the VM.
4	Disk image	check option with custom fields	Select <i>Create new disk image</i> to create a new zvol on an existing dataset. This is used as a virtual hard drive for the VM. Select <i>Use existing disk image</i> to <i>Browse</i> to an existing zvol or file for the VM.
4	Zvol size (GiB)	integer	Allocate the amount of storage in gibibytes (https://simple.wikipedia.org/wiki/Gibibyte) for the new zvol. Only appears if <i>Create new disk image</i> is selected.
4	Select pool or dataset	drop-down menu	Select a pool or dataset for the new zvol.
4	Select disk type	drop-down menu	Select the disk type. Choices are <i>AHCI</i> and <i>VirtIO</i> . Refer to Disk Devices (page 313) for more information about these disk types.
5	Adapter Type	drop-down menu	<i>Intel e82545 (e1000)</i> emulates the same Intel Ethernet card. This provides compatibility with most operating systems. <i>VirtIO</i> provides better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.
5	MAC Address	string	Enter the desired MAC address to override the auto-generated randomized MAC address.
5	Attach NIC	drop-down menu	Select the physical interface to associate with the VM.
6	Choose installation media	browse button	<i>Browse</i> to the operating system installation media.
6	Upload ISO	checkbox and buttons	Set to display image upload options.

The final screen of the Wizard displays the chosen options for the new VM. Click *SUBMIT* to create the VM or *BACK* to change any settings.

This example creates a FreeBSD VM. *ADD* is clicked on the *Virtual Machines* page to start the VM wizard.

1. *Wizard type* is set to *Virtual Machine (VM)*.
2. *Guest Operating System* is set to *FreeBSD*. *Name* is set to *samplevm*. Other options are left at default values.
3. *Virtual CPUs* is set to 2 and *Memory Size* is set to 2048 MiB.
4. *Create new disk image* is selected. The zvol size is set to 20 GiB and stored on the *pool1* pool.
5. Network settings are left at default values.
6. A FreeBSD ISO installation image has been selected and uploaded to the FreeNAS® system. The *Choose installation media image* field is populated when the upload completes.
7. After verifying the *VM Summary* is correct, *SUBMIT* is clicked.

Figure 16.2 shows the confirmation step and basic settings for the new virtual machine:

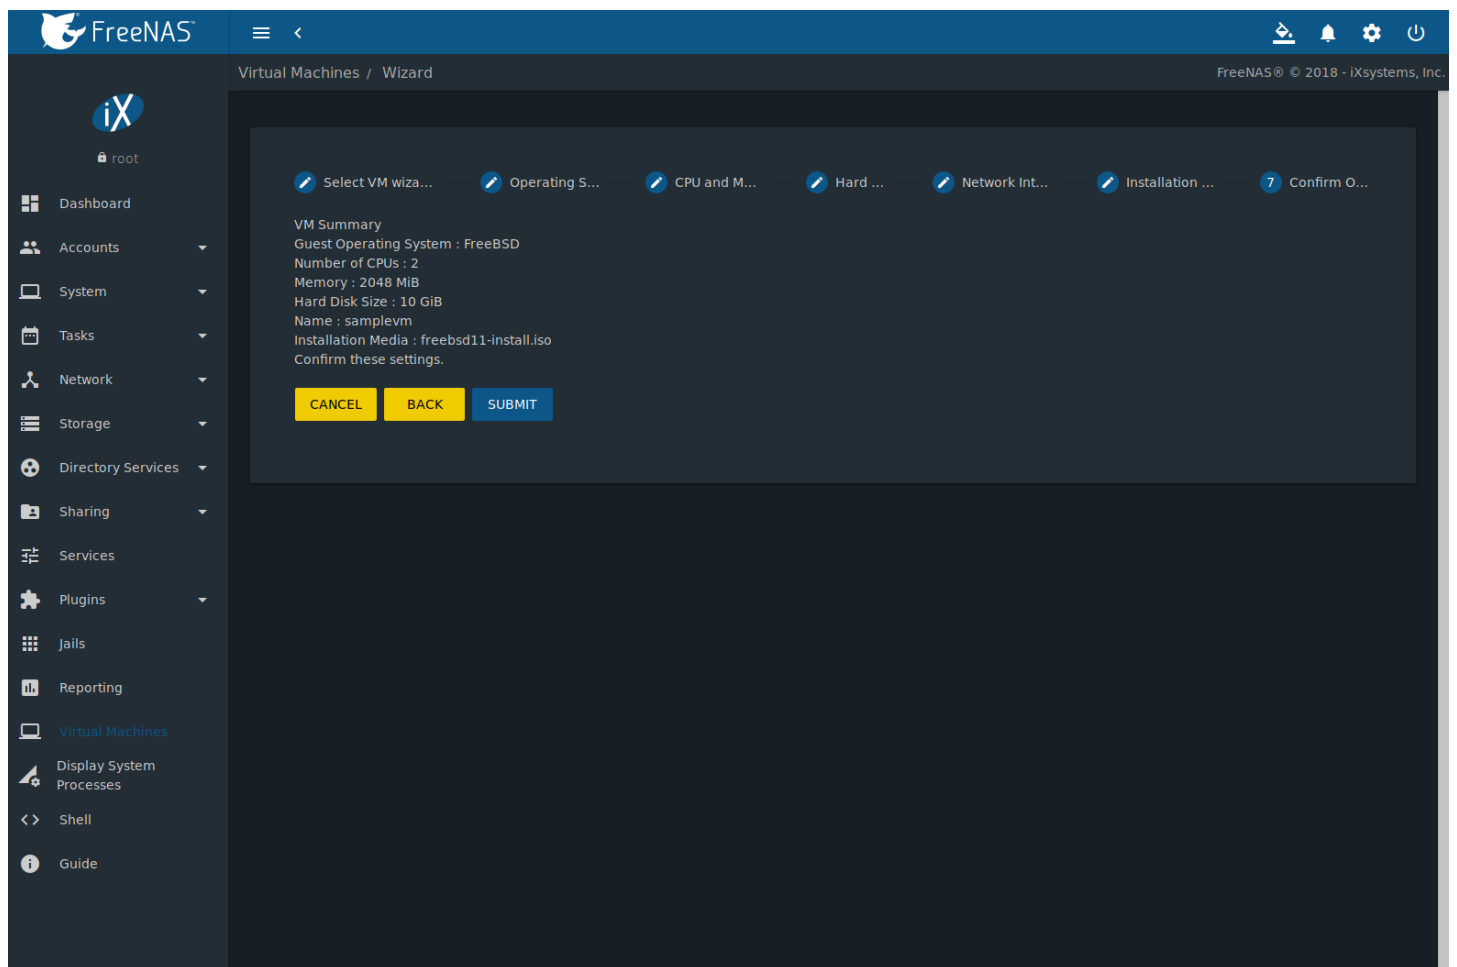


Fig. 16.2: Creating a Sample Virtual Machine

16.2 Adding Devices to a VM

To add a device to a VM, click **Options** → *Devices* on the VM card. Then click **ADD**.

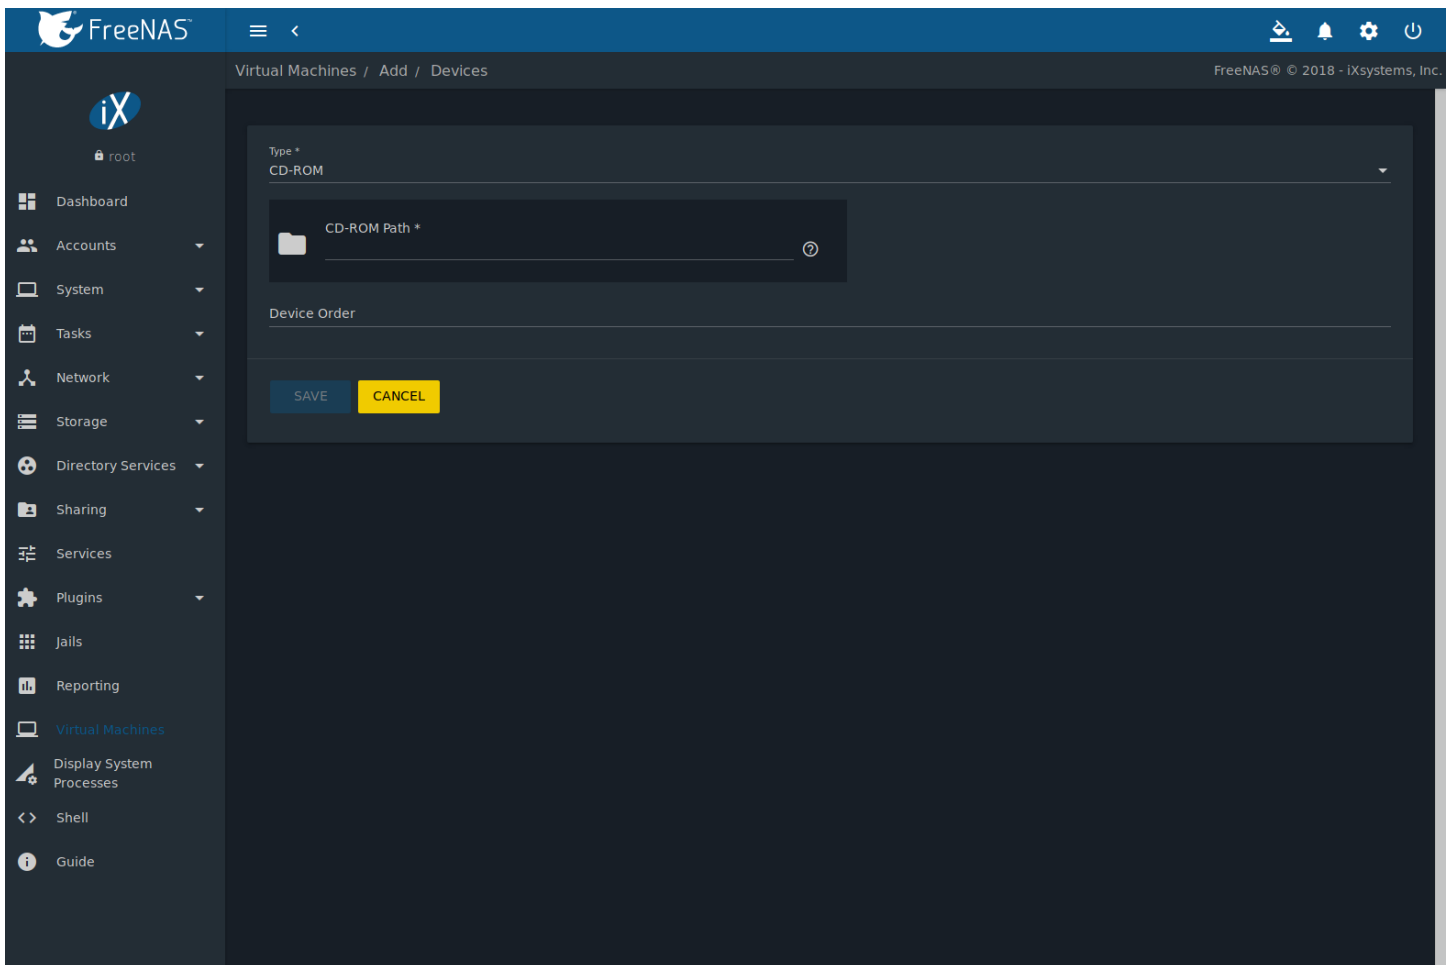


Fig. 16.3: VM Devices

Select the new device from the *Type* field. These devices are available:

- *CD-ROM* (page 311)
- *NIC (Network Interface Card)* (page 312)
- *Disk Device* (page 313)
- *Raw File* (page 314)
- *VNC Interface* (page 315)

Virtual Machines → ⓘ (Options) → *Devices* is also used to edit or delete an existing device for a specific VM. Click ⓘ (Options) for a device to display *Edit*, *Delete*, *Change Device Order*, and *Details*.

16.2.1 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:

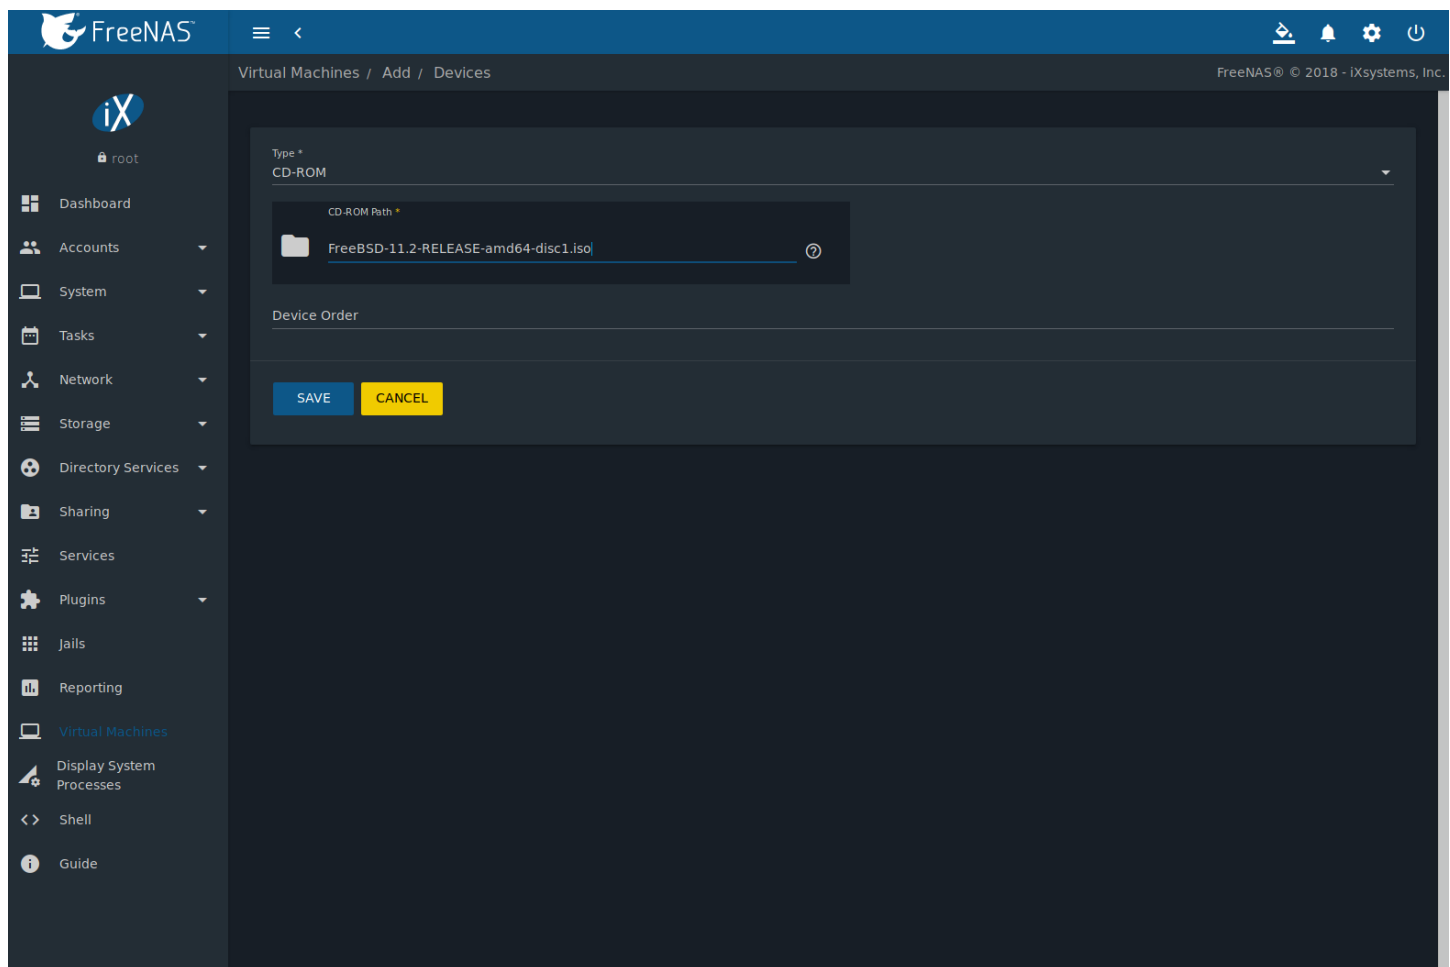


Fig. 16.4: CD-ROM Device

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 169) 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.

16.2.2 NIC (Network Interfaces)

Figure 16.5 shows the fields that appear after going to *Virtual Machines* → *Options* → *Devices*, clicking *ADD*, and selecting *NIC* as the *Type*.

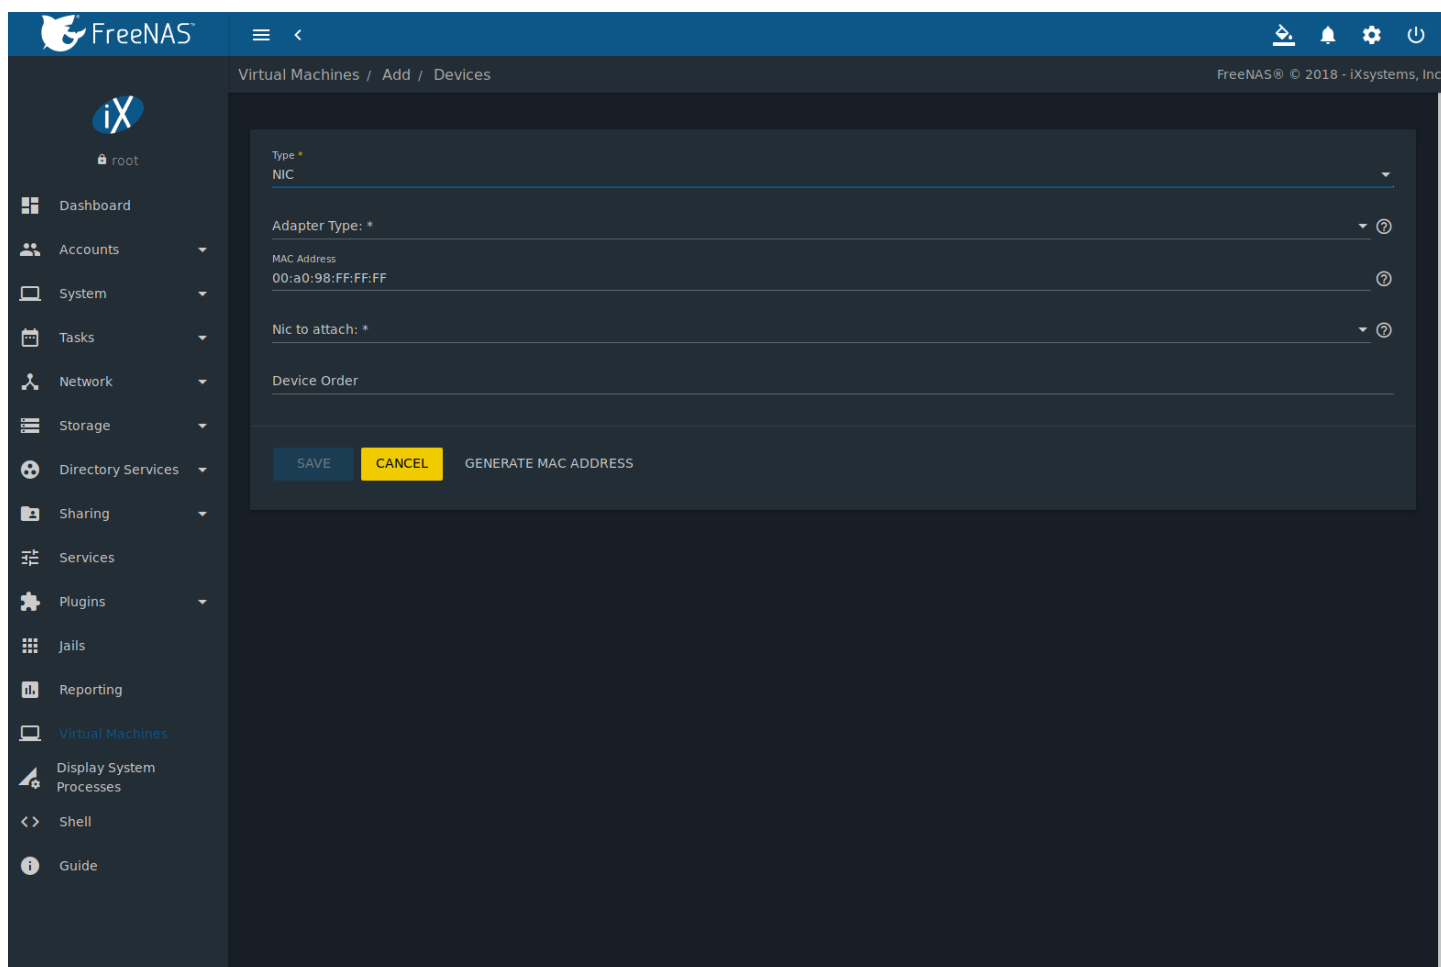


Fig. 16.5: Network Interface Device

The *Adapter Type* can emulate an Intel e82545 (e1000) Ethernet card for compatibility with most operating systems. *VirtIO* can provide better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.

By default, the VM receives an auto-generated random MAC address. To override the default with a custom value, enter the desired address in *MAC Address*. Click *GENERATE MAC ADDRESS* to automatically populate *MAC Address* with a new randomized MAC address.

If the system has multiple physical network interface cards, use the *NIC to attach* drop-down menu to specify which physical interface to associate with the VM.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.

16.2.3 Disk Devices

Zvols (page 169) are typically used as virtual hard drives. After *creating a zvol* (page 169), associate it with the VM by clicking *Virtual Machines* → *Options* → *Devices*, clicking *ADD*, and selecting *Disk* as the *Type*.

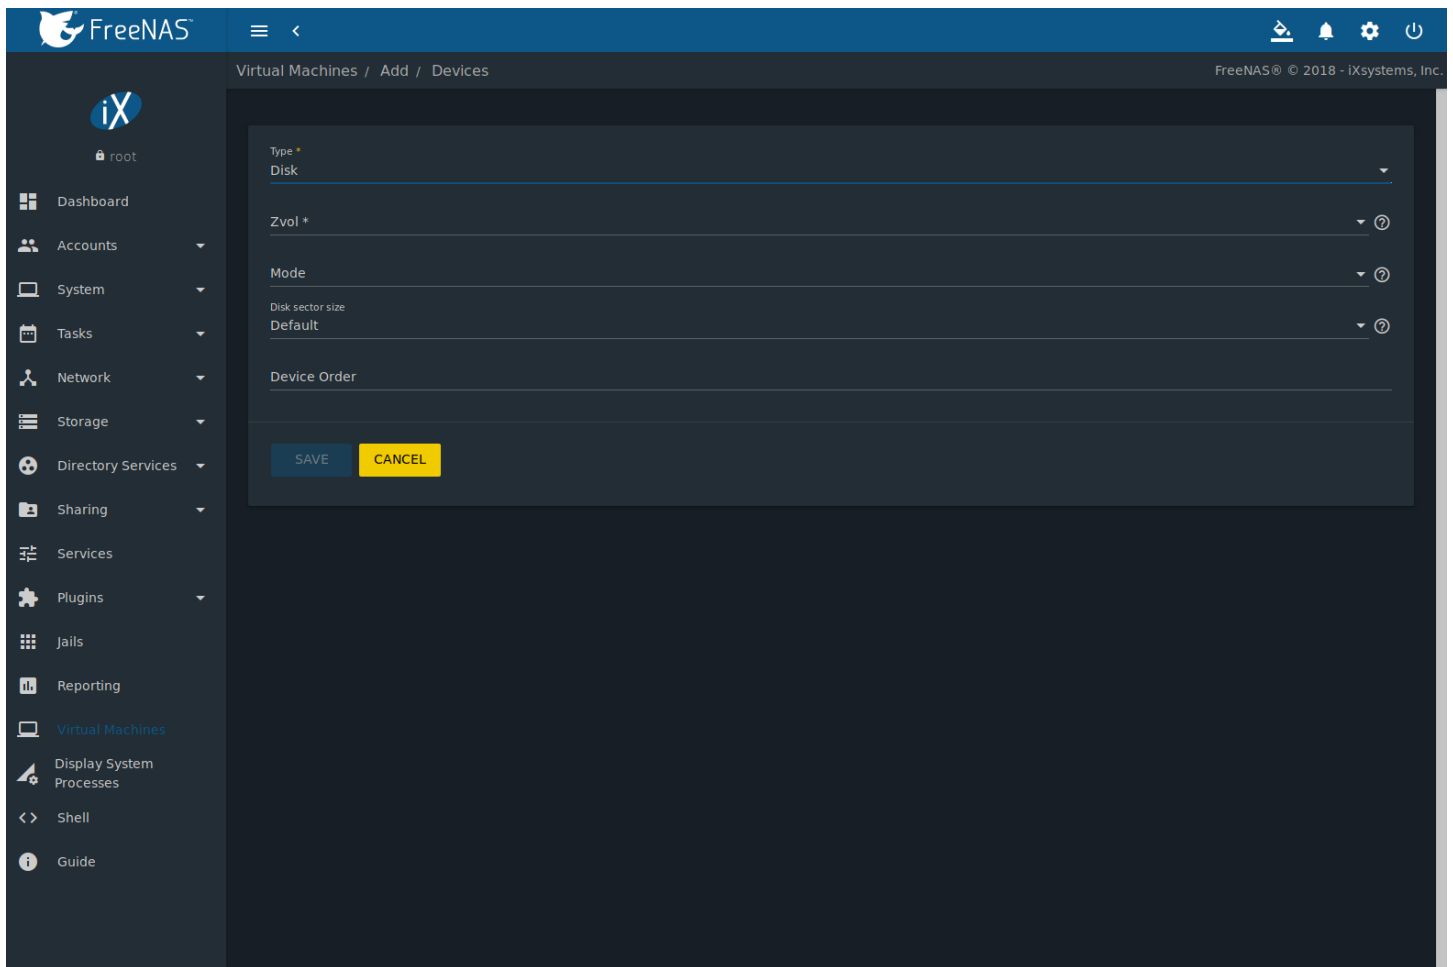


Fig. 16.6: Disk Device

Open the drop-down menu to select a created *Zvol*, then set the disk *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 in *Disk sector size*. The default of 0 uses an autotune script to determine the best sector size for the zvol.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.

16.2.4 Raw Files

Raw Files are similar to *Zvol* (page 169) 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, click *Virtual Machines* → ⚙️ (Options) → *Devices*, click *ADD*, then set the *Type* to *Raw File*.

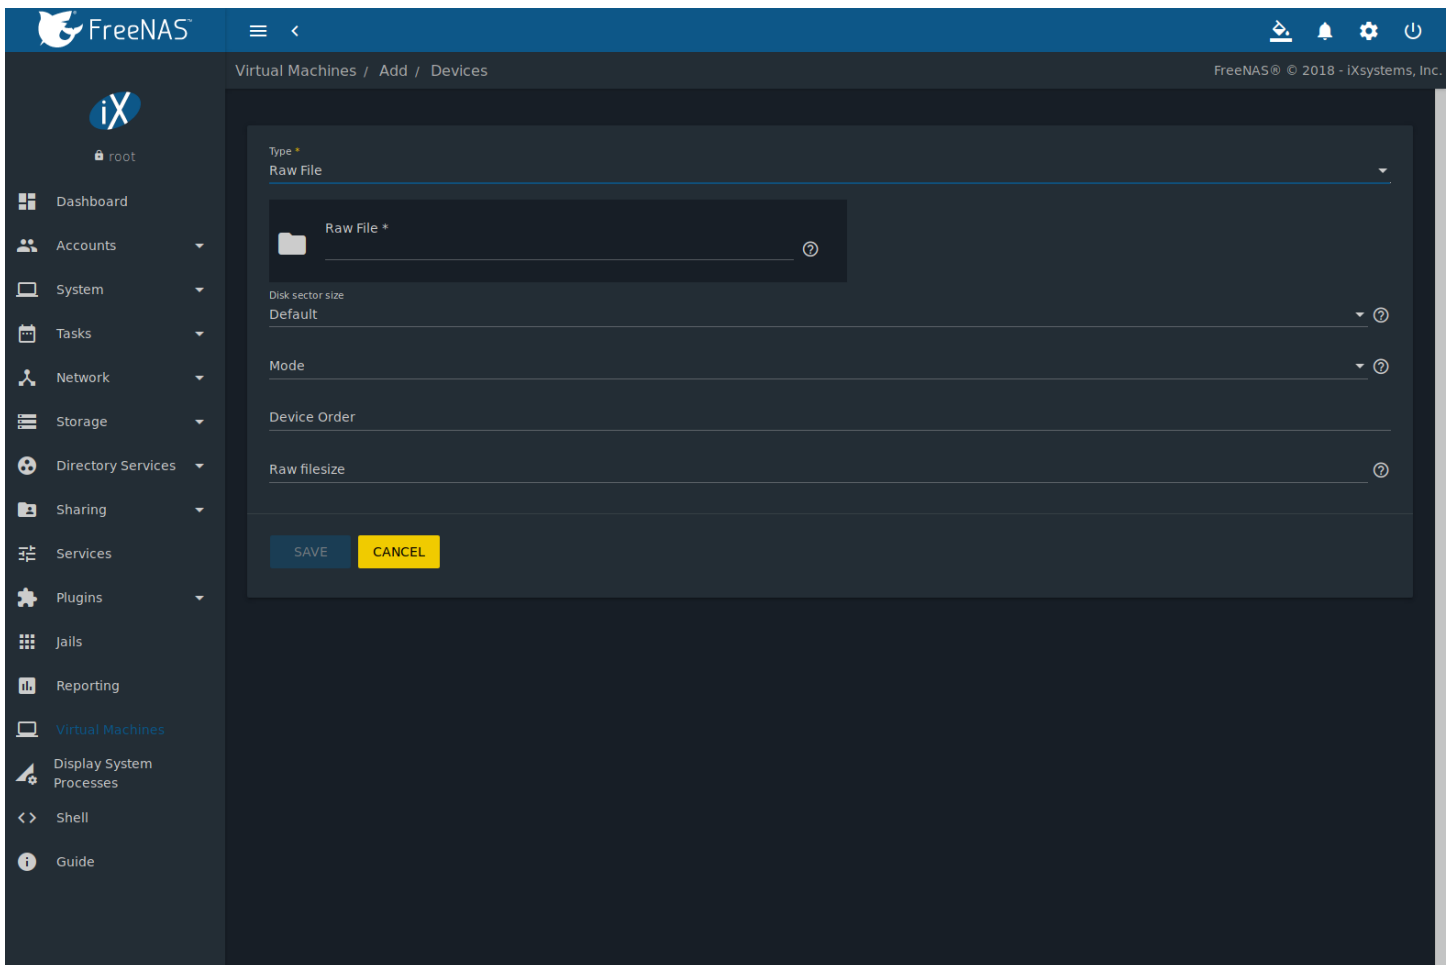


Fig. 16.7: Raw File Disk Device

Browse to select the image file. If a specific sector size is required, enter the number of bytes in *Disk sector size*. The default value of 0 uses an autotuner to find and set the best sector size for the file.

Setting the disk *Mode* to *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.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.

16.2.5 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) (https://en.wikipedia.org/wiki/Virtual_Network_Computing) client can connect to the VM to provide screen output and keyboard and mouse input.

Note: Each VM can only have a single VNC device. An existing VNC interface can be changed by clicking ⓘ (Options) and *Edit*.

Note: *Docker Hosts* (page 319) are not compatible with VNC connections and cannot have a VNC interface.

Note: Using a non-US keyboard via VNC is not yet supported. As a workaround, select the US keymap on the system running

the VNC client, then configure the operating system running in the VM to use a keymap that matches the physical keyboard. This will enable passthrough of all keys regardless of the keyboard layout.

Figure 16.8 shows the fields that appear after going to *Virtual Machines* → ⓘ (Options) → *Devices*, and clicking ⓘ (Options) → *Edit* for VNC.

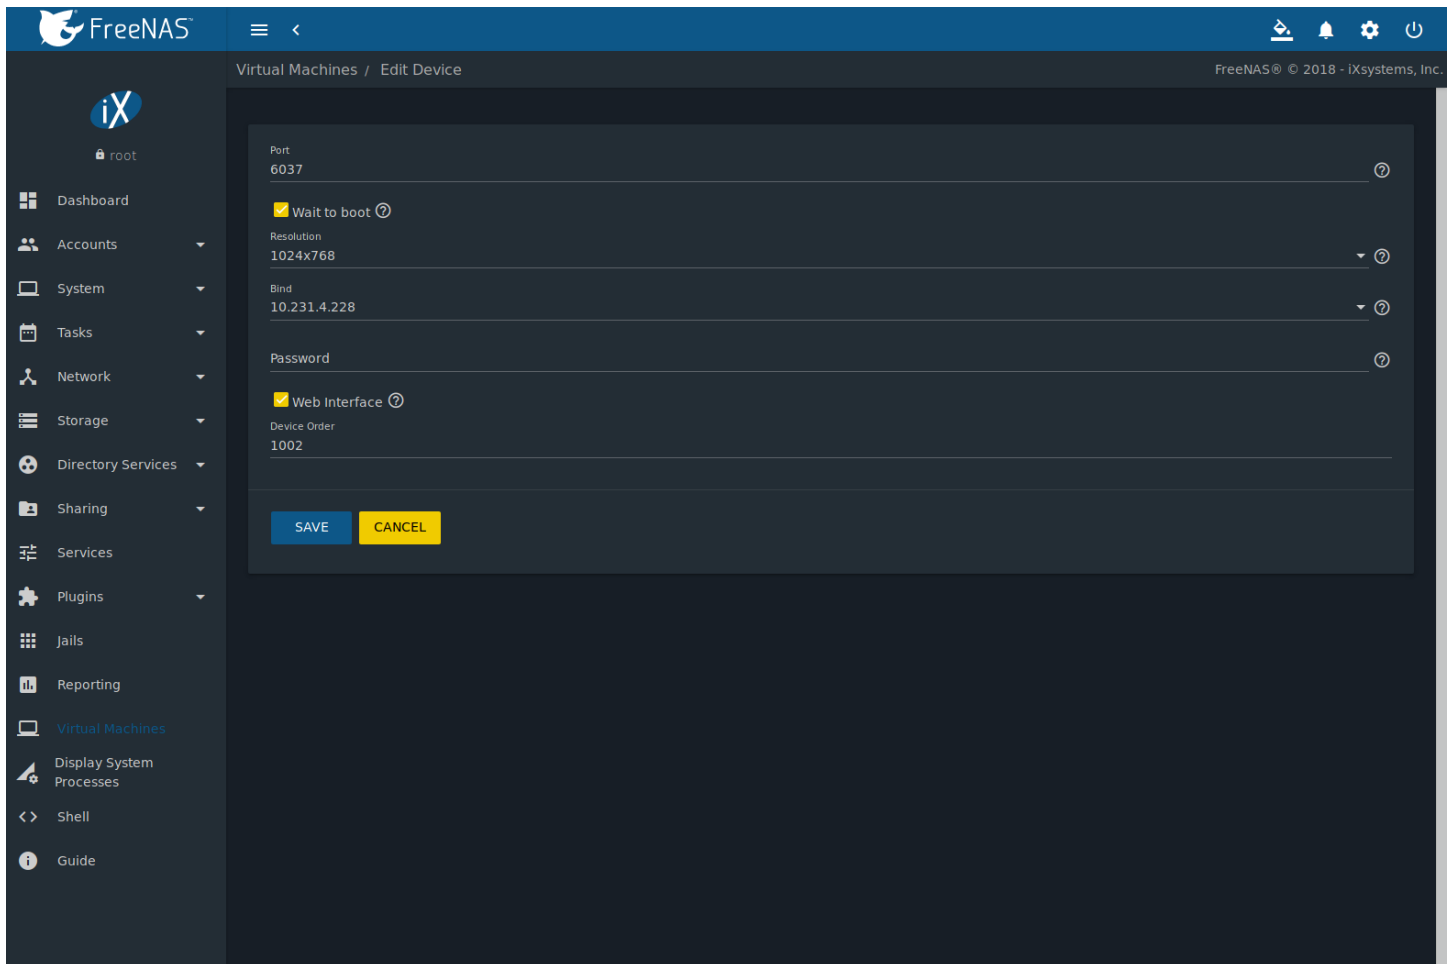


Fig. 16.8: VNC Device

The *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.

Set *Wait to boot* to force the VNC client to wait until the VM has booted before attempting the connection.

Use the *Resolution* drop-down menu to choose the default screen resolution used by the VNC session.

Select the IP address for VNC to listen on with the *Bind* drop-down menu.

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, set *Web Interface*.

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, Protocol/Version*, then select 4.1 from the drop-down menu.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.

16.2.6 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 325).

Tip: The `nmdm` (<https://www.freebsd.org/cgi/man.cgi?query=nmdm>) device is dynamically created. The actual `nmdm` name can differ on each system.

To connect to the first VM:

```
cu -l /dev/nmdm1B -s 9600
```

See [cu\(1\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=cu>) for more information on using `cu`.

16.3 Running VMs

Go to *Virtual Machines* to see a card for each installed VM. There are different view options at the top of the screen for this page:

- *Cards*: Default view. Shows a large “card” for each VM with all options for that VM contained on that card.
- *Slim*: Similar to the default, but reduces card size to fit more cards on the screen.
- *Table*: Removes cards entirely to show all VMs in a space efficient table. This view also has a *Virtual Machines Summary* which displays VM memory usage.

The default *Cards* view is described and shown in this section, but VM options are the same between each view. Each VM card has a *CONNECT* button at the bottom and will flip to show more options after clicking ⓘ (Options) and *Edit*.

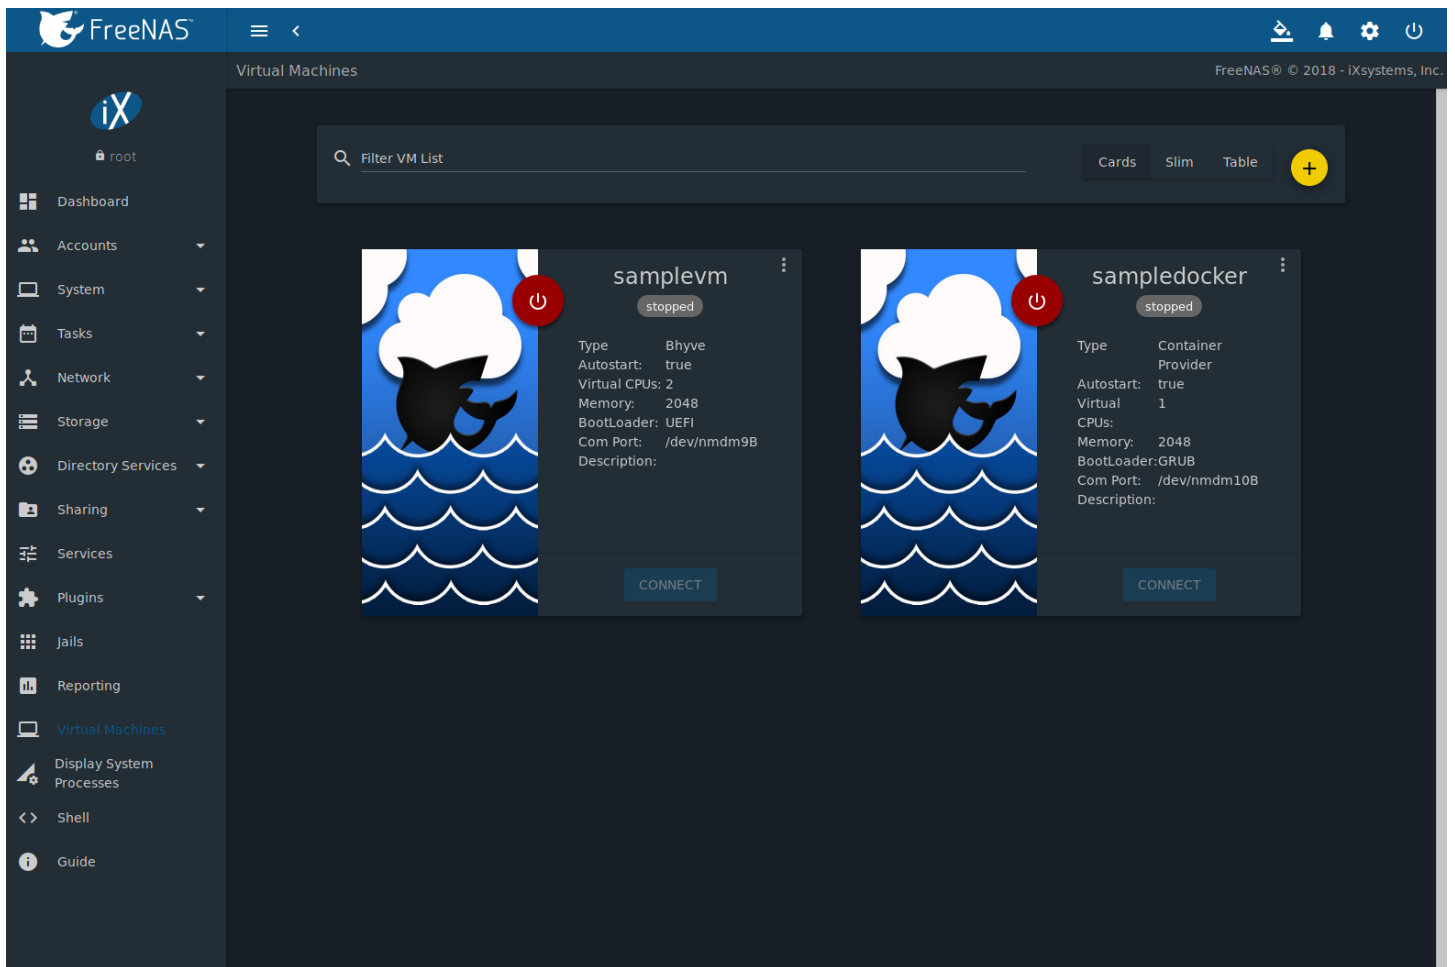




Fig. 16.9: VM Cards

The name, description, running state, com port (if present), and other configuration values are shown on the card. Click  (Options) for the *Start*, *Stop*, *Power Off*, *Restart*, *Edit*, *Devices*, and *Delete* options.


Some buttons are available for all VMs:

- *Delete* **removes the VM** (page 319).
- *Devices* is used to add, remove, and edit the VM devices.

When a VM is not running, these buttons are available:

-  (Power) (Red) starts the VM. A confirmation dialog appears and offers the option to *Overcommit Memory*. Memory overcommitment allows multiple VMs to be launched when there is not enough free memory for configured RAM of all VMs. Use with caution. The web interface displays a substantive error message when a VM fails to start.
- *Edit* changes VM settings, and includes an option to *Clone* an existing VM. Cloning makes a copy of the VM. The new clone has `_cloneN` added to the name, where *N* is the number of the clone. FreeNAS® will display a substantive error message if cloning a VM fails.

When a VM is already running, these buttons are available:

-  (Power) (Green) shuts down the VM.
- *Stop* also shuts down the VM.
- *Power off* immediately halts the VM. This is equivalent to disconnecting the power to a physical computer.
- *Restart* shuts down and immediately starts the VM.

- *CONNECT* displays options to connect to the VM. This can include *VNC* or a *Serial* shell. For *VNC*, the VM must have a VNC device with *Web Interface* enabled.

Popups ask to confirm the choice when starting, shutting down, powering off, or deleting a VM.

16.4 Deleting VMs

A VM is deleted by clicking  (Options) on the desired VM card, then *Delete*. A dialog prompts for confirmation.

Tip: *Zvols* (page 169) used in *disk devices* (page 313) and image files used in *raw file* (page 314) 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.

16.5 Docker Hosts

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 (<https://rancher.com/>) is a web interface tool for managing Docker containers.

FreeNAS® runs the Rancher web interface within the Docker Host.

16.5.1 Docker Host Requirements

The system BIOS **must** have virtualization support enabled for a Docker Host to run properly after installation. On Intel systems this is typically an option called *VT-x*. AMD systems generally have an *SVM* option.

20 GiB of storage space is required for the Docker Host.

For setup, the *SSH* (page 263) service must be enabled.

The Docker Host requires 2 GiB of RAM while running.

16.5.2 Creating Docker Host

Figure 16.10 shows the Wizard that appears after going to *Virtual Machines*, clicking *ADD*, and selecting *Docker Host* as the *Virtual Machine (VM) Wizard* type.

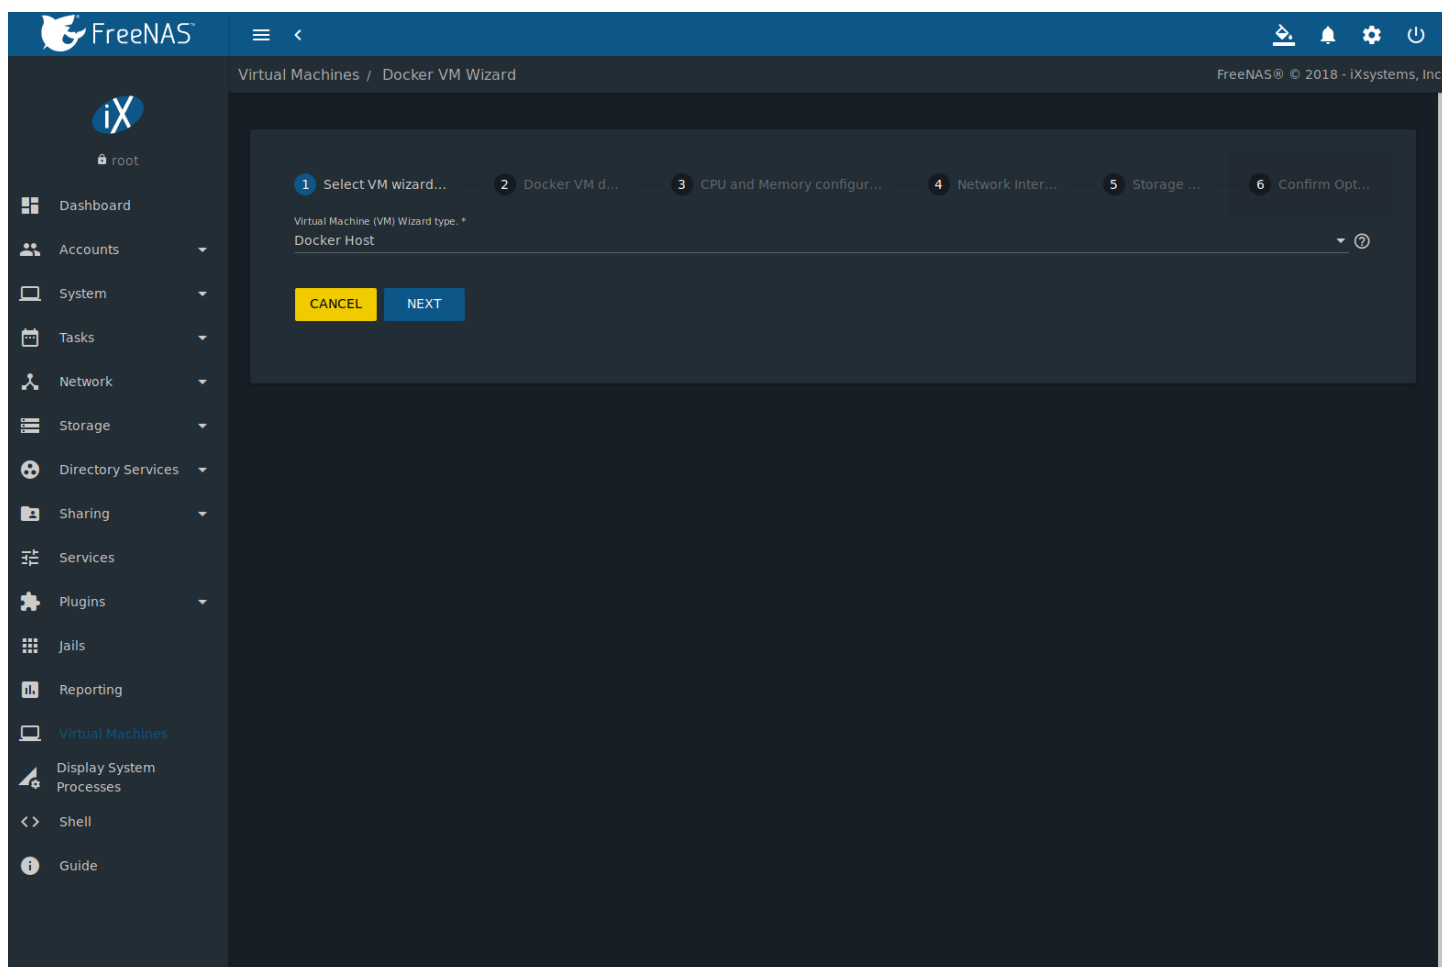


Fig. 16.10: Add Docker Host

Docker Host configuration options are described in [Table 16.2](#).

Table 16.2: Docker Host Options

Screen #	Setting	Value	Description
1	Virtual Machine (VM) Wizard type	drop-down menu	Choose to create either a standard <i>VM</i> or a <i>Docker Host</i> .
2	Name	string	Enter a descriptive name for the Docker Host.
2	Start on Boot	checkbox	Set to start this Docker Host when the FreeNAS® system boots.
3	Virtual CPUs	integer	Enter the number of virtual CPUs to allocate to the Docker Host. The maximum is 16 unless the host CPU also limits the maximum. The VM operating system can also have operational or licensing restrictions on the number of CPUs.
3	Memory Size (MiB)	integer	Allocate the amount of RAM in MiB for the Docker Host. A minimum 2048 MiB of RAM is required.
4	Adapter Type	drop-down menu	<i>Intel e82545 (e1000)</i> emulates the same Intel Ethernet card. This provides compatibility with most operating systems. <i>VirtIO</i> provides better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.
4	MAC Address	string	Enter the desired MAC address to override the auto-generated randomized MAC address.

Continued on next page

Table 16.2 – continued from previous page

Screen #	Setting	Value	Description
4	Attach NIC	drop-down menu	Select the physical interface to associate with the VM.
5	Raw filename	string	Enter a name for the new raw file.
5	Raw file size	integer	Allocate storage size in GiB for the new raw file.
5	Raw file location	browse button	Select a directory to store the new raw file.
5	Disk sector size	integer	Define the disk sector size in bytes. <i>Default</i> leaves the sector size un-set.

Choose the base options for the VM at each step of the wizard. *Virtual CPUs* is set to 1. *Memory Size* must be set to at least 2048 MiB.

The *Network Interface* step is automatically populated with system defaults. Customize these fields as necessary and press *NEXT* to continue.

The *Storage Files* section of the wizard contains options to create and store a raw file. Add a filename by typing an *.img* name in the *Raw filename* field. Enter a number of gigabytes for the *Raw file size*. Set the raw file location with the folder button or by typing a directory in the field.

The final screen of the Wizard displays the chosen options for the new Docker Host. Click *SUBMIT* to create the Host or *BACK* to change any settings. Click *CANCEL* at any time to return to the *Virtual Machines* page.

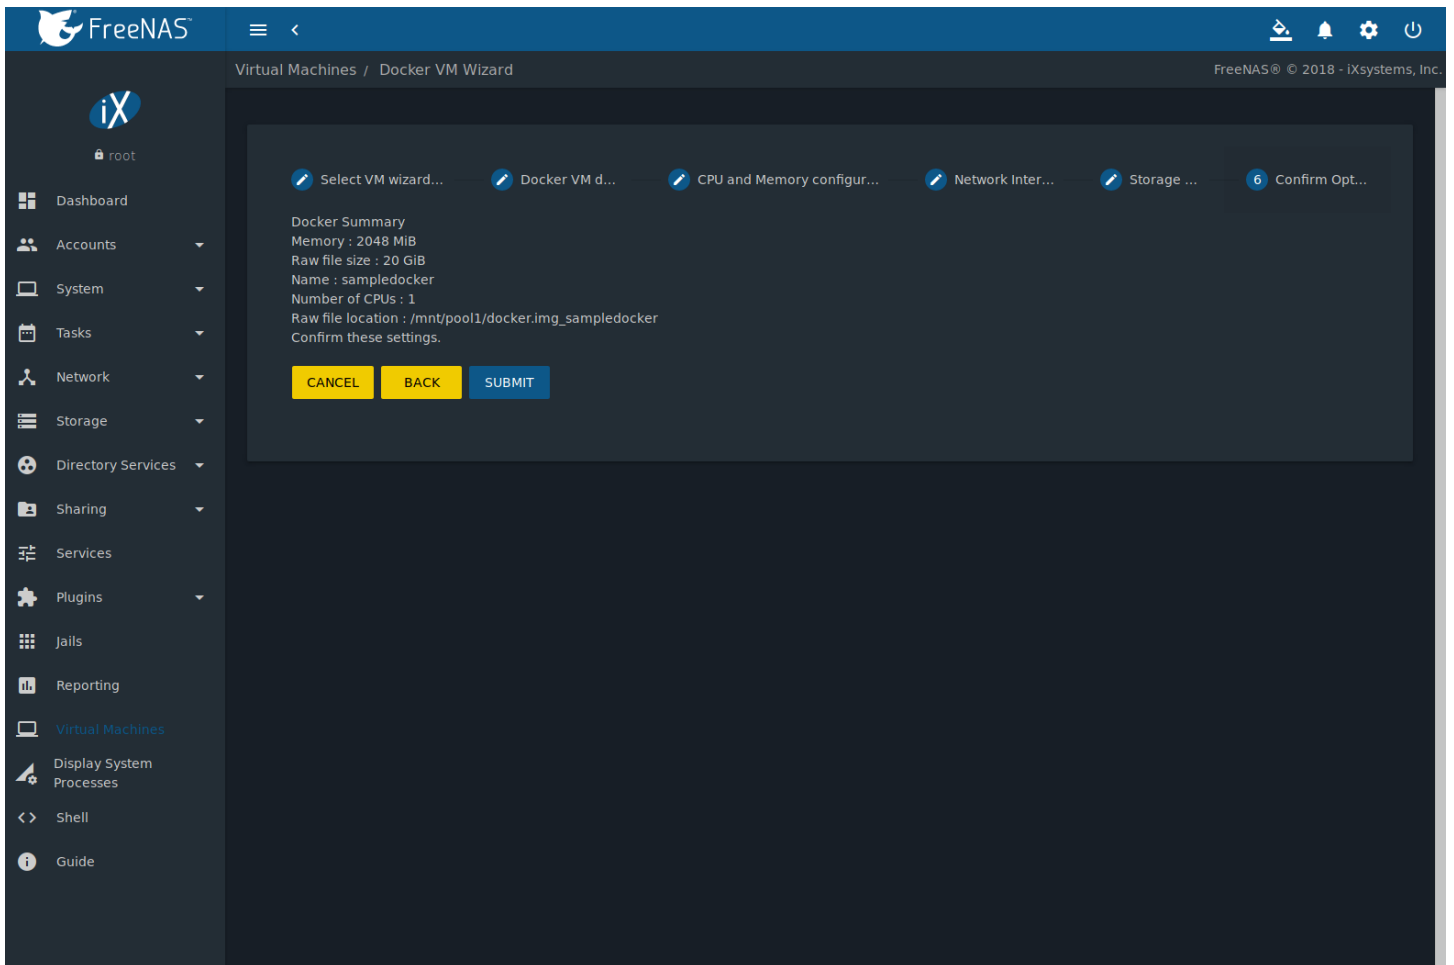





Fig. 16.11: Docker Host Configuration

Click  (Power), *CONNECT*, and *Serial* to log in to the Docker Host. Enter `rancher` for the user name and `docker` for the password.

The default password is changed in the *Devices* by stopping the Docker Host, clicking  (Options), and *Devices*. Click  (Options) and *Edit* for the *RAW* device and enter a new value in the *password* field. Passwords cannot contain a space.

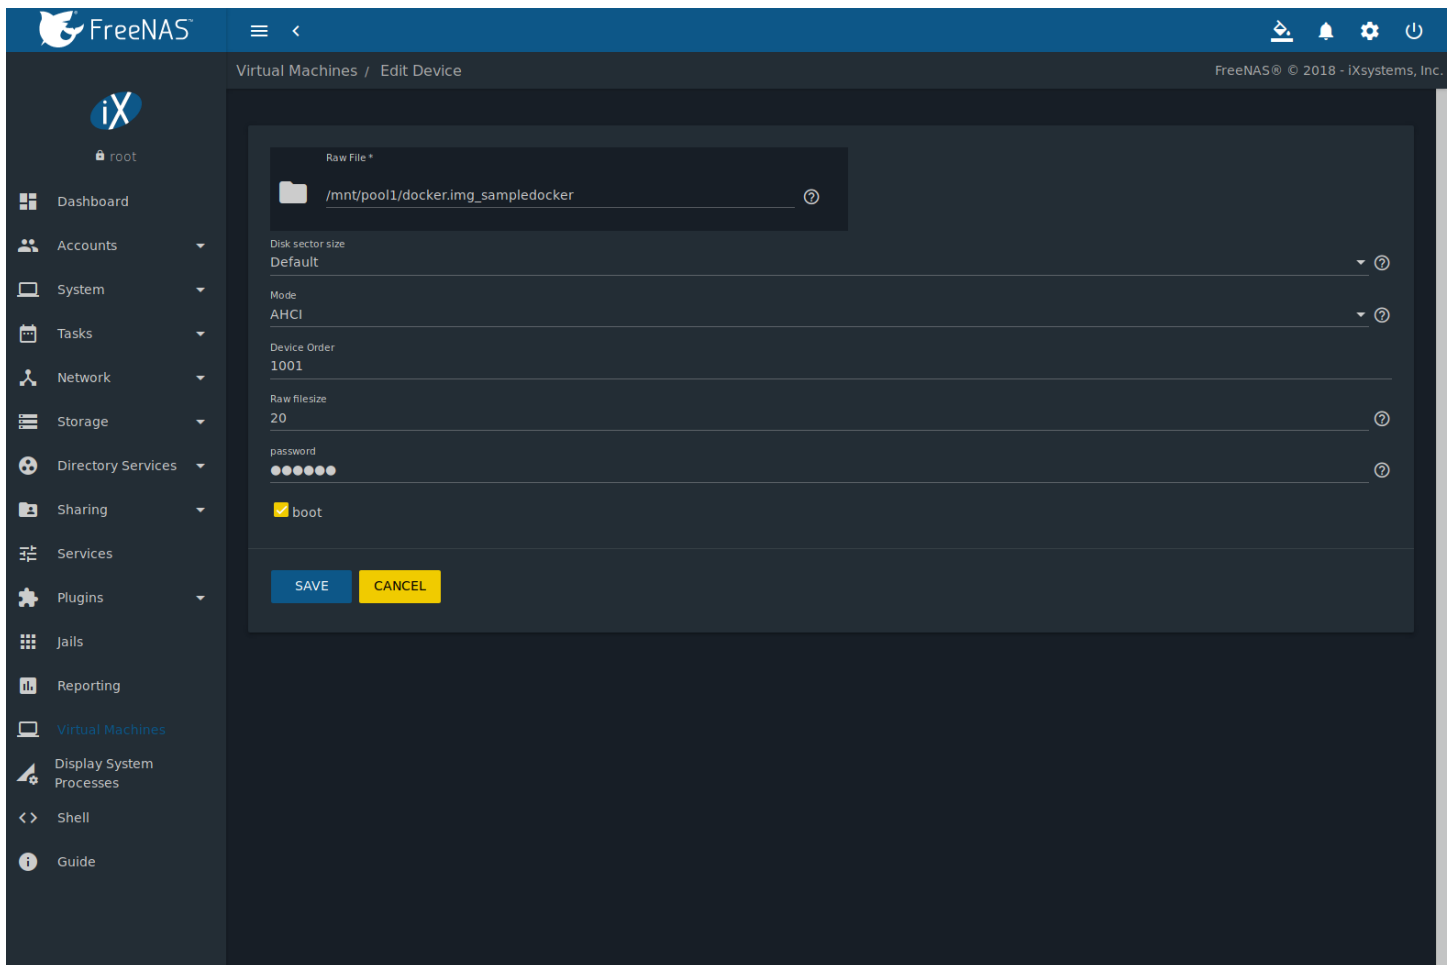


Fig. 16.12: Edit Rancher Password in Raw File Device

16.5.3 Start the Docker Host

Go to *Virtual Machines*, then click on the red  (Power) button of the Docker Host to start it.

Starting a Docker Host can take some time. Connecting to the Serial Shell is possible during the startup process to view the activity of the Docker Host. When a message about `RancherOS` starting appears and the shell stops posting new messages, press `Enter` to see the `ClientHost login:` text and continue to log in.

16.5.4 SSH into the Docker Host

It is possible to SSH into a running Docker Host. Go to the *Virtual Machines* page and find the card for the Docker Host. The card shows the *Com Port* for the Docker Host. In this example, `/dev/nmdm12B` is used.

Use an SSH client to connect to the FreeNAS® server. Remember this also requires the [SSH](#) (page 263) service to be running. Depending on the FreeNAS® system configuration, it might also require changes to the *SSH* service settings, like setting *Login as Root with Password*.

At the FreeNAS® console prompt, connect to the Rancher Host with `cu` (<https://www.freebsd.org/cgi/man.cgi?query=cu>), replacing `/dev/nmdm12B` with the value from the Docker Host *Com Port*:

```
cu -l /dev/nmdm12B -s 9600
```

If the terminal does not show a `rancher login: prompt`, press `Enter`. The Docker Host can take some time to start and display the login prompt.

16.5.5 Installing and Configuring the Rancher Server

Go to *Virtual Machines* and locate the card for the Docker Host. Start the Host and click *Connect* and *Serial* to open the Host Serial shell.

Continuing to set up the Rancher Host is done from the command line. Enter *rancher* as the username, press `Enter`, then enter either the default password *docker* or the custom password created by editing the raw file. Press `Enter` again. After logging in, a `[rancher@ClientHost ~]$` prompt is displayed.

Ensure Rancher has functional networking and can `ping` an outside website. Adjust the VM *Network Interface* (page 312) and reboot the VM if necessary.

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 web interface 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 web interface, click *Add a host*, ensure the radial *This site's address* button is set, and click *Save*. Follow the instructions that now display and run the `sudo docker run --rm --privileged -v` command in the Docker Host Serial shell. After the command runs a message displays `Launched Rancher Agent:.` Refresh or go to the *Hosts* page of the Rancher web interface to confirm the Docker Host displays in the web interface. Rancher is now configured and ready for use.

For more information on using RancherOS, see the RancherOS [Documentation](https://rancher.com/docs/os/v1.x/en/) (<https://rancher.com/docs/os/v1.x/en/>).

DISPLAY SYSTEM PROCESSES

Clicking *Display System Processes* opens a screen showing the output of `top(1)` (<https://www.freebsd.org/cgi/man.cgi?query=top>). An example is shown in Figure 17.1.

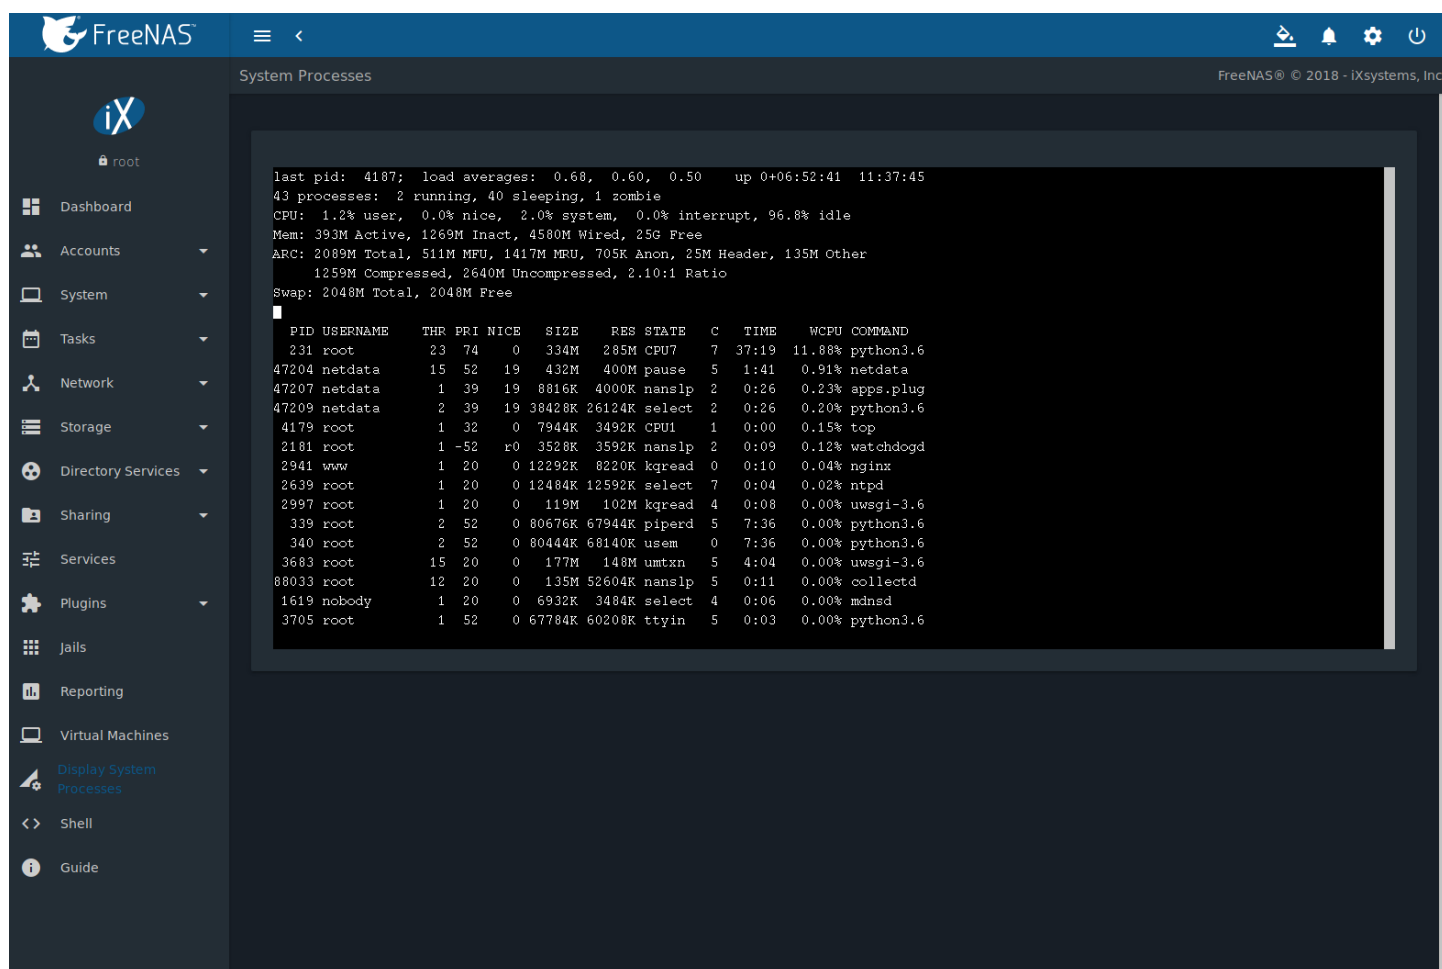


Fig. 17.1: System Processes Running on FreeNAS®

The display automatically refreshes itself. The display is read-only.

SHELL

Beginning with version 8.2.0, the FreeNAS® web interface 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 second entry from the bottom of the menu tree. In Figure 18.1, the link has been clicked and Shell is open.

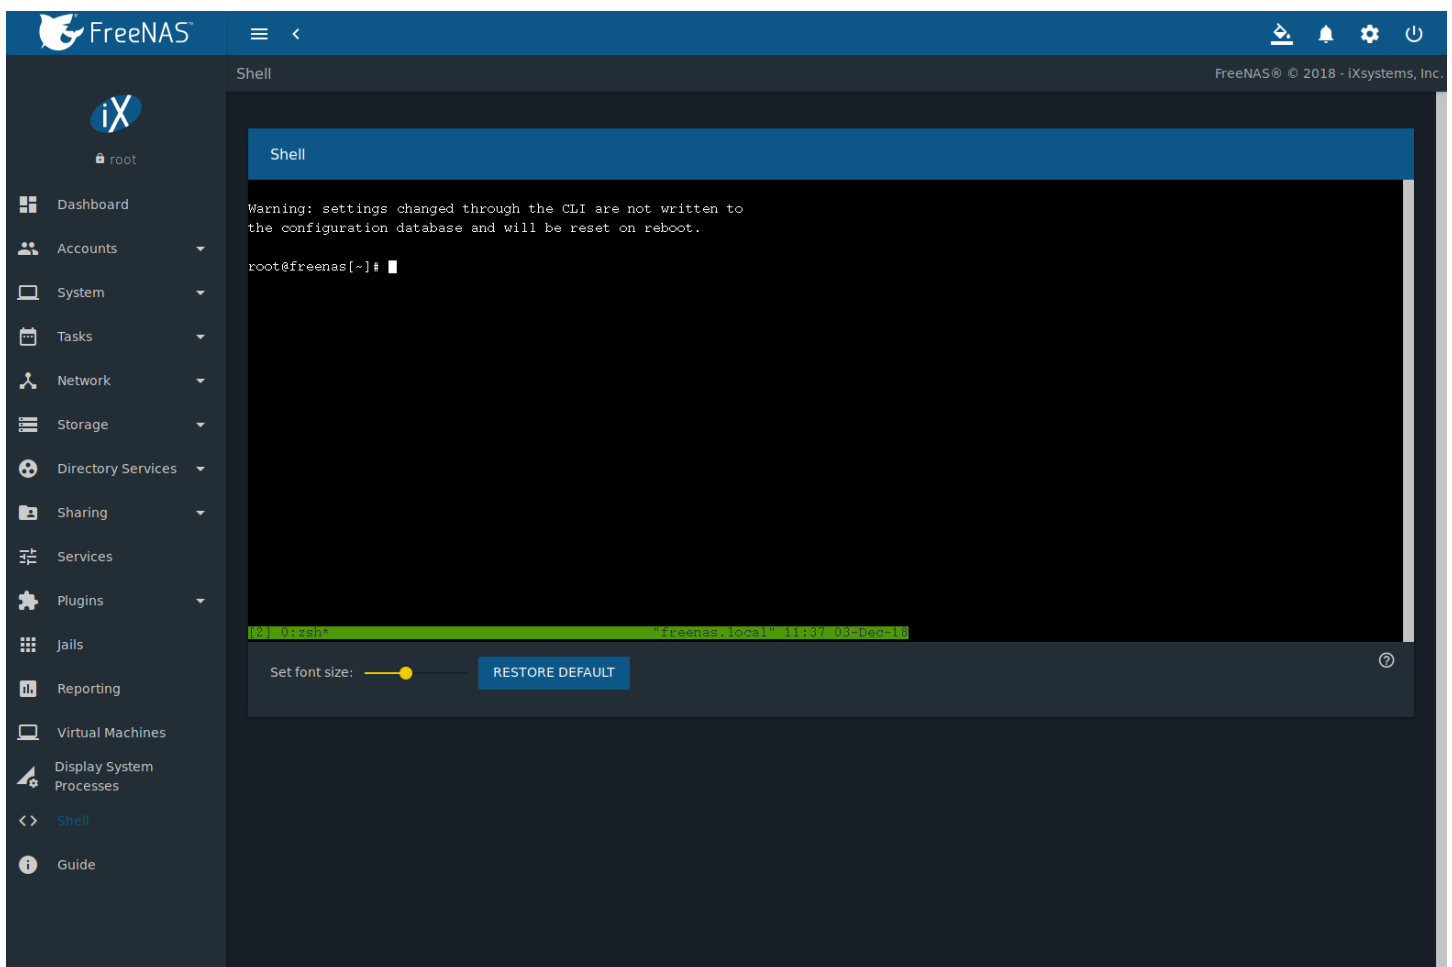


Fig. 18.1: Web Shell

The prompt indicates that the current user is *root*, the hostname is *freenas*, and the current working directory is *~* (*root* home directory).

Note: The default shell for a new install of FreeNAS® is *zsh*. FreeNAS® systems that are upgraded from an earlier version will continue to use *osh* as the default shell. The default shell can be changed by clicking *Accounts* → *Users*, clicking the edit icon.

(Options) of the desired user, then selecting *Edit*. Choose the desired shell from the *Shell* drop-down.

The *Set font size* slider adjusts the size of text displayed in the Shell.


Shell provides a history of commands used. Use the arrow keys to see previously entered commands and press `Enter` to repeat the command. The keys `Home`, `End`, and `Delete` are also supported in the shell. The shell also provides tab completion. Type a few letters and press `tab` to complete a command name or filename in the current directory. Type `exit` to leave the session.

Web interface menus cannot be accessed while using the *Shell*. To access a prompt while using the web interface menus, use *tmux* (page 351) as it supports multiple shell sessions and the detachment and reattachment of sessions.

Note: Not all Shell features render correctly in Chrome. Firefox is the recommended browser for using Shell.

Most FreeBSD command line utilities are available in the *Shell*. Additional troubleshooting utilities that are provided by FreeNAS® are described in *Command Line Utilities* (page 338).


LOG OUT, RESTART, OR SHUT DOWN

The  (Power) button is used to log out of the web interface or restart or shut down the FreeNAS[®] system.

19.1 Log Out

To log out, click  (Power), then *Log Out*. After logging out, the login screen is displayed.

19.2 Restart

To restart the system, click  (Power), then *Restart*. A confirmation screen asks for verification of the restart. [Figure 19.1](#). Click *Confirm* to confirm the action, then click *RESTART* to restart the system.

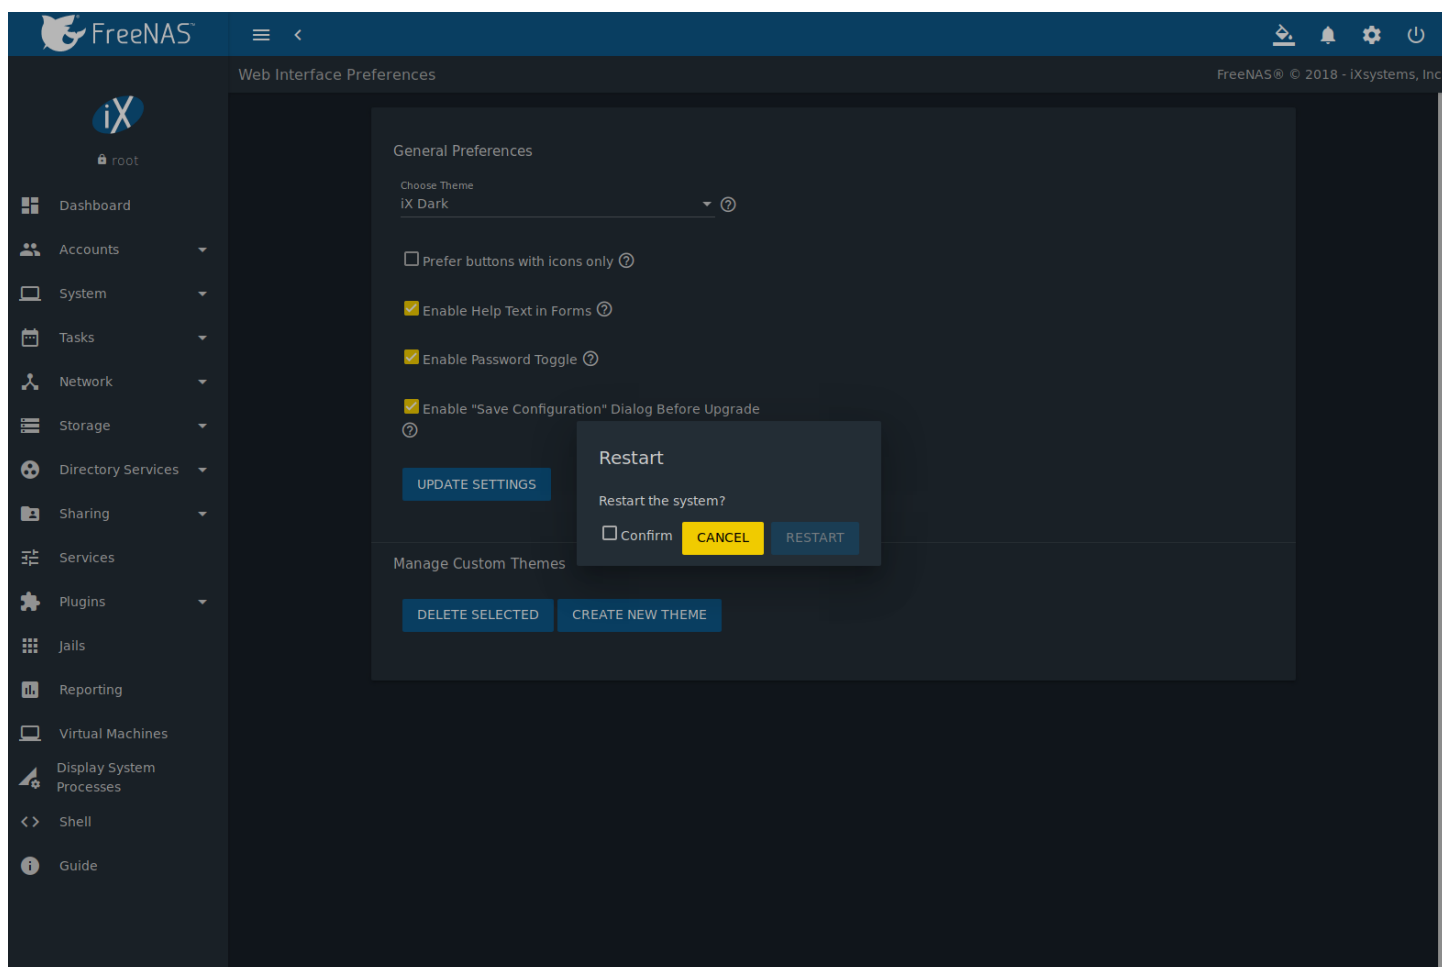



Fig. 19.1: Restart Warning Message

An additional warning message appears when a restart is attempted when a scrub or resilver is in progress. When that warning appears, the recommended steps are to *CANCEL* the restart request and to periodically run `zpool status` from *Shell* (page 325) until it shows that the scrub or verify has completed. Then the restart request can be entered again.

To complete the restart request, click the *Confirm* checkbox and then the *RESTART* button. Restarting the system disconnects all clients, including the web administration interface. Wait a few minutes for the system to boot, then use the back button in the browser to return to the IP address of the FreeNAS® system. The login screen appears after a successful reboot. If the login screen does not appear, using a monitor and keyboard to physically access the FreeNAS® system is required to determine the issue preventing the system from resuming normal operation.

19.3 Shut Down

Click  (Power) and *Shut Down* to shut down the system. The warning message shown in Figure 19.2 is displayed.

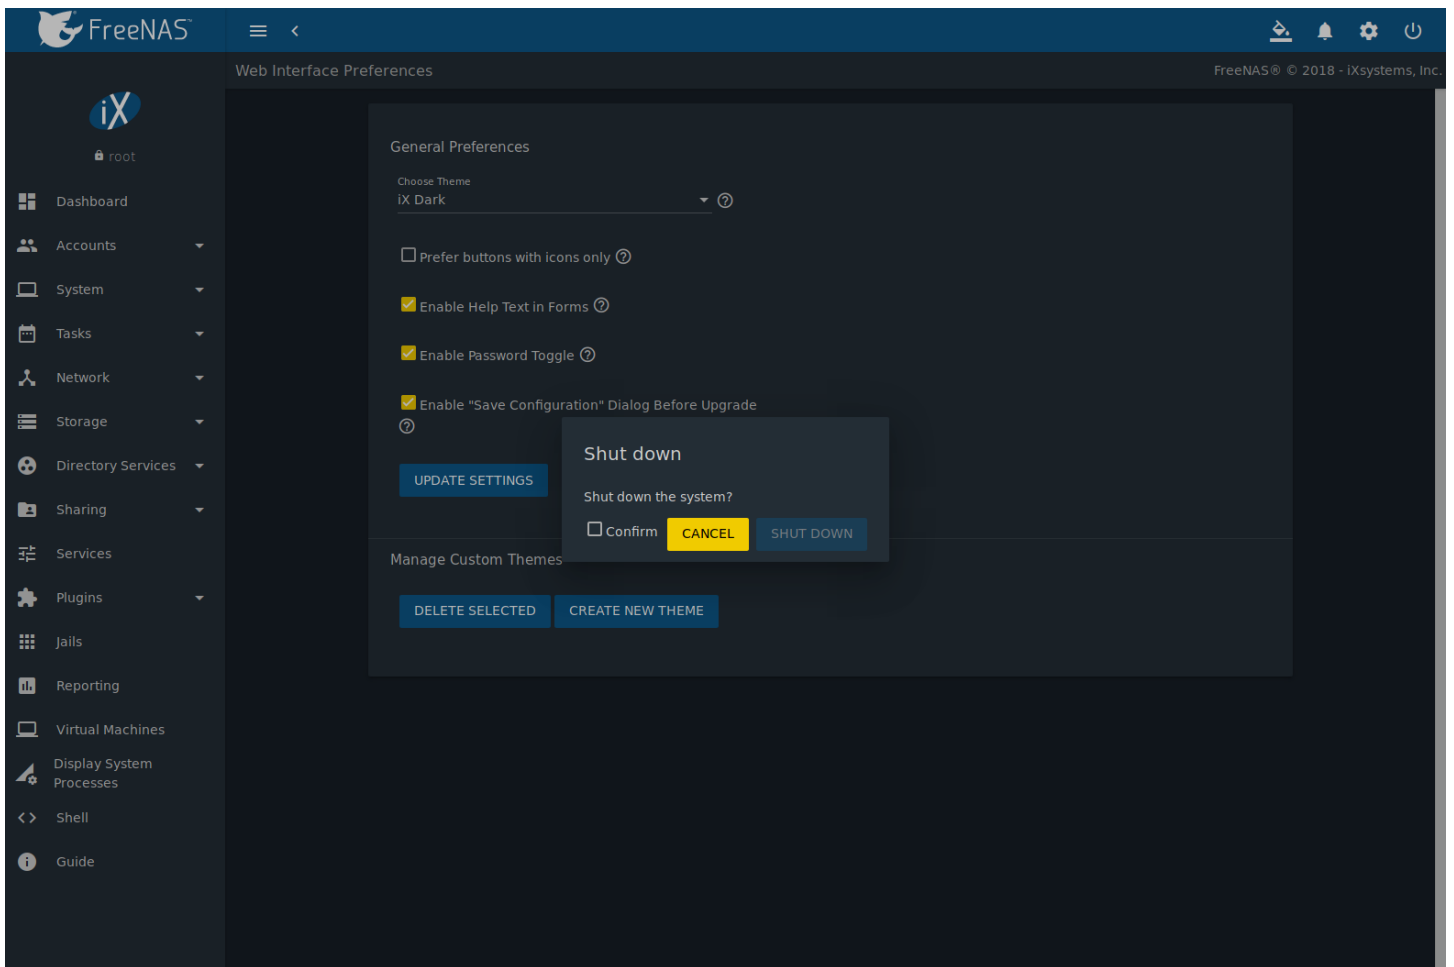


Fig. 19.2: Shut Down Warning Message

Click *Confirm* and then *SHUT DOWN* to shut down the system. Shutting down the system disconnects all clients, including the web interface. Physical access to the FreeNAS® system is required to turn it back on.

ALERT

The FreeNAS® alert system provides a visual warning of any conditions that require administrative attention. The *Alert* icon in the upper right corner has a notification badge that displays the total number of unread alerts. In the example alert shown in Figure 20.1, the system is warning that a pool is degraded.

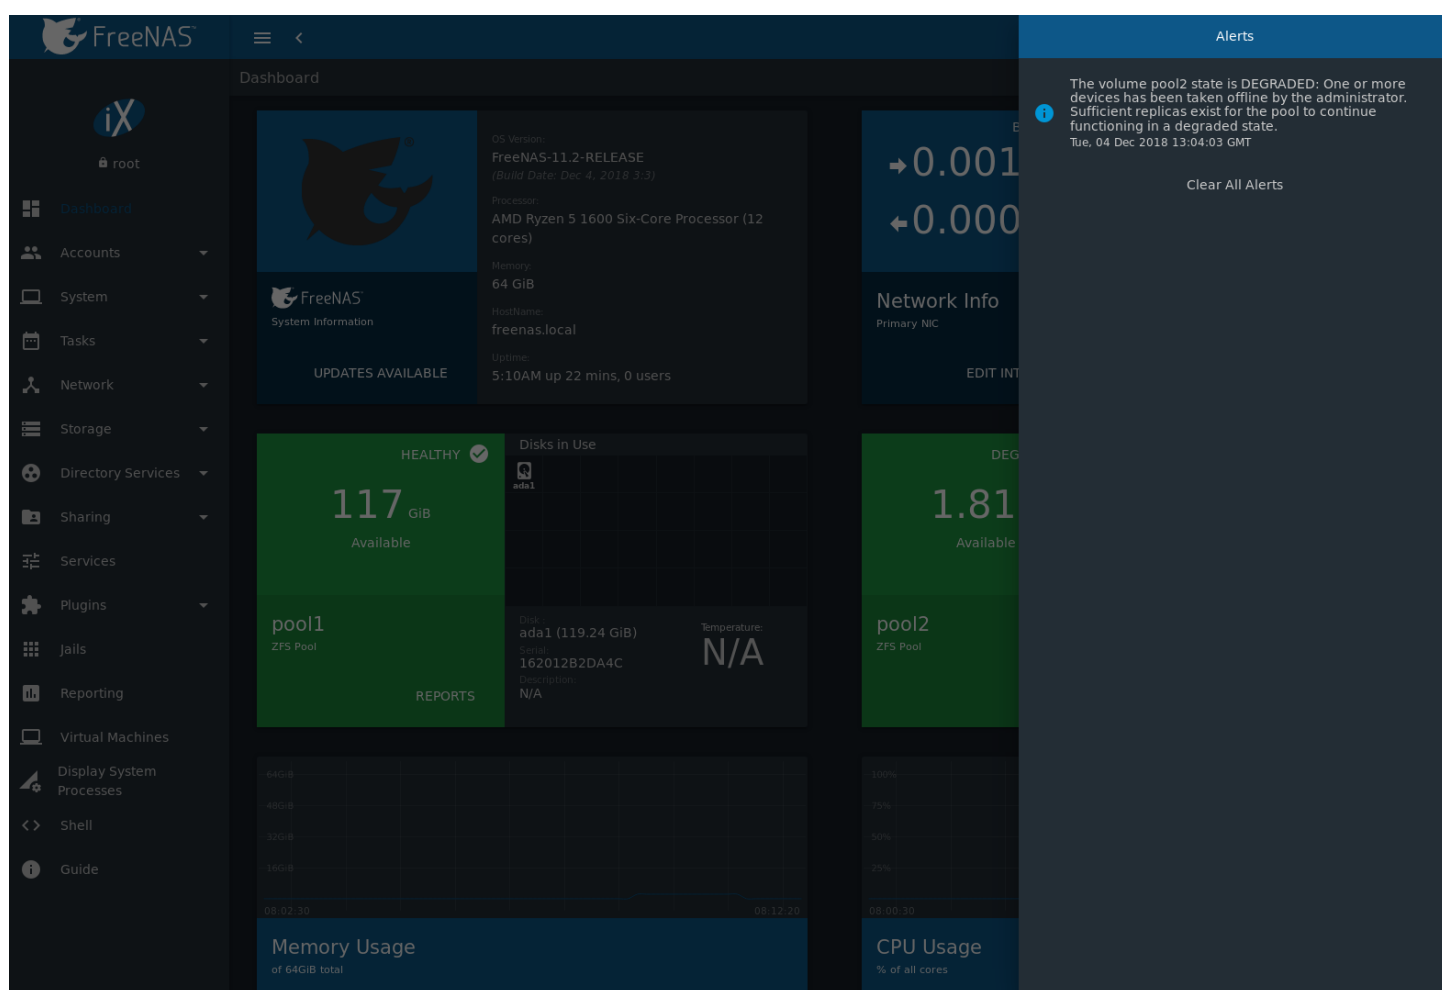


Fig. 20.1: Example Alert Message

Table 20.1 shows the icons that indicate notification, warning, and critical alerts. Critical messages are also emailed to the root account.

Table 20.1: FreeNAS® Alert Icons

Alert Level	Icon
Notification	🔔
Warning	⚠️
Critical	🔴

Close an alert message by hovering over it until *Click to Dismiss* appears. There is also an option to *CLEAR ALL ALERTS*. Close all messages to remove any notification badge from the alerts icon.

Behind the scenes, an alert daemon checks for various alert conditions, such as pool and disk status, and writes the current conditions to the system RAM. These messages are flushed to the SQLite database periodically and then published to the user interface.

Current alerts are viewed from the Shell option of the Console Setup Menu (Figure 3.1) or the Web Shell (Figure 18.1) by running `midclt call alert.list`.

Notifications for specific alerts are adjusted in the *Alert Settings* menu. An alert message can be set to publish *IMMEDIATELY*, *HOURLY*, *DAILY*, or *NEVER*.

Some of the conditions that trigger an alert include:

- used space on a pool, 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 adjusted in *Alert Settings* if a pool upgrade is not desired at present
- a new update is available
- ZFS pool status changes from *HEALTHY*
- a S.M.A.R.T. error occurs
- 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
- the NTP server cannot be contacted
- a replication task fails
- a VMware login or a *VMware-Snapshots* (page 175) 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\(8\)](https://www.freebsd.org/cgi/man.cgi?query=mfiutil) (<https://www.freebsd.org/cgi/man.cgi?query=mfiutil>) is included for managing these devices
- a scrub is paused

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 an issue occurs while using FreeNAS®, it can be helpful to spend a few minutes searching the Internet for the word *FreeNAS* with some keywords that describe the error message or function that is being implemented.

The section discusses resources available to FreeNAS® users:

- *Website and Social Media* (page 332)
- *Forums* (page 332)
- *IRC* (page 333)
- *Videos* (page 333)
- *Professional Support* (page 333)

21.1 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](https://www.linkedin.com/groups/3903140/profile) (<https://www.linkedin.com/groups/3903140/profile>)
- [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://mobile.twitter.com/freenas) (<https://mobile.twitter.com/freenas>)

21.2 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 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**.

- [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/) (<https://forums.freenas.org/index.php?forums/dutch-nederlands.35/>)
- [French - Francais](https://forums.freenas.org/index.php?forums/french-francais.29/) (<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 forums, 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. 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 a questions that will get a response.

21.3 IRC

To ask a question in real time, use the `#freenas` channel on IRC [Freenode](http://freenode.net/) (<http://freenode.net/>). Depending on the time of day and the 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](https://en.wikipedia.org/wiki/Comparison_of_Internet_Relay_Chat_clients) (https://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.
- Do not ask a question and leave. Users who know the answer cannot help you if you disappear.
- If no one answers, the question may be difficult to answer or it has been asked before. Research other resources while waiting for the question to be answered.
- Do not post error messages in the channel. Instead, use a pasting service such as [pastebin](https://pastebin.com/) (<https://pastebin.com/>) and paste the resulting URL into the IRC discussion.

21.4 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 \(https://www.youtube.com/watch?v=RxxggaE935PM\)](https://www.youtube.com/watch?v=RxxggaE935PM)
- [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 \(https://www.youtube.com/watch?v=R3f-Sr6v-c4\)](https://www.youtube.com/watch?v=R3f-Sr6v-c4)

21.5 Professional Support

In addition to free community resources, support might be available in your area through third-party consultants. Submit a support inquiry using the form at <https://www.ixsystems.com/freenas-commercial-support/>.

CONTRIBUTING TO FREENAS®

FreeNAS® is an open source community, relying on the input and expertise of users to grow and improve. When users take time to assist the community, their contributions benefit everyone.

This section describes how to participate and contribute to FreeNAS®. It is by no means an exhaustive list. If you have an idea that will benefit the community, bring it up on one of the resources mentioned in [Support Resources](#) (page 332).

This section demonstrates how to:

- [Help with Translation](#) (page 334)

22.1 Translation

FreeNAS® is developed and documented in English. Having complete translations of the user interface into other languages helps make FreeNAS® much more useful to communities around the world.

FreeNAS® uses .po files stored in the [webui GitHub repository](https://github.com/freenas/webui/tree/master/src/assets/i18n) (https://github.com/freenas/webui/tree/master/src/assets/i18n) to manage the translation of text shown in the FreeNAS® graphical administrative interface. GitHub provides an easy to use web-based editor, making it possible for individuals to assist with translation or comment on existing translations.

To view translation files, open the `/src/assets/i18n` directory of the FreeNAS® [webui repository](https://github.com/freenas/webui/tree/master/src/assets/i18n) (https://github.com/freenas/webui/tree/master/src/assets/i18n), as shown in [Figure 22.1](#).

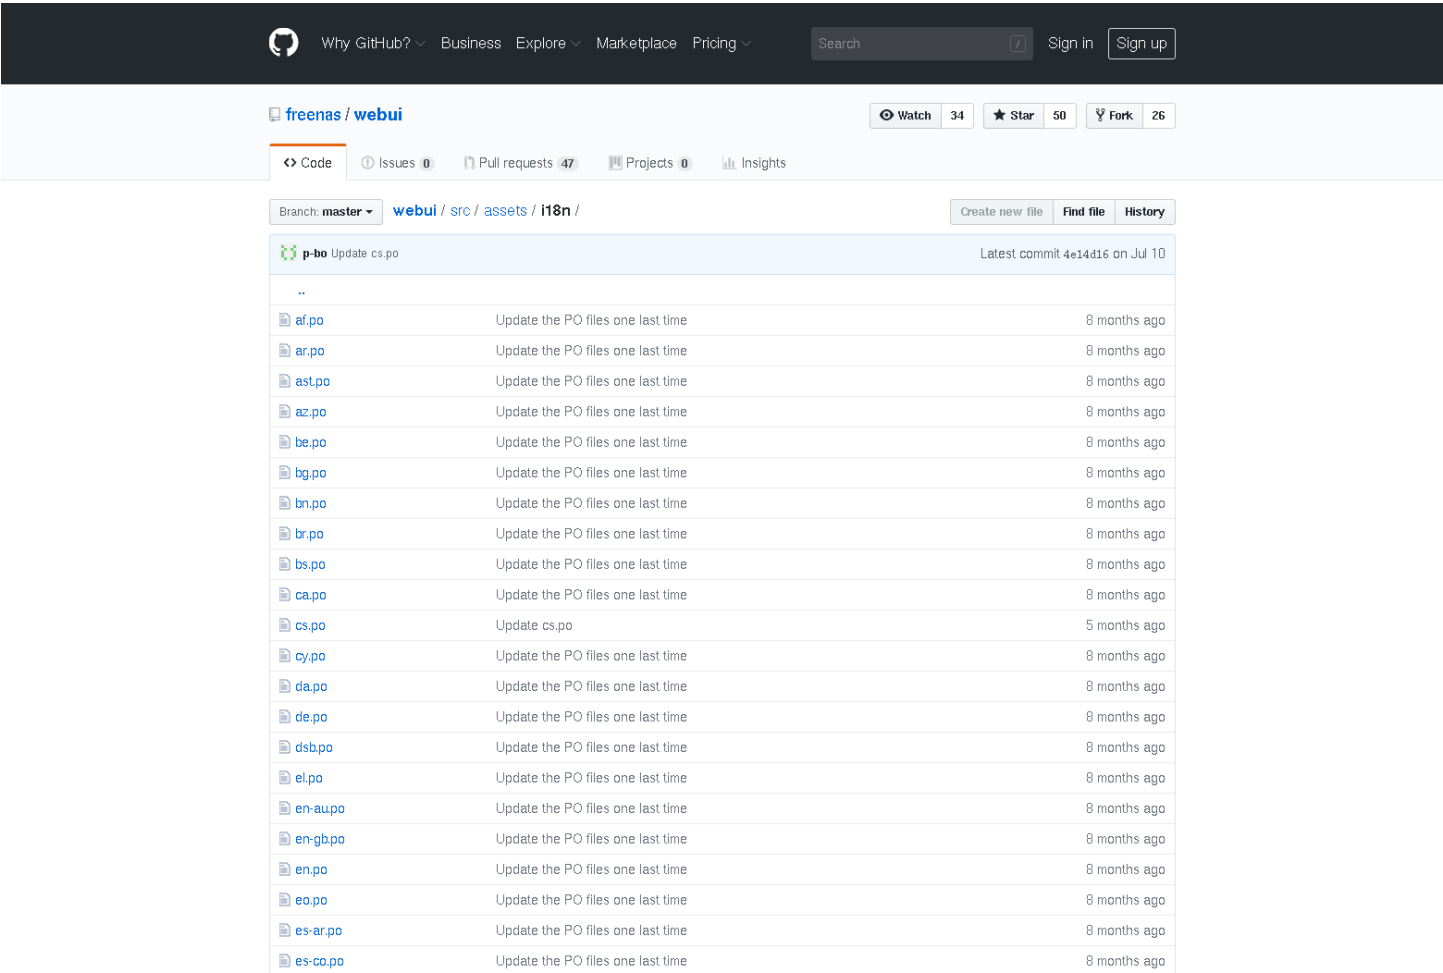


Fig. 22.1: FreeNAS® Translation Files

To assist with translating FreeNAS®, first create an account with [GitHub](https://github.com/) (<https://github.com/>).

There are two methods to contribute translations to the project:

- 1. Edit `po` files and submit pull requests through the GitHub website.
- OR
- 2. Copy the `freenas/webui` (<https://github.com/freenas/webui>) repository, make changes using a `po` editor, and submit these changes back “upstream” to the `freenas/webui` repository.

22.1.1 Translate with GitHub

Go to the `freenas/webui` (<https://github.com/freenas/webui>) repository. Select `src` → `assets` and click the `i18n` (<https://github.com/freenas/webui/tree/master/src/assets/i18n>) directory. Click on the desired language `po` file to begin translating.

Tip: Here is a list of [common language abbreviations](https://www.abbreviations.com/acronyms/LANGUAGES2L) (<https://www.abbreviations.com/acronyms/LANGUAGES2L>)

Click the *Pencil* icon in the upper right area to open the online file editor. [Figure 22.2](#) shows the page that appears:

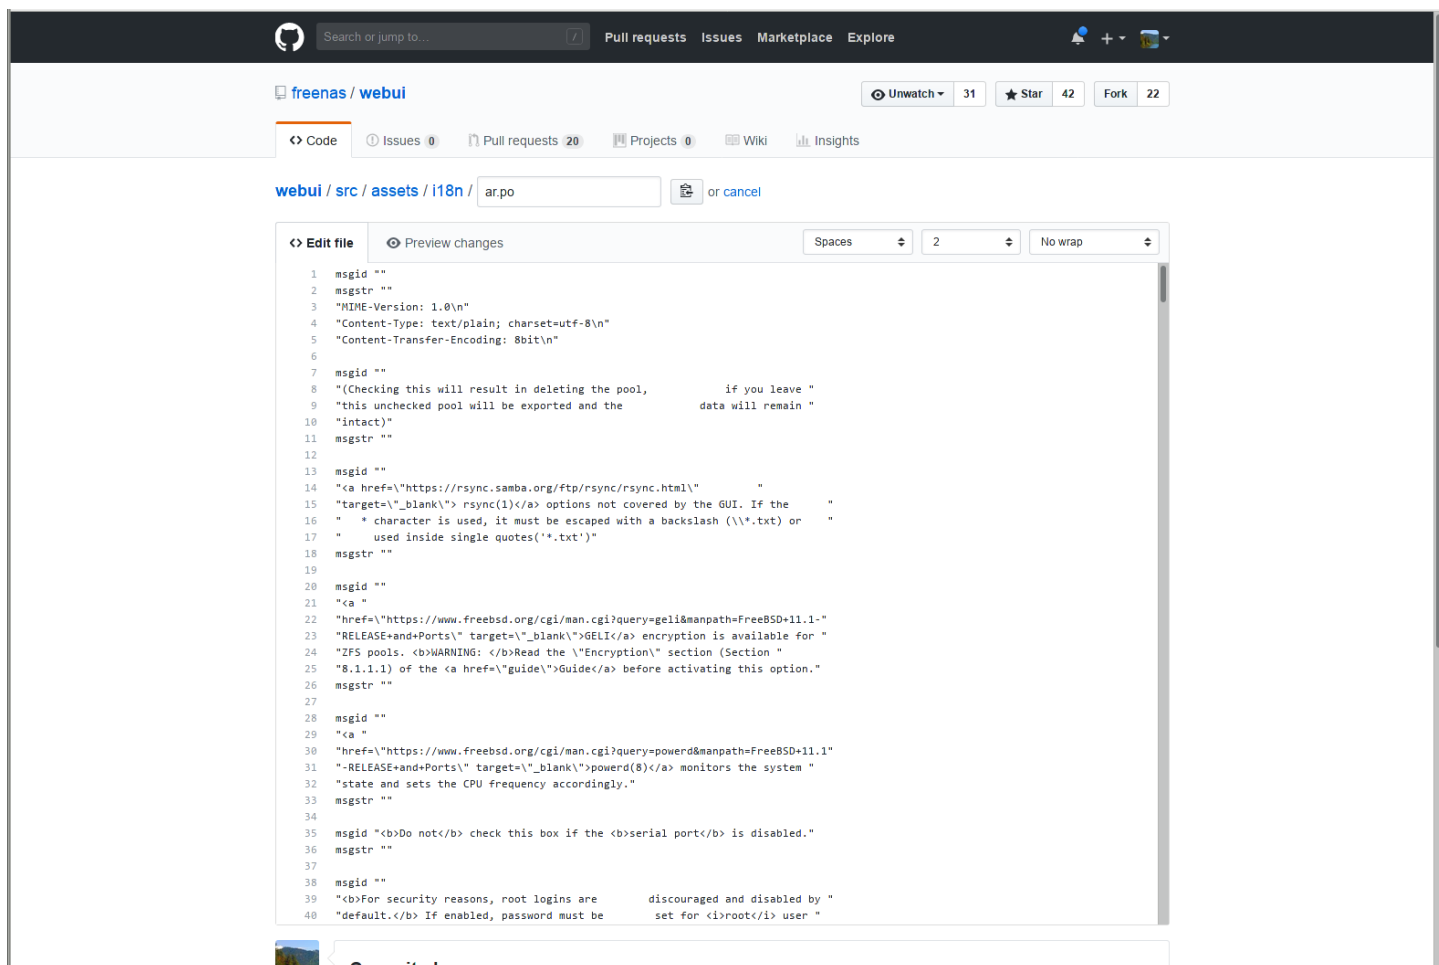


Fig. 22.2: GitHub Online Editor

There are numerous `msgid ""` and `msgstr ""` entries in the file. Read the `msgid` text and enter the translation between the `msgstr` quotes.

Scroll to the bottom of the page when finished entering translations. Enter a descriptive title and summary of changes for the edits and set *Create a new branch*. Click *Propose file change* to submit the translations to the FreeNAS® project.

22.1.2 Download and Translate Offline

Install Git (<https://git-scm.com/book/en/v2/Getting-Started-Installing-Git>). There are numerous examples in these instructions of using git, but full documentation for git is [available online](https://git-scm.com/doc) (<https://git-scm.com/doc>).

Open a Command Line Interface (CLI). Navigate to or create a suitable location to store the local copy of the `webui` repository. Download the repository with `git clone`:

```
% git clone https://github.com/freenas/webui.git
```

The download can take several minutes, depending on connection speed.

`cd` into the `webui` directory and create a new branch of the repository to store the translation changes:

```
% git checkout -b new_translations
```

Tip: Type `git status` at any time to see which branch of the repository is active.

Navigate to the `i18n` directory:

```
% cd src/assets/i18n/
```

Use a `po` editor to add translations to the desired language file. Any capable editor will work, but [poedit](https://poedit.net/) (<https://poedit.net/>) and [gtranslator](https://wiki.gnome.org/Apps/Gtranslator) (<https://wiki.gnome.org/Apps/Gtranslator>) are two common options.

Commit any file changes with `git commit`:

```
% git commit ar.po
```

Enter a descriptive message about the changes and save the commit.

When finished making commits to the branch, `git push` the branch to the online `freenas/webui` repository:

```
% git push origin new_translations
Username for `https://github.com`:
Password for `https://account@github.com`:
Counting objects: 6, done.
Delta compression using up to 4 threads.
Compressing objects: 100% (6/6), done.
Writing objects: 100% (6/6), 8.56 KiB | 4.28 MiB/s, done.
Total 6 (delta 5), reused 0 (delta 0)
remote: Resolving deltas: 100% (5/5), completed with 5 local objects.
To https://github.com/freenas/webui.git
* [new branch]      new_translations -> new_translations
```

Open a web browser and navigate to the [freenas/webui](https://github.com/freenas/webui) (<https://github.com/freenas/webui>) repository. GitHub automatically detects when a new branch is pushed to the repository and displays a message with an option to *Compare & pull request*. Click this, review the changes, and click *Create pull request*.

22.1.3 Translation Pull Requests

The FreeNAS® project automatically tests pull requests for compatibility. If there any issues with a pull request, either the automated system will update the request or a FreeNAS® team member will leave a message in the comment section of the request.

All assistance with translations helps to benefit the FreeNAS® community. Thank you!

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 338): used for measuring maximum TCP and UDP bandwidth performance
- *Netperf* (page 341): a tool for measuring network performance
- *IOzone* (page 342): filesystem benchmark utility used to perform a broad filesystem analysis
- *arcstat* (page 344): used to gather ZFS ARC statistics

The following utilities are specific to RAID controllers:

- *tw_cli* (page 349): used to monitor and maintain 3ware RAID controllers
- *MegaCli* (page 350): used to configure and manage Broadcom MegaRAID SAS family of RAID controllers

This section also describes these utilities:

- *freenas-debug* (page 351): the backend used to dump FreeNAS® debugging information
- *tmux* (page 351): a terminal multiplexer similar to GNU screen
- *Dmidecode* (page 352): reports information about system hardware as described in the system's BIOS

23.1 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, it is used to test the speed of different types of shares to determine which type performs best on the network.

FreeNAS® includes the iperf server. To perform network testing, install an iperf client on a desktop system that has network access to the FreeNAS® system. This section demonstrates how to use the [xjperf user interface client](https://code.google.com/archive/p/xjperf/downloads) (<https://code.google.com/archive/p/xjperf/downloads>) as it works on Windows, macOS, 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 the package management system for their operating 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 macOS, Linux, or BSD, unzip the downloaded file, `cd` to the unzipped directory, type `chmod u+x jperf.sh`, and run `./jperf.sh`.

Start the iperf server on FreeNAS® when the client is ready.

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: The daemon process stops when [Shell](#) (page 325) closes. Set up the environment with shares configured and started **before** starting the Iperf process.

From the 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 23.1](#) shows an example of the client running on a Windows system while an SFTP transfer is occurring on the network.




Fig. 23.1: Viewing Bandwidth Statistics Using xjperf

Check the type of traffic before testing UDP or TCP. The iperf server is used to get additional details for services using TCP `iperf -sD` or UDP `iperf -sDu`. The startup message indicates when the server is listening for TCP or UDP. The `sockstat -4 | more` command gives an overview of the services running on the FreeNAS® system. This helps to determine if the traffic to test is UDP or TCP.

```
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 testing is finished, either type `killall iperf` or close Shell to terminate the iperf server process.

23.2 Netperf

Netperf is a benchmarking utility that can be used to measure the performance of unidirectional throughput and end-to-end latency.

Before using the `netperf` command, start its server process with this command:

```

netserver
Starting netserver with host 'IN(6)ADDR_ANY' port '12865' and family AF_UNSPEC

```

The following command displays the available options for performing tests with the `netperf` command. The [Netperf Manual](https://hewlettpackard.github.io/netperf/) (<https://hewlettpackard.github.io/netperf/>) 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 the results. When testing is finished, 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, and specifying a value with a trailing comma will set 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 behavior.

23.3 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 so it can be run from Shell. When using IOzone on FreeNAS®, `cd` to a directory in a pool that you have permission to write to, otherwise an error about being unable to write the temporary file will occur.

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 the results.

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](https://www.cyberciti.biz/tips/linux-filesystem-benchmarking-with-iozone.html) (<https://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](https://www.thegeekstuff.com/2011/05/iozone-examples/) (<https://www.thegeekstuff.com/2011/05/iozone-examples/>)

Type the following command to receive a summary of the available switches. IOzone is comprehensive so 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 `-+w 1 -+y 1 -+C 1`.

Alternatively, consider temporarily disabling compression on the ZFS pool or dataset when running IOzone benchmarks.

Note: If a visual representation of the collected data is preferred, 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.

```

23.4 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 then ported as a Python script for use on FreeNAS®.

Watching ARC hits/misses and percentages shows how well the ZFS pool is fetching from the ARC rather than using disk I/O. Ideally, there will be as many things fetching from cache as possible. Keep the load in mind while reviewing 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 made 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 network may not benefit another.

In particular, the value of pre-fetching depends upon the amount of memory and the type of workload, as seen in [Understanding ZFS: Prefetch](http://cuddletech.com/?page_id=834&id=1040) (http://cuddletech.com/?page_id=834&id=1040)

FreeNAS® provides two command line scripts which can be manually run from *Shell* (page 325):

- `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 provide real time information, whereas the current web interface 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
[-havxp] [-f fields] [-o file] [-s string] [interval [count]]

-h : Print this help message
-a : Print all possible stats
-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
-p : Disable auto-scaling of numerical fields
```

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%  arcz  c
06:19:03    7    0    0    0    0    0    0    0    0    153M 6.6G
06:19:04   257    0    0    0    0    0    0    0    0    153M 6.6G
06:19:05   193    0    0    0    0    0    0    0    0    153M 6.6G
06:19:06   193    0    0    0    0    0    0    0    0    153M 6.6G
06:19:07   255    0    0    0    0    0    0    0    0    153M 6.6G
```

Table 23.1 briefly describes the columns in the output.

Table 23.1: 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
arcz	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:
      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_enabled1
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_threshold70
vfs.zfs.metaslab.gang_bang 16777217
vfs.zfs.free_bpojb_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_recordsz 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 reading man pages and searching for the meaning of unfamiliar acronyms is required. **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 the workload.

If any of the ZFS tunables are changed, continue to monitor the system to determine the effect of the change. It is recommended that the changes are tested 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, it can be made permanent by creating a `sysctl` using the instructions in [Tunables](#) (page 91).

23.5 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\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=twe>) and [twa\(4\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=twa>) drivers.

Before using this command, read its [man page](#) (https://www.cyberciti.biz/files/tw_cli.8.html) as it describes the terminology and provides some usage examples.

When `tw_cli` in Shell is entered, the prompt will change, indicating that interactive mode is enabled where all sorts of maintenance commands on the controller and its arrays can be run.

Alternately, one command can be specified 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 the disks added to the array do not appear in the web interface, 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` to make them available in the FreeNAS® web interface.

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 give error notifications.

23.6 MegaCli

MegaCli is the command line interface for the Broadcom :MegaRAID SAS family of RAID controllers. FreeNAS® also includes the [mfiutil\(8\)](https://www.freebsd.org/cgi/man.cgi?query=mfiutil) (<https://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.

23.7 freenas-debug

The FreeNAS® web interface 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 325) 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:

-A Dump all debug information
-B Dump System Configuration Database
-C Dump SMB Configuration
-D Dump Domain Controller Configuration
-I Dump IPMI Configuration
-M Dump SATA DOMs Information
-N Dump NFS Configuration
-S Dump SMART Information
-T Loader Configuration Information
-Z Remove old debug information
-a Dump Active Directory Configuration
-c Dump (AD|LDAP) Cache
-e Email debug log to this comma-delimited list of email addresses
-f Dump AFP Configuration
-g Dump GEOM Configuration
-h Dump Hardware Configuration
-i Dump iSCSI Configuration
-j Dump Jail Information
-l Dump LDAP Configuration
-n Dump Network Configuration
-s Dump SSL Configuration
-t Dump System Information
-v Dump Boot System File Verification Status and Inconsistencies
-y Dump Sysctl Configuration
-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
```

23.8 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 325), `tmux` provides access to a command prompt while still giving access to the graphical administration screens.

To start a session, simply type `tmux`. As seen in Figure 23.2, 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.

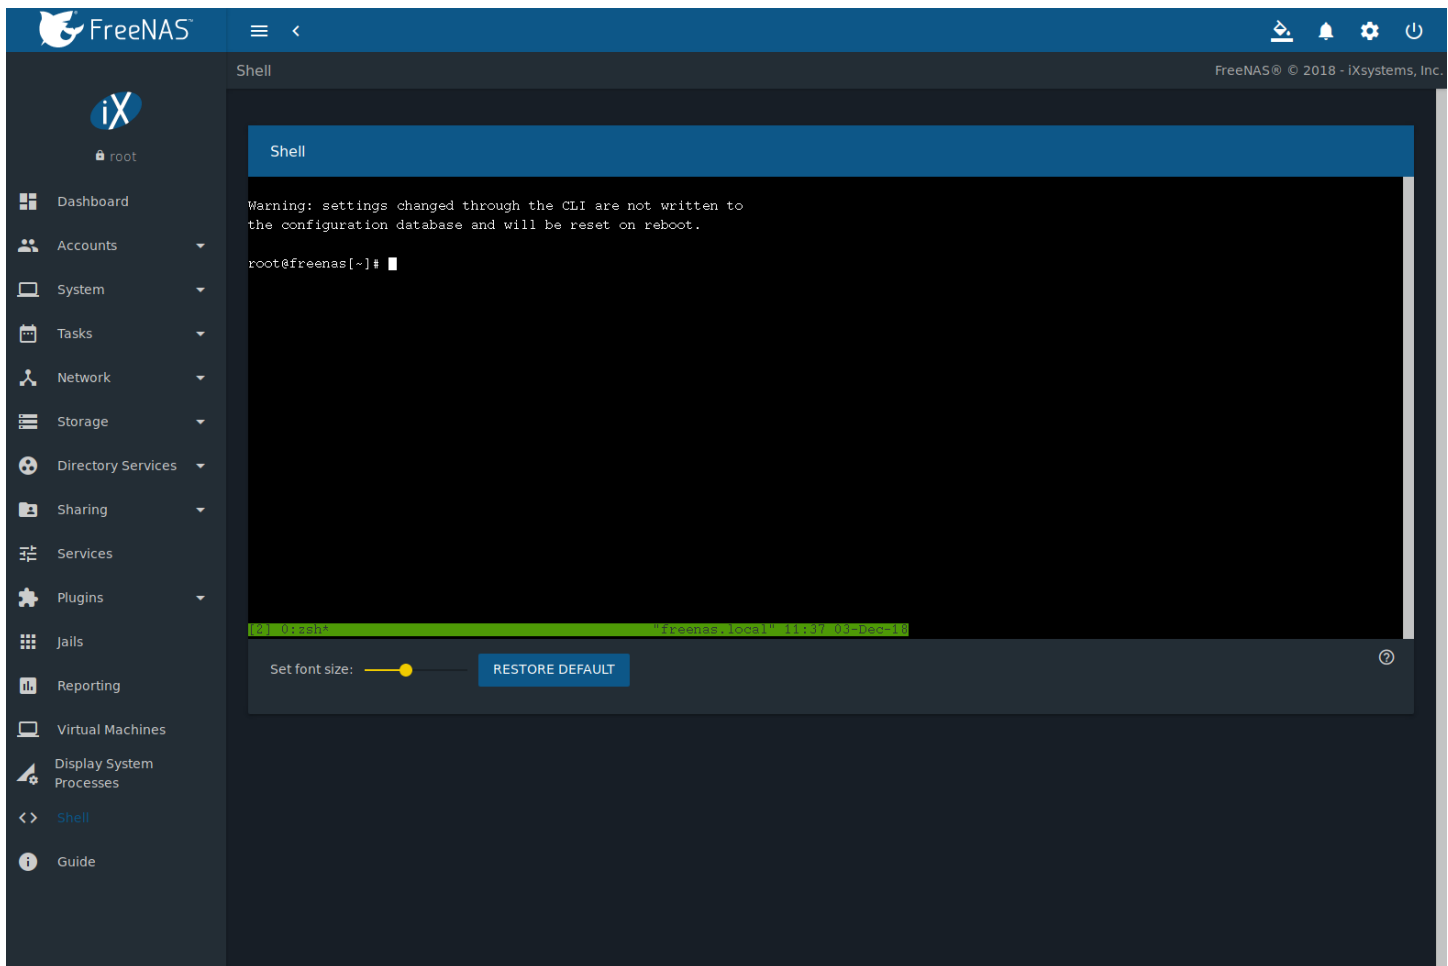


Fig. 23.2: tmux Session

To create a second window, press `Ctrl+b` then `"`. To close a window, type `exit` within the window.

`tmux(1)` (<http://man.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 `Shell` (page 325) is closed while `tmux` is running, it will detach its session. The next time `Shell` is open, run `tmux attach` to return to the previous session. To leave the `tmux` session entirely, type `exit`. If multiple windows are running, it is required 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/>)

23.9 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\(8\)](https://linux.die.net/man/8/dmidecode) (<https://linux.die.net/man/8/dmidecode>) describes the supported strings and types.

23.10 Midnight Commander

Midnight Commander is a program used to manage files from the shell. Open the application by running `mc`. The arrow keys are used to navigate and select files. Function keys are used to perform operations such as renaming, editing, and copying files. These resources provide more information about using Midnight Commander:


- [Midnight Commander wikipedia page](https://en.wikipedia.org/wiki/Midnight_Commander) (https://en.wikipedia.org/wiki/Midnight_Commander)
- [Midnight Commander website](https://midnight-commander.org/) (<https://midnight-commander.org/>)
- `mc(1)` (<https://www.freebsd.org/cgi/man.cgi?query=mc>)
- [Basic Tutorial](http://linuxcommand.org/lc3_adv_mc.php) (http://linuxcommand.org/lc3_adv_mc.php)

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/raid-z) (<https://blogs.oracle.com/bonwick/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 [Pools](#) (page 154), clicking  (Settings), then the *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®, [Pools](#) (page 154) is used to create or extend 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. ZFS pools created on FreeNAS® version 9.2.1 or later 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, store them for a period of time, 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, within 15 minutes of the data loss, for example. 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, pool 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). 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 325) 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 GiB 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. The ZFS pool version is checked from the *Shell* (page 325) 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)` (<https://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** (<http://www.brendangregg.com/blog/2008-07-22/zfs-l2arc.html>). 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 a 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 GiB 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 web interface 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](https://www.delphix.com/blog/delphix-engineering/zfs-raidz-stripe-width-or-how-i-learned-stop-worrying-and-love-raidz) (<https://www.delphix.com/blog/delphix-engineering/zfs-raidz-stripe-width-or-how-i-learned-stop-worrying-and-love-raidz>) for details.

These resources can also help determine the RAID configuration best suited to the specific storage requirements:

- [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](https://constantin.glez.de/2010/06/04/a-closer-look-zfs-vdevs-and-performance/) (<https://constantin.glez.de/2010/06/04/a-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 118) and [Replication Tasks](#) (page 120) 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 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 TiB in size.
- Using drives of equal sizes is recommended 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](https://docs.oracle.com/cd/E19253-01/819-5461/index.html) (<https://docs.oracle.com/cd/E19253-01/819-5461/index.html>)
- [Becoming a ZFS Ninja \(video\)](https://www.youtube.com/watch?v=6_K55Ira1Cs) (https://www.youtube.com/watch?v=6_K55Ira1Cs)
- [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-for-noobs.7775/) (<https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-for-noobs.7775/>)
- [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.

25.1 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://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)) ([https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628\(v=ws.11\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11))).
- *LUN Reporting* allows a hypervisor to query the NAS to determine whether a LUN is using thin provisioning.
- *Stun* pauses virtual machines when a pool 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 is configured at the pool level when using zvols (see [Table 11.6](#)) or at the extent level (see [Table 11.11](#)) 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 26.1](#), an online version of the API is available at api.freenas.org (http://api.freenas.org).

FreeNAS API 1.0 documentation » next | routing table | index

Table Of Contents
Welcome to FreeNAS API's documentation!
Indices and tables

Next topic
Release Notes

This Page
Show Source

Quick search
Go

Welcome to FreeNAS API's documentation!

Contents:

- Release Notes
 - Resources
 - Resources
 - Fields
 - Backwardly incompatible changes
- Authentication
 - Client
 - Examples
 - Python
- Resources
 - Account
 - Users
 - Groups
 - All Users
 - All Groups
 - Directory Service
 - ActiveDirectory
 - LDAP
 - NIS
 - Idmap
 - Jails
 - Configuration
 - Jails
 - MountPoints
 - Templates
 - Network
 - Global Configuration
 - Interface
 - VLAN
 - LAGG
 - LAGG Members
 - Static Route
 - Plugins
 - Plugins
 - Services
 - Services
 - AFP
 - CIFS
 - Domain Controller
 - DynamicDNS
 - FTP
 - iSCSI
 - LLD
 - NFS
 - Rsyncd
 - RsyncMod
 - SMART
 - SNMP
 - SSH
 - TFTP
 - UPS

Fig. 26.1: API Documentation

The rest of this section shows code examples to illustrate the use of the API.

Note: A new API was released with FreeNAS® 11.1. The previous API is still present and in use because it is feature-complete. 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. Work is under way to make the new API feature-complete. The new APIv2 uses [WebSockets](https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API) (https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API). This advanced technology makes it possible to open interactive communication sessions between web browsers and servers, allowing event-driven responses without the need to poll the server for a reply. When APIv2 is feature-complete, the FreeNAS® documentation will include relevant examples that make use of the new API.

26.1 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 pool or dataset for the user *Home Directory*. After creating the user, start the SSH service in *Services* → *SSH*. 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. Do not change the text in black. After saving changes, run the script by typing `python scriptname.py`. The new user account will appear in *Accounts* → *Users* in the FreeNAS® web interface.

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 *Network* → *Global Configuration*. Note that the script will fail if the machine running it is unable to resolve that hostname. Go to *System* → *General* and set the *Protocol* to *HTTP*.

Line 5: replace *freenas* with the password used to access the FreeNAS® system.

Line 7: to force validation of the SSL certificate while using HTTPS, 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 26.1](#) summarizes acceptable values. This resource uses JSON, so the boolean values are *True* or *False*.

Table 26.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
bsdusr_username	string	Maximum 32 characters, though a maximum of 8 is recommended for interoperability. Can include numerals but cannot include a space.
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 <i>bsdusr_creategroup</i> is set to <i>False</i> , specify the numeric ID of the group to create.
bsdusr_creategroup	boolean	Set <i>True</i> to automatically create a primary group with the same numeric ID as <i>bsdusr_uid</i> .
bsdusr_mode	string	Sets default numeric UNIX permissions of a user home directory.
bsdusr_shell	string	Specify the full path to a UNIX shell that is installed on the system.
bsdusr_password_disabled	boolean	Set to <i>True</i> to disable user login.
bsdusr_locked	boolean	Set to <i>True</i> to disable user login.
bsdusr_sudo	boolean	Set to <i>True</i> to enable <i>sudo</i> 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*.

26.2 A More Complex Example

This section provides a walk-through of a more complex example found in the `startup.py` script. Use the search bar within the API documentation to quickly locate the JSON parameters used here. This example defines a class and several methods to create a ZFS pool, 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 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 through 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, 152

A

Accessing a Jail Using SSH, 291

Add Dataset, 166

Add Group, 65

Add Jail, 282

Add Storage, 293

Add User, 68

Adding Devices to a VM, 310

Adding Storage, 293

Additional Storage, 293

Advanced Jail Creation, 283

AFP, 196, 238

Alert, 329

Alert Services, 85

Alert Settings, 87

API, 358

Apple Filing Protocol, 196, 238

arcstat, 344

Autotune, 81

B

Boot Environments, 75

Burn ISO, 19

C

CA, 97

Certificate Authority, 97

Certificates, 101

Change Theme, 60

Checksum, 19

CIFS, 211, 258

Cloud Credentials, 88

Cloud Sync, 134

Compression, 168

Create Group, 65

Create Jail, 282

Create New Themes, 61

Create User, 68

Creating VMs, 307

Cron Jobs, 108

D

DC, 240

DDNS, 241

Deduplication, 168

Delete Group, 67

Delete User, 70

Deleting VMs, 319

Dell PERC H330, 17

Dell PERC H730, 17

Disks, 176

Dmidecode, 352

Docker Host, 319

Docker Host Requirements, 319

Docker Hosts, 319

Domain Controller, 240

Download, 19

Dynamic DNS, 241

E

Email, 82

Encryption, 157

EtherChannel, 146

F

File Transfer Protocol, 243

Forums, 332

freenas-debug, 351

FTP, 243

G

Getting FreeNAS\ :sup:'@', 19

Groups, 64

H

Hardware Recommendations, 15

Highpoint RAID, 17

Hot Spares, 160

I

Install, 21

Internet Small Computer System Interface, 222

iocage, 302

IOzone, 342

Iperf, 338

IRC, 333

iSCSI, 222

ISO, 19

J

Jail Storage, 281
Jail Wizard, 282
Jails, 280

L

LACP, 146
LAGG, 146
Link Aggregation, 146
Link Layer Discovery Protocol, 248
LLDP, 248
Localize, 334
Log Out, 326, 327

M

Managing Jails, 289
MegaCli, 350
Midnight Commander, 353
Minio, 254
Mirroring the Boot Device, 77
Multiple Boot Environments, 75

N

Netdata, 249
Netperf, 341
Network File System, 203, 250
Network Settings, 139
New Group, 65
New Jail, 282
New User, 68
NFS, 203, 250
NTP Servers, 73

O

or Shut Down, 326

P

Periodic Snapshot, 118
Plugin, 270
Pools, 154
Professional Support, 333

R

Remove Group, 67
Remove User, 70
Replace Failed Drive, 178
Replication, 120
Reporting, 304
Resilver Priority, 131
Restart, 326, 327
RFC
 RFC 2616, 359
 RFC 3721, 224
Route, 151
Rsync, 252
Rsync Tasks, 111

Running VMs, 317

S

S.M.A.R.T., 256
S.M.A.R.T. Tests, 116
S3, 254
Samba, 211, 258
SCP, 264
Scrub, 132
Scrubs, 165
Secure Copy, 264
Secure Shell, 263
Self-Encrypting Drives, 81
Services, 236
Settings, 58
Shadow Copies, 221
Shell, 324
Shutdown, 328
Simple Network Management Protocol, 261
SMB, 211, 258
Snapshot, 118
Snapshots, 172
SNMP, 261
Spares, 160
SSH, 263, 291
Start Service, 237
Static Route, 151
Stop Service, 237
Support, 105
System Dataset, 84

T

Tasks, 107
TFTP, 265
Time Machine, 200
tmux, 351
Translate, 334
Translation, 334
Trivial File Transfer Protocol, 265
Trunking, 152
Tunables, 91
tw_cli, 349

U

Uninterruptible Power Supply, 266
Updating a Jail, 291
Upgrade, 29
Upgrade ZFS Pool, 35
Upgrading a Jail, 291
UPS, 266
USB Stick, 19
Users, 67

V

VAAI, 357
VAAI for iSCSI, 358
Virtual Machines, 306

Virtualization, [36](#)
VLAN, [152](#)
VM, [36](#)
VMs, [306](#)
VMware Snapshot, [175](#)

W

Web Interface Preferences, [59](#)
WebDAV, [210](#), [269](#)
Windows File Share, [258](#)
Windows Shares, [211](#)

Z

ZVOL, [169](#)