

TrueNAS® 11.3-U5 User Guide

Note: Starting with version 12.0, FreeNAS and TrueNAS are unifying (<https://www.ixsystems.com/blog/freenas-truenas-unification/>) into “TrueNAS”. Documentation for TrueNAS 12.0 and later releases has been unified and moved to the [TrueNAS Documentation Hub](https://www.truenas.com/docs/) (<https://www.truenas.com/docs/>).

Warning: To avoid the potential for data loss, iXsystems must be contacted **before** replacing a controller or upgrading to High Availability.

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 8:00AM to 5:00PM Pacific Standard Time: <ul style="list-style-type: none">• 1 (855) 473-7449 option 2 (US-only toll-free)• 1 (408) 943-4100 option 2 (local and international)
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none">• 1 (855) 499-5131 (US-only toll-free)• 1 (678) 835-6101 (local and international)

Copyright iXsystems 2011-2020

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

CONTENTS

Welcome	8
Typographic Conventions	9
1 Introduction	10
1.1 Contacting iXsystems	10
1.2 Path and Name Lengths	10
1.3 Using the Web Interface	12
1.3.1 Tables and Columns	12
1.3.2 Advanced Scheduler	12
1.3.3 Schedule Calendar	13
1.3.4 Changing TrueNAS® Settings	13
1.3.5 Web Interface Troubleshooting	14
1.3.6 Help Text	14
1.3.7 Humanized Fields	14
1.3.8 File Browser	14
2 Initial Setup	15
2.1 Hardware Setup	15
2.2 Console Setup Menu	15
2.3 Accessing the Web Interface	17
2.3.1 Web Interface Troubleshooting	19
3 Settings	20
3.1 Change Password	20
3.2 Preferences	20
3.2.1 Web Interface Preferences	20
3.2.2 Themes	21
3.2.2.1 Create New Themes	21
3.3 API Documentation	23
3.4 About	23
3.5 Legacy Web Interface	23
4 Accounts	24
4.1 Groups	24
4.2 Users	27
5 System	32
5.1 General	32
5.2 NTP Servers	35
5.3 Boot	37
5.3.1 Operating System Device Mirroring	38
5.4 Advanced	39
5.4.1 Autotune	41
5.4.2 Self-Encrypting Drives	41

5.4.2.1	Deploying SEDs	42
5.4.2.2	Check SED Functionality	43
5.4.2.3	Managing SED Passwords and Data	44
5.5	View Enclosure	45
5.6	Email	47
5.7	System Dataset	48
5.8	Reporting	50
5.9	Alert Services	50
5.10	Alert Settings	52
5.11	Cloud Credentials	53
5.12	SSH Connections	57
5.12.1	Manual Setup	59
5.12.2	Semi-Automatic Setup	60
5.13	SSH Keypairs	61
5.14	Tunables	61
5.15	Update	64
5.15.1	Preparing for Updates	64
5.15.2	Updates and Trains	65
5.15.3	Checking for Updates	65
5.15.4	Saving the Configuration File	67
5.15.5	Applying Updates	68
5.15.6	Manual Updates	68
5.15.7	Update in Progress	69
5.15.8	Updating from the Shell	69
5.15.9	Updating an HA System	69
5.15.10	If Something Goes Wrong	71
5.15.11	Upgrading a ZFS Pool	71
5.16	CAs	72
5.17	Certificates	76
5.18	Failover	81
5.19	Support	84
5.19.1	License Information	85
5.19.2	Proactive Support	85
5.19.3	Contact Support	86
6	Tasks	87
6.1	Cron Jobs	87
6.2	Init/Shutdown Scripts	89
6.3	Rsync Tasks	90
6.3.1	Rsync Module Mode	93
6.3.2	Rsync over SSH Mode	93
6.4	S.M.A.R.T. Tests	96
6.5	Periodic Snapshot Tasks	97
6.5.1	Snapshot Autoremoval	99
6.6	Replication	100
6.6.1	Replication Creation Wizard	100
6.6.2	Advanced Replication Creation	103
6.6.3	Replication Tasks	107
6.6.4	Limiting Replication Times	107
6.6.5	Replication Topologies and Scenarios	108
6.6.5.1	Star Replication	108
6.6.5.2	Tiered Replication	108
6.6.5.3	N-way Replication	108
6.6.5.4	Disaster Recovery	108
6.6.6	Troubleshooting Replication	109
6.6.6.1	SSH	109
6.6.6.2	Compression	109

6.6.6.3	Manual Testing	109
6.7	Resilver Priority	110
6.8	Scrub Tasks	111
6.9	Cloud Sync Tasks	113
6.9.1	Cloud Sync Example	116
7	Network	119
7.1	Global Configuration	119
7.2	Interfaces	121
7.2.1	Network Bridges	124
7.2.2	Link Aggregations	124
7.2.2.1	LACP, MPIO, NFS, and ESXi	125
7.2.2.2	Creating a Link Aggregation	125
7.2.2.3	Link Aggregation Options	125
7.2.3	VLANs	126
7.3	IPMI	126
7.4	Network Summary	128
7.5	Static Routes	128
8	Storage	130
8.1	Swap Space	130
8.2	Pools	130
8.2.1	Creating Pools	131
8.2.2	Managing Encrypted Pools	133
8.2.2.1	Encryption and Recovery Keys	135
8.2.2.2	Encryption Operations	135
8.2.3	Adding Cache or Log Devices	136
8.2.4	Removing Cache or Log Devices	137
8.2.5	Adding Spare Devices	137
8.2.6	Extending a Pool	137
8.2.7	Export/Disconnect a Pool	137
8.2.8	Importing a Pool	139
8.2.9	Viewing Pool Scrub Status	141
8.2.10	Adding Datasets	142
8.2.10.1	Compression	145
8.2.11	Adding Zvols	145
8.2.12	Setting Permissions	147
8.2.13	ACL Management	148
8.3	Snapshots	152
8.3.1	Browsing a Snapshot Collection	153
8.3.2	Creating a Single Snapshot	154
8.4	VMware-Snapshots	155
8.5	Disks	156
8.5.1	Replacing a Failed Disk	158
8.5.1.1	Removing a Log or Cache Device	161
8.5.2	Replacing Disks to Grow a Pool	161
8.6	Importing a Disk	162
8.7	Multipaths	163
9	Overprovisioning	164
10	Directory Services	165
10.1	Active Directory	165
10.1.1	Leaving the Domain	170
10.1.2	Troubleshooting Tips	170
10.2	LDAP	170
10.3	NIS	173
10.4	Kerberos Realms	174

10.5	Kerberos Keytabs	175
10.6	Kerberos Settings	176
11	Sharing	177
11.1	Apple (AFP) Shares	178
11.1.1	Creating AFP Guest Shares	180
11.2	Block (iSCSI)	183
11.2.1	iSCSI Wizard	184
11.2.2	Target Global Configuration	185
11.2.3	Portals	186
11.2.4	Initiators	187
11.2.5	Authorized Access	189
11.2.6	Targets	190
11.2.7	Extents	192
11.2.8	Associated Targets	194
11.2.9	Fibre Channel Ports	196
11.2.10	Connecting to iSCSI	198
11.2.11	Growing LUNs	199
11.2.11.1	zvol Based LUN	199
11.2.11.2	File Extent Based LUN	200
11.3	Unix (NFS) Shares	201
11.3.1	Example Configuration	204
11.3.2	Connecting to the Share	204
11.3.2.1	From BSD or Linux	204
11.3.2.2	From Microsoft	205
11.3.2.3	From macOS	205
11.3.3	Troubleshooting NFS	206
11.4	WebDAV Shares	207
11.5	Windows (SMB) Shares	209
11.5.1	Configuring Unauthenticated Access	215
11.5.2	Configuring Authenticated Access With Local Users	216
11.5.3	User Quota Administration	218
11.5.4	Configuring Shadow Copies	218
11.6	Creating Authenticated and Time Machine Shares	218
11.6.1	Setting SMB and AFP Share Quotas	220
11.6.2	Client Time Machine Configuration	221
12	Services	222
12.1	Configure Services	222
12.2	AFP	223
12.2.1	Troubleshooting AFP	225
12.3	Dynamic DNS	225
12.4	FTP	226
12.4.1	Anonymous FTP	229
12.4.2	FTP in chroot	230
12.4.3	Encrypting FTP	231
12.4.4	Troubleshooting FTP	231
12.5	iSCSI	231
12.6	LLDP	231
12.7	NFS	232
12.8	Rsync	234
12.8.1	Configure Rsyncd	234
12.8.2	Rsync Modules	235
12.9	S3	237
12.10	S.M.A.R.T.	238
12.11	SMB	240
12.11.1	Troubleshooting SMB	242

12.12SNMP	242
12.13SSH	244
12.13.1 SCP Only	246
12.13.2 Troubleshooting SSH	246
12.14TFTP	247
12.15UPS	248
12.15.1 Multiple Computers with One UPS	251
12.16WebDAV	251
13 Plugins	254
13.1 Installing Plugins	254
13.2 Updating Plugins	256
13.3 Uninstalling Plugins	257
13.4 Asigra Plugin	258
14 Jails	260
14.1 Jail Storage	260
14.2 Creating Jails	261
14.2.1 Jail Wizard	261
14.2.2 Advanced Jail Creation	263
14.2.2.1 Creating Template Jails	271
14.3 Managing Jails	271
14.3.1 Jail Updates and Upgrades	273
14.3.2 Accessing a Jail Using SSH	273
14.3.3 Additional Storage	275
14.4 Jail Software	278
14.4.1 Installing FreeBSD Packages	278
14.4.2 Compiling FreeBSD Ports	279
14.4.3 Starting Installed Software	282
15 Reporting	285
16 Virtual Machines	288
16.1 Creating VMs	290
16.2 Installing Docker	293
16.3 Adding Devices to a VM	293
16.3.1 CD-ROM Devices	294
16.3.2 NIC (Network Interfaces)	295
16.3.3 Disk Devices	296
16.3.4 Raw Files	297
16.3.5 VNC Interface	298
17 vCenter Plugin	300
18 Additional Options	301
18.1 Display System Processes	301
18.2 Shell	302
18.3 Log Out, Restart, or Shut Down	303
18.3.1 Log Out	303
18.3.2 Restart	303
18.3.3 Shut down	304
18.4 Alert	305
19 Task Manager	308
20 ZFS Primer	309
20.1 ZFS Feature Flags	312

21 VMware Recommendations	313
21.1 TrueNAS® as a VMware Guest	313
21.2 Hosting VMware Storage with TrueNAS®	313
21.3 VAAI for iSCSI	314
22 Using the API	315
22.1 A Simple API Example	316
22.2 A More Complex Example	317
23 User Guide	319
24 Appendix A: End-User License Agreement	320
25 Appendix B: TrueNAS® Product Catalog	324
25.1 TrueNAS® Unified Storage Arrays	324
25.1.1 X-Series	324
25.1.2 M-Series	325
25.2 Expansion Shelves	325
25.2.1 ES12	325
25.2.2 ES24	326
25.2.3 ES60	326

Welcome

Welcome to the TrueNAS® User Guide.

TrueNAS® and the TrueNAS® 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.

Asigra Inc. Asigra, the Asigra logo, Asigra Cloud Backup, Recovery is Everything, Recovery Tracker and Attack-Loop are trademarks of Asigra Inc.

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.

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

Linux® is a registered trademark of Linus Torvalds.

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

UNIX® is a registered trademark of The Open Group.

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 TrueNAS® Administrator 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: TrueNAS® Icons

Icon	Usage
<i>ADD</i>	Add a new item.
⚙️ (Settings)	Show a settings menu.
⋮ (Options)	Show an Options menu.
📁 (Browse)	Shows an expandable view of system directories.
⏻ (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.
🔒 (Encryption Options)	Encryption options for a pool.
📌 (Pin)	Pin a help box to the screen.
✕ (Close)	Close a help box.

INTRODUCTION

This Guide provides information about configuring and managing the TrueNAS® Unified Storage Array. Your iXsystems support engineer will assist with the initial setup and configuration of the array. After becoming familiar with the configuration workflow, this document can be used as a reference guide to the many features provided by TrueNAS®.

1.1 Contacting iXsystems

For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 6:00AM to 6:00PM Pacific Standard Time: <ul style="list-style-type: none">• US-only toll-free: 855-473-7449 option 2• Local and international: 408-943-4100 option 2
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none">• US-only toll-free: 855-499-5131• International: 408-878-3140 (international calling rates will apply)

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 1.2: Path and Name Lengths

Type	Maximum Length	Description
File Paths	1023 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. If a path with long directory and file names exceeds the 1023-byte limit, it prevents 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>SPECNAMELEN</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 302) 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 302) 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 Using the Web Interface

1.3.1 Tables and Columns

Tables show a subset of all available columns. Additional columns can be shown or hidden with the *COLUMNS* button. Set a checkmark by the fields to be shown in the table. Column settings are remembered from session to session.

The original columns can be restored by clicking *Reset to Defaults* in the column list.

Each row in a table can be expanded to show all the information by clicking the > (Expand) button.

1.3.2 Advanced Scheduler

When choosing a schedule for different TrueNAS® *Tasks* (page 87), clicking *Custom* opens the custom schedule dialog.

The screenshot displays the 'Schedule Preview' dialog in TrueNAS. On the left, a calendar for January 2020 is shown with the 3rd day highlighted. Below the calendar, a list of scheduled events is visible: 'Sun Jan 5 2020 12:00 GMT-5', 'Sun Jan 12 2020 12:00 GMT-5', 'Sun Jan 19 2020 12:00 GMT-5', and 'Sun Jan 26 2020 12:00 GMT-5'. On the right, the 'Presets' section shows 'Daily' selected. Below this, the 'Minutes/Hours/Days' section has input fields for 'Minutes *' (0), 'Hours *' (0), and 'Days *' (*). The 'Months' section shows checkboxes for all months (Jan-Dec), all of which are currently unchecked. The 'Days of Week' section shows checkboxes for all days (Sun-Sat), with 'Sun' checked. A 'DONE' button is located at the bottom right.

Fig. 1.1: Creating a Custom Schedule

Choosing a preset schedule fills in the rest of the fields. To customize a schedule, enter *crontab* (<https://www.freebsd.org/cgi/man.cgi?query=crontab&sektion=5>) values for the *Minutes/Hours/Days*.

These fields accept standard `cron` values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering `10` means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means “match all values”.

Specific time ranges are set by entering hyphenated number values. For example, entering `30–35` in the *Minutes* field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

Lists of values can also be entered. Enter individual values separated by a comma (,). For example, entering `1, 14` in the *Hours* field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, while entering * in *Days* means the task runs every day of the month, `*/2` means the task runs every other day.

Combining all these examples together creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

There is an option to select which *Months* the task will run. Leaving each month unset is the same as selecting every month.

The *Days of Week* schedules the task to run on specific days. This is in addition to any listed *Days*. For example, entering `1` in *Days* and setting *W* for *Days of Week* creates a schedule that starts a task on the first day of the month **and** every Wednesday of the month.

Schedule Preview shows when the current schedule settings will cause the task to run.

1.3.3 Schedule Calendar

The *Schedule* column has a calendar icon (📅). Clicking this icon opens a dialog showing scheduled dates and times for the related task to run.

The screenshot shows a table titled "Scrub Tasks" with columns: Pool, Threshold days, Description, Schedule, Next Run, and Enabled. A task row for "pool1" is selected, showing a threshold of 35 days and a schedule of "00 00 * * *". A calendar icon is visible in the Schedule column. A popup titled "Upcoming tasks" is displayed, listing five scheduled runs: Sun Sep 08 2019 00:00:00 GMT-0400, Sun Sep 15 2019 00:00:00 GMT-0400, Sun Sep 22 2019 00:00:00 GMT-0400, Sun Sep 29 2019 00:00:00 GMT-0400, and Sun Oct 06 2019 00:00:00 GMT-0400.

Fig. 1.2: Example Schedule Popup

Scrub Tasks (page 111) can have a number of *Threshold days* set. The configured scrub task continues to follow the displayed calendar schedule, but it does not run until the configured number of threshold days have elapsed.

1.3.4 Changing TrueNAS® Settings

It is important to use the web interface or the Console Setup menu for all configuration changes. TrueNAS® stores configuration settings in a database. Commands entered at the command line **do not modify the settings database**. This means that changes made at the command line will be lost after a restart and overwritten by the values in the settings database.

1.3.5 Web Interface Troubleshooting

If the web interface is shown but seems unresponsive or incomplete:

- Make sure the browser allows cookies, Javascript, and custom fonts from the TrueNAS® system.
- Try a different browser. [Firefox](https://www.mozilla.org/en-US/firefox/all/) (<https://www.mozilla.org/en-US/firefox/all/>) is recommended.

If a web browser cannot connect to the TrueNAS® system by IP address, DNS hostname, or mDNS name:

- Check or disable proxy settings in the browser.
- Verify the network connection by pinging the TrueNAS® system by IP address from another computer on the same network. For example, if the TrueNAS® system is at IP address 192.168.1.19, enter `ping 192.168.1.19` on the command line of the other computer. If there is no response, check network configuration.

1.3.6 Help Text

Most fields and settings in the web interface have a ⓘ (Help Text) icon. Additional information about the field or setting can be shown by clicking ⓘ (Help Text). The help text window can be dragged to any location, and will remain there until ✕ (Close) or ⓘ (Help Text) is clicked to close the window.

1.3.7 Humanized Fields

Some numeric value fields accept *humanized* values. This means that the field accepts numbers or numbers followed by a unit, like `M` or `MiB` for megabytes or `G` or `GiB` for gigabytes. Entering `1048576` or `1M` are equivalent. Units of KiB, MiB, GiB, TiB, and PiB are available, and decimal values like `1.5 GiB` are supported when the field allows them. Some fields have minimum or maximum limits on the values which can restrict the units available.

1.3.8 File Browser

Certain sections of the web interface have a built in file browser. The file browser is used to traverse through directories and choose datasets on the system. Datasets that have [complex ACL permissions](#) (page 148) are tagged so they can be distinguished from non-ACL datasets.

INITIAL SETUP

2.1 Hardware Setup

Basic Setup Guides for TrueNAS® systems and expansion shelves are included with the hardware and also available in the [iX Information Library](https://www.ixsystems.com/blog/knowledgebase_category/truenas/) (https://www.ixsystems.com/blog/knowledgebase_category/truenas/). These guides provide detailed instructions on included components, controls, ports, rack installation, drive loading, and cable connections.

Complete hardware installation before continuing.

Note: Always perform the initial TrueNAS® setup in consultation with your iXsystems Support Representative. iXsystems Support can be contacted at truenas-support@ixsystems.com. Be sure to have all TrueNAS® hardware serial numbers on hand. The serial numbers are located on the back of each chassis.

2.2 Console Setup Menu

The Console Setup menu, shown in [Figure 2.1](#), appears at the end of the boot process. If the TrueNAS® system has a keyboard and monitor, this Console Setup menu can be used to administer the system.

Note: When connecting to the TrueNAS® system with SSH or the web [Shell](#) (page 302), 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
12) Toggle automatic support alerts to iXsystems

The web user interface is at:

http://10.215.5.2
http://10.215.5.4
https://10.215.5.2
https://10.215.5.4

Enter an option from 1-12: 
```

Fig. 2.1: Console Setup Menu

Note: On HA systems, some of these menu options are not available unless HA has been administratively disabled.

The menu provides these options:

- 1) *Configure Network Interfaces* provides a configuration wizard to set up the system's network interfaces. If the system has been licensed for High Availability (HA), the wizard prompts for IP addresses for both "This Controller" and "TrueNAS Controller 2".
- 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 TrueNAS® 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.

During boot, TrueNAS® automatically attempts to connect to a DHCP server from all live interfaces. If it successfully receives an IP address, the address is displayed so it can be used to access the graphical user interface. In the example seen in [Figure 2.1](#), the TrueNAS® system is accessible at <http://10.0.0.102>.

Some TrueNAS® 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 *truenas.local*.

If the TrueNAS® 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 TrueNAS® system has one network interface, *em0*.

```
Enter an option from 1-12: 1
1) em0
Select an interface (q to quit): 1
Remove the current settings of this interface? (This causes a momentary disconnection of the network.) (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
```

2.3 Accessing the Web Interface

The IP address of the TrueNAS® graphical web interface is provided on the TrueNAS® sales order or configuration sheet. Please [contact iX Support](#) (page 10) if the TrueNAS® web interface IP address has not been provided with these documents or cannot be identified from the TrueNAS® system console.

On a computer that can access the same network as the TrueNAS® system, enter the IP address in a web browser to connect to the web interface. The password for the root user is requested.

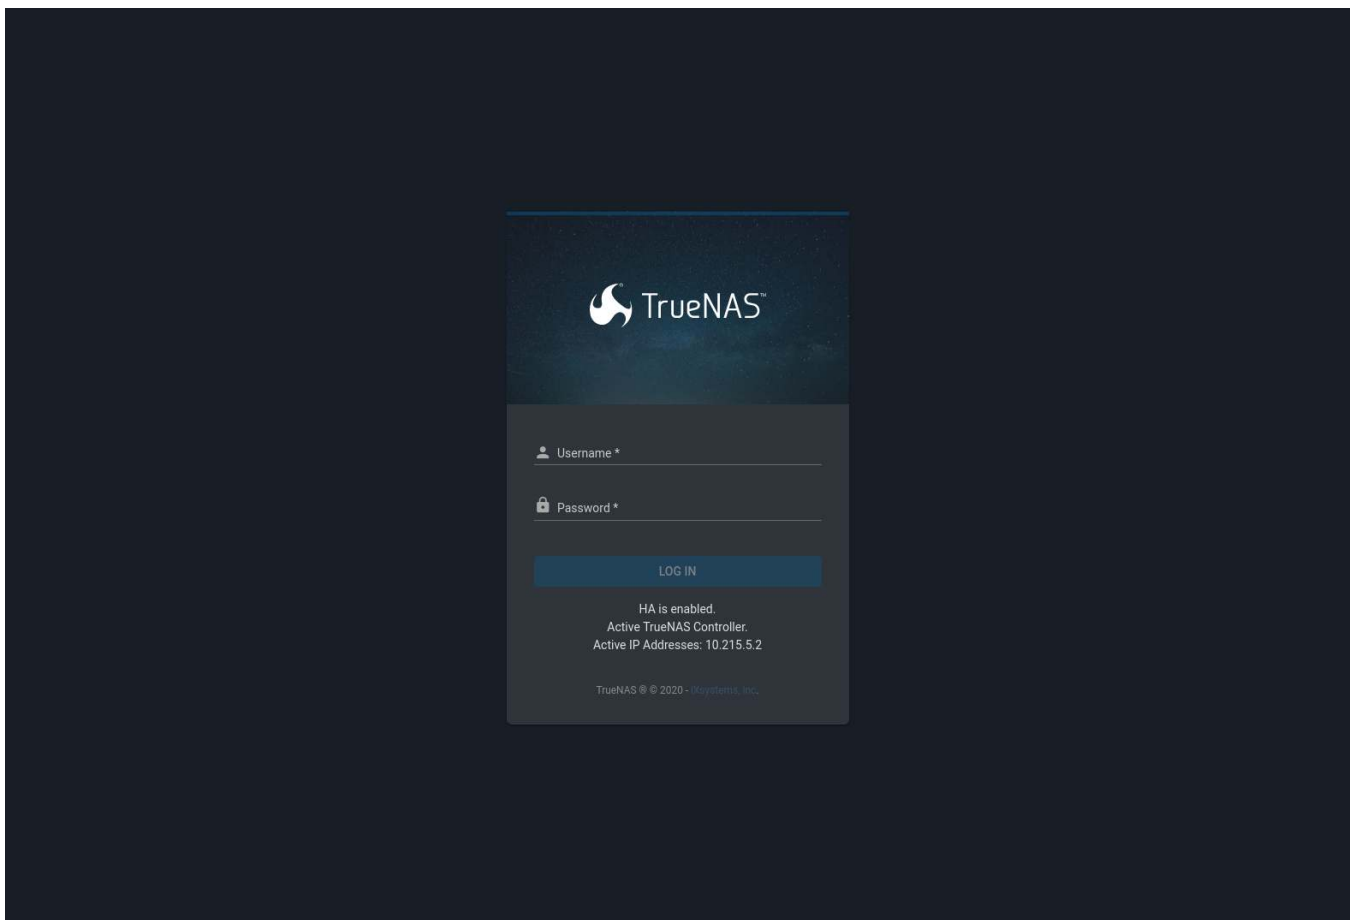


Fig. 2.2: Login Screen

The [High Availability \(HA\)](#) (page 81) status and information about the active TrueNAS controller is displayed on this screen. Log in with:

- **Username:** root
- **Password:** abcd1234

Note: The default *root* password can be changed to a more secure value by going to *Accounts* → *Users*. Expand the entry for *root* and click *EDIT*. Enter the new password in the *Password* and *Confirm Password* fields and click *SAVE*. The new password is used for subsequent logins.

On the first login, the EULA found in [Appendix A: End-User License Agreement](#) (page 320) is displayed, along with a box where the license key for the TrueNAS[®] array can be pasted. Read the EULA and paste in the license key. High Availability (HA) systems must have both active and standby TrueNAS controllers running before the license key for the HA TrueNAS[®] system can be entered. The key is entered on the active TrueNAS controller. Click *OK* to save the license key and access the web interface.

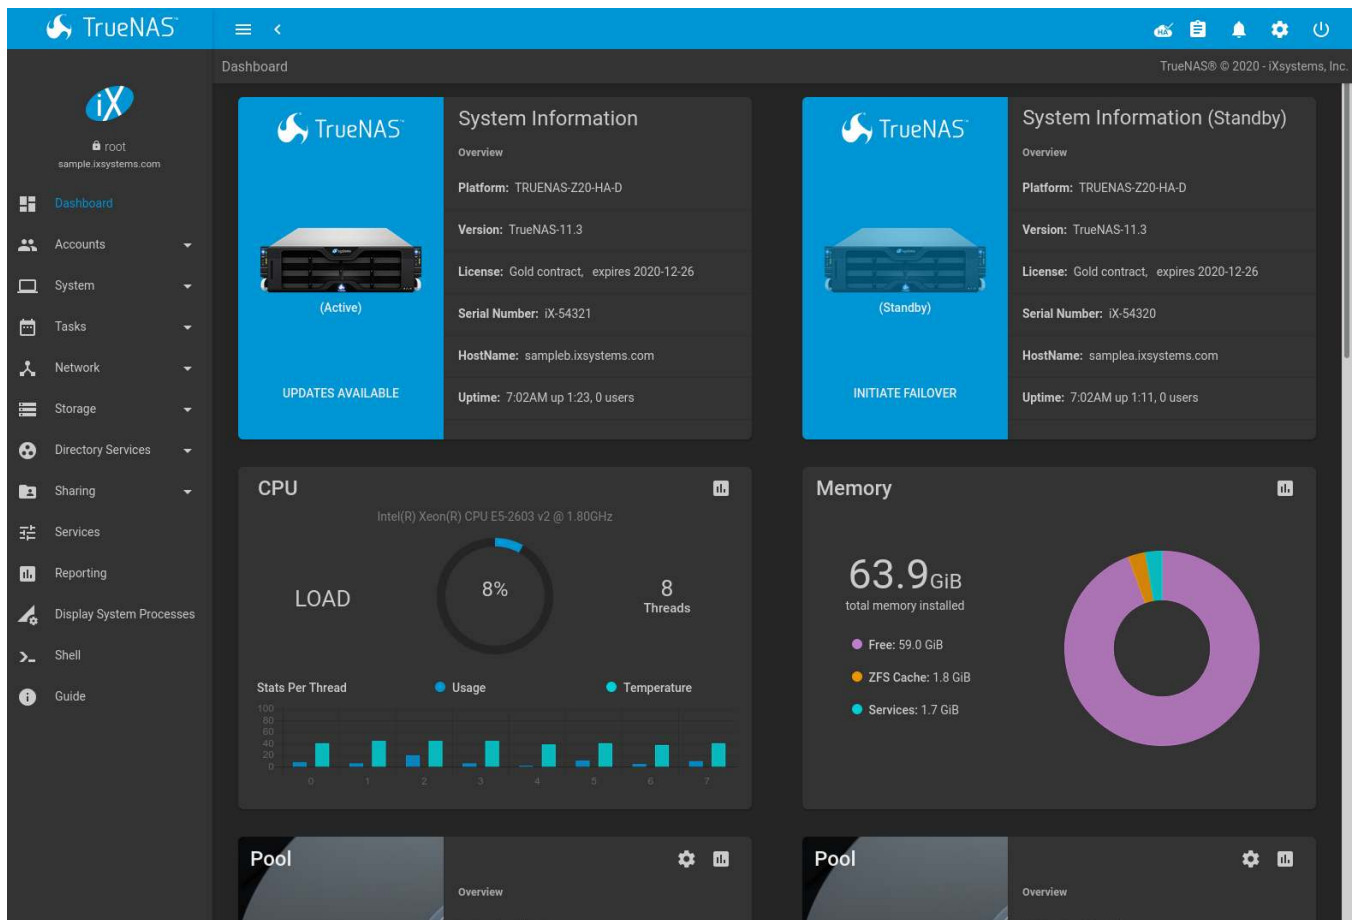


Fig. 2.3: Dashboard

The *Dashboard* shows details about the system. These details are grouped into sections about the hardware components, networking, storage, and other categories.

2.3.1 Web Interface Troubleshooting

If the user interface is not accessible by IP address from a browser, check these things:

- Are proxy settings enabled in the browser configuration? If so, disable the settings and try connecting again.
- If the page does not load, make sure that a `ping` reaches the TrueNAS® system's IP address. If the address is in a private IP address range, it is only accessible from within that private network.

If the UI becomes unresponsive after an upgrade or other system operation, clear the site data and refresh the browser.

The rest of this User Guide describes the TrueNAS® 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.

Please [contact iXsystems Support](#) (page 10) for initial setup and configuration assistance.

Warning: It is important to use the web interface or the console setup menu for all configuration changes. Do not make changes from the command line unless directed by an iXsystems Support Engineer.

SETTINGS

The ⚙️ (Settings) menu provides options to change the administrator password, set preferences, and view system information.

3.1 Change Password

To change the `root` account password, click ⚙️ (Settings) and *Change Password*. The current `root` password must be entered before a new password can be saved.

3.2 Preferences

The TrueNAS® 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*.

3.2.1 Web Interface Preferences

This page has options to adjust global settings in the web interface, manage custom themes, and create new themes. [Figure 3.1](#) shows the different options:

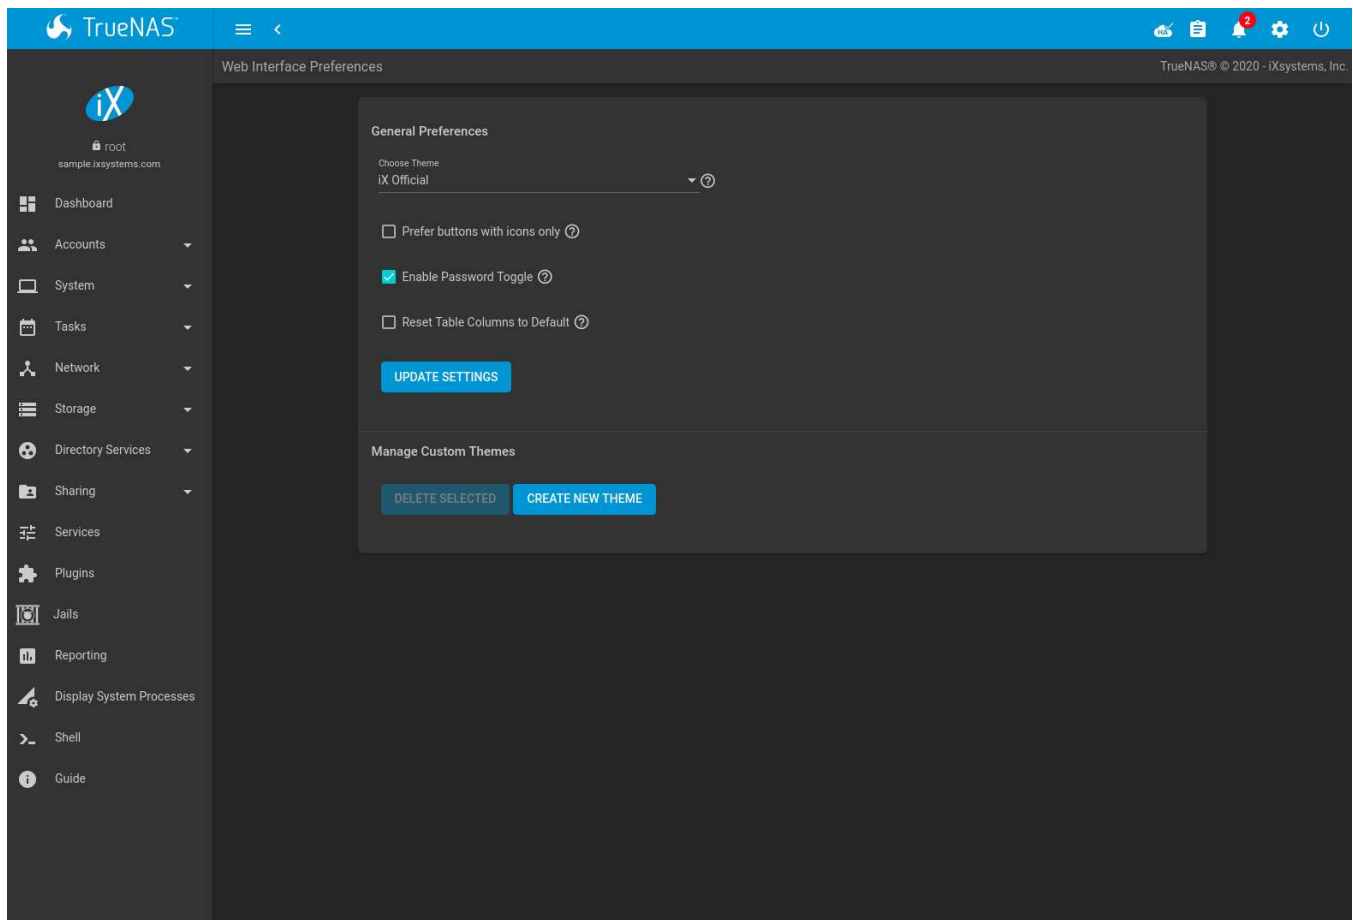


Fig. 3.1: Web Interface Preferences

These options are applied to the entire web interface:

- *Choose Theme*: Change the active theme. Custom themes are added to this list.
- *Prefer buttons with icons only*: Set to preserve screen space and only display icons and tooltips instead of text labels.
- *Enable Password Toggle*: When set, an eye icon appears next to password fields. Clicking the icon reveals the password. Clicking it again hides the password.
- *Reset Table Columns to Default*: Set to reset all tables to display default columns.

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

3.2.2 Themes

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

3.2.2.1 Create New Themes

This page is used to create and preview custom TrueNAS® themes. [Figure 3.2](#) shows many of the theming and preview options:

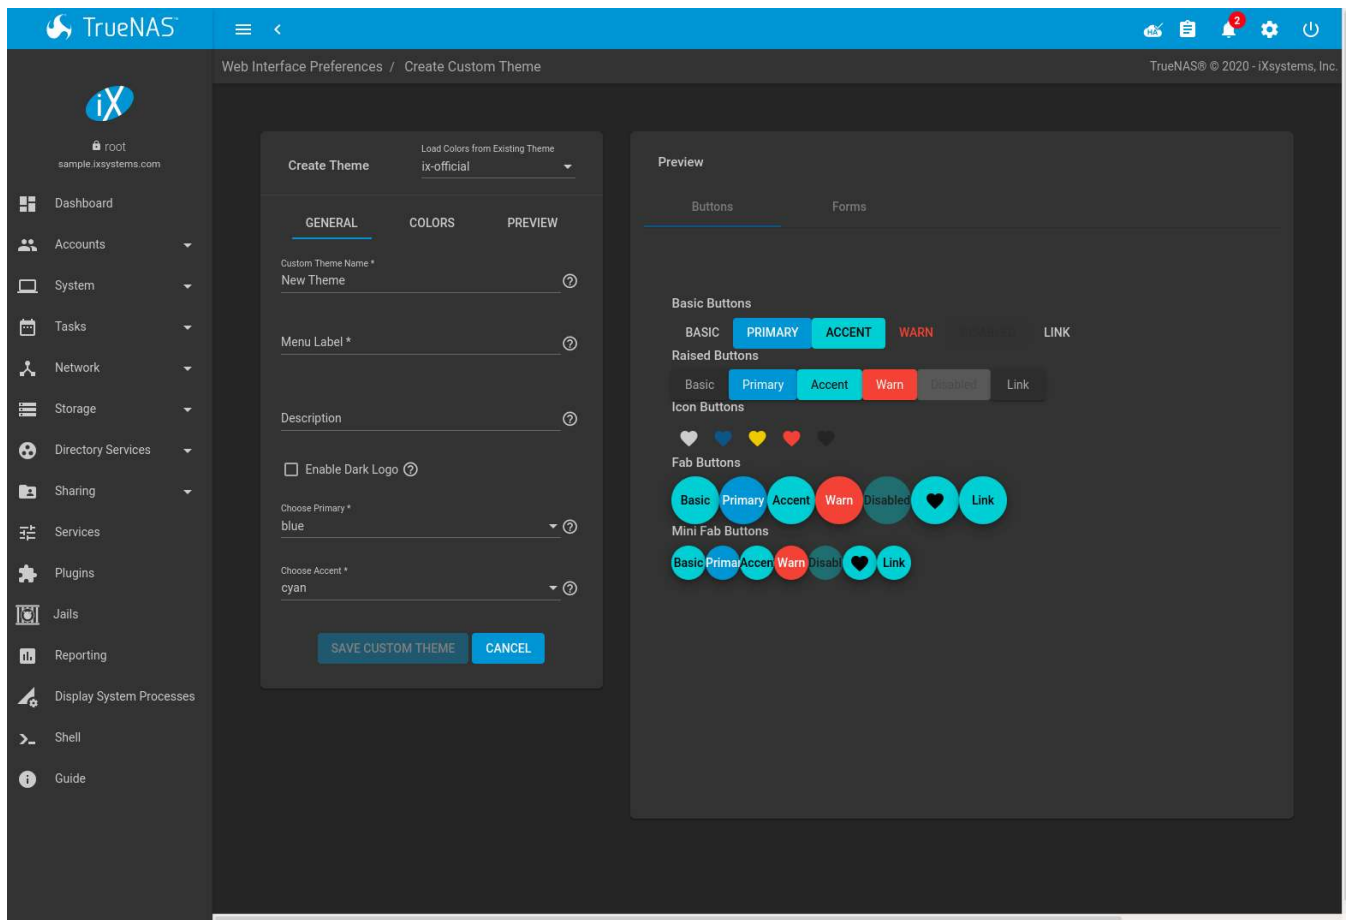


Fig. 3.2: Create and Preview a Custom Theme

To create a new custom theme, click *CREATE NEW THEME*. Colors from an existing theme can be used when creating a new custom theme. Select a theme from the *Load Colors from Theme* drop-down to use the colors from that theme for the new custom theme. Table 3.1 describes each option:

Table 3.1: 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 TrueNAS® 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.
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 is added to the list of available themes in *Web Interface Preferences*.

Click *PREVIEW* → *Global Preview* to apply the unsaved custom theme to the current session of the TrueNAS® 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.

3.3 API Documentation

Click *API* to see documentation for the [websocket protocol API](https://en.wikipedia.org/wiki/WebSocket) (<https://en.wikipedia.org/wiki/WebSocket>) used in TrueNAS®.

3.4 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*.

3.5 Legacy Web Interface

This option is only available when *Enable Legacy User Interface* is set in the [advanced system settings](#) (page 40).

Click ⚙ (Settings) and *Legacy Web Interface* to switch to the previous TrueNAS® 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 24): used to manage UNIX-style groups on the TrueNAS® system.
- [Users](#) (page 27): used to manage UNIX-style accounts on the TrueNAS® system.

Each entry is described in more detail in this section.

4.1 Groups

The Groups interface provides management of UNIX-style groups on the TrueNAS® 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 TrueNAS®. Refer to [Directory Services](#) (page 165) for details.

This section describes how to create a group and assign user accounts to it. The *Groups* page lists all groups, including those built in and used by the operating system.

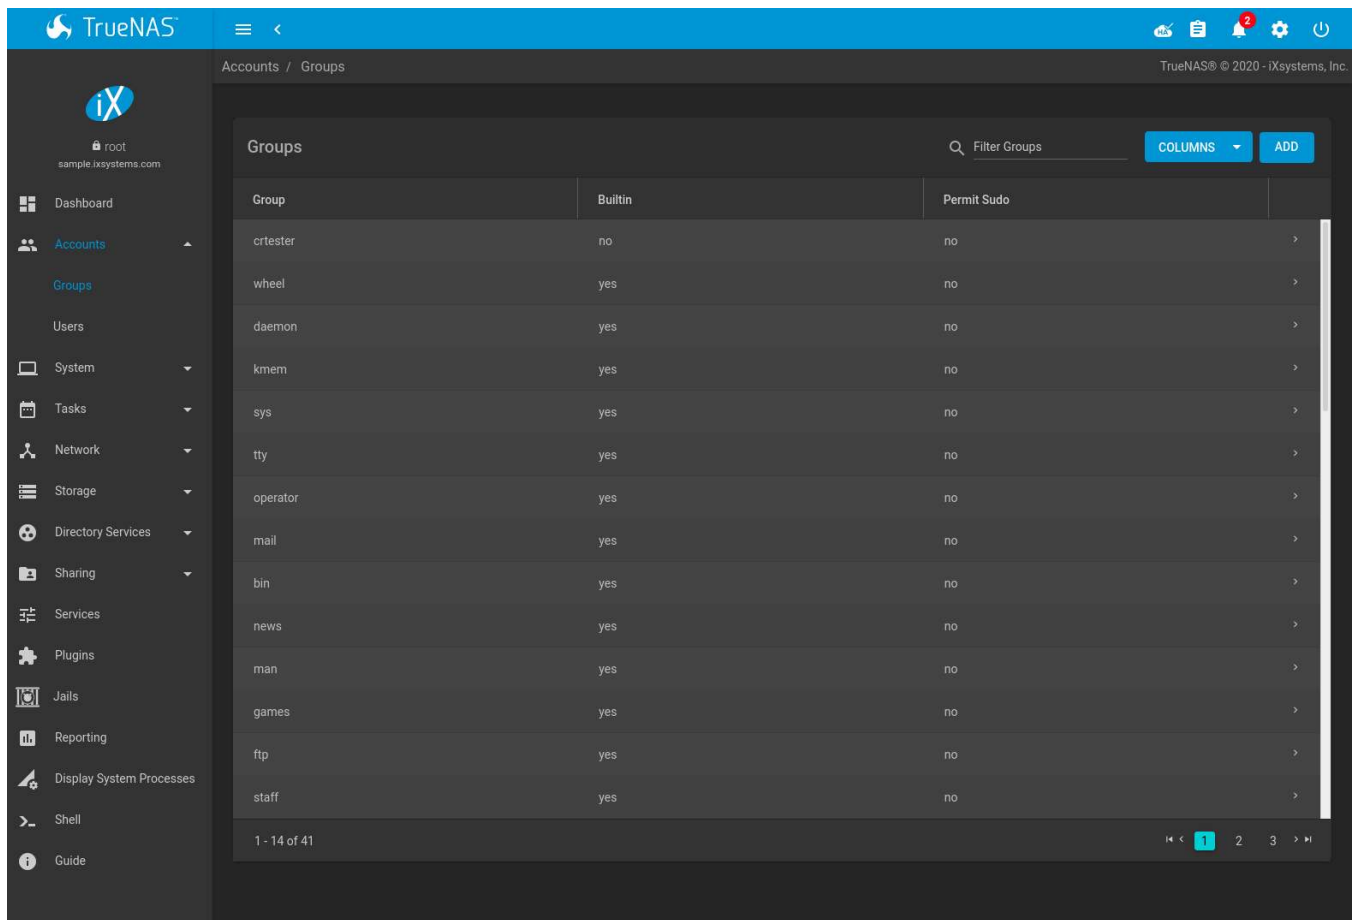


Fig. 4.1: Group Management

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 TrueNAS® system and cannot be edited or deleted.

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

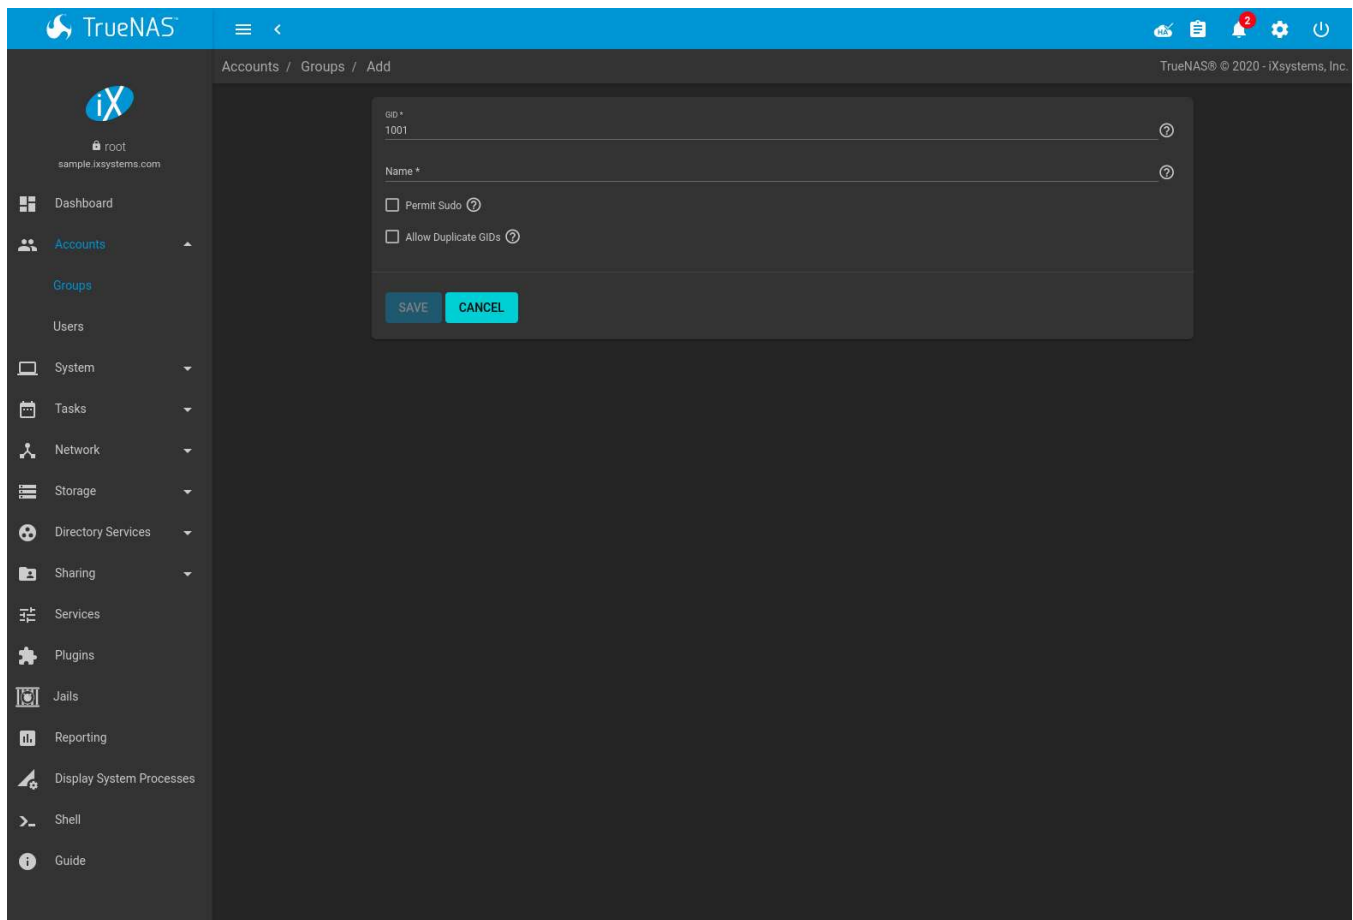


Fig. 4.2: Creating a New Group

Table 4.1: 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. This setting cannot be edited once the group is created.
Name	string	Enter an alphanumeric name for the new group. Group names cannot begin with a hyphen (-) or contain a space, tab, or these characters: , : + & # % ^ () ! @ ~ * ? < > = . \$ can only be used as the last character of the group name.
Permit Sudo	checkbox	Set to allow group members to use sudo (https://www.sudo.ws/). When using <code>sudo</code> , a user is prompted for their own password.
Allow Duplicate GIDs	checkbox	Not recommended. Allow more than one group to have the same group ID.

To change which users are members of a group, expand the group from the list and click *Members*. To add users to the group, select users in the left frame and click `->`. To remove users from the group, select users in the right frame and click `<-`. Click *SAVE* when finished changing the group members.

Figure 4.3, shows adding a user as a member of a group.

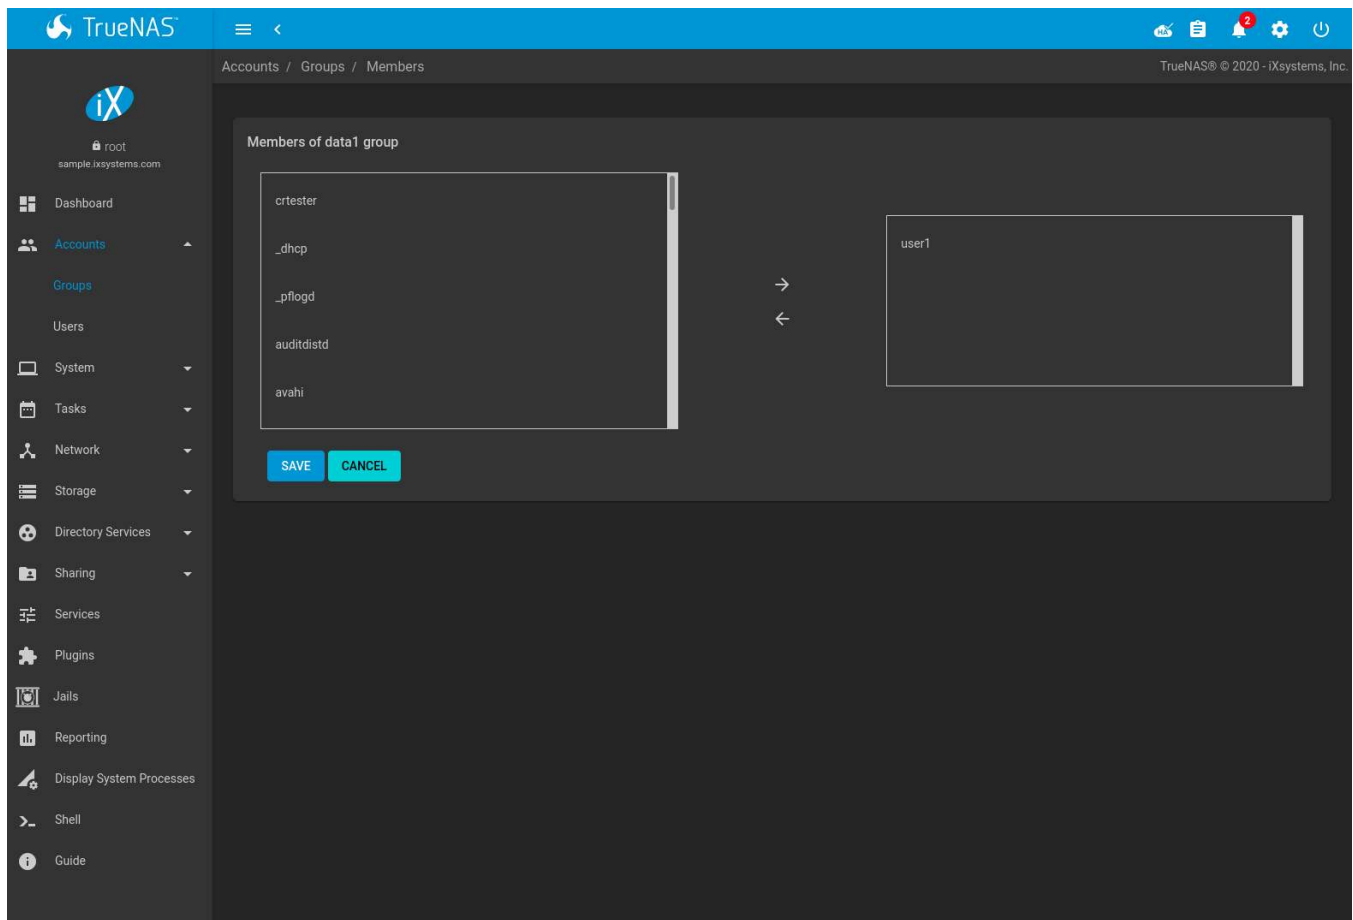


Fig. 4.3: Assigning a User to a Group

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.

4.2 Users

TrueNAS® supports users, groups, and permissions, allowing flexibility in configuring which users have access to the data stored on TrueNAS®. 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 TrueNAS®. 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 165).

Accounts → *Users* lists all system accounts installed with the TrueNAS® operating system, as shown in [Figure 4.4](#).

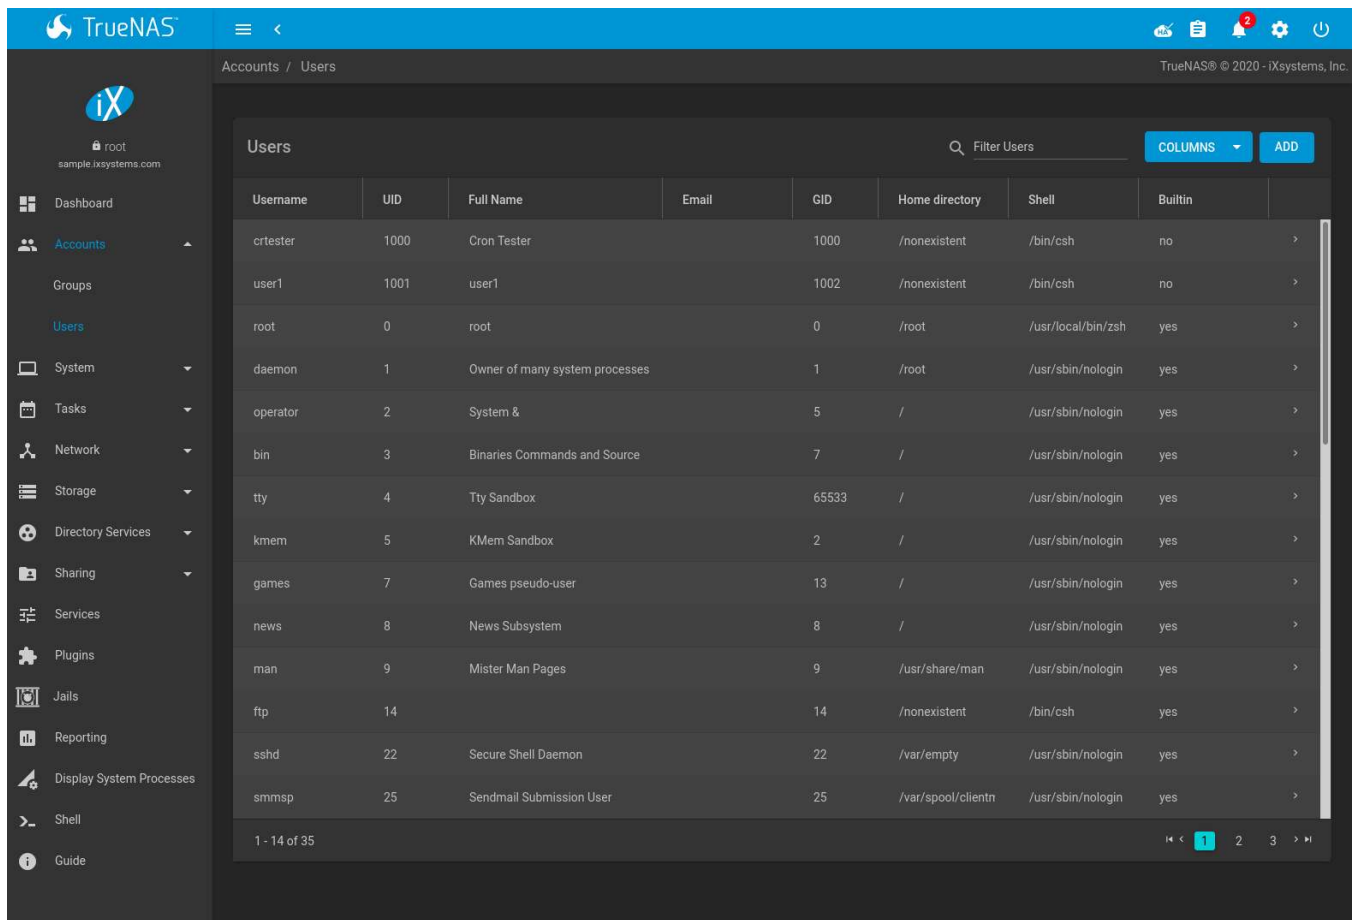


Fig. 4.4: Managing User Accounts

By default, each user entry displays the username, User ID (UID), whether the user is built into TrueNAS®, and full name. This table is adjustable by clicking *COLUMNS* and setting the desired columns.

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 TrueNAS® 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 4.5. Table 4.2 summarizes the options that are available when user accounts are created or modified.

Warning: When using *Active Directory* (page 165), Windows user passwords must be set from within Windows.

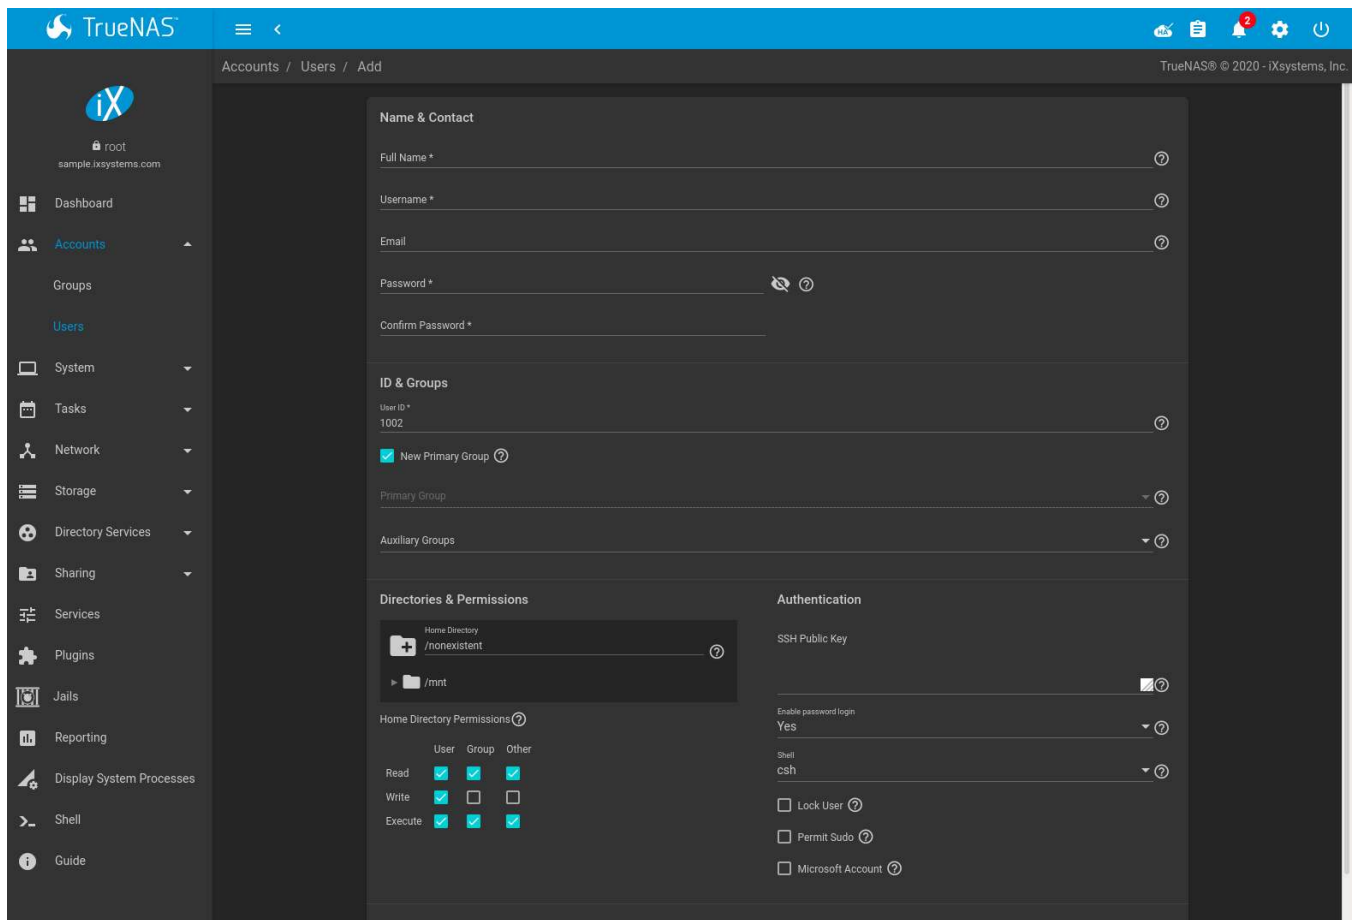



Fig. 4.5: Adding or Editing a User Account

Table 4.2: User Account Configuration

Setting	Value	Description
Username	string	Usernames can be up to 16 characters long. When using NIS or other legacy software with limited username lengths, keep usernames to eight characters or less for compatibility. Usernames cannot begin with a hyphen (-) or contain a space, tab, or these characters: , : + & # % ^ () ! @ ~ * ? < > = . \$ can only be used as the last character of the username.
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>Disable Password</i> is Yes. 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 4.2 – 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 not 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	Choose a path to the user's home directory. If the directory exists and matches the username, it is set as the user's home directory. When the path does not end with a subdirectory matching the username, a new subdirectory is created. The full path to the user's home directory is shown here when editing a user.
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!
Disable Password	drop-down	<i>Yes</i> : Disables the <i>Password</i> fields and removes the password from the account. The account cannot use password-based logins for services. For example, disabling the password prevents using account credentials to log in to an SMB share or open an SSH session on the system. The <i>Lock User</i> and <i>Permit Sudo</i> options are also removed. <i>No</i> : Requires adding a <i>Password</i> to the account. The account can use the saved <i>Password</i> to authenticate with password-based services.
Shell	drop-down menu	Select the shell to use for local and SSH logins. The <i>root</i> user shell is used for web interface <i>Shell</i> (page 302) sessions. See Table 4.3 for an overview of available shells.
Lock User	checkbox	Prevent the user from logging in or using password-based services until this option is unset. Locking an account is only possible when <i>Disable Password</i> is <i>No</i> and a <i>Password</i> has been created for the account.
Permit Sudo	checkbox	Give this user permission to use <code>sudo</code> (https://www.sudo.ws/). When using <code>sudo</code> , a user is prompted for their account <i>Password</i> .
Microsoft Account	checkbox	Set if the user is connecting from a Windows 8 or newer system or when using a Microsoft cloud service.

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

Table 4.3: Available Shells

Shell	Description
csh	C shell (https://en.wikipedia.org/wiki/C_shell)
sh	Bourne shell (https://en.wikipedia.org/wiki/Bourne_shell)
tcsh	Enhanced C shell (https://en.wikipedia.org/wiki/Tcsh)
bash	Bourne Again shell (https://en.wikipedia.org/wiki/Bash_%28Unix_shell%29)
ksh93	Korn shell (http://www.kornshell.com/)
mksh	mirBSD Korn shell (https://www.mirbsd.org/mksh.htm)
rbash	Restricted bash (http://www.gnu.org/software/bash/manual/html_node/The-Restricted-Shell.html)
rzsh	Restricted zsh (http://www.csse.uwa.edu.au/programming/linux/zsh-doc/zsh_14.html)

Continued on next page

Table 4.3 – continued from previous page

Shell	Description
scponly	Select scponly (https://github.com/scponly/scponly/wiki) to restrict the user's SSH usage to only the <code>scp</code> and <code>sftp</code> commands.
zsh	Z shell (http://www.zsh.org/)
git-shell	restricted git shell (https://git-scm.com/docs/git-shell)
nologin	Use when creating a system account or to create a user account that can authenticate with shares but which cannot login to the FreeNAS system using <code>ssh</code> .

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 web interface contains these entries:

- [General](#) (page 32) configures general settings such as HTTPS access, the language, and the timezone
- [NTP Servers](#) (page 35) adds, edits, and deletes Network Time Protocol servers
- [Boot](#) (page 37) creates, renames, and deletes boot environments. It also shows the condition of the Boot Pool.
- [Advanced](#) (page 39) configures advanced settings such as the serial console, swap space, and console messages
- [View Enclosure](#) (page 45): view status of disk enclosures.
- [Email](#) (page 47) configures the email address to receive notifications
- [System Dataset](#) (page 48) configures the location where logs and reporting graphs are stored
- [Alert Services](#) (page 50) configures services used to notify the administrator about system events.
- [Alert Settings](#) (page 52) lists the available [Alert](#) (page 305) conditions and provides configuration of the notification frequency for each alert.
- [Cloud Credentials](#) (page 53) is used to enter connection credentials for remote cloud service providers
- [SSH Connections](#) (page 57) manages connecting to a remote system with SSH.
- [SSH Keypairs](#) (page 61) manages all private and public SSH key pairs.
- [Tunables](#) (page 61) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- [Update](#) (page 64) performs upgrades and checks for system updates
- [CAs](#) (page 72): import or create internal or intermediate CAs (Certificate Authorities)
- [Certificates](#) (page 76): import existing certificates or create self-signed certificates.
- [Failover](#) (page 81): manage High Availability.
- [Support](#) (page 84): view licensing information or create a support ticket.
- [Proactive Support](#) (page 85): enable and configure automatic proactive support (Silver or Gold support coverage only).

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

5.1 General

System → *General* contains options for configuring the web interface and other basic system settings.

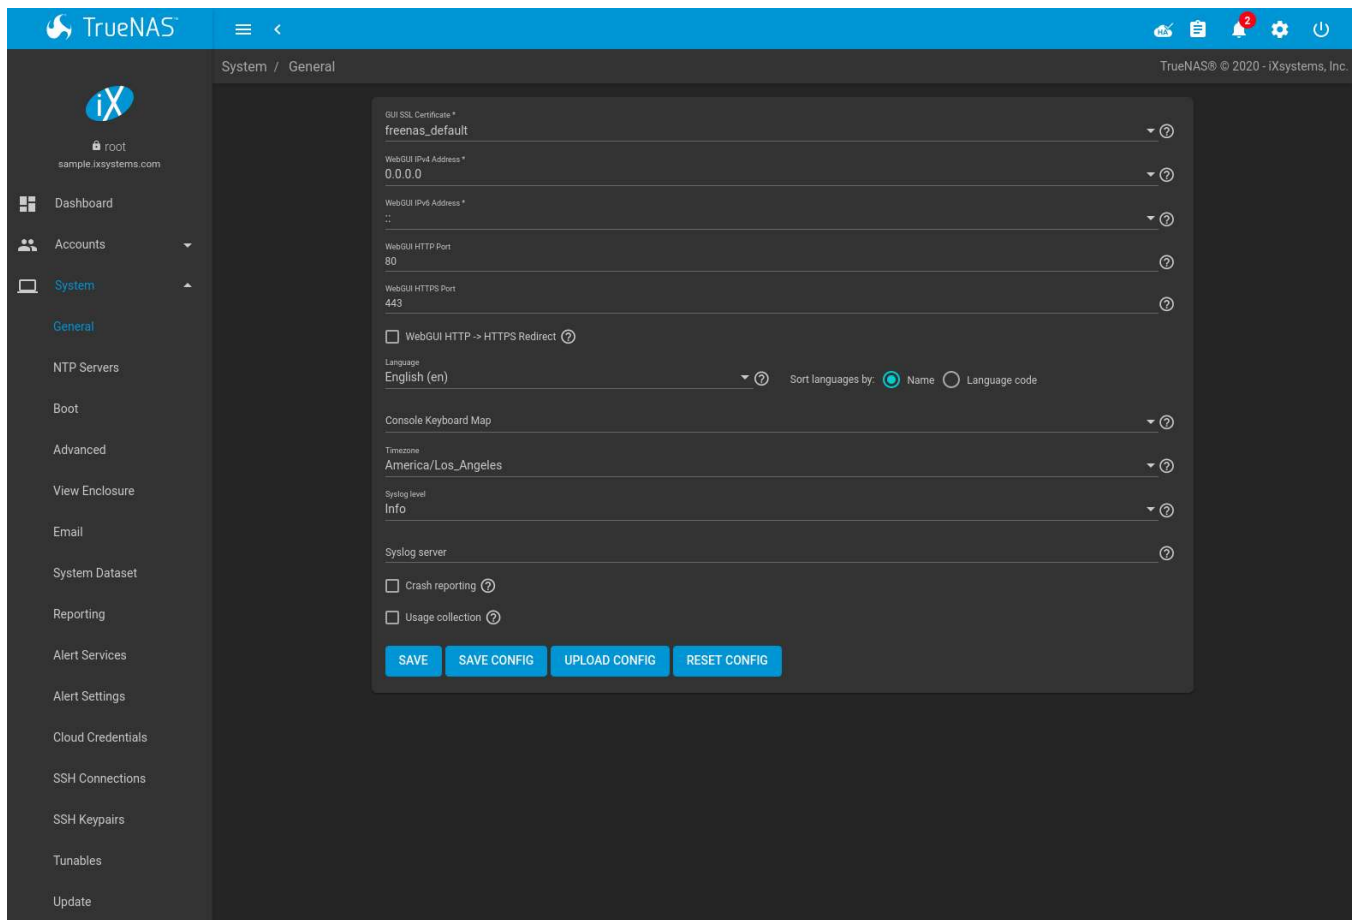


Fig. 5.1: General System Options

Table 5.1: General Configuration Settings

Setting	Value	Description
GUI SSL Certificate	drop-down menu	The system uses a self-signed certificate (page 76) to enable encrypted web interface connections. To change the default certificate, select a different created or imported certificate.
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 the wildcard address of <i>0.0.0.0</i> (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 might 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 to access the web interface over HTTPS.

Continued on next page

Table 5.1 – continued from previous page

Setting	Value	Description
WebGUI HTTP -> HTTPS Redirect	checkbox	Redirect <i>HTTP</i> connections to <i>HTTPS</i> . A <i>GUI SSL Certificate</i> is required for <i>HTTPS</i> . Activating this also sets the HTTP Strict Transport Security (HSTS) (https://en.wikipedia.org/wiki/HTTP_Strict_Transport_Security) maximum age to 31536000 seconds (one year). This means that after a browser connects to the TrueNAS® web interface for the first time, the browser continues to use <i>HTTPS</i> and renews this setting every year.
Language	combo box	Select a language from the drop-down menu. The list can be sorted by <i>Name</i> or Language code (https://en.wikipedia.org/wiki/List_of_ISO_639-1_codes). View the translated status of a language in the webui GitHub repository (https://github.com/freenas/webui/tree/master/src/assets/i18n).
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	Remote syslog server DNS hostname or IP address. Nonstandard port numbers can be used by adding a colon and the port number to the hostname, like <code>mysyslogserver:1928</code> . Log entries are written to local logs and sent to the remote syslog server.
Crash reporting	checkbox	Send failed HTTP request data which can include client and server IP addresses, failed method call tracebacks, and middleware log file contents to iXsystems.
Usage Collection	checkbox	Enable sending anonymous usage statistics to iXsystems.

After making any changes, click *SAVE*. Changes to any of the *GUI* fields can interrupt web interface connectivity while the new settings are applied.

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 web interface. Saving the configuration after making any configuration changes is highly recommended. TrueNAS® 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 244) keys are not stored in the configuration database and must be backed up separately. System host keys are files with names beginning with `ssh_host_` in `/usr/local/etc/ssh/`. The root user keys are stored in `/root/.ssh`.

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 operating system device. When *Save Config* is chosen, a dialog gives two options. *Export Password Secret Seed* includes passwords in the configuration file which allows the configuration file to be restored to a different operating system device where the decryption seed is not already present. Configuration backups containing the seed must be physically

secured to prevent decryption of passwords and unauthorized access.

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

Export Pool Encryption Keys includes the encryption keys of encrypted pools in the configuration file. The encryption keys are restored if the configuration file is uploaded to the system using *UPLOAD CONFIG*.

- *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.

5.2 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, TrueNAS® is pre-configured to use three public NTP servers. If the network is using a directory service, ensure that the TrueNAS® 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 TrueNAS® system.

Click *System* → *NTP Servers* and *ADD* to add an NTP server. [Figure 5.2](#) shows the configuration options. [Table 5.2](#) summarizes the options available when adding or editing an NTP server. [ntp.conf\(5\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=ntp.conf>) explains these options in more detail.

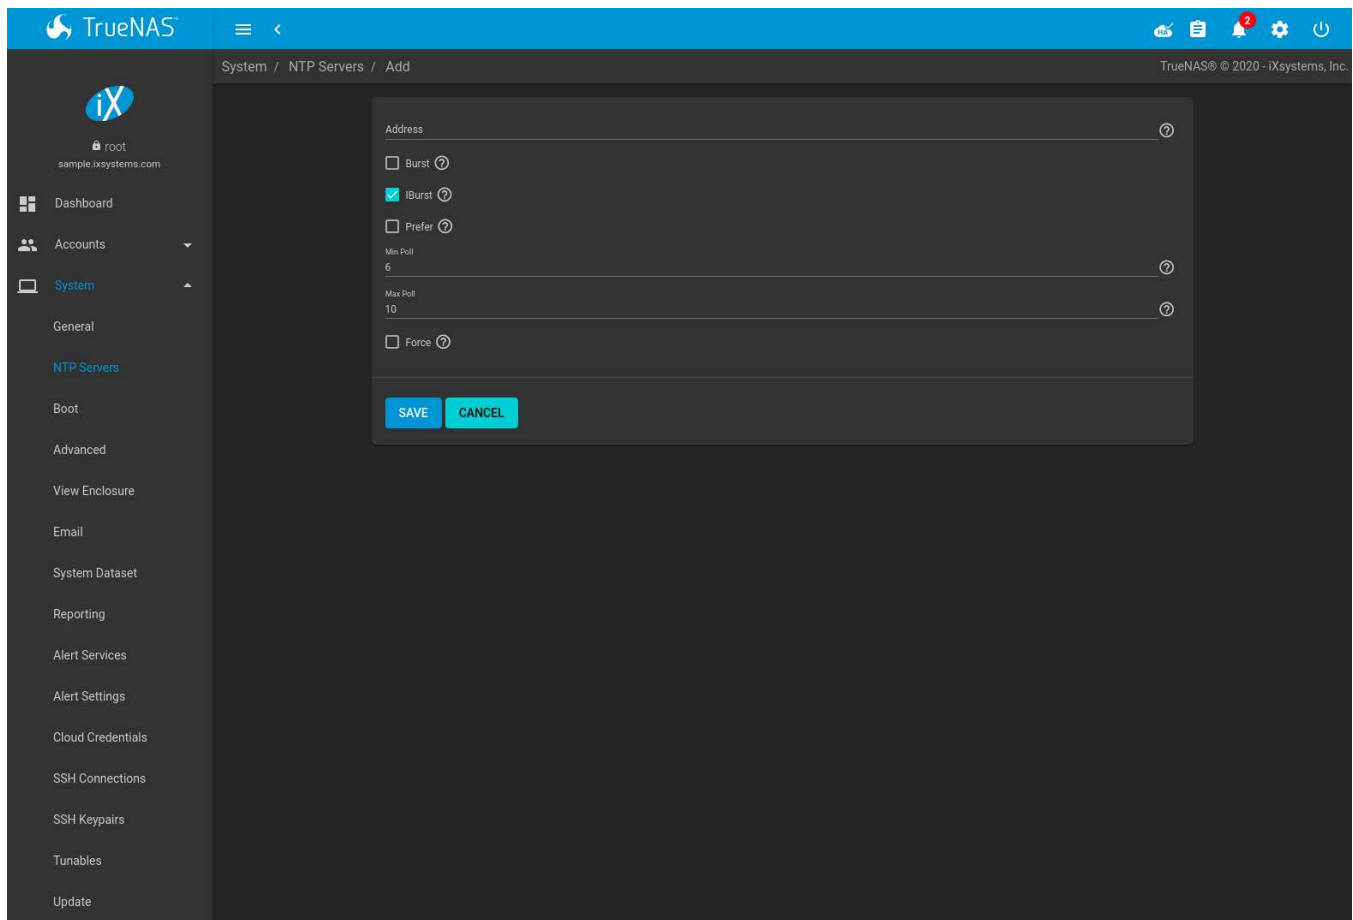


Fig. 5.2: Add an NTP Server

Table 5.2: NTP Servers Configuration Options

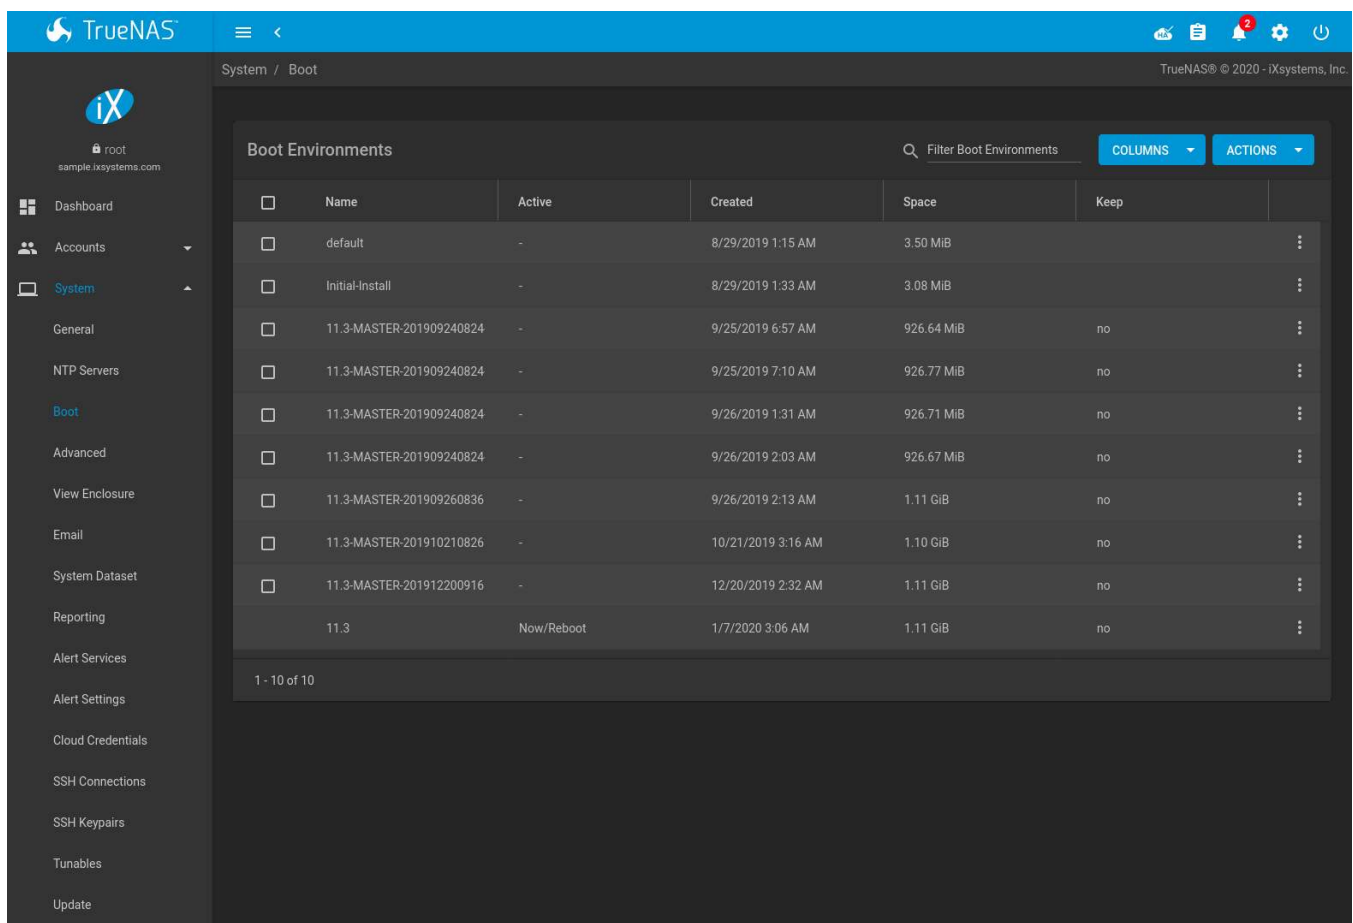
Setting	Value	Description
Address	string	Enter the hostname or IP address of the NTP server.
Burst	checkbox	Recommended when <i>Max. Poll</i> is greater than 10. Only use on personal servers. Do not use with a public NTP server.
IBurst	checkbox	Speed up the initial synchronization, taking seconds rather than minutes.
Prefer	checkbox	This option is only recommended for highly accurate NTP servers, such as those with time monitoring hardware.
Min Poll	integer	The minimum polling interval, in seconds, as a power of 2. For example, 6 means 2^6 , or 64 seconds. The default is 6, minimum value is 4.
Max Poll	integer	The maximum polling interval, in seconds, as a power of 2. For example, 10 means 2^{10} , or 1,024 seconds. The default is 10, maximum value is 17.
Force	checkbox	Force the addition of the NTP server, even if it is currently unreachable.

5.3 Boot

TrueNAS® 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.

Note: Boot environments are separate from the configuration database. Boot environments are a snapshot of the *operating system* at a specified time. When a TrueNAS® 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 the instructions in [System -> General](#) (page 32).

The example shown in [Figure 5.3](#), includes the two boot environments that are created when TrueNAS® is installed. The *Initial-Install* boot environment can be booted into if the system needs to be returned to a non-configured version of the installation.




<input type="checkbox"/>	Name	Active	Created	Space	Keep	
<input type="checkbox"/>	default	-	8/29/2019 1:15 AM	3.50 MiB		⋮
<input type="checkbox"/>	Initial-Install	-	8/29/2019 1:33 AM	3.08 MiB		⋮
<input type="checkbox"/>	11.3-MASTER-201909240824	-	9/25/2019 6:57 AM	926.64 MiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201909240824	-	9/25/2019 7:10 AM	926.77 MiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201909240824	-	9/26/2019 1:31 AM	926.71 MiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201909240824	-	9/26/2019 2:03 AM	926.67 MiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201909260836	-	9/26/2019 2:13 AM	1.11 GiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201910210826	-	10/21/2019 3:16 AM	1.10 GiB	no	⋮
<input type="checkbox"/>	11.3-MASTER-201912200916	-	12/20/2019 2:32 AM	1.11 GiB	no	⋮
	11.3	Now/Reboot	1/7/2020 3:06 AM	1.11 GiB	no	⋮


1 - 10 of 10

Fig. 5.3: 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. Alphanumeric characters, dashes (-), underscores (_), and periods (.) are allowed.
- **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 access actions specific to that entry:

- **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.
- **Clone:** makes a new boot environment from the selected boot environment. When prompted for the name of the clone, alphanumeric characters, dashes (-), underscores (_), and periods (.) are allowed.
- **Rename:** used to change the name of the boot environment. Alphanumeric characters, dashes (-), underscores (_), and periods (.) are allowed.
- **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. 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.
- **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.

Click *ACTIONS* to:

- **Add:** make 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*.
- **Stats/Settings:** display statistics for the operating system device: condition, total and used size, and date and time of the last scrub. By default, the operating system device is scrubbed every 7 days. To change the default, input a different number in the *Automatic scrub interval (in days)* field and click *UPDATE INTERVAL*.
- **Boot Pool Status:** display the status of each device in the operating system device, including any read, write, or checksum errors.
- **Scrub Boot Pool:** perform a manual scrub of the operating system device.

5.3.1 Operating System Device Mirroring

System → *Boot* → *Boot Pool Status* is used to manage the devices comprising the operating system device. An example is seen in [Figure 5.4](#).

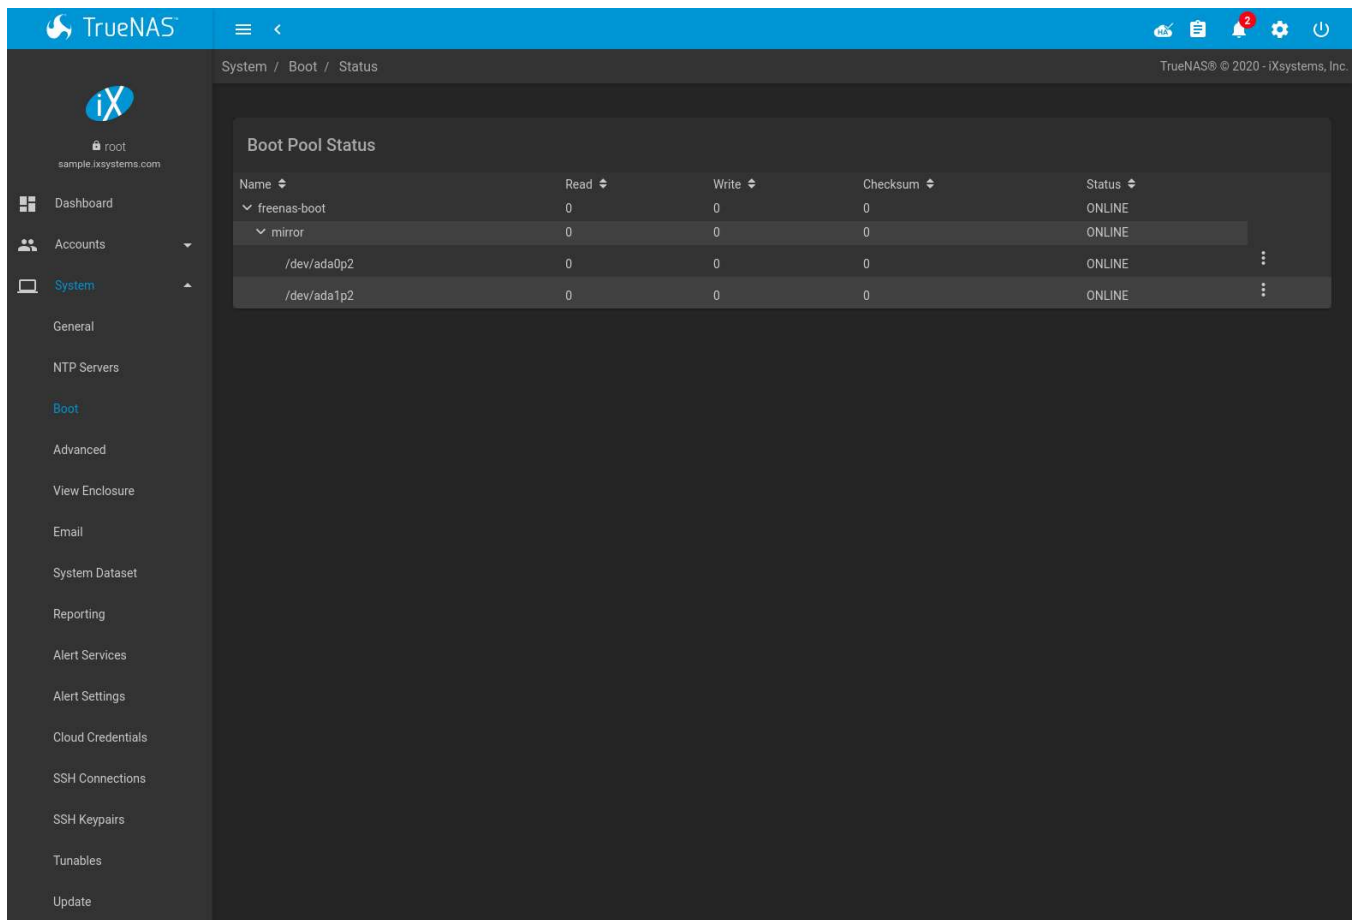



Fig. 5.4: Viewing the Status of the Operating System Device

TrueNAS® supports 2-device mirrors for the operating system device. In a mirrored configuration, a failed device can be detached and replaced.

Click  (Options) on a device entry to access actions specific to that device:

- **Attach:** use to add a second device to create a mirrored operating system device. If another device is available, it appears in the *Member disk* drop-down menu. Select the desired device. The *Use all disk space* option controls the capacity made available to the operating system device. By default, the new device is partitioned to the same size as the existing device. When *Use all disk space* is enabled, the entire capacity of the new device is used. If the original operating system device fails and is detached, 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.
- **Detach:** remove the failed device from the mirror so that it can be replaced.
- **Replace:** once the failed device has been detached, select the new replacement device from the *Member disk* drop-down menu to rebuild the mirror.

5.4 Advanced

System → *Advanced* is shown in [Figure 5.5](#). The configurable settings are summarized in [Table 5.3](#).

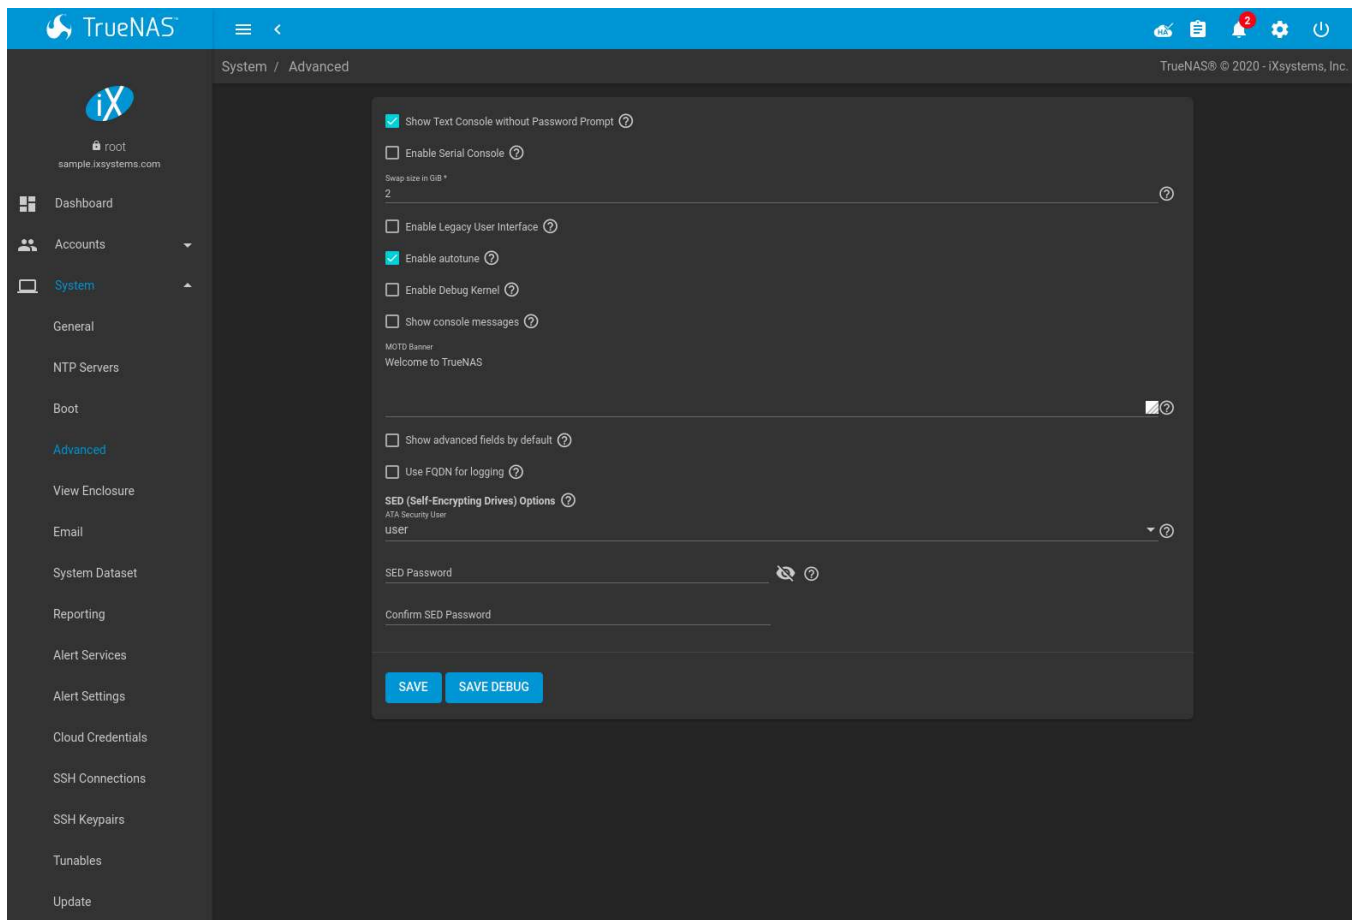


Fig. 5.5: Advanced Screen

Table 5.3: Advanced Configuration Settings

Setting	Value	Description
Show Text Console without Password Prompt	checkbox	Set for the text console to be available without entering a password.
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.
Enable Legacy User Interface	checkbox	WARNING: The legacy user interface is deprecated. All management should be performed through the new user interface. Shows legacy UI login buttons on the web interface log in screen and Settings (page 20) menu. These buttons allow switching to the interface that was available with TrueNAS® 11.2 and earlier.
Enable autotune	checkbox	Enable the Autotune (page 41) 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 5.3 – continued from previous page

Setting	Value	Description
Show console messages	checkbox	Display console messages from <code>/var/log/console.log</code> in real time at bottom of browser window. Click the console to bring up a scrollable screen. Set the <i>Stop refresh</i> option in the scrollable screen to pause updates. Unset to continue watching messages as they occur. When this option is set, a button to show the console log appears on busy spinner dialogs.
MOTD banner	string	This message is shown when a user logs in with SSH.
Show advanced fields by default	checkbox	Show <i>Advanced Mode</i> fields by default.
Use FQDN for logging	checkbox	Include the Fully-Qualified Domain Name (FQDN) in logs to precisely identify systems with similar hostnames.
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 <i>Self-Encrypting Drives</i> (page 41).
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 `.tar` file.

5.4.1 Autotune

TrueNAS® provides an autotune script which optimizes the system. The *Enable autotune* option in *System* → *Advanced* is enabled by default, so this script runs automatically. Leaving autotune enabled is recommended unless advised otherwise by an iXsystems support engineer.

If the autotune script adjusts any settings, the changed values appear in *System* → *Tunables*. While these values can be modified and overridden, speak to a support engineer first. Manual changes can have a negative impact on system performance. Note that deleting tunables that were created by autotune only affects the current session, as autotune-set tunables are recreated at boot.

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

5.4.2 Self-Encrypting Drives

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

These SED specifications are supported:

- Legacy interface for older ATA devices. **Not recommended for security-critical environments**
- **TCG Opal 1** (https://trustedcomputinggroup.org/wp-content/uploads/Opal_SSC_1.00_rev3.00-Final.pdf) legacy specification
- **TCG OPAL 2** (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Opal_SSC_v2.01_rev1.00.pdf) standard for newer consumer-grade devices
- **TCG Opalite** (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Opalite_SSC_FAQ.pdf) is a reduced form of OPAL 2
- **TCG Pyrite Version 1** (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Pyrite_SSC_v1.00_r1.00.pdf) and **Version 2** (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Pyrite_SSC_v2.00_r1.00_PUB.pdf) are similar to Opalite, but hardware

encryption is removed. Pyrite provides a logical equivalent of the legacy ATA security for non-ATA devices. Only the drive firmware is used to protect the device.

Danger: Pyrite Version 1 SEDs do not have PSID support and **can become unusable if the password is lost.**

- **TCG Enterprise** (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-SSC_Enterprise-v1.01_r1.00.pdf) is designed for systems with many data disks. These SEDs do not have the functionality to be unlocked before the operating system boots.

See this Trusted Computing Group® and NVM Express® [joint white paper](https://nvmexpress.org/wp-content/uploads/TCGandNVMe_Joint_White_Paper-TCG_Storage_Opal_and_NVMe_FINAL.pdf) (https://nvmexpress.org/wp-content/uploads/TCGandNVMe_Joint_White_Paper-TCG_Storage_Opal_and_NVMe_FINAL.pdf) for more details about these specifications.

TrueNAS® implements the security capabilities of **camcontrol** (<https://www.freebsd.org/cgi/man.cgi?query=camcontrol>) for legacy devices and **sedutil-cli** (<https://www.mankier.com/8/sedutil-cli>) for TCG devices. When managing a SED from the command line, it is recommended to use the **sedhelper** wrapper script for **sedutil-cli** to ease SED administration and unlock the full capabilities of the device. Examples of using these commands to identify and deploy SEDs are provided below.

A SED can be configured before or after assigning the device to a [pool](#) (page 130).

By default, SEDs are not locked until the administrator takes ownership of them. Ownership is taken by explicitly configuring a global or per-device password in the TrueNAS® web interface and adding the password to the SEDs. Adding SED passwords to TrueNAS® also allows TrueNAS® to automatically unlock SEDs.

A password-protected SED protects the data stored on the device when the device is physically removed from the TrueNAS® system. This allows secure disposal of the device without having to first wipe the contents. Repurposing a SED on another system requires the SED password.

5.4.2.1 Deploying SEDs

Run **sedutil-cli --scan** in the [Shell](#) (page 302) to detect and list devices. The second column of the results identifies the drive type:

- **no** indicates a non-SED device
- **1** indicates a legacy TCG OPAL 1 device
- **2** indicates a modern TCG OPAL 2 device
- **L** indicates a TCG Opalite device
- **p** indicates a TCG Pyrite 1 device
- **P** indicates a TCG Pyrite 2 device
- **E** indicates a TCG Enterprise device

Example:

```
root@truenas1:~ # sedutil-cli --scan
Scanning for Opal compliant disks
/dev/ada0 No 32GB SATA Flash Drive SFDK003L
/dev/ada1 No 32GB SATA Flash Drive SFDK003L
/dev/da0 No HGST HUS726020AL4210 A7J0
/dev/da1 No HGST HUS726020AL4210 A7J0
/dev/da10 E WDC WUSTR1519ASS201 B925
/dev/da11 E WDC WUSTR1519ASS201 B925
```

TrueNAS® supports setting a global password for all detected SEDs or setting individual passwords for each SED. Using a global password for all SEDs is strongly recommended to simplify deployment and avoid maintaining separate passwords for each SED.

Setting a global password for SEDs

Go to *System* → *Advanced* → *SED Password* and enter the password. **Record this password and store it in a safe place!**

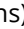
Now the SEDs must be configured with this password. Go to the *Shell* (page 302) and enter `sedhelper setup password`, where *password* is the global password entered in *System* → *Advanced* → *SED Password*.

`sedhelper` ensures that all detected SEDs are properly configured to use the provided password:

```
root@truenas1:~ # sedhelper setup abcd1234
da9                [OK]
da10               [OK]
da11               [OK]
```

Rerun `sedhelper setup password` every time a new SED is placed in the system to apply the global password to the new SED.

Creating separate passwords for each SED

Go to *Storage* → *Disks*. Click  (Options) for the confirmed SED, then *Edit*. Enter and confirm the password in the *SED Password* and *Confirm SED Password* fields.

The *Storage* → *Disks* screen shows which disks have a configured SED password. The *SED Password* column shows a mark when the disk has a password. Disks that are not a SED or are unlocked using the global password are not marked in this column.

The SED must be configured to use the new password. Go to the *Shell* (page 302) and enter `sedhelper setup -disk da1 password`, where *da1* is the SED to configure and *password* is the created password from *Storage* → *Disks* → *Edit Disks* → *SED Password*.

This process must be repeated for each SED and any SEDs added to the system in the future.

Danger: Remember SED passwords! If the SED password is lost, SEDs cannot be unlocked and their data is unavailable. Always record SED passwords whenever they are configured or modified and store them in a secure place!

5.4.2.2 Check SED Functionality

When SED devices are detected during system boot, TrueNAS® checks for configured global and device-specific passwords.

Unlocking SEDs allows a pool to contain a mix of SED and non-SED devices. Devices with individual passwords are unlocked with their password. Devices without a device-specific password are unlocked using the global password.

To verify SED locking is working correctly, go to the *Shell* (page 302). Enter `sedutil-cli --listLockingRange 0 password dev/da1`, where *da1* is the SED and *password* is the global or individual password for that SED. The command returns `ReadLockEnabled: 1, WriteLockEnabled: 1, and LockOnReset: 1` for drives with locking enabled:

```
root@truenas1:~ # sedutil-cli --listLockingRange 0 abcd1234 /dev/da9
Band[0]:
  Name:                Global_Range
  CommonName:          Locking
  RangeStart:          0
  RangeLength:         0
  ReadLockEnabled:     1
  WriteLockEnabled:    1
```

```
ReadLocked:      0
WriteLocked:     0
LockOnReset:     1
```

5.4.2.3 Managing SED Passwords and Data

This section contains command line instructions to manage SED passwords and data. The command used is `sedutil-cli(8)` (<https://www.mankier.com/8/sedutil-cli>). Most SEDs are TCG-E (Enterprise) or TCG-Opal ([Opal v2.0](https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Opal_SSC_v2.01_rev1.00.pdf) (https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage-Opal_SSC_v2.01_rev1.00.pdf)). Commands are different for the different drive types, so the first step is identifying which type is being used.

Warning: These commands can be destructive to data and passwords. Keep backups and use the commands with caution.

Check SED version on a single drive, `/dev/da0` in this example:

```
root@truenas:~ # sedutil-cli --isValidSED /dev/da0
/dev/da0 SED --E--- Micron_5N/A U402
```

All connected disks can be checked at once:

```
root@truenas:~ # sedutil-cli --scan
Scanning for Opal compliant disks
/dev/ada0 No 32GB SATA Flash Drive SFDK003L
/dev/ada1 No 32GB SATA Flash Drive SFDK003L
/dev/da0 E Micron_5N/A U402
/dev/da1 E Micron_5N/A U402
/dev/da12 E SEAGATE XS3840TE70014 0103
/dev/da13 E SEAGATE XS3840TE70014 0103
/dev/da14 E SEAGATE XS3840TE70014 0103
/dev/da2 E Micron_5N/A U402
/dev/da3 E Micron_5N/A U402
/dev/da4 E Micron_5N/A U402
/dev/da5 E Micron_5N/A U402
/dev/da6 E Micron_5N/A U402
/dev/da9 E Micron_5N/A U402
No more disks present ending scan
root@truenas:~ #
```

TCG-Opal Instructions

Reset the password without losing data: `sedutil-cli --revertNoErase oldpassword /dev/device`

Use **both** of these commands to change the password without destroying data:

```
sedutil-cli --setSIDPassword oldpassword newpassword /dev/device
sedutil-cli --setPassword oldpassword Admin1 newpassword /dev/device
```

Wipe data and reset password to default MSID: `sedutil-cli --revertPer oldpassword /dev/device`

Wipe data and reset password using the PSID: `sedutil-cli --yesIreallywanttoERASEALLmydatausingthePSID PSINODASHED /dev/device` where *PSINODASHED* is the PSID located on the physical drive with no dashes (-).

TCG-E Instructions

Use **all** of these commands to reset the password without losing data:

```
sedutil-cli --setSIDPassword oldpassword "" /dev/device
sedutil-cli --setPassword oldpassword EraseMaster "" /dev/device
sedutil-cli --setPassword oldpassword BandMaster0 "" /dev/device
sedutil-cli --setPassword oldpassword BandMaster1 "" /dev/device
```

Use **all** of these commands to change the password without destroying data:

```
sedutil-cli --setSIDPassword oldpassword newpassword /dev/device
sedutil-cli --setPassword oldpassword EraseMaster newpassword /dev/device
sedutil-cli --setPassword oldpassword BandMaster0 newpassword /dev/device
sedutil-cli --setPassword oldpassword BandMaster1 newpassword /dev/device
```

Wipe data and reset password to default MSID:

```
sedutil-cli --eraseLockingRange 0 password /dev/<device>
sedutil-cli --setSIDPassword oldpassword "" /dev/<device>
sedutil-cli --setPassword oldpassword EraseMaster "" /dev/<device>
```

Wipe data and reset password using the PSID: `sedutil-cli --yesIreallywanttoERASEALLmydatausingthePSID PSINODASHED /dev/device` where *PSINODASHED* is the PSID located on the physical drive with no dashes (-).

5.5 View Enclosure

Click *System* → *View Enclosure* to display the status of connected disks and hardware.

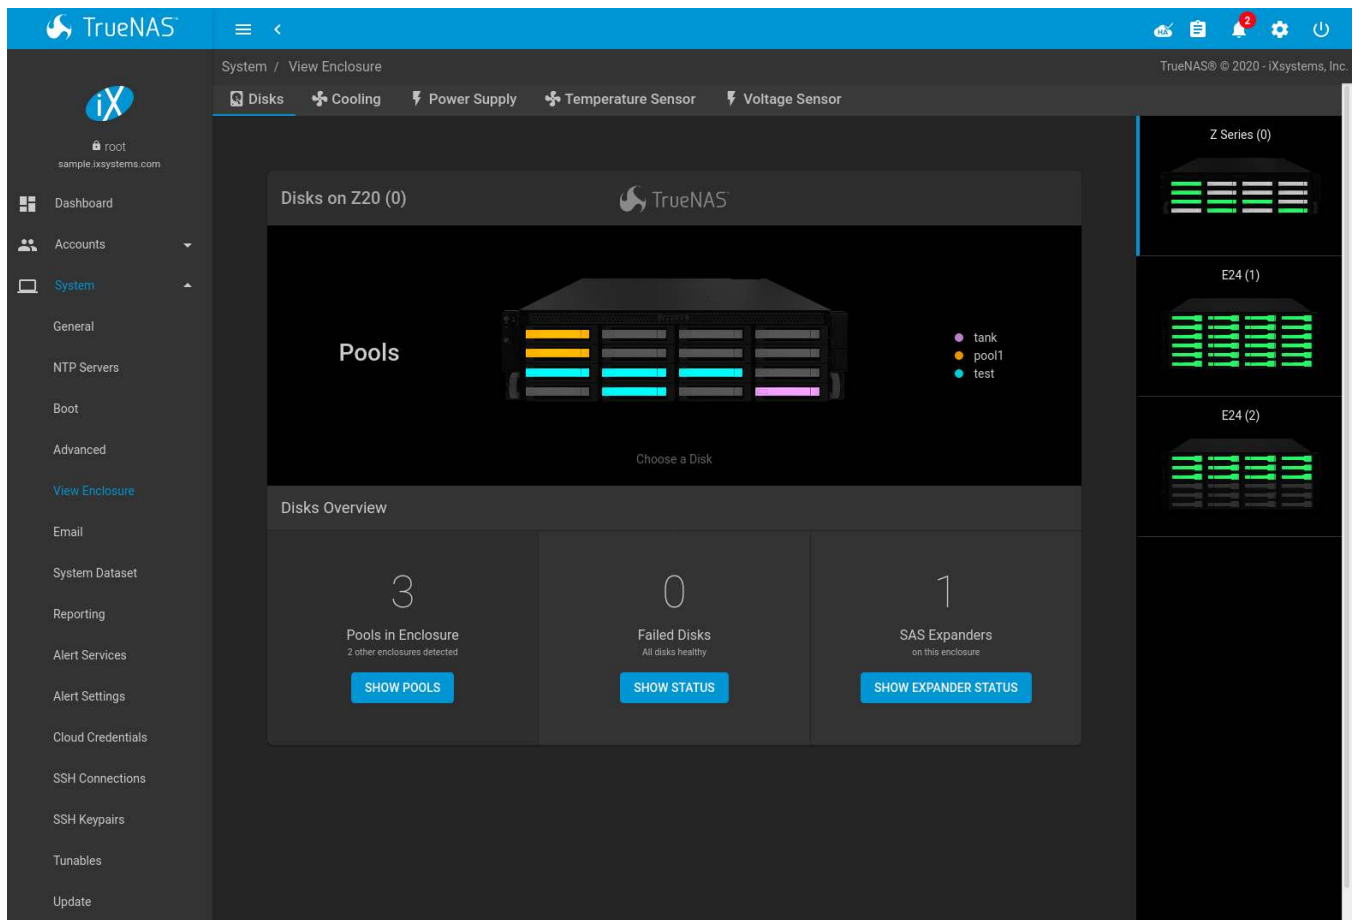


Fig. 5.6: View Enclosure

Detected TrueNAS® hardware is added to a column on the right side of the screen. Click an enclosure to show details about that hardware.

The screen is divided into different tabs. These tabs reflect the sensors that are active in the chosen hardware.

Disks shows a graphic representation of the TrueNAS® hardware and details about connected disks. Click any disk slot to see specific details about the disk like the FreeBSD device name, vdev assignment and function, serial number, and current drive settings. The *IDENTIFY DRIVE* button flashes the identification LED for the chosen drive.

The *Disks Overview* shows statistics about the enclosure pools, status, and detected expanders. There are options to show more details about pools in the enclosure, disk status, and expansion shelf status. Clicking any of the buttons changes the graphic to show the requested details.

Cooling has an entry for each fan with status and RPM.

Enclosure Services Controller Electronics shows the enclosure status.

Power Supply shows the status of each power supply.

SAS Connector shows the status of the expansion shelf.

Temperature Sensor shows the current temperature of each expansion shelf and the disk chassis.

Voltage Sensor shows the current voltage for each sensor, VCCP, and VCC.

5.6 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 305) events are also emailed to the *root* user account. Problems with *Scrub Tasks* (page 111) are reported separately in an email sent at 03:00AM.

Note: *S.M.A.R.T.* (page 238) reports are mailed separately to the address configured in that service.

The administrator typically does not read email directly on the TrueNAS® 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 5.7.

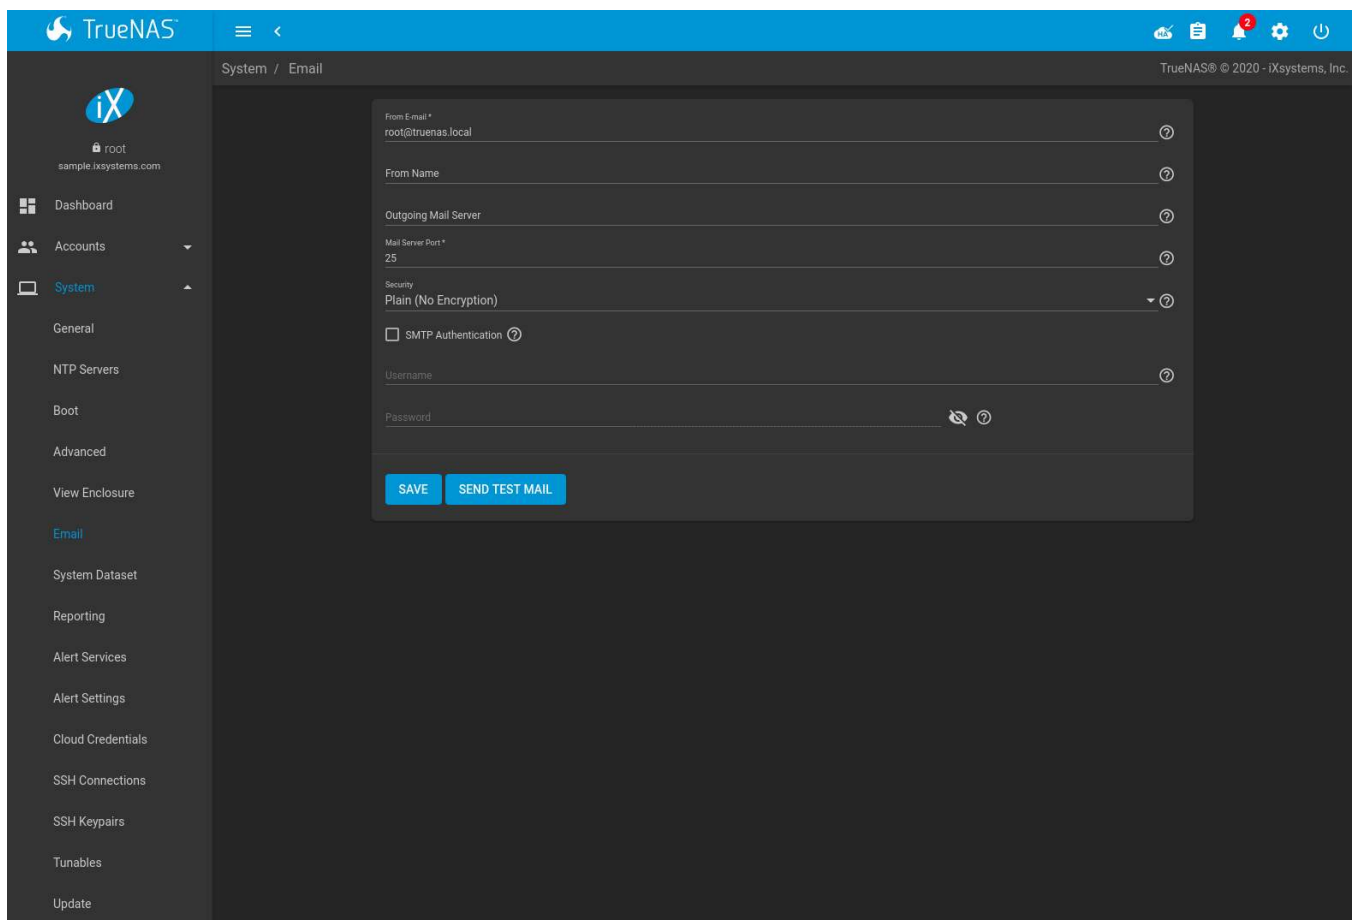


Fig. 5.7: Email Screen

Table 5.4: Email Configuration Settings

Setting	Value	Description
From E-mail	string	The envelope From address shown in the email. This can be set to make filtering mail on the receiving system easier.
From Name	string	The friendly name to show in front of the sending email address.
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 (No Encryption)</i> , <i>SSL (Implicit TLS)</i> , or <i>TLS (STARTTLS)</i> .
SMTP Authentication	checkbox	Enable or disable SMTP AUTH (https://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN SASL. Setting this enables the required <i>Username</i> and optional <i>Password</i> fields.
Username	string	Enter the SMTP username when the SMTP server requires authentication.
Password	string	Enter the SMTP account password if needed for authentication. Only plain text characters (7-bit ASCII) are allowed in passwords. UTF or composed characters are not allowed.

Click the *SEND TEST 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/>).

5.7 System Dataset

System → *System Dataset*, shown in [Figure 5.8](#), is used to select the pool which contains the persistent system dataset. The system dataset stores debugging core files, [encryption keys](#) (page 135) for encrypted pools, and Samba4 metadata such as the user/group cache and share level permissions.

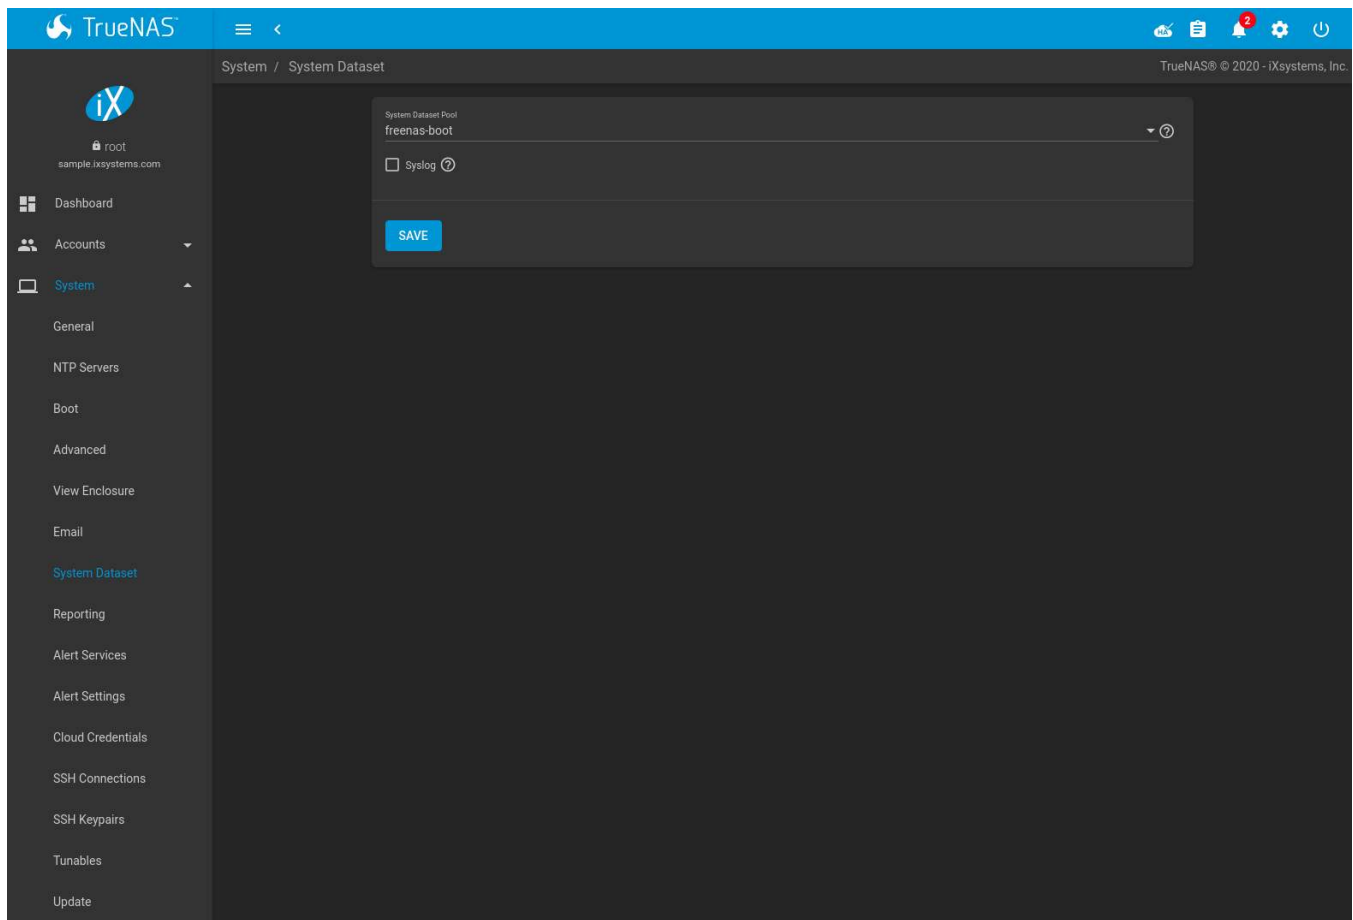


Fig. 5.8: System Dataset Screen

Use the *System Dataset Pool* drop-down menu to select the volume (pool) to contain the system dataset. The system dataset can be moved to unencrypted volumes (pools) or encrypted volumes which do not have passphrases. If the system dataset is moved to an encrypted volume, that volume is no longer allowed to be locked or have a passphrase set.

Moving the system dataset also requires rebooting the standby TrueNAS controller for *High Availability* (page 81) TrueNAS® systems and restarting the *SMB* (page 240) service. A dialog warns that the SMB service must be restarted, causing a temporary outage of any active SMB connections.

System logs can also be stored on the system dataset. Storing this information on the system dataset is recommended when large amounts of data is being generated and the system has limited memory or a limited capacity operating system device.

Set *Syslog* to store system logs on the system dataset. Leave unset to store system logs in `/var` on the operating system device.

Click *SAVE* to save changes.

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

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

5.8 Reporting

This section contains settings to customize some of the reporting tools. These settings are described in [Table 5.5](#)

Table 5.5: Reporting Settings

Setting	Value	Description
Report CPU usage in percent	checkbox	Report CPU usage in percent instead of units of kernel time.
Remote Graphite Server Host-name	string	Hostname or IP address of a remote Graphite (http://graphiteapp.org/) server.
Graph Age in Months	integer	Maximum time a graph is stored in months (allowed values are 1 - 60). Changing this value causes the <i>Confirm RRD Destroy</i> dialog to appear. Changes do not take effect until the existing reporting database is destroyed.
Number of Graph Points	integer	Number of points for each hourly, daily, weekly, monthly, or yearly graph (allowed values are 1 - 4096). Changing this value causes the <i>Confirm RRD Destroy</i> checkbox to appear. Changes do not take effect until the existing reporting database is destroyed.

Changes to [Reporting settings](#) (page 50) clear the report history. To keep history with the old settings, cancel the warning dialog. Click *RESET TO DEFAULTS* to restore the original settings.

5.9 Alert Services

TrueNAS® can use a number of methods to notify the administrator of system events that require attention. These events are system [Alerts](#) (page 305).

Available alert services:

- [AWS-SNS](https://aws.amazon.com/sns/) (<https://aws.amazon.com/sns/>)
- E-mail
- [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 5.9](#).

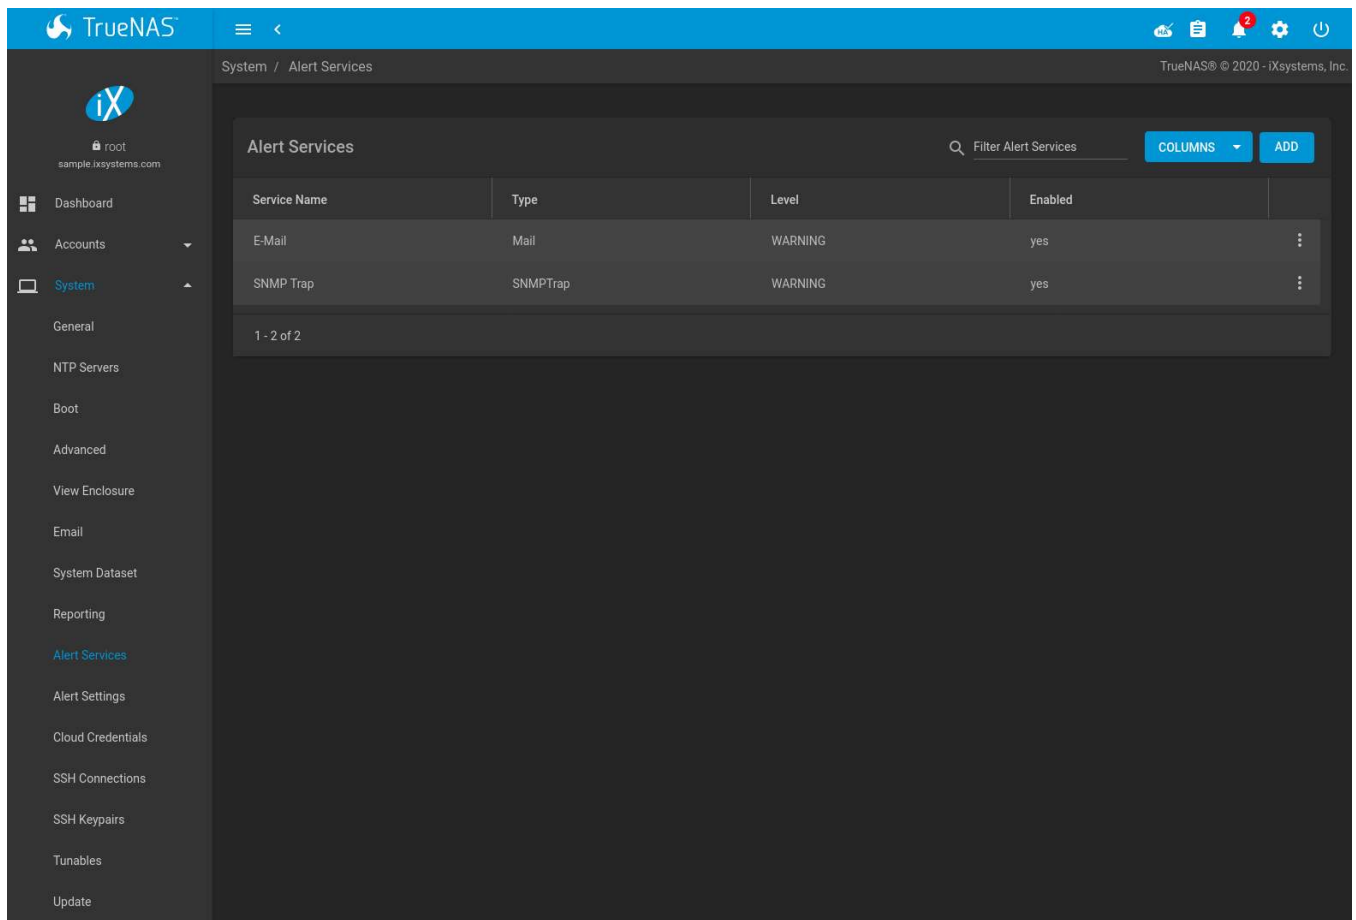


Fig. 5.9: Alert Services

Click **ADD** to display the *Add Alert Service* form, [Figure 5.10](#).

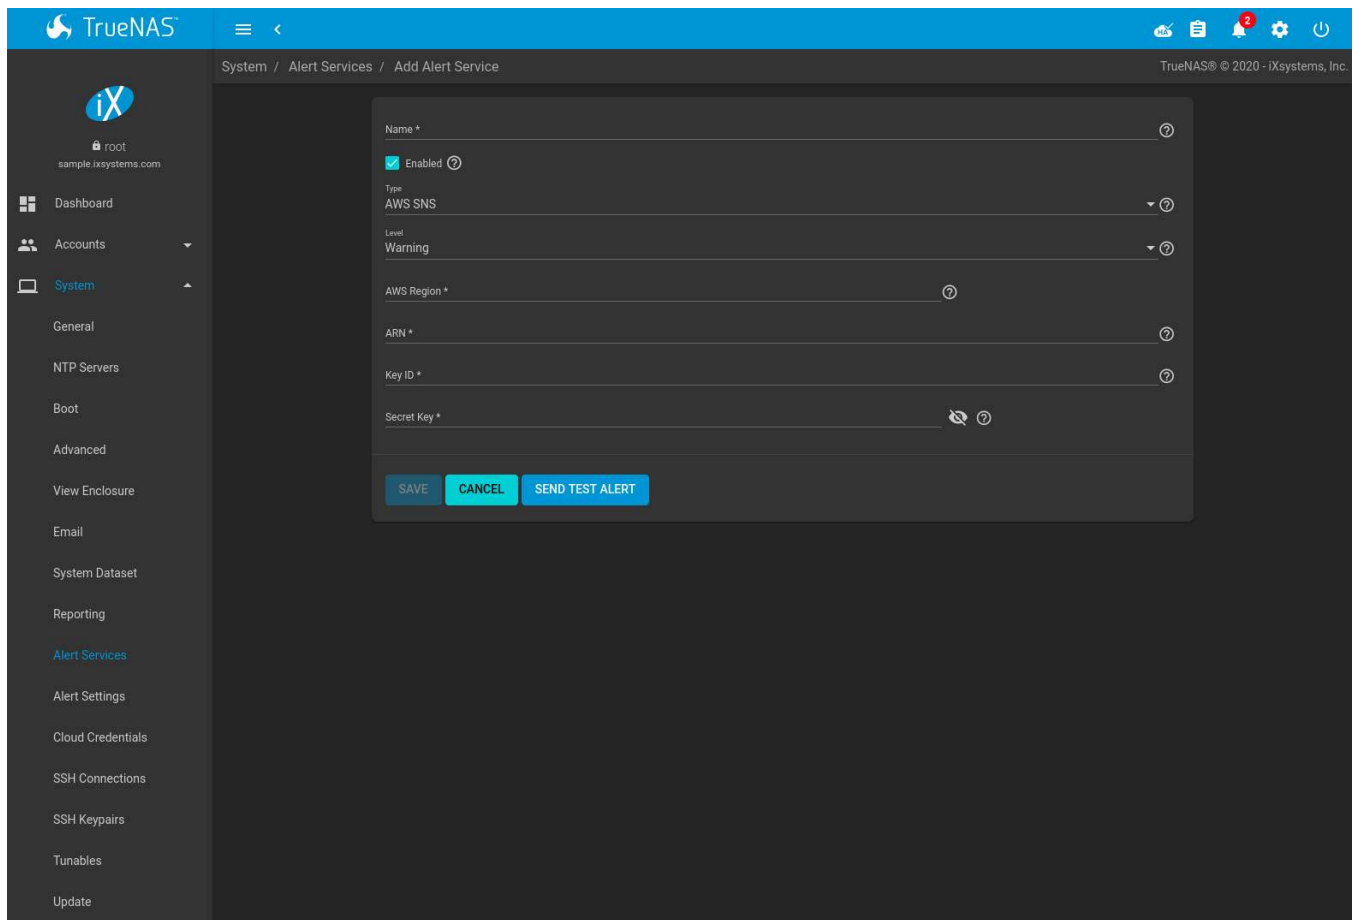


Fig. 5.10: Add Alert Service

Select the *Type* to choose an alert service to configure.

Alert services can be set for a particular severity *Level*. All alerts of that level are then sent out with that alert service. For example, if the *E-Mail* alert service *Level* is set to *Info*, any *Info* level alerts are sent by that service. Multiple alert services can be set to the same level. For instance, *Critical* alerts can be sent both by email and PagerDuty by setting both alert services to the *Critical* level.

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*.

Click *SEND TEST ALERT* to test the chosen alert 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.

5.10 Alert Settings

System → *Alert Settings* has options to configure each TrueNAS® *Alert* (page 305).

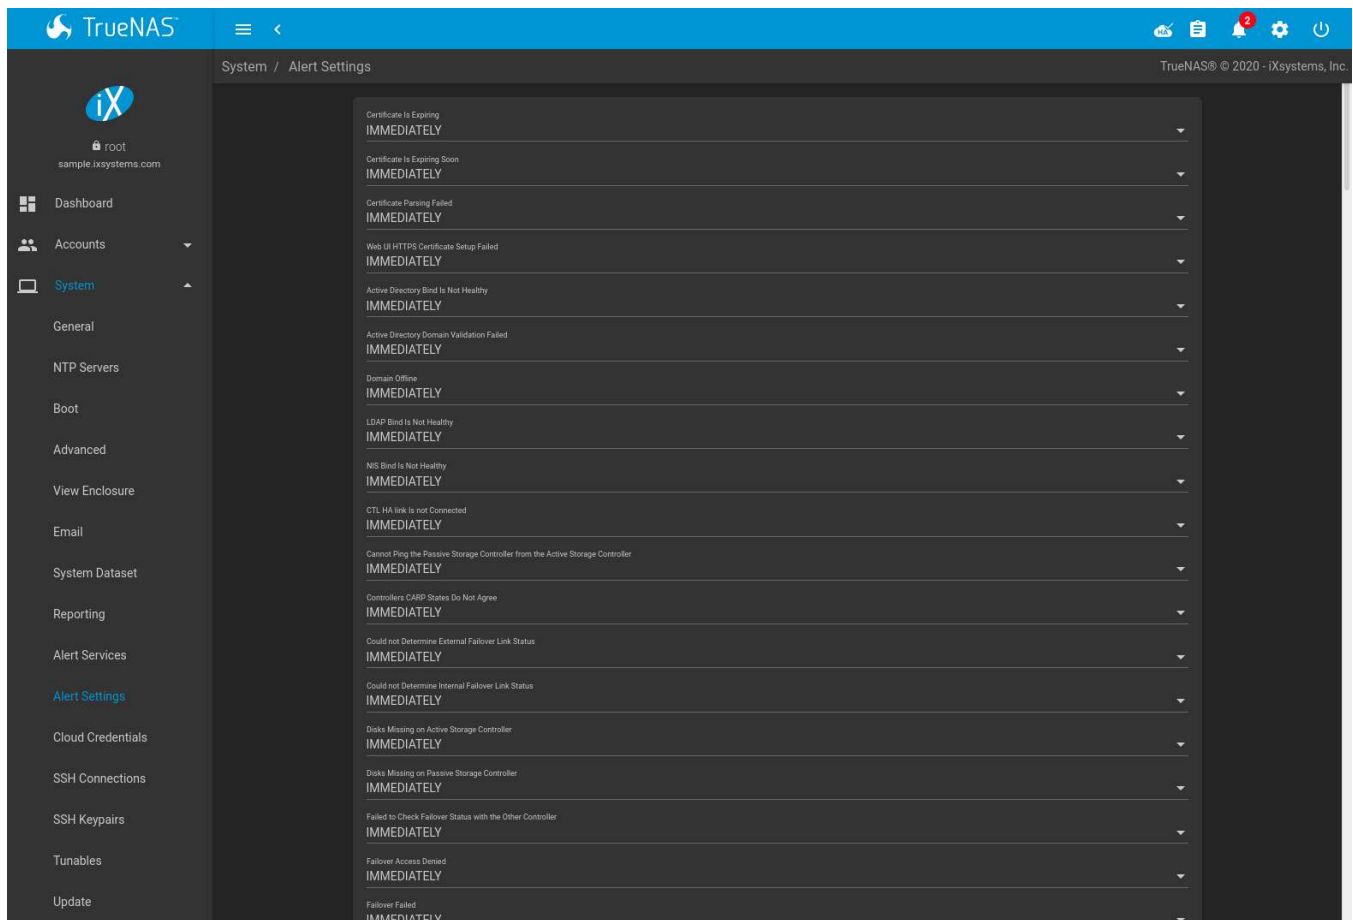


Fig. 5.11: Alert Settings

Alerts are grouped by web interface feature or service monitor. To customize alert importance, use the *Warning Level* drop-down. To adjust how often alert notifications are sent, use the *Frequency* drop-down. Setting the *Frequency* to *NEVER* prevents that alert from being added to alert notifications, but the alert can still show in the web interface if it is triggered.

To configure where alert notifications are sent, use [Alert Services](#) (page 50).

5.11 Cloud Credentials

TrueNAS® can use cloud services for features like [Cloud Sync Tasks](#) (page 113). The [rclone](https://rclone.org/) credentials to provide secure connections with cloud services are entered here. 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 available.

Note: The hubiC cloud service has [suspended creation of new accounts](https://www.ovh.co.uk/subscriptions-hubic-ended/) (https://www.ovh.co.uk/subscriptions-hubic-ended/).

Warning: Cloud Credentials are stored in encrypted form. To be able to restore Cloud Credentials from a [saved configuration](#) (page 32), “Export Password Secret Seed” must be set when saving that configuration.

Click *System* → *Cloud Credentials* to see the screen shown in [Figure 5.12](#).

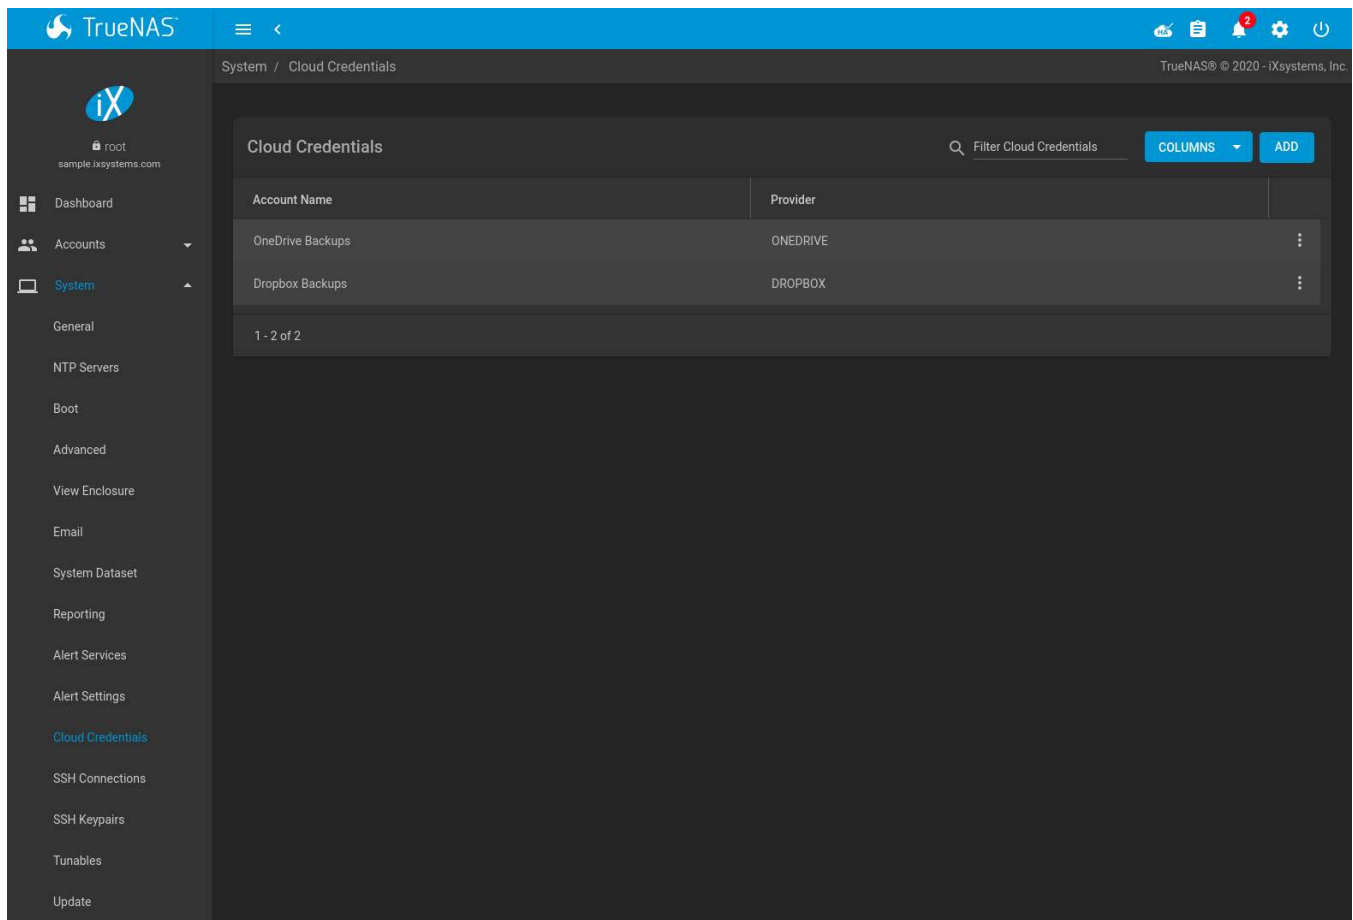


Fig. 5.12: Cloud Credentials List

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

Click **ADD** to add a new cloud credential. Choose a *Provider* to display any specific options for that provider. [Figure 5.13](#) shows an example configuration:

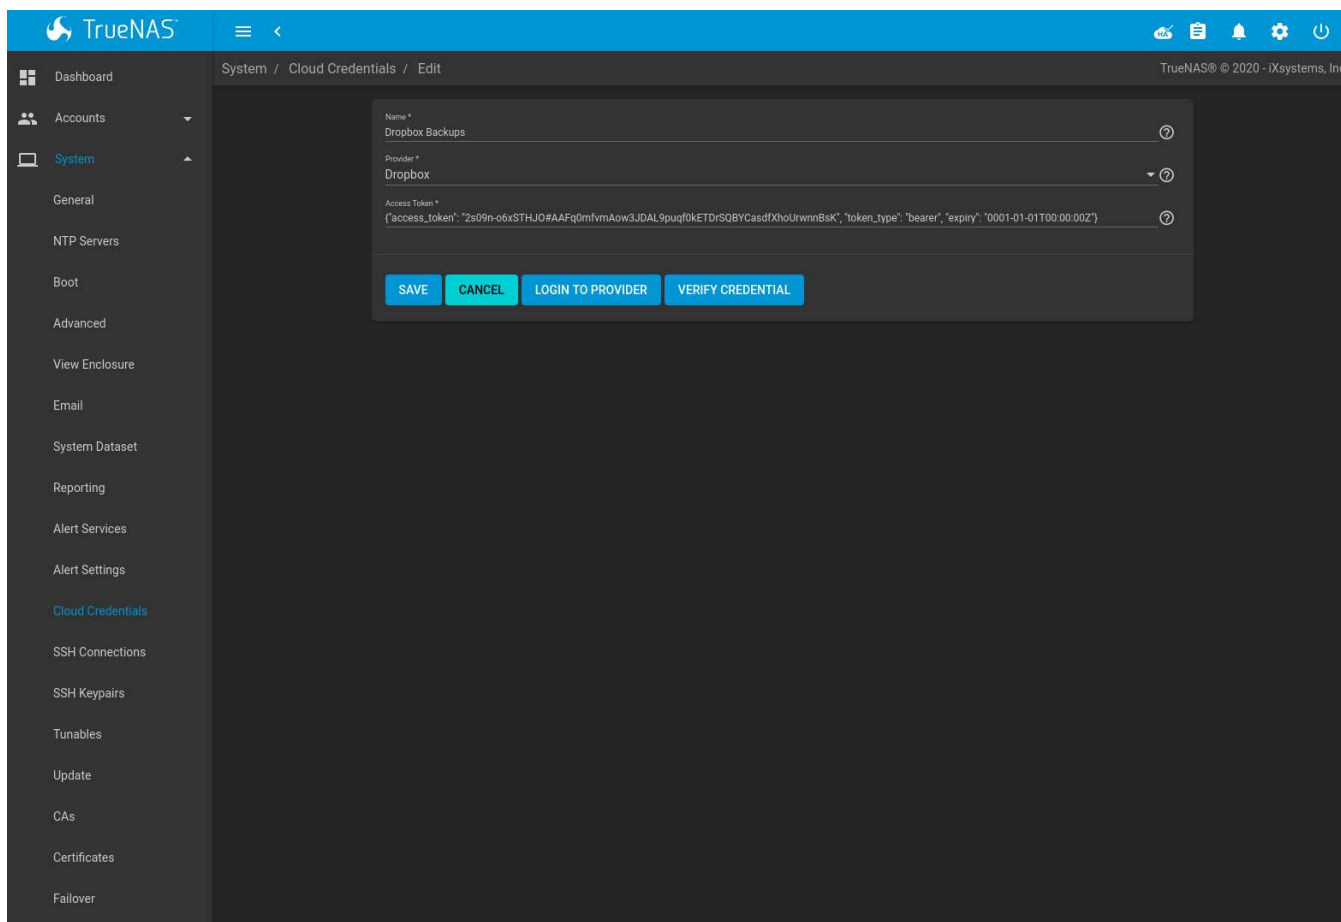


Fig. 5.13: Add Amazon S3 Credential

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 5.6](#). Clicking a provider name opens a new browser tab to the [rclone documentation](https://rclone.org/docs/) (<https://rclone.org/docs/>) for that provider.

Table 5.6: Cloud Credential Options

Provider	Setting	Description
Amazon S3 (https://rclone.org/s3/)	Access Key ID	Enter the Amazon Web Services Key ID. This is found on Amazon AWS (https://aws.amazon.com) by going through <i>My Account</i> → <i>Security Credentials</i> → <i>Access Keys</i> . Must be alphanumeric and between 5 and 20 characters.
Amazon S3 (https://rclone.org/s3/)	Secret Access Key	Enter the Amazon Web Services password. If the Secret Access Key cannot be found or remembered, go to <i>My Account</i> → <i>Security Credentials</i> → <i>Access Keys</i> and create a new key pair. Must be alphanumeric and between 8 and 40 characters.
Amazon S3 (https://rclone.org/s3/)	Endpoint URL	Set <i>Advanced Settings</i> to access this option. S3 API endpoint URL (https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteEndpoints.html). When using AWS, the endpoint field can be empty to use the default endpoint for the region, and available buckets are automatically fetched. Refer to the AWS Documentation for a list of Simple Storage Service Website Endpoints (https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_website_region_end).

Continued on next page

Table 5.6 – continued from previous page

Provider	Setting	Description
Amazon S3 (https://rclone.org/s3/)	Region	AWS resources in a geographic area (https://docs.aws.amazon.com/general/latest/gr/rande-manage.html). Leave empty to automatically detect the correct public region for the bucket. Entering a private region name allows interacting with Amazon buckets created in that region. For example, enter <code>us-gov-east-1</code> to discover buckets created in the eastern AWS GovCloud (https://docs.aws.amazon.com/govcloud-us/latest/UserGuide/whatis.html) region.
Amazon S3 (https://rclone.org/s3/)	Disable Endpoint Region	Set <i>Advanced Settings</i> to access this option. Skip automatic detection of the <i>Endpoint URL</i> region. Set this when configuring a custom <i>Endpoint URL</i> .
Amazon S3 (https://rclone.org/s3/)	Use Signature Version 2	Set <i>Advanced Settings</i> to access this option. Force using <i>Signature Version 2</i> (https://docs.aws.amazon.com/general/latest/gr/signature-version-2.html) to sign API requests. Set this when configuring a custom <i>Endpoint URL</i> .
Backblaze B2 (https://rclone.org/b2/)	Key ID, Application Key	Alphanumeric Backblaze B2 (https://www.backblaze.com/b2/cloud-storage.html) application keys. To generate a new application key, log in to the Backblaze account, go to the <i>App Keys</i> page, and add a new application key. Copy the <code>keyID</code> and <code>applicationKey</code> strings into the TrueNAS® web interface fields.
Box (https://rclone.org/box/)	Access Token	Configured with <i>Open Authentication</i> (page 57).
Dropbox (https://rclone.org/dropbox/)	Access Token	Configured with <i>Open Authentication</i> (page 57). The access token can be manually created by going to the Dropbox <i>App Console</i> (https://www.dropbox.com/developers/apps). After creating an app, go to <i>Settings</i> and click <i>Generate</i> under the Generated access token field.
FTP (https://rclone.org/ftp/)	Host, Port	Enter the FTP host and port.
FTP (https://rclone.org/ftp/)	Username, Password	Enter the FTP username and password.
Google Cloud Storage (https://rclone.org/googlecloudstorage/)	JSON Service Account Key	Upload a Google <i>Service Account credential file</i> (https://rclone.org/googlecloudstorage/#service-account-support). The file is created with the <i>Google Cloud Platform Console</i> (https://console.cloud.google.com/apis/credentials).
Google Drive (https://rclone.org/drive/)	Access Token, Team Drive ID	The <i>Access Token</i> is configured with <i>Open Authentication</i> (page 57). <i>Team Drive ID</i> is only used when connecting to a <i>Team Drive</i> (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 (https://rclone.org/http/)	URL	Enter the HTTP host URL.
hubiC (https://rclone.org/hubic/)	Access Token	Enter the access token. See the <i>HubiC guide</i> (https://api.hubic.com/sandbox/) for instructions to obtain an access token.
Mega (https://rclone.org/mega/)	Username, Password	Enter the Mega (https://mega.nz/) username and password.
Microsoft Azure Blob Storage (https://rclone.org/azureblob/)	Account Name, Account Key	Enter the Azure Blob Storage account name and key.

Continued on next page

Table 5.6 – continued from previous page

Provider	Setting	Description
Microsoft OneDrive (https://rclone.org/onedrive/)	Access Token, Drives List, Drive Account Type, Drive ID	The <i>Access Token</i> is configured with <i>Open Authentication</i> (page 57). Authenticating a Microsoft account adds the <i>Drives List</i> and selects the correct <i>Drive Account Type</i> . The <i>Drives List</i> shows all the drives and IDs registered to the Microsoft account. Selecting a drive automatically fills the <i>Drive ID</i> field.
pCloud (https://rclone.org/pcloud/)	Access Token	Configured with <i>Open Authentication</i> (page 57).
SFTP (https://rclone.org/sftp/)	Host, Port, User-name, Password, Private Key ID	Enter the SFTP host and port. Enter an account user name that has SSH access to the host. Enter the password for that account <i>or</i> import the private key from an existing <i>SSH keypair</i> (page 61). To create a new SSH key for this credential, open the <i>Private Key ID</i> drop-down and select <i>Generate New</i> .
WebDAV (https://rclone.org/webdav/)	URL, WebDAV service	Enter the URL and use the dropdown to select the WebDAV service.
WebDAV (https://rclone.org/webdav/)	Username, Password	Enter the username and password.
Yandex (https://rclone.org/yandex/)	Access Token	Configured with <i>Open Authentication</i> (page 57).

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 TrueNAS® 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. *Open Authentication* (OAuth) (<https://openauthentication.org/>) is used with some cloud providers. These providers have a *LOGIN TO PROVIDER* button that opens a dialog to log in to that provider and fill the *Access Token* field with valid credentials.

Enter the information and click *VERIFY CREDENTIAL*. The *Credential is valid.* displays when the credential information is verified.

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

5.12 SSH Connections

Secure Socket Shell (SSH) (<https://searchsecurity.techtarget.com/definition/Secure-Shell>) is a network protocol that provides a secure method to access and transfer files between two hosts while using an unsecure network. SSH can use user account credentials to establish secure connections, but often uses key pairs shared between host systems for authentication.

TrueNAS® uses *System* → *SSH Connections* to quickly create SSH connections and show any saved connections. These connections are required when creating a new *replication* (page 107) to back up dataset snapshots.

The remote system must be configured to allow SSH connections. Some situations can also require allowing root account access to the remote system. For TrueNAS® systems, go to *Services* and edit the *SSH* (page 244) service to allow SSH connections and root account access.

To add a new SSH connection, go to *System* → *SSH Connections* and click *ADD*.

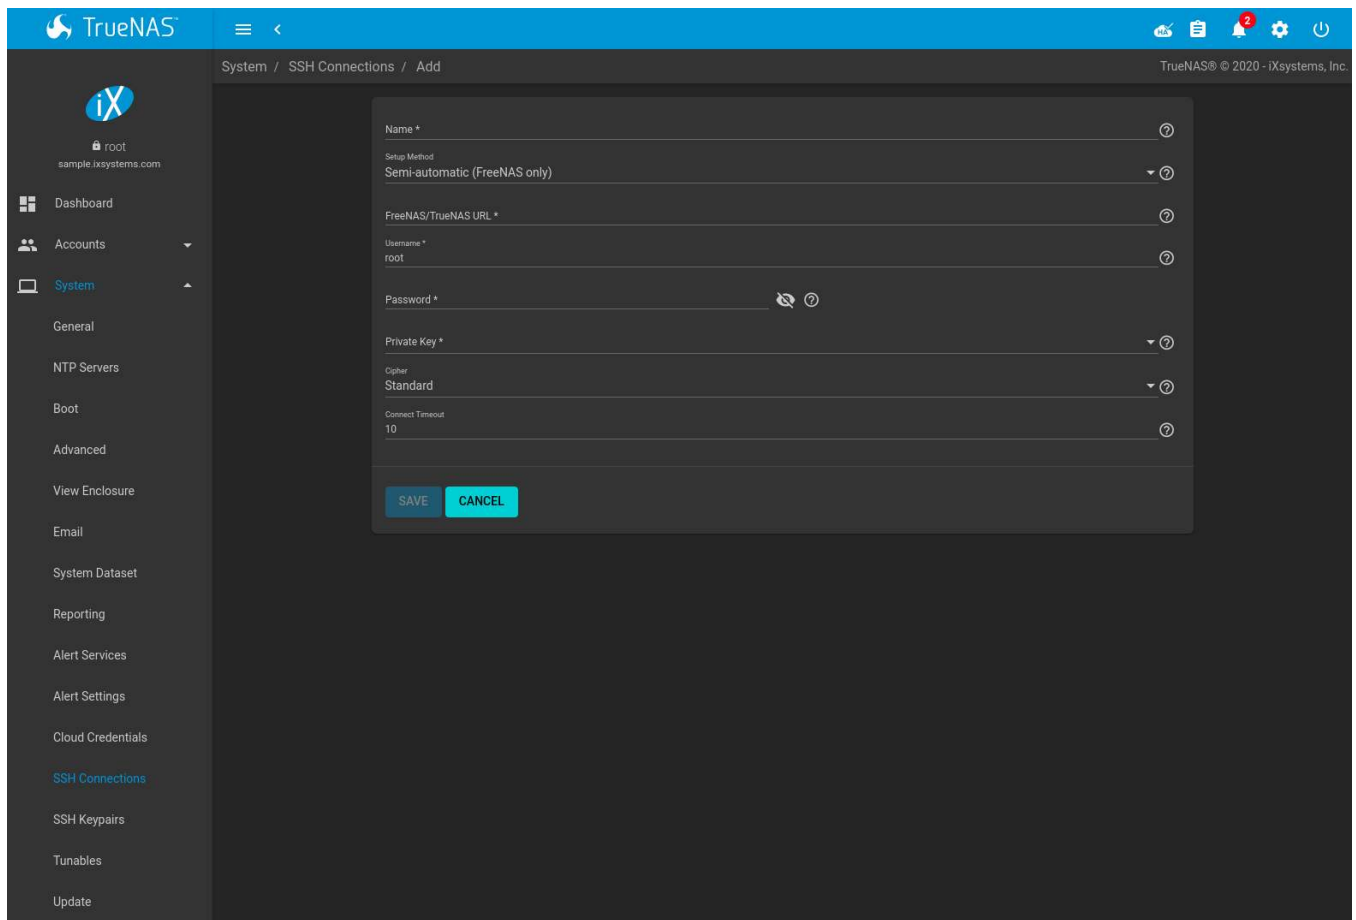


Table 5.7: SSH Connection Options

Setting	Value	Description
Name	string	Descriptive name of this SSH connection. SSH connection names must be unique.
Setup Method	drop-down menu	How to configure the connection: <i>Manual</i> requires configuring authentication on the remote system. This can require copying SSH keys and modifying the <i>root</i> user account on that system. See Manual Setup (page 59). <i>Semi-automatic</i> is only functional when configuring an SSH connection between TrueNAS® systems. After authenticating the connection, all remaining connection options are automatically configured. See Semi-Automatic Setup (page 60).
Host	string	Enter the hostname or IP address of the remote system. Only available with <i>Manual</i> configurations.
Port	integer	Port number on the remote system to use for the SSH connection. Only available with <i>Manual</i> configurations.
FreeNAS URL	string	Hostname or IP address of the remote TrueNAS® system. Only available with <i>Semi-automatic</i> configurations. A valid URL scheme is required. Example: <code>https://10.231.3.76</code>
Username	string	User account name to use for logging in to the remote system
Password	string	User account password used to log in to the TrueNAS® system. Only available with <i>Semi-automatic</i> configurations.
Private Key	drop-down menu	Choose a saved SSH Keypair (page 61) or select <i>Generate New</i> to create a new keypair and apply it to this connection.
Remote Host Key	string	Remote system SSH key for this system to authenticate the connection. Only available with <i>Manual</i> configurations. When all other fields are properly configured, click <i>DISCOVER REMOTE HOST KEY</i> to query the remote system and automatically populate this field.
Cipher	drop-down menu	Connection security level: <ul style="list-style-type: none"> • <i>Standard</i> is most secure, but has the greatest impact on connection speed. • <i>Fast</i> is less secure than <i>Standard</i> but can give reasonable transfer rates for devices with limited cryptographic speed. • <i>Disabled</i> removes all security in favor of maximizing connection speed. Disabling the security should only be used within a secure, trusted network.
Connect Timeout	integer	Time (in seconds) before the system stops attempting to establish a connection with the remote system.

Saved connections can be edited or deleted. Deleting an SSH connection also deletes or disables paired [SSH Keypairs](#) (page 61), [Replication Tasks](#) (page 107), and [Cloud Credentials](#) (page 53).

5.12.1 Manual Setup

Choosing to manually set up the SSH connection requires copying a public encryption key from the local to remote system. This allows a secure connection without a password prompt.

The examples here and in [Semi-Automatic Setup](#) (page 60) refer to the TrueNAS® system that is configuring a new connection in *System* → *SSH Connections* as *Host 1*. The TrueNAS® system that is receiving the encryption key is *Host 2*.

On *Host 1*, go to *System* → *SSH Keypairs* and create a new [SSH Keypair](#) (page 61). Highlight the entire *Public Key* text, right-click in the highlighted area, and click *Copy*.

Log in to *Host 2* and 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 5.14.

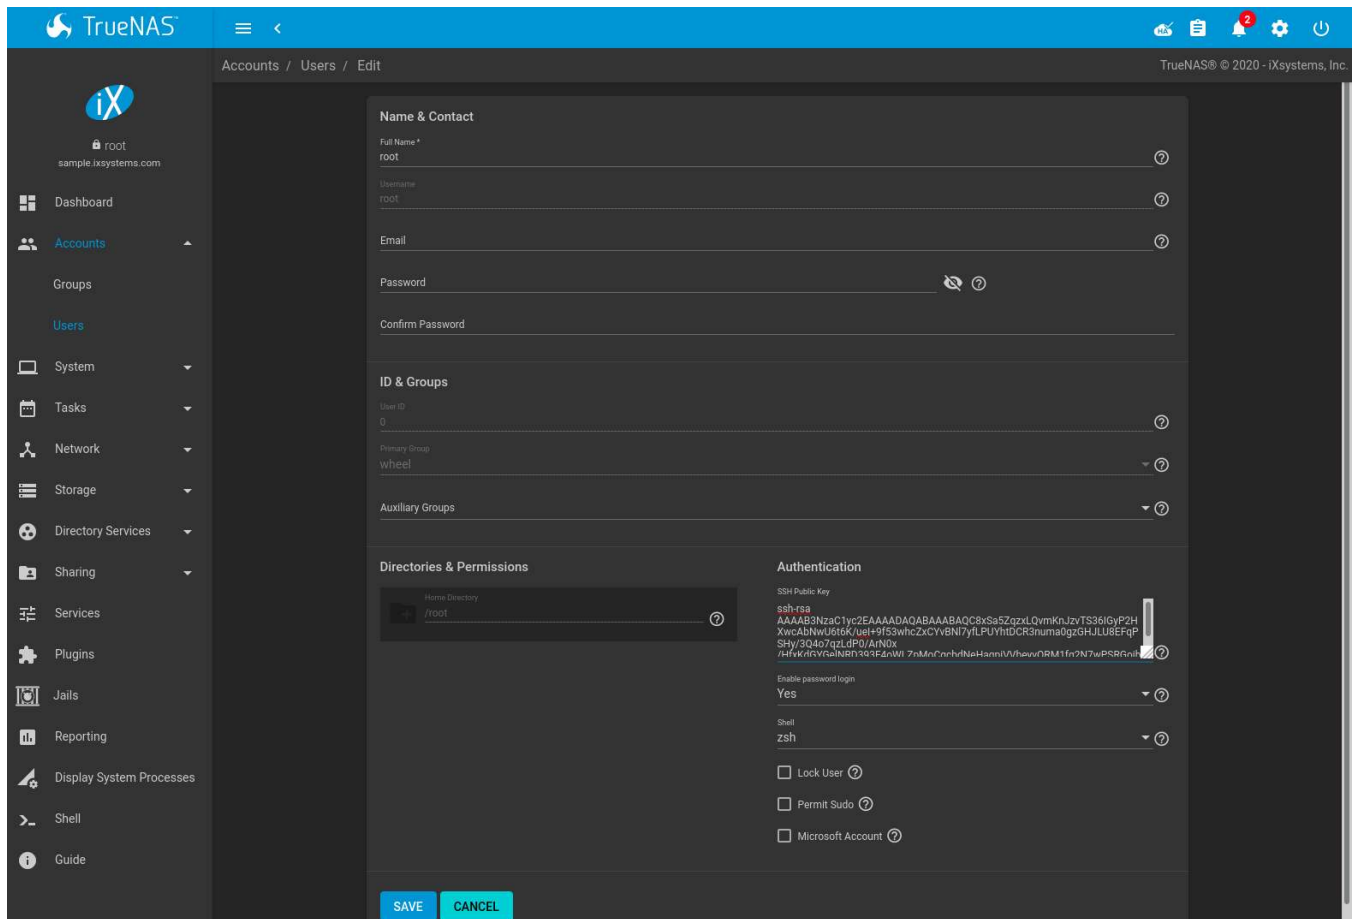


Fig. 5.14: Paste the Replication Key

Switch back to *Host 1* and go to *System* → *SSH Connections* and click *ADD*. Set the *Setup Method* to *Manual*, select the previously created keypair as the *Private Key*, and fill in the rest of the connection details for *Host 2*. Click *DISCOVER REMOTE HOST KEY* to obtain the remote system key. Click *SAVE* to store this SSH connection.

5.12.2 Semi-Automatic Setup

TrueNAS® offers a semi-automatic setup mode that simplifies setting up an SSH connection with another FreeNAS or TrueNAS system. When administrator account credentials are known for *Host 2*, semi-automatic setup allows configuring the SSH connection without logging in to *Host 2* to transfer SSH keys.

In *Host 1*, go to *System* → *SSH Keypairs* and create a new *SSH Keypair* (page 61). Go to *System* → *SSH Connections* and click *ADD*.

Choose *Semi-automatic* as the *Setup Method*. Enter the *Host 2* URL in *FreeNAS URL* using the format `http://freenas.remote`, where *freenas.remote* is the *Host 2* hostname or IP address.

Enter credentials for an *Host 2* user account that can accept SSH connection requests and modify *Host 2*. This is typically the *root* account.

Select the SSH keypair that was just created for the *Private Key*.

Fill in the remaining connection configuration fields and click *SAVE*. *Host 1* can use this saved configuration to establish a connection to *Host 2* and exchange the remaining authentication keys.

5.13 SSH Keypairs

TrueNAS® generates and stores [RSA-encrypted](https://en.wikipedia.org/wiki/RSA_%28cryptosystem%29) (https://en.wikipedia.org/wiki/RSA_%28cryptosystem%29) SSH public and private keypairs in *System* → *SSH Keypairs*. These are generally used when configuring [SSH Connections](#) (page 57) or [SFTP Cloud Credentials](#) (page 53). Encrypted keypairs or keypairs with passphrases are not supported.

To generate a new keypair, click **ADD**, enter a name, and click **GENERATE KEYPAIR**. The *Private Key* and *Public Key* fields fill with the key strings. SSH key pair names must be unique.

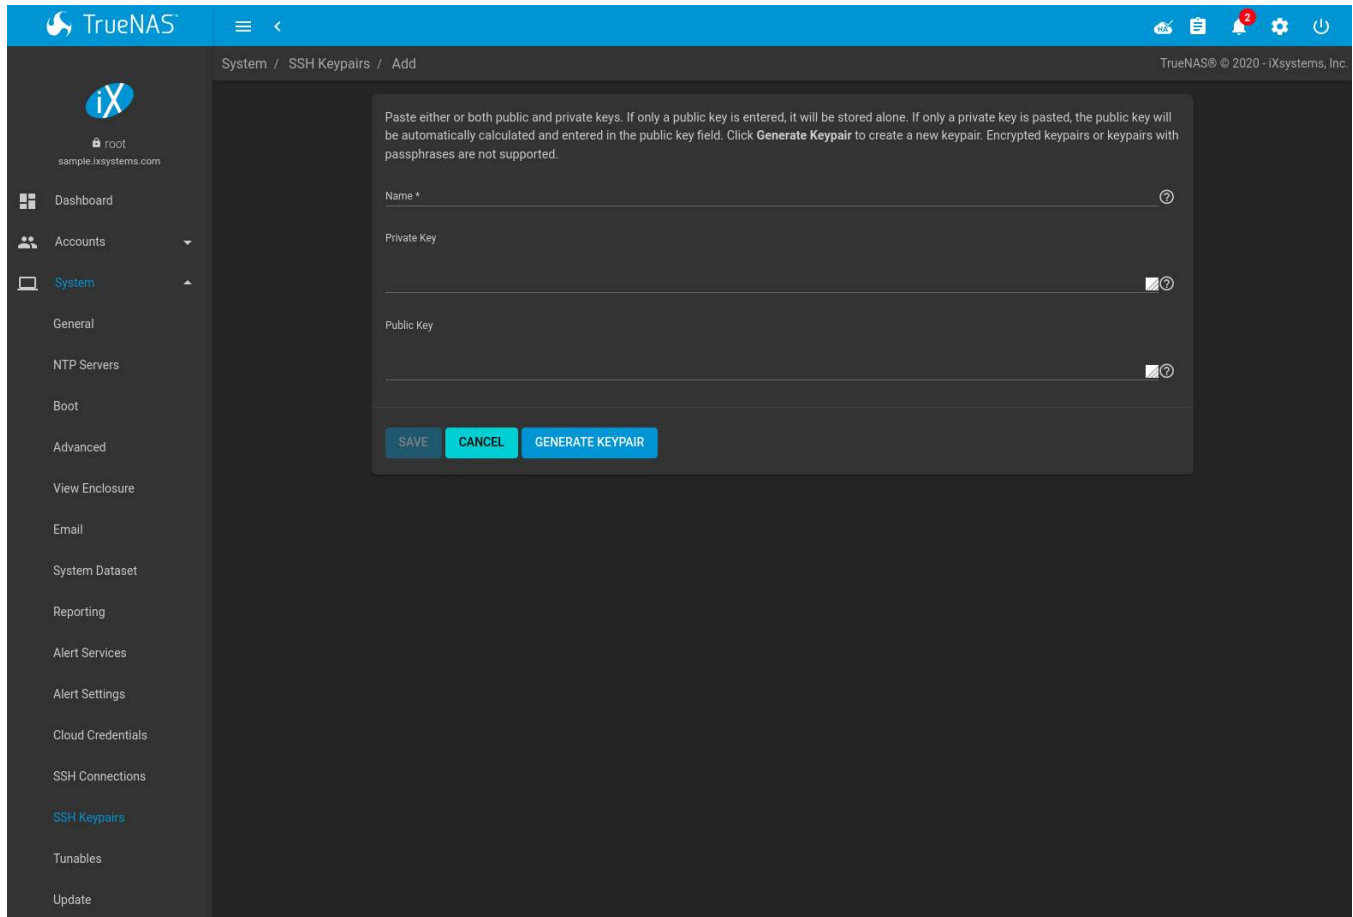
The screenshot shows the TrueNAS web interface. The top navigation bar is blue with the TrueNAS logo and a hamburger menu. Below it, a breadcrumb trail reads 'System / SSH Keypairs / Add'. The left sidebar is dark grey with a list of system settings: Dashboard, Accounts, System (selected), General, NTP Servers, Boot, Advanced, View Enclosure, Email, System Dataset, Reporting, Alert Services, Alert Settings, Cloud Credentials, SSH Connections, SSH Keypairs (highlighted), Tunables, and Update. The main content area is dark grey and contains a form for adding a new SSH keypair. At the top of the form is a text box with instructions: 'Paste either or both public and private keys. If only a public key is entered, it will be stored alone. If only a private key is pasted, the public key will be automatically calculated and entered in the public key field. Click Generate Keypair to create a new keypair. Encrypted keypairs or keypairs with passphrases are not supported.' Below this are three input fields: 'Name *' (with a question mark icon), 'Private Key', and 'Public Key' (with a copy icon and a question mark icon). At the bottom of the form are three buttons: 'SAVE' (grey), 'CANCEL' (cyan), and 'GENERATE KEYPAIR' (blue).

Fig. 5.15: Example Keypair

Click **SAVE** to store the new keypair. These saved keypairs can be selected later in the web interface without having to manually copy the key values.

Keys are viewed or modified by going to *System* → *SSH Keypairs* and clicking **⋮** (Options) and *Edit* for the keypair name.

Deleting an SSH Keypair also deletes any associated [SSH Connections](#) (page 57), [Replication Tasks](#) (page 107) or [SFTP Cloud Credentials](#) (page 53) that use this keypair are disabled but not removed.

5.14 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 TrueNAS® 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 TrueNAS® has been optimized for storage, not all of the services mentioned in `rc.conf(5)` are available for configuration. Note that in TrueNAS®, 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 TrueNAS® system and a loader could adversely affect the ability of the TrueNAS® 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 5.16.

The screenshot displays the TrueNAS web interface. On the left is a dark sidebar with a navigation menu. The 'System' section is expanded, and 'Tunables' is highlighted. The main content area shows the 'Add' form for a new tunable. The form has a light gray background and includes the following fields: 'Variable *' (text input), 'Value *' (text input), 'Type' (dropdown menu with 'loader' selected), and 'Description' (text input). Below these fields is a checkbox labeled 'Enabled' which is checked. At the bottom of the form are two buttons: 'SAVE' and 'CANCEL'. The top of the interface features a blue header with the TrueNAS logo, a hamburger menu, and user information. The breadcrumb trail at the top reads 'System / Tunables / Add'.

Fig. 5.16: Adding a Tunable

Table 5.8 summarizes the options when adding a tunable.

Table 5.8: Adding a Tunable

Setting	Value	Description
Variable	string	The name of the <i>sysctl</i> 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> .
Description	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 TrueNAS® 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 TrueNAS® is installed. TrueNAS® 11.3 ships with the *sysctls* set:

```
kern.metadelays=3
kern.dirdelay=4
kern.filedelay=5
kern.coredump=1
net.inet.carp.preempt=1
debug.ddb.textdump.pending=1
vfs.nfsd.tpcachetimeo=300
vfs.nfsd.tcphighwater=150000
vfs.zfs.vdev.larger_ashift_minimal=0
net.inet.carp.senderr_demotion_factor=0
net.inet.carp.ifdown_demotion_factor=0
```

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 TrueNAS® is installed. TrueNAS® 11.3 ships with these loaders set:

```
autoboot_delay="2"
loader_logo="truenas-logo"
loader_menu_title="Welcome to TrueNAS"
loader_brand="truenas-brand"
loader_version=" "
kern.cam.boot_delay="10000"
debug.debugger_on_panic=1
debug.ddb.textdump.pending=1
hw.hpttr.attach_generic=0
ispfw_load="YES"
freenas_sysctl_load="YES"
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
kern.geom.label.ufs.enable=0
kern.geom.label.ufsid.enable=0
kern.geom.label.reiserfs.enable=0
kern.geom.label.ntfs.enable=0
kern.geom.label.msdosfs.enable=0
kern.geom.label.ext2fs.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
vfs.nfsd.fha.write=0
vfs.nfsd.fha.max_nfsds_per_fh=32
kern.ipc.nmbclusters="262144"
kern.hwpmc.nbuffers="4096"
kern.hwpmc.nsamples="4096"
hw.memtest.tests="0"
vfs.zfs.trim.enabled="0"
kern.cam.cttl.ha_mode=2
hint.ntb_hw.0.config="ntb_pmem:1:4:0,ntb_transport"
hint.ntb_transport.0.config=":3"
hw.ntb.msix_mw_idx="-1"
```

Do not add or edit the default tunables. Changing the default tunables can make the system unusable.

The ZFS version used in 11.3 deprecates these tunables:

```
kvfs.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 TrueNAS[®], these tunables are automatically deleted. Please do not manually add them back.

5.15 Update

TrueNAS[®] has an integrated update system to make it easy to keep up to date.

5.15.1 Preparing for Updates

An update usually takes between thirty minutes and an hour. A reboot is required after the update, so it is recommended to schedule updates during a maintenance window, allowing two to three hours to update, test, and possibly roll back if issues appear. On very large systems, a proportionally longer maintenance window is recommended.

For individual support during an upgrade, please open a ticket at <https://support.ixsystems.com>, or call 408-943-4100 to schedule one. Scheduling at least two days in advance of a planned upgrade gives time to make sure a specialist is available for assistance.

Updates from older versions of TrueNAS[®] before 9.3 must be scheduled with support.

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, go to [Boot](#) (page 37) to remove unneeded boot environments.

Operating system updates only modify the operating system devices and do not affect end-user data on storage drives.

Available ZFS version upgrades are indicated by an [Alert](#) (page 305) in the web interface. However, upgrading the ZFS version on storage drives is not recommended until after verifying that rolling back to previous versions of the operating system will not be necessary, and that interchanging the devices with some other system using an older ZFS version is not needed. After a ZFS version upgrade, the storage devices will not be accessible by older versions of TrueNAS®.

5.15.2 Updates and Trains

Cryptographically signed update files are used to update TrueNAS®. Update files provide flexibility in deciding when to upgrade the system. Go to [Boot](#) (page 71) to test an update.

TrueNAS® defines software branches, known as *trains*. There are several trains available for updates:

For Production Use

After new bugfixes and security updates have been tested as production-ready, they are added to these trains. It is recommended to select the update train that matches the currently installed TrueNAS® feature release:

- **TrueNAS-11-STABLE**
- **TrueNAS-11.2-STABLE**
- **TrueNAS-11.3-STABLE**

Legacy Versions

- **TrueNAS-9.10-STABLE**

Maintenance-only updates for the previous branch of TrueNAS®.

- **TrueNAS-9.3-STABLE**

Maintenance-only updates for the older 9.3 branch of TrueNAS®. Use this train only at the recommendation of an iXsystems support engineer.

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

5.15.3 Checking for Updates

Figure 5.17 shows an example of the *System* → *Update* screen.

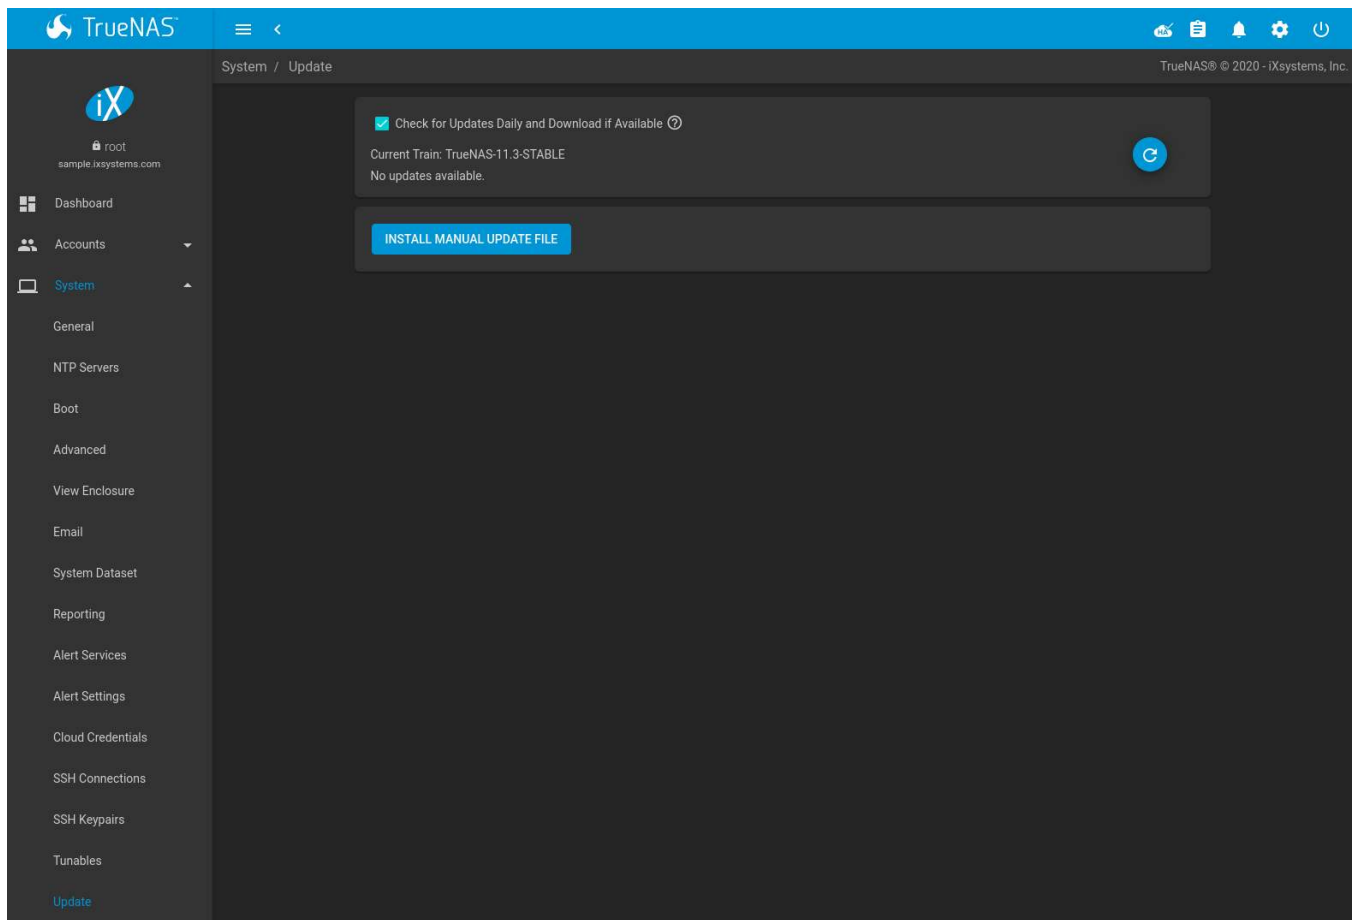



Fig. 5.17: Update Options

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*. Click  (Refresh) to perform another check for updates.

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 5.18](#), 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.

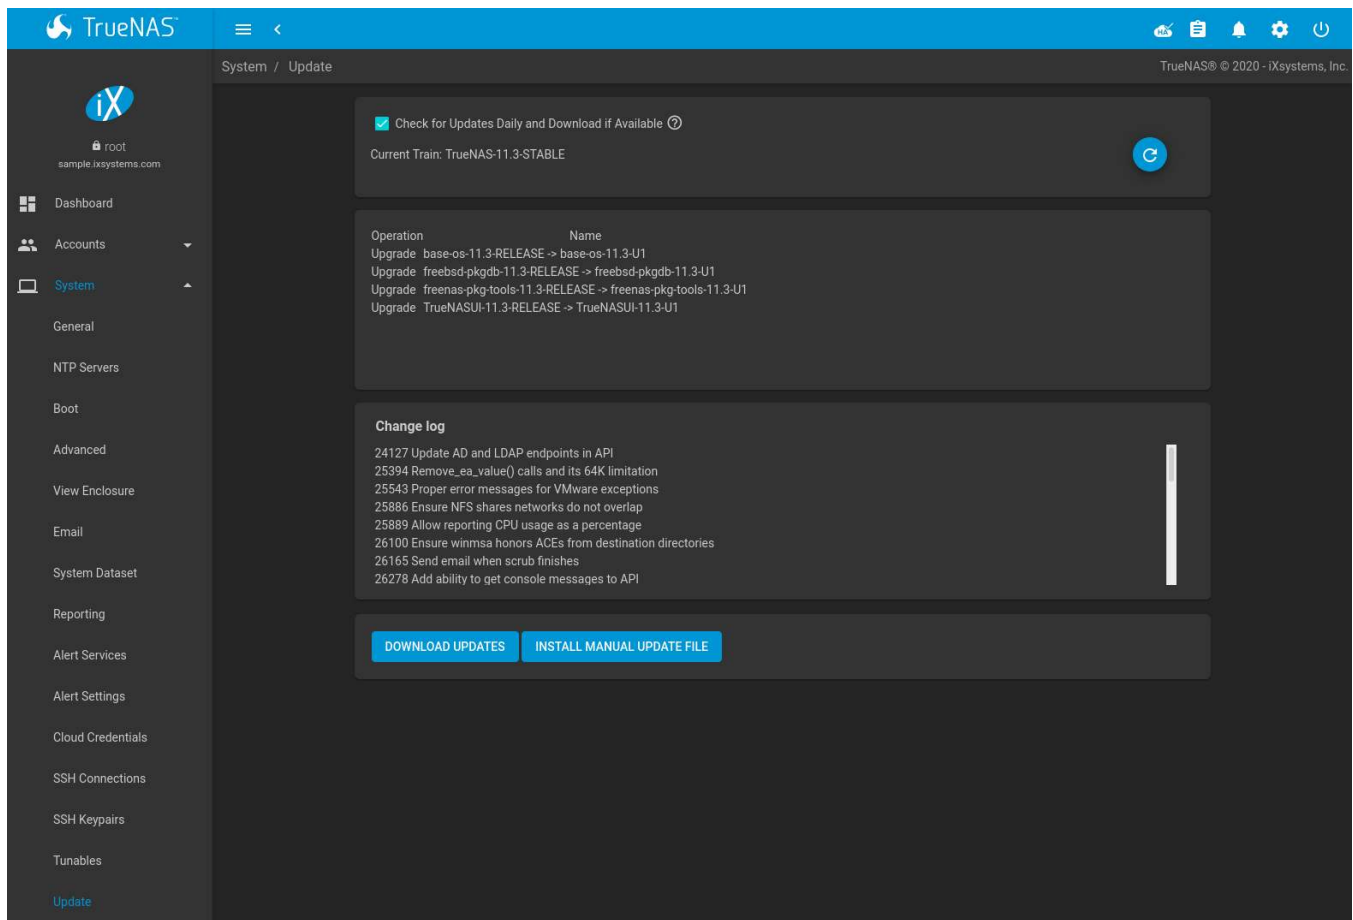
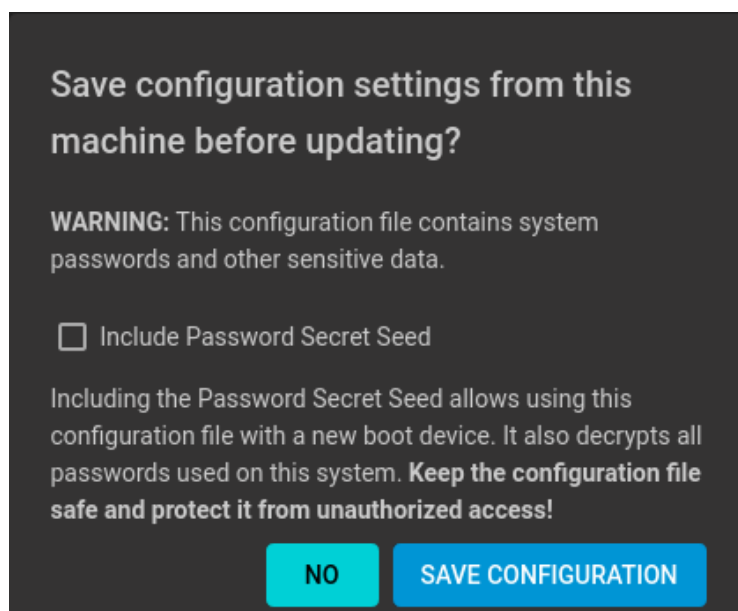


Fig. 5.18: Reviewing Updates

5.15.4 Saving the Configuration File

A dialog to save the system *configuration file* (page 34) appears before installing updates.



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 TrueNAS® system.

5.15.5 Applying Updates

Make sure the system is in a low-usage state as described above in *Preparing for Updates* (page 64).

Click *DOWNLOAD UPDATES* to immediately download and install an update.

The *Save Configuration* (page 67) 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* (page 37) are not removed. If space for a new boot environment is not available, the upgrade fails. Space on the operating system device can be manually freed using *System → Boot*. Review the boot environments and remove the *Keep* attribute or delete any boot environments that are no longer needed.

5.15.6 Manual Updates

Updates can also be manually downloaded and applied in *System → Update*.

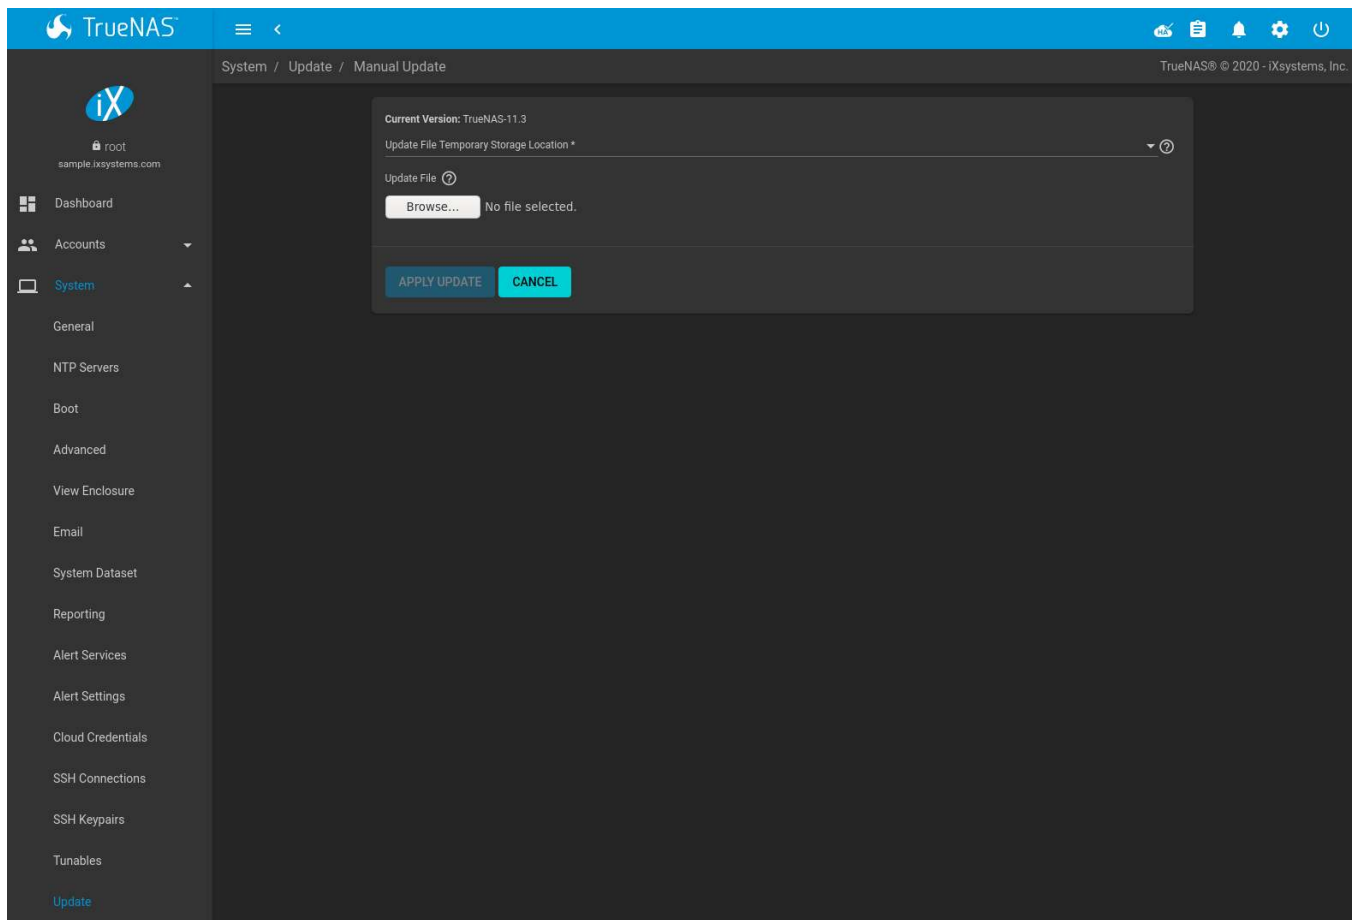
Note: Manual updates cannot be used to upgrade from older major versions.

Go to <https://download.freenas.org/> and find an update file of the desired version. Manual update file names end with `-manual-update-unsigned.tar`.

Download the file to a desktop or laptop computer. Connect to TrueNAS® with a browser and go to *System → Update*. Click *INSTALL MANUAL UPDATE FILE*.

The *Save Configuration* (page 67) dialog opens. This makes it possible to save a copy of the current configuration to external media for backup in case of an update problem.


After the dialog closes, the manual update screen is shown:



The current version of TrueNAS® is shown for verification.

Select the manual update file with the *Browse* button. Set *Reboot After Update* to reboot the system after the update has been installed. Click *APPLY UPDATE* to begin the update.

5.15.7 Update in Progress

Starting an update shows a progress dialog. When an update is in progress, the web interface shows an  icon in the top row. Dialogs also appear in every active web interface session to warn that a system update is in progress. **Do not** interrupt a system update.

5.15.8 Updating from the Shell

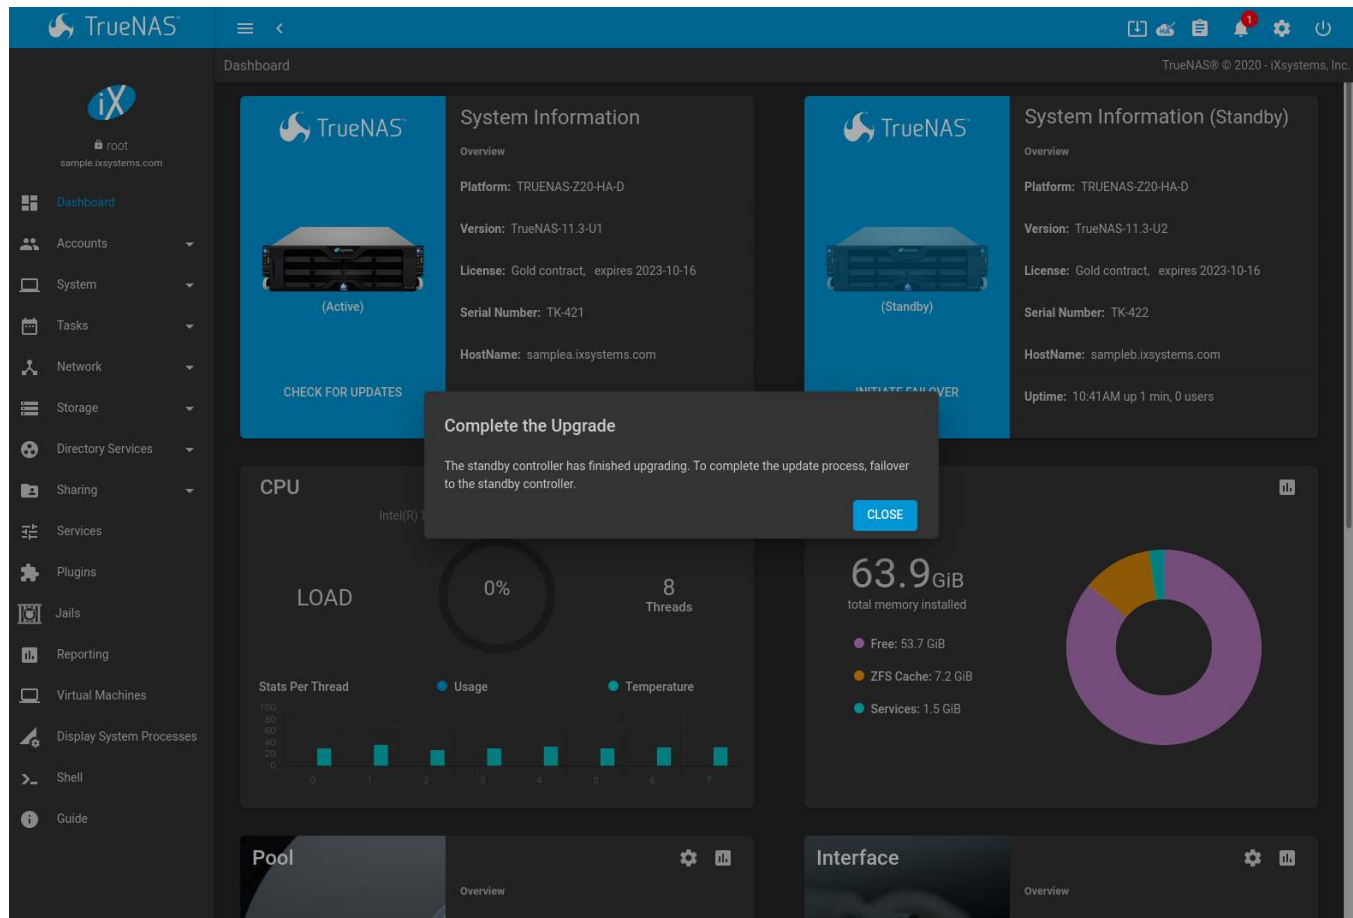
Updates can also be performed from the [Shell](#) (page 302) with an update file. Make the update file available by copying it to the TrueNAS® system, then run the update program, giving it the path to the file: `freenas-update update_file`.

5.15.9 Updating an HA System

On the *Dashboard* of the active TrueNAS controller, click *CHECK FOR UPDATES*. This button changes to *UPDATES AVAILABLE* when there is an available update. Clicking the button goes to *System* → *Update*. When *DOWNLOAD UPDATES* is clicked, it first gives an opportunity to *save the current system configuration* (page 34). Backing up the system configuration is strongly recommended before starting the update. Click *CONTINUE* to start updating both TrueNAS controllers.

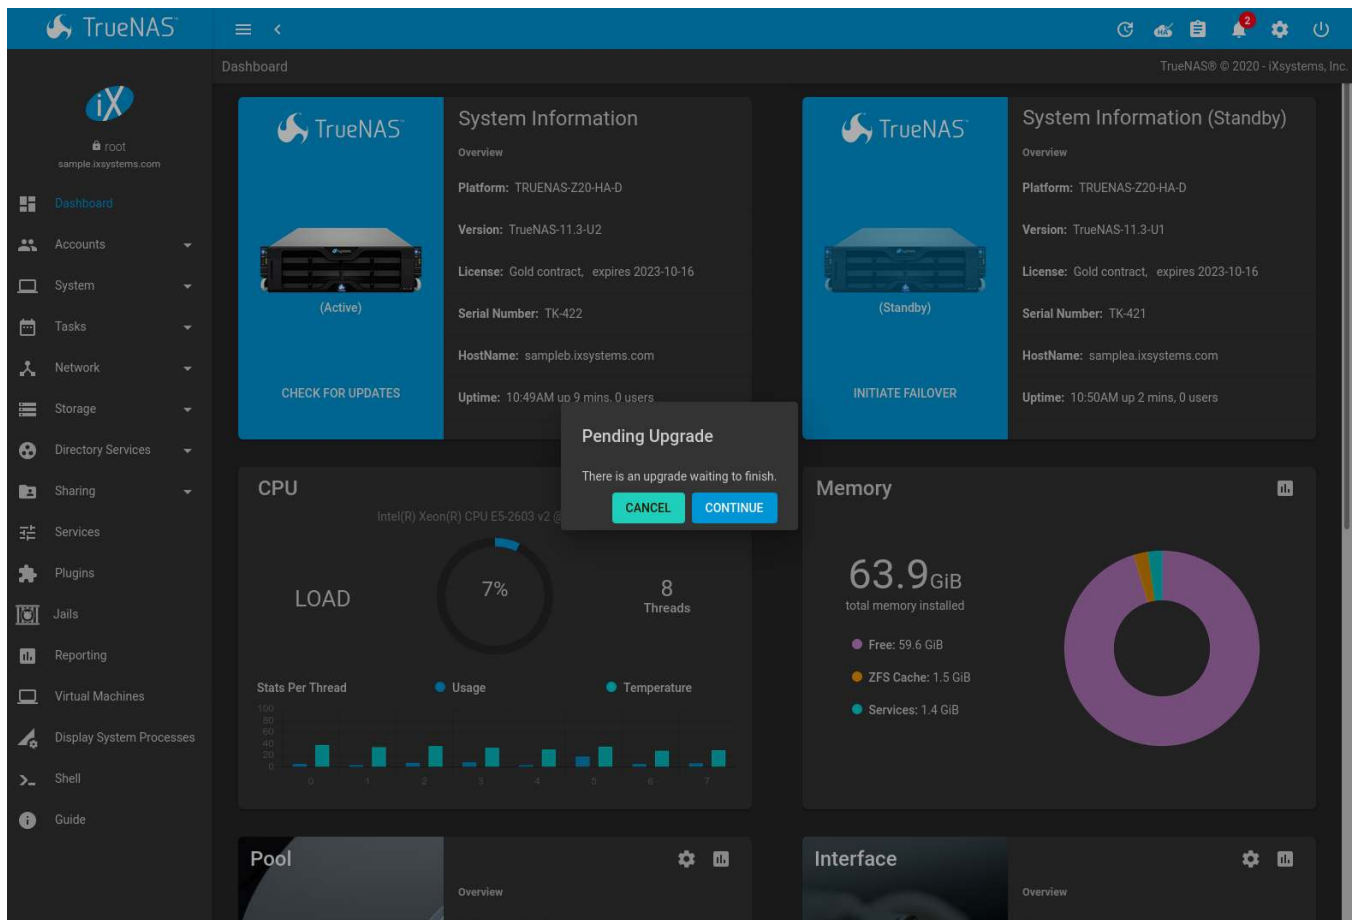
A warning dialog appears for any other user that is logged into the web interface and a “System Updating” icon is shown in the top bar while the update is in progress.

Update progress is shown for both TrueNAS controllers. The standby TrueNAS controller reboots when it is finished updating. To finish updating the active TrueNAS controller, the system must *fail over* (page 81) and deactivate the active TrueNAS controller.



To deactivate the active TrueNAS controller and finish the update, go to the *Dashboard* and click *INITIATE FAILOVER*. This will temporarily interrupt TrueNAS® services and availability. To start the failover, confirm the action and click *FAILOVER*. The browser logs out of the web interface while the active TrueNAS controller deactivates and the other TrueNAS controller is brought online.

The browser shows the web interface login screen when the other TrueNAS controller finishes activating. Log in to the web interface and check the *HA status icon* (page 82) in the top toolbar. This icon shows that HA is unavailable while the previously active TrueNAS controller reboots. The icon changes to show HA is available when the TrueNAS controller is back online. Click *CONTINUE* to finish updating the previously active TrueNAS controller and reboot it again.



When both TrueNAS controllers are online, verify that the update is complete by going to *Dashboard* and confirming that *Version* is the same on both TrueNAS controllers.

5.15.10 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 TrueNAS® console is required. Reboot the system and press the space bar when the boot menu appears, pausing the boot. Select an entry with a date prior to the update, then press `Enter` to boot into that version of the operating system before the update was applied.

5.15.11 Upgrading a ZFS Pool

In TrueNAS®, 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 ZFS features.**
- 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 TrueNAS® 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](#) (page 312). 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 5.19](#).

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.

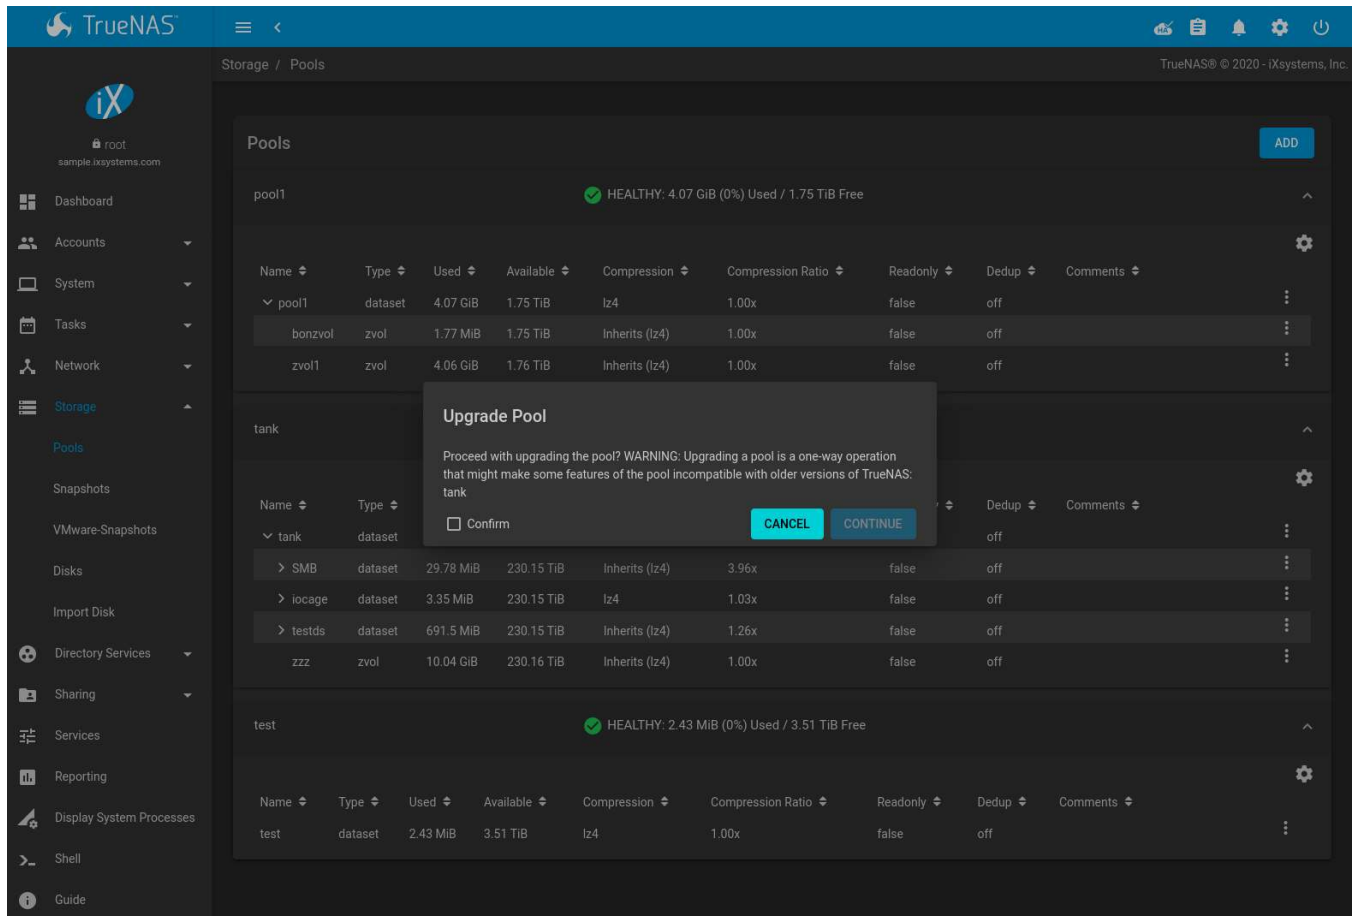


Fig. 5.19: Upgrading a Pool

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.

5.16 CAs

TrueNAS® can act as a Certificate Authority (CA). When encrypting SSL or TLS connections to the TrueNAS® system, either import an existing certificate, or create a CA on the TrueNAS® 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 TrueNAS® system and used on the LDAP server also.

Figure 5.20 shows the screen after clicking *System* → *CAs*.

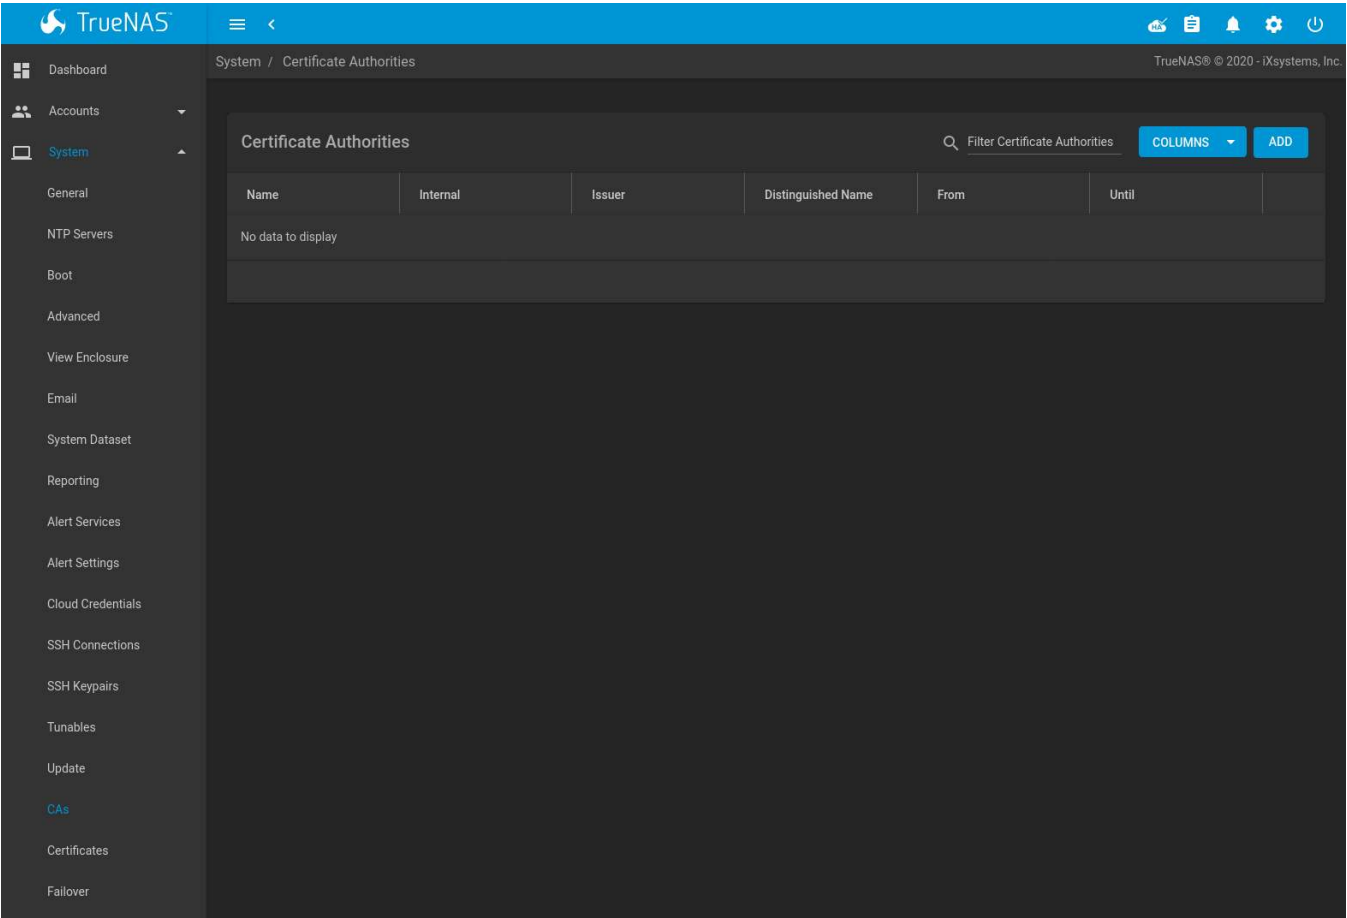


Fig. 5.20: Initial CA Screen

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 5.21. The configurable options are summarized in Table 5.9.

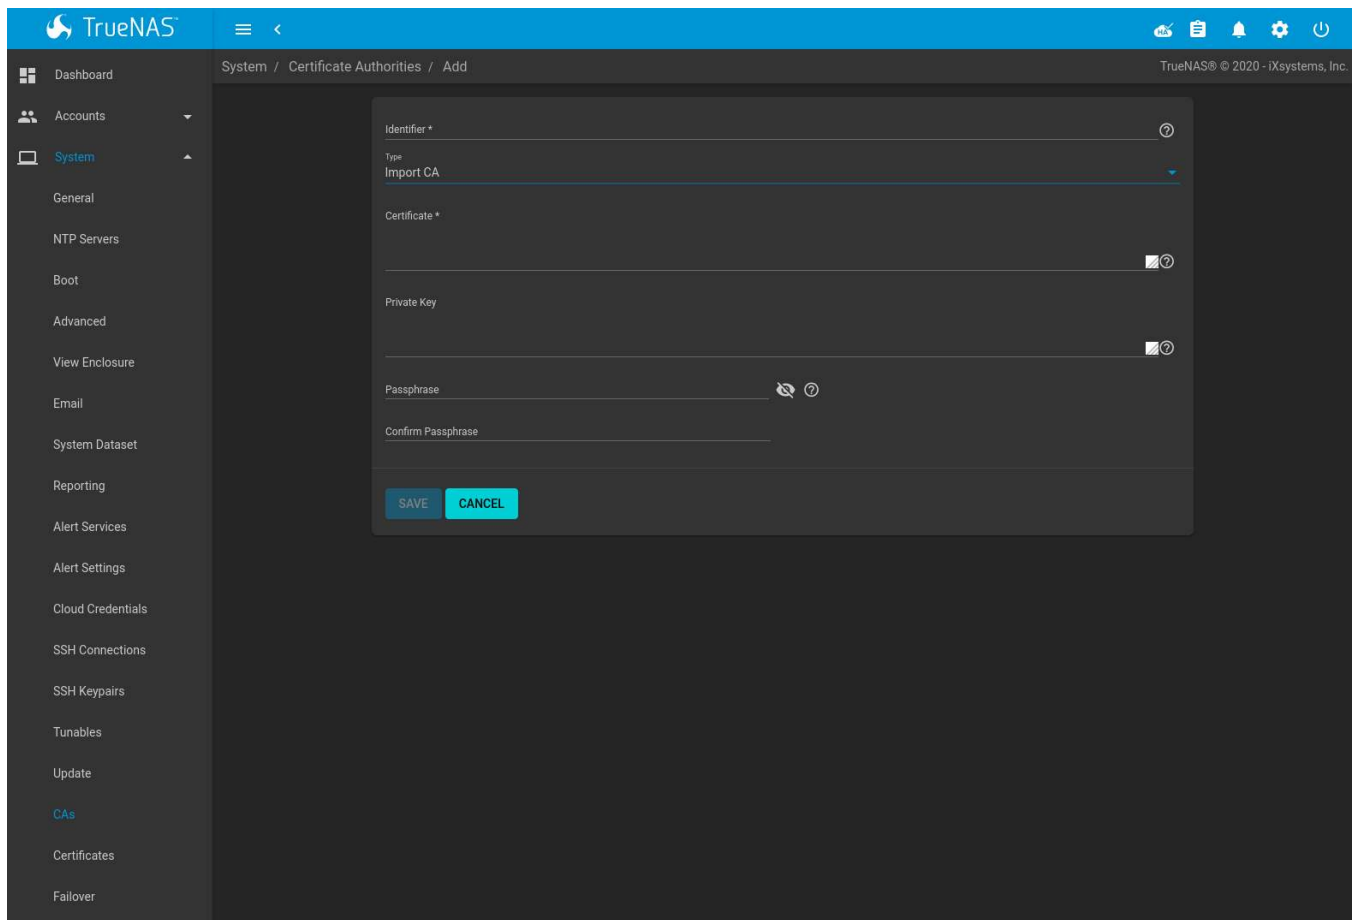


Fig. 5.21: Importing a CA

Table 5.9: 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. Private keys must be at least 1024 bits long.
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 5.22](#) shows the available options.

Fig. 5.22: Creating an Internal CA

The configurable options are described in Table 5.10. When completing the fields for the certificate authority, supply the information for the organization.

Table 5.10: 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 Type	drop-down menu	Cryptosystem for the certificate authority key. Choose between <i>RSA</i> (<i>Rivest-Shamir-Adleman</i> (<i>https://en.wikipedia.org/wiki/RSA_(cryptosystem)</i>)) and <i>EC</i> (<i>Elliptic-curve</i> (<i>https://en.wikipedia.org/wiki/Elliptic-curve_cryptography</i>)) encryption.
EC Curve	drop-down menu	Elliptic curve to apply to the certificate authority key. Choose from different <i>Brainpool</i> or <i>SEC</i> curve parameters. See <i>RFC 5639</i> (<i>https://tools.ietf.org/html/rfc5639</i>) and <i>SEC 2</i> (<i>http://www.secg.org/sec2-v2.pdf</i>) for more details. Applies to <i>EC</i> keys only.
Key Length	drop-down menu	For security reasons, a minimum of <i>2048</i> is recommended. Applies to <i>RSA</i> keys only.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algorithm.

Continued on next page


Table 5.10 – continued from previous page

Setting	Value	Description
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.
Organizational Unit	string	Organizational unit of the entity.
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 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 5.22](#):

- **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*. Signing a request adds a new certificate to *System* → *Certificates*.
- **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 TrueNAS® 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 TrueNAS® system. This option only appears if the CA has a private key.
- **Delete:** prompts for confirmation before deleting the CA.

5.17 Certificates

TrueNAS® can import existing certificates or certificate signing requests, 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 72).

Go to *System* → *Certificates* to add or view certificates.

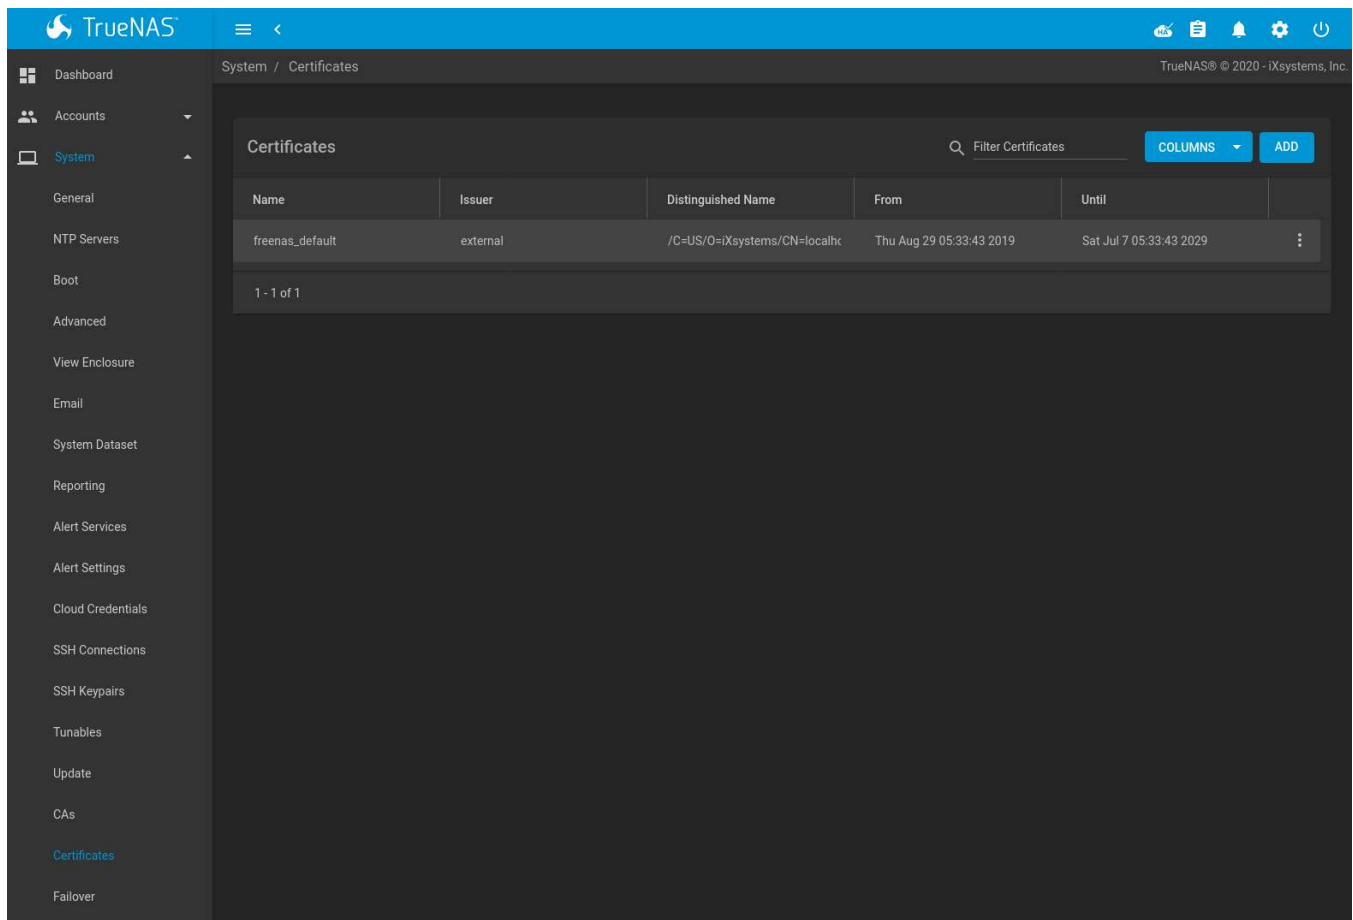


Fig. 5.23: Certificates

TrueNAS® uses a self-signed certificate to enable encrypted access to the web interface. This certificate is generated at boot and cannot be deleted until a different certificate is chosen as the [GUI SSL Certificate](#) (page 33).

To import an existing certificate, click *ADD* and set the *Type* to *Import Certificate*. [Figure 5.24](#) shows the options. When importing a certificate chain, paste the primary certificate, followed by any intermediate certificates, followed by the root CA certificate.

On TrueNAS® [High Availability \(HA\)](#) (page 81) systems, the imported certificate must include the IP addresses or DNS hostnames of both TrueNAS controllers and the CARP virtual IP address. These IP addresses or DNS hostnames can be placed in the *Subject Alternative Name (SAN)* x509 extension field.

The configurable options are summarized in [Table 5.11](#).

The screenshot shows the TrueNAS web interface with the 'System / Certificates / Add' path selected. The 'Import Certificate' form is displayed with the following fields:

- Identifier ***: A text input field.
- Type**: A dropdown menu set to 'Import Certificate'.
- CSR exists on this system**: An unchecked checkbox.
- Signing Certificate Authority**: A dropdown menu.
- Certificate ***: A large text area for pasting the certificate content.
- Private Key**: A large text area for pasting the private key.
- Passphrase**: A text input field with a toggle for visibility.
- Confirm Passphrase**: A text input field for re-entering the passphrase.

At the bottom of the form are 'SAVE' and 'CANCEL' buttons.

Fig. 5.24: Importing a Certificate

Table 5.11: 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> , <i>Import Certificate</i> , and <i>Import Certificate Signing Request</i> .
CSR exists on this system	checkbox	Set when the certificate being imported already has a Certificate Signing Request (CSR) on the system.
Signing Certificate Authority	drop-down menu	Select a previously created or imported CA. Active when <i>CSR exists on this system</i> is set.
Certificate	string	Paste the contents of the certificate.
Private Key	string	Paste the private key associated with the certificate. Private keys must be at least 1024 bits long. Active when <i>CSR exists on this system</i> is unset.
Passphrase	string	If the private key is protected by a passphrase, enter it here and repeat it in the <i>Confirm Passphrase</i> field. Active when <i>CSR exists on this system</i> is unset.

Importing a certificate signing request requires copying the contents of the signing request and key files into the form. Having the signing request `CERTIFICATE REQUEST` and `PRIVATE KEY` strings visible in a separate window simplifies the import process.

Table 5.12: Certificate Signing Request 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> , <i>Import Certificate</i> , and <i>Import Certificate Signing Request</i> .
Signing Request	drop-down menu	Paste the CERTIFICATE REQUEST string from the signing request.
Private Key	string	Paste the private key associated with the certificate signing request. Private keys must be at least 1024 bits long.
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 5.25](#). The configurable options are summarized in [Table 5.13](#). 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 72) as the signing authority.

The screenshot shows the TrueNAS web interface for adding a new certificate. The left sidebar contains navigation links: Dashboard, Accounts, System (selected), General, NTP Servers, Boot, Advanced, View Enclosure, Email, System Dataset, Reporting, Alert Services, Alert Settings, Cloud Credentials, SSH Connections, SSH Keypairs, Tunables, Update, CAs, Certificates (highlighted), and Failover. The main content area is titled 'System / Certificates / Add'. The form includes the following fields:

- Identifier *
- Type: Internal Certificate (dropdown)
- Signing Certificate Authority *
- Key Type *: RSA (dropdown)
- EC Curve: BrainpoolP384R1 (dropdown)
- Key Length *: 2048 (dropdown)
- Digest Algorithm *: SHA256 (dropdown)
- Lifetime *: 3650
- Country *: United States (dropdown)
- State *
- Locality *
- Organization *
- Organizational Unit
- Email *
- Common Name *
- Subject Alternate Names

At the bottom of the form are 'SAVE' and 'CANCEL' buttons.

Fig. 5.25: Creating a New Certificate

Table 5.13: 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 <i>CAs</i> (page 72).
Key Type	drop-down menu	Cryptosystem for the certificate key. Choose between <i>RSA</i> (<i>Rivest-Shamir-Adleman</i> (<i>https://en.wikipedia.org/wiki/RSA_(cryptosystem)</i>)) and <i>EC</i> (<i>Elliptic-curve</i> (<i>https://en.wikipedia.org/wiki/Elliptic-curve_cryptography</i>)) encryption.
EC Curve	drop-down menu	Elliptic curve to apply to the certificate key. Choose from different <i>Brainpool</i> or <i>SEC</i> curve parameters. See <i>RFC 5639</i> (<i>https://tools.ietf.org/html/rfc5639</i>) and <i>SEC 2</i> (<i>http://www.secg.org/sec2-v2.pdf</i>) for more details. Applies to <i>EC</i> keys only.
Key Length	drop-down menu	For security reasons, a minimum of <i>2048</i> is recommended. Applies to <i>RSA</i> keys only.
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.
Organizational Unit	string	Organizational unit of the entity.
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 5.25](#) display, but without the *Signing Certificate Authority* and *Lifetime* fields.

Certificates that are imported, self-signed, or for which a certificate signing request is created are added as entries to *System* → *Certificates*. In the example shown in [Figure 5.26](#), 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.

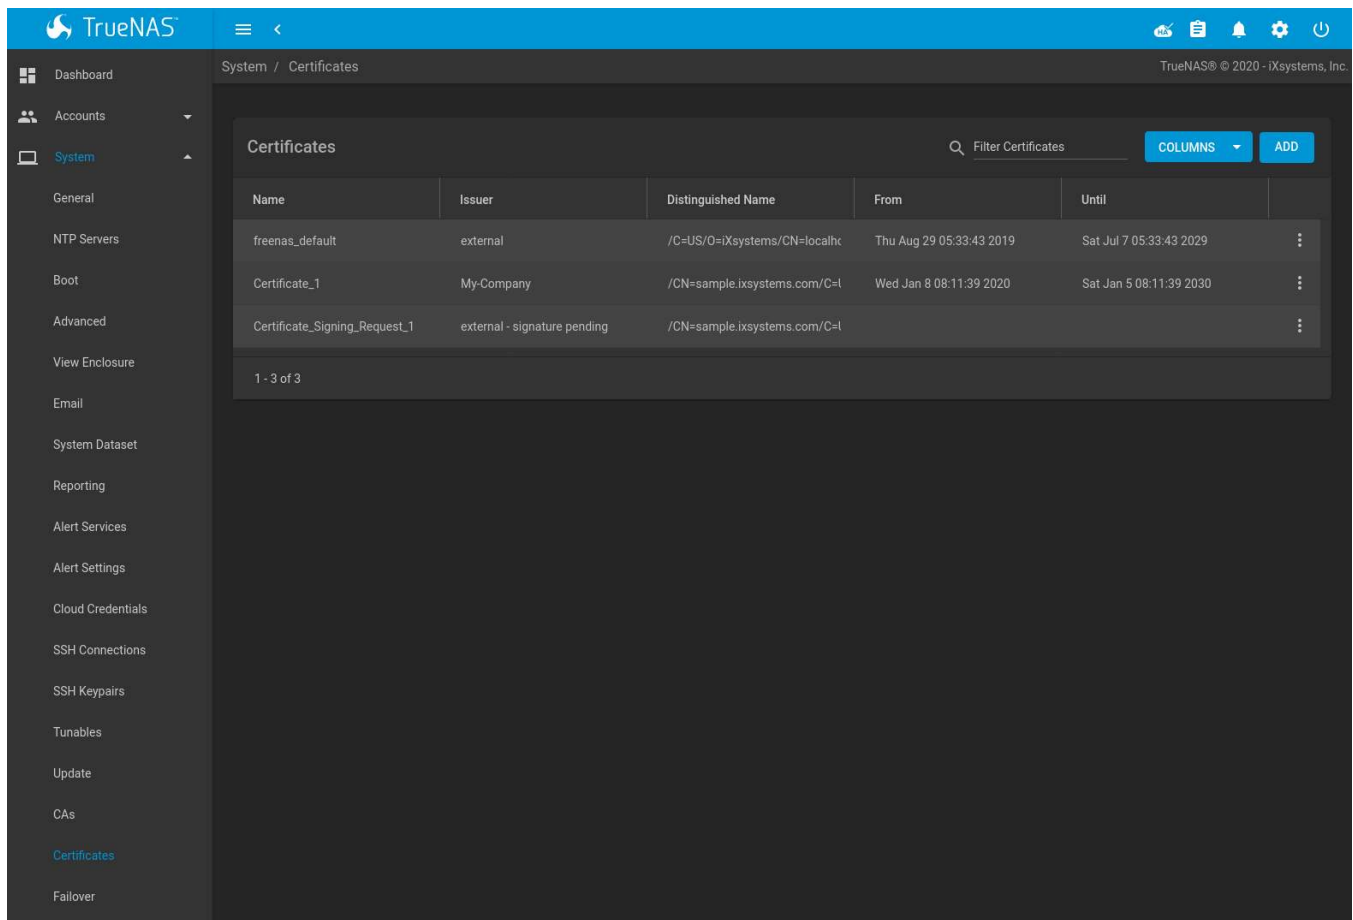



Fig. 5.26: Managing Certificates

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 TrueNAS® 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 TrueNAS® system.
- **Delete** is used to delete a certificate or certificate signing request.

5.18 Failover

Warning: to avoid the potential for data loss, iXsystems must be contacted **before** replacing a controller or upgrading to High Availability.

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 8:00AM to 5:00PM Pacific Standard Time: <ul style="list-style-type: none"> • 1 (855) 473-7449 option 2 (US-only toll-free) • 1 (408) 943-4100 option 2 (local and international)
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none"> • 1 (855) 499-5131 (US-only toll-free) • 1 (678) 835-6101 (local and international)

When the TrueNAS® array has been licensed for High Availability (HA), a *Failover* option appears in *System*.

TrueNAS® uses an active/standby configuration of dual TrueNAS controllers for HA. Dual-ported disk drives are connected to both TrueNAS controllers simultaneously. One TrueNAS controller is active, the other standby. The active TrueNAS controller sends periodic announcements to the network. If a fault occurs and the active TrueNAS controller stops sending the announcements, the standby TrueNAS controller detects this and initiates a failover. Storage and cache devices are imported on the standby TrueNAS controller, then I/O operations switch over to it. The standby TrueNAS controller then becomes the active TrueNAS controller. This failover operation can happen in seconds rather than the minutes of other configurations, significantly reducing the chance of a client timeout.

Note: Seamless failover is only available with iSCSI or NFSv4. Other system services do fail over, but the connections are briefly disrupted by the event.

The Common Address Redundancy Protocol ([CARP](http://www.openbsd.org/faq/pf/carp.html) (<http://www.openbsd.org/faq/pf/carp.html>)) is used to provide high availability and failover. CARP was originally developed by the OpenBSD project and provides an open source, non patent-encumbered alternative to the VRRP and HSRP protocols.

To configure HA, turn on both TrueNAS controllers. Use the IP address shown in the *Console Setup Menu* (page 15) to access the web interface of one of the TrueNAS controllers units. Either TrueNAS controller can be used to configure HA. The *Upload License* dialog is shown on the first login. Otherwise, go to *System* → *Support* → *Upload License*.

Paste the HA license received from iXsystems and press *SAVE LICENSE* to activate it. The license contains the serial numbers for both units in the chassis.

Activating the license adds the *Failover* option to *System*. Some fields are modified in *Network* so that the peer IP address, peer hostname, and virtual IP can be configured. An extra drop-down is added to *IPMI* to allow configuring *IPMI* (page 126) for each TrueNAS controller. The *Dashboard* also updates to add an entry for the standby TrueNAS controller. This entry includes a button to manually initiate a failover.

Fields modified by activating the HA license use 1, 2, or active/standby to identify the TrueNAS controllers. The numbers correspond to the TrueNAS controller labels on the TrueNAS® chassis.

To [configure HA networking](#) (page 119), go to *Network* → *Global Configuration*. The *Hostname* field is replaced by two fields:

- *Hostname*: enter the hostname to use for TrueNAS Controller 1.
- *Hostname (TrueNAS Controller 2)*: enter the hostname to use for TrueNAS controller 2.

Next, go to *Network* → *Interfaces* and click *ADD*. The HA license adds several fields to the *Interfaces* (page 121) screen:

- *Critical*: set this option when a failover should occur if this interface becomes unavailable. How many seconds it takes for the failover to occur depends on the *Timeout* value, as described in [Table 5.14](#). This option is interface-specific, allowing different settings for a management network and a data network. Setting this option requires the *Virtual IP* to be set and that at least one interface needs to be set as *Critical* to configure failover.
- *Failover Group*: allows grouping multiple, critical-for-failover interfaces. Groups apply to single systems. A failover occurs when every interface in the group fails. Groups with a single interface trigger a failover when that interface fails. Configuring the system to failover when any interface fails requires marking each interface as critical and placing them in separate groups.
- *Failover VHID*: use a unique Virtual Host ID (VHID) on the broadcast segment of the network. Configuring multiple Virtual IP addresses requires a separate VHID for each address.
- *IP Address (This Controller)*: a static IP address and netmask. Required when TrueNAS controller 1 is not using DHCP.
- *Failover IP Address (TrueNAS Controller 2)*: a static IP address and netmask. Required when TrueNAS controller 2 is not using DHCP.
- *Virtual IP Address*: enter the IP address to use for administrative access to the array. The netmask /32 is reserved for this value and cannot be changed.

After the network configuration is complete, log out and log back in, this time using the virtual IP address. Pools and shares can now be configured as usual and configuration automatically synchronizes between the active and standby TrueNAS controllers.

All subsequent logins should use the virtual IP address. Connecting directly to the standby TrueNAS controller with a browser does not allow web interface logins. The screen shows the HA status, TrueNAS controller state, and the configuration management virtual IP address. After HA is configured, an *HA Enabled* icon appears in the upper-right section of the web interface

When HA is disabled by the system administrator, the status icon changes to *HA Disabled*. If the standby TrueNAS controller is not available because it is powered off, still starting up, disconnected from the network, or if failover has not been configured, the status icon changes to *HA Unavailable*.

The remaining failover options are found in *System* → *Failover*.

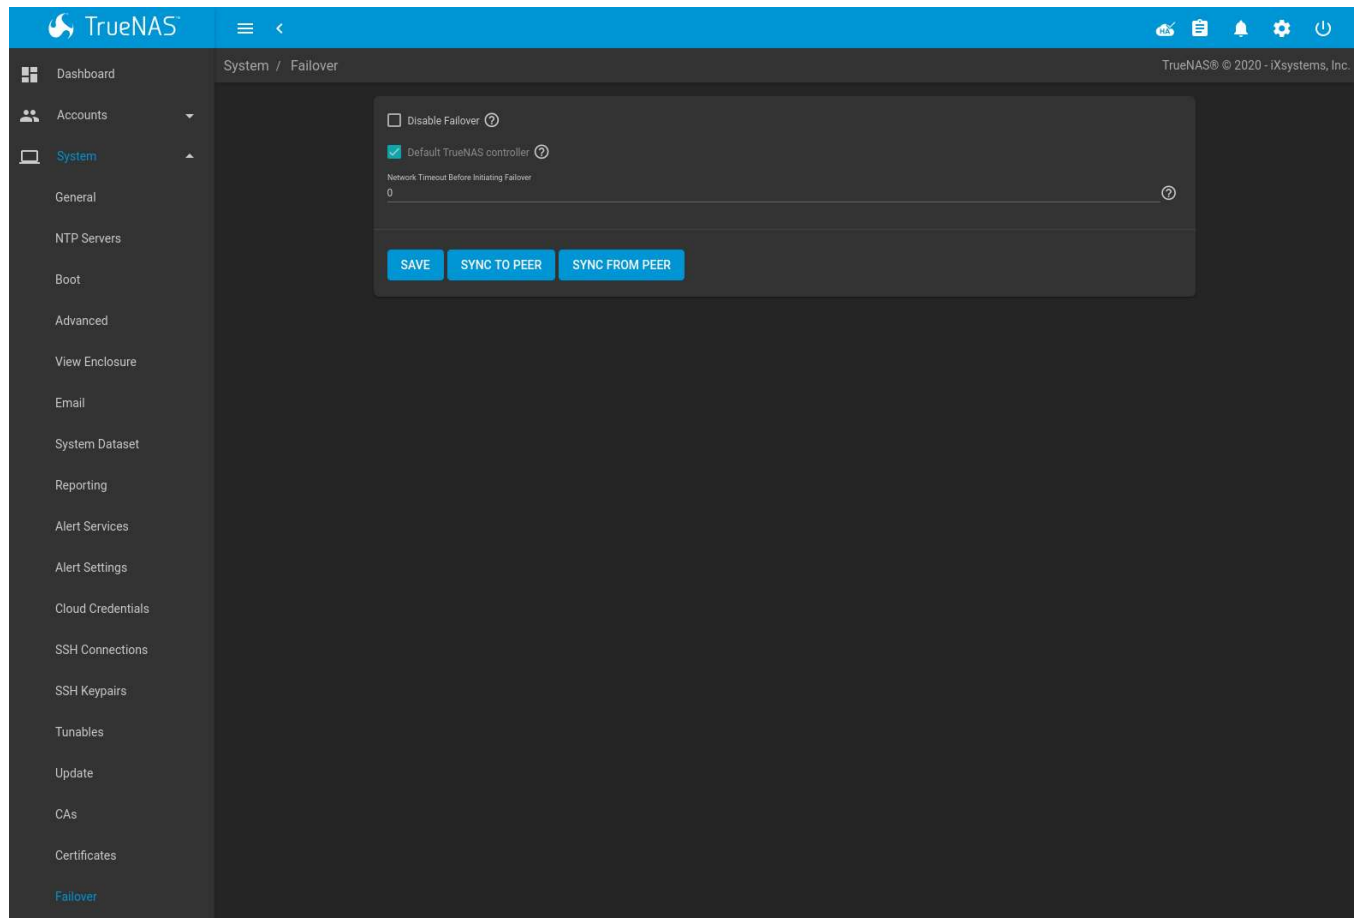


Table 5.14: Failover Options

Setting	Value	Description
Disabled	checkbox	Disables failover. Activates the <i>Master</i> checkbox. The <i>HA Enabled</i> icon changes to <i>HA Disabled</i> . An error message is generated if the standby TrueNAS controller is not responding or failover is not configured.
Master	checkbox	Only available when <i>Disabled</i> is set. Set to mark the current active TrueNAS controller as <i>primary</i> . The <i>primary</i> TrueNAS controller is the default active TrueNAS controller when both TrueNAS controllers are online and HA is enabled. To change which TrueNAS controller is <i>primary</i> , unset this option and allow TrueNAS® to fail over. This will briefly disrupt system services.
Timeout	integer	Number of seconds to wait after a network failure before triggering a failover. <i>0</i> indicates that a failover either occurs immediately or after two seconds when the system is using a link aggregation.
SYNC TO PEER	button	Force synchronizing the TrueNAS® configuration from the active TrueNAS controller to the standby TrueNAS controller. The standby TrueNAS controller must be rebooted after the synchronization is complete to load the new configuration. Synchronization occurs automatically in TrueNAS® and this option is only used when troubleshooting HA configurations. Do not use this unless requested by an iXsystems Support Engineer.
SYNC FROM PEER	button	Force synchronizing the TrueNAS® configuration from the standby TrueNAS controller to the active TrueNAS controller. Synchronization occurs automatically in TrueNAS® and this option is only used when troubleshooting HA configurations. Do not use this unless requested by an iXsystems Support Engineer.

Notes about High Availability and failovers:

Booting an HA pair with failover disabled causes both TrueNAS controllers to come up in standby mode. The web interface shows an additional *Force Takeover* button which can be used to force that TrueNAS controller to take control.

Failover is not allowed if both TrueNAS controllers have the same CARP state. A critical [Alert](#) (page 305) is generated and the HA icon shows *HA Unavailable*.

The TrueNAS® version of the `ifconfig` command adds two additional fields to the output to help with failover troubleshooting: `CriticalGroupn` and `Interlink`.

If both TrueNAS controllers reboot simultaneously, the GELI passphrase for an [encrypted](#) (page 133) pool must be entered at the web interface login screen.

If there are a different number of disks connected to each TrueNAS controller, an [Alert](#) (page 305) is generated and the HA icon switches to *HA Unavailable*.

5.19 Support

The TrueNAS® *Support* page, shown in [Figure 5.27](#), is used to view or update the system license information, activate [Proactive Support](#) (page 85), or generate [Support requests](#) (page 86).

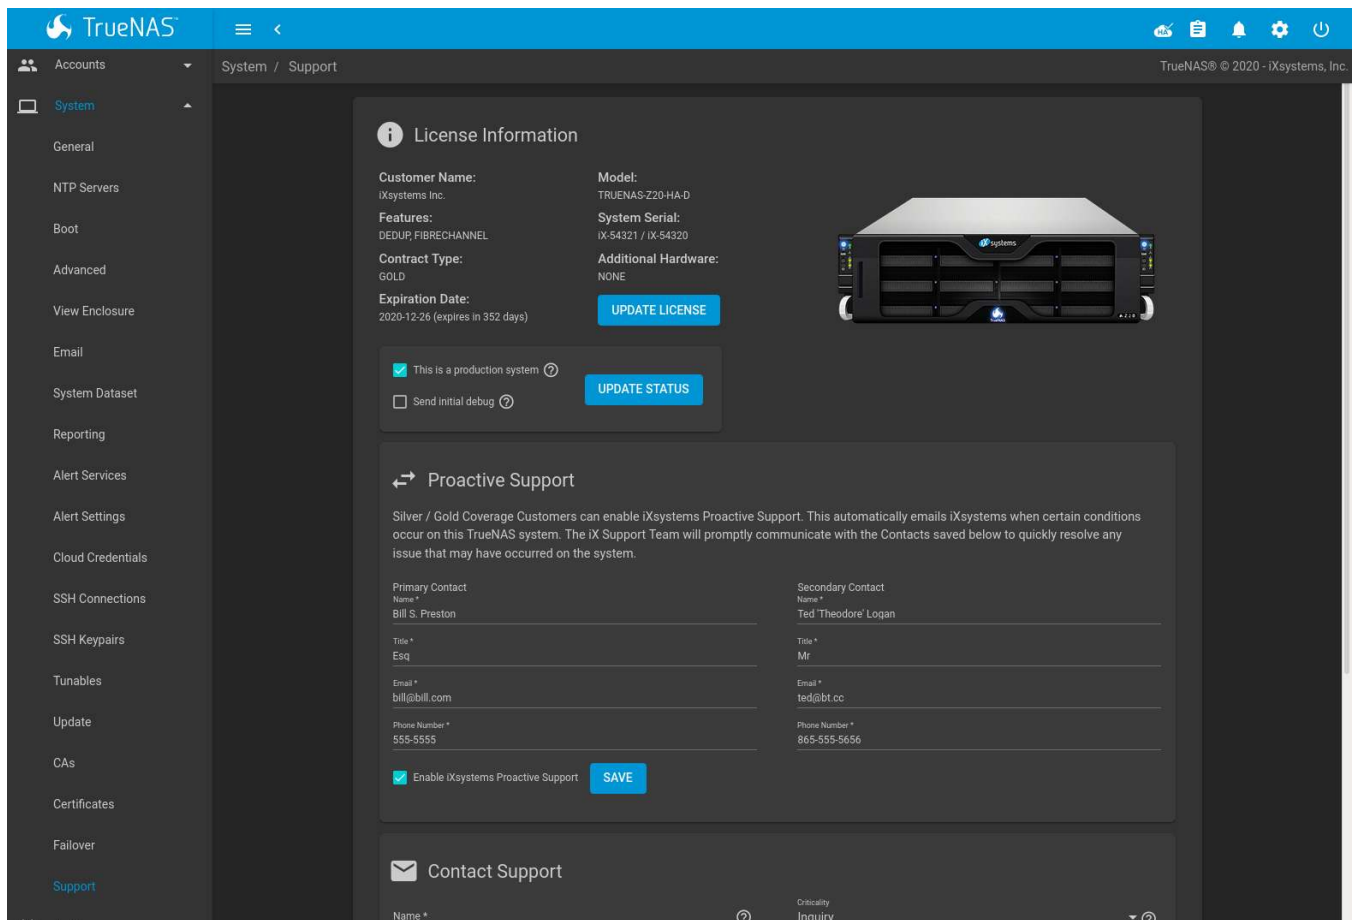


Fig. 5.27: Support Options

5.19.1 License Information

Systems with a valid license display the hardware model, system serial number, support contract type, licensed period, customer name, licensed features, and additional supported hardware.

If the license expires or additional hardware, features, or contract type are required, [contact iXsystems Support](#) (page 10). After a new license has been provided, click *UPDATE LICENSE*, paste in the new license, and click *SAVE LICENSE*. An additional dialog prompts to reload the web interface and show the new license details.

There are also options to mark the system for production use or to send an initial debug to iXsystems. To update the status, set either option and click *UPDATE STATUS*.

5.19.2 Proactive Support

The Proactive Support feature can notify iXsystems by email when hardware conditions on the system require attention.

Note: The fields on this tab are only enabled for Silver and Gold support coverage level customers. Please [contact iXsystems](#) (page 10) for information on upgrading from other support levels.

Before enabling proactive support, provide primary and secondary contact information. This ensures iX Support can promptly communicate and quickly resolve any issues.

To enable proactive support, make sure all contact information is correct, set *Enable iXsystems Proactive Support*, and click *SAVE*.

5.19.3 Contact Support

To generate a support ticket, fill in the fields:

- *Name* is the name of the person the iXsystems Support Representative should contact to assist with the issue.
- *Email* is the email address of the person to contact.
- *Phone* is the phone number of the person to contact.
- *Type* is a drop-down menu to select the ticket type: a software bug, a hardware failure, a request for help with installing or configuring the system, or a request for help with diagnosing a performance bottleneck.
- *Environment* is a drop-down menu to indicate the role of the affected system.

Table 5.15: Environment Options

Environment	Description
Production	This is a production system in daily use.
Staging	The system is being prepared for production.
Testing	This system is only being used for testing purposes.
Prototyping	The system is unique. It is likely to be a proof of concept.
Initial Deployment/ Setup	This is a new system being prepared for deployment into production.

- *Criticality* is a drop-down menu to indicate how the issue has affected the TrueNAS® system. Choices are *Inquiry*, *Loss of Functionality*, or *Total Down*.
- *Attach Debug* is an option to include an overview of the system hardware, build string, and configuration with the ticket. Generating and attaching a debug to the ticket can take some time.

Debug file attachments are limited to 20 MiB. If the debug file is too large to include, unset the option to generate the debug file and let the system create an issue ticket as shown below. Manually create a debug file by going to *System* → *Advanced* and clicking *SAVE DEBUG*.

Go to the ticket at [iXsystems Support](https://support.ixsystems.com/) (<https://support.ixsystems.com/>) and upload the debug file.

- *Subject* is a descriptive title for the ticket.
- *Description* is a one- to three-paragraph summary of the issue that describes the problem, and if applicable, steps to reproduce it.
- *Attach screenshots* is an optional field where screenshots of any errors or tracebacks can be included.

Click *SUBMIT* to generate and send the support ticket to iXsystems. This process can take several minutes while information is collected and sent. TrueNAS® sends an email alert if ticket creation fails while Proactive Support is active.

After the new ticket is created, the URL is shown for viewing or updating with more information. An [iXsystems Support](https://support.ixsystems.com/) (<https://support.ixsystems.com/>) account is required to view the ticket. Click the URL to log in or register with the support portal. Use the same e-mail address submitted with the ticket when registering.

UPDATE LICENSE functions identically to the button in the [License Information](#) (page 85) section.

USER GUIDE opens a new browser tab to the iXsystems TrueNAS® [Information Library](https://www.ixsystems.com/blog/knowledgebase_category/truenas/) (https://www.ixsystems.com/blog/knowledgebase_category/truenas/). The TrueNAS® User Guide, product datasheets, TrueNAS® hardware setup guides, and task assistance articles are all available in this library.

EULA shows the TrueNAS® End User License Agreement.

TASKS

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

- [Cron Jobs](#) (page 87) schedules a command or script to automatically execute at a specified time
- [Init/Shutdown Scripts](#) (page 89) configures a command or script to automatically execute during system startup or shutdown
- [Rsync Tasks](#) (page 90) schedules data synchronization to another system
- [S.M.A.R.T. Tests](#) (page 96) schedules disk tests
- [Periodic Snapshot Tasks](#) (page 97) schedules automatic creation of filesystem snapshots
- [Replication Tasks](#) (page 107) automate the replication of snapshots to a remote system
- [Resilver Priority](#) (page 110) controls the priority of resilvers
- [Scrub Tasks](#) (page 111) schedules scrubs as part of ongoing disk maintenance
- [Cloud Sync Tasks](#) (page 113) schedules data synchronization to cloud providers

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

Note: By default, [Scrub Tasks](#) (page 111) are run once a month by an automatically-created task. [S.M.A.R.T. Tests](#) (page 96) and [Periodic Snapshot Tasks](#) (page 97) must be set up manually.

6.1 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.

Go to *Tasks* → *Cron Jobs* and click *ADD* to create a cron job.

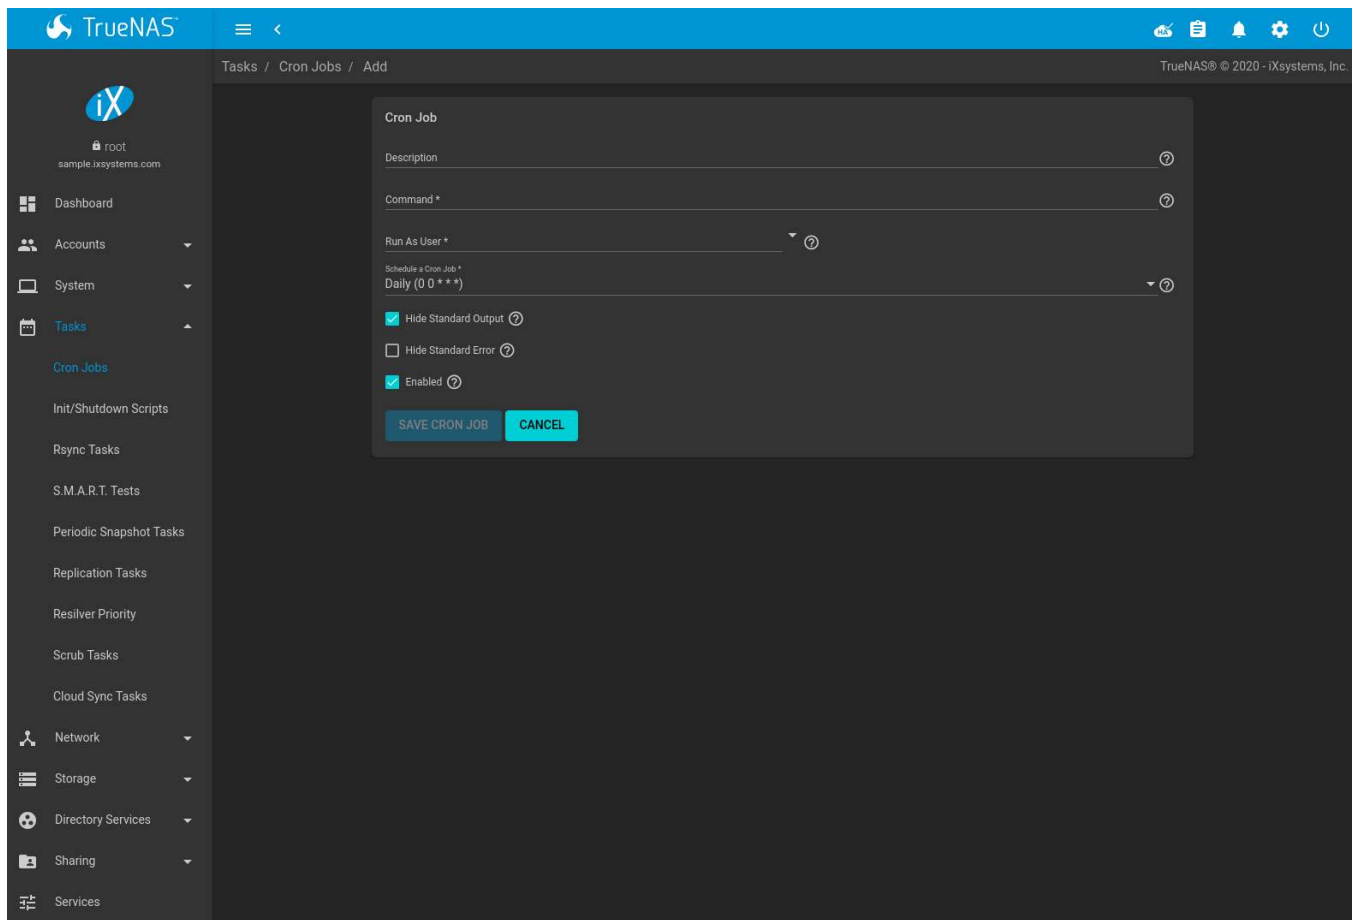


Fig. 6.1: Cron Job Settings

Table 6.1 lists the configurable options for a cron job.

Table 6.1: 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 first is recommended.
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. Output from executing a cron task is emailed to this user if <i>Email</i> has been configured for that <i>user account</i> (page 27).
Schedule	drop-down menu	Select a schedule preset or choose <i>Custom</i> to open the advanced scheduler. Note that an in-progress cron task postpones any later scheduled instance of the same task until the running task is complete.
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	Set to allow this scheduled cron task to activate. Unsetting this option disables the cron task 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.

6.2 Init/Shutdown Scripts

TrueNAS® provides the ability to schedule commands or scripts to run at system startup or shutdown. Go to *Tasks* → *Init/Shutdown Scripts* and click *ADD*.

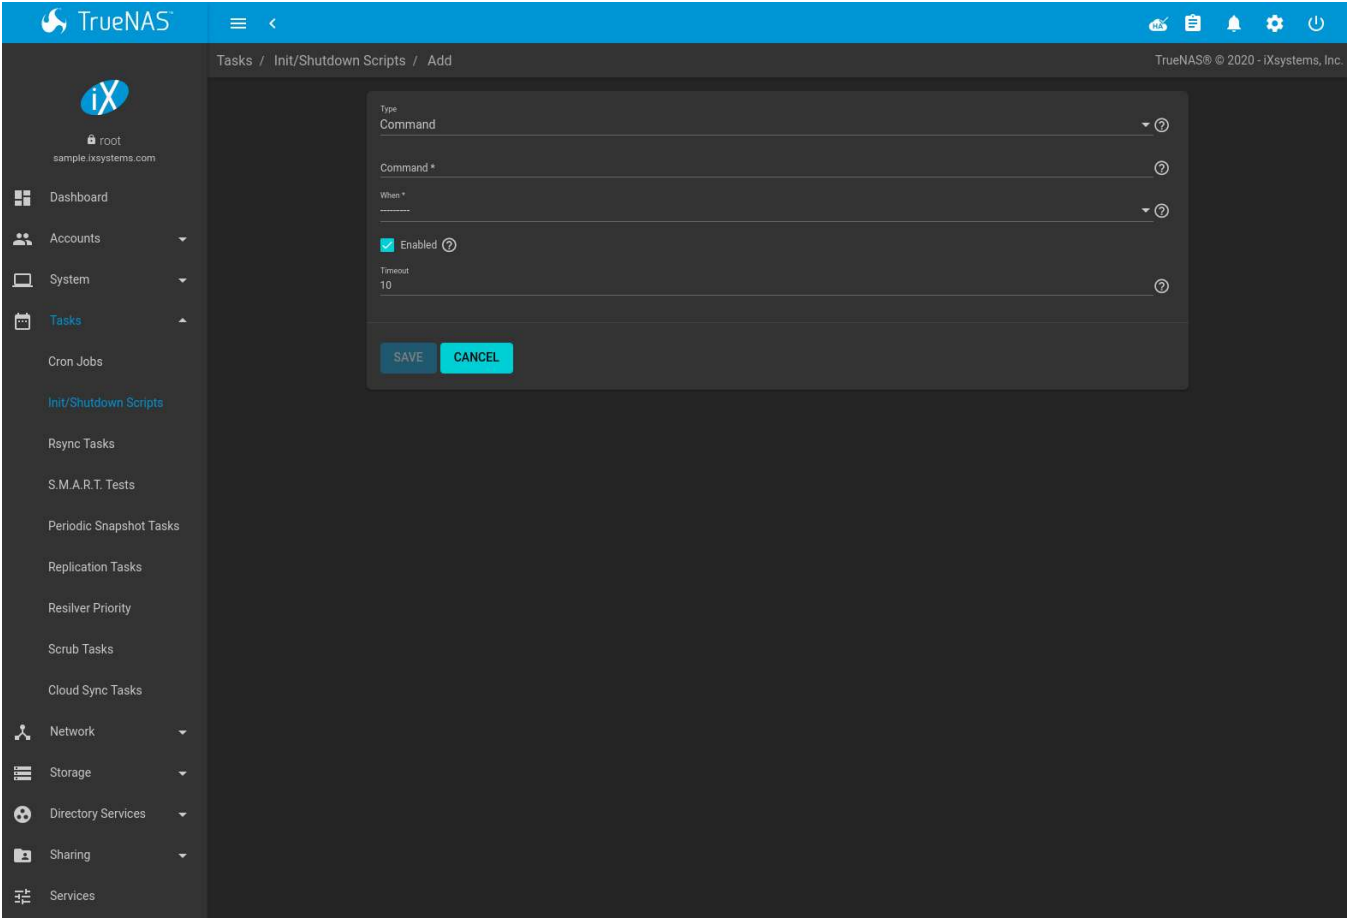


Fig. 6.2: Add an Init/Shutdown Command or Script

Table 6.2: Init/Shutdown Command or Script Options

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 with any options. When <i>Script</i> is selected, click ■ (Browse) to select the script from an existing pool.

Continued on next page

Table 6.2 – continued from previous page

Setting	Value	Description
When	drop-down menu	Select when the <i>Command</i> or <i>Script</i> runs: <ul style="list-style-type: none"> • <i>Pre Init</i>: early in the boot process, after mounting filesystems and starting networking • <i>Post Init</i>: at the end of the boot process, before TrueNAS® services start • <i>Shutdown</i>: during the system power off process.
Enabled	checkbox	Enable this task. Unset to disable the task without deleting it.
Timeout	integer	Automatically stop the script or command after the specified number of seconds.

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 with `which {commandname}` in the *Shell* (page 302). When available, the path to the command is shown:

```
[root@freenas ~]# which ls
/bin/ls
```

When scheduling a script, test the script first to verify it is executable and achieves the desired results.

Note: Init/shutdown scripts are run with `sh`.

Init/Shutdown tasks are shown in *Tasks* → *Init/Shutdown Scripts*. Click **:** (Options) for a task to *Edit* or *Delete* that task.

6.3 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 107) 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 302).

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.

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

TrueNAS® supports two modes of *rsync* operation:

- **Module:** 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 TrueNAS® 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) (<https://www.samba.org/ftp/rsync/rsyncd.conf.html>).

- **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 TrueNAS® systems for each mode of rsync operation.

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

Figure 6.3 shows the screen that appears after navigating to *Tasks* → *Rsync Tasks* and clicking *ADD*. Table 6.3 summarizes the configuration options available when creating an rsync task.

Fig. 6.3: Adding an Rsync Task

Table 6.3: Rsync Configuration Options


Setting	Value	Description
Path	browse button	<i>Browse</i> to the path to be copied. TrueNAS® verifies that the remote path exists. <i>FreeBSD path length limits</i> (page 10) apply on the TrueNAS® system. Other operating systems can have different limits which might affect how they can be used as sources or destinations.
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.

Continued on next page

Table 6.3 – continued from previous page

Setting	Value	Description
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 user-name differs on the remote host.
Remote SSH Port	integer	Only available in <i>SSH</i> mode. Allows specifying an SSH port other than the default of 22.
Rsync mode	drop-down menu	The choices are <i>Module</i> mode or <i>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>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> .
Direction	drop-down menu	Direct the flow of the data to the remote host. Choices are <i>Push</i> or <i>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</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Selecting <i>Custom</i> opens the Advanced Scheduler (page 12).
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	Suppress rsync task status alerts (page 305).
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 rsync(1) (http://rsync.samba.org/ftp/rsync/rsync.html) options to include. Note: The * character must be escaped with a backslash (<code>*.txt</code>) 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*.

The *Status* column shows the status of the rsync task. To view the detailed rsync logs for a task, click the *Status* entry when the task is running or finished.

Rsync tasks also generate an [Alert](#) (page 305) on task completion. The alert shows if the task succeeded or failed.

6.3.1 Rsync Module Mode

This configuration example configures rsync module mode between the two following TrueNAS® 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 *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

Descriptions of the configurable options can be found in [Rsync Modules](#) (page 235).

- *Hosts allow* is set to 192.168.2.2, the IP address of the rsync client

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/`.

6.3.2 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 302) 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        |
|-----|
|
|-----|

```

TrueNAS® 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 AAAAB3NzaC1yc2EAAAADAQABAAQAC1lBEXRgw1W8y8k+lXP1VR3xsmVSjtsoyIzV/PlQPoSrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4hdcD7Y5mvU3MAEeDC1t02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmasti00qmDDcp/kOxT+S6DFNDBY6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+l9RaEXMRuTyQgqJB/rsRcmJX5fApdDmNfwrRSxLjDvUzfywnjFh1Kk/+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 6.4](#). When pasting the key, ensure that it is pasted as one long line and, if necessary, remove any extra spaces representing line breaks.

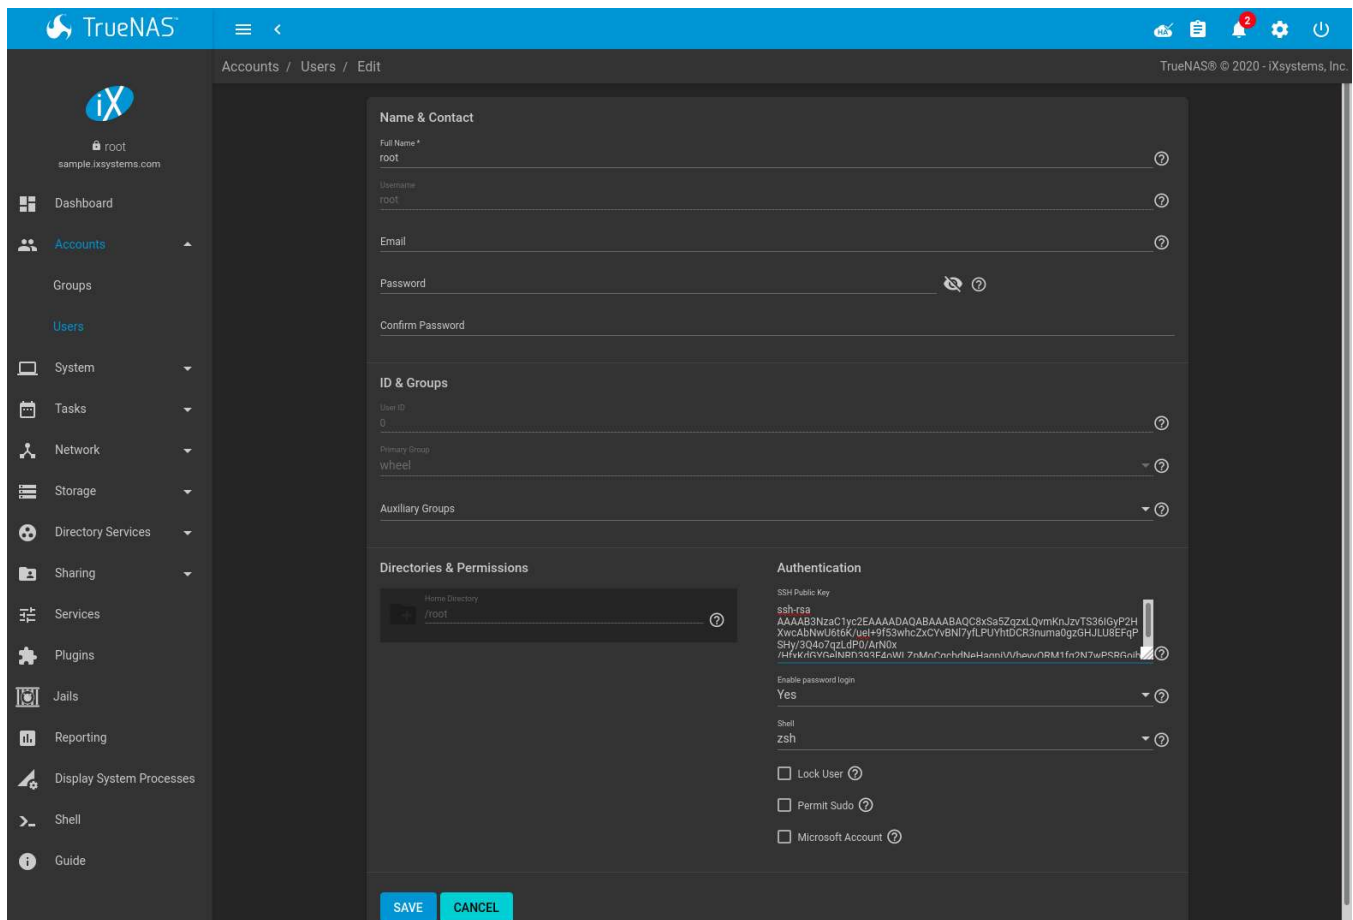


Fig. 6.4: Pasting the User SSH Public Key

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 *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.

6.4 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 6.5](#) 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.

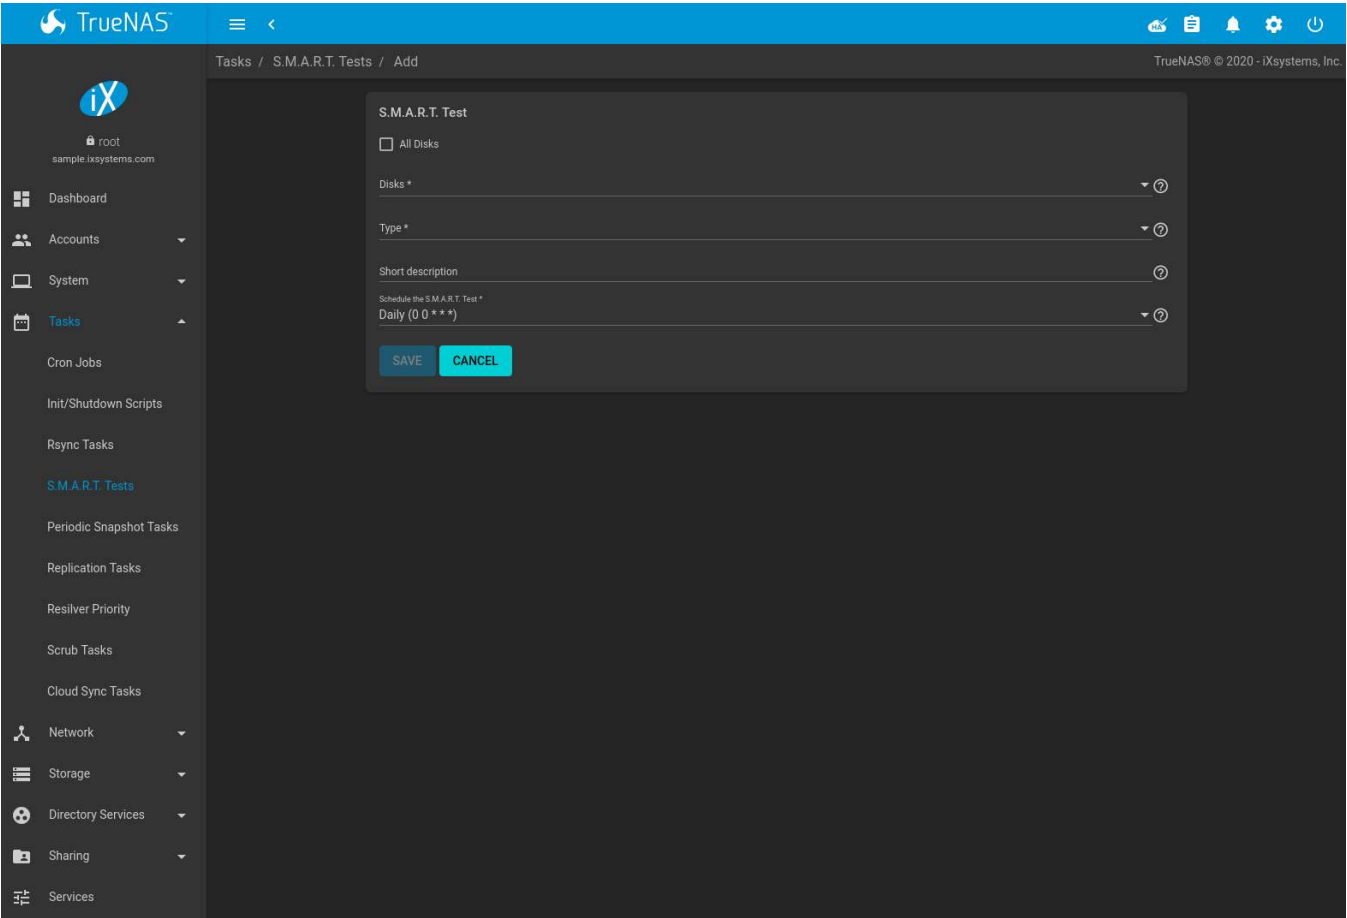


Fig. 6.5: Adding a S.M.A.R.T. Test

[Table 6.4](#) summarizes the configurable options when creating a S.M.A.R.T. test.

Table 6.4: S.M.A.R.T. Test Options

Setting	Value	Description
All Disks	checkbox	Set to monitor all disks.
Disks	drop-down menu	Select the disks to monitor. Available when <i>All Disks</i> is unset.
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> . Selecting <i>Custom</i> opens the Advanced Scheduler (page 12).

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 302).

The results of a test can be checked from [Shell](#) (page 302) by specifying the name of the drive. For example, to see the results for disk *ada0*, type:

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

6.5 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 107). 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 130).

View the list of periodic snapshot tasks by going to *Tasks* → *Periodic Snapshot Tasks*. If a periodic snapshot task encounters an error, the status column will show *ERROR*. Click the status to view the logs of the task.

To create a periodic snapshot task, navigate to *Tasks* → *Periodic Snapshot Tasks* and click *ADD*. This opens the screen shown in [Figure 6.6](#). [Table 6.5](#) describes the fields in this screen.

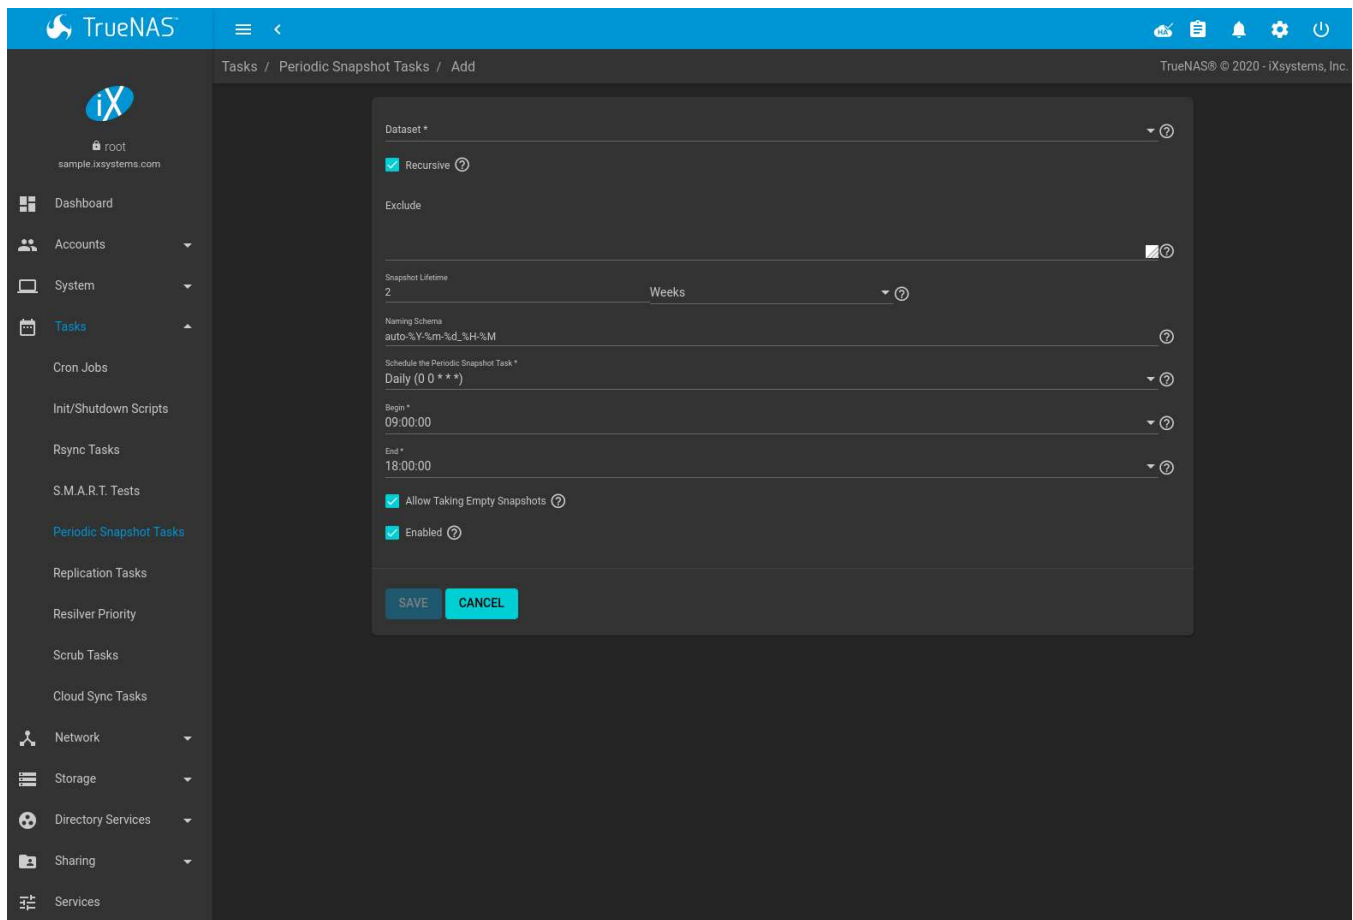


Fig. 6.6: Creating a Periodic Snapshot

Table 6.5: Periodic Snapshot Options

Setting	Value	Description
Dataset	drop-down menu	Select a pool, dataset, or zvol.
Recursive	checkbox	Set to take separate snapshots of the dataset and each of its child datasets. Leave unset to take a single snapshot only of the specified dataset <i>without</i> child datasets.
Exclude	string	Exclude specific child datasets from the snapshot. Use with recursive snapshots. Comma-separated list of paths to any child datasets to exclude. Example: <code>pool1/dataset1/child1</code> . A recursive snapshot of <code>pool1/dataset1</code> will include all child datasets except <code>child1</code> .
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 which have been replicated to other systems are not affected.
Snapshot Life-time Unit	drop-down	Select a unit of time to retain the snapshot on this system.
Naming Schema	string	Snapshot name format string. The default is <code>auto-%Y-%m-%d_%H-%M</code> . Must include the strings <code>%Y</code> , <code>%m</code> , <code>%d</code> , <code>%H</code> , and <code>%M</code> , which are replaced with the four-digit year, month, day of month, hour, and minute as defined in <code>strftime(3)</code> (https://www.freebsd.org/cgi/man.cgi?query=strftime). For example, snapshots of <code>pool1</code> with a Naming Schema of <code>customsnap-%Y%m%d.%H%M</code> have names like <code>pool1@customsnap-20190315.0527</code> .

Continued on next page

Table 6.5 – continued from previous page

Setting	Value	Description
Schedule the Periodic Snapshot Task	drop-down menu	When the periodic snapshot task runs. Choose one of the preset schedules or choose <i>Custom</i> to use the Advanced Scheduler (page 12).
Begin	drop-down menu	Hour and minute when the system can begin taking snapshots.
End	drop-down menu	Hour and minute the system must stop creating snapshots. Snapshots already in progress will continue until complete.
Allow Taking Empty Snapshots	checkbox	Creates dataset snapshots even when there have been no changes to the dataset from the last snapshot. Recommended for creating long-term restore points, multiple snapshot tasks pointed at the same datasets, or to be compatible with snapshot schedules or replications created in TrueNAS® 11.2 and earlier. For example, allowing empty snapshots for a monthly snapshot schedule allows that monthly snapshot to be taken, even when a daily snapshot task has already taken a snapshot of any changes to the dataset.
Enabled	checkbox	To activate this periodic snapshot schedule, set this option. To disable this task without deleting it, unset this option.

Setting *Recursive* adds child datasets to the snapshot. Creating separate snapshots for each child dataset is not needed.

The *Naming Schema* can be manually adjusted to include more information. For example, after configuring a periodic snapshot task with a lifetime of two weeks, it could be helpful to define a *Naming Schema* that shows the lifetime: `autosnap-%Y-%m-%d.%H-%M-2w`.

Click **SAVE** when finished customizing the task. Defined tasks are listed alphabetically in *Tasks* → *Periodic Snapshot Tasks*.

Click **:** (Options) for a periodic snapshot task to see options to *Edit* or *Delete* the scheduled task.

Deleting a dataset does not delete snapshot tasks for that dataset. To re-use the snapshot task for a different dataset, *Edit* the task and choose the new *Dataset*. The original dataset is shown in the drop-down, but cannot be selected.

Deleting the last periodic snapshot task used by a replication task is not permitted while that replication task remains active. The replication task must be disabled before the related periodic snapshot task can be deleted.

6.5.1 Snapshot Autoremoval

The periodic snapshot task autoremoval process (which removes snapshots after their configured *Snapshot Lifetime*) is run whenever any *Enabled* periodic snapshot task runs.

When the autoremoval process runs, all snapshots on the system are checked for removal. First, each snapshot is matched with a periodic snapshot task according to the following criteria:

- *Dataset/Recursive*: To match a task, a snapshot must be on the same *Dataset* as the task, or on a child dataset if the task is marked *Recursive*.
- *Naming Schema*: To match a task, a snapshot's name must match the *Naming Schema* defined in that task.
- *Schedule*: To match a task, the time at which the snapshot was created (according to its name and naming schema) must match the schedule defined in the task (*Schedule the Periodic Snapshot Task*).
- *Enabled*: To match a task, the periodic snapshot task must be *Enabled*.

At this point, if the snapshot does not match any periodic snapshot tasks then it is not considered for autoremoval. However, if it does match one (or possibly more than one) periodic snapshot task, it is deleted if its creation time (according to its name and naming schema) is older than the longest *Snapshot Lifetime* of any of the tasks it was matched with.

One notable detail of this process is that there is no saved memory of which task created which snapshot, or what the parameters of the periodic snapshot task were at the time a snapshot was created. All checks for autoremoval are based on the current state of the system.

These details become important when existing periodic snapshot tasks are edited, disabled, or deleted. When editing a periodic snapshot task, if the *Naming Schema* is changed, *Recursive* is unchecked, or the task is rescheduled (*Schedule the Periodic Snapshot Task*), previously created snapshots may not be automatically removed as expected since the previously created snapshots may no longer match any periodic snapshot tasks. Similarly, if a periodic snapshot task is deleted or marked not *Enabled*, snapshots previously created by that task will no longer be automatically removed.

In these cases, the user must manually remove unneeded snapshots that were previously created by the modified or deleted periodic snapshot task.

6.6 Replication

Replication is the process of copying *ZFS dataset snapshots* (page 309) from one storage pool to another. Replications can be configured to copy snapshots to another pool on the local system or send copies to a remote system that is in a different physical location.

Replication schedules are typically paired with *Periodic Snapshot Tasks* (page 97) to generate local copies of important data and replicate these copies to a remote system.

Replications require a source system with dataset snapshots and a destination that can store the copied data. Remote replications require a saved *SSH Connection* (page 57) on the source system and the destination system must be configured to allow *SSH* (page 244) connections. Local replications do not use SSH.

Snapshots are organized and sent to the destination according to the creation date included in the snapshot name. When replicating manually created snapshots, make sure snapshots are named according to their actual creation date.

First-time replication tasks can take a long time to complete as the entire dataset snapshot must be copied to the destination system. Replicated data is not visible on the receiving system until the replication task is complete.

Later replications only send incremental snapshot changes to the destination system. This reduces both the total space required by replicated data and the network bandwidth required for the replication to complete.

The replication task asks to destroy destination dataset snapshots when those snapshots are not related to the replication snapshots. Verify that the snapshots in the destination dataset are unneeded or are backed up in a different location! Allowing the replication task to continue destroys the current snapshots in the destination dataset and replicates a full copy of the source snapshots.

The target dataset on the destination system is created in *read-only* mode to protect the data. To mount or browse the data on the destination system, use a clone of the snapshot. Clones are created in *read/write* mode, making it possible to browse or mount them. See *Snapshots* (page 152) for more details.

Replications run in parallel as long as they do not conflict with each other. Completion time depends on the number and size of snapshots and the bandwidth available between the source and destination computers.

Examples in this section refer to the TrueNAS® system with the original datasets for snapshot and replication as *Primary* and the TrueNAS® system that is storing replicated snapshots as *Secondary*.

6.6.1 Replication Creation Wizard

To create a new replication, go to *Tasks* → *Replication Tasks* and click *ADD*.

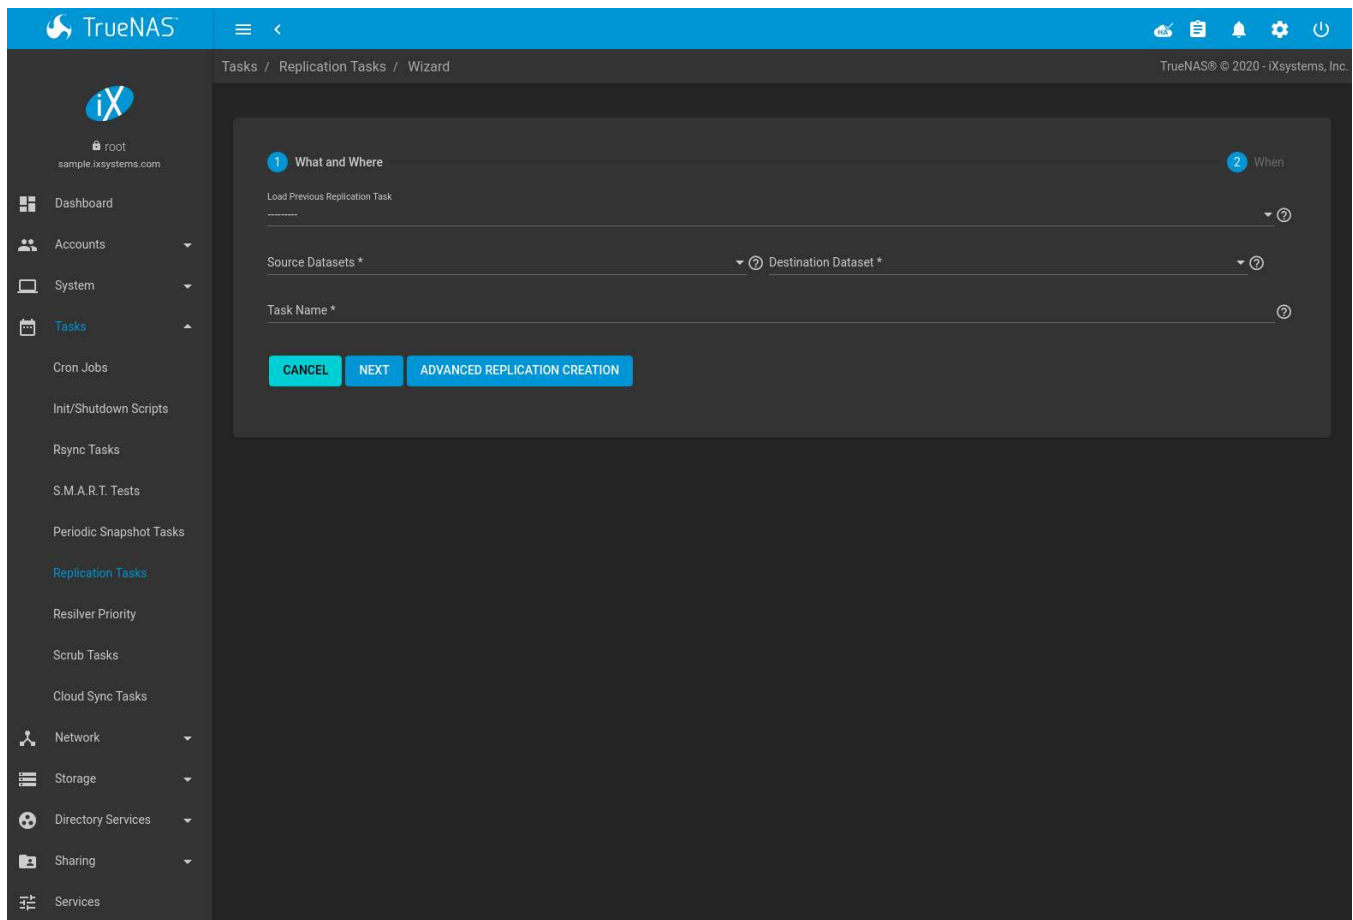



Fig. 6.7: Replication Wizard: What and Where

The wizard allows loading previously saved replication configurations and simplifies many replication settings. To see all possible *replication creation options* (page 103), click *ADVANCED REPLICATION CREATION*.

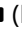
Using the wizard to create a new replication task begins by defining what is being replicated and where. Choosing *On a Different System* for either the *Source Location* or *Destination Location* requires an *SSH Connection* (page 57) to the remote system. Open the drop-down menu to choose an SSH connection or click *Create New* to add a new connection.

Start by selecting the *Source* datasets to be replicated. To choose a dataset, click  (Browse) and select the dataset from the expandable tree. The path of the dataset can also be typed into the field. Multiple snapshot sources can be chosen using a comma (,) to separate each selection. *Recursive* replication will include all snapshots of any descendant datasets of the chosen *Source*.

Source datasets on the local system are replicated using existing snapshots of the chosen datasets. When no snapshots exist, TrueNAS® automatically creates snapshots of the chosen datasets before starting the replication. To manually define which dataset snapshots to replicate, set *Replicate Custom Snapshots* and define a snapshot *Naming Schema*.

Source datasets on a remote system are replicated by defining a snapshot *Naming Schema*. The schema is a pattern of the name and [strftime\(3\)](https://www.freebsd.org/cgi/man.cgi?query=strftime) (<https://www.freebsd.org/cgi/man.cgi?query=strftime>) %Y, %m, %d, %H, and %M strings that match names of the snapshots to include in the replication. For example, to replicate a snapshot named `auto-2019-12-18.05-20` from a remote source, enter `auto-%Y-%m-%d.%H-%M` as the replication task *Naming Schema*.

The number of snapshots that will be replicated is shown. There is also a *Recursive* option to include child datasets with the selected datasets.

Now choose the *Destination* to receive the replicated snapshots. To choose a destination path, click  (Browse)

and select the dataset from the expandable tree or type a path to the location in the field. Only a single *Destination* path can be defined.

Using an SSH connection for replication adds the *SSH Transfer Security* option. This sets the data transfer security level. The connection is authenticated with SSH. Data can be encrypted during transfer for security or left unencrypted to maximize transfer speed. **WARNING:** Encryption is recommended, but can be disabled for increased speed on secure networks.

A suggested replication *Task Name* is shown. This can be changed to give a more meaningful name to the task. When the source and destination have been set, click *NEXT* to choose when the replication will run.

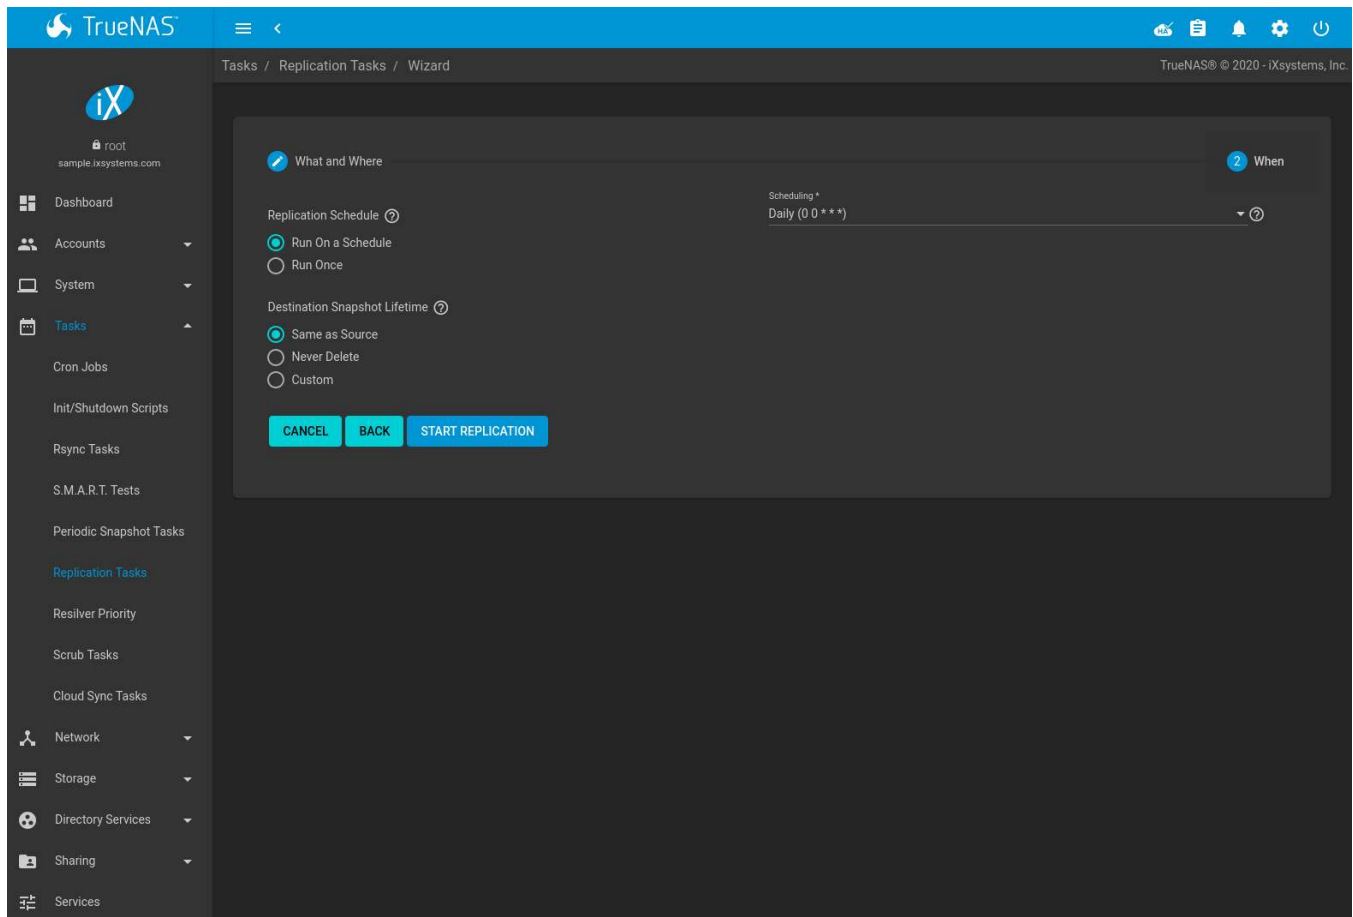


Fig. 6.8: Replication Wizard: When

The replication task can be configured to run on a schedule or left unscheduled and manually activated. Choosing *Run On a Schedule* adds the *Scheduling* drop-down to choose from preset schedules or define a *Custom* replication schedule. Choosing *Run Once* removes all scheduling options.

Destination Snapshot Lifetime determines when replicated snapshots are deleted from the destination system:

- *Same as Source*: duplicate the configured *Snapshot Lifetime* value from the source dataset *periodic snapshot task* (page 97).
- *Never Delete*: never delete snapshots from the destination system.
- *Custom*: define how long a snapshot remains on the destination system. Enter a number and choose a measure of time from the drop-down menus.

Clicking *START REPLICATION* saves the replication configuration and activates the schedule. When the replication configuration includes a source dataset on the local system and has a schedule, a *periodic snapshot task* (page 97) of that dataset is also created.

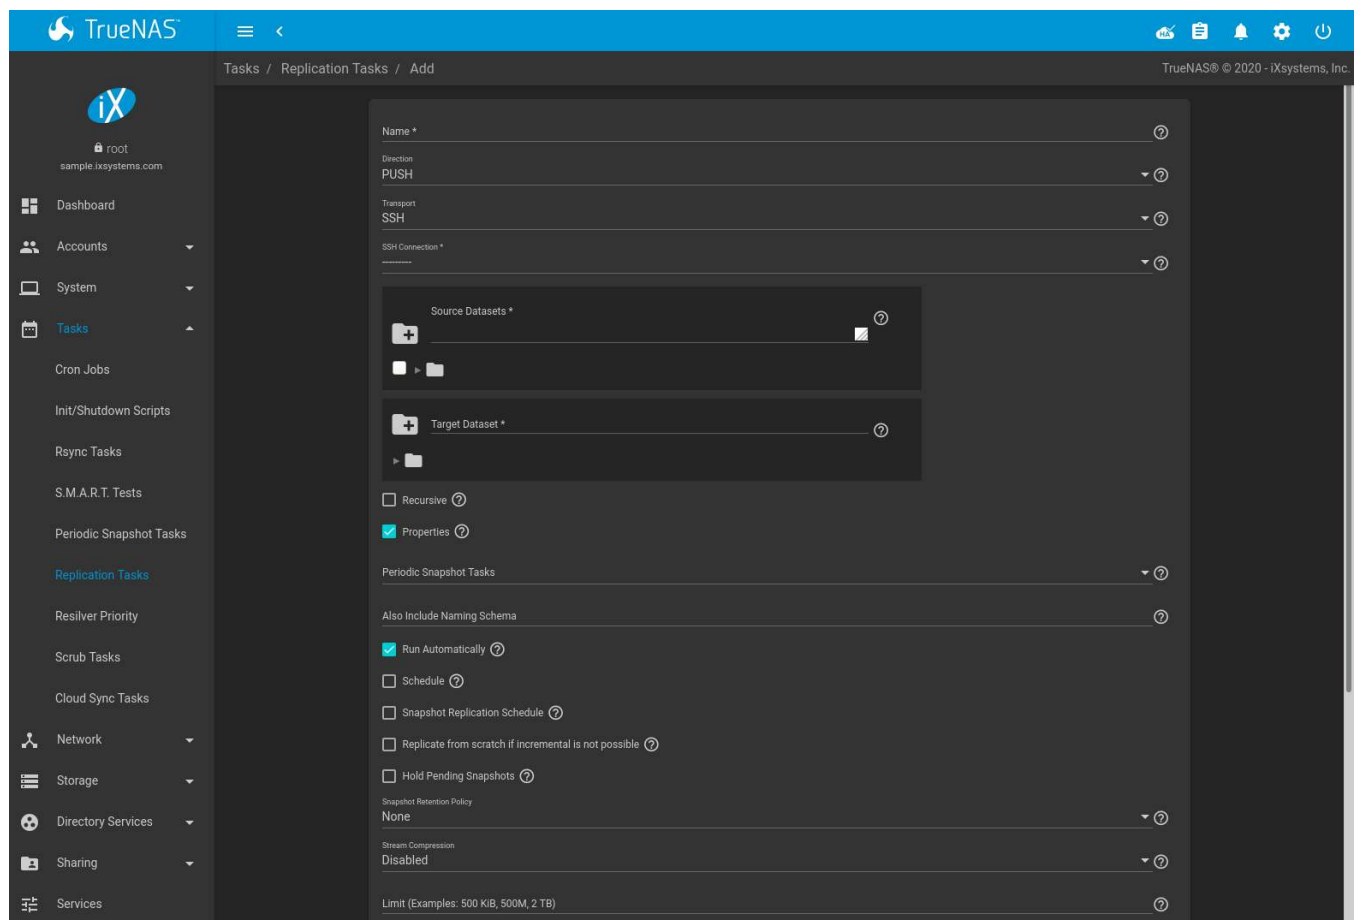
Tasks set to *Run Once* will start immediately. If a one-time replication has no valid local system source dataset snapshots, TrueNAS® will snapshot the source datasets and immediately replicate those snapshots to the destination dataset.

All replication tasks are displayed in *Tasks → Replication Tasks*. The task settings that are shown by default can be adjusted by opening the *COLUMNS* drop-down. To see more details about the last time the replication task ran, click the entry under the *State* column. Tasks can also be expanded by clicking *>* (Expand) for that task. Expanded tasks show all replication settings and have *RUN NOW*, *EDIT*, and *DELETE* buttons.

6.6.2 Advanced Replication Creation

The advanced replication creation screen has more options for fine-tuning a replication. It also allows creating local replications, legacy engine replications from TrueNAS® 11.1 or earlier, or even creating a one-time replication that is not linked to a periodic snapshot task.

Go to *System → Replication Tasks*, click *ADD* and *ADVANCED REPLICATION CREATION* to see these options. This screen is also displayed after clicking *:* (Options) and *Edit* for an existing replication.



The *Transport* value changes many of the options for replication. Table 6.6 shows abbreviated names of the *Transport* methods in the *Transport* column to identify fields which appear when that method is selected.

- ALL: All *Transport* methods
- SSH: *SSH*
- NCT: *SSH+NETCAT*
- LOC: *LOCAL*

- LEG: *LEGACY*

Table 6.6: Replication Task Options

Setting	Transport	Value	Description
Name	All	string	Descriptive name for the replication.
Direction	SSH, NCT, LEG	drop-down menu	<i>PUSH</i> sends snapshots to a destination system. <i>PULL</i> connects to a remote system and retrieves snapshots matching a <i>Naming Schema</i> .
Transport	All	drop-down menu	Method of snapshot transfer: <ul style="list-style-type: none"> • <i>SSH</i> is supported by most systems. It requires a previously created <i>SSH connection</i> (page 57). • <i>SSH+NETCAT</i> uses SSH to establish a connection to the destination system, then uses <i>py-libzfs</i> (https://github.com/freenas/py-libzfs) to send an unencrypted data stream for higher transfer speeds. By default, this is supported by TrueNAS® systems with 11.2 or later installed (11.3 or later is recommended). Destination systems that do not have TrueNAS® 11.2 or later installed might have to manually install <i>py-libzfs</i>. • <i>LOCAL</i> efficiently replicates snapshots to another dataset on the same system. • <i>LEGACY</i> uses the legacy replication engine from TrueNAS® 11.2 and earlier.
SSH Connection	SSH, NCT, LEG	drop-down menu	Choose the <i>SSH connection</i> (page 57).
Netcat Active Side	NCT	drop-down menu	Establishing a connection requires that one of the connection systems has open TCP ports. Choose which system (<i>LOCAL</i> or <i>REMOTE</i>) will open ports. Consult your IT department to determine which systems are allowed to open ports.
Netcat Active Side Listen Address	NCT	string	IP address on which the connection <i>Active Side</i> listens. Defaults to 0.0.0.0.
Netcat Active Side Min Port	NCT	integer	Lowest port number of the active side listen address that is open to connections.
Netcat Active Side Max Port	NCT	integer	Highest port number of the active side listen address that is open to connections. The first available port between the minimum and maximum is used.
Netcat Active Side Connect Address	NCT	string	Hostname or IP address used to connect to the active side system. When the active side is <i>LOCAL</i> , this defaults to the <i>SSH_CLIENT</i> environment variable. When the active side is <i>REMOTE</i> , this defaults to the SSH connection hostname.
Source	All	<input type="button" value="Browse"/> , string	Define the path to a system location that has snapshots to replicate. Click the <input type="button" value="Browse"/> (Browse) to see all locations on the source system or click in the field to manually type a location (Example: pool1/dataset1). Multiple source locations can be selected or manually defined with a comma (literal;) separator.

Continued on next page

Table 6.6 – continued from previous page

Setting	Transport	Value	Description
Destination	All	<input type="button" value="Browse"/> (Browse), string	<p>Define the path to a system location that will store replicated snapshots. Click the <input type="button" value="Browse"/> (Browse) to see all locations on the destination system or click in the field to manually type a location path (Example: <code>pool1/dataset1</code>). Selecting a location defines the full path to that location as the destination. Appending a name to the path will create new zvol at that location.</p> <p>For example, selecting <code>pool1/dataset1</code> will store snapshots in <code>dataset1</code>, but clicking the path and typing <code>/zvol1</code> after <code>dataset1</code> will create <code>zvol1</code> for snapshot storage.</p>
Recursive	All	checkbox	Replicate all child dataset snapshots. When set, <i>Exclude Child Datasets</i> becomes visible.
Exclude Child Datasets	SSH, NCT, LOC	string	Exclude specific child dataset snapshots from the replication. Use with <i>Recursive</i> replications. List child dataset names to exclude. Separate multiple entries with a comma (,). Example: <code>pool1/dataset1/child1</code> . A recursive replication of <code>pool1/dataset1</code> snapshots includes all child dataset snapshots except <code>child1</code> .
Properties	SSH, NCT, LOC	checkbox	Include dataset properties with the replicated snapshots.
Periodic Snapshot Tasks	SSH, NCT, LOC	drop-down menu	Snapshot schedule for this replication task. Choose from configured <i>Periodic Snapshot Tasks</i> (page 97). This replication task must have the same <i>Recursive</i> and <i>Exclude Child Datasets</i> values as the chosen periodic snapshot task. Selecting a periodic snapshot schedule removes the <i>Schedule</i> field.
Naming Schema	SSH, NCT, LOC	string	Visible with <i>PULL</i> replications. Pattern of naming custom snapshots to be replicated. Enter the name and <code>strftime(3)</code> (https://www.freebsd.org/cgi/man.cgi?query=strftime) <code>%Y</code> , <code>%m</code> , <code>%d</code> , <code>%H</code> , and <code>%M</code> strings that match the snapshots to include in the replication.
Also Include Naming Schema	SSH, NCT, LOC	string	<p>Visible with <i>PUSH</i> replications. Pattern of naming custom snapshots to include in the replication with the periodic snapshot schedule. Enter the <code>strftime(3)</code> (https://www.freebsd.org/cgi/man.cgi?query=strftime) strings that match the snapshots to include in the replication.</p> <p>When a periodic snapshot is not linked to the replication, enter the naming schema for manually created snapshots. Has the same <code>%Y</code>, <code>%m</code>, <code>%d</code>, <code>%H</code>, and <code>%M</code> string requirements as the <i>Naming Schema</i> in a <i>periodic snapshot task</i> (page 98).</p>
Run Automatically	SSH, NCT, LOC	checkbox	Set to either start this replication task immediately after the linked periodic snapshot task completes or continue to create a separate <i>Schedule</i> for this replication.
Schedule	SSH, NCT, LOC	checkbox and drop-down menu	Start time for the replication task. Select a preset schedule or choose <i>Custom</i> to use the advanced scheduler. Adds the <i>Begin</i> and <i>End</i> fields.
Begin	SSH, NCT, LOC	drop-down menu	Start time for the replication task.
End	SSH, NCT, LOC	drop-down menu	End time for the replication task. A replication that is already in progress can continue to run past this time.

Continued on next page

Table 6.6 – continued from previous page

Setting	Transport	Value	Description
Replicate Specific Snapshots	SSH, NCT, LOC	checkbox and drop-down menu	Only replicate snapshots that match a defined creation time. To specify which snapshots will be replicated, set this checkbox and define the snapshot creation times that will be replicated. For example, setting this time frame to <i>Hourly</i> will only replicate snapshots that were created at the beginning of each hour.
Begin	SSH, NCT, LOC	drop-down menu	Daily time range for the specific periodic snapshots to replicate, in 15 minute increments. Periodic snapshots created before the <i>Begin</i> time will not be included in the replication.
End	SSH, NCT, LOC	drop-down menu	Daily time range for the specific periodic snapshots to replicate, in 15 minute increments. Snapshots created after the <i>End</i> time will not be included in the replication.
Only Replicate Snapshots Matching Schedule	SSH, NCT, LOC	checkbox	Set to use the <i>Schedule</i> in place of the <i>Replicate Specific Snapshots</i> time frame. The <i>Schedule</i> values are read over the <i>Replicate Specific Snapshots</i> time frame.
Replicate from scratch if incremental is not possible	SSH, NCT, LOC	checkbox	If the destination system has snapshots but they do not have any data in common with the source snapshots, destroy all destination snapshots and do a full replication. Warning: enabling this option can cause data loss or excessive data transfer if the replication is misconfigured.
Hold Pending Snapshots	SSH, NCT, LOC	checkbox	Prevent source system snapshots that have failed replication from being automatically removed by the <i>Snapshot Retention Policy</i> .
Snapshot Retention Policy	SSH, NCT, LOC	drop-down menu	When replicated snapshots are deleted from the destination system: <ul style="list-style-type: none"> • <i>Same as Source</i>: use <i>Snapshot Lifetime</i> value from the source <i>periodic snapshot task</i> (page 97). • <i>Custom</i>: define a <i>Snapshot Lifetime</i> for the destination system. • <i>None</i>: never delete snapshots from the destination system.
Snapshot Lifetime	All	integer and drop-down menu	Added with a <i>Custom</i> retention policy. How long a snapshot remains on the destination system. Enter a number and choose a measure of time from the drop-down.
Stream Compression	SSH	drop-down menu	Select a compression algorithm to reduce the size of the data being replicated. Only appears when <i>SSH</i> is chosen for <i>Transport</i> .
Limit (Examples: 500 KiB, 500M, 2 TB)	SSH	integer	Limit replication speed to this number of bytes per second. Zero means no limit. This is a <i>humanized field</i> (page 14).
Send Deduplicated Stream	SSH, NCT, LOC	checkbox	Deduplicate the stream to avoid sending redundant data blocks. The destination system must also support deduplicated streams. See zfs(8) (https://www.freebsd.org/cgi/man.cgi?query=zfs).
Allow Blocks Larger than 128KB	SSH, NCT, LOC	checkbox	Allow sending large data blocks. The destination system must also support large blocks. See zfs(8) (https://www.freebsd.org/cgi/man.cgi?query=zfs).
Allow Compressed WRITE Records	SSH, NCT, LOC	checkbox	Use compressed WRITE records to make the stream more efficient. The destination system must also support compressed WRITE records. See zfs(8) (https://www.freebsd.org/cgi/man.cgi?query=zfs).

Continued on next page

Table 6.6 – continued from previous page

Setting	Transport	Value	Description
Number of retries for failed replications	SSH, NCT, LOC	integer	Number of times the replication is attempted before stopping and marking the task as failed.
Logging Level	All	drop-down menu	Message verbosity level in the replication task log.
Enabled	All	checkbox	Activates the replication schedule.

6.6.3 Replication Tasks

Saved replications are shown on the *Replication Tasks* page.

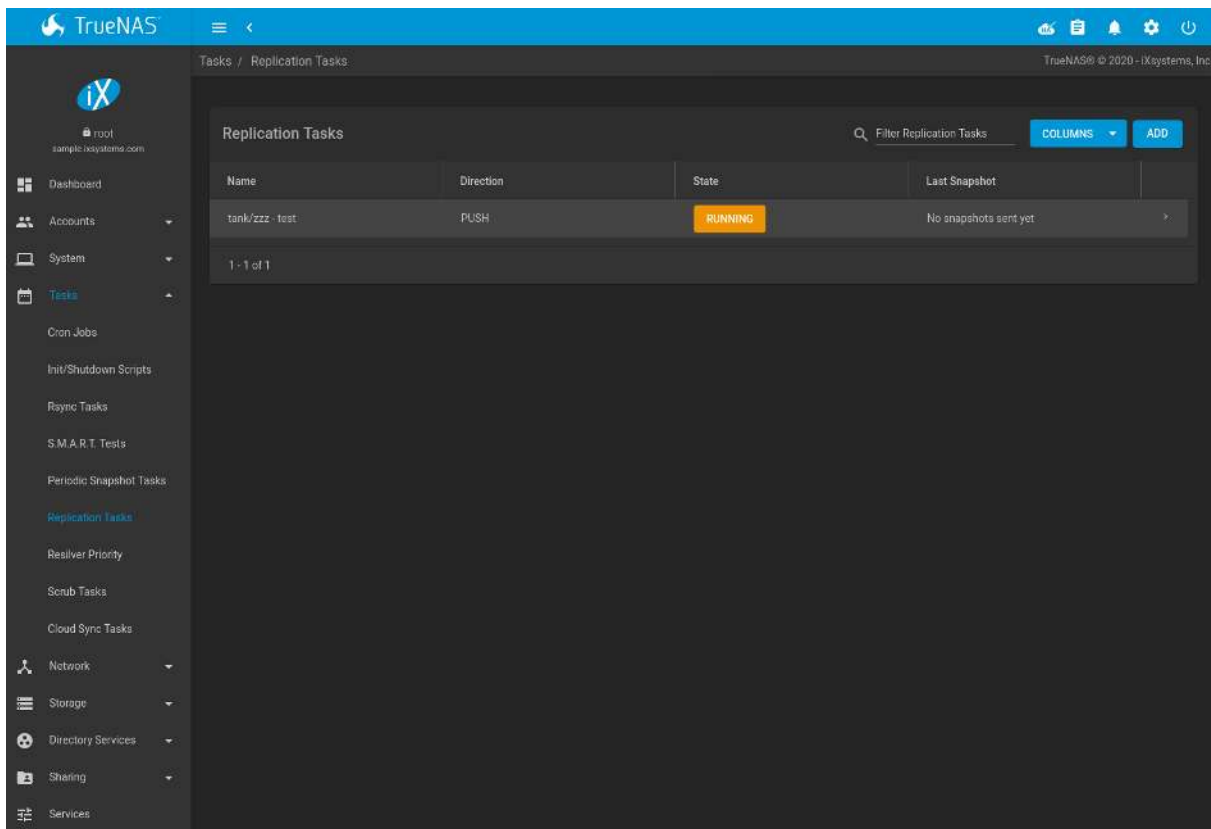


Fig. 6.9: Replication Task List

The replication name and configuration details are shown in the list. To adjust the default table view, open the **COLUMNS** menu and select the replication details to show in the normal table view.

The *State* column shows the status of the replication task. To view the detailed replication logs for a task, click the *State* entry when the task is running or finished.

Expanding an entry shows additional buttons for starting or editing a replication task.

6.6.4 Limiting Replication Times

The *Schedule*, *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 ac-

tivity will not slow down other operations like snapshots or *Scrub Tasks* (page 111). 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.

6.6.5 Replication Topologies and Scenarios

The replication examples shown above are known as *simple* or *A to B* replication, where one machine replicates data to one other machine. Replication can also be set up in more sophisticated topologies to suit various purposes and needs.

6.6.5.1 Star Replication

In a *star* topology, a single TrueNAS® computer replicates data to multiple destination computers. This provides data redundancy with the multiple copies of data, and geographical redundancy if the destination computers are located at different sites.

An *Alpha* computer with three separate replication tasks to replicate data to *Beta*, then *Gamma*, and finally *Delta* computers demonstrates this arrangement. *A to B* replication is really just a star arrangement with only one target computer.

The star topology is simple to configure and manage, but it can place relatively high I/O and network loads on the source computer, which must run an individual replication task for each target computer.

6.6.5.2 Tiered Replication

In *tiered* replication, the data is replicated from the source computer onto one or a few destination computers. The destination computers then replicate the same data onto other computers. This allows much of the network and I/O load to be shifted away from the source computer.

For example, consider both *Alpha* and *Beta* computers to be located inside the same data center. Replicating data from *Alpha* to *Beta* does not protect that data from events that would involve the whole data center, like flood, fire, or earthquake. Two more computers, called *Gamma* and *Delta*, are set up. To provide geographic redundancy, *Gamma* is in a data center on the other side of the country, and *Delta* is in a data center on another continent. A single periodic snapshot replicates data from *Alpha* to *Beta*. *Beta* then replicates the data onto *Gamma*, and again onto *Delta*.

Tiered replication shifts most of the network and I/O overhead of repeated replication off the source computer onto the target computers. The source computer only replicates to the second-tier computers, which then handle replication to the third tier, and so on. In this example, *Alpha* only replicates data onto *Beta*. The I/O and network load of repeated replications is shifted onto *Beta*.

6.6.5.3 N-way Replication

N-way replication topologies recognize that hardware is sometimes idle, and computers can be used for more than a single dedicated purpose. An individual computer can be used as both a source and destination for replication. For example, the *Alpha* system can replicate a dataset to *Beta*, while *Beta* can replicate datasets to both *Alpha* and *Gamma*.

With careful setup, this topology can efficiently use I/O, network bandwidth, and computers, but can quickly become complex to manage.

6.6.5.4 Disaster Recovery

Disaster recovery is the ability to recover complete datasets from a replication destination computer. The replicated dataset is replicated back to new hardware after an incident caused the source computer to fail.

Recovering data onto a replacement computer is done manually with the `zfs send` and `zfs recv` commands, or a replication task can be defined on the target computer containing the backup data. This replication task would normally be disabled. If a disaster damages the source computer, the target computer replication task is temporarily enabled, replicating the data onto the replacement source computer. After the disaster recovery replication completes, the replication task on the target computer is disabled again.

6.6.6 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.

Replication logs are saved in `var/log/zettarepl.log`. Logs of individual replication tasks can be viewed by clicking the replication *State*.

6.6.6.1 SSH

SSH (page 244) must be able to connect from the source system to the destination system with an encryption key. This is tested from *Shell* (page 302) by making an *SSH* (page 244) connection from the source system to the destination system. For example, this is a connection from *Alpha* to *Beta* at `10.0.0.118`. Start the *Shell* (page 302) on the source machine (*Alpha*), then enter this command:

```
ssh -vv 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 244) shell connection is open to the destination system, *Beta*.

If a password is requested, SSH authentication is not working. An SSH key value must be present in the destination system `/root/.ssh/authorized_keys` file. `/var/log/auth.log` file can show diagnostic errors for login problems on the destination computer also.

6.6.6.2 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 TrueNAS®, 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 TrueNAS® system for errors.

6.6.6.3 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 302) 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 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 302) 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 302) with `zfs list -t snapshot` or from *Storage* → *Snapshots*.

Error messages here can indicate any remaining problems.

6.7 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 6.10](#). [Table 6.7](#) describes the fields on this screen.

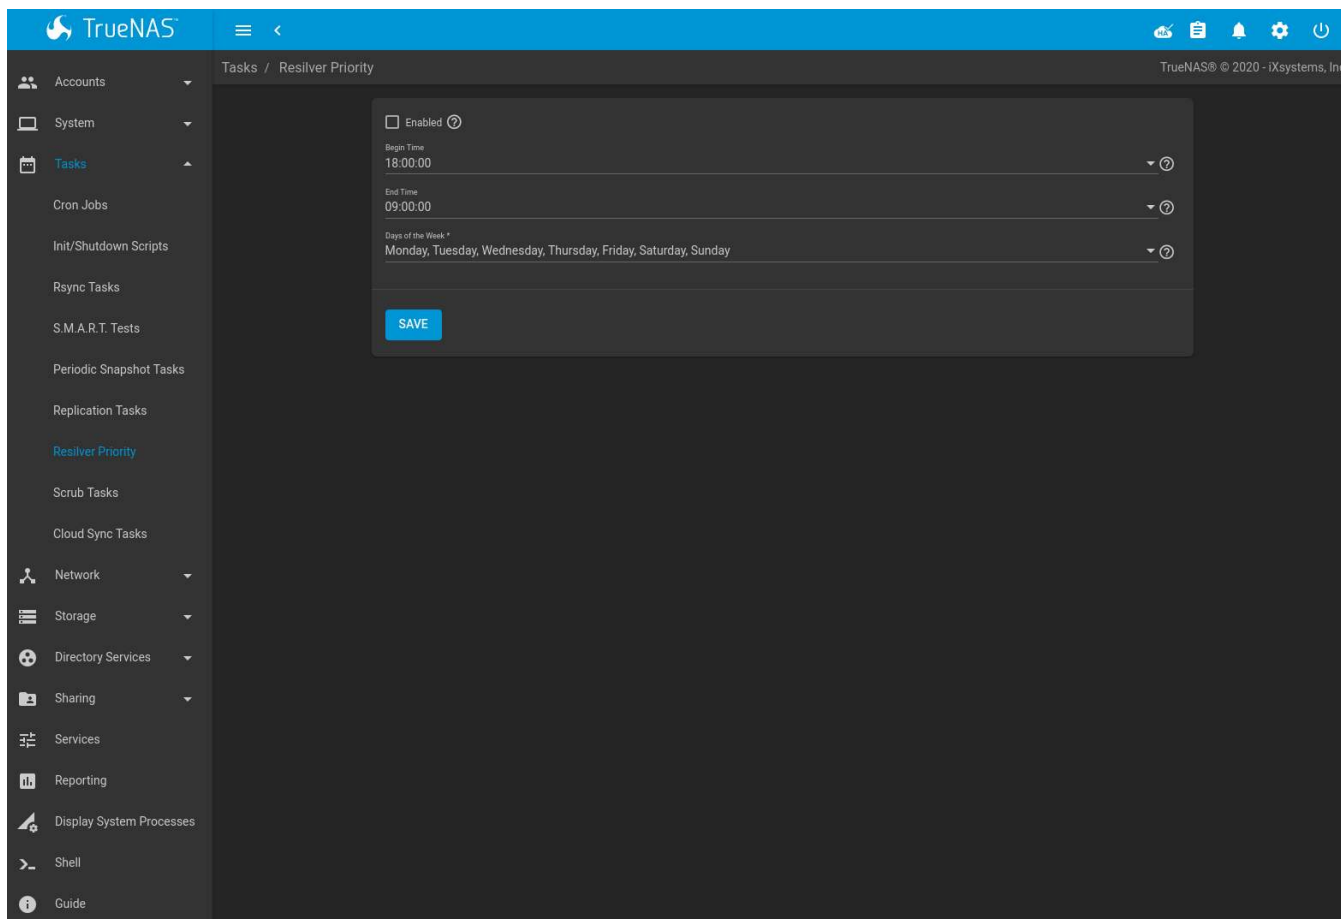


Fig. 6.10: Resilver Priority

Table 6.7: 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.

Continued on next page

Table 6.7 – continued from previous page

Setting	Value	Description
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.

6.8 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. TrueNAS® 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 minimize impact to users. Make certain that scrubs and other disk-intensive activity like [S.M.A.R.T. Tests](#) (page 96) 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 96) 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 6.11](#) displays the default settings for the pool named `pool11`. In this example,  (Options) and *Edit* for a pool is clicked to display the *Edit* screen. [Table 6.8](#) summarizes the options in this screen.

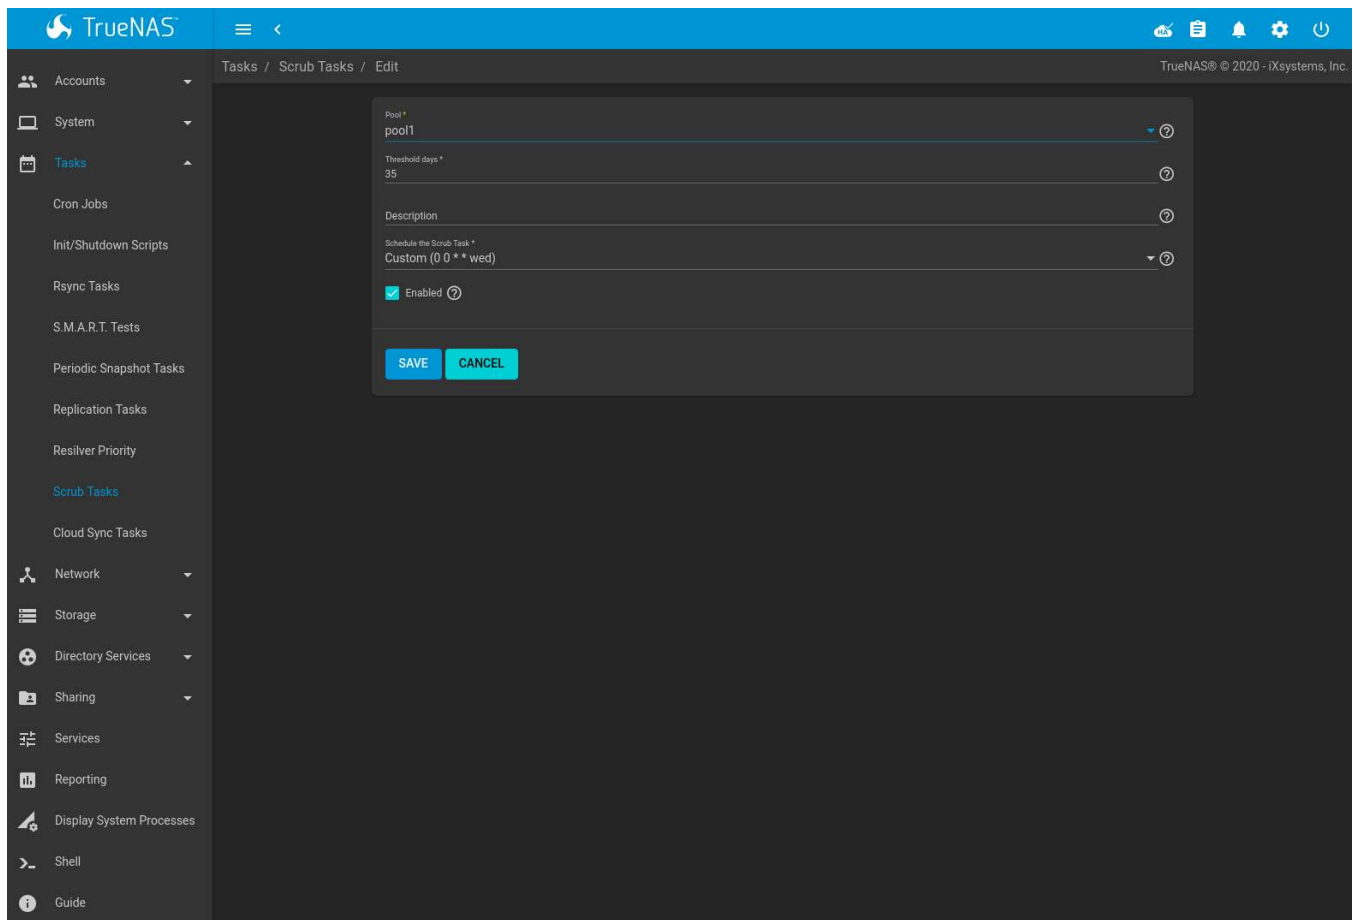


Fig. 6.11: Viewing Pool Default Scrub Settings

Table 6.8: ZFS Scrub Options

Setting	Value	Description
Pool	drop-down menu	Choose a pool to scrub.
Threshold days	string	Days before a completed scrub is allowed to run again. This controls the task schedule. For example, scheduling a scrub to run daily and setting <i>Threshold days</i> to 7 means the scrub attempts to run daily. When the scrub is successful, it continues to check daily but does not run again until seven days have elapsed. Using a multiple of seven ensures the scrub always occurs on the same weekday.
Description	string	Describe the scrub task.
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> . Selecting <i>Custom</i> opens the Advanced Scheduler (page 12).
Enabled	checkbox	Unset to disable the scheduled scrub without deleting it.

Review the default selections and, if necessary, modify them to meet the needs of the environment. Scrub tasks cannot run for locked or unmounted pools.

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.

6.9 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 53) must be defined before a cloud sync is created. One set of credentials can be used for more than one cloud sync. For example, a single set of credentials for Amazon S3 can be used for separate cloud syncs that push different sets of files or directories.

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

After the cloud credentials have been configured, *Tasks* → *Cloud Sync Tasks* is used to define the schedule for running a cloud sync task. The time selected is when the Cloud Sync task is allowed to begin. An in-progress cloud sync must complete before another cloud sync can start. The cloud sync runs until finished, even after the selected ending time. To stop the cloud sync task before it is finished, click ⋮ (Options) → *Stop*.

An example is shown in [Figure 6.12](#).

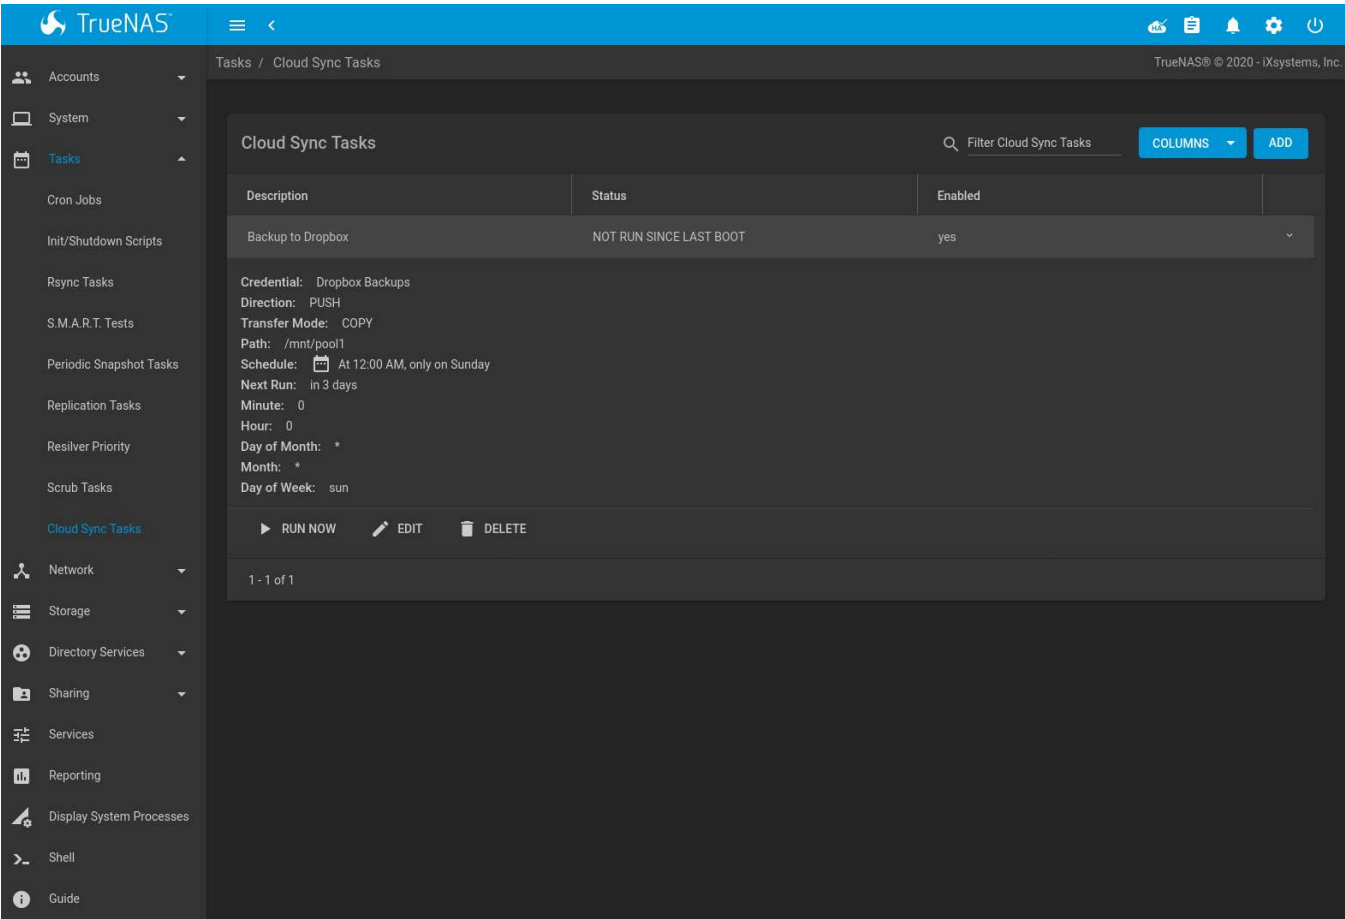


Fig. 6.12: Cloud Sync Status

The cloud sync *Status* indicates the state of most recent cloud sync. Clicking the *Status* entry shows the task logs and includes an option to download them.

Click **ADD** to display the *Add Cloud Sync* menu shown in Figure 6.13.

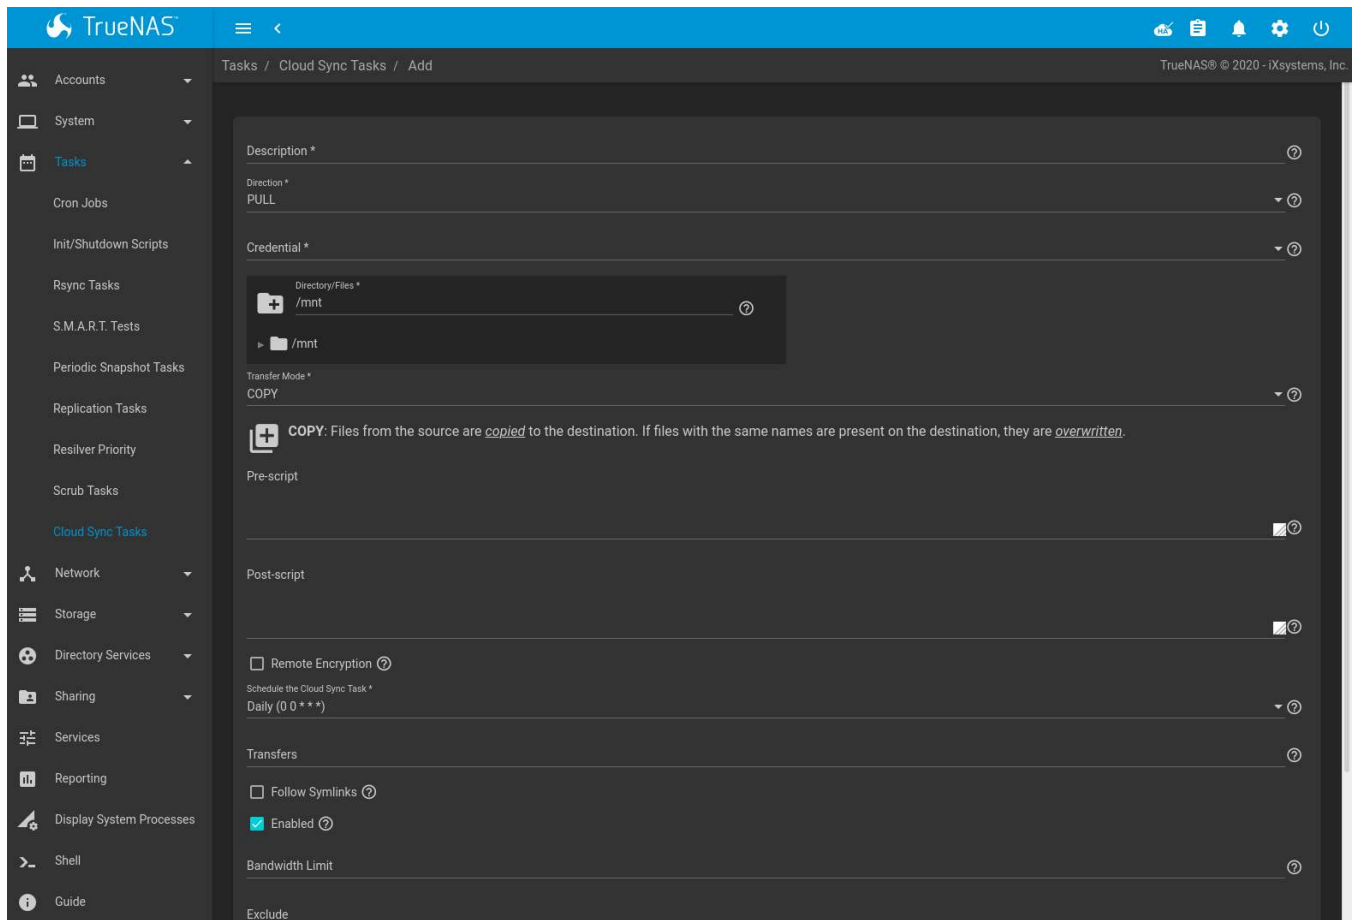
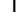


Fig. 6.13: Adding a Cloud Sync

Table 6.9 shows the configuration options for Cloud Syncs.

Table 6.9: Cloud Sync Options

Setting	Value Type	Description
Description	string	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. Changing the direction resets the <i>Transfer Mode</i> to <i>COPY</i> .
Credential	drop-down menu	Select the cloud storage provider credentials from the list of available Cloud Credentials (page 53). The credential is tested and an error is displayed if a connection cannot be made. Click <i>Fix Credential</i> to go to the configuration page for that Cloud Credential (page 53). <i>SAVE</i> is disabled until a valid credential is selected.
Bucket/Container	drop-down menu	<i>Bucket</i> : Only appears when an S3 credential is the <i>Provider</i> . Select the predefined S3 bucket to use. <i>Container</i> : The pre-configured container name. Only appears when a <i>AZUREBLOB</i> or <i>hubsC</i> credential is selected as the <i>Credential</i> .
Folder	browse button	The name of the predefined folder within the selected bucket or container. Type the name or click  (Browse) to list the remote filesystem and choose the folder.
Server Side Encryption	drop-down menu	Active encryption on the cloud provider account. Choose <i>None</i> or <i>AES-256</i> . Only visible when the cloud provider supports encryption.

Continued on next page

Table 6.9 – continued from previous page

Setting	Value Type	Description
Storage Class	drop-down menu	Classification for each S3 object. Choose a class based on the specific use case or performance requirements. See Amazon S3 Storage Classes (https://docs.aws.amazon.com/AmazonS3/latest/dev/storage-class-intro.html) for more information on which storage class to choose. <i>Storage Class</i> only appears when an S3 credential is the <i>Provider</i> .
Upload Chunk Size (MiB)	integer	Files are split into chunks of this size before upload. The number of chunks that can be simultaneously transferred is set by the <i>Transfers</i> number. The single largest file being transferred must fit into no more than 10,000 chunks.
Use -fast-list	checkbox	Use fewer transactions in exchange for more RAM (https://rclone.org/docs/#fast-list). Modifying this setting can speed up or slow down the transfer. Only appears with a compatible <i>Credential</i> .
Directory/Files	browse button	Select 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.
Transfer Mode	drop-down menu	SYNC : Files on the destination are changed to match those on the source. If a file does not exist on the source, it is also deleted from the destination. There are exceptions (page 116) to this behavior. COPY : Files from the source are copied to the destination. If files with the same names are present on the destination, they are overwritten . MOVE : After files are copied from the source to the destination, they are deleted from the source. Files with the same names on the destination are overwritten .
Take Snapshot	checkbox	Take a snapshot of the dataset before a <i>PUSH</i> . This cannot be enabled when the chosen dataset to <i>PUSH</i> has nested datasets.
Pre-script	string	A script to execute before the Cloud Sync Task is run.
Post-script	string	A script to execute after the Cloud Sync Task is run.
Remote Encryption	checkbox	Use rclone crypt (https://rclone.org/crypt/) to manage data encryption during <i>PUSH</i> or <i>PULL</i> transfers: <i>PUSH</i> : Encrypt files before transfer and store the encrypted files on the remote system. Files are encrypted using the <i>Encryption Password</i> and <i>Encryption Salt</i> values. <i>PULL</i> : Decrypt files that are being stored on the remote system before the transfer. Transferring the encrypted files requires entering the same <i>Encryption Password</i> and <i>Encryption Salt</i> that was used to encrypt the files. Adds the <i>Filename Encryption</i> , <i>Encryption Password</i> , and <i>Encryption Salt</i> options. Additional details about the encryption algorithm and key derivation are available in the rclone crypt File formats documentation (https://rclone.org/crypt/#file-formats).
Filename Encryption	checkbox	Encrypt (<i>PUSH</i>) or decrypt (<i>PULL</i>) file names with the rclone “ Standard ” file name encryption mode (https://rclone.org/crypt/#file-name-encryption-modes). The original directory structure is preserved. A filename with the same name always has the same encrypted filename. <i>PULL</i> tasks that have <i>Filename Encryption</i> enabled and an incorrect <i>Encryption Password</i> or <i>Encryption Salt</i> will not transfer any files but still report that the task was successful. To verify that files were transferred successfully, click the finished task status (page 113) to see a list of transferred files.


Continued on next page

Table 6.9 – continued from previous page

Setting	Value Type	Description
Encryption Password	string	Password to encrypt and decrypt remote data. Warning: Always securely back up this password! Losing the encryption password will result in data loss.
Encryption Salt	string	Enter a long string of random characters for use as salt (https://searchsecurity.techtarget.com/definition/salt) for the encryption password. Warning: Always securely back up the encryption salt value! Losing the salt value will result in data loss.
Schedule the Cloud Sync Task	drop-down menu	Choose how often or at what time to start a sync. Choices are <i>Hourly</i> , <i>Daily</i> , <i>Weekly</i> , <i>Monthly</i> , or <i>Custom</i> . Selecting <i>Custom</i> opens the Advanced Scheduler (page 12).
Transfers	integer	Number of simultaneous file transfers. Enter a number based on the available bandwidth and destination system performance. See rclone -transfers (https://rclone.org/docs/#transfers-n).
Follow Sym-links	checkbox	Include symbolic link targets in the transfer.
Enabled	checkbox	Enable this Cloud Sync Task. Unset to disable this Cloud Sync Task without deleting it.
Bandwidth Limit	string	A single bandwidth limit or bandwidth limit schedule in rclone format. Example: <i>08:00,512 12:00,10MB 13:00,512 18:00,30MB 23:00,off</i> . Units can be specified with the beginning letter: b, k (default), M, or G. See rclone -bwlimit . (https://rclone.org/docs/#bwlimit-bandwidth-spec)
Exclude	string	List of files and directories to exclude from sync, one per line. See https://rclone.org/filtering/ .

There are specific circumstances where a *SYNC* task does not delete files from the destination:

- If [rclone sync](https://rclone.org/commands/rclone_sync/) (https://rclone.org/commands/rclone_sync/) encounters any errors, files are not deleted in the destination. This includes a common error when the Dropbox [copyright detector](https://techcrunch.com/2014/03/30/how-dropbox-knows-when-youre-sharing-copyrighted-stuff-without-actually-looking-at-your-stuff/) (https://techcrunch.com/2014/03/30/how-dropbox-knows-when-youre-sharing-copyrighted-stuff-without-actually-looking-at-your-stuff/) flags a file as copyrighted.
- Syncing to a [B2 bucket](#) (page 55) does not delete files from the bucket, even when those files have been deleted locally. Instead, files are tagged with a version number or moved to a hidden state. To automatically delete old or unwanted files from the bucket, adjust the [Backblaze B2 Lifecycle Rules](https://www.backblaze.com/blog/backblaze-b2-lifecycle-rules/) (https://www.backblaze.com/blog/backblaze-b2-lifecycle-rules/)
- Files stored in Amazon S3 Glacier or S3 Glacier Deep Archive cannot be deleted by [rclone sync](https://rclone.org/s3/#glacier-and-glacier-deep-archive/) (https://rclone.org/s3/#glacier-and-glacier-deep-archive/). These files must first be restored by another means, like the [Amazon S3 console](https://docs.aws.amazon.com/AWSAmazonS3/latest/user-guide/restore-archived-objects.html) (https://docs.aws.amazon.com/AWSAmazonS3/latest/user-guide/restore-archived-objects.html).

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

6.9.1 Cloud Sync Example

This example shows a *Push* cloud sync that copies files from a TrueNAS® pool to a cloud service provider.

The cloud service provider was configured with a location to store data received from the TrueNAS® system.

In the TrueNAS® web interface, go to *System* → *Cloud Credentials* and click *ADD* to configure the cloud service provider credentials:




Fig. 6.14: Example: Adding Cloud Credentials

Go to *Tasks* → *Cloud Sync* and click *ADD* to create a cloud sync job. The *Description* is filled with a simple note describing the job. 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 folder was configured in the cloud provider account.

The *Directory/Files* is set to the file or directory to copy to the cloud provider.

The *Transfer Mode* is set to *COPY* so that only the files stored by the cloud provider are modified.

The remaining requirement is to schedule the task. The default is to send the data to cloud storage daily, but the schedule can be *customized* (page 12) to fine-tune when the task runs.

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

An example of a completed cloud sync task is shown in [Figure 6.15](#):

TrueNAS

Tasks / Cloud Sync Tasks

TrueNAS © 2020 - iXsystems, Inc.

Accounts

System

Tasks

Cron Jobs

Init/Shutdown Scripts

Rsync Tasks

S.M.A.R.T. Tests

Periodic Snapshot Tasks

Replication Tasks

Resilver Priority

Scrub Tasks

Cloud Sync Tasks

Network

Storage

Directory Services

Sharing

Services

Reporting

Display System Processes

Shell

Guide

Cloud Sync Tasks

Filter Cloud Sync Tasks

COLUMNS

ADD

Description	Status	Enabled
Backup to Dropbox	SUCCESS	yes

Credential: Dropbox Backups

Direction: PUSH

Transfer Mode: COPY

Path: /mnt/pool1

Schedule: At 12:00 AM, only c

Next Run: in 3 days

Minute: 0

Hour: 0

Day of Month: *

Month: *

Day of Week: sun

RUN NOW

EDIT

1 - 1 of 1

Logs

2020/01/08 09:58:43 INFO : Dropbox root 'test1': Waiting for checks to finish

2020/01/08 09:58:43 INFO : Dropbox root 'test1': Waiting for transfers to finish

2020/01/08 09:58:43 INFO :

Transferred: 0 / 0 Bytes, -, 0 Bytes/s, ETA -

Errors: 0

Checks: 0 / 0, -

Transferred: 0 / 0, -

Elapsed time: 0s

CANCEL

DOWNLOAD LOGS

Fig. 6.15: Example: Successful Cloud Sync

NETWORK

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

- *Global Configuration* (page 119): general network settings.
- *Interfaces* (page 121): settings for each network interface and options to configure *Bridge* (page 124), *Link Aggregation* (page 124), and *VLAN* (page 126) interfaces.
- *IPMI* (page 126): settings controlling connection to the appliance through the hardware side-band management interface if the user interface becomes unavailable.
- *Static Routes* (page 128): add static routes.

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

Note: When any network changes are made an animated icon appears in the upper-right web interface panel to show there are pending network changes. When the icon is clicked it prompts to review the recent network changes. Reviewing the network changes goes to *Network → Interfaces* where the changes can be permanently applied or discarded.

When *APPLY CHANGES* is clicked the network changes are temporarily applied for 60 seconds by default. This value can be changed by entering a positive integer in the seconds field. This feature is nice because the network settings preview can automatically roll back any configuration errors that are accidentally saved.

If the network settings applied work as intended, click *KEEP CHANGES*. Otherwise, the changes can be discarded by clicking *DISCARD CHANGES*.

7.1 Global Configuration

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

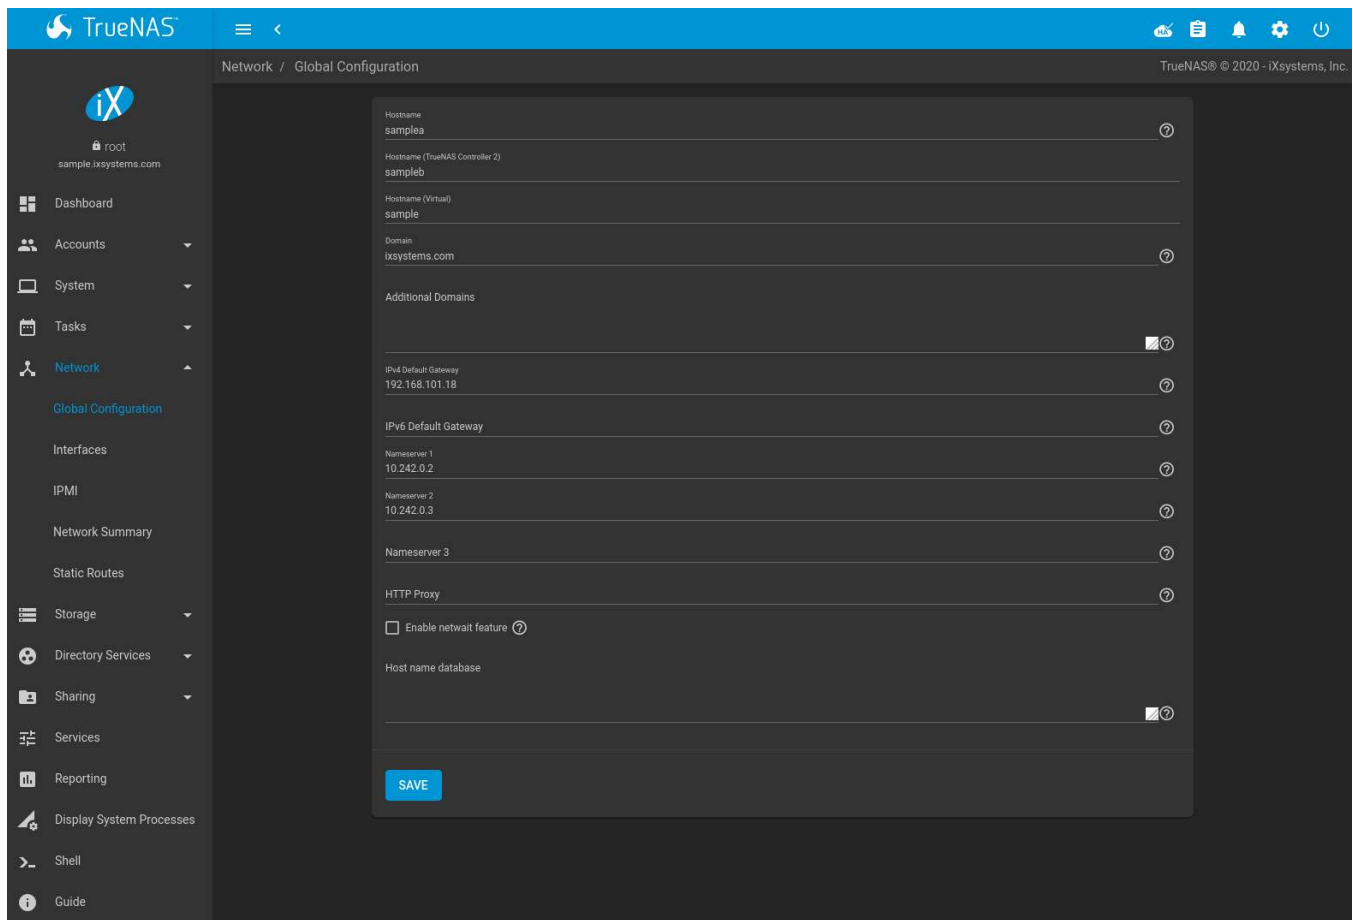


Fig. 7.1: Global Network Configuration

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

Table 7.1: Global Configuration Settings

Setting	Value	Description
Hostname	string	Host name of first TrueNAS controller. Upper and lower case alphanumeric, ., and – characters are allowed.
Hostname (TrueNAS Controller 2)	string	Host name of second TrueNAS controller. Upper and lower case alphanumeric, ., and – characters are allowed.
Hostname (Virtual)	string	Virtual host name. When using a virtualhost, this is also used as the Kerberos principal name. Enter the fully qualified hostname plus the domain name. Upper and lower case alphanumeric, ., and – characters are allowed.
Domain	string	System domain name. The <i>Hostname</i> and <i>Domain</i> are also displayed under the iXsystems logo at the top left of the main screen.
Additional Domains	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 121). 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 121).
Nameserver 1	IP address	Primary DNS server.

Continued on next page

Table 7.1 – continued from previous page

Setting	Value	Description
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 <i>http://my.proxy.server:3128</i> or <i>http://user:password@my.proxy.server:3128</i> .
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> .
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 <i>/etc/hosts</i> . 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 TrueNAS® system in the *Host name database* field.

Note: In many cases, a TrueNAS® 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 TrueNAS® 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 128) 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 TrueNAS® system is protected by a properly configured firewall.

7.2 Interfaces

Network → Interfaces shows all physical Network Interface Controllers (NICs) connected to the TrueNAS® system. These can be edited or new *bridge*, *link aggregation*, or *Virtual LAN (VLAN)* interfaces can be created and added to the interface list.

Be careful when configuring the network interface that controls the TrueNAS® web interface or *web connectivity can be lost* (page 119).

To configure a new network interface, go to *Network → Interfaces* and click *ADD*.

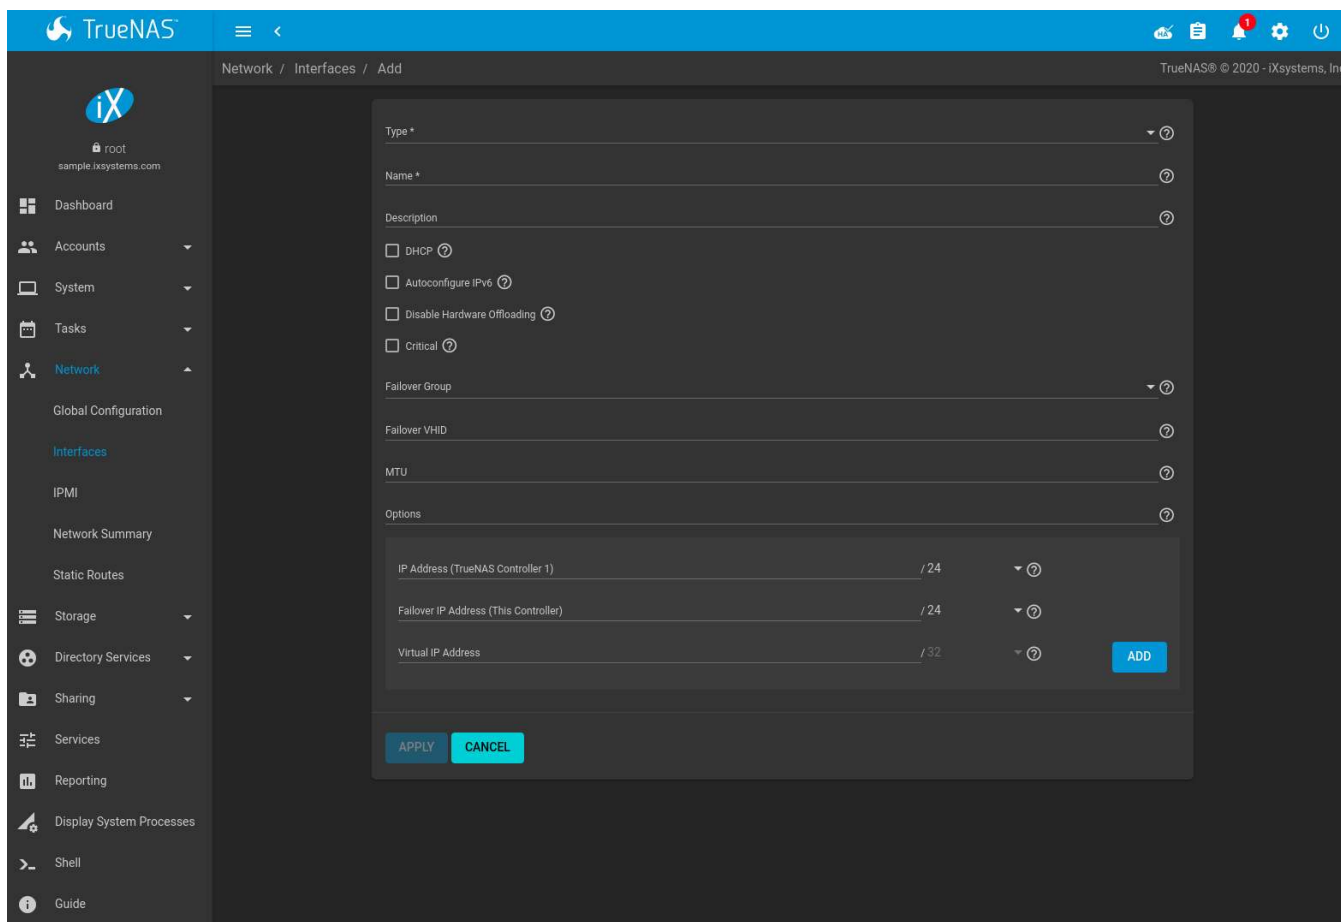


Fig. 7.2: Adding a Network Interface

Each *Type* of configurable network interface changes the available options. Table 7.2 shows which settings are available with each interface type.

Table 7.2: Interface Configuration Options

Setting	Value	Type	Description
Type	drop-down menu	All	Choose the type of interface. <i>Bridge</i> creates a logical link between multiple networks. <i>Link Aggregation</i> combines multiple network connections into a single interface. A virtual LAN (<i>VLAN</i>) partitions and isolates a segment of the connection.
Name	string	All	Enter a name to use for the the interface. Use the format laggX, vlanX, or bridgeX where X is a number representing a non-parent interface.
Description	string	All	Notes or explanatory text about this interface.
DHCP	checkbox	All	Enable DHCP (https://en.wikipedia.org/wiki/Dynamic_Host_Configuration_Protocol) to auto-assign an IPv4 address to this interface. Leave unset to create a static IPv4 or IPv6 configuration. Only one interface can be configured for DHCP.
Autoconfigure IPv6	drop-down menu	All	Automatically configure the IPv6 address with rtsol(8) (https://www.freebsd.org/cgi/man.cgi?query=rtsol). Only one interface can be configured this way.

Continued on next page

Table 7.2 – continued from previous page

Setting	Value	Type	Description
Disable Hardware Offloading	checkbox	All	Turn off hardware offloading for network traffic processing. WARNING: disabling hardware offloading can reduce network performance and is only recommended when the interface is managing <i>jails</i> (page 260), <i>plugins</i> (page 254), or <i>virtual machines (VMs)</i> (page 288).
Bridge Members	drop-down menu	Bridge	Network interfaces to include in the bridge.
Lagg Protocol	drop-down menu	Link Aggregation	Select the <i>Protocol Type</i> (page 124). <i>LACP</i> is the recommended protocol if the network switch is capable of active LACP. <i>Failover</i> is the default protocol choice and should only be used if the network switch does not support active LACP.
Lagg Interfaces	drop-down menu	Link Aggregation	Select the interfaces to use in the aggregation. Warning: Lagg creation fails when the selected interfaces have manually assigned IP addresses.
Parent Interface	drop-down menu	VLAN	Select the VLAN Parent Interface. Usually an Ethernet card connected to a switch port configured for the VLAN. A <i>bridge</i> cannot be selected as a parent interface. New <i>Link Aggregations</i> (page 124) are not available until the system is restarted.
Vlan Tag	integer	VLAN	The numeric tag provided by the switched network.
Priority Code Point	drop-down menu	VLAN	Select the <i>Class of Service</i> (https://en.wikipedia.org/wiki/Class_of_service). The available 802.1p Class of Service ranges from <i>Best effort (default)</i> to <i>Network control (highest)</i> .
MTU	integer	All	Maximum Transmission Unit, the largest protocol data unit that can be communicated. The largest workable MTU size varies with network interfaces and equipment. <i>1500</i> and <i>9000</i> are standard Ethernet MTU sizes. Leaving blank restores the field to the default value of <i>1500</i> .
Options	string	All	Additional parameters from <i>ifconfig(8)</i> (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 <i>this note</i> (page 125) about MTU and lagg interfaces.
IP Address	integer and drop-down menu	All	Static IPv4 or IPv6 address and subnet mask. Example: <i>10.0.0.3</i> and <i>/24</i> . Click <i>ADD</i> to add another IP address. Clicking <i>DELETE</i> removes that <i>IP Address</i> .

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

Saving a new interface adds an entry to the list in *Network* → *Interfaces*.

Expanding an entry in the list shows further details for that interface.

Editing an interface allows changing all the *interface options* (page 122) except the interface *Type* and *Name*.

Note: Interfaces cannot be edited or deleted when *High Availability (HA)* (page 81) has been enabled.

7.2.1 Network Bridges

A network bridge allows multiple network interfaces to function as a single interface.

To create a bridge, go to *Network* → *Interfaces* and click *ADD*. Choose *Bridge* as the *Type* and continue to configure the interface. See the *Interface Configuration Options table* (page 122) for descriptions of each option.

Enter `bridgeX` for the *Name*, where *X* is a unique interface number. Open the *Bridge Members* drop-down menu and select each interface that will be part of the bridge. Click *SAVE* to add the new bridge to *Network* → *Interfaces* and show options to confirm or revert the new network settings.

7.2.2 Link Aggregations

TrueNAS® 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 determines 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. TrueNAS® 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*.

Note: The *Failover* `lagg` protocol can interfere with HA (High Availability) systems and is disabled on those systems.

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.

7.2.2.1 LACP, MPIO, NFS, and ESXi

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

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

MPIO operates at the iSCSI protocol level. For example, if four IP addresses are created and there are four simultaneous TCP connections, MPIO will send the data over all available links. When configuring MPIO, make sure that the IP addresses on the interfaces are configured to be on separate subnets with non-overlapping netmasks, or configure static routes to do point-to-point communication. Otherwise, all packets will pass through one interface.

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

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

7.2.2.2 Creating a Link Aggregation

Before creating a link aggregation, see this [warning](#) (page 119) about changing the interface that the web interface uses.

To create a link aggregation, go to *Network → Interfaces* and click **ADD**. Choose *Link Aggregation* as the *Type* and continue to fill in the remaining configuration options. See the [Interface Configuration Options table](#) (page 122) for descriptions of each option.

Enter `laggX` for the *Name*, where *X* is a unique interface number. There are several *Lagg Protocol* options, but *LACP* is preferred. Choose *Failover* when the network switch does not support LACP. Open the *Lagg Interfaces* drop-down menu to associate NICs with the lagg device. Click **SAVE** to add the new aggregation to *Network → Interfaces* and show options to confirm or revert the new network settings.

7.2.2.3 Link Aggregation Options

Options are set at the lagg level from *Network → Interfaces*. Find the lagg interface, expand the entry with **>** (Expand), and click **EDIT**. 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. Setting options at the individual interface level is done by editing the parent interface in the same way as the lagg 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>).

7.2.3 VLANs

TrueNAS® uses [vlan\(4\)](https://www.freebsd.org/cgi/man.cgi?query=vlan) (<https://www.freebsd.org/cgi/man.cgi?query=vlan>) 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.

To add a new VLAN interface, go to *Network* → *Interfaces* and click *ADD*. Choose *VLAN* as the *Type* and continue filling in the remaining fields. See the [Interface Configuration Options table](#) (page 122) for descriptions of each option.

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 interface, then select *Network* → *Interfaces*, and click *:* (Options) and *Edit* for the parent interface. Enter `up` in the *Options* field and click *SAVE*. This brings up the parent interface. If an IP address is required, configure it using the rest of the options in the edit screen.

Warning: Creating a VLAN causes network connectivity to be interrupted and, if [Failover](#) (page 81) is configured, a failover event. The web interface requires confirming the new network configuration before it is permanently applied to the TrueNAS® system.

7.3 IPMI

The TrueNAS® Storage Array provides a built-in out-of-band management port which can be used to provide side-band management should the system become unavailable through the graphical administrative interface. 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. It can also be used to allow another person remote access to the system to assist with a configuration or troubleshooting issue.

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 7.3](#), provides a shortcut to the most basic IPMI configuration. Those already familiar with IPMI management tools can use them instead. [Table 7.3](#) summarizes the options available when configuring IPMI with the TrueNAS® web interface.

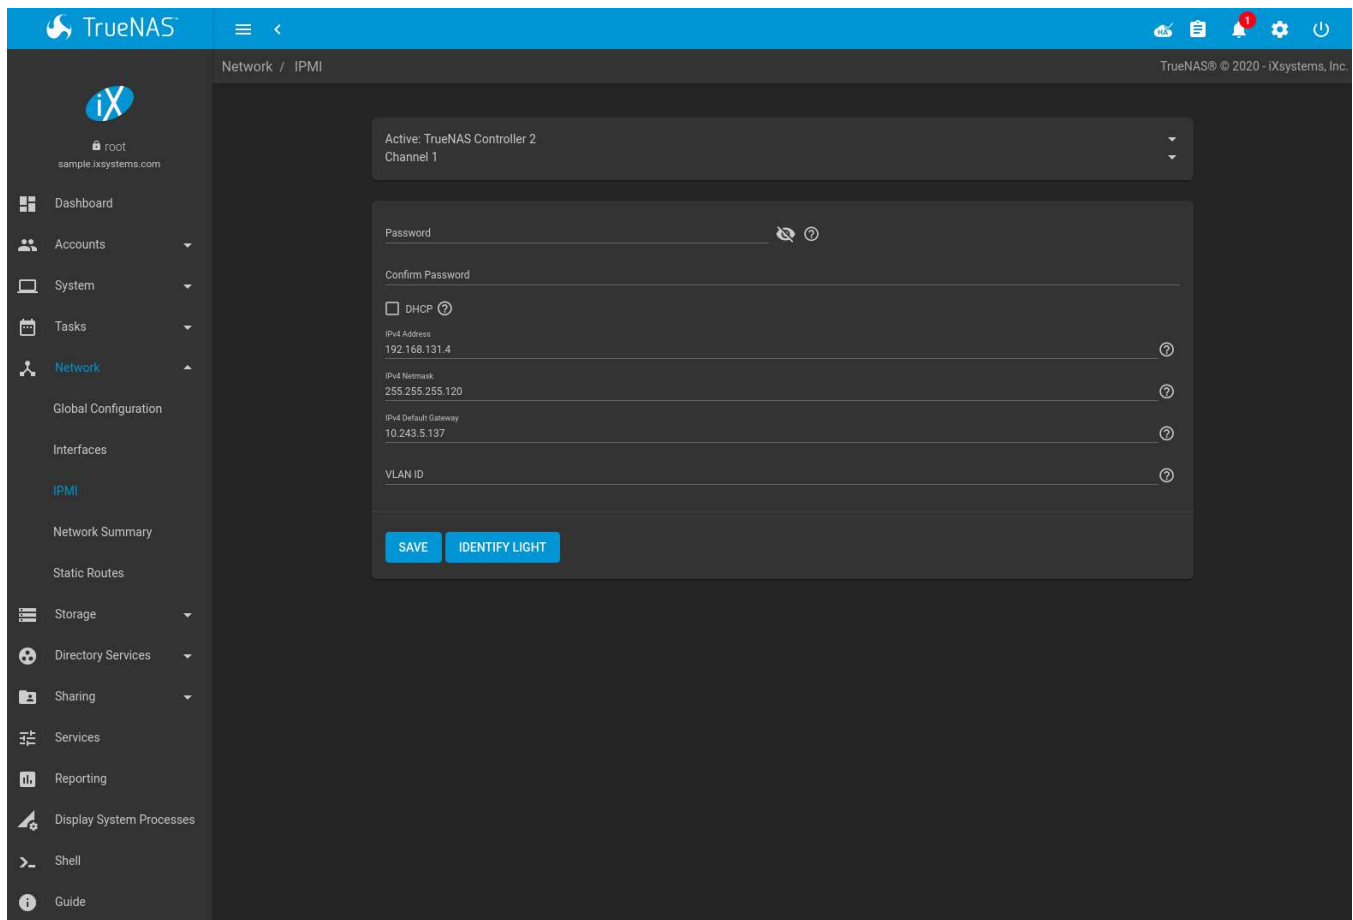


Fig. 7.3: IPMI Configuration

Table 7.3: IPMI Options

Setting	Value	Description
TrueNAS Controller	drop-down menu	Select a TrueNAS controller. All IPMI changes are applied to that TrueNAS controller.
Channel	drop-down menu	Select the communications channel (https://www.thomas-krenn.com/en/wiki/IPMI_Basics#Channel_Model) to use. Available channel numbers vary by hardware.
Password	string	Enter the password used to connect to the IPMI interface from a web browser. The maximum length accepted in the UI is 20 characters, but different hardware might require shorter passwords.
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.
IDENTIFY LIGHT	button	Show a dialog to activate an IPMI identify light on the compatible connected hardware.

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.

7.4 Network Summary

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

7.5 Static Routes

No static routes are defined on a default TrueNAS® 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 7.4](#).





Fig. 7.4: Adding a Static Route

The available options are summarized in [Table 7.4](#).

Table 7.4: 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.

STORAGE

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

- [Swap Space](#) (page 130): Change the swap space size.
- [Pools](#) (page 130): create and manage storage pools.
- [Snapshots](#) (page 152): manage local snapshots.
- [VMware-Snapshots](#) (page 155): coordinate OpenZFS snapshots with a VMware datastore.
- [Disks](#) (page 156): view and manage disk options.
- [Importing a Disk](#) (page 162): import a **single** disk that is formatted with the UFS, NTFS, MSDOS, or EXT2 filesystem.
- [Multipaths](#) (page 163): View multipath information for systems with compatible hardware.

Note: When using an HA (High Availability) TrueNAS® system, connecting to the web interface on the standby TrueNAS controller only shows a screen indicating that it is the standby TrueNAS controller. All of the options discussed in this chapter can only be configured on the active TrueNAS controller.

8.1 Swap Space

Swap is space on a disk set aside to be used as memory. When the TrueNAS® system runs low on memory, less-used data can be “swapped” onto the disk, freeing up main memory.

For reliability, TrueNAS® creates swap space as mirrors of swap partitions on pairs of individual disks. For example, if the system has three hard disks, a swap mirror is created from the swap partitions on two of the drives. The third drive is not used, because it does not have redundancy. On a system with four drives, two swap mirrors are created.

Swap space is allocated when drives are partitioned before being added to a [vdev](#) (page 309). A 2 GiB partition for swap space is created on each data drive by default. The size of space to allocate can be changed in *System* → *Advanced* in the *Swap size in Gib* field. Changing the value does not affect the amount of swap on existing disks, only disks added after the change. This does not affect log or cache devices, which are created without swap. Swap can be disabled by entering 0, but that is **strongly discouraged**.

8.2 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.**

8.2.1 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.

Go to *Storage* → *Pools* and click *ADD*. Select *Create new pool* and click *CREATE POOL* to open the screen shown in Figure 8.1.

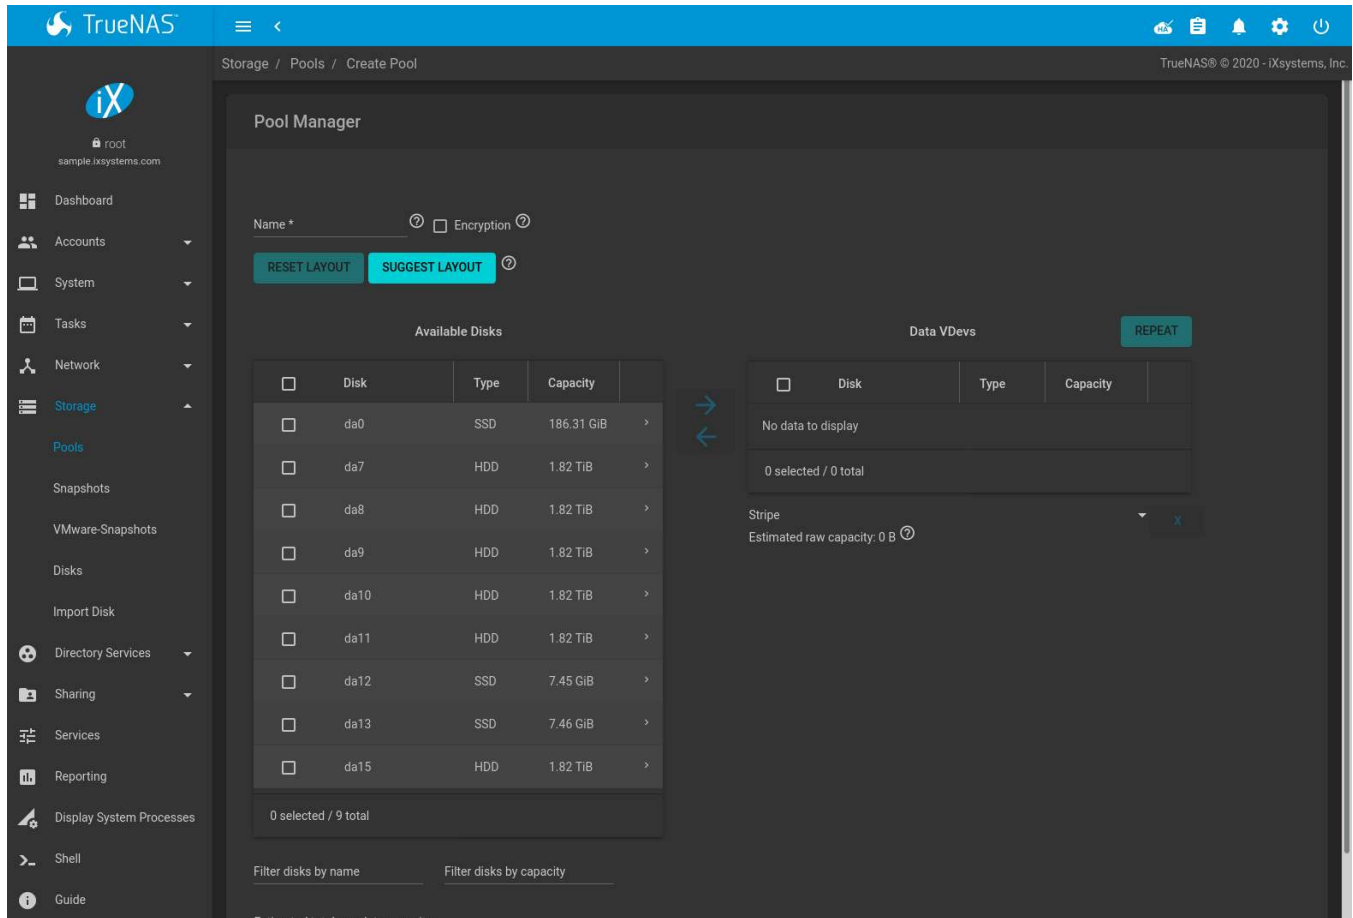


Fig. 8.1: Creating a Pool

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 dialog displays a reminder to back up the *encryption key* (page 135). The data on the disks is inaccessible without the key. Select *Confirm* then click *I UNDERSTAND*.

Warning: Refer to the warnings in [Managing Encrypted Pools](#) (page 133) before enabling encryption!

From the *Available Disks* section, select disks to add to the pool. Enter a value in *Filter disks by name* or *Filter disks by capacity* to change the displayed disk order. These fields support [PCRE regular expressions](http://php.net/manual/en/reference.pcre.pattern.syntax.php) (<http://php.net/manual/en/reference.pcre.pattern.syntax.php>) for filtering. For example, to show only *da* and *nvd* disks in *Available Disks*, type `^(da)|(nvd)` in *Filter disks by name*.

Type and maximum capacity is displayed for available disks. To show the disk *Rotation Rate*, *Model*, and *Serial*, click *>* (Expand).

After selecting disks, click the right arrow to add them to the *Data VDevs* section. The usable space of each disk in a vdev is limited to the size of the smallest disk in the vdev. Additional data vdevs must have the same configuration as the initial vdev.

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

After adding one data vdev, additional data vdevs can be added with *REPEAT*. This creates additional vdevs of the same layout as the initial vdev. Select the number of additional vdevs and click *REPEAT VDEV*.

RESET LAYOUT returns all disks to the *Available Disks* area and closes all but one *Data VDevs* table.

SUGGEST LAYOUT arranges all disks in 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 309) 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 dialog shows 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 162) and *Importing a Pool* (page 139) 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 re-store 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 dialog provides a link to *Download Recovery Key*. Click the link and save the key to a safe location. When finished, click *DONE*.

Figure 8.2 shows the new *pool1*. Select the pool to see more information. The first entry in the list represents the root dataset and has the same name as the pool.

The *Available* column shows the estimated storage space before [compression](https://en.wikipedia.org/wiki/Data_compression) (https://en.wikipedia.org/wiki/Data_compression). The *Used* column shows the estimated space used after compression. These numbers come from `zfs list`.

Other utilities can report different storage estimates. For example, the available space shown in `zpool status` is the cumulative space of all drives in the pool, regardless of pool configuration or compression.

Other information shown is 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.

Pool status is indicated by one of these symbols:

Table 8.1: Pool Status

Symbol	Color	Meaning
✔ HEALTHY	Green	The pool is healthy.
⚠ DEGRADED	Orange	The pool is in a degraded state.
❓ UNKNOWN	Blue	Pool status cannot be determined.
🔒 LOCKED	Yellow	The pool is locked.
✖ Pool Fault	Red	The pool has a critical error.

There is an option to *Upgrade Pool*. This upgrades the pool to the latest *ZFS Feature Flags* (page 312). See the warnings in *Upgrading a ZFS Pool* (page 71) before selecting this option. This button does not appear when the pool is running the latest version of the feature flags.

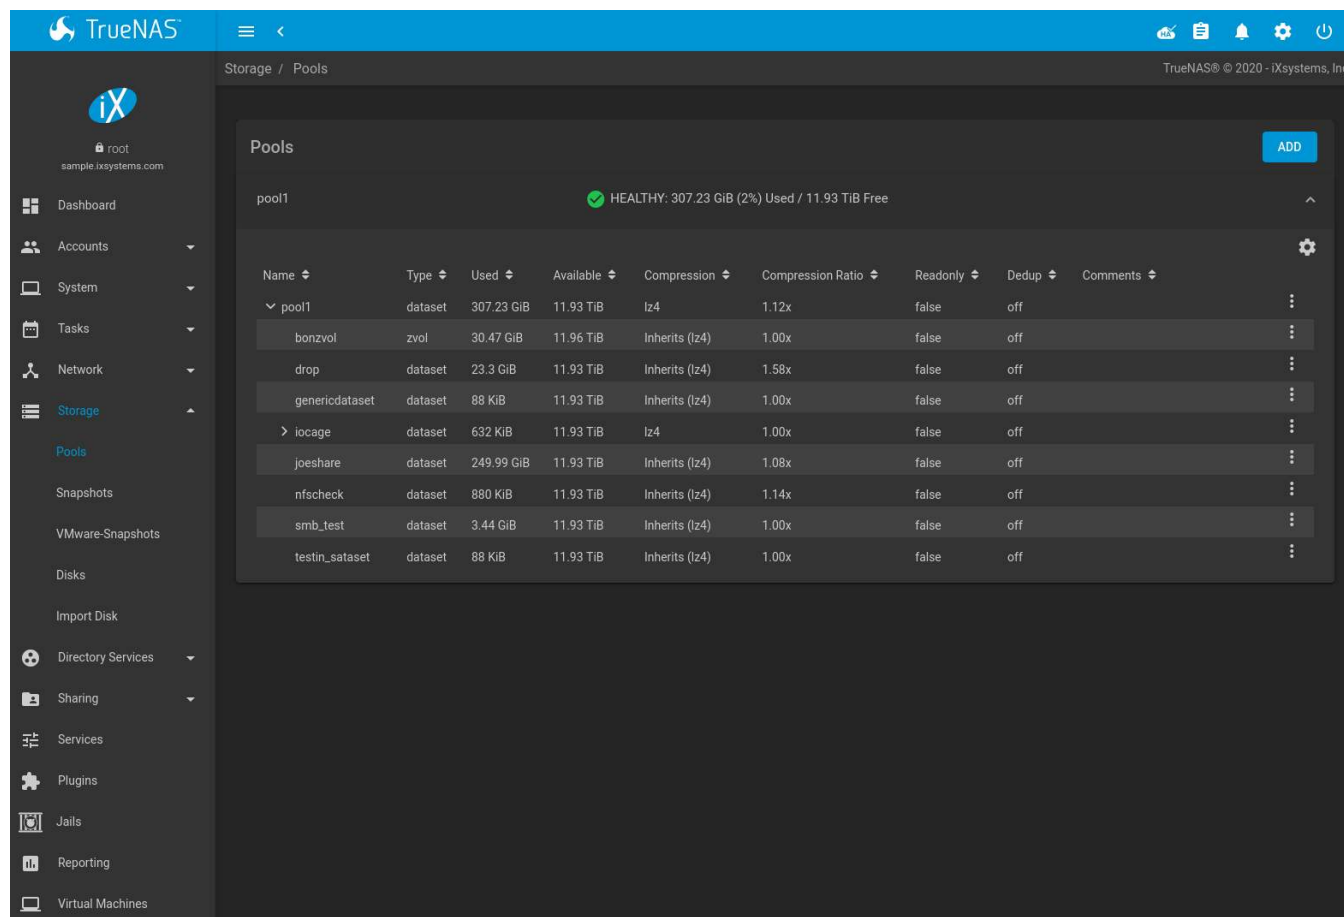


Fig. 8.2: Viewing Pools

Creating a pool adds a card to the *Dashboard*. Available space, disk details, and pool status is shown on the card. The background color of the card indicates the pool status:

- Green: healthy or locked
- Yellow: unknown, offline, or degraded
- Red: faulted or removed

8.2.2 Managing Encrypted Pools

TrueNAS® uses [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. This type of encryption is intended to protect against the risks of data being read or copied when the system is powered

down, when the pool is locked, or when disks are physically stolen.

TrueNAS® encrypts disks and pools, not individual filesystems. The partition table on each disk is not encrypted, but only identifies the location of partitions on the disk. On an encrypted pool, the data in each partition is encrypted. These are generally called “encrypted drives”, even though the partition table is not encrypted. To use drive firmware to completely encrypt the drive, see [Self-Encrypting Drives](#) (page 41).

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. Data in memory, including ARC, is not encrypted. ZFS data on disk, including L2ARC and SLOG, are encrypted if the underlying disks are encrypted. Swap data on disk is always encrypted.

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.** Please benchmark encrypted pools before using them in production.

Creating an encrypted pool means GELI encrypts the data on the disk and generates a *master key* to decrypt this data. This master key is also encrypted. 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.

There are two *user keys* that can be used to unlock the master key and then decrypt the disks. In TrueNAS®, these user keys are named the **encryption key** and the **recovery key**. Because data cannot be read without first providing a key, encrypted disks containing sensitive data can be safely removed, reused, or discarded without secure wiping or physical destruction of the media.

When discarding disks that still contain encrypted sensitive data, the encryption and recovery keys should also be destroyed or securely deleted. Keys that are not destroyed must be stored securely and kept physically separate from the discarded disks. Data is vulnerable to decryption when the encryption key is present with the discarded disks or can be obtained by the same person who gains access to the disks.

This encryption method is **not** designed to protect against unauthorized access when the pool is already unlocked. Before sensitive data is stored on the system, ensure that only authorized users have access to the web interface and that permissions with appropriate restrictions are set on shares.

Here are some important points about TrueNAS® behavior to remember when creating or using an encrypted pool:

- At present, there is no one-step way to encrypt an existing pool. The data must be copied to an existing or new encrypted pool. After that, the original pool and any unencrypted backup should be destroyed to prevent unauthorized access and any disks that contained unencrypted data should be wiped.
- Hybrid pools are not supported. Added vdevs must match the existing encryption scheme. [Extending a Pool](#) (page 137) automatically encrypts a new vdev being added to an existing encrypted pool.
- TrueNAS® encryption differs from the encryption used in the Oracle proprietary version of ZFS. To convert between these formats, both pools must be unlocked, and the data copied between them.
- Each pool has a separate encryption key. Pools can also add a unique recovery key to use if the passphrase is forgotten or encryption key invalidated.
- Encryption applies to a pool, not individual users. The data from an unlocked pool is accessible to all users with permissions to access it. Encrypted pools with a passphrase can be locked on demand by users that know the passphrase. Pools are automatically locked when the system is shut down.
- 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.
- Encrypted pools that have no passphrase are unlocked at startup. Pools with a passphrase remain locked until a user enters the passphrase to unlock them.

8.2.2.1 Encryption and Recovery Keys

TrueNAS® generates a randomized *encryption key* whenever a new encrypted pool is created. This key is stored in the *system dataset* (page 48). It is the primary key used to unlock the pool each time the system boots. Creating a passphrase for the pool adds a passphrase component to the encryption key and allows the pool to be locked.

A pool encryption key backup can be downloaded to allow disk decryption on a different system in the event of failure or to allow the TrueNAS® stored key to be deleted for extra security. The combination of encryption key location and passphrase usage 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 TrueNAS® administrator.
- *Key not stored locally*: the encrypted pool is not accessible until the TrueNAS® administrator uploads the key file. When the key also has a passphrase, it must be provided with the key file.

Encrypted pools cannot be locked in the web interface until a passphrase is created for the encryption key.


The recovery key is an optional keyfile that is generated by TrueNAS®, provided for download, and wiped from the system. It is designed as an emergency backup to unlock or import an encrypted pool if the passphrase is forgotten or the encryption key is somehow invalidated. This file is not stored anywhere on the TrueNAS® system and only one recovery key can exist for each encrypted pool. Adding a new recovery key invalidates any previously downloaded recovery key file for that pool.

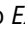
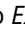
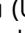
Existing encryption or recovery keys can be invalidated in several situations:

- An encryption re-key invalidates all encryption and recovery keys as well as an existing passphrase.
- Using a recovery key file to import an encrypted pool invalidates the existing encryption key and passphrase for that pool. TrueNAS® generates a new encryption key for the imported pool, but a new passphrase must be created before the pool can be locked.
- Creating or changing a passphrase invalidates any existing recovery key.
- Adding a new recovery key invalidates any existing recovery key files for the pool.
- *Extending a Pool* (page 137) invalidates all encryption and recovery keys as well as an existing passphrase.

Be sure to download and securely store copies of the most current encryption and recovery keys. Protect and backup encryption key passphrases. **Losing the encryption and recovery keys or the passphrase can result in irrevocably losing all access to the data stored in the encrypted pool!**

8.2.2.2 Encryption Operations

Encryption operations are seen by clicking  (Encryption Options) for the encrypted pool in *Storage* → *Pools*. These options are available:

- *Lock*: Only appears after a passphrase is created. Locking a pool restricts data accessibility in TrueNAS® until the pool is unlocked. Selecting this action requires entering the passphrase. The pool status changes to **LOCKED**, *Pool Operations* are limited to *Export/Disconnect*, and  (Encryption Options) changes to  (Unlock).
- *Unlock*: Decrypt the pool by clicking  (Unlock) and entering the passphrase *or* uploading the recovery key file. Only the passphrase is used when both a passphrase and a recovery key are entered. The services listed in *Restart Services* restart when the pool is unlocked. This enables TrueNAS® to begin accessing the decrypted data. Individual services can be prevented from restarting by opening *Restart Services* and deselecting them. Deselecting services can prevent them from properly accessing the unlocked pool.
- *Encryption Key/Passphrase*: Create or change the encryption key passphrase and download a backup of the encryption key. Unlike a password, a passphrase can contain spaces and is typically a series of words. A good passphrase is easy to remember but hard to guess.

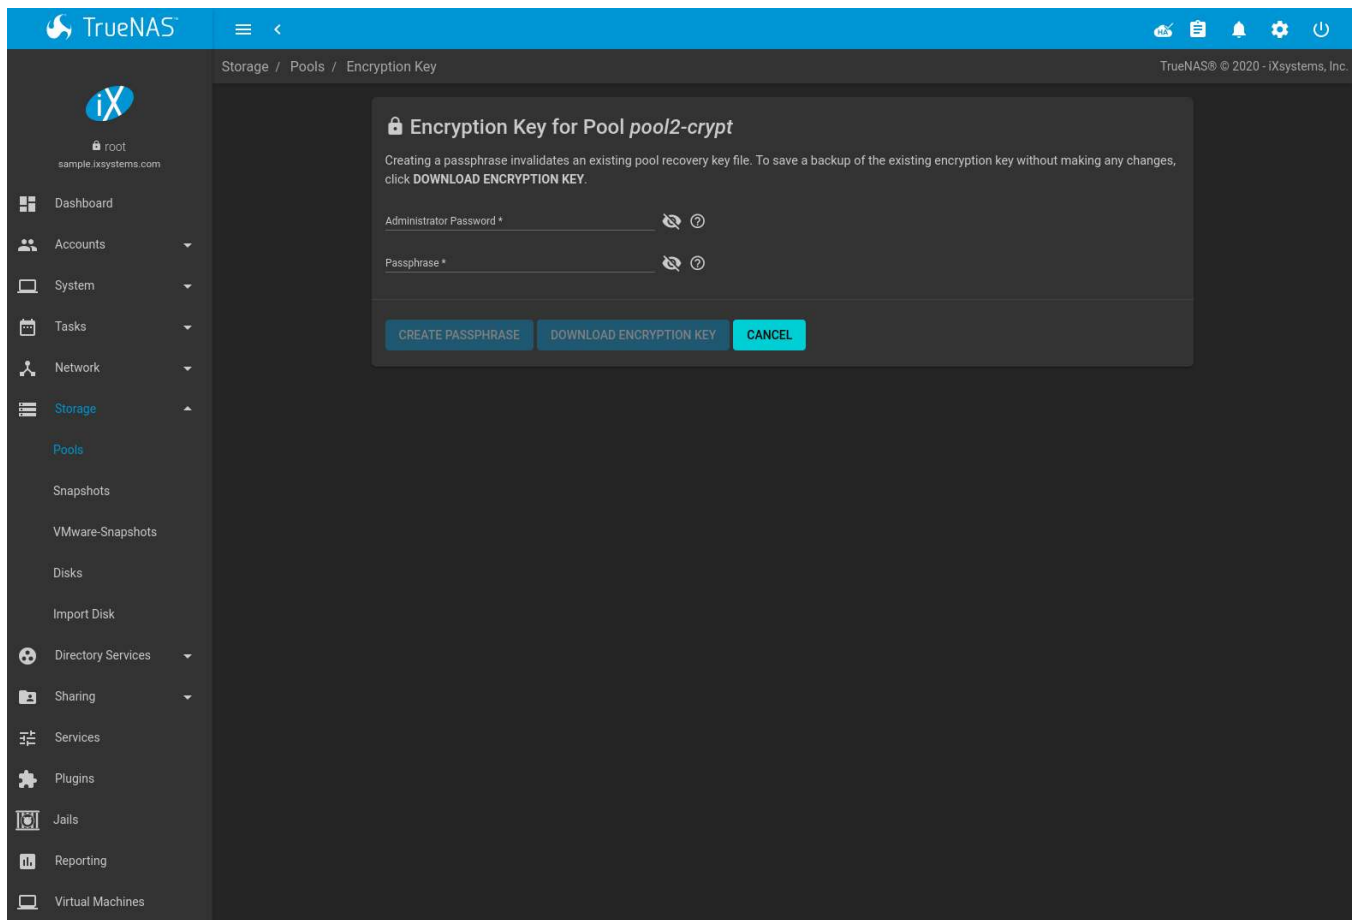


Fig. 8.3: Encryption Key/Passphrase Options

The administrator password is required for encryption key changes. Setting *Remove Passphrase* invalidates the current pool passphrase. Creating or changing a passphrase invalidates the pool recovery key.

- *Recovery Key*: Generate and download a new recovery key file or invalidate an existing recovery key. The TrueNAS® administrative password is required. Generating a new recovery key file invalidates previously downloaded recovery key files for the pool.
- *Reset Keys*: Reset the encryption on the pool GELI master key and invalidate all encryption keys, recovery keys, and any passphrase for the pool. A dialog opens to save a backup of the new encryption key. A new passphrase can be created and a new pool recovery key file can be downloaded. The administrator password is required to reset pool encryption.

If a key reset 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.

Note: A key reset is not allowed if *Failover* (page 81) (High Availability) has been enabled and the standby TrueNAS controller is down.

8.2.3 Adding Cache or Log Devices

Pools (page 130) 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 309) 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, *Extend* (page 137) that pool.

8.2.4 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*.

8.2.5 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 detached 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 TrueNAS®, 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, *Extend* (page 137) that pool.

8.2.6 Extending a Pool

To increase the capacity of an existing pool, click the pool name, ⚙ (Settings), then *Extend*.

If the existing pool is *encrypted* (page 133), an additional warning message shows a reminder that **extending a pool resets the passphrase and recovery key**. Extending an encrypted pool opens a dialog to download the new encryption key file. Remember to use the *Encryption Operations* (page 135) to set a new passphrase and create a new recovery key file.

When adding disks to increase the capacity of a pool, ZFS supports the addition of virtual devices, or *vdevs*, to an existing ZFS pool. **After a vdev is created, more drives cannot be added to that vdev**, but a new vdev can be striped with another of the **same type** to increase the overall size of the pool. To extend a pool, the vdev being added must be the same type as existing vdevs. The *EXTEND* button is only enabled when the vdev being added is the same type as the existing vdevs. Some vdev extending examples:

- to extend a ZFS mirror, add the same number of drives. The result is a striped mirror. For example, if ten new drives are available, a mirror of two drives could be created initially, then extended by adding another mirror of two drives, and repeating three more times until all ten drives have been added.
- to extend a three-drive RAIDZ1, add another three drives. The resulting pool is a stripe of two RAIDZ1 vdevs, similar to RAID 50 on a hardware controller.
- to extend a four-drive RAIDZ2, add another four drives. The result is a stripe of RAIDZ2 vdevs, similar to RAID 60 on a hardware controller.

8.2.7 Export/Disconnect a Pool

Export/Disconnect is used to cleanly disconnect a pool from the system. This is used before physically disconnecting the pool so it can be imported on another system, or to optionally detach and erase the pool so the disks can be reused.

To export or destroy an existing pool, click the pool name, ⚙ (Settings), then *Export/Disconnect*. A dialog shows which system *Services* (page 222) will be disrupted by exporting the pool and additional warnings for encrypted pools. Keep or erase the contents of the pool by setting the options shown in *Figure 8.4*.

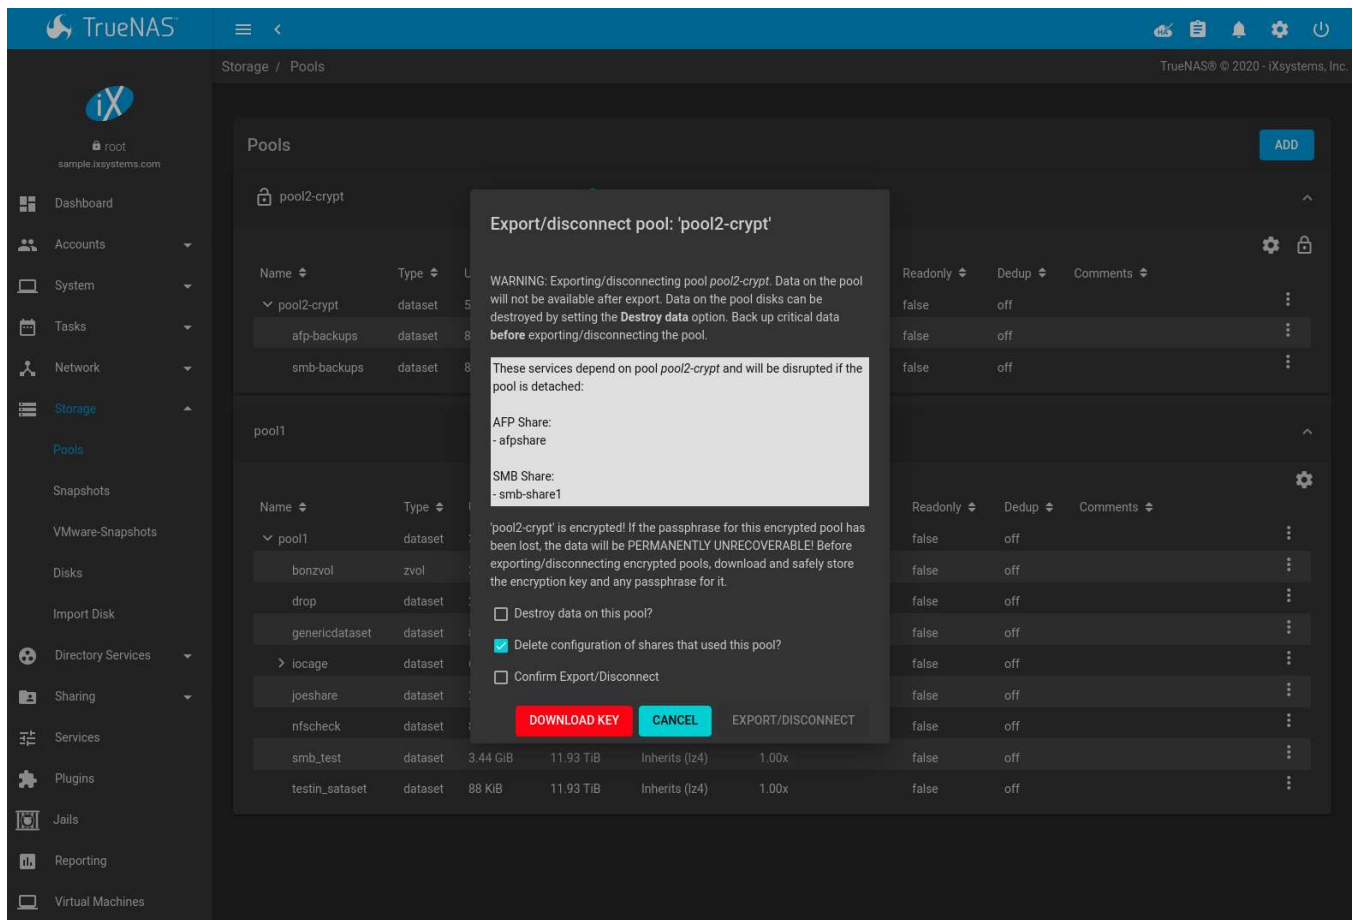


Fig. 8.4: Export/Disconnect a Pool

Note: At least one pool is required for *High Availability (HA)* (page 81). If HA is enabled and only one pool is connected, HA must be disabled before that pool can be removed.

Warning: Do not export/disconnect an encrypted pool if the passphrase has not been set! **An encrypted pool cannot be reimported without a passphrase!** When in doubt, use the instructions in *Managing Encrypted Pools* (page 133) to set a passphrase.

The *Export/Disconnect Pool* screen provides these options:

Table 8.2: Export/Disconnect Pool Options

Setting	Description
Destroy data on this pool?	Destroy all data on the disks in the pool. This action cannot be undone.
Delete configuration of shares	Delete any share configurations set up on the pool.
Confirm export/disconnect	Confirm the export/disconnect operation.

If the pool is encrypted, *DOWNLOAD KEY* is also shown to download the *encryption key* (page 135) for that pool.

To *Export/Disconnect* the pool and keep the data and configurations of shares, set **only** *Confirm export/disconnect* and click *EXPORT/DISCONNECT*.

To instead destroy the data and share configurations on the pool, also set the *Destroy data on this pool?* option. To verify that data on the pool is to be destroyed, type the name of the pool and click *EXPORT/DISCONNECT*. Data on the pool is destroyed, including share configuration, zvols, datasets, and the pool itself. The disk is returned to a raw state.

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

8.2.8 Importing a Pool

A pool that has been exported and disconnected from the system can be reconnected with *Storage* → *Pools* → *Add*, then selecting *Import an existing pool*. This works for pools that were exported/disconnected from the current system, created on another system, or to reconnect a pool after reinstalling the TrueNAS® 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 TrueNAS® system. This prevents an “in use by another machine” error during the import to TrueNAS®.

Existing ZFS pools can be imported by clicking *Storage* → *Pools* and *ADD*. Select *Import an existing pool*, then click *NEXT* as shown in Figure 8.5.

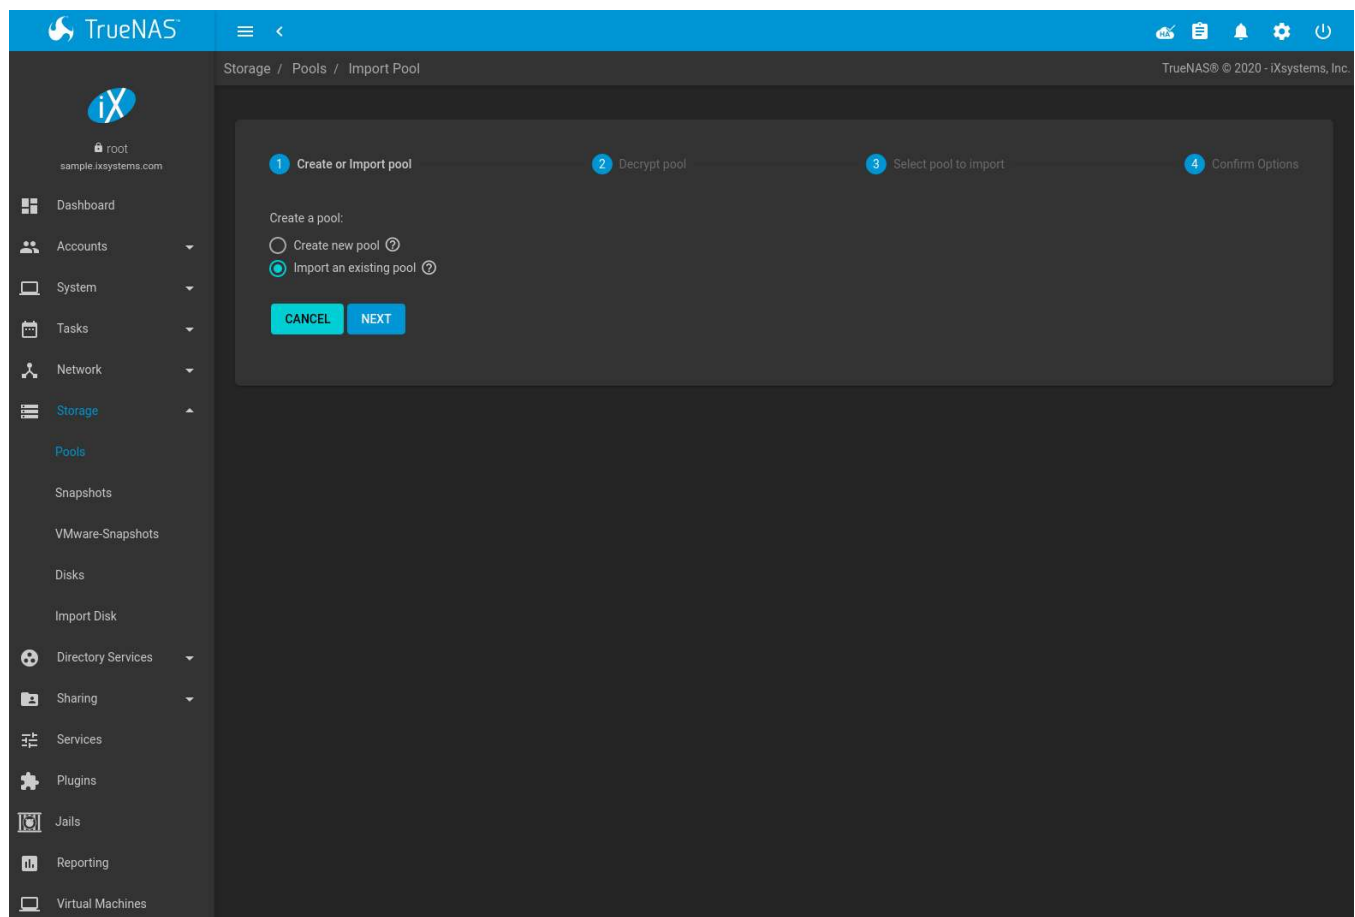


Fig. 8.5: Pool Import

To import a pool, click *No, continue with import* then *NEXT* as shown in Figure 8.6.

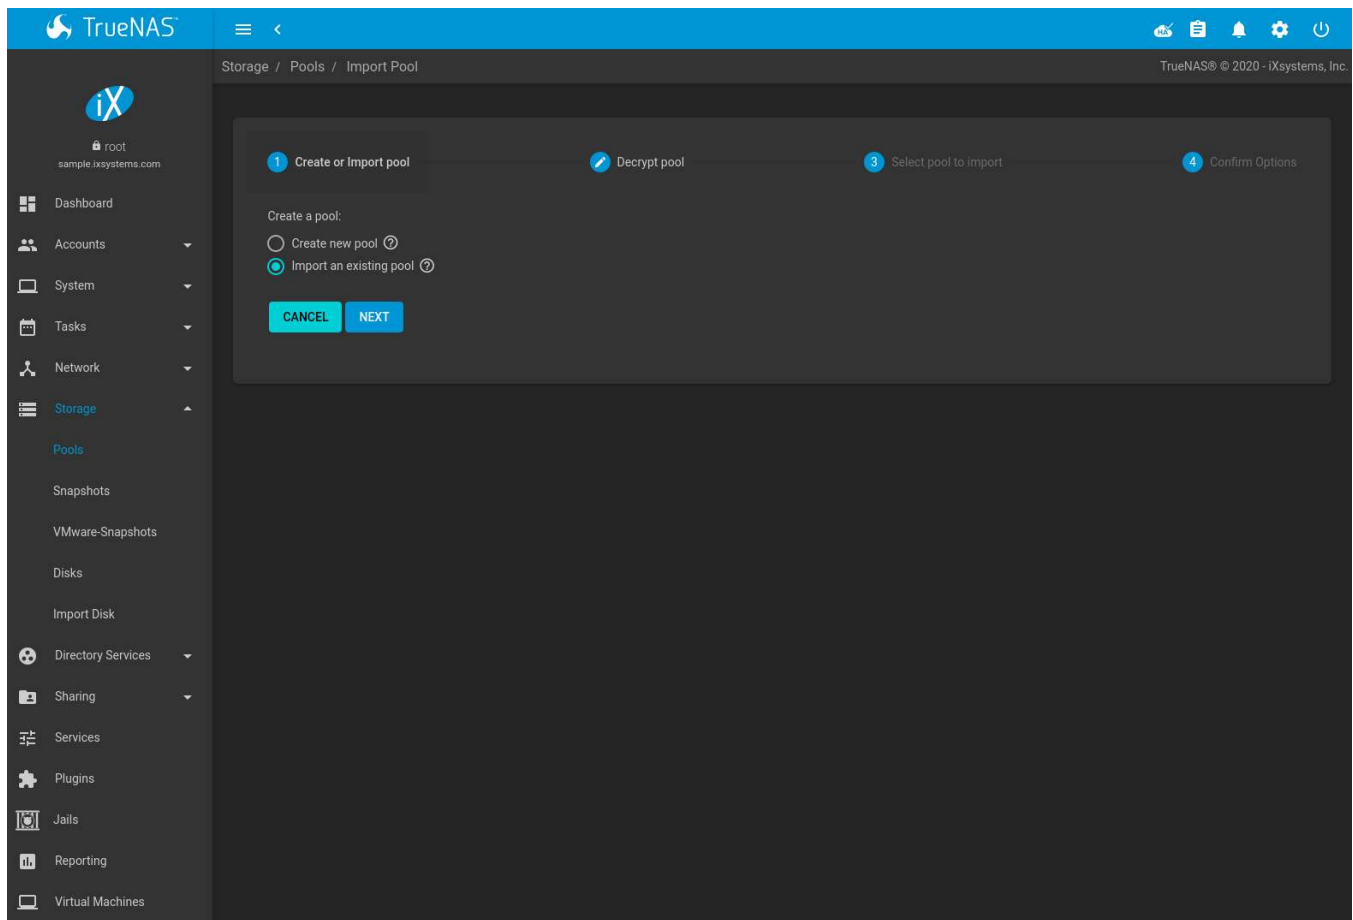


Fig. 8.6: Importing a Pool

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

Before importing an *encrypted pool* (page 133), disks must first be decrypted. Click *Yes, decrypt the disks*. This is shown in [Figure 8.7](#).

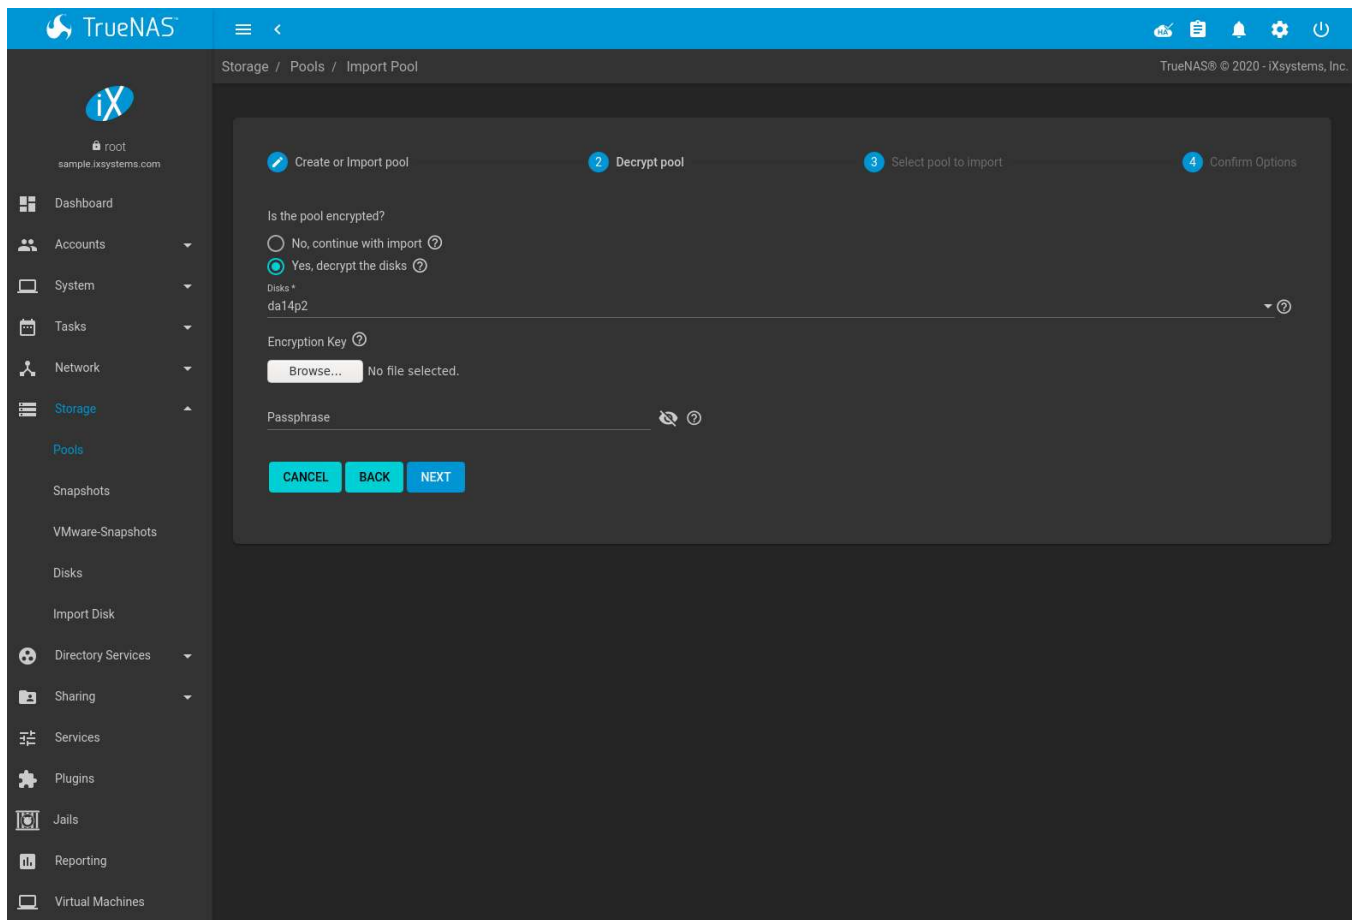


Fig. 8.7: Decrypting Disks Before Importing a Pool

Use the *Disks* dropdown menu to select the disks to decrypt. Click *Browse* to select the encryption key file stored on the client system. Enter the *Passphrase* associated with the encryption key, then click *NEXT* to continue importing the pool.

Danger: The encryption key file and passphrase are 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 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 133) for instructions on managing keys.

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

Note: For security reasons, encrypted pool keys are not saved in a configuration backup file. When TrueNAS® has been installed to a new device and a saved configuration file restored to it, the keys for encrypted disks will not be present, and the system will not request them. To correct this, export the encrypted pool with (Configure) → *Export/Disconnect*, making sure that *Destroy data on this pool?* is **not** set. Then import the pool again. During the import, the encryption keys can be entered as described above.

8.2.9 Viewing Pool Scrub Status

Scrubs and how to set their schedule are described in more detail in [Scrub Tasks](#) (page 111).

To view the scrub status of a pool, click the pool name, ⚙️ (Settings), then *Status*. The resulting screen will display the status and estimated time remaining for 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.

8.2.10 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 8.8](#).

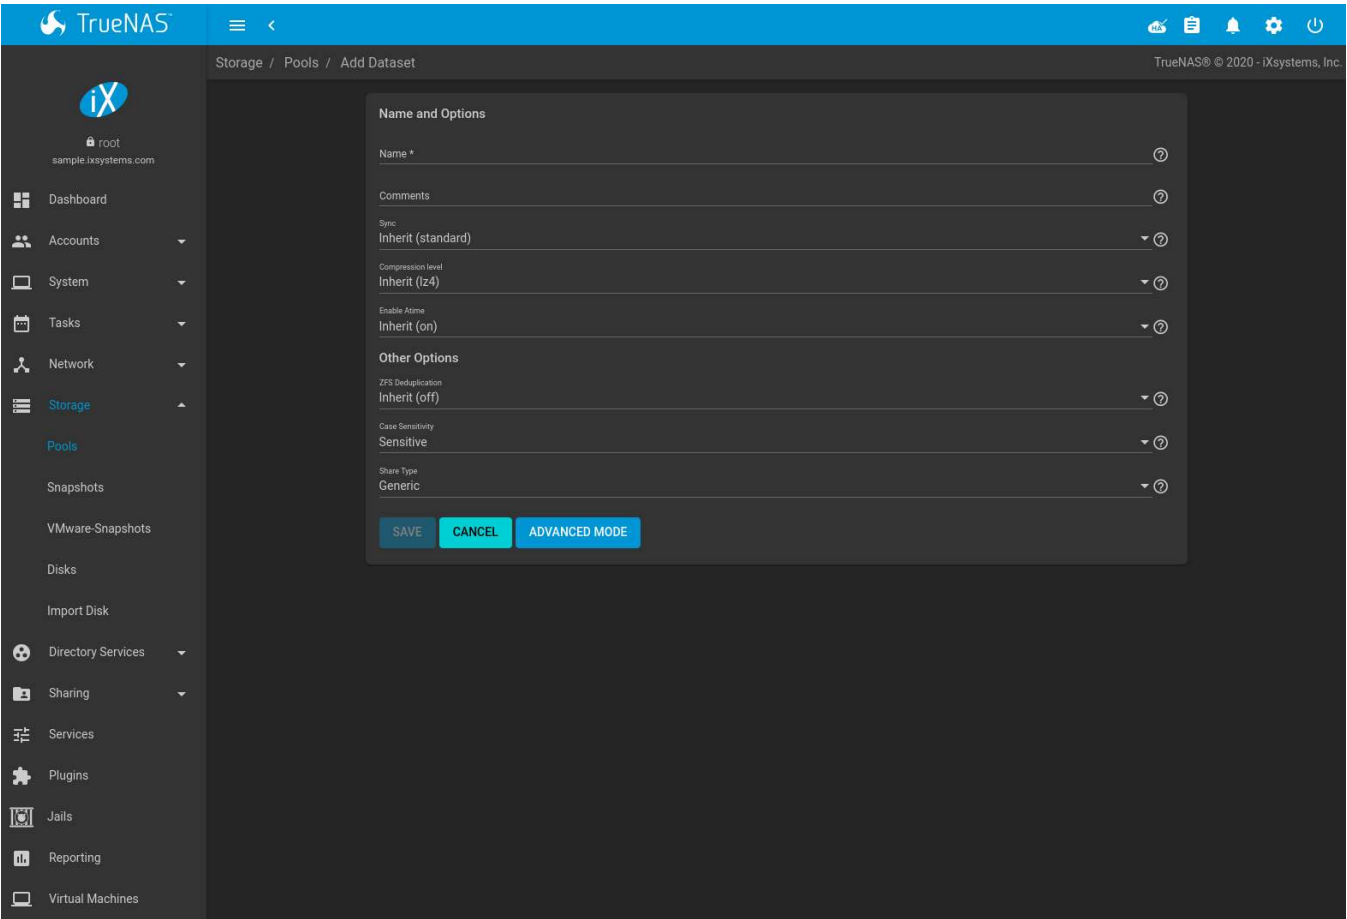


Fig. 8.8: Creating a ZFS Dataset

Table 8.3 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 8.3: Dataset Options

Setting	Value	Advanced Mode	Description
Name	string		Required. Enter a unique name for the dataset.
Comments	string		Enter any additional comments or user notes about this dataset.
Sync	drop-down menu		Set 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 145) for a description of the available algorithms.
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 warning alert at, %	integer	✓	Set Inherit to apply the same quota warning alert settings as the parent dataset.
Quota critical alert at, %	integer	✓	Set Inherit to apply the same quota critical alert settings as the parent dataset.
Quota for this dataset and all children	integer	✓	A specified value applies to both this dataset and any child datasets.
Quota warning alert at, %	integer	✓	Set Inherit to apply the same quota warning alert settings as the parent dataset.
Quota critical alert at, %	integer	✓	Set Inherit to apply the same quota critical alert settings as the parent dataset.
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		Do not change this setting unless instructed to do so by your iXsystems support engineer.
Read-only	drop-down menu	✓	Choices are <i>Inherit</i> , <i>On</i> , or <i>Off</i> .
Exec	drop-down menu	✓	Choices are <i>Inherit</i> , <i>On</i> , or <i>Off</i> . Setting to <i>Off</i> prevents the installation of Plugins (page 254) or Jails (page 260).
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.

Continued on next page

Table 8.3 – continued from previous page

Setting	Value	Advanced Mode	Description
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.
ACL Mode	drop-down menu	✓	Determine how <code>chmod(2)</code> (https://www.freebsd.org/cgi/man.cgi?query=chmod) behaves when adjusting file ACLs. See the zfs(8) aclmode property (https://www.freebsd.org/cgi/man.cgi?query=zfs). <i>Passthrough</i> only updates ACL entries that are related to the file or directory mode. <i>Restricted</i> does not allow <code>chmod</code> to make changes to files or directories with a non-trivial ACL. An ACL is trivial if it can be fully expressed as a file mode without losing any access rules. Setting the <i>ACL Mode</i> to <i>Restricted</i> is typically used to optimize a dataset for <i>SMB sharing</i> (page 209), but can require further optimizations. For example, configuring an <i>rsync</i> (page 90) with this dataset could require adding <code>--no-perms</code> in the task <i>Extra options</i> field.
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). This can only be set when creating a new dataset.
Share Type	drop-down menu		Select the type of share that will be used on the dataset. Choose between <i>Generic</i> for most sharing options or <i>SMB</i> for a <i>SMB share</i> (page 209). Choosing <i>SMB</i> sets the <i>ACL Mode</i> to <i>Restricted</i> and <i>Case Sensitivity</i> to <i>Insensitive</i> . This field is only available when creating a new dataset.

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 145) for more information about zvols.

Edit Options: edit the pool properties described in [Table 8.8](#). Note that *Dataset Name* and *Case Sensitivity* are read-only as they cannot be edited after dataset creation.

Edit Permissions: refer to [Setting Permissions](#) (page 147) for more information about permissions.

Danger: Removing a dataset is a permanent action and results in data loss!

Edit ACL: see [ACL Management](#) (page 148) for details about modifying an Access Control List (ACL).

Delete Dataset: removes the dataset, snapshots of that dataset, and any objects stored within the dataset. To remove the dataset, set *Confirm*, click *DELETE DATASET*, verify that the correct dataset to be deleted has been chosen by entering the dataset name, and click *DELETE*. When the dataset has active shares or is still being used by other parts of the system, the dialog shows what is still using it and allows forcing the deletion anyway. **Caution:** forcing the deletion of an in-use dataset can cause data loss or other problems.

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. A dialog opens to name the snapshot. Options to include child datasets in the snapshot and synchronize with VMware can also be shown. To schedule snapshot creation, use [Periodic Snapshot Tasks](#) (page 97).

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.

8.2.10.1 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.

8.2.11 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 231) 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 8.9](#).

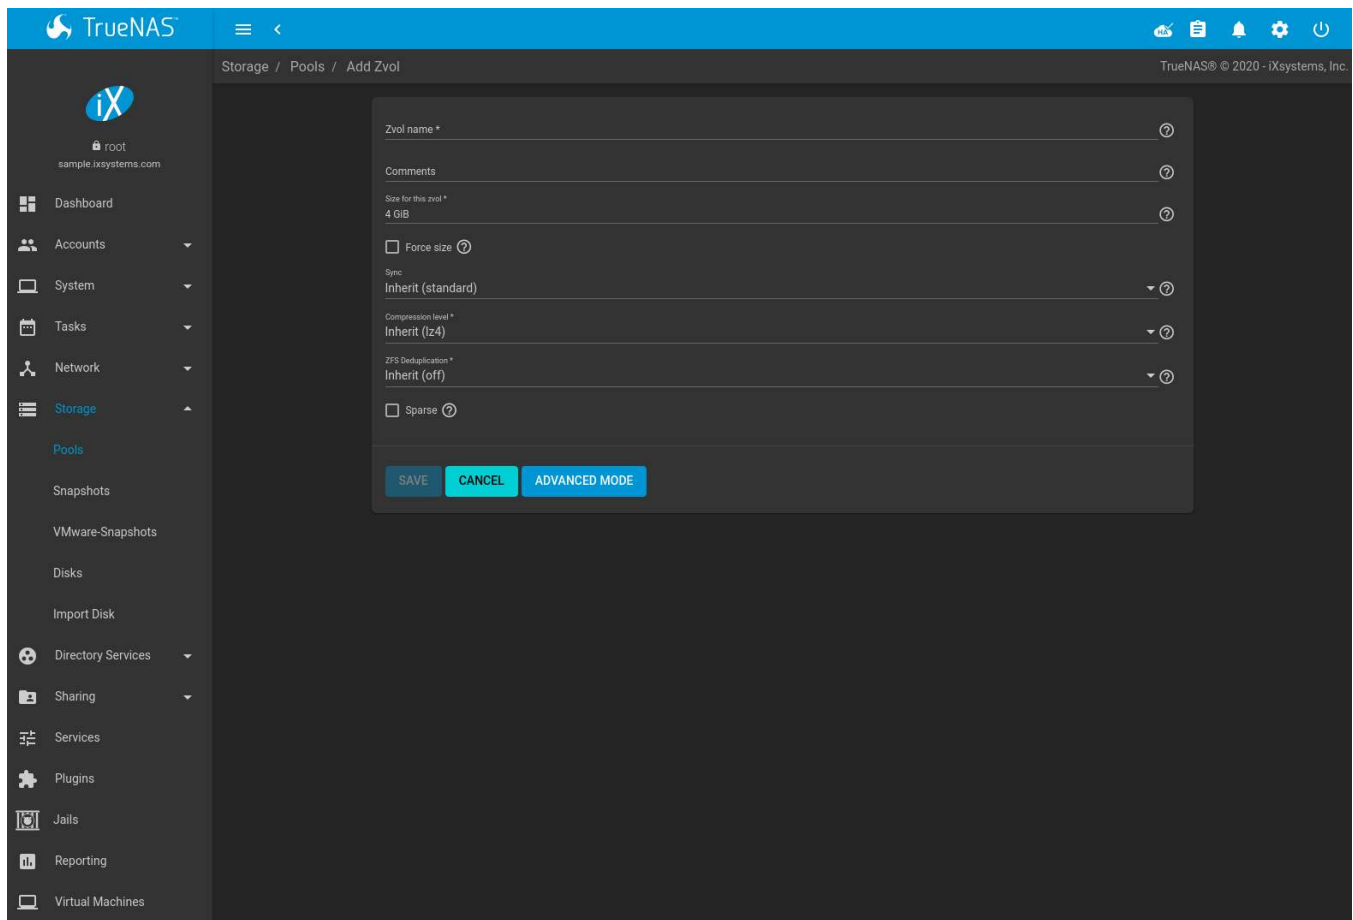


Fig. 8.9: Adding a Zvol

The configuration options are described in [Table 8.4](#).


Table 8.4: 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. Units like <code>t</code> , <code>TiB</code> , and <code>G</code> can be used. The size of the zvol can be increased later, but cannot be reduced. 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.

Continued on next page

Table 8.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Compression level	drop-down menu		Compress data to save space. Refer to Compression (page 145) for a description of the available algorithms.
ZFS Deduplication	drop-down menu		Do not change this setting unless instructed to do so by your iXsystems support engineer.
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*, *Edit Zvol*, *Create Snapshot*, and, for an existing zvol snapshot, *Promote Dataset* options.

Similar to datasets, a zvol name cannot be changed.


Choosing a zvol for deletion shows a warning that all snapshots of that zvol will also be deleted.

8.2.12 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. When a share is made available, the client operating system and [ACL manager](#) (page 148) is used to fine-tune the permissions of the files and directories that are created by the client.

[Sharing](#) (page 177) 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 24) or imported from a directory service using the instructions in [Directory Services](#) (page 165). The drop-down menus described in this section are automatically truncated to 50 entries for performance reasons. To find an unlisted entry, begin typing the desired user or group name for the drop-down menu to show matching results.

To set the permissions on a dataset, select it in *Storage* → *Pools*, click  (Options), then *Edit Permissions*. [Table 8.5](#) describes the options in this screen.

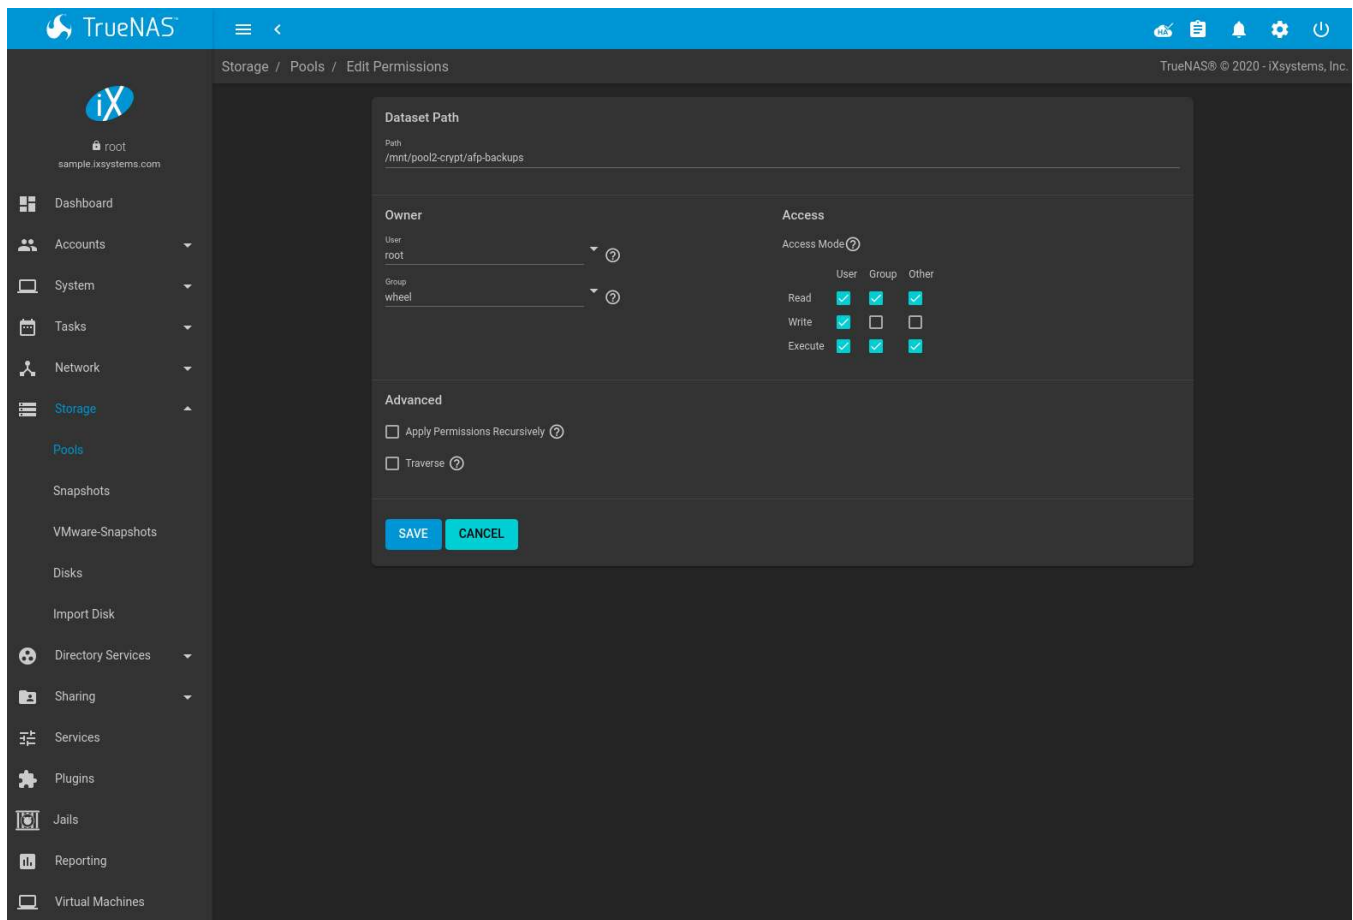


Fig. 8.10: Editing Dataset Permissions

Table 8.5: Permission Options

Setting	Value	Description
Path	string	Displays the path to the dataset or zvol directory.
User	drop-down menu	Select the user to control the dataset. Users created manually or imported from a directory service appear in the drop-down menu.
Group	drop-down menu	Select the group to control the dataset. Groups created manually or imported from a directory service appear in the drop-down menu.
Access Mode	checkboxes	Set the read, write, and execute permissions for the dataset.
Apply Permissions Recursively	checkbox	Apply permissions recursively to all directories and files within the current dataset.
Traverse	checkbox	Movement permission for this dataset. Allows users to view or interact with child datasets even when those users do not have permission to view or manage the contents of this dataset.

8.2.13 ACL Management

An Access Control List (ACL) is a set of account permissions associated with a dataset and applied to directories or files within that dataset. These permissions control the actions users can perform on the dataset contents. ACLs are typically used to manage user interactions with [shared datasets](#) (page 177). Datasets with an ACL have (ACL) appended to their name in the directory browser.

The ACL for a new file or directory is typically determined by the parent directory ACL. An exception is when there are no *File Inherit* or *Directory Inherit* *flags* (page 151) in the parent ACL `owner@`, `group@`, or `everyone@` entries. These non-inheriting entries are appended to the ACL of the newly created file or directory based on the [Samba create and directory masks](https://www.samba.org/samba/docs/using_samba/ch08.html) (https://www.samba.org/samba/docs/using_samba/ch08.html) or the `umask` (https://www.freebsd.org/cgi/man.cgi?query=umask&sektion=2) value.

By default, a file ACL is preserved when it is moved or renamed within the same dataset. The *SMB winmsa module* (page 214) can override this behavior to force an ACL to be recalculated whenever the file moves, even within the same dataset.

Datasets optimized for SMB sharing can restrict ACL changes. See *ACL Mode* in the *Dataset Options table* (page 143).

ACLs are modified by adding or removing Access Control Entries (ACEs) in *Storage* → *Pools*. Find the desired dataset, click `:` (Options), and select *Edit ACL*. The *ACL Manager* opens. The ACL manager must be used to modify permissions on a dataset with an ACL.

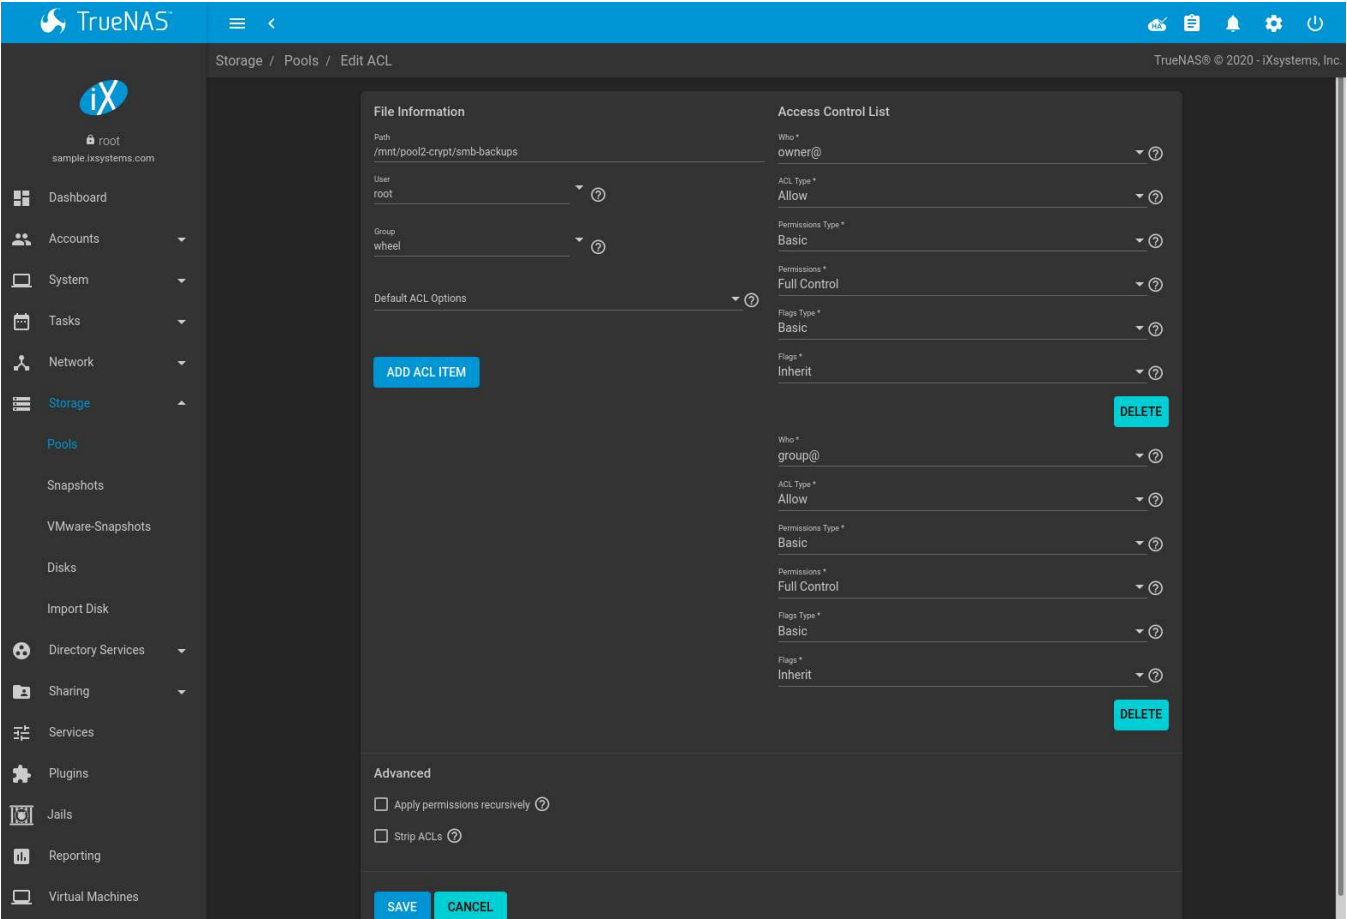


Fig. 8.11: ACL Manager

The ACL Manager options are split into the *File Information*, *Access Control List*, and *Advanced* sections. [Table 8.6](#) sorts these options by their section.

Table 8.6: ACL Options

Setting	Section	Value	Description
Path	File Information	string	Location of the dataset that is being modified. Read-only.

Continued on next page

Table 8.6 – continued from previous page

Setting	Section	Value	Description
User	File Information	drop-down menu	User who controls the dataset. This user always has permissions to read or write the ACL and read or write attributes. Users created manually or imported from a directory service (page 165) appear in the drop-down menu.
Apply User	File Information	checkbox	Confirm changes to User. To prevent errors, changes to the User are submitted only when this box is set.
Group	File Information	drop-down menu	The group which controls the dataset. This group has all permissions that are granted to the <i>@group Tag</i> . Groups created manually or imported from a directory service (page 165) appear in the drop-down menu.
Apply Group	File Information	checkbox	Confirm changes to Group. To prevent errors, changes to the Group are submitted only when this box is set.
Default ACL Options	File Information	drop-down menu	Default ACLs. Choosing an entry loads a preset ACL that is configured to match general permissions situations.
Who	Access Control List	drop-down menu	Access Control Entry (ACE) user or group. Select a specific <i>User</i> or <i>Group</i> for this entry, <i>owner@</i> to apply this entry to the selected <i>User</i> , <i>group@</i> to apply this entry to the selected <i>Group</i> , or <i>everyone@</i> to apply this entry to all users and groups. See setfacl(1) NFSv4 ACL ENTRIES (https://www.freebsd.org/cgi/man.cgi?query=setfacl).
User	Access Control List	drop-down menu	User account to which this ACL entry applies. Only visible when <i>User</i> is the chosen <i>Tag</i> .
Group	Access Control List	drop-down menu	Group to which this ACL entry applies. Only visible when <i>Group</i> is the chosen <i>Tag</i> .
ACL Type	Access Control List	drop-down menu	How the <i>Permissions</i> are applied to the chosen <i>Who</i> . Choose <i>Allow</i> to grant the specified permissions and <i>Deny</i> to restrict the specified permissions.
Permissions Type	Access Control List	drop-down menu	Choose the type of permissions. <i>Basic</i> shows general permissions. <i>Advanced</i> shows each specific type of permission for finer control.
Permissions	Access Control List	drop-down menu	Select permissions to apply to the chosen <i>Tag</i> . Choices change depending on the <i>Permissions Type</i> . See the permissions list (page 151) for descriptions of each permission.
Flags Type	Access Control List	drop-down menu	Select the set of ACE inheritance <i>Flags</i> to display. <i>Basic</i> shows un-specific inheritance options. <i>Advanced</i> shows specific inheritance settings for finer control.
Flags	Access Control List	drop-down menu	How this ACE is applied to newly created directories and files within the dataset. <i>Basic</i> flags enable or disable ACE inheritance. <i>Advanced</i> flags allow further control of how the ACE is applied to files and directories in the dataset. See the inheritance flags list (page 151) for descriptions of <i>Advanced</i> inheritance flags.
Apply permissions recursively	Advanced	checkbox	Apply permissions recursively to all directories and files in the current dataset.
Apply permissions to child datasets	Advanced	checkbox	Apply permissions recursively to all child datasets of the current dataset. Only visible when <i>Apply permissions recursively</i> is set.

Continued on next page

Table 8.6 – continued from previous page

Setting	Section	Value	Description
Strip ACLs	Advanced	checkbox	Set to remove all ACLs from the current dataset. ACLs are also recursively stripped from directories and child datasets when <i>Apply permissions recursively</i> and <i>Apply permissions to child datasets</i> are set.

Additional ACEs are created by clicking *ADD ACL ITEM* and configuring the added fields. One ACE is required in the ACL.

See [setfacl\(1\)](https://www.freebsd.org/cgi/man.cgi?query=setfacl) (<https://www.freebsd.org/cgi/man.cgi?query=setfacl>), [nfs4_acl\(5\)](https://linux.die.net/man/5/nfs4_acl) (https://linux.die.net/man/5/nfs4_acl), and [NFS Version 4 ACLs memo](https://tools.ietf.org/html/draft-falkner-nfsv4-acls-00) (<https://tools.ietf.org/html/draft-falkner-nfsv4-acls-00>) for more details about Access Control Lists, permissions, and inheritance flags. The following lists show each permission or flag that can be applied to an ACE with a brief description. An ACE can have a variety of basic or advanced permissions:

Basic Permissions

- *Read* : view file or directory contents, attributes, named attributes, and ACL. Includes the *Traverse* permission.
- *Modify* : adjust file or directory contents, attributes, and named attributes. Create new files or subdirectories. Includes the *Traverse* permission. Changing the ACL contents or owner is not allowed.
- *Traverse* : Execute a file or move through a directory. Directory contents are restricted from view unless the *Read* permission is also applied. To traverse and view files in a directory, but not be able to open individual files, set the *Traverse* and *Read* permissions, then add the advanced *Directory Inherit* flag.
- *Full Control* : Apply all permissions.

Advanced Permissions

- *Read Data* : View file contents or list directory contents.
- *Write Data* : Create new files or modify any part of a file.
- *Append Data* : Add new data to the end of a file.
- *Read Named Attributes* : view the named attributes directory.
- *Write Named Attributes* : create a named attribute directory. Must be paired with the *Read Named Attributes* permission.
- *Execute* : Execute a file, move through, or search a directory.
- *Delete Children* : delete files or subdirectories from inside a directory.
- *Read Attributes* : view file or directory non-ACL attributes.
- *Write Attributes* : change file or directory non-ACL attributes.
- *Delete* : remove the file or directory.
- *Read ACL* : view the ACL.
- *Write ACL* : change the ACL and the ACL mode.
- *Write Owner* : change the user and group owners of the file or directory.
- *Synchronize* : synchronous file read/write with the server. This permission does not apply to FreeBSD clients.

Basic inheritance flags only enable or disable ACE inheritance. Advanced flags offer finer control for applying an ACE to new files or directories.

- *File Inherit* : The ACE is inherited with subdirectories and files. It applies to new files.
- *Directory Inherit* : new subdirectories inherit the full ACE.
- *No Propagate Inherit* : The ACE can only be inherited once.
- *Inherit Only* : Remove the ACE from permission checks but allow it to be inherited by new files or subdirectories. *Inherit Only* is removed from these new objects.

- *Inherited* : set when the ACE has been inherited from another dataset.

8.3 Snapshots

To view and manage the listing of created snapshots, use *Storage* → *Snapshots*. An example is shown in [Figure 8.12](#).

Note: If snapshots do not appear, check that the current time configured in *Periodic Snapshot Tasks* (page 97) 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 302).

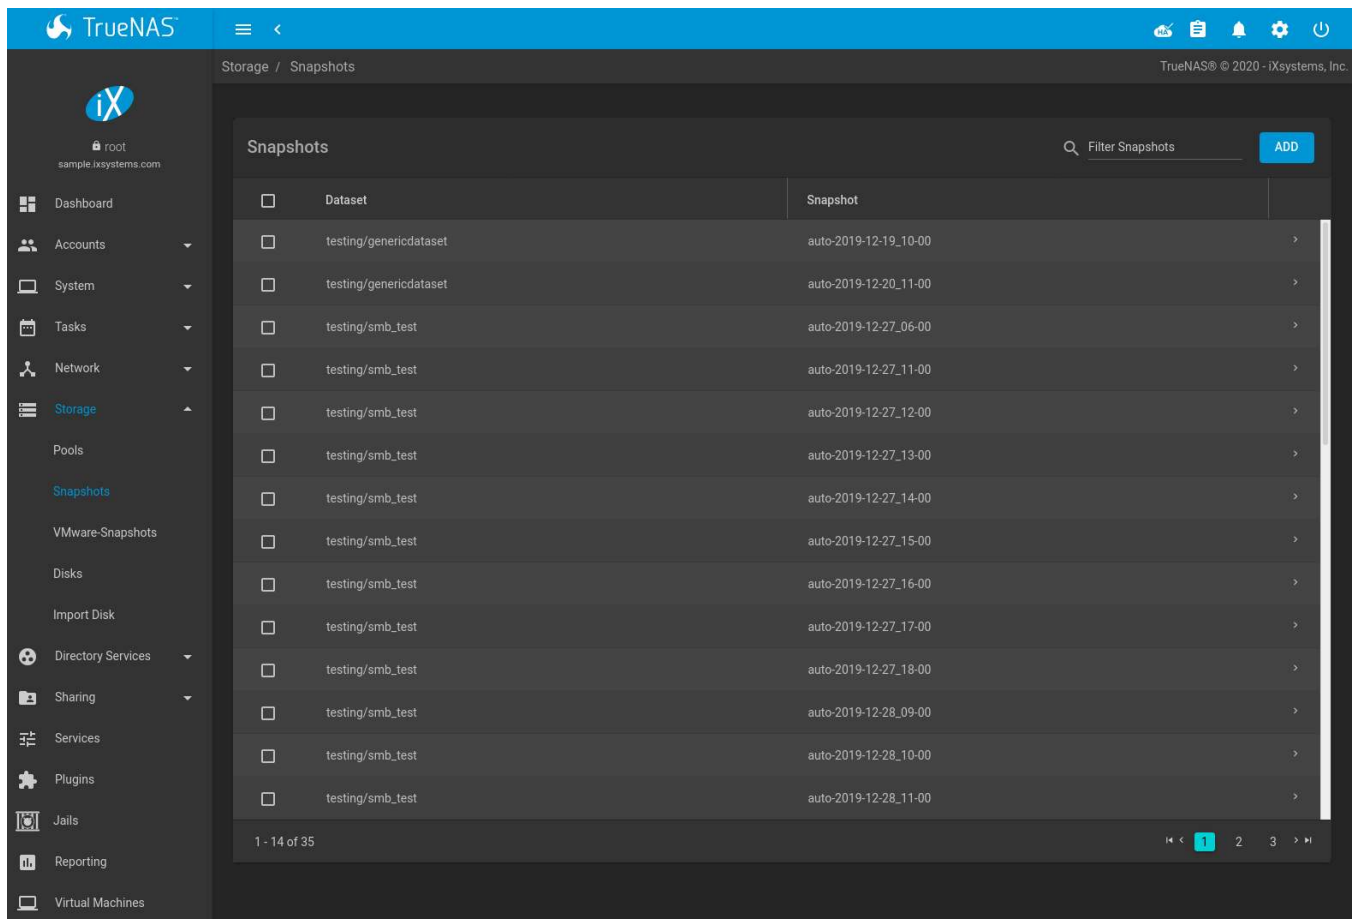


Fig. 8.12: Viewing Available Snapshots

Each entry in the list includes the name of the dataset and snapshot. Click > (Expand) to view these options:

DATE CREATED shows the exact time and date of the snapshot creation.

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 302).

REFERENCED 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.

DELETE shows a confirmation dialog. 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 TO NEW DATASET prompts for the name of the new dataset created from the cloned snapshot. A default name is provided based on the name of the original snapshot. 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: Clicking `:` (Options) → *Rollback* asks for confirmation before rolling back to the chosen snapshot state. Clicking *Yes* causes all files in the dataset to revert to the state they were in when the snapshot was created.

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 TrueNAS® 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 218). Users can access the files in the shadow copy using Explorer without requiring any interaction with the TrueNAS® 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.

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 10).

8.3.1 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 auxillary 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.

8.3.2 Creating a Single Snapshot

To create a snapshot separately from a *periodic snapshot schedule* (page 97), go to *Storage* → *Snapshots* and click *ADD*.

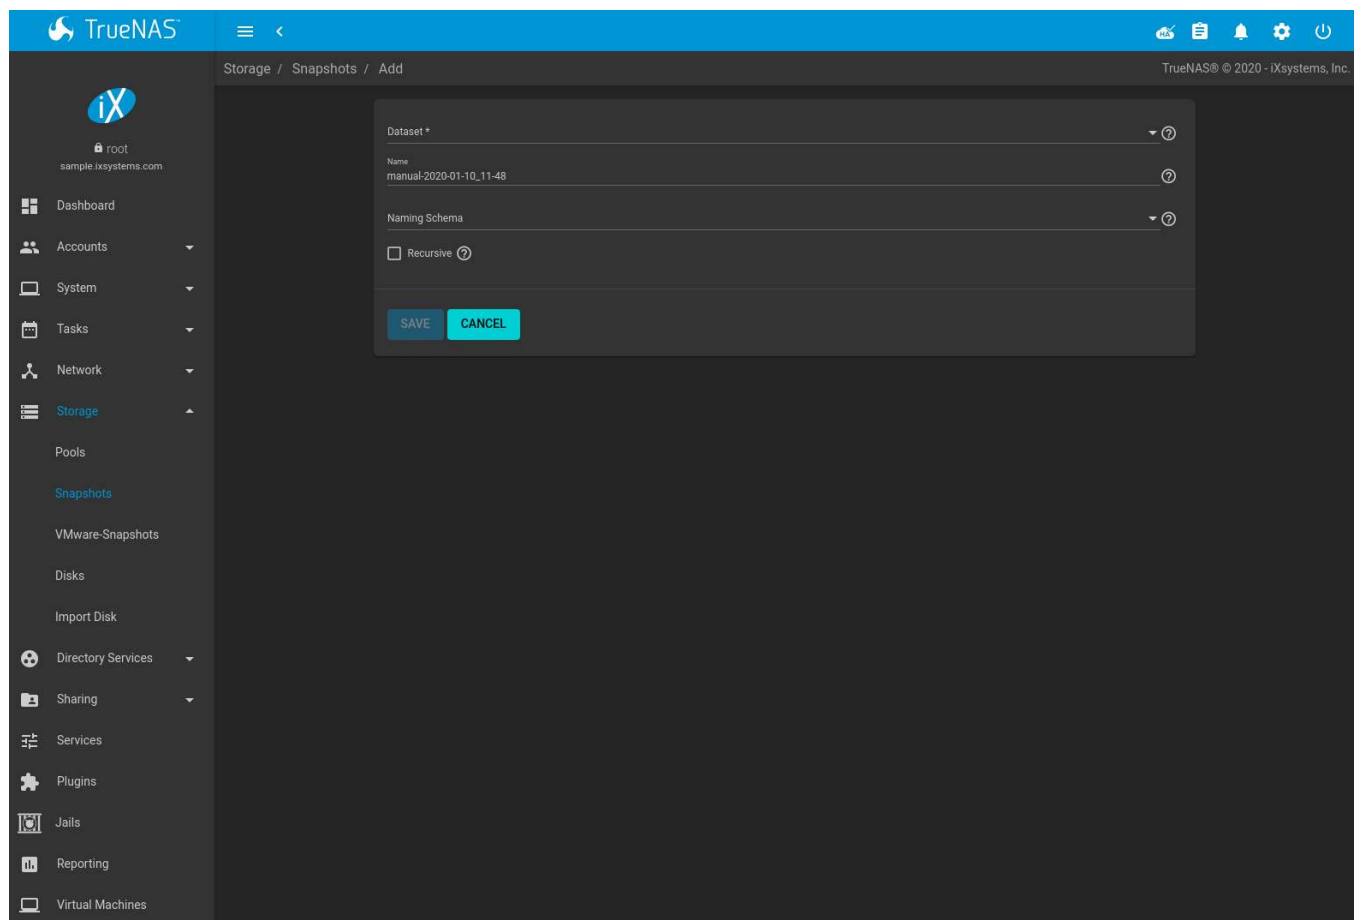


Fig. 8.13: Single Snapshot Options

Select an existing ZFS pool, dataset, or zvol to snapshot. To include child datasets with the snapshot, set *Recursive*.

The snapshot can have a custom *Name* or be automatically named by a *Naming Schema*. Using a *Naming Schema* allows the snapshot to be included in [Replication Tasks](#) (page 107). The *Naming Schema* drop-down is populated with previously created schemas from [Periodic Snapshot Tasks](#) (page 97).

8.4 VMware-Snapshots

Storage → *VMware-Snapshots* is used to coordinate ZFS snapshots when using TrueNAS® as a VMware datastore. When a ZFS snapshot is created, TrueNAS® automatically snapshots any running VMware virtual machines before taking a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore. Virtual machines **must be powered on** for TrueNAS® snapshots to be copied to VMware. 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 are listed in [Snapshots](#) (page 152).

Figure 8.14 shows the menu for adding a VMware snapshot and Table 8.7 summarizes the available options.

The screenshot shows the TrueNAS web interface. The left sidebar contains a navigation menu with options like Dashboard, Accounts, System, Tasks, Network, Storage, Pools, Snapshots, VMware-Snapshots (highlighted), Disks, Import Disk, Directory Services, Sharing, Services, Plugins, Jails, Reporting, and Virtual Machines. The main content area is titled 'Storage / VMware Snapshots / Add'. It features a form with the following fields: 'Hostname *', 'Username *', 'Password *' (with a toggle for visibility), 'ZFS Filesystem *', and 'Datastore *'. At the bottom of the form are three buttons: 'SAVE', 'CANCEL', and 'FETCH DATASTORES'. The top of the interface shows the TrueNAS logo and a user profile for 'root'.

Fig. 8.14: Adding a VMware Snapshot

Table 8.7: VMware Snapshot Options

Setting	Value	Description
Hostname	string	Enter the IP address or hostname of the VMware host. When clustering, use the IP address or hostname of the vCenter server for the cluster.
Username	string	Enter a user account name created on the VMware host. The account must have permission to snapshot virtual machines.

Continued on next page

Table 8.7 – continued from previous page

Setting	Value	Description
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.

TrueNAS® connects to the VMware host after the credentials are entered. The *ZFS Filesystem* and *Datastore* drop-down menus are populated with information from the VMware host. Choosing a datastore also selects any previously mapped dataset.

8.5 Disks

To view all of the disks recognized by the TrueNAS® system, use *Storage* → *Disks*. As seen in the example in Figure 8.15, each disk entry displays its device name, serial number, size, advanced power management settings, acoustic level settings, and whether *S.M.A.R.T.* (page 238) tests are enabled. The pool associated with the disk is displayed in the *Pool* column. *Unused* is displayed if the disk is not being used in a pool. Click *COLUMNS* and select additional information to be shown as columns in the table. Additional information not shown in the table can be seen by clicking > (Expand).

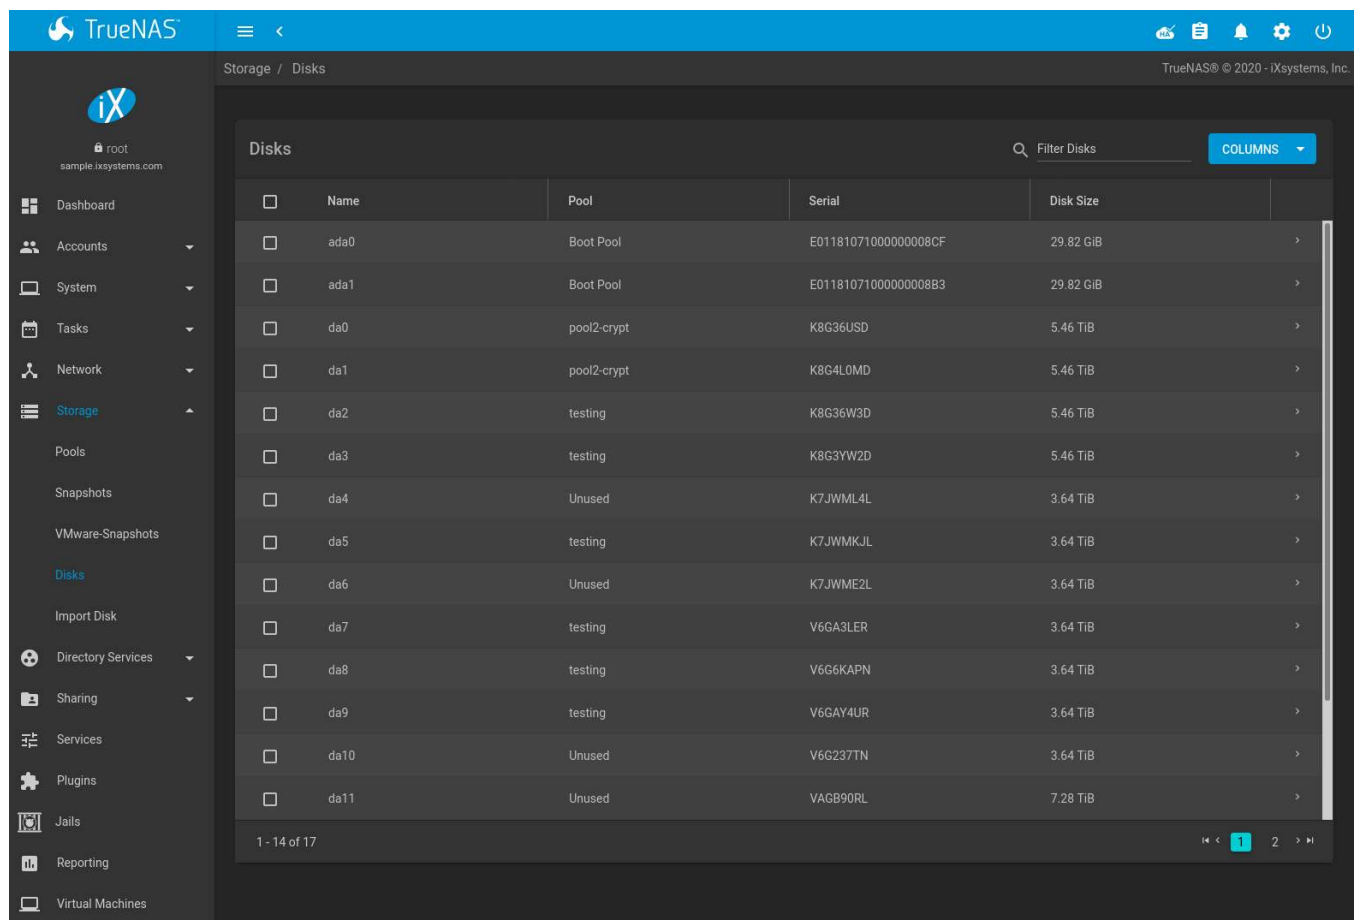



Fig. 8.15: Viewing Disks

To edit the options for a disk, click ⋮ (Options) on a disk, then *Edit* to open the screen shown in Figure 8.16. Table 8.8 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 157) indicates the options available when editing multiple disks.

To offline, online, or replace the device, see [Replacing a Failed Disk](#) (page 158).

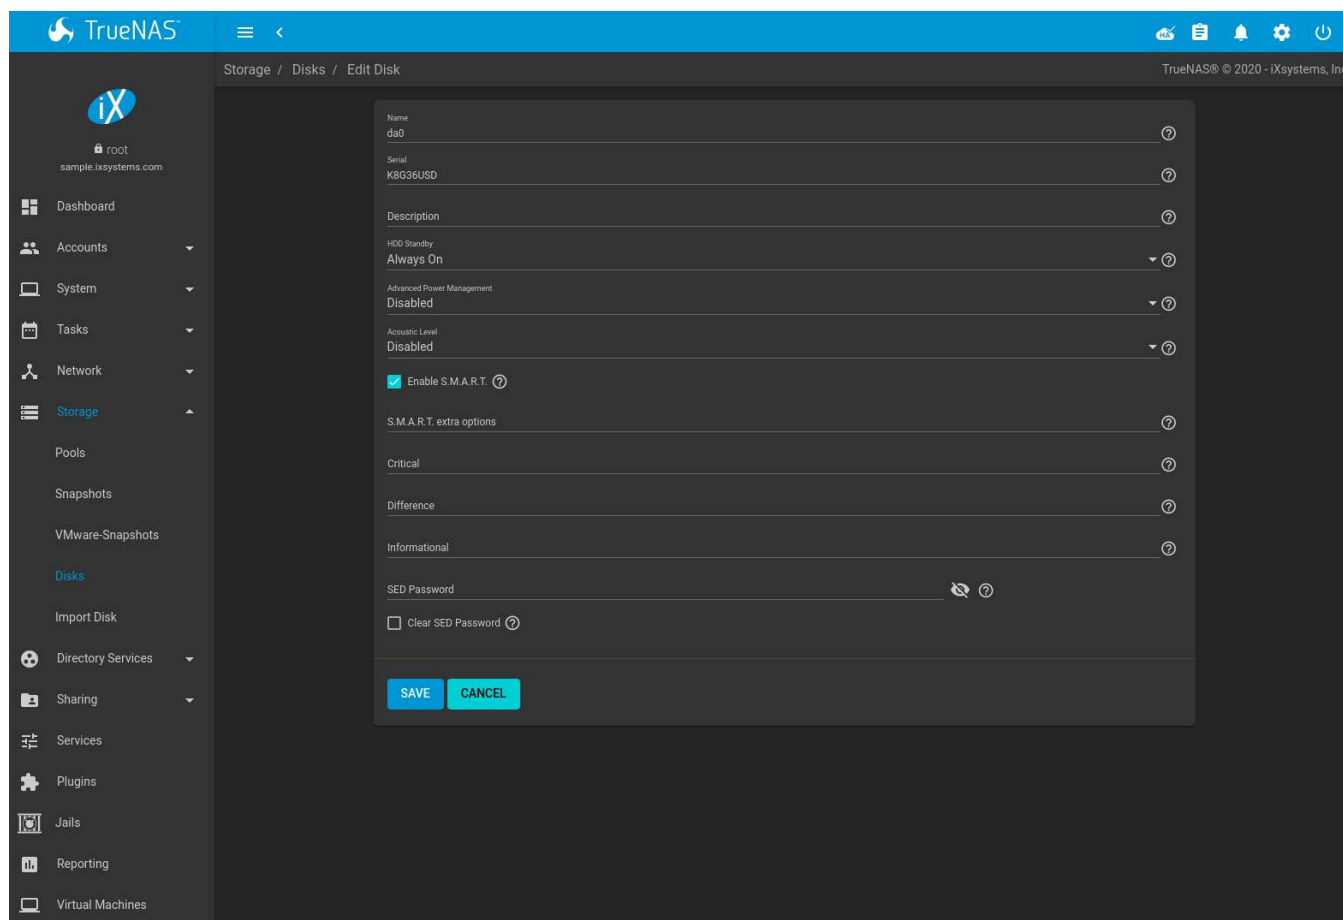


Fig. 8.16: Editing a Disk

Table 8.8: 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	✓	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/) shows how to determine if a drive has spun down. Temperature monitoring is disabled if the disk is set to enter standby.
Advanced Power Management	drop-down menu	✓	Select a power management profile from the menu. The default value is <i>Disabled</i> .

Continued on next page

Table 8.8 – continued from previous page

Setting	Value	Bulk Edit	Description
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/Automatic_acoustic_management).
Enable S.M.A.R.T.	checkbox	✓	Enabled by default when the disk supports S.M.A.R.T. Disabling S.M.A.R.T. tests prevents collecting new temperature data for this disk. Historical temperature data is still displayed in Reporting (page 285).
S.M.A.R.T. extra options	string	✓	Enter additional smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) options.
Critical	string		Threshold temperature in Celsius. If the drive temperature is higher than this value, a <code>LOG_CRIT</code> level log entry is created and an email is sent. 0 disables this check.
Difference	string		Report if the temperature of a drive has changed by this many degrees Celsius since the last report. 0 disables the report.
Informational	string		Report if drive temperature is at or above this temperature in Celsius. 0 disables the report.
SED Password	string		Set or change the password of this SED. This password is used instead of the global SED password in <i>System</i> → <i>Advanced</i> . See Self-Encrypting Drives (page 41).
Clear SED Password	checkbox		Clear the SED password for this disk.

Tip: If the serial number for a disk is not displayed in this screen, use the `smartctl` command from [Shell](#) (page 302). 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 TrueNAS® system, and wipe the disk in a separate computer.

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

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

8.5.1 Replacing a Failed Disk

Replace failed drives as soon as possible to repair the degraded state of the RAID.

Striping (RAID0) does not provide redundancy. Disk failure in a stripe results in losing the pool. The pool must be recreated and data stored in the failed stripe will have to be restored from backups.

Warning: Encrypted pools must have a valid passphrase to replace a failed disk. Set a passphrase and back up the encryption key using the pool [Encryption Operations](#) (page 135) **before** attempting to replace the failed drive.

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. *Warning:* encrypted disks that are set *OFFLINE* cannot be set back *ONLINE*. Click *Offline* and pull the disk. If there is no *Offline* but only *Replace*, 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. 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. Encrypted pools require entering the [encryption key passphrase](#) (page 135) when choosing a replacement disk. Clicking *REPLACE DISK* begins the process to reformat the replacement, apply the current pool encryption algorithm, and resilver the pool. The current pool encryption key and passphrase remains valid, but any pool recovery key file is invalidated by the replacement process. To maximize pool security, it is recommended to [reset pool encryption](#) (page 136).
3. After the drive replacement process is complete, re-add the replaced disk in the [S.M.A.R.T. Tests](#) (page 96) screen.

To refresh the screen with updated entries, click *REFRESH*. If any problems occur during a disk replacement process, one of the disks can be detached. To detach a disk in the replacement process, find the disk to be replaced and click ⋮ (Options) → *Detach*.

[Figure 8.17](#) shows an example of going to *Storage* → *Pools* → *Status* and replacing a disk in an active pool.

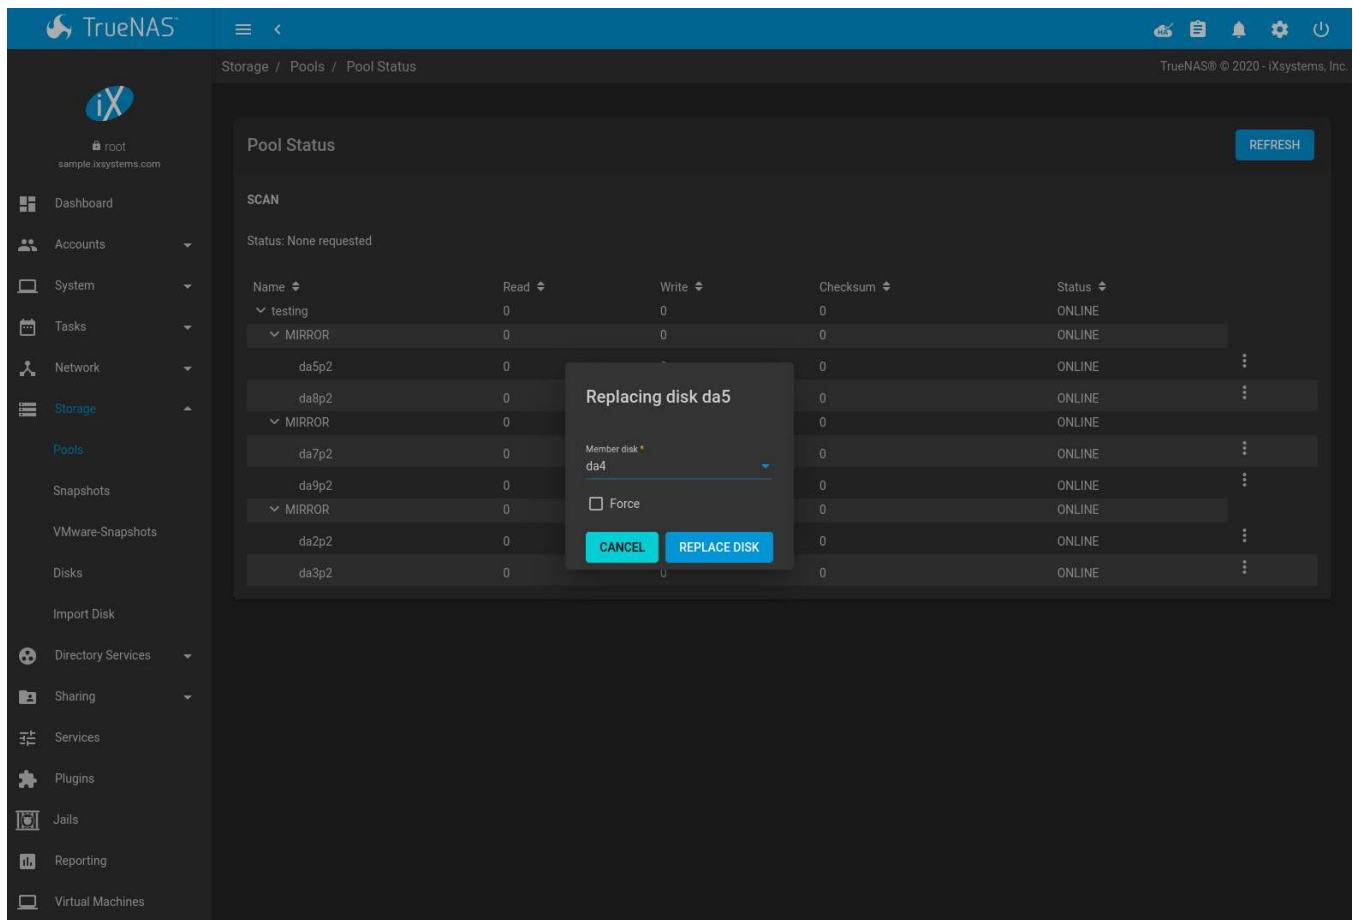


Fig. 8.17: Replacing a Failed Disk

After the resilver is complete, the pool status shows a *Completed* resilver status and indicates any errors. [Figure 8.18](#) 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.

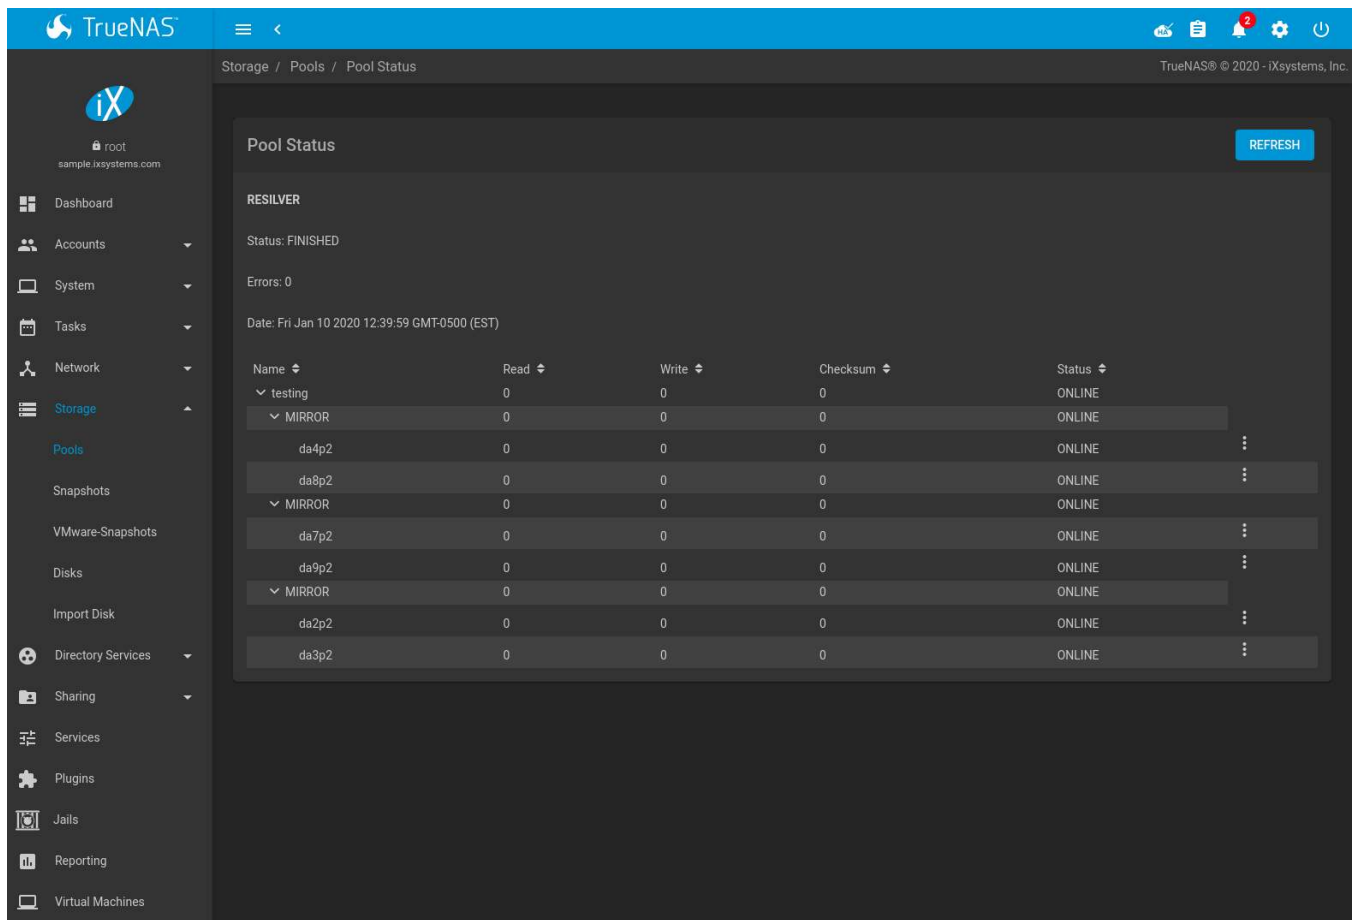


Fig. 8.18: Disk Replacement is Complete

8.5.1.1 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.

8.5.2 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 from *Pools* (page 130) as additional capacity is needed.

But adding vdevs is not an option if there are not enough unused disk ports. If there is at least one unused disk port or drive bay, a single disk at a time can be replaced with a larger disk, waiting for the resilvering process to include the new disk into the pool, removing the old disk, then repeating with another disk until all of the original disks have been replaced. At that point, the pool capacity automatically increases to include the new space.

One advantage of this method is that disk redundancy is present during the process.

Note: A pool that is configured as a *stripe* (https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_0) can only be increased by following the steps in *Extending a Pool* (page 137).

1. Connect the new, larger disk to the unused disk port or drive bay.

-
2. Go to *Storage* → *Pools*.
 3. Select the pool and click ⚙ (Settings) → *Status*.
 4. Select one of the old, smaller disks in the pool. Click ⋮ (Options) → *Replace*. Choose the new disk as the replacement.

The status of the resilver process is shown on the screen, or can be viewed with `zpool status`. When the new disk has resilvered, the old one is automatically offlined. It can then be removed from the system, and that port or bay used to hold the next new disk.

If a unused disk port or bay is not available, a drive can be replaced with a larger one as shown in [Replacing a Failed Disk](#) (page 158). 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 appears in the pool.

8.6 Importing a Disk

The *Storage* → *Import Disk* screen, shown in [Figure 8.19](#), is used to import disks that are formatted with UFS (BSD Unix), FAT(MSDOS) or NTFS (Windows), or EXT2 (Linux) filesystems. This is a designed to be used as a one-time import, copying the data from that disk into a dataset on the TrueNAS® system. Only one disk can be imported at a time.

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




Fig. 8.19: Importing a Disk

Use the drop-down menu to select the disk to import, confirm the detected filesystem is correct, 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 is displayed. 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.

After importing a disk, a dialog allows viewing or downloading the disk import log.

8.7 Multipaths

This option is only displayed on systems that contain multipath-capable hardware like a chassis equipped with a dual SAS expander backplane or an external JBOD that is wired for multipath.

TrueNAS® uses [gmultipath\(8\)](https://www.freebsd.org/cgi/man.cgi?query=gmultipath) (<https://www.freebsd.org/cgi/man.cgi?query=gmultipath>) to provide [multipath I/O](https://en.wikipedia.org/wiki/Multipath_I/O) (https://en.wikipedia.org/wiki/Multipath_I/O) support on systems containing multipath-capable hardware.

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

TrueNAS® automatically detects active/active and active/passive multipath-capable hardware. Discovered multipath-capable devices are placed in multipath units with the parent devices hidden. The configuration is displayed in *Storage* → *Multipaths*.

OVERPROVISIONING

Overprovisioning SSDs can be done using the `disk_resize` command in the *Shell* (page 302). This can be useful for many different scenarios. Perhaps the most useful benefit of overprovisioning is that it can extend the life of an SSD greatly. Overprovisioning an SSD distributes the total number of writes and erases across more flash blocks on the drive. Read more about overprovisioning SSDs [here](https://www.seagate.com/tech-insights/ssd-over-provisioning-benefits-master-ti/) (<https://www.seagate.com/tech-insights/ssd-over-provisioning-benefits-master-ti/>).

The command to overprovision an SSD is `disk_resize device size`, where *device* is the device name of the SSD and *size* is the desired size of the provision in *GB* or *TB*. Here is an example of the command: `disk_resize ada5 16GB`. When no size is specified, it reverts the provision back the full size of the device.

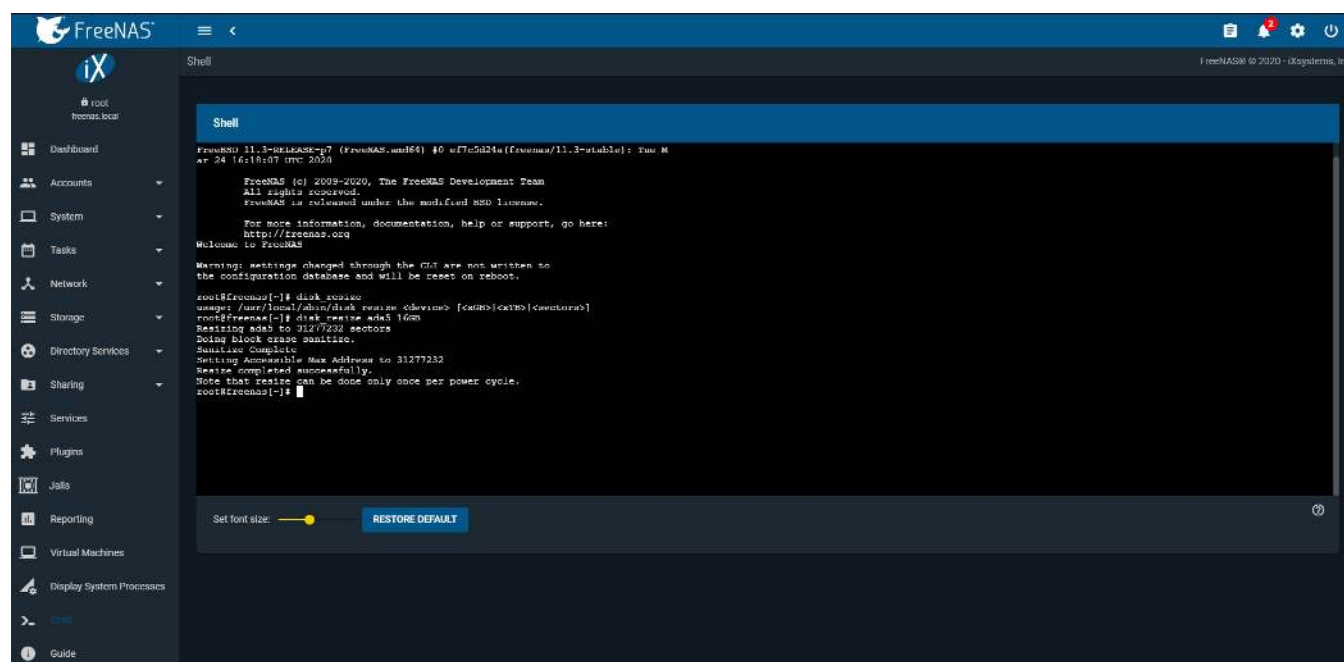


Fig. 9.1: `disk_resize` Command



Note: Some SATA devices may be limited to one resize per power cycle. Some BIOS may block resize during boot and require a live power cycle.

DIRECTORY SERVICES

TrueNAS® supports integration with these directory services:

- [Active Directory](#) (page 165) (for Windows 2000 and higher networks)
- [LDAP](#) (page 170)
- [NIS](#) (page 173)

TrueNAS® also supports [Kerberos Realms](#) (page 174), [Kerberos Keytabs](#) (page 175), and the ability to add more parameters to [Kerberos Settings](#) (page 176).

This section summarizes each of these services and the available configuration options within the TrueNAS® web interface. After successfully enabling a directory service,  appears in the top toolbar row. Click  to show the *Directory Services Monitor* menu. This menu shows the name and status of each directory service.

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_Directory_Domain_Controller)

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 TrueNAS® system. Instead, configure the Active Directory service so account information and imported users can be authorized to access the SMB shares on the TrueNAS® system.

Many changes and improvements have been made to Active Directory support within TrueNAS®. It is strongly recommended to update the system to the latest TrueNAS® 11.3 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 302) on the TrueNAS® system. If the `ping` fails, check the DNS server and default gateway settings in *Network* → *Global Configuration* on the TrueNAS® system.

By default, *Allow DNS updates* in the [Active Directory options](#) (page 166) is enabled. This adds TrueNAS® *SMB 'Bind IP Addresses'* (page 240) DNS records to the Active Directory DNS when the domain is joined. Disabling *Allow DNS updates* means that the Active Directory DNS records must be updated manually.

Active Directory relies on Kerberos, a time-sensitive protocol. During the domain join process the [PDC emulator FSMO role](#) (https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-adts/f96ff8ec-c660-4d6c-924f-c0dbbcac1527) server is added as the preferred NTP server. The time on the TrueNAS® system and the Active Directory Domain Controller cannot be out of sync by more than five minutes in a default Active Directory environment. An [Alert](#) (page 305) is sent when the time is out of sync.

To ensure both systems are set to the same time:

- use the same NTP server (set in *System* → *NTP Servers* on the TrueNAS® 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. Hidden when a <i>Kerberos Principal</i> is selected.
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. Hidden when a <i>Kerberos Principal</i> is selected.
Domain Account Password	string		Password for the Active Directory administrator account. Required the first time a domain is configured. After initial configuration, the password is not needed to edit, start, or stop the service.

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Encryption Mode	drop-down	✓	Choices are <i>Off</i> , <i>SSL (LDAPS protocol port 636)</i> , or <i>TLS (LDAP protocol port 389)</i> . See http://info.ssl.com/article.aspx?id=10241 and https://hpbnc.co/transport-layer-security-tls/ for more information about SSL and TLS.
Certificate	drop-down menu	✓	Select the Active Directory server certificate if SSL connections are used. If a certificate does not exist, create or import a <i>Certificate Authority</i> (page 72), then create a certificate on the Active Directory server. Import the certificate to the TrueNAS® system using the <i>Certificates</i> (page 76) menu. It is recommended to leave this drop-down unset when configuring LDAPs. To clear a saved certificate, choose the blank entry and click <i>SAVE</i> .
Validate Certificate	checkbox	✓	Check server certificates in a TLS session.
Verbose logging	checkbox	✓	Set to log attempts to join the domain to <code>/var/log/messages</code> .
Allow Trusted Domains	checkbox	✓	Do not set this unless the network has active <i>domain/forest trusts</i> (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc757352(v=ws.10)) and managing files on multiple domains is required. Setting this option generates more winbindd traffic and slows down filtering with user and group information. If enabled, also configuring the idmap ranges and a backend for each trusted domain in the environment is recommended.
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	✓	Disable caching AD users and groups. Setting this hides all AD users and groups from web interface drop-down menus and auto-completion suggestions, but manually entering names is still allowed. This can help when unable to bind to a domain with a large number of users or groups.
Site Name	string	✓	Auto-detected site name. Do not change this unless the detected site name is incorrect for the particular AD environment.
Kerberos Realm	drop-down menu	✓	Select the realm created using the instructions in <i>Kerberos Realms</i> (page 174).
Kerberos Principal	drop-down menu	✓	Select a keytab created using the instructions in <i>Kerberos Keytabs</i> (page 175). Selecting a principal hides the <i>Domain Account Name</i> and <i>Domain Account Password</i> fields. An existing account name is not overwritten by the principal.
Computer Account OU	string	✓	The OU in which new computer accounts are created. The OU string is read from top to bottom without RDNs. Slashes (/) are used as delimiters, like <i>Computers/Servers/NAS</i> . The backslash (\) is used to escape characters but not as a separator. Backslashes are interpreted at multiple levels and might require doubling or even quadrupling to take effect. When this field is blank, new computer accounts are created in the Active Directory default OU.

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Advanced Mode	Description
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.
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. This should be set to <i>PLAIN</i> when using Microsoft Active Directory. This can be set to <i>SIGN</i> or <i>SEAL</i> when using Samba Active Directory if <i>allow sasl over tls</i> has been explicitly enabled in the Samba Domain Controller configuration.
Enable (requires password or Kerberos principal)	checkbox		Activate the Active Directory service.
NetBIOS Name	string	✓	Name for the computer object generated in AD. Automatically populated with the active TrueNAS controller host-name from the Global Configuration (page 119). Limited to 15 characters. It must be different from the <i>Workgroup</i> name.
NetBIOS Name (TrueNAS Controller 1/2)	string	✓	Name for the computer object generated in AD. Automatically populated with the standby TrueNAS controller host-name from the Global Configuration (page 119). Limited to 15 characters. When using Failover (page 81), set a unique NetBIOS name for the standby TrueNAS controller.
NetBIOS Alias	string	✓	Limited to 15 characters. When using Failover (page 81), this is the NetBIOS name that resolves to either TrueNAS controller.

[Table 10.2](#) summarizes the backends which are available in the *Idmap backend* drop-down menu. Each backend has its own [man page](http://samba.org.ru/samba/docs/man/manpages/) (<http://samba.org.ru/samba/docs/man/manpages/>) that gives implementation details.

Changing idmap backends automatically refreshes the `windbind` resolver cache by sending `SIGHUP` (signal hang up) to the parent `windbindd` process. To find this parent process, start an [SSH](#) (page 244) session with the TrueNAS® system and enter `service samba_server status`. To manually send the `SIGHUP`, enter `kill -HUP pid`, where `pid` is the parent process ID.

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 <code>uidNumber</code> attributes for users and <code>gidNumber</code> attributes for groups in the AD.

Continued on next page

Table 10.2 – continued from previous page

Value	Description
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.
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	IDs for AD users stored as RFC2307 (https://tools.ietf.org/html/rfc2307) ldap schema extensions. This module can either look up the IDs in the AD LDAP servers or an external (non-AD) LDAP server.
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.

REBUILD DIRECTORY SERVICE CACHE immediately refreshes the web interface directory service cache. This occurs automatically once a day as a cron job.

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. Active Directory does not allow \$ characters in Domain or NetBIOS names. The length of those names is also limited to 15 characters. The Administrator account password cannot contain the \$ character.

It can take a few minutes after configuring the Active Directory service for the AD information to be populated to the TrueNAS® system. To check the AD join progress, open the web interface Task Manager in the upper-right corner. Any errors during the join process are also displayed in the Task Manager.

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 TrueNAS® system are shown by typing commands in the TrueNAS® *Shell* (page 302):

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

In addition, `wbinfo -m` shows the domains and `wbinfo -t` tests the connection. When successful, `wbinfo -t` 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, open the *Shell* (page 302) and enter `smbclient//127.0.0.1/SHARE -U DOMAIN\username`, where *SHARE* is the SMB share name, *DOMAIN* is the name of the trusted domain, and *username* is the user account for authentication testing.

`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 TrueNAS® system taking longer than the default ten seconds to join Active Directory. Increase the value of *AD timeout* to 60 seconds.

10.1.1 Leaving the Domain

A *Leave Domain* button appears on the service dialog when a domain is connected. To leave the domain, click the button and enter credentials with privileges sufficient to permit leaving.

10.1.2 Troubleshooting Tips

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 Technet 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 TrueNAS® 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.

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*.

If any of the commands fail or result in a traceback, create a bug report at <https://bugs.ixsystems.com>. Include the commands in the order in which they were run and the exact wording of the error message or traceback.

10.2 LDAP

TrueNAS® 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 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 TrueNAS® LDAP service so network users can authenticate to the LDAP server and have authorized access to the data stored on the TrueNAS® 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) (https://manuals.info.apple.com/MANUALS/0/MA954/en_US/Open_Directory_Admin_v10.5_3rd_Ed) is an LDAP-compatible directory service into which TrueNAS® 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*.

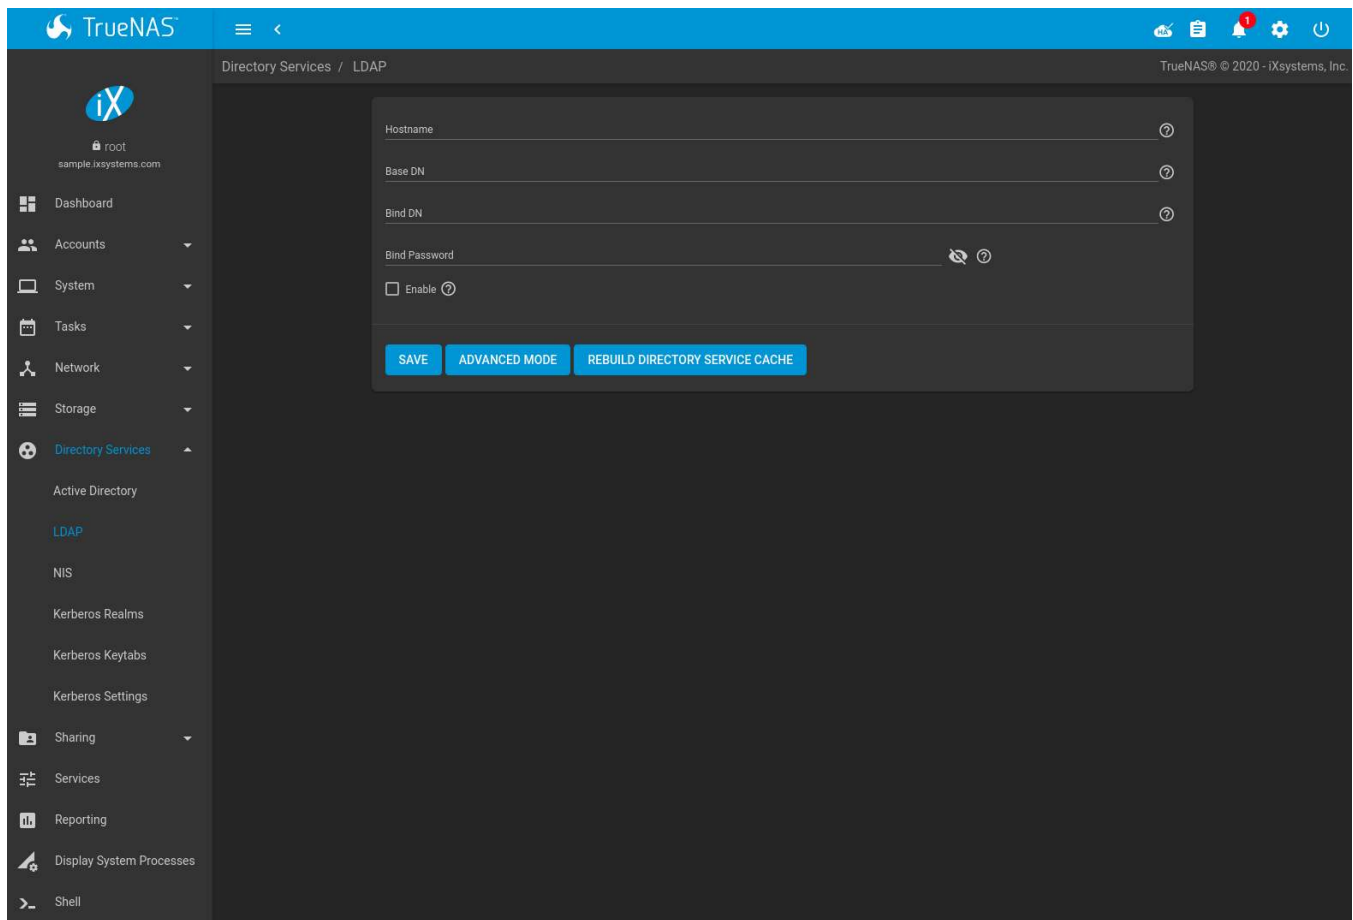


Fig. 10.2: Configuring LDAP

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		LDAP server hostnames or IP addresses. Separate entries with an empty space. Multiple hostnames or IP addresses can be entered to create an LDAP failover priority list. If a host does not respond, the next host in the list is tried until a new connection is established.
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.

Continued on next page

Table 10.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Kerberos Realm	drop-down menu	✓	The realm created using the instructions in Kerberos Realms (page 174).
Kerberos Principal	drop-down menu	✓	The location of the principal in the keytab created as described in Kerberos Keytabs (page 175).
Encryption Mode	drop-down menu	✓	Options for encrypting the LDAP connection: <ul style="list-style-type: none"> • <i>OFF</i>: do not encrypt the LDAP connection. • <i>ON</i>: encrypt the LDAP connection with SSL on port 636. • <i>START_TLS</i>: encrypt the LDAP connection with START-TLS on the default LDAP port 389.
Certificate	drop-down menu	✓	Certificate (page 76) to use when performing LDAP certificate-based authentication. To configure LDAP certificate-based authentication, create a Certificate Signing Request for the LDAP provider to sign. A certificate is not required when using username/password or Kerberos authentication.
Validate Certificate	checkbox	✓	Verify certificate authenticity.
Disable LDAP User/Group Cache	checkbox	✓	Disable caching LDAP users and groups in large LDAP environments. When caching is disabled, LDAP users and groups do not appear in dropdown menus, but are still accepted when manually entered.
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	✓	Backend used to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. See Table 10.2 for a summary of the available backends. To configure the selected backend, click <i>EDIT IDMAP</i> .
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 nslcd.conf (https://arthurdejong.org/nss-pam-ldapd/nslcd.conf.5).
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.

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 TrueNAS® *Shell* (page 302) to verify the users have been imported. Type `getent group` to verify the groups have been imported. When the *Samba Schema* is enabled, LDAP users also appear in the output of `pdbedit -L`.

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.

Any LDAP bind errors are displayed during the LDAP bind process. When troubleshooting LDAP, you can open the TrueNAS® *Shell* (page 302) and find `nslcd.conf` errors in `/var/log/messages`. When *Samba schema* is enabled, any Samba errors are recorded in `/var/log/samba4/log.smbd`. Additional details are saved in `/var/log/middlewared.log`.

To clear LDAP users and groups from TrueNAS®, 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, hostnames, email aliases, and other text-based tables of information. If an NIS server is running on the network, the TrueNAS® 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 TrueNAS®. 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.

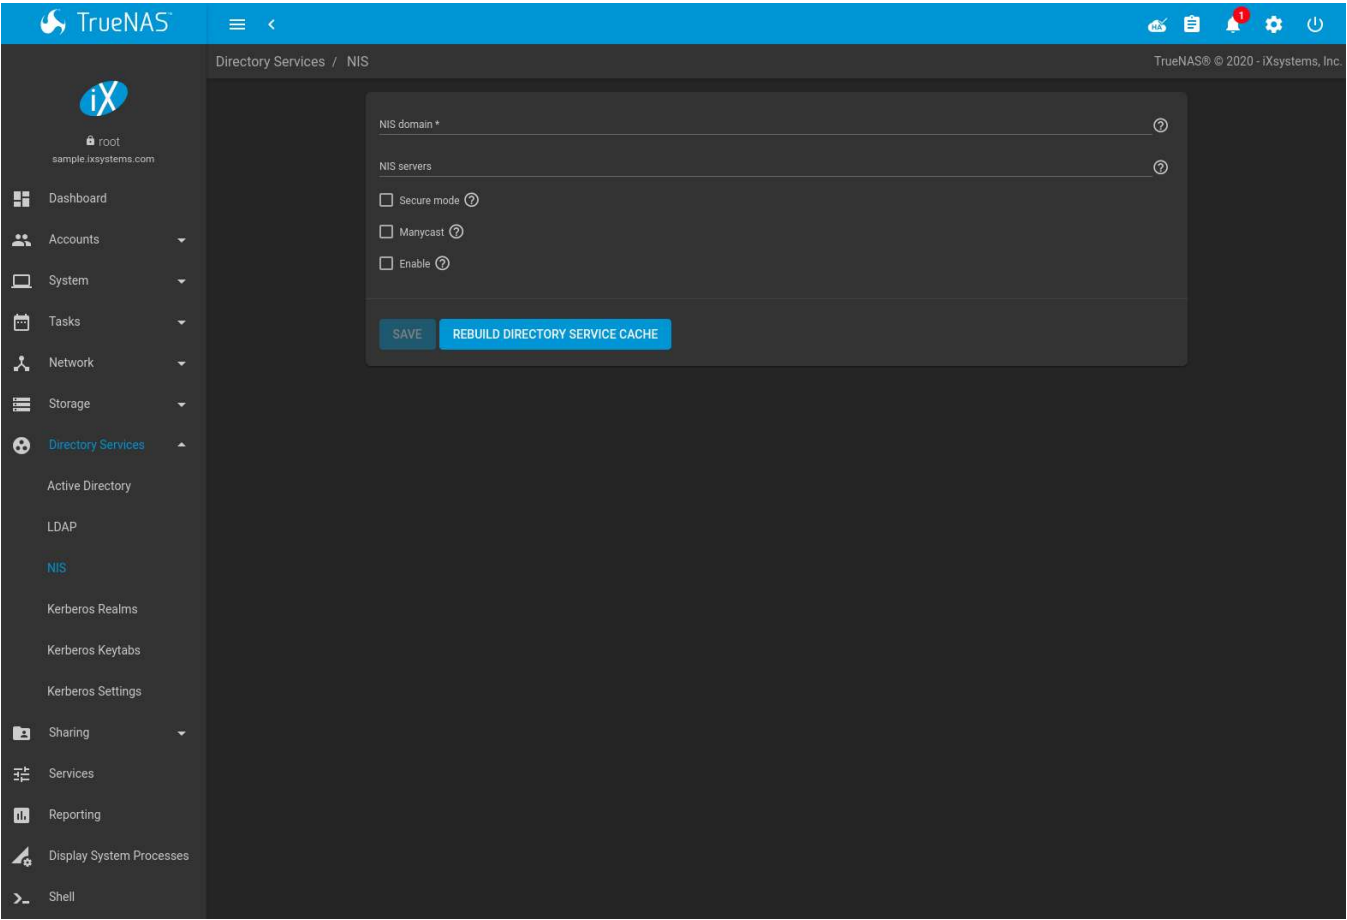


Fig. 10.3: NIS Configuration

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 sub-net.
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 TrueNAS®. *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 TrueNAS® 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 TrueNAS® 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 TrueNAS® server.
- `useraccount` is the name of the user account for the TrueNAS® 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 174) 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 TrueNAS® 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 provide and control access to an area of storage. Consider factors like operating system, security, transfer speed, and user access before creating a new share. This information can help determine the type of share, if multiple datasets are needed to divide the storage into areas with different access and permissions, and the complexity of setting up permissions.

Note that shares are only used to provide access to data. Deleting a share configuration does not affect the data that was being shared.

These types of shares and services are available:

- **AFP** (page 178): Apple Filing Protocol shares are used when the client computers all run macOS. Apple has deprecated AFP in favor of **SMB** (page 209). Using AFP in modern networks is no longer recommended.
- **Unix (NFS)** (page 201): 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 207): 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 209): 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 or Mac systems. However, it is a poor choice if the CPU on the TrueNAS® system is limited. If it is maxed out, upgrade the CPU or consider a different type of share.
- **Block (iSCSI)** (page 183): 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 226) 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 244) 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 222).

11.1 Apple (AFP) Shares

TrueNAS® 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 TrueNAS® 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 TrueNAS® 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.
Time Machine Quota	integer		Appears when <i>Time Machine</i> is set. Enter a storage quota for each Time Machine backup on this share. The share must be remounted for any changes to this value to take effect.
Use as home share	checkbox		Allows the share to host user home directories. Each user is given a personal home directory when connecting to the share which is not accessible by other users. This allows for a personal, dynamic share. 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 preexec script.
AFP3 UNIX Privs	checkbox	✓	Set to enable Unix privileges supported by Mac OS X 10.5 and higher. Do not enable if the network has Mac OS X 10.4 or lower clients. Those systems 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. Please see the <i>note</i> (page ??) for more information.
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Separate entries with a comma, space, or tab. Please see the <i>note</i> (page ??) for more information.

Continued on next page

Table 11.1 – continued from previous page

Setting	Value	Advanced Mode	Description
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.

Note: If neither *Hosts Allow* or *Hosts Deny* contains an entry, then AFP share access is allowed for any host.

If there is a *Hosts Allow* list but no *Hosts Deny* list, then only allow hosts on the *Hosts Allow* list.


If there is a *Hosts Deny* list but no *Hosts Allow* list, then allow all hosts that are not on the *Hosts Deny* list.

If there is both a *Hosts Allow* and *Hosts Deny* list, then allow all hosts that are on the *Hosts Allow* list. If there is a host not on the *Hosts Allow* and not on the *Hosts Deny* list, then allow it.

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 TrueNAS® 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*.

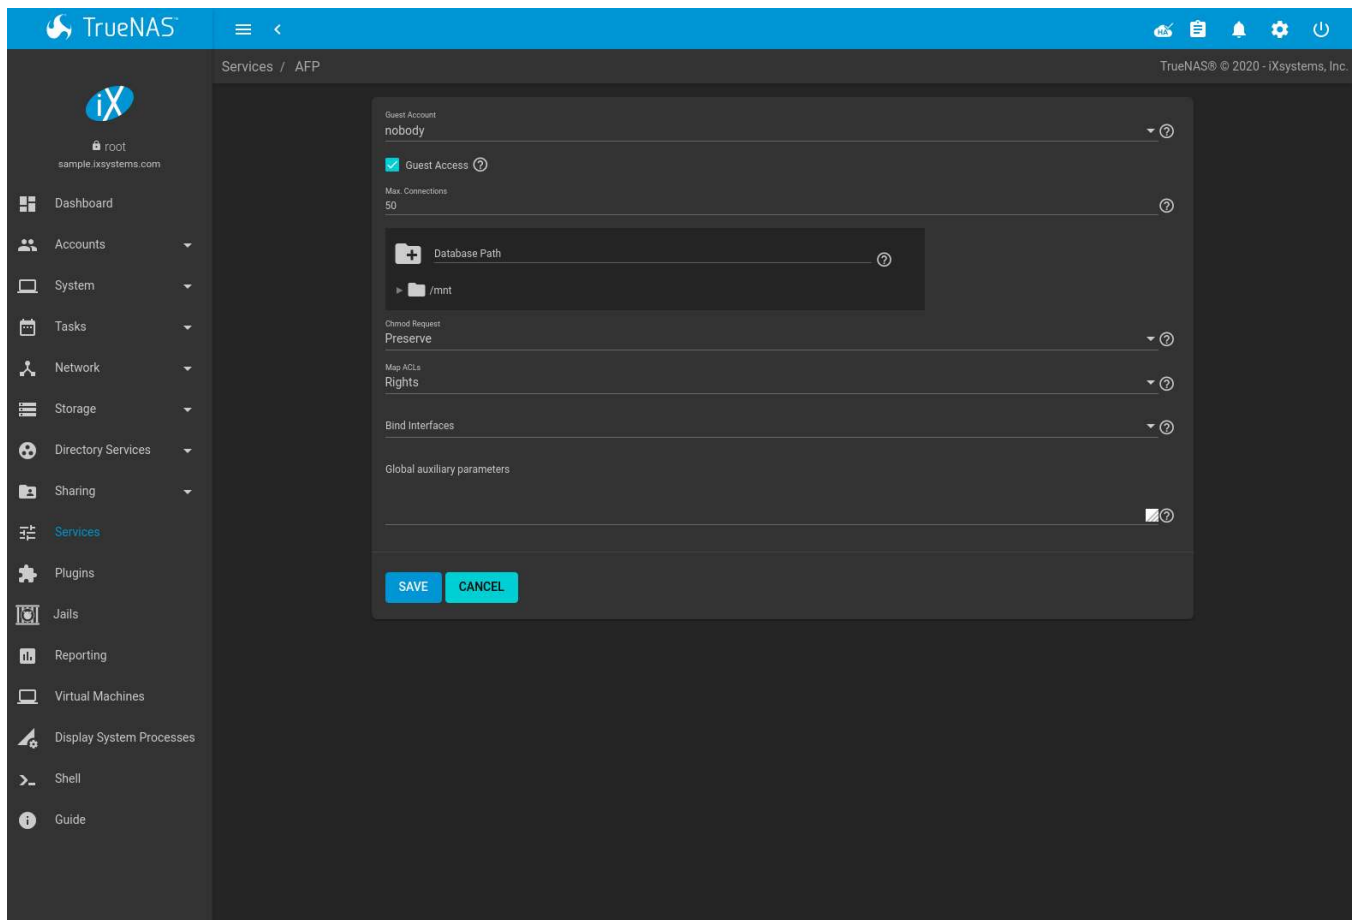


Fig. 11.2: Creating a Guest AFP Share

Next, create a dataset for the guest share. Refer to [Adding Datasets](#) (page 142) 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. **User:** Use the drop-down to select *Nobody*.
2. Click *SAVE*.

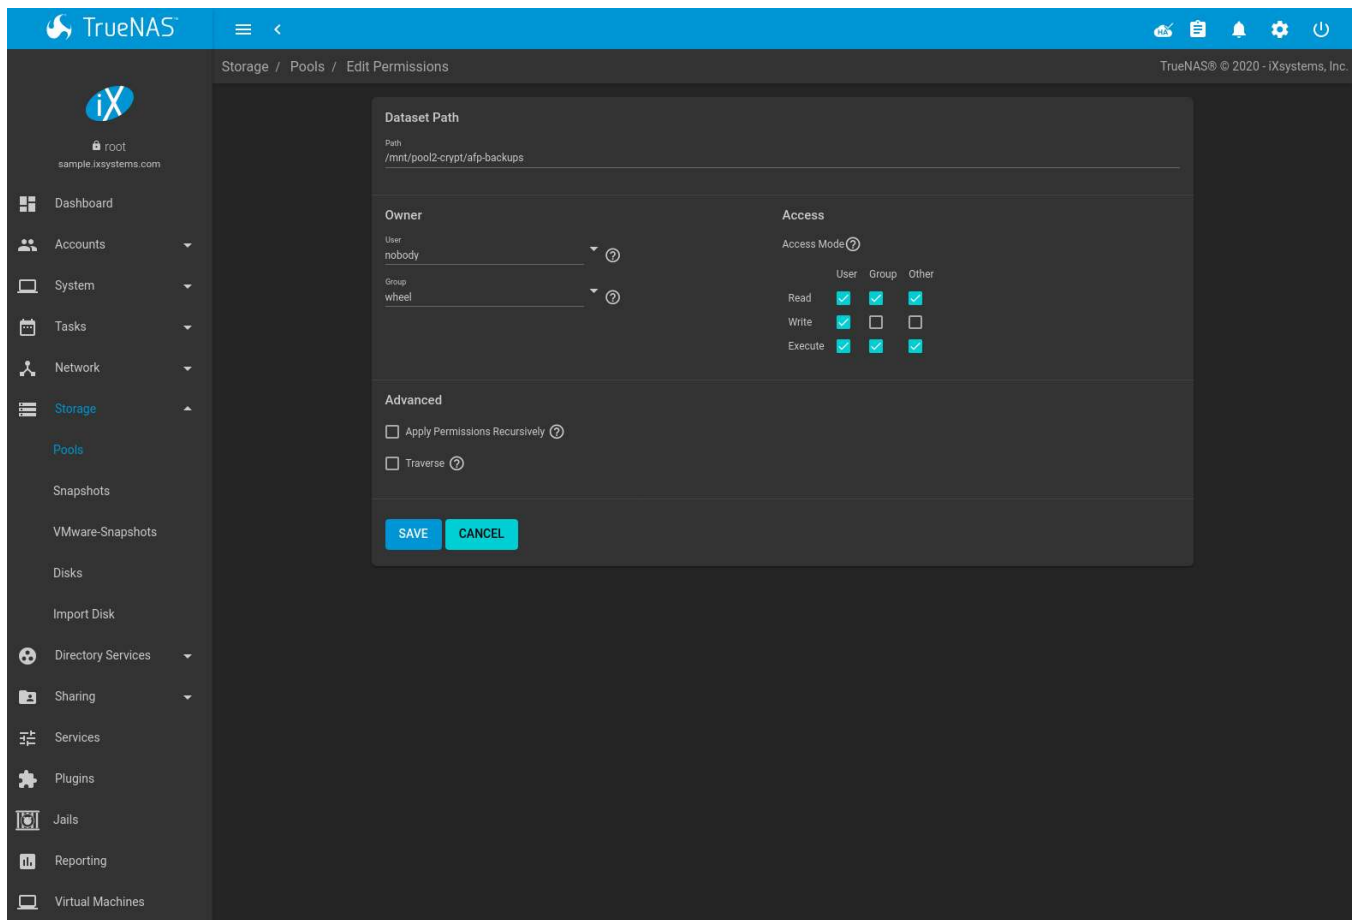


Fig. 11.3: Editing Dataset Permissions for Guest AFP Share

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 use Finder to 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 TrueNAS® 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.

11.2 Block (iSCSI)

iSCSI is a protocol standard for the consolidation of storage data. iSCSI allows TrueNAS® 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 TrueNAS® system. The client requires initiator software to initiate the connection to the iSCSI share.

Target: a storage resource on the TrueNAS® 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. TrueNAS® supports up to 1024 LUNs.

ALUA: *Asymmetric Logical Unit Access* allows a client computer to discover the best path to the storage on a TrueNAS® system. HA storage clusters can provide multiple paths to the same storage. For example, the disks are directly connected to the primary computer and provide high speed and bandwidth when accessed through that primary computer. The same disks are also available through the secondary computer, but because they are not directly connected to it, speed and bandwidth are restricted. With ALUA, clients automatically ask for and use the best path to the storage. If one of the TrueNAS® HA computers becomes inaccessible, the clients automatically switch to the next best alternate path to the storage. When a better path becomes available, as when the primary host becomes available again, the clients automatically switch back to that better path to the storage.

Note: Do not enable ALUA on TrueNAS® unless it is supported by and enabled on the client computers also. ALUA only works properly when enabled on both the client and server.

In TrueNAS®, 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 314) (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](#) (page 145) and use it to [create a device extent](#) (page 192).

11.2.1 iSCSI Wizard

To configure iSCSI, click *WIZARD* and follow each step:

1. Create or Choose Block Device:

- **Name:** Enter a name for the block device. Keeping the name short is recommended. Using a name longer than 63 characters can prevent access to the block device.
- **Type:** Select *File* or *Device* as the type of block device. *Device* provides virtual storage access to zvols, zvol snapshots, or physical devices. *File* provides virtual storage access to an individual file.
- **Device:** Select the unformatted disk, controller, zvol, or zvol snapshot. Select *Create New* for options to create a new zvol. If *Create New* is selected, use the browser to select an existing pool or dataset to store the new zvol. Enter the desired size of the zvol in *Size*. Only displayed when *Type* is set to *Device*.
- **File:** Browse to an existing file. Create a new file by browsing to a dataset and appending the file name to the path. When the file already exists, enter a size of 0 to use the actual file size. For new files, enter the size of the file to create. Only displayed when *Type* is set to *File*.
- **What are you using this for:** Choose the platform that will use this share. The associated options are applied to this share.

2. Portal

- **Portal:** Select an existing portal or choose *Create New* to configure a new portal.
- **Discovery Auth Method:** *NONE* allows anonymous discovery while *CHAP* and *Mutual CHAP* require authentication.
- **Discovery Auth Group:** Choose an existing [Authorized Access](#) (page 189) group ID or create a new authorized access. This is required when the *Discovery Auth Method* is set to *CHAP* or *Mutual CHAP*.
- **IP:** Select IP addresses to be listened on by the portal. Click *ADD* to add IP addresses with a different network port. The address 0.0.0.0 can be selected to listen on all IPv4 addresses, or :: to listen on all IPv6 addresses.
- **Port:** TCP port used to access the iSCSI target. Default is 3260.

3. Initiator

- **Initiators:** Leave blank to allow all or enter a list of initiator hostnames separated by spaces.
- **Authorized Networks:** Network addresses allowed to use this initiator. Leave blank to allow all networks or list network addresses with a CIDR mask. Separate multiple addresses with a space: 192.168.2.0/24 192.168.2.1/12.

4. Confirm Options

- Review the configuration and click *SUBMIT* to set up the iSCSI share.

The rest of this section describes iSCSI configuration in more detail.

Note: If the system has been licensed for Fibre Channel, the screens will vary slightly from those found in the rest of this section. Refer to the section on *Fibre Channel Ports* (page 196) for details.

11.2.2 Target Global Configuration

Sharing → *Block (iSCSI)* → *Target Global Configuration* contains settings that apply to all iSCSI shares. [Table 11.2](#) describes each option.

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.

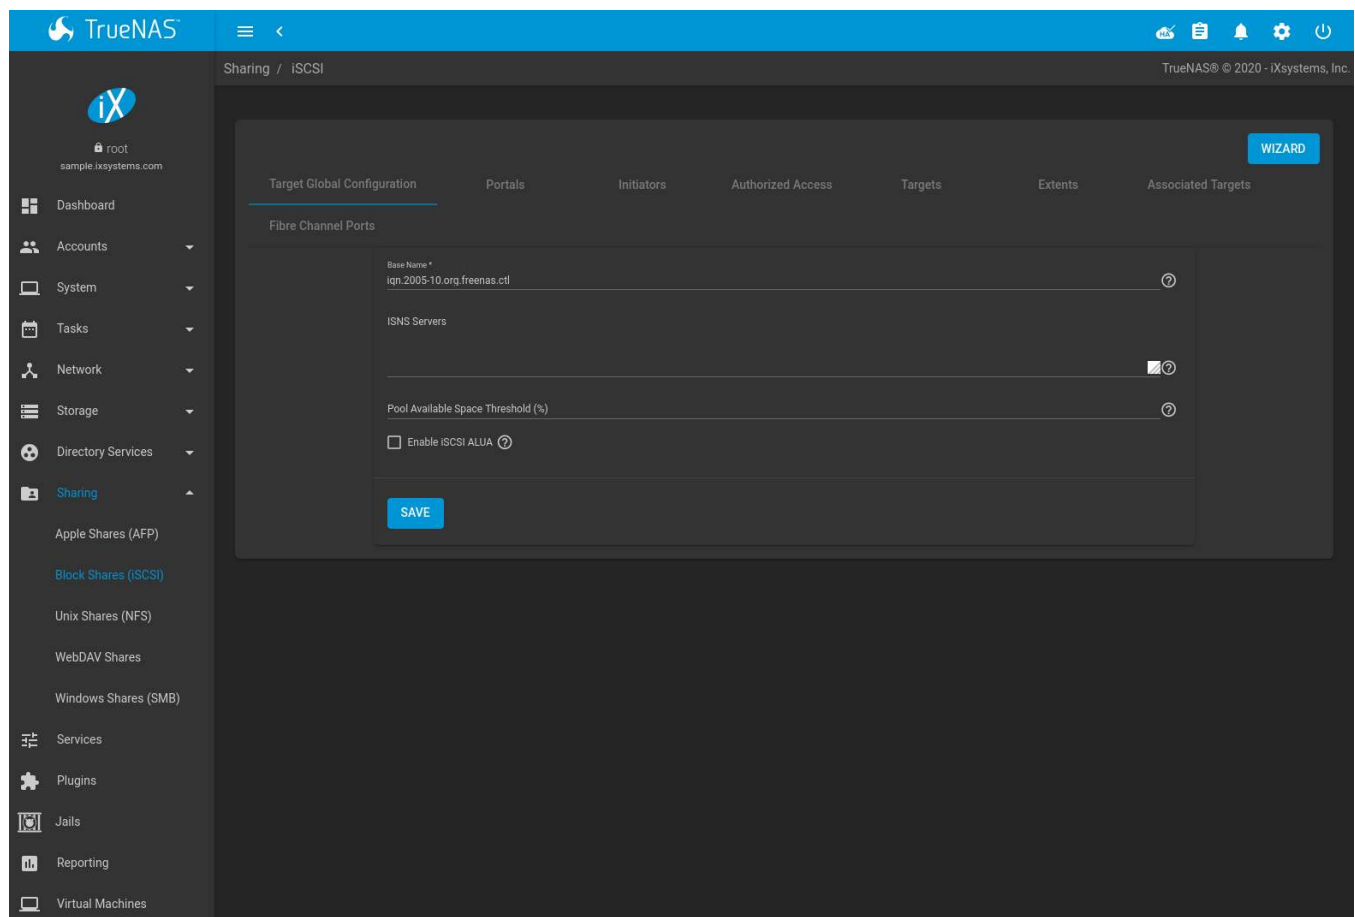


Fig. 11.5: iSCSI Target Global Configuration Variables

Table 11.2: Target Global Configuration Settings

Setting	Value	Description
Base Name	string	Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the “Constructing iSCSI names using the iqn. format” section of RFC 3721 (https://tools.ietf.org/html/rfc3721.html).
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 remain in the pool. When this percentage is reached, the system issues an alert, but only if zvols are used. See VAAI (page 314) Threshold Warning for more information.
Enable iSCSI ALUA	checkbox	Allow initiator to discover paths to both TrueNAS controllers on the target and increase storage traffic efficiency. Requires ALUA-capable, High Availability (HA) hardware.

11.2.3 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 11.6](#).

Table 11.6 summarizes the settings that can be configured when adding a portal.

The screenshot displays the TrueNAS web interface for adding a new iSCSI portal. The breadcrumb navigation at the top indicates the path: *Sharing* / *iSCSI* / *Portals* / *Add*. The left sidebar contains various system management options. The main configuration area includes the following fields:

- Description:** A text input field with a help icon.
- Discovery Auth Method:** A dropdown menu currently set to *NONE*.
- Discovery Auth Group:** A dropdown menu.
- IP Address:** A text input field with a help icon.
- Port:** A dropdown menu currently set to *3260*.

At the bottom of the form are two buttons: **SAVE** and **CANCEL**.

Fig. 11.6: Adding an iSCSI Portal

Table 11.3: Portal Configuration Settings

Setting	Value	Description
Description	string	Optional description. Portals are automatically assigned a numeric group.
Discovery Auth Method	drop-down menu	<i>iSCSI</i> (page 231) 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 Group ID 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 IP addresses to be listened on by the portal. Click <i>ADD</i> to add IP addresses with a different network port. The address <code>0.0.0.0</code> can be selected to listen on all IPv4 addresses, or <code>::</code> to listen on all IPv6 addresses. Choose only physical interface IP addresses when configuring iSCSI ALUA. Do not use Virtual IP addresses with an ALUA configuration.
Port	integer	TCP port used to access the iSCSI target. Default is 3260.

TrueNAS® 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 TrueNAS® 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.

11.2.4 Initiators

The next step is to configure authorized initiators, or the systems which are allowed to connect to the iSCSI targets on the TrueNAS® system. To configure which systems can connect, go to *Sharing* → *Block (iSCSI)* → *Initiators* and click *ADD* as shown in Figure 11.7.






Fig. 11.7: Adding an iSCSI Initiator

Table 11.4 summarizes the settings that can be configured when adding an initiator.

Table 11.4: Initiator Configuration Settings

Setting	Value	Description
Allow All Initiators	checkbox	Accept all detected initiators. When set, all other initiator fields are disabled.
Connected Initiators	string	Initiators currently connected to the system. Shown in IQN format with an IP address. Set initiators and click an  to add the initiators to either the <i>Allowed Initiators</i> or <i>Authorized Networks</i> lists. Clicking <i>REFRESH</i> updates the <i>Connected Initiators</i> list.
Allowed Initiators (IQN)	string	Initiators allowed access to this system. Enter an iSCSI Qualified Name (IQN) (https://tools.ietf.org/html/rfc3720#section-3.2.6) and click + to add it to the list. Example: <code>iqn.1994-09.org.freebsd:freenas.local</code>
Authorized Networks	string	Network addresses allowed to use this initiator. Each address can include an optional CIDR (https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing) netmask. Click + to add the network address to the list. Example: <code>192.168.2.0/24</code>
Description	string	Any notes about initiators.

Click  (Options) on an initiator entry for options to *Edit* or *Delete* it.

11.2.5 Authorized Access

When using CHAP or mutual CHAP to provide authentication, creating authorized access is recommended. Do this by going to *Sharing* → *Block (iSCSI)* → *Authorized Access* and clicking *ADD*. The screen is shown in [Figure 11.8](#).

Note: This screen sets login authentication. This is different from discovery authentication which is set in [Global Configuration](#) (page 119).

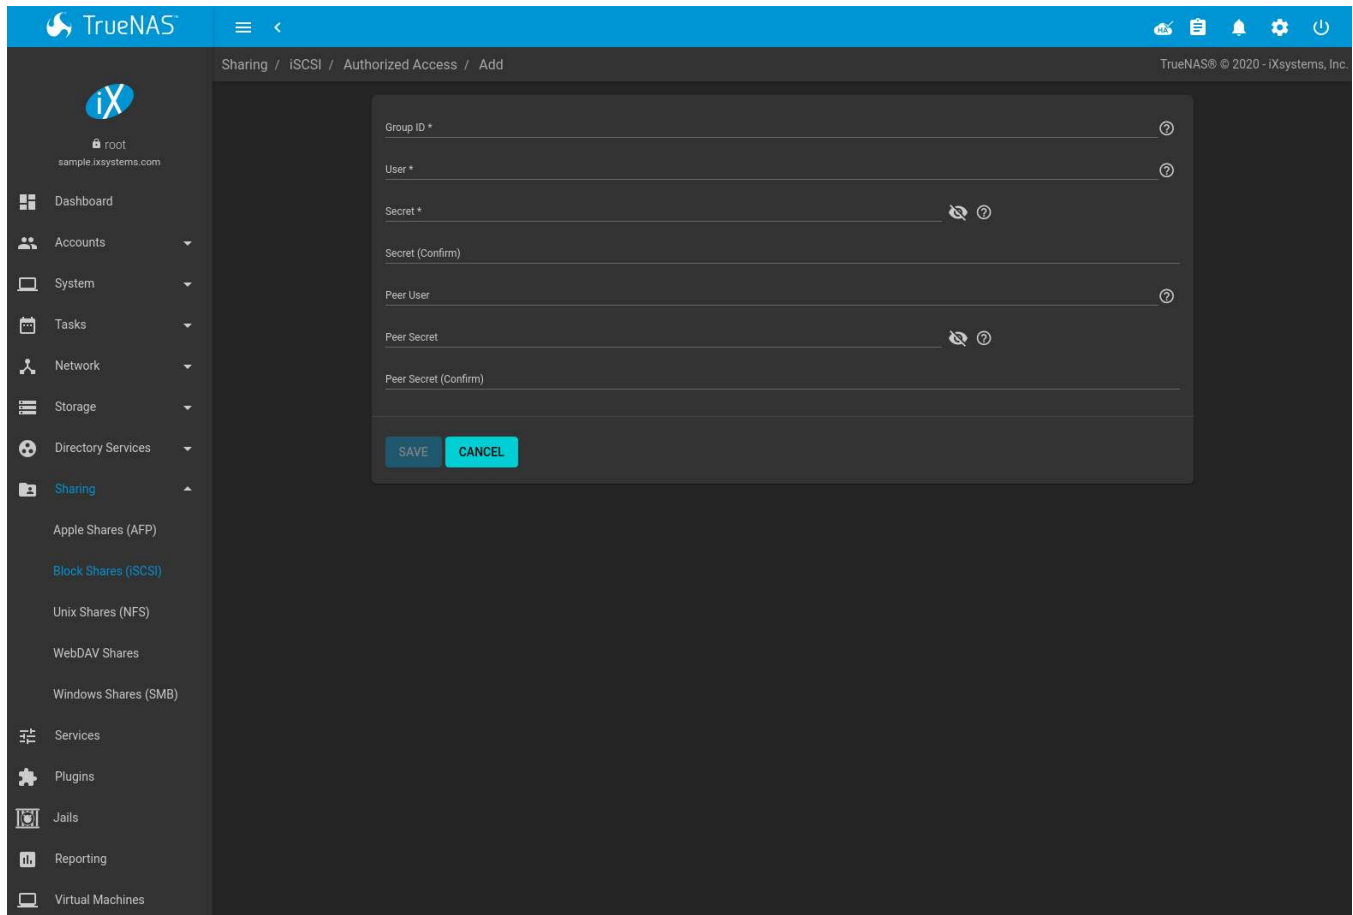


Fig. 11.8: Adding an iSCSI Authorized Access

Table 11.5 summarizes the settings that can be configured when adding an authorized access:

Table 11.5: Authorized Access Configuration Settings

Setting	Value	Description
Group ID	integer	Allow different groups to be configured with different authentication profiles. Example: enter <i>1</i> for all users in Group <i>1</i> to inherit the Group <i>1</i> authentication profile. Group IDs that are already configured with authorized access cannot be reused.
User	string	User account to create for CHAP authentication with the user on the remote system. Many initiators use the initiator name as the user name.
Secret	string	User password. Must be at least <i>12</i> and no more than <i>16</i> characters long.

Continued on next page

Table 11.5 – continued from previous page

Setting	Value	Description
Peer User	string	Only entered when configuring mutual CHAP. Usually the same value as <i>User</i> .
Peer Secret	string	Mutual secret password. Required when <i>Peer User</i> is set. Must be different than the <i>Secret</i> . Must be at least 12 and no more than 16 characters long.

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 11.9, 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.

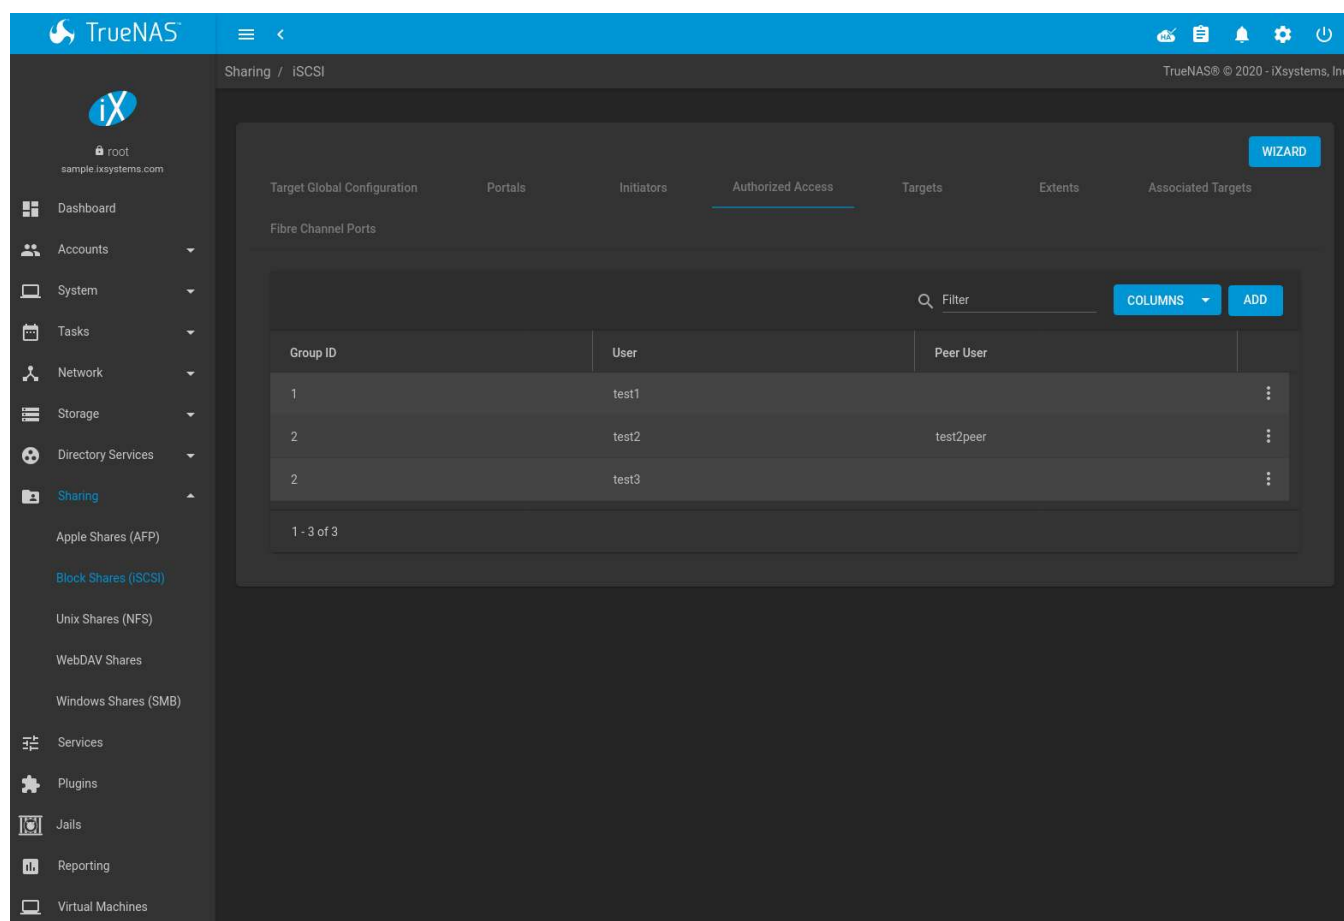


Fig. 11.9: Viewing Authorized Accesses

11.2.6 Targets

Next, create a Target by going to *Sharing* → *Block (iSCSI)* → *Targets* and clicking *ADD* as shown in Figure 11.10. A target combines a portal ID, allowed initiator ID, and an authentication method. Table 11.6 summarizes the settings that can be configured when creating a Target.

Note: An iSCSI target creates a block device that may be accessible to multiple initiators. A clustered filesystem is required on the block device, such as VMFS used by VMware ESX/ESXi, in order for multiple initiators to mount the

block device read/write. If a traditional filesystem such as EXT, XFS, FAT, NTFS, UFS, or ZFS is placed on the block device, care must be taken that only one initiator at a time has read/write access or the result will be filesystem corruption. If multiple clients need access to the same data on a non-clustered filesystem, use SMB or NFS instead of iSCSI, or create multiple iSCSI targets (one per client).

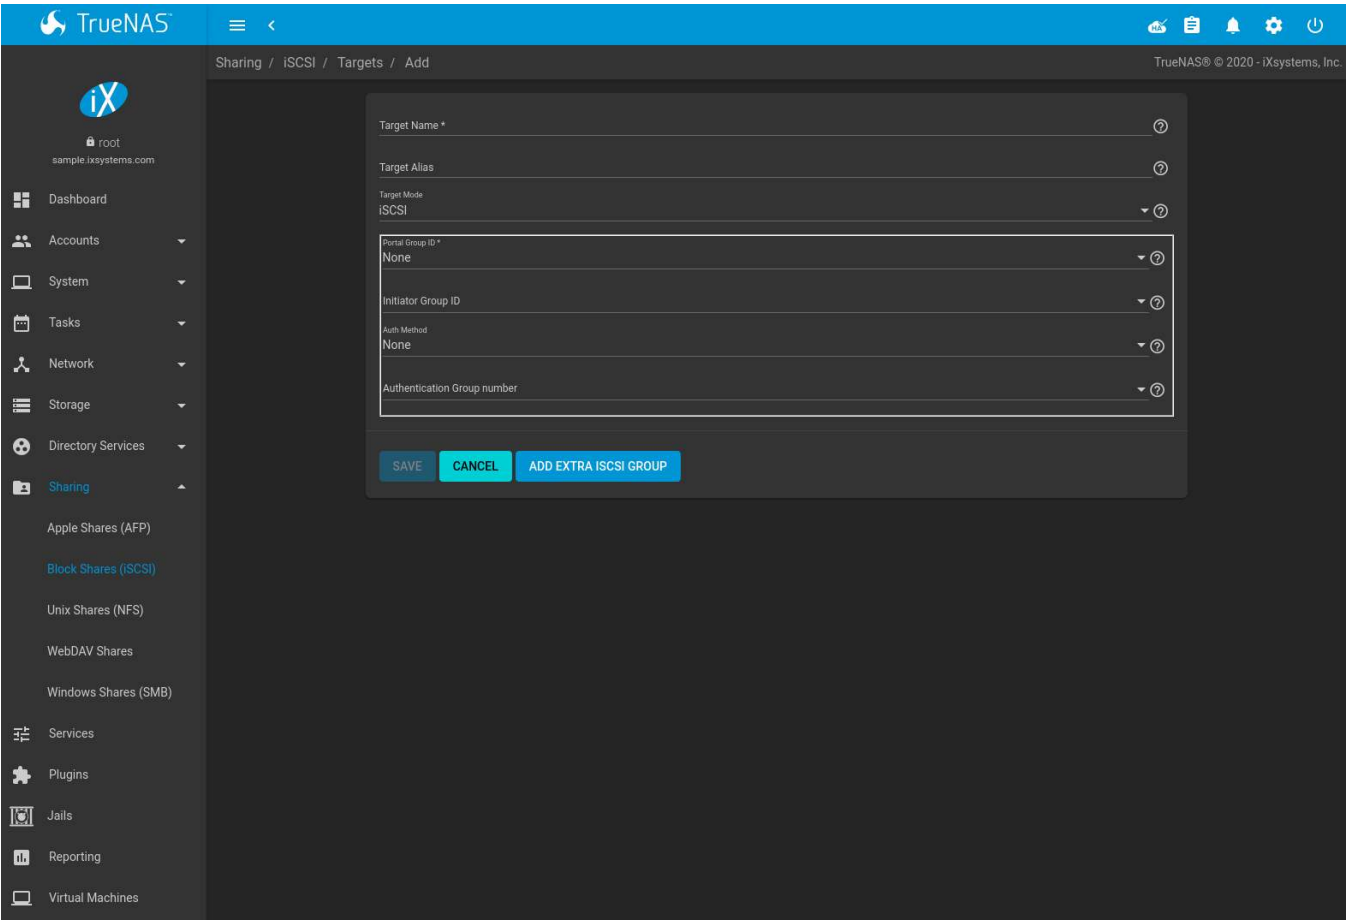


Fig. 11.10: Adding an iSCSI Target

Table 11.6: Target Settings

Setting	Value	Description
Target Name	string	Required. The base name is automatically prepended if the target name does not start with <i>iqn</i> . Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the “Constructing iSCSI names using the iqn. format” section of RFC 3721 (https://tools.ietf.org/html/rfc3721.html).
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	<i>None, Auto, CHAP, or Mutual CHAP.</i>

Continued on next page

Table 11.6 – continued from previous page

Setting	Value	Description
Authentication Group number	drop-down menu	Select <i>None</i> or an integer. This number represents the number of existing authorized accesses.

11.2.7 Extents

iSCSI targets provide virtual access to resources on the TrueNAS® 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, or a hardware RAID volume.

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 TrueNAS® [VAAI](#) (page 314) 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 80% when using iSCSI. The capacity of an existing extent can be increased as shown in [Growing LUNs](#) (page 199).

To add an extent, go to *Sharing* → *Block (iSCSI)* → *Extents* and click *ADD*. In the example shown in [Figure 11.11](#), the device extent is using the `export` zvol that was previously created from the `/mnt/pool1` pool.

[Table 11.7](#) 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**.

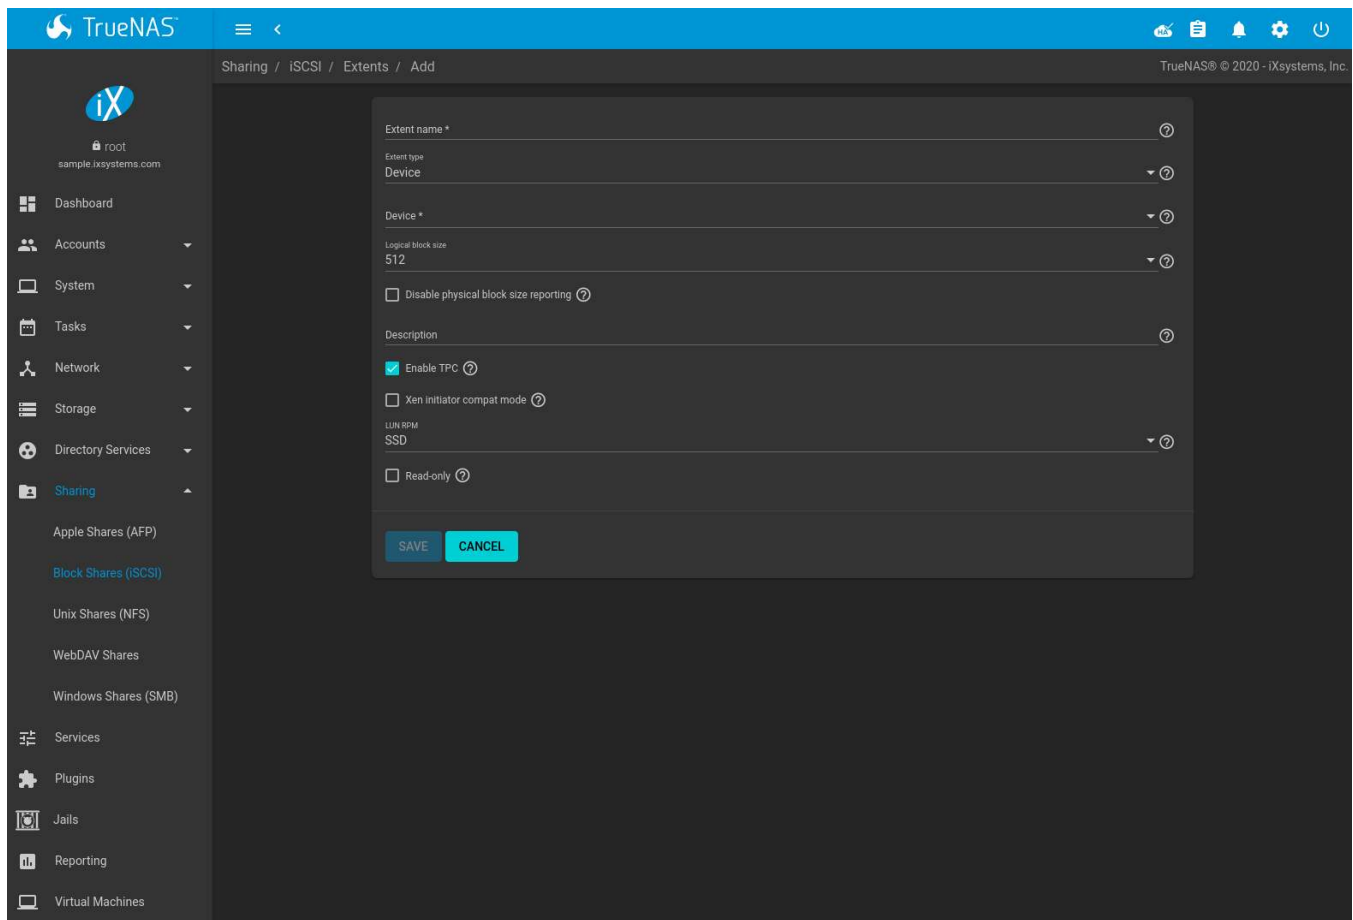


Fig. 11.11: Adding an iSCSI Extent

Table 11.7: 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	<i>File</i> shares the contents of an individual file. <i>Device</i> shares an entire device.
Path to the extent	browse button	Only appears when <i>File</i> is selected. Browse to an existing file. Create a new file by browsing to a dataset and appending the file name to the path. Extents cannot be created inside a jail root directory.
Extent size	integer	Only appears when <i>File</i> is selected. Entering 0 uses the actual file size and requires that the file already exists. Otherwise, specify the file size for the new file.
Device	drop-down menu	Only appears when <i>Device</i> is selected. Select the unformatted disk, controller, zvol, or zvol snapshot.
Logical block size	drop-down menu	Leave at the default of 512 unless the initiator requires a different block size.

Continued on next page

Table 11.7 – continued from previous page

Setting	Value	Description
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.
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 314) Threshold Warning.
Description	string	Notes about this extent.
Enable TPC	checkbox	Set to allow an initiator to bypass normal access control and access any scannable target. This allows xcopy (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc771254(v=ws.11)) operations which are otherwise blocked by access control.
Xen initiator compat mode	checkbox	Set 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 to prevent the initiator from initializing this LUN.
Enable	checkbox	Set to enable the iSCSI extent.

New extents have been added to *Sharing* → *Block (iSCSI)* → *Extents*. The associated *Serial* and Network Address Authority (NAA) are shown along with the extent name.

11.2.8 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 11.12](#). Use the drop-down menus to select the existing target and extent. Click *SAVE* to add an entry for the LUN.




Fig. 11.12: Associating a Target With an Extent

Table 11.8 summarizes the settings that can be configured when associating targets and extents.

Table 11.8: Associated Target Configuration Settings

Setting	Value	Description
Target	drop-down menu	Select an existing target.
LUN ID	integer	Select or enter a value between 0 and 1023. Some initiators expect a value less than 256. Leave this field blank to automatically 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.

11.2.9 Fibre Channel Ports

If the TrueNAS® system has Fibre Channel ports, *Sharing* → *Block (iSCSI)* appears as *Sharing* → *Block (iSCSI/FC)* and an extra *Fibre Channel Ports* tab is added. An example is shown in Figure 11.13.

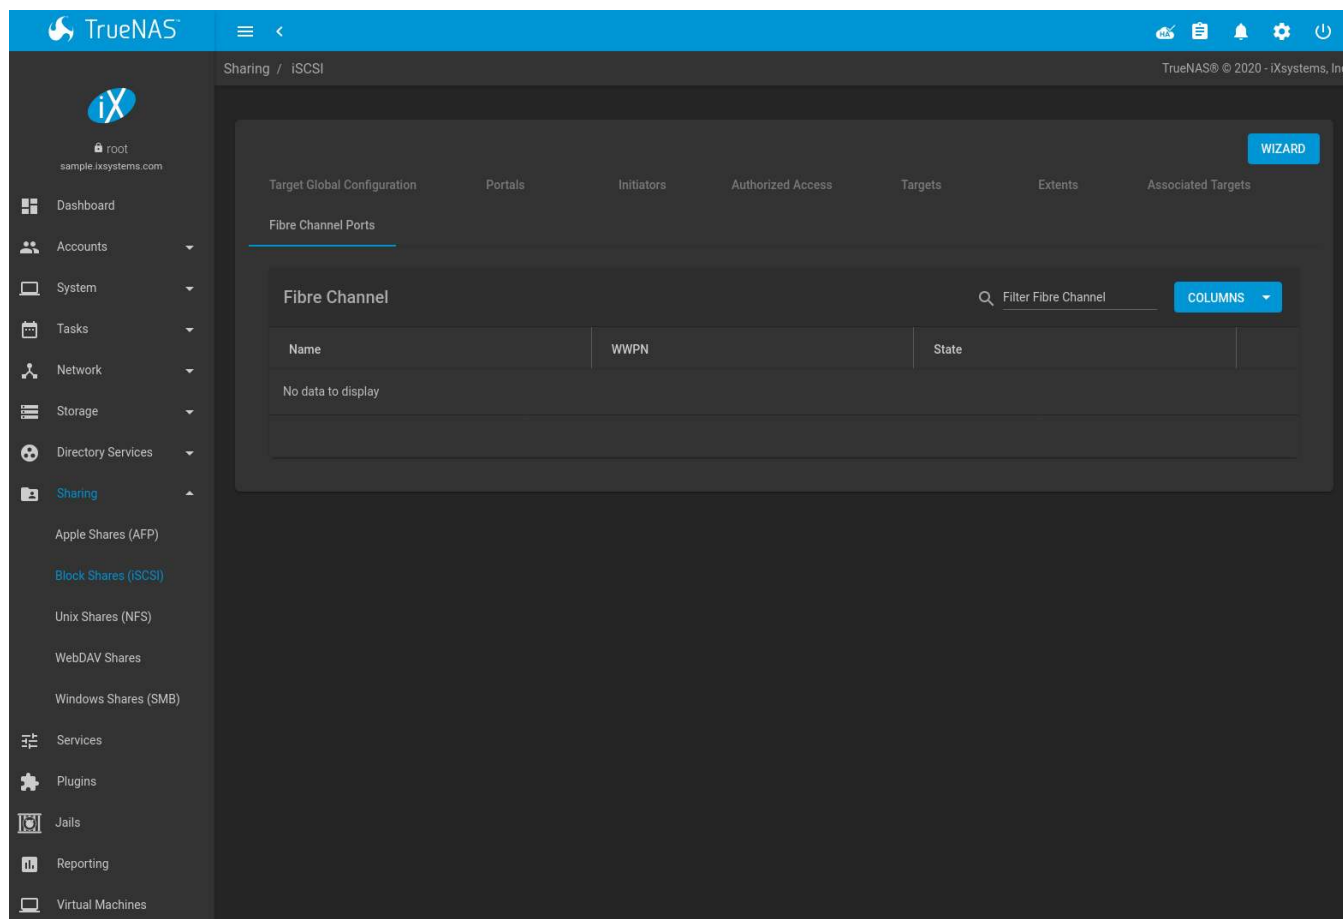


Fig. 11.13: Block (iSCSI) Screen

Since the *Portals*, *Initiators*, and *Authorized Access* screens only apply to iSCSI, they are marked as such and can be ignored when configuring Fibre Channel.

As shown in Figure 11.14, an extra *Target Mode* option appears after going to *Targets* and clicking *ADD*. This new option is to select whether the target to create is iSCSI, Fibre Channel, or both.

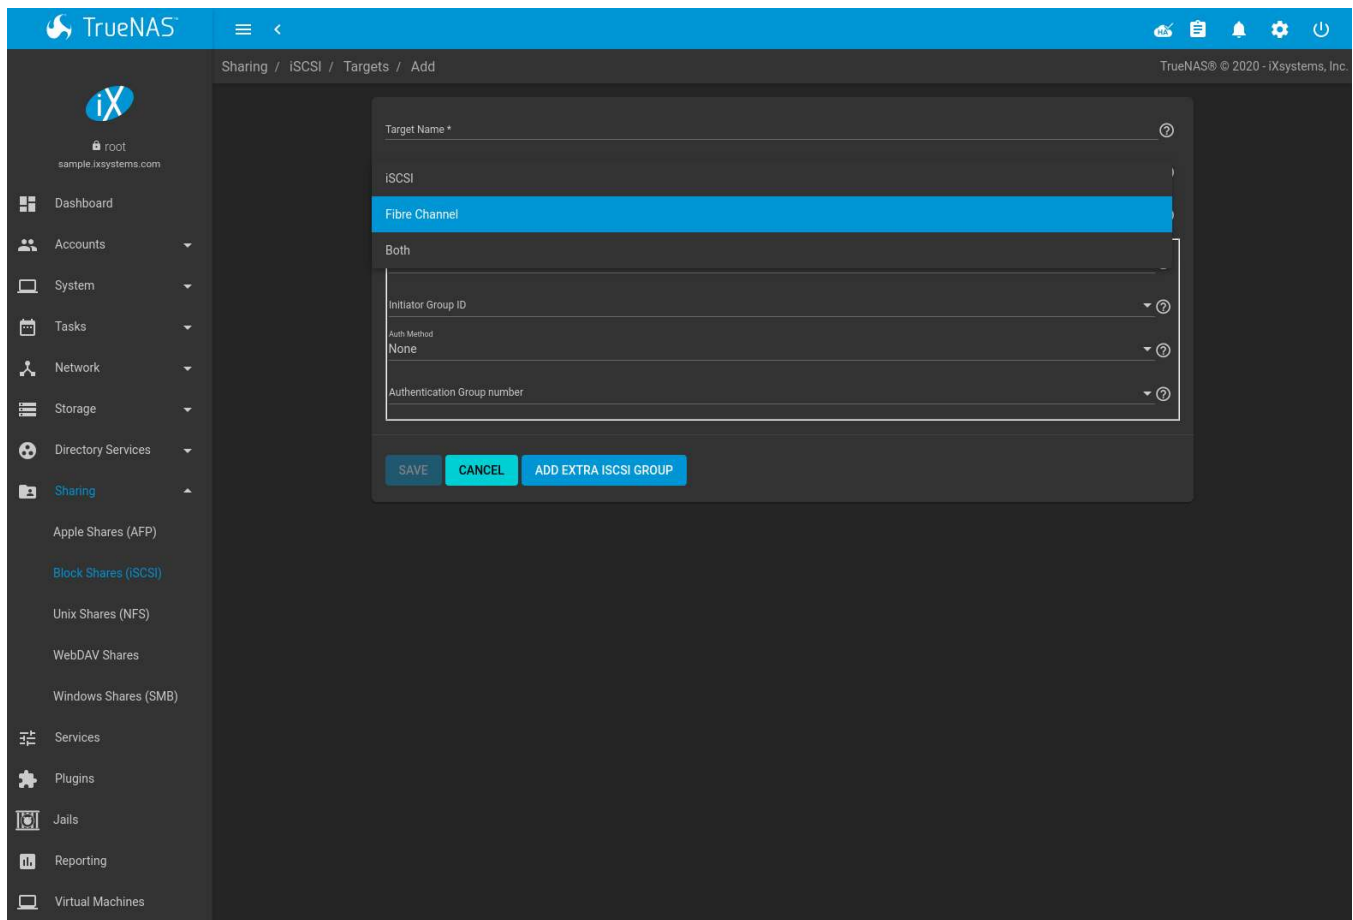


Fig. 11.14: Add Target Screen

The screens for adding an extent and associating a target are the same as described in [Extents](#) (page 192) and [Associated Targets](#) (page 194).

Note: The *Target* tab of [Reporting](#) (page 285) provides Fibre Channel port bandwidth graphs.

Fibre Channel can be configured for NPIV (N_Port ID Virtualization). NPIV allows the administrator to use switch zoning to configure each virtual port as if it was a physical port in order to provide access control. This is important in an environment with a mix of Windows systems and virtual machines in order to prevent automatic or accidental reformatting of targets containing unrecognized filesystems. It can also be used to segregate data; for example, to prevent the engineering department from accessing data from the human resources department. Refer to the switch documentation for details on how to configure zoning of virtual ports.

To create the virtual ports on the TrueNAS® system, go to *System* → *Tunables*, click *ADD*, and enter these options:

- **Variable:** input *hint.isp.X.vports*, replacing X with the number of the physical interface.
- **Value:** input the number of virtual ports to create. Note that there cannot be more than 125 SCSI target ports and that number includes all physical Fibre Channel ports, all virtual ports, and all configured combinations of iSCSI portals and targets.
- **Type:** make sure *loader* is selected.

In the example shown in [Figure 11.15](#), two physical interfaces were each assigned 4 virtual ports. Note that two tunables were required, one for each physical interface. After the tunables are created, the configured number of virtual ports appears in the *Fibre Channel Ports* screen so they can be associated with targets. They will also be advertised to the switch so zoning can be configured on the switch. After a virtual port has been associated with a target, it is added to the *Target* tab of [Reporting](#) (page 285) where its bandwidth usage can be viewed.

Variable	Value	Type	Description	Enabled
hint.isp.1.vports	4	loader	for NPIV	yes
hint.isp.0.vports	4	loader	for NPIV	yes
vfs.zfs.fetch.max_distance	33554432	sysctl	Generated by autotune	yes
vfs.zfs.metastab.lba_weighting_r	1	sysctl	Generated by autotune	yes
vfs.zfs.l2arc.write_max	10000000	sysctl	Generated by autotune	yes
vfs.zfs.l2arc.write_boost	40000000	sysctl	Generated by autotune	yes
vfs.zfs.l2arc.norw	0	sysctl	Generated by autotune	yes
vfs.zfs.l2arc.noprefetch	0	sysctl	Generated by autotune	yes
vfs.zfs.l2arc.headroom	2	sysctl	Generated by autotune	yes
vfs.zfs.arc_max	57365710848	sysctl	Generated by autotune	yes
net.inet.tcp.sendspace	262144	sysctl	Generated by autotune	yes
net.inet.tcp.sendbuf_max	16777216	sysctl	Generated by autotune	yes
net.inet.tcp.sendbuf_inc	16384	sysctl	Generated by autotune	yes
net.inet.tcp.recvspace	262144	sysctl	Generated by autotune	yes

Fig. 11.15: Adding Virtual Ports

11.2.10 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.

macOS 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 ESX\(i\)](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 TrueNAS® 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*.

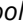
If the LUN is not discovered by ESXi, make sure that promiscuous mode is set to *Accept* in the vSwitch.

11.2.11 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 re-size 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.

11.2.11.1 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 11.16](#), the current size of the zvol named *zvol1* is 4 GiB.

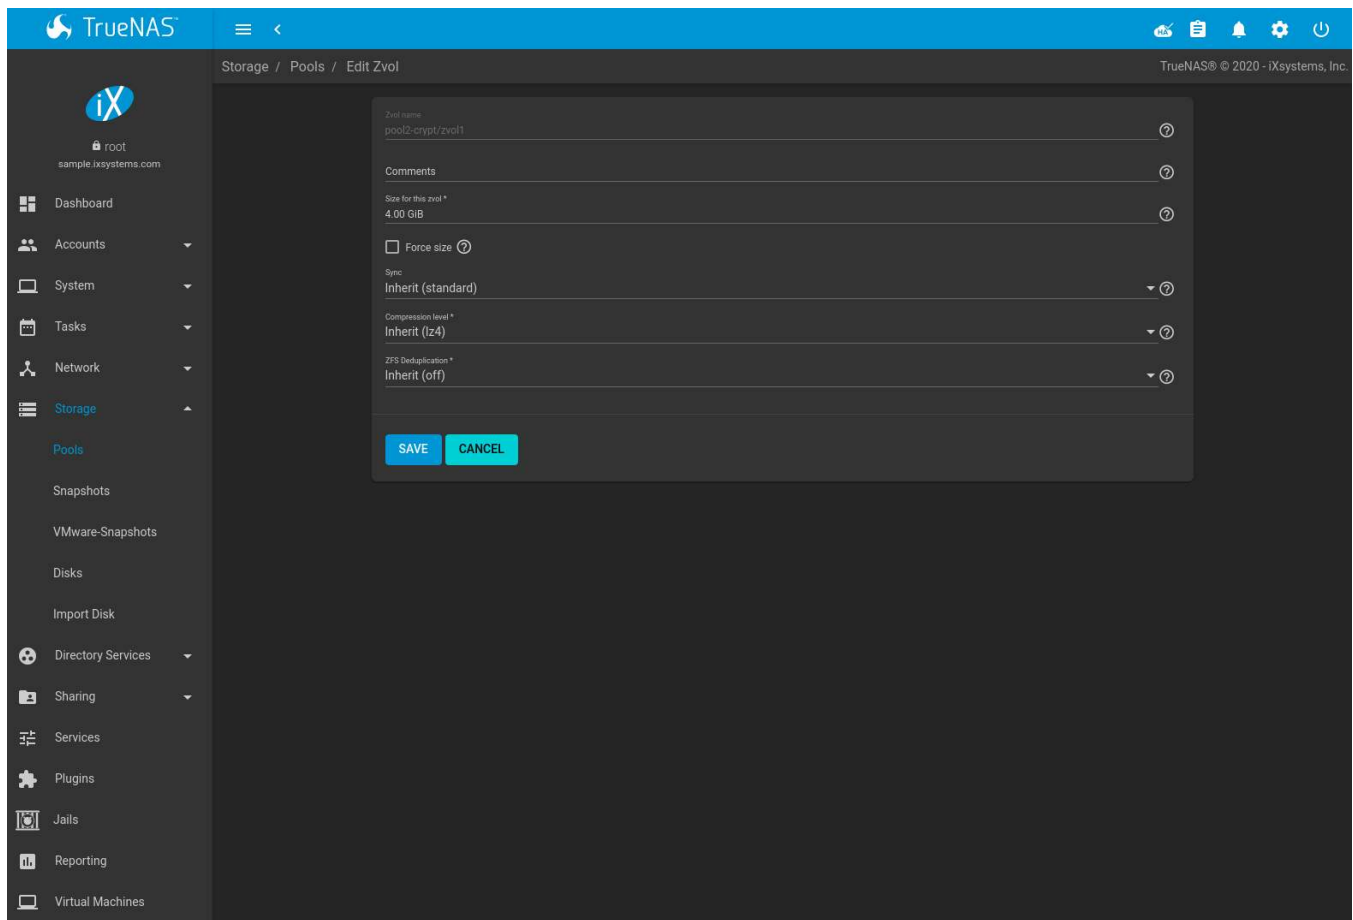


Fig. 11.16: Editing an Existing Zvol

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.

11.2.11.2 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 302) to grow the file extent. This example grows `/mnt/pool11/data` by 2 GiB:

```
truncate -s +2g /mnt/pool11/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.

11.3 Unix (NFS) Shares

TrueNAS® supports sharing pools, datasets, and directories over the Network File System (NFS). Clients use the `mount` command to mount the share. Mounted NFS shares appear as another directory on the client system. Some Linux distros require the installation of additional software to mount an NFS share. Windows systems must enable Services for NFS in the Ultimate or Enterprise editions or install an NFS client application.

Create an NFS share by going to *Sharing* → *Unix (NFS) Shares* and clicking *ADD*. Figure 11.17 shows an example of creating an NFS share.




Fig. 11.17: NFS Share Creation

Remember these points when creating NFS shares:

1. Clients specify the *Path* when mounting the share.
2. The *Maproot* and *Mapall* options cannot both be enabled. 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 a unique filesystem. Individual NFS shares cannot cross filesystem boundaries. Adding paths to share more directories only works if those directories are within the same filesystem.
4. The network and host must be unique to both each created share and the filesystem or directory included in that share. Because `/etc/exports` is not an access control list (ACL), the rules contained in `/etc/exports` become undefined with 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 scenarios where there are:

- two networks, `10.0.0.0/8` and `20.0.0.0/8`
- a ZFS pool named `pool1` with a dataset named `dataset1`
- `dataset1` contains directories named `directory1`, `directory2`, and `directory3`

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 the dataset `/mnt/pool1/dataset1`. An additional path to directory `/mnt/pool1/dataset1/directory1` is added.

The correct method to configure this share is to set the *Path* to `/mnt/pool1/dataset1` and set the *All dirs* box. This allows the client to also mount `/mnt/pool1/dataset1/directory1` when `/mnt/pool1/dataset1` is mounted.

Additional paths are used to define specific directories to be shared. For example, `dataset1` has three directories. To share only `/mnt/pool1/dataset1/directory1` and `/mnt/pool1/dataset1/directory2`, create paths for `directory1` and `directory2` within the share. This excludes `directory3` from the share.

Restricting a specific directory to a single network is done by creating a share for the volume or dataset and a share for the directory within that volume or dataset. Define the authorized networks for both shares.

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`

This requires the creation of two shares. It cannot be done with only one share.

Table 11.9 summarizes the available configuration options in the *Sharing/NFS/Add* screen. Click *ADVANCED MODE* to see all settings.

Table 11.9: NFS Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		Browse to the dataset or directory to be shared. Click <i>ADD</i> to specify multiple paths.
Comment	string		Text describing the share. Typically used to name the share. If left empty, this shows the <i>Path</i> entries of the share.
All dirs	checkbox		Allow the client to also mount any subdirectories of the selected pool or dataset.
Read only	checkbox		Prohibit writing to the share.
Quiet	checkbox	✓	Restrict some syslog diagnostics to avoid some 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: <code>1.2.3.0/24</code> . 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.
Maproot User	drop-down menu	✓	When a user is selected, the <i>root</i> user is limited to permissions of that user.

Continued on next page

Table 11.9 – continued from previous page

Setting	Value	Advanced Mode	Description
Maproot Group	drop-down menu	✓	When a group is selected, the <i>root</i> user is also limited to permissions of that group.
Mapall User	drop-down menu	✓	TrueNAS® user or user imported with <i>Active Directory</i> (page 165). The specified permissions of that user are used by all clients.
Mapall Group	drop-down menu	✓	TrueNAS® group or group imported with <i>Active Directory</i> (page 165). 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> . 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). If multiple security mechanisms are added to the <i>Selected</i> column using the arrows, use the <i>Up</i> or <i>Down</i> buttons to list in order of preference.

Go to *Sharing* → *Unix (NFS)* and click ⓘ (Options) and *Edit* to edit an existing share. Figure 11.18 shows the configuration screen for the existing *nfs_share1* share. Options are the same as described in *NFS Share Options* (page 202).




Fig. 11.18: NFS Share Settings

11.3.1 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.

11.3.2 Connecting to the Share

The following examples share this configuration:

1. The TrueNAS® 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*

11.3.2.1 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 TrueNAS® 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.

Successfully mounting 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/) (<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* write to the TrueNAS® system */mnt/pool1/nfs_share1* dataset.

NFS share settings cannot be changed when the share is mounted on a client computer. 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
```

11.3.2.2 From Microsoft

Windows NFS client support varies with versions and releases. For best results, use [Windows \(SMB\) Shares](#) (page 209).

11.3.2.3 From macOS

A macOS client uses Finder to mount the NFS volume. Go to *Go* → *Connect to Server*. In the *Server Address* field, enter `nfs://` followed by the IP address of the TrueNAS® system, and the name of the pool or dataset being shared by NFS. The example shown in [Figure 11.19](#) continues with the example of `192.168.2.2:/mnt/pool1/nfs_share1`.

Finder opens automatically after connecting. The IP address of the TrueNAS® system displays in the *SHARED* section of the left frame and the contents of the share display in the right frame. [Figure 11.20](#) shows an example where `/mnt/data` has one folder named `images`. The user can now copy files to and from the share.

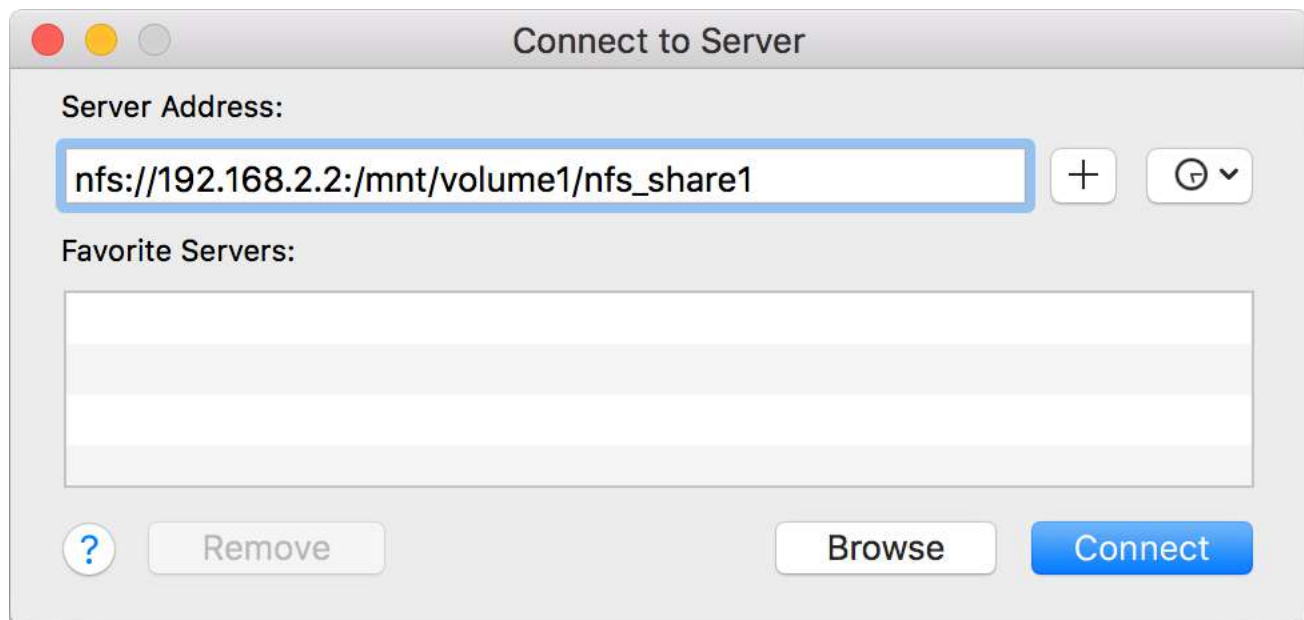


Fig. 11.19: Mounting the NFS Share from macOS




Fig. 11.20: Viewing the NFS Share in Finder

11.3.3 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 TrueNAS® and restart the NFS service after the upgrade to clear the NFS cache.

If clients see “reverse DNS” errors, add the TrueNAS® 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, TrueNAS® 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 302). A high proportion of retries and timeouts compared to reads usually indicates network problems.

11.4 WebDAV Shares

In TrueNAS®, 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 TrueNAS® 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 TrueNAS® 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 11.21](#).




Fig. 11.21: Adding a WebDAV Share

Table 11.10 summarizes the available options.

Table 11.10: WebDAV Share Options

Setting	Value	Description
Share Path Name	string	Enter a name for the share.
Comment	string	Optional.
Path	browse button	Enter the path or <i>Browse</i> to the pool or dataset to share. Appending a new name to the path creates a new dataset. Example: <i>/mnt/pool1/newdataset</i> .
Read Only	checkbox	Set to prohibit users from writing to the share.
Change User & Group	checkbox	Ownership of all files in the share will be changed to user <code>webdav</code> and group <code>webdav</code> . Existing permissions will not be changed, but the ownership change might make files inaccessible to their original owners. This operation cannot be undone! If unset, ownership of files to be accessed through WebDAV must be manually set to the <code>webdav</code> or <code>www</code> user/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 251).

11.5 Windows (SMB) Shares

TrueNAS® 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 [Robo-copy](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\(v=ws.11\)](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.

[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/>).

Warning: SMB1 is disabled by default for security (<https://www.ixsystems.com/blog/library/do-not-use-smb1/>). If necessary, SMB1 can be enabled in *Services* → *SMB Configure*.

Figure 11.22 shows the configuration screen that appears after clicking *Sharing* → *Windows (SMB Shares)*, then *ADD*.




Fig. 11.22: Adding an SMB Share

Table 11.11 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) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf>) provides more details for each configurable option.

Table 11.11: SMB Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		Select the pool, dataset, or directory to share. The same path can be used by more than one share.
Name	string		Name the new share. Each share name must be unique. The names <i>global</i> , <i>homes</i> , and <i>printers</i> are reserved and cannot be used.

Continued on next page

Table 11.11 – continued from previous page

Setting	Value	Advanced Mode	Description
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 <code>john</code> . 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.
Description	string		Description of the share or notes on how it is used.
Time Machine	checkbox		Enable Time Machine (https://developer.apple.com/library/archive/releasenotes/Networking/CH1-SW1) backups for this share. The process to configure a Time Machine backup is shown in Creating Authenticated and Time Machine Shares (page 218). Changing this setting on an existing share requires an <i>SMB</i> (page 240) service restart.
Export Read Only	checkbox	✓	Prohibit write access to this share.
Browsable to Network Clients	checkbox	✓	Determine whether this share name is included when browsing shares. Home shares are only visible to the owner regardless of this setting.
Export Recycle Bin	checkbox	✓	Files that are deleted from the same dataset are moved to the Recycle Bin and do not take any additional space. This is only applies over the SMB protocol. Deleting files over NFS will remove the files permanently. When the files are in a different dataset or a child dataset, they are copied to the dataset where the Recycle Bin is located. To prevent excessive space usage, files larger than 20 MiB are deleted rather than moved. Adjust the <i>Auxiliary Parameter</i> <code>crossrename:sizelimit=</code> setting to allow larger files. For example, <code>crossrename:sizelimit=50</code> allows moves of files up to 50 MiB in size. The recycle bin has read-write functionality. This means files can be permanently deleted or moved from the recycle bin. This is not a replacement for <i>ZFS Snapshots</i> (page 152).
Show Hidden Files	checkbox	✓	Disable the Windows <i>hidden</i> attribute on a new Unix hidden file. Unix hidden filenames start with a dot: <code>.foo</code> . Existing files are not affected.
Allow Guest Access	checkbox		Privileges are the same as the guest account. Guest access is disabled by default in Windows 10 version 1709 and Windows Server version 1903. Additional client-side configuration is required to provide guest access to these clients. MacOS clients: Attempting to connect as a user that does not exist in TrueNAS® <i>does not</i> automatically connect as the guest account. The <i>Connect As: Guest</i> option must be specifically chosen in MacOS to log in as the guest account. See the Apple documentation (https://support.apple.com/guide/mac-help/connect-mac-shared-computers-servers-mchlp1140/) for more details.

Continued on next page

Table 11.11 – continued from previous page

Setting	Value	Advanced Mode	Description
Only Allow Guest Access	checkbox	✓	Requires <i>Allow guest access</i> to also be enabled. Forces guest access for all connections.
Access Based Share Enumeration	checkbox	✓	Restrict share visibility to users with a current Windows Share ACL access of read or write. Use Windows administration tools to adjust the share permissions. See <code>smb.conf(5)</code> (https://www.freebsd.org/cgi/man.cgi?query=smb.conf).
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma (,), space, or tab. Please see the <i>note</i> (page ??) for more information.
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Specify <code>ALL</code> and list any hosts from <i>Hosts Allow</i> to have those hosts take precedence. Separate entries with a comma (,), space, or tab. Please see the <i>note</i> (page ??) for more information.
VFS Objects	selection	✓	Add virtual file system objects to enhance functionality. Table 11.12 summarizes the available objects.
Enable Shadow Copies	checkbox		Expose ZFS snapshots as Windows Shadow Copies (https://docs.microsoft.com/en-us/windows/desktop/vss/shadow-copies-and-shadow-copy-sets).
Auxiliary Parameters	string	✓	Additional <code>smb4.conf</code> (https://www.freebsd.org/cgi/man.cgi?query=smb.conf) parameters not covered by other option fields.


Note: If neither *Hosts Allow* or *Hosts Deny* contains an entry, then SMB share access is allowed for any host.

If there is a *Hosts Allow* list but no *Hosts Deny* list, then only allow hosts on the *Hosts Allow* list.

If there is a *Hosts Deny* list but no *Hosts Allow* list, then allow all hosts that are not on the *Hosts Deny* list.

If there is both a *Hosts Allow* and *Hosts Deny* list, then allow all hosts that are on the *Hosts Allow* list. If there is a host not on the *Hosts Allow* and not on the *Hosts Deny* list, then allow it.

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* →  (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) (<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) (<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* →  (Configure).

Table 11.12 provides an overview of the available VFS objects. Be sure to research each object **before** adding or deleting it from the *Selected* column of the *VFS Objects* field of the share. Some objects 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 11.12: Available VFS Objects

Value	Description
audit	Log share access, connects/disconnects, directory opens/creates/removes, and file opens/closes/renames/unlinks/chmods to syslog.
catia	Improve Mac interoperability by translating characters that are unsupported by Windows.
crossrename	Allow server side rename operations even if source and target are on different physical devices. Required for the recycle bin to work across dataset boundaries. Automatically added when <i>Export Recycle Bin</i> is enabled.
dirsort	Sort directory entries alphabetically before sending them to the client.
fruit	Enhance macOS support by providing the SMB2 AAPL extension and Netatalk interoperability. Automatically loads <i>catia</i> and <i>streams_xattr</i> , but see the warning (page 215) below.
full_audit	Record selected client operations to the system log.
ixnas	<p>Improves ACL compatibility with Windows, stores DOS attributes as file flags, optimizes share case sensitivity to improve performance, and enables User Quota Administration (page 218) from Windows. Enabled by default. Several <i>Auxiliary Parameters</i> are available with <i>ixnas</i>.</p> <p>Userspace Quota Settings:</p> <ul style="list-style-type: none"> • <i>ixnas:base_user_quota</i> = sets a ZFS user quota on every user that connects to the share. Example: <i>ixnas:base_user_quota</i> = 80G sets the quota to 80 GiB. • <i>ixnas:zfs_quota_enabled</i> = enables support for userspace quotas. Choices are <i>True</i> or <i>False</i>. Default is <i>True</i>. Example: <i>ixnas:zfs_quota_enabled</i> = <i>True</i>. <p>Home Dataset Settings:</p> <ul style="list-style-type: none"> • <i>ixnas:chown_homedir</i> = changes the owner of a created home dataset to the currently authenticated user. <i>ixnas:zfs_auto_homedir</i> must be set to <i>True</i>. Choices are <i>True</i> or <i>False</i>. Example: <i>ixnas:chown_homedir</i> = <i>True</i>. • <i>ixnas:homedir_quota</i> = sets a quota on new ZFS datasets. <i>ixnas:zfs_auto_homedir</i> must be set to <i>True</i>. Example: <i>ixnas:homedir_quota</i> = 20G sets the quota to 20 GiB. • <i>ixnas:zfs_auto_homedir</i> = creates new ZFS datasets for users connecting to home shares instead of folders. Choices are <i>True</i> or <i>False</i>. Default is <i>False</i>. Example: <i>ixnas:zfs_auto_homedir</i> = <i>False</i>.
media_harmony	Allow Avid editing workstations to share a network drive.
noacl	Disable NT ACL support. If an extended ACL is present in the share connection path, all access to this share will be denied. When the Read-only attribute (https://www.oreilly.com/openbook/samba/book/ch05_03.html)

Warning: 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 210).

Creating or editing an SMB share on a dataset with a [trivial Access Control List \(ACL\)](#) (<https://www.ixsystems.com/community/threads/methods-for-fine-tuning-samba-permissions.50739/>) prompts to [configure the ACL](#) (page 148) for the dataset.

To view all active SMB connections and users, enter `smbstatus` in the *Shell* (page 302). To log more details for clients that are attempting to authenticate to an SMB share, open the *Service* → *SMB* options and add `log_level = 1`, `auth_audit:5` to the *Auxiliary Parameters*.

11.5.1 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 TrueNAS® 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.

Note: Windows 10, Windows Server 2016 version 1709, and Windows Server 2019 disable SMB2 guest access. Read the [Microsoft security notice](#) (<https://support.microsoft.com/en-hk/help/4046019/guest-access-in-smb2-disabled-by-default-in-windows-10-and-windows-ser>) for details about security vulnerabilities with SMB2 guest access and instructions to re-enable guest logins on these Microsoft systems.

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 11.23](#).
3. Enable *Allow Guest Access*.
4. Press *SAVE*.

Note: If a dataset for the share has not been created, refer to [Adding Datasets](#) (page 142) to find out more about dataset creation.

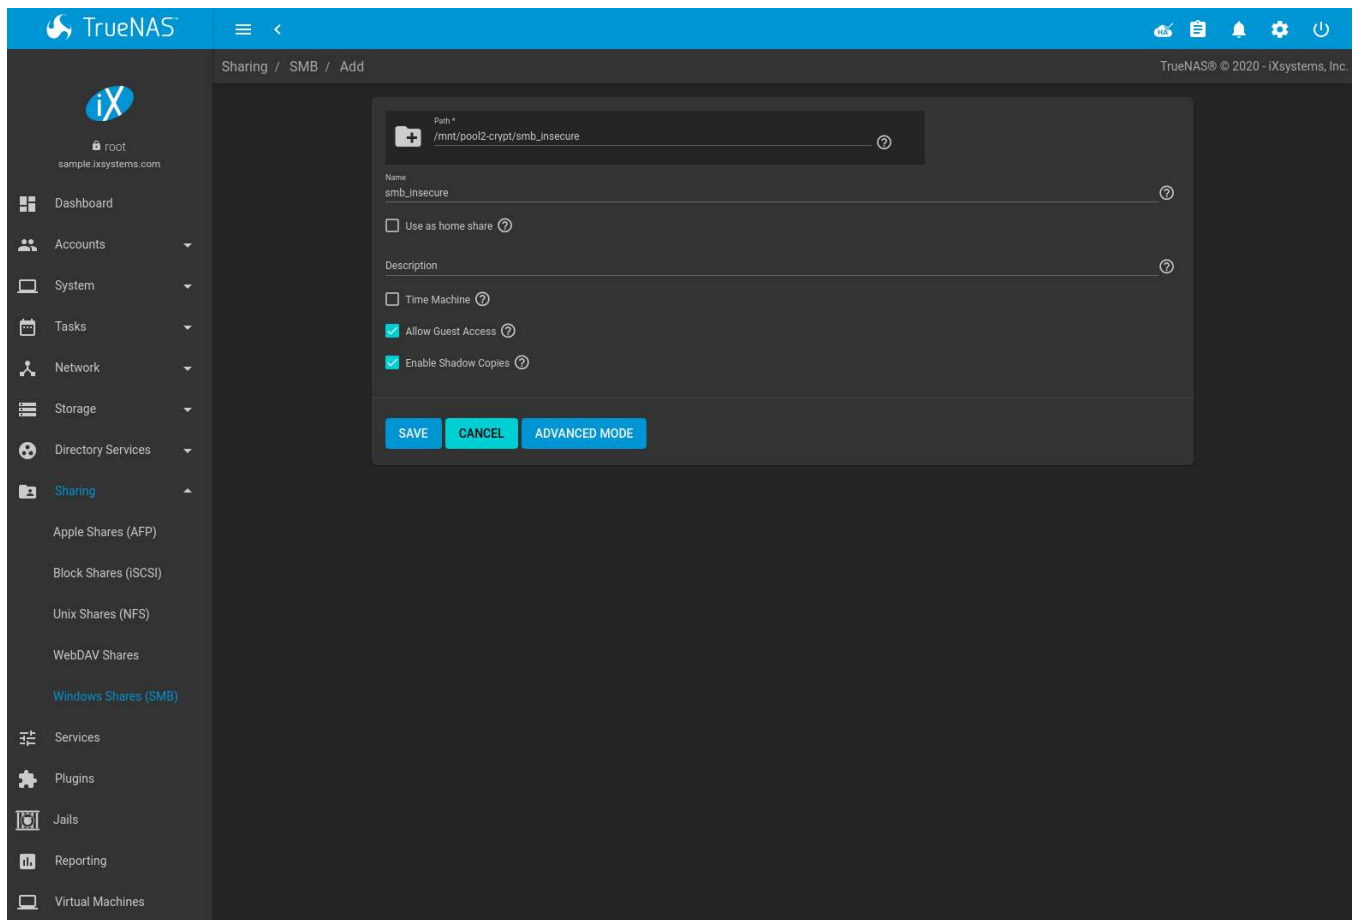


Fig. 11.23: Creating an Unauthenticated SMB Share

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

By default, users that access the share from an SMB client will not be prompted for a username or password. For example, to access the share from a Windows system, open Explorer and click on *Network*. In this example, a system named *FREENAS* appears with a share named *p2ds2-smb*. The user can copy data to and from this share.

The guest account can be changed by opening the *Services* → *SMB* options and selecting a different account from the *Guest Account* dropdown.

The guest account can also have an [Access Control Entry \(ACE\)](#) (page 148) that governs the permissions of the guest account to access the different pools and datasets on the system. To change the guest account permissions, edit the dataset Access Control List (ACL) and add a new item with the *Who* set to *User* and *User* set to the account used for guest access (*nobody* by default). The ACE can then be adjusted to define the access level required for guest sessions. See [ACL Management](#) (page 148) for more details about each available setting.

Changing the Guest Account permissions will not grant access for anonymous sessions. This is best accomplished by creating or editing the *everyone@* ACE in the dataset ACL. Note that anonymous sessions also do not have the guest SID in the security token.

11.5.2 Configuring Authenticated Access With Local Users

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, because there is no domain controller to provide authentication for the network, each user account must be created on

the TrueNAS[®] system. This type of configuration scenario is often used in home and small networks as it does not scale well if many user accounts are needed.

To configure authenticated access for an SMB share, first create a [group](#) (page 24) for all the SMB user accounts in TrueNAS[®]. Go to *Accounts* → *Groups* and click *ADD*. Use a descriptive name for the group like `local_smb_users`.

Configure the SMB share dataset with permissions for this new group. When [creating a new dataset](#) (page 142), set the *Share Type* to *SMB*. After the dataset is created, open the dataset [Access Control List \(ACL\)](#) (page 148) and add a new entry. Set *Who* to *Group* and select the SMB group for the *Group*. Finish [defining the permissions](#) (page 151) for the SMB group. Any [members of this group](#) (page 24) now have access to the dataset.

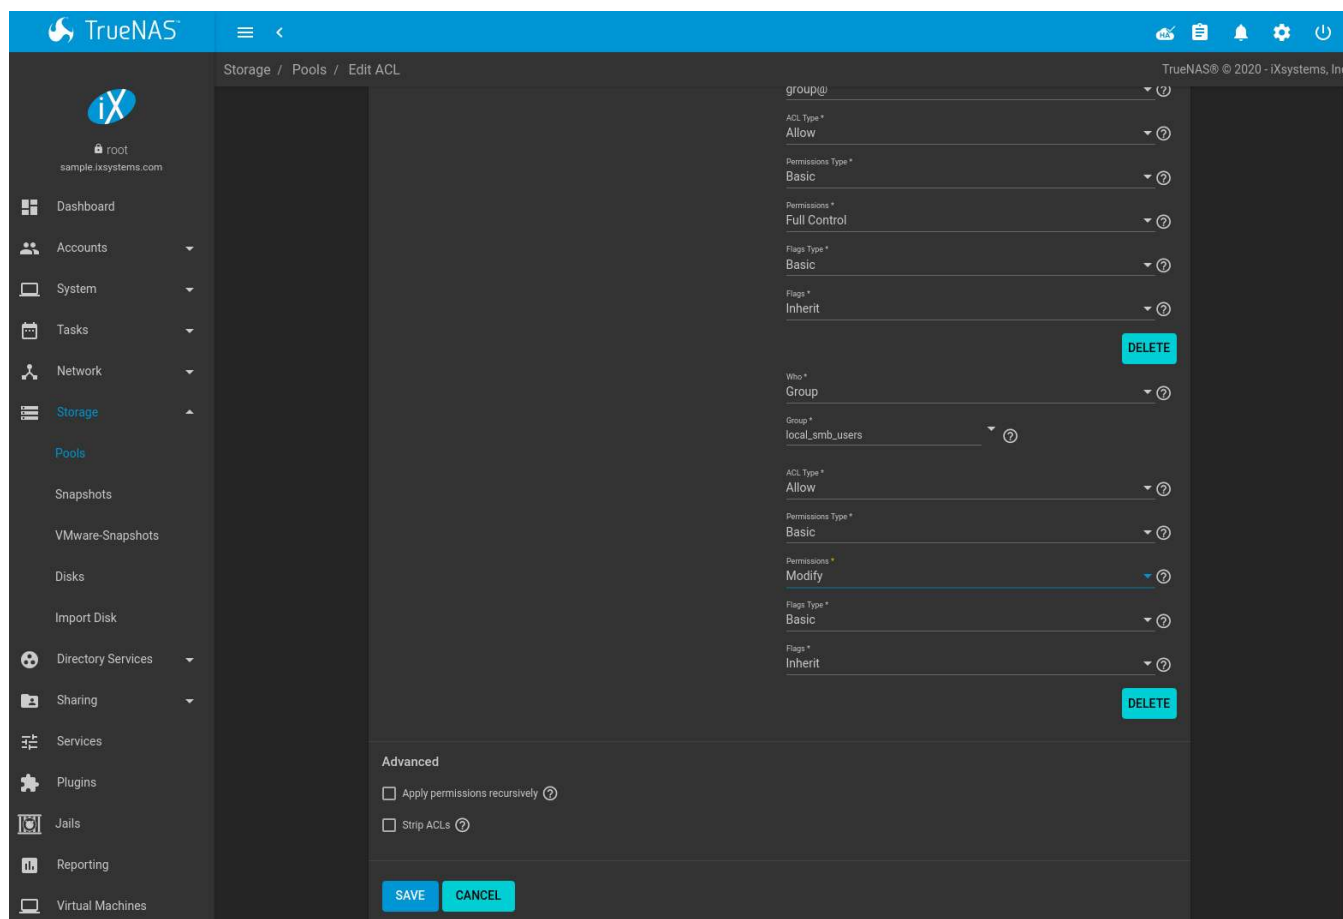


Fig. 11.24: Defining Permissions for a Group

Determine which users need authenticated access to the dataset and [create new accounts](#) (page 27) in TrueNAS[®]. It is recommended to use the same username and password from the client system for the associated TrueNAS[®] user account. Add the SMB group to the *Auxiliary Groups* list during account creation.

Finally, [create the SMB share](#) (page 209). Make sure the *Path* is pointed to the dataset that has defined permissions for the SMB group and that the [SMB](#) (page 240) service is active.

Testing the Share

The authenticated share can be tested from any SMB client. For example, to test an authenticated share from a Windows system with network discovery enabled, open Explorer and click on *Network*. If network discovery is disabled, open Explorer and enter `\HOST` in the address bar, where *HOST* is the IP address or hostname of the share system. This example shows a system named *FREENAS* with a share named *smb_share*.

After clicking *smb_share*, a Windows Security dialog prompts for the username and password of the user associated with *smb_share*. After authenticating, the user can copy data to and from the SMB share.

Map the share as a network drive to prevent Windows Explorer from hanging when accessing the share. Right-click the share and select *Map network drive...*. Choose a drive letter from the drop-down menu and click *Finish*.

Windows caches user account credentials with the authenticated share. This sometimes prevents connection to a share, even when the correct username and password are provided. Logging out of Windows clears the cache. The authentication dialog reappears the next time the user connects to an authenticated share.

11.5.3 User Quota Administration

File Explorer can manage quotas on SMB shares connected to an *Active Directory* (page 165) server. Both the share and dataset being shared must be configured to allow this feature:

- Create an authenticated share with `domain admins` as both the user and group name in *Ownership*.
- Edit the SMB share and add *ixnas* to the list of selected *VFS Object* (page 214).
- In Windows Explorer, connect to and map the share with a user account which is a member of the `domain admins` group. The *Quotas* tab becomes active.

11.5.4 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 TrueNAS®, it is automatically configured to support shadow copies.

Before using shadow copies with TrueNAS®, 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 TrueNAS® 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 With Local Users* (page 216) to create the desired number of shares.

To enable shadow copies, check the *Enable Shadow Copies* setting when creating an *smb share* (page 209).

11.6 Creating Authenticated and Time Machine Shares

macOS includes the *Time Machine* (<https://support.apple.com/en-us/HT201250>) feature which performs automatic backups. TrueNAS® supports Time Machine backups for both *SMB* (page 209) and *AFP* (page 178) shares.

The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

Create Time Machine or authenticated shares on a [new dataset](#) (page 142).

Change permissions on the new dataset by going to *Storage* → *Pools*. Select the dataset, click **:** (Options), *Change Permissions*.

Enter these settings:

1. **User:** Use the drop-down to select the desired user account. If the user does not yet exist on the TrueNAS® system, create one with *Accounts* → *Users*. See [users](#) (page 27) for more information.
2. **Group:** Select the desired group name. If the group does not yet exist on the TrueNAS® system, create one with *Accounts* → *Groups*. See [groups](#) (page 24) for more information.
3. Click *SAVE*.

Create the authenticated or Time Machine share:

1. Go to *Sharing* → *Windows (SMB) Shares* or *Sharing* → *Apple (AFP) Shares* and click *ADD*. Apple [deprecated the AFP protocol](https://support.apple.com/en-us/HT207828) (<https://support.apple.com/en-us/HT207828>) and recommends using SMB.
2. *Browse* to the dataset created for the share.
3. When creating a Time Machine share, set the *Time Machine* option.
4. Fill out the other required fields.
5. Click *SAVE*.

When creating multiple authenticated or Time Machine shares, repeat this process for each user. [Figure 11.25](#) shows creating a Time Machine Share in *Sharing* → *Apple (AFP) Shares*.




Fig. 11.25: Creating an Authenticated or Time Machine Share

Configuring a quota for each Time Machine share helps prevent backups from using all available space on the TrueNAS® system. Time Machine waits two minutes before creating a full backup. It then creates ongoing hourly, daily, weekly, and monthly backups. **The oldest backups are deleted when a Time Machine share fills up, so make sure that the quota size is large enough to hold the desired number of backups.** Note that a default installation of macOS is over 20 GiB.

Configure 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/) or create individual share quotas.

11.6.1 Setting SMB and AFP Share Quotas

SMB Quota

Go to *Sharing* → *Windows (SMB) Shares*, click **⋮** (Options) on the Time Machine share, and *Edit*. Click *Advanced Mode* and enter a `vfs_fruit(8)` (https://www.samba.org/samba/docs/current/man-html/vfs_fruit.8.html) parameter in the *Auxiliary Parameters*. Time Machine quotas use the `fruit:time machine max size` parameter. For example, to set a quota of 500 GiB, enter `fruit:time machine max size = 500 G`.

AFP Quota

Go to *Sharing* → *Apple (AFP) Shares*, click **⋮** (Options) on the Time Machine share, and *Edit*. In the example shown in Figure 11.26, the Time Machine share name is *backup_user1*. Enter a value in the *Time Machine Quota* field, and click *SAVE*. In this example, the Time Machine share is restricted to 200 GiB.

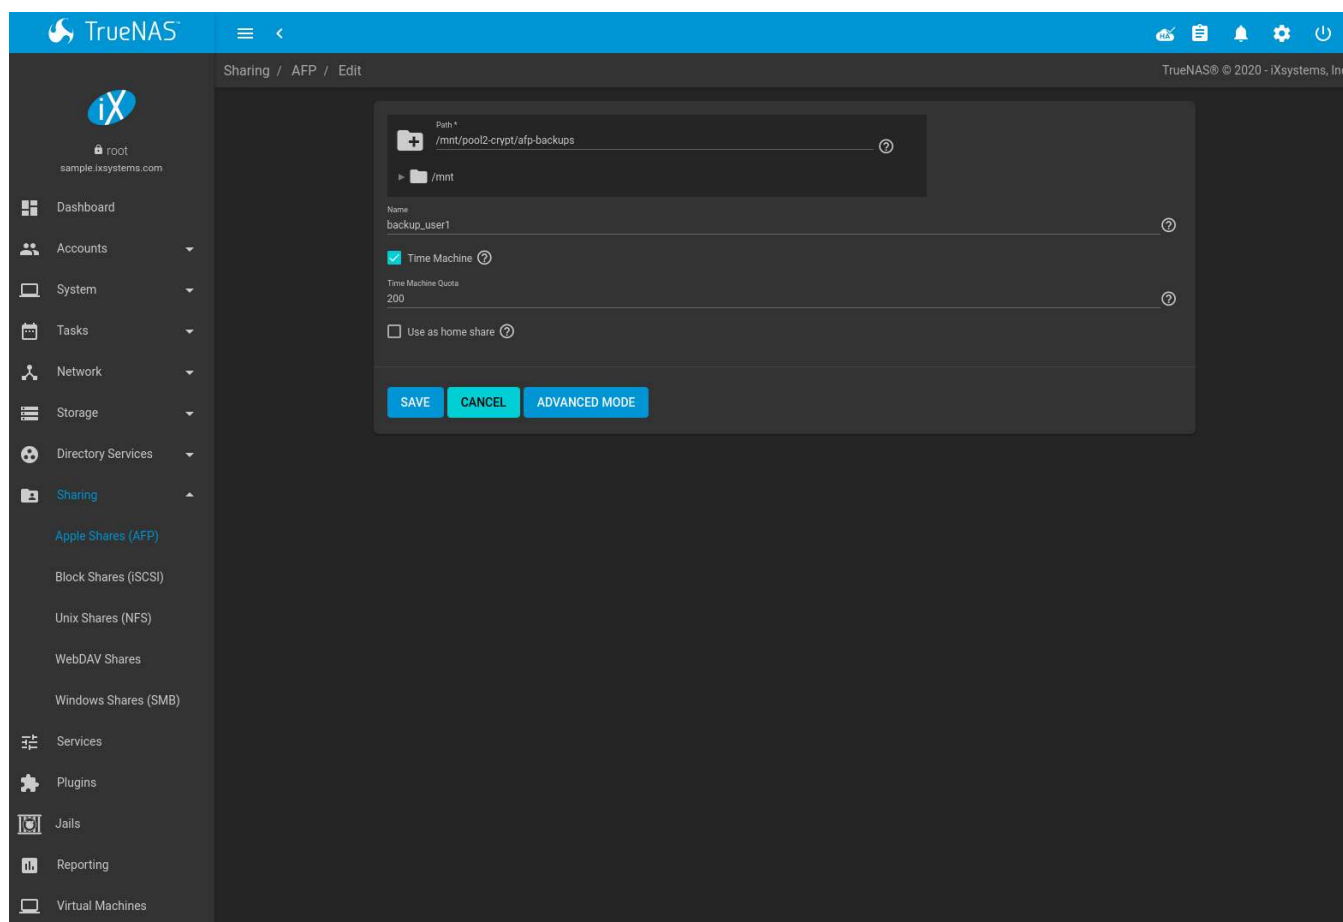


Fig. 11.26: Setting an AFP Share Quota

11.6.2 Client Time Machine Configuration

Note: The example shown here is intended to show the general process of adding a TrueNAS® share in Time Machine. The example might not reflect the exact process to configure Time Machine on a specific version of macOS. See the [Apple documentation](https://support.apple.com/en-us/HT201250) (<https://support.apple.com/en-us/HT201250>) for detailed Time Machine configuration instructions.

To configure Time Machine on the macOS client, go to *System Preferences* → *Time Machine*, and click *ON* in the left panel.



Fig. 11.27: Configuring Time Machine on macOS

Click *Select Disk...* in the right panel to find the TrueNAS® system with the share. Highlight the share and click *Use Backup Disk*. A connection dialog prompts to log in to the TrueNAS® system.

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 TrueNAS® 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.

SERVICES

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

- [AFP](#) (page 223)
- [Dynamic DNS](#) (page 225)
- [FTP](#) (page 226)
- [iSCSI](#) (page 231)
- [LLDP](#) (page 231)
- [NFS](#) (page 232)
- [Rsync](#) (page 234)
- [S3](#) (page 237)
- [S.M.A.R.T.](#) (page 238)
- [SMB](#) (page 240)
- [SNMP](#) (page 242)
- [SSH](#) (page 244)
- [TFTP](#) (page 247)
- [UPS](#) (page 248)
- [WebDAV](#) (page 251)

This section demonstrates starting a TrueNAS® service and the available configuration options for each TrueNAS® 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.) (<https://en.wikipedia.org/wiki/S.M.A.R.T.>). Other services default to *off* until started.

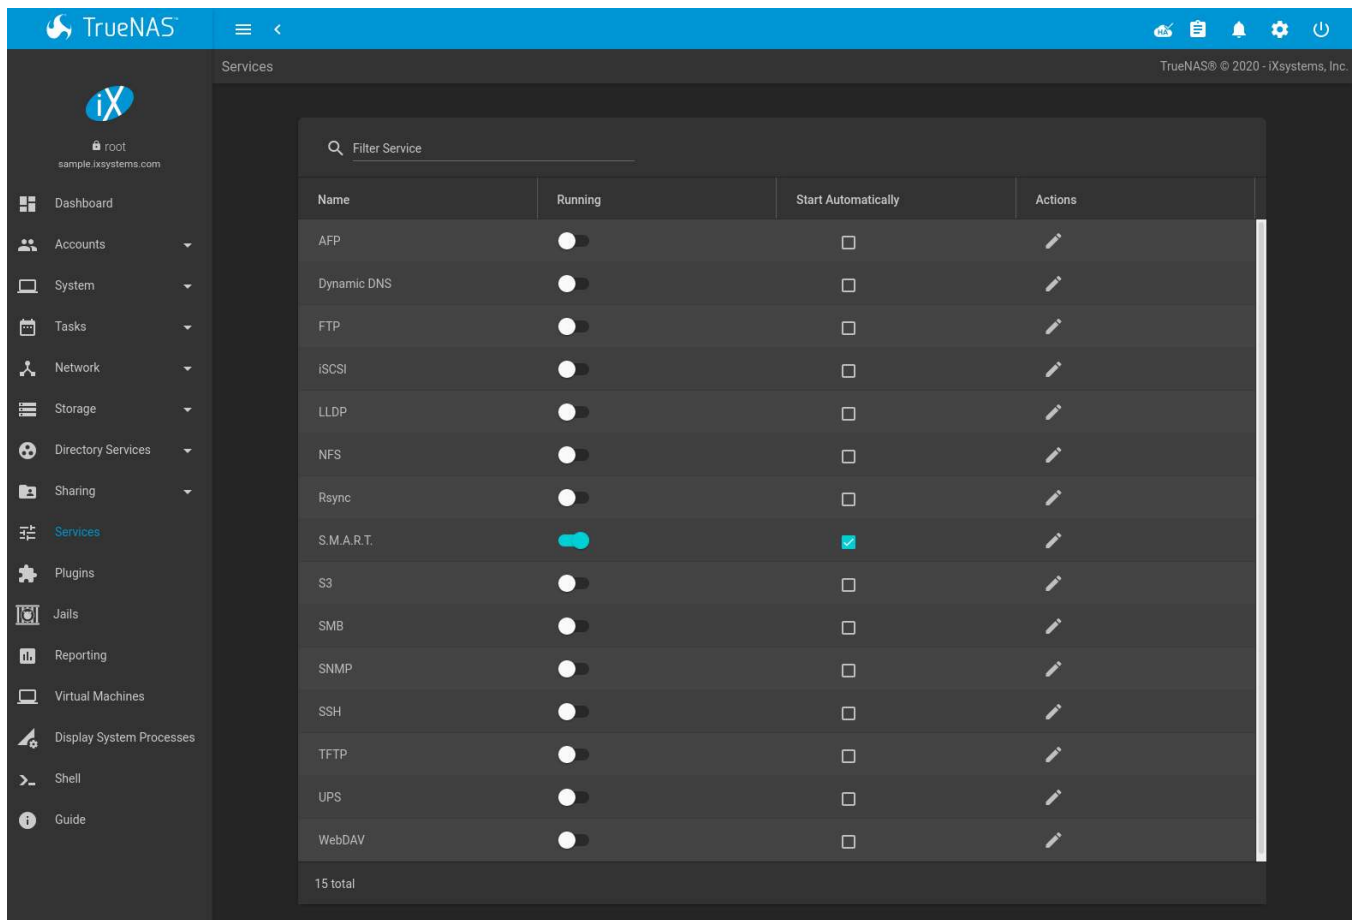


Fig. 12.1: Configure Services

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

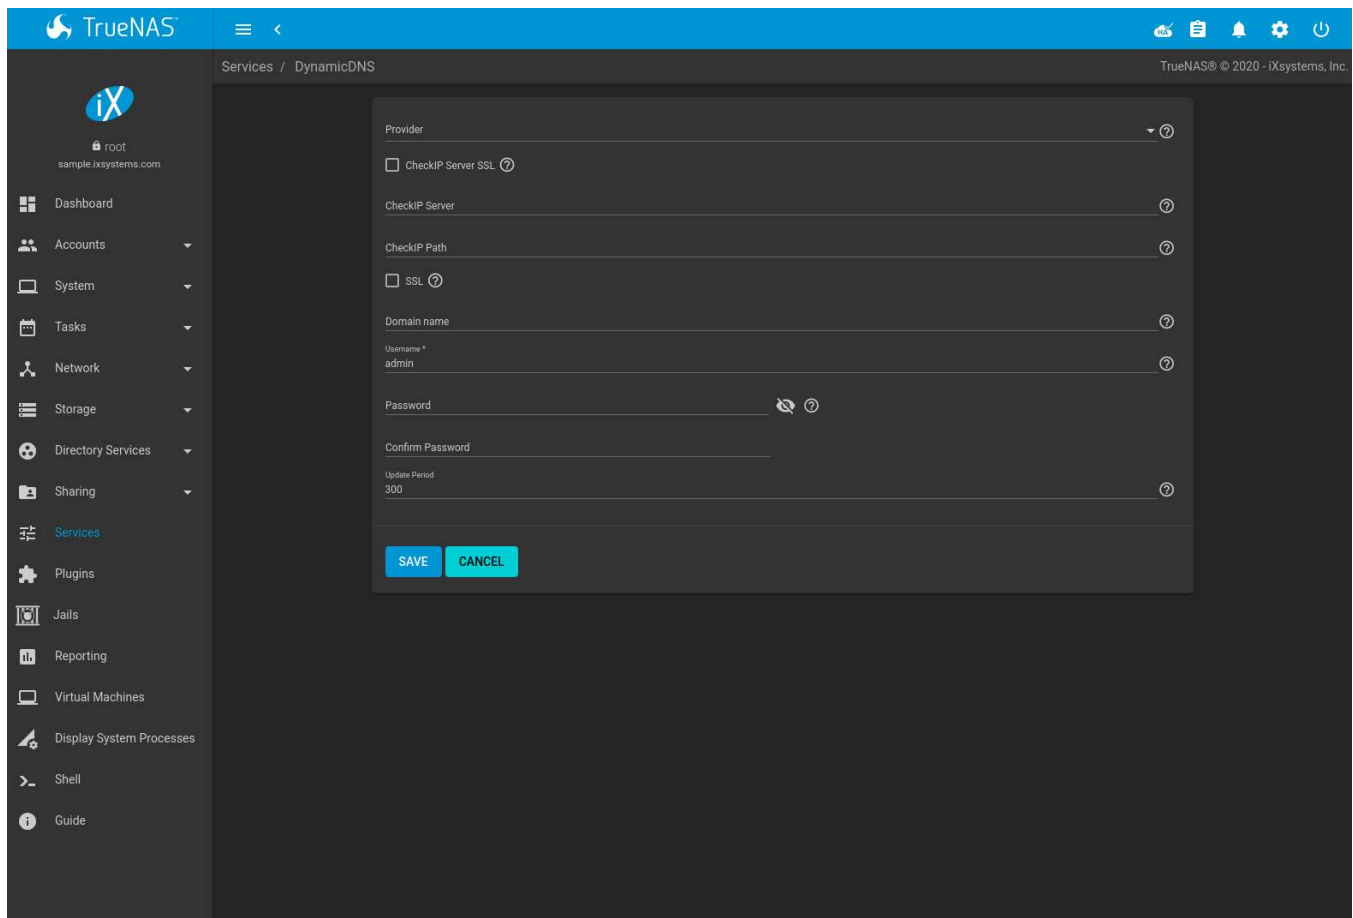
```
dbd -rf /path/to/share
```

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 Dynamic DNS

Dynamic DNS (DDNS) is useful if the TrueNAS® 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 TrueNAS® 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 12.3 shows the DDNS configuration screen and Table 12.2 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*.



The screenshot shows the TrueNAS web interface for configuring Dynamic DNS. The left sidebar contains navigation links: Dashboard, Accounts, System, Tasks, Network, Storage, Directory Services, Sharing, Services (highlighted), Plugins, Jails, Reporting, Virtual Machines, Display System Processes, Shell, and Guide. The main content area is titled 'Services / DynamicDNS'. The configuration form includes the following fields: 'Provider' (a dropdown menu), 'CheckIP Server SSL' (a checkbox), 'CheckIP Server' (a text field), 'CheckIP Path' (a text field), 'SSL' (a checkbox), 'Domain name' (a text field), 'Username *' (a text field with the value 'admin'), 'Password' (a text field with a toggle icon), 'Confirm Password' (a text field), and 'Update Period' (a text field with the value '300'). At the bottom of the form are 'SAVE' and 'CANCEL' buttons. The top right of the interface shows the TrueNAS logo, a hamburger menu, and system status icons. The bottom right corner of the interface displays the copyright notice 'TrueNAS® © 2020 - iXsystems, Inc.'

Fig. 12.3: Configuring DDNS

Table 12.2: 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	checkbox	Use HTTPS for the connection to the <i>CheckIP Server</i> .
CheckIP Server	string	Name and port of the server that reports the external IP address. For example, entering <code>checkip.dyndns.org:80</code> uses Dyn IP detection (https://help.dyn.com/remote-access-api/checkip-tool/) to discover the remote socket IP address.
CheckIP Path	string	Path to the <i>CheckIP Server</i> . For example, <i>no-ip.com</i> uses a <i>CheckIP Server</i> of <code>dynamic.zoneedit.com</code> and <i>CheckIP Path</i> of <code>/checkip.html</code> .
SSL	checkbox	Use HTTPS for the connection to the server that updates the DNS record.
Custom Server	string	DDNS server name. For example, <code>members.dyndns.org</code> denotes a server similar to <code>dyndns.org</code> .
Custom Path	string	DDNS server path. Path syntax varies by provider and must be obtained from that provider. For example, <code>/update?hostname=</code> is a simple path for the <code>update.twodns.de</code> <i>Custom Server</i> . The hostname is automatically appended by default. More examples are in the In-A-Dyn documentation (https://github.com/troglobit/inadyn#custom-ddns-providers).
Domain name	string	Fully qualified domain name of the host with the dynamic IP address. Separate multiple domains with a space, comma (,), or semicolon (;). Example: <code>myname.dyndns.org</code> ; <code>myothername.dyndns.org</code>
Username	string	Username for logging in to the provider and updating the record.
Password	string	Password for logging in to the provider and updating the record.
Update period	integer	How often the IP is checked in seconds.

When using the `he.net` *Provider*, 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*.

12.4 FTP

TrueNAS® 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 TrueNAS® 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 231).

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 12.4](#) 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*.

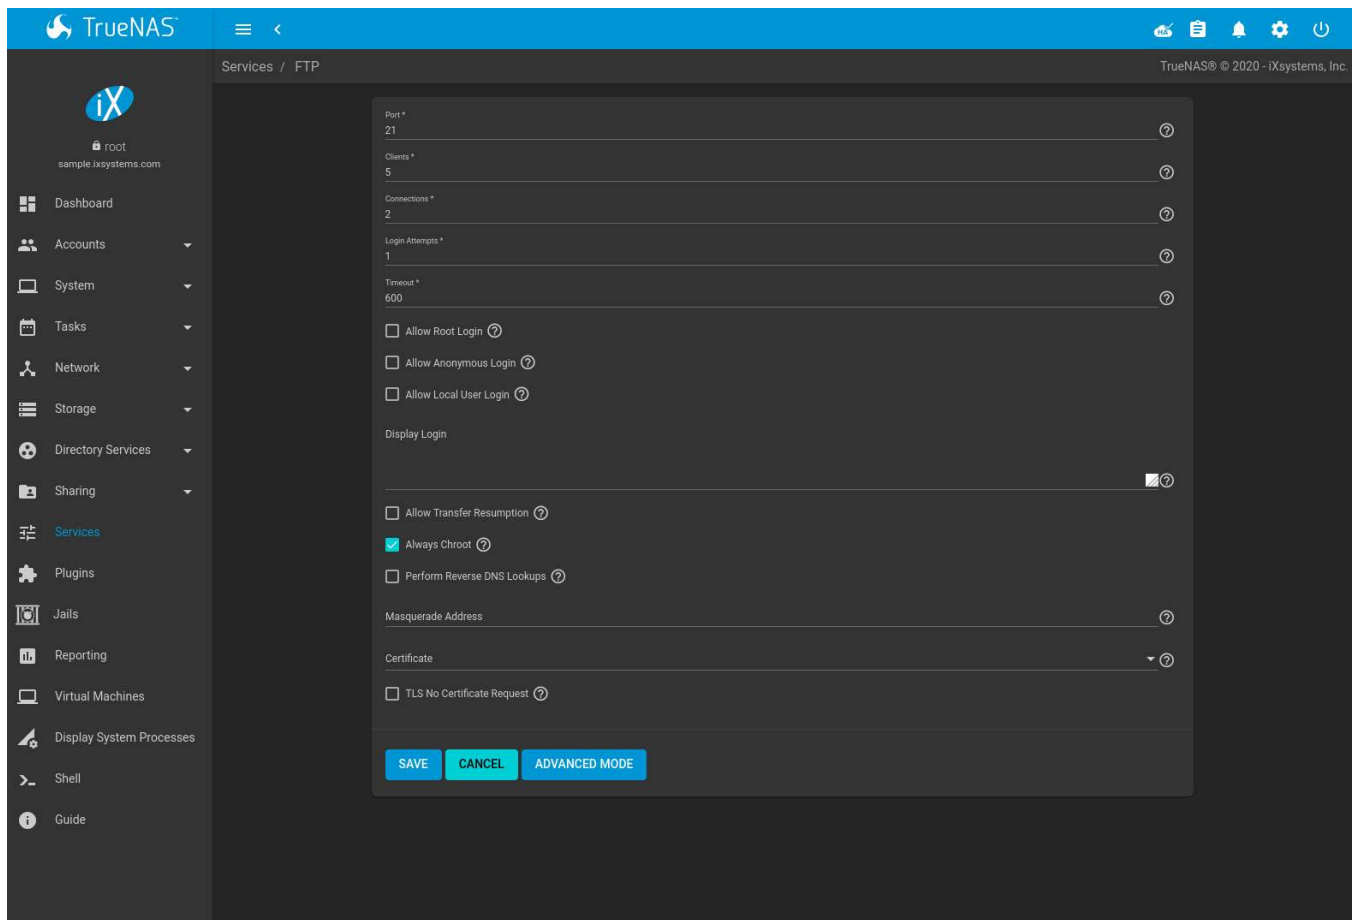


Fig. 12.4: Configuring FTP

Table 12.3 summarizes the available options when configuring the FTP server.

Table 12.3: 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		Allow anonymous FTP logins with access to the directory specified in the <i>Path</i> .
Path	browse button		Set the root directory for anonymous FTP connections.
Allow Local User Login	checkbox		Allow any local user to log in. By default, only members of the <code>ftp</code> group are allowed to log in.
Display Login	string		Specify the message displayed to local login users after authentication. Not displayed to anonymous login users.

Continued on next page

Table 12.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Allow Transfer Re- sumption	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 ad- dress	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 Permis- sion	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 Au- thentication	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 Down- load 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 Band- width	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 76).
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.

Continued on next page

Table 12.3 – continued from previous page

Setting	Value	Advanced Mode	Description
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.
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) 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>
```

12.4.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 12.5](#), The user has entered this information into the Filezilla client:

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




Fig. 12.5: 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 TrueNAS® system).

12.4.2 FTP in chroot

If users are required to authenticate before accessing the data on the TrueNAS® system, either create a user account for each user or import existing user accounts using [Active Directory](#) (page 165) or [LDAP](#) (page 170). 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 165) or [LDAP](#) (page 170) 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® system).

12.4.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

1. Import or create a certificate authority using the instructions in [CAs](#) (page 72). Then, import or create the certificate to use for encrypted connections using the instructions in [Certificates](#) (page 76).
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 TrueNAS® 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 TrueNAS® system. Click *SAVE* to accept the certificate and negotiate an encrypted connection.
4. To force encrypted connections, select *On* for the *TLS Policy*.

12.4.4 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 302) 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 TrueNAS® 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 TrueNAS® system hostname and IP address, or add an entry for the IP address of the TrueNAS® system in the *Network* → *Global Configuration Host name database* field.

12.5 iSCSI

Refer to [Block \(iSCSI\)](#) (page 183) 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 302) and type `ctladm islist` to determine the names of the connected initiators.

12.6 LLDP

The Link Layer Discovery Protocol (LLDP) is used by network devices to advertise their identity, capabilities, and neighbors on an Ethernet network. TrueNAS® uses the `ladvd` (<https://github.com/sspan/ladvd>) LLDP implemen-

tation. If the network contains managed switches, configuring and starting the LLDP service will tell the TrueNAS® system to advertise itself on the network.

Figure 12.6 shows the LLDP configuration screen and Table 12.4 summarizes the configuration options for the LLDP service.

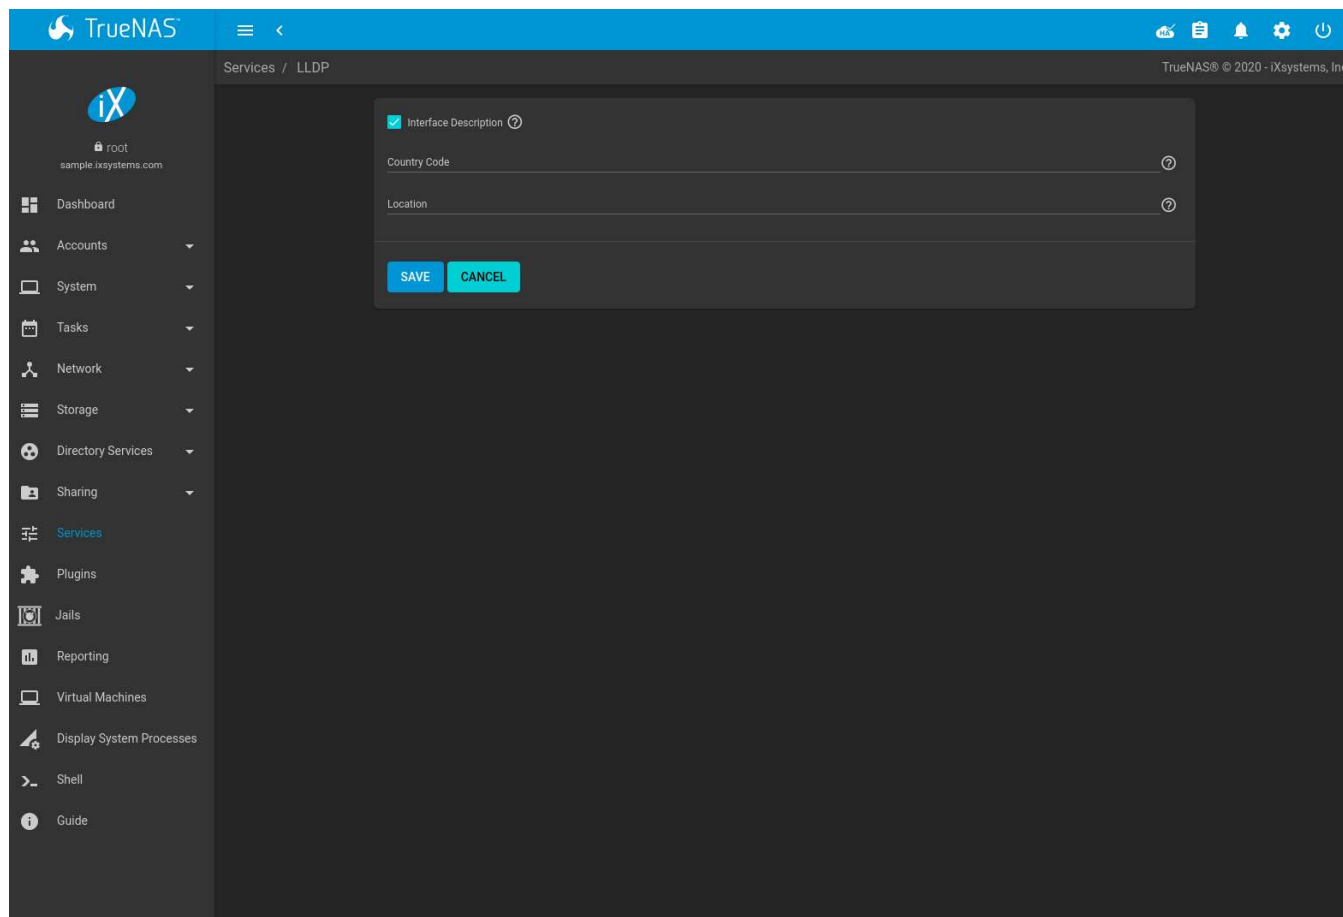


Fig. 12.6: Configuring LLDP

Table 12.4: 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.

12.7 NFS

The settings that are configured when creating NFS shares in 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 12.7 shows the configuration screen and Table 12.5 summarizes the configuration options for the NFS service.

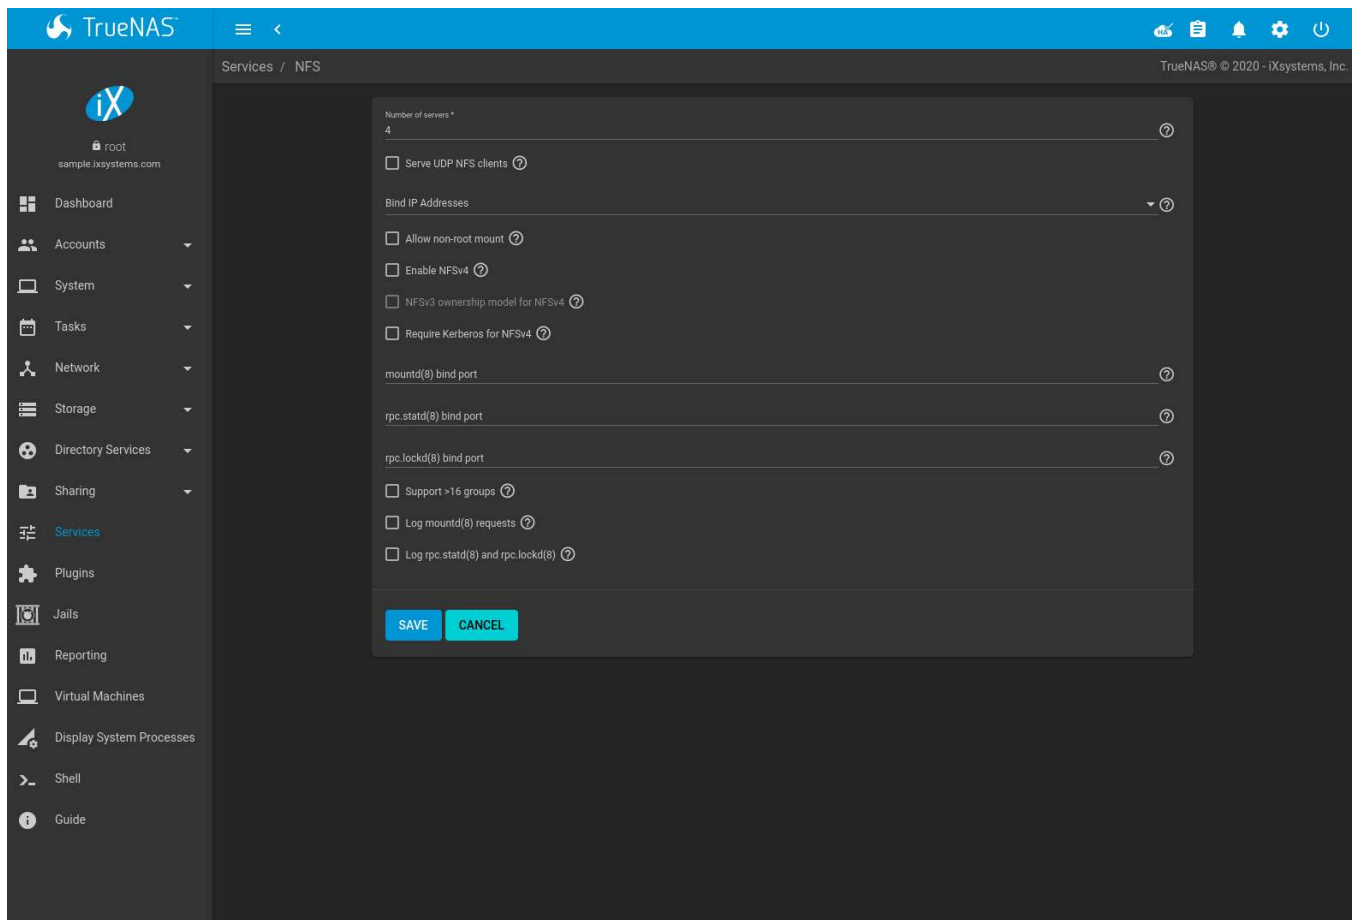


Fig. 12.7: Configuring NFS

Table 12.5: 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. Disabling this option allows using either default NFS or Kerberos authentication.

Continued on next page

Table 12.5 – continued from previous page

Setting	Value	Description
mountd(8) bind port	integer	Optional. Specify the port that mountd(8) (https://www.freebsd.org/cgi/man.cgi?query=mountd) binds to.
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.

12.8 Rsync

Services → *Rsync* is used to configure an rsync server when using rsync module mode. Refer to [Rsync Module Mode](#) (page 93) for a configuration example.

This section describes the configurable options for the `rsyncd` service and rsync modules.

12.8.1 Configure Rsyncd

To configure the `rsyncd` server, go to *Services* and click  *EDIT* for the *Rsync* service.





Fig. 12.8: Rsyncd Configuration

Table 12.6 summarizes the configuration options for the rsync daemon:

Table 12.6: Rsyncd Configuration Options

Setting	Value	Description
TCP Port	integer	<code>rsyncd</code> listens on this port. The default is 873.
Auxiliary parameters	string	Enter any additional parameters from <code>rsyncd.conf(5)</code> (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

12.8.2 Rsync Modules

To add a new Rsync module, go to *Services*, click  *EDIT* for the *Rsync* service, select the *Rsync Module* tab, and click *ADD*.




Fig. 12.9: Adding an Rsync Module

Table 12.7 summarizes the configuration options available when creating a rsync module.

Table 12.7: Rsync Module Configuration Options

Setting	Value	Description
Name	string	Module name that matches the name requested by the rsync client.
Comment	string	Describe this module.
Path	file browser	Browse to the pool or dataset to store received data.
Access Mode	drop-down menu	Choose permissions for this rsync module.
Maximum connections	integer	Maximum connections to this module. 0 is unlimited.
User	drop-down menu	User to run as during file transfers to and from this module.
Group	drop-down menu	Group to run as during file transfers to and from this module.
Hosts Allow	string	From rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf). A list of patterns to match with the hostname and IP address of a connecting client. The connection is rejected if no patterns match. Separate patterns with whitespace or a comma.
Hosts Deny	string	From rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf). A list of patterns to match with the hostname and IP address of a connecting client. The connection is rejected when the patterns match. Separate patterns with whitespace or a comma.
Auxiliary parameters	string	Enter any additional parameters from rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

12.9 S3

S3 is a distributed or clustered filesystem protocol compatible with Amazon S3 cloud storage. The TrueNAS® S3 service uses [Minio](https://minio.io/) (<https://minio.io/>) to provide S3 storage hosted on the TrueNAS® system itself. Minio also provides features beyond the limits of the basic Amazon S3 specifications.

Figure 12.10 shows the S3 service configuration screen and Table 12.8 summarizes the configuration options. After configuring the S3 service, start it in *Services*.




Fig. 12.10: Configuring S3

Table 12.8: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	Enter the IP address to run the S3 service. <i>0.0.0.0</i> sets the server to listen on all addresses.
Port	string	Enter the TCP port on which to provide the S3 service. Default is <i>9000</i> .
Access Key	string	Enter the S3 access ID. See Access keys (https://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html#access-keys-and-secret-access-keys) for more information.
Secret Key	string	Enter the S3 secret access key. See Access keys (https://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html#access-keys-and-secret-access-keys) for more information.

Continued on next page

Table 12.8 – continued from previous page

Setting	Value	Description
Confirm Secret Key	string	Re-enter the S3 password to confirm.
Disk	browse	Directory where the S3 filesystem will be mounted. Ownership of this directory and all subdirectories is set to <i>minio:minio</i> . Create a separate dataset (page 142) for Minio to avoid issues with conflicting directory permissions or ownership.
Enable Browser	checkbox	Set to enable the web user interface for the S3 service. Access the minio web interface by entering the IP address and port number separated by a colon in the browser address bar.
Certificate	drop-down menu	Add the SSL certificate (page 76) to be used for secure S3 connections.

12.10 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 111).

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.

TrueNAS® 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, including disk temperature. 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 12.11](#) shows the configuration screen that appears after going to *Services* → *S.M.A.R.T* and clicking  (Configure).




Fig. 12.11: 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 12.9 summarizes the options in the S.M.A.R.T configuration screen.

Table 12.9: 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 only performed when <i>Never</i> is selected. 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.
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.

12.11 SMB

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) (<https://www.samba.org/samba/docs/old/Samba3-HOWTO/NetworkBrowsing.html#id2581357>) to occur and for the TrueNAS® system to become available in Windows Explorer.

Figure 12.12 shows the global configuration options which apply to all SMB shares. This configuration screen displays the configurable options from [smb4.conf](https://www.freebsd.org/cgi/man.cgi?query=smb4.conf) (<https://www.freebsd.org/cgi/man.cgi?query=smb4.conf>). These options are described in Table 12.10.

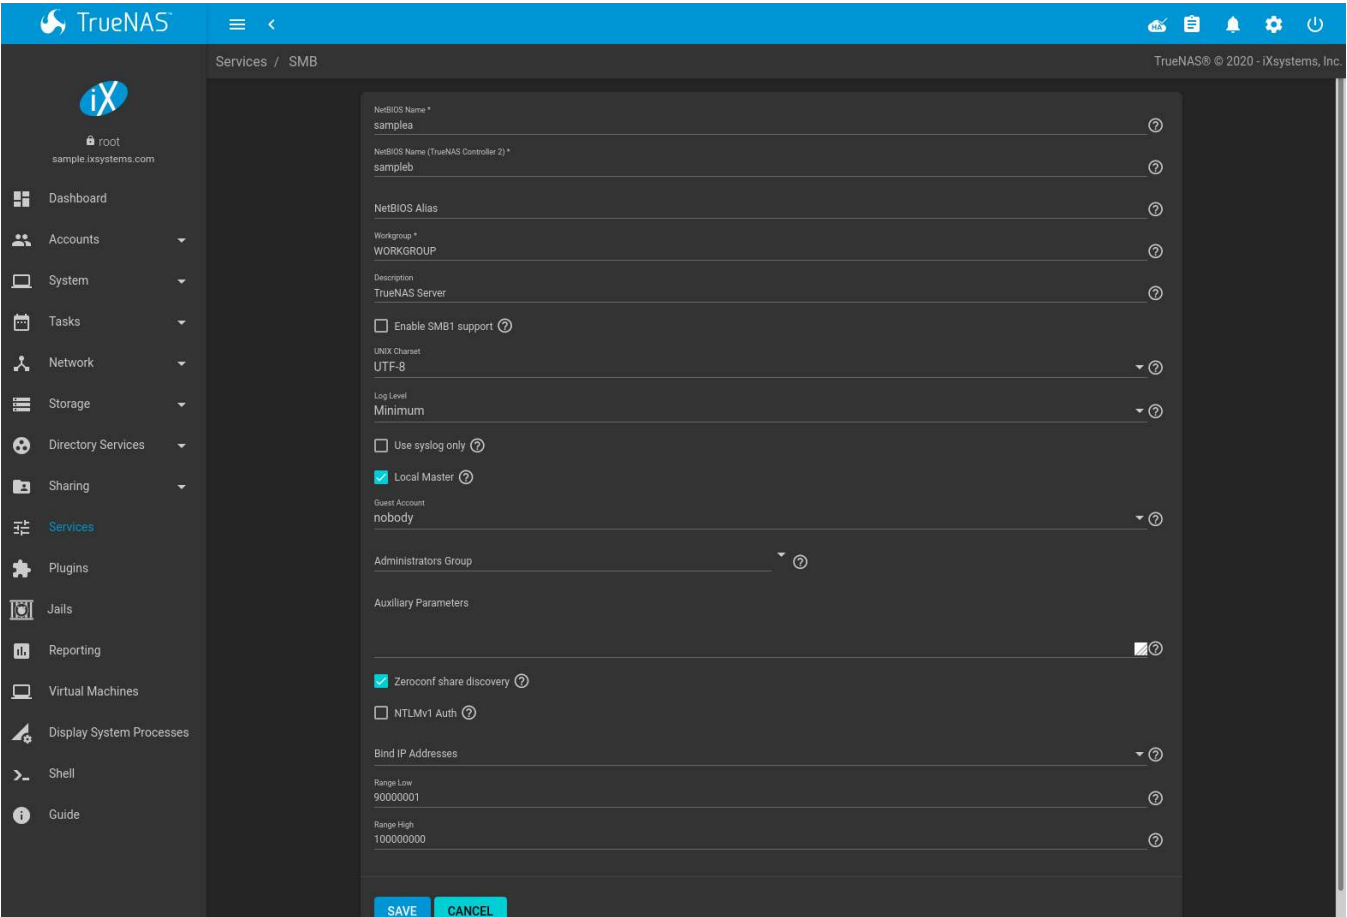


Fig. 12.12: Global SMB Configuration

Table 12.10: Global SMB Configuration Options

Setting	Value	Description
NetBIOS Name	string	Automatically populated with the active TrueNAS controller host-name from the Global Configuration (page 119). Limited to 15 characters. It must be different from the <i>Workgroup</i> name.
NetBIOS Name (TrueNAS Controller 1/2)	string	Automatically populated with the standby TrueNAS controller host-name from the Global Configuration (page 119). Limited to 15 characters. When using Failover (page 81), set a unique NetBIOS name for the standby TrueNAS controller.

Continued on next page

Table 12.10 – continued from previous page

Setting	Value	Description
NetBIOS Alias	string	Limited to 15 characters. When using <i>Failover</i> (page 81), this is the NetBIOS name that resolves to either TrueNAS controller.
Workgroup	string	Must match the Windows workgroup name. This setting is ignored if the <i>Active Directory</i> (page 165) or <i>LDAP</i> (page 170) service is running.
Description	string	Enter a server description. Optional.
Enable SMB1 support	checkbox	Allow legacy SMB clients to connect to the server. Warning: SMB1 is not secure and has been deprecated by Microsoft. See <i>Do Not Use SMB1</i> (https://www.ixsystems.com/blog/library/do-not-use-smb1/).
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.
Guest Account	drop-down menu	Account to be used for guest access. Default is <i>nobody</i> . The chosen account is required to have permissions to the shared pool or dataset. To adjust permissions, edit the dataset Access Control List (ACL), add a new entry for the chosen guest account, and configure the permissions in that entry. If the selected <i>Guest Account</i> is deleted the field resets to <i>nobody</i> .
Administrators Group	drop-down menu	Members of this group are local admins and automatically have privileges to take ownership of any file in an SMB share, reset permissions, and administer the SMB server through the Computer Management MMC snap-in.
Auxiliary Parameters	string	Enter additional <code>smb.conf</code> (https://www.freebsd.org/cgi/man.cgi?query=smb.conf) options. See the <i>Samba Guide</i> (http://www.oreilly.com/openbook/samba/book/appb_02.html) for more information on the available settings. To log more details when a client attempts to authenticate to the share, add <code>log_level = 1, auth_audit:5</code> .
Zeroconf share discovery	checkbox	Enable if Mac clients will be connecting to the SMB share.
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	Static IP addresses which SMB listens on for connections. Leaving all unselected defaults to listening on all active interfaces.
Range Low	integer	Range Low and Range High set the range of UID/GID numbers which this IDMap backend translates. If an external credential like a Windows SID maps to a UID or GID number outside this range, the external credential is ignored.
Range High	integer	

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 240) cannot be disabled while *Active Directory* (page 165) is enabled.

12.11.1 Troubleshooting SMB

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 the SMB service will not start, run this command from *Shell* (page 302) to see if there is an error in the configuration:

```
testparm /usr/local/etc/smb4.conf
```

Using a dataset for SMB sharing is recommended. When creating the dataset, make sure that the *Share type* is set to *SMB*.

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 use the *ACL manager* (page 148).

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.**

- *Log Level* can also have a performance penalty. When not needed, it can be disabled or reduced in the *global SMB service options* (page 240).
- Create as SMB-style dataset and enable the `ixnas` auxiliary parameter
- Disable as many *VFS Objects* as possible in the *share settings* (page 210). Many have performance overhead.

12.12 SNMP

SNMP (Simple Network Management Protocol) is used to monitor network-attached devices for conditions that warrant administrative attention. TrueNAS® uses *Net-SNMP* (http://net-snmp.sourceforge.net/) to provide SNMP. When starting the SNMP service, this port will be enabled on the TrueNAS® system:

- UDP 161 (listens here for SNMP requests)

Available MIBS are located in `/usr/local/share/snmp/mibs`.

[Figure 12.13](#) shows the *Services* → *SNMP* → *Configure* screen. [Table 12.11](#) summarizes the configuration options.




Fig. 12.13: Configuring SNMP

Table 12.11: SNMP Configuration Options

Setting	Value	Description
Location	string	Enter the location of the system.
Contact	string	Enter an email address to receive messages from the SNMP service.
Community	string	Change from <i>public</i> to increase system security. Can only contain alphanumeric characters, underscores, dashes, periods, and spaces. This can be left empty for SNMPv3 networks.
SNMP v3 Support	checkbox	Set to enable support for SNMP version 3 (https://tools.ietf.org/html/rfc3410). See 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.
Username	string	Only applies if <i>SNMP v3 Support</i> is set. Enter a username to register with this service.
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	Enter a separate privacy passphrase. <i>Password</i> is used when this is left empty.

Continued on next page

Table 12.11 – 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.
Expose zillstat via SNMP	checkbox	Enabling this option may have pool performance implications.
Log Level	drop-down menu	Choose how many log entries to create. Choices range from the least log entries (Emergency) to the most (Debug).

[Zenoss](https://www.zenoss.com/) (<https://www.zenoss.com/>) provides a seamless monitoring service through SNMP for TrueNAS® called [TrueNAS ZenPack](https://www.zenoss.com/product/zenpacks/truenas) (<https://www.zenoss.com/product/zenpacks/truenas>).

12.13 SSH

Secure Shell (SSH) is used to transfer files securely over an encrypted network. When a TrueNAS® 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 TrueNAS® SSH configuration options, demonstrates an example configuration that restricts users to their home directory, and provides some troubleshooting tips.

[Figure 12.14](#) 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.

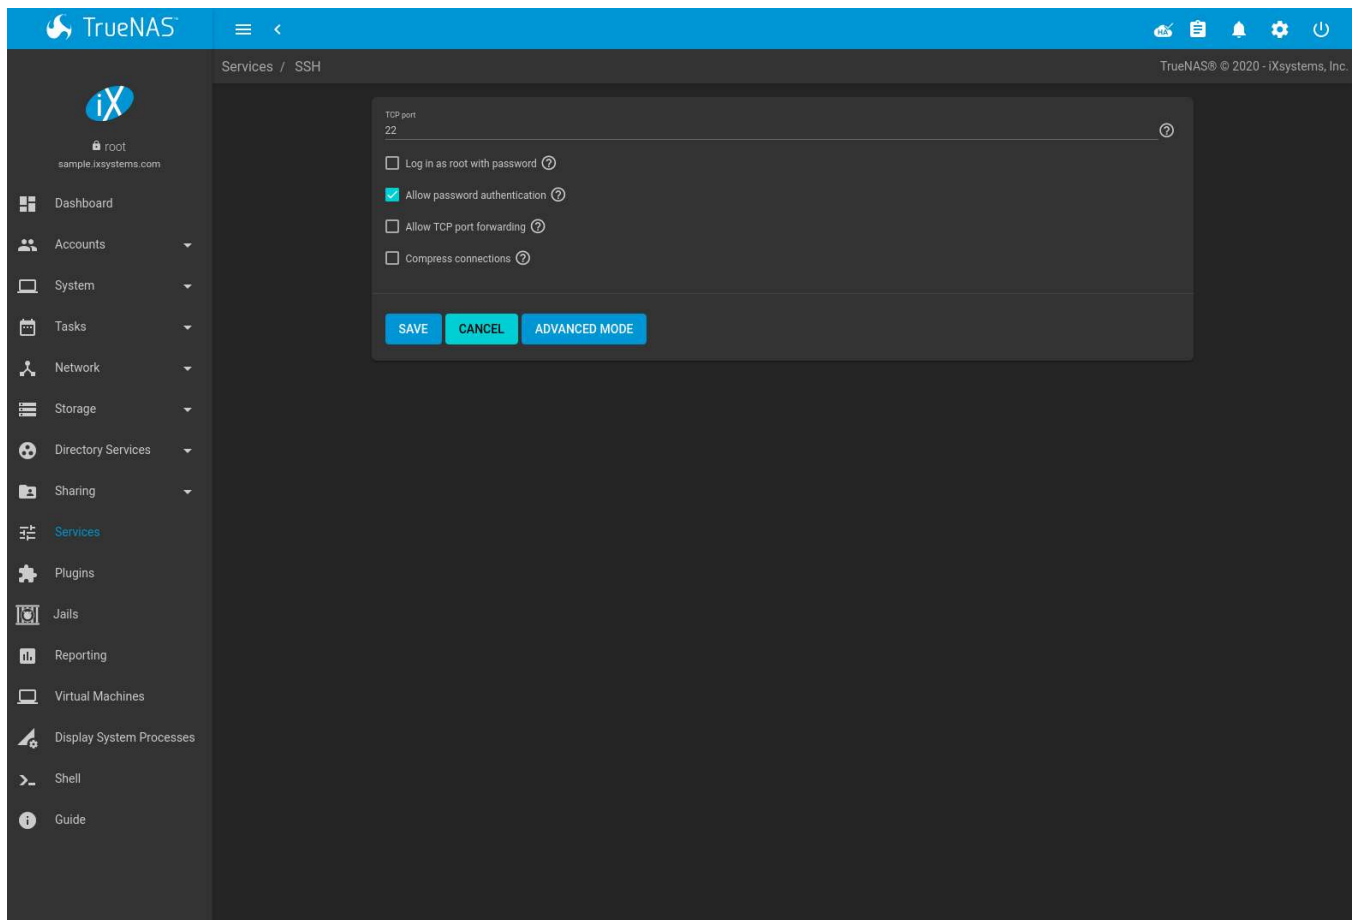


Fig. 12.14: SSH Configuration

Table 12.12 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 12.12: 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/html/doc/Chapter8.html) on both the SSH client and server.
Allow kerberos authentication	checkbox	✓	Ensure Kerberos Realms (page 174) and Kerberos Keytabs (page 175) are configured and TrueNAS® can communicate with the Kerberos Domain Controller (KDC) before enabling this option.

Continued on next page

Table 12.12 – continued from previous page

Setting	Value	Advanced Mode	Description
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.

Here are some recommendations for the *Extra options*:

- Add `NoneEnabled no` to disable the insecure `none` cipher.
- Increase the `ClientAliveInterval` if SSH connections tend to drop.
- `ClientMaxStartup` defaults to `10`. Increase this value when more concurrent SSH connections are required.

12.13.1 SCP Only

When SSH is configured, authenticated users with a user account can use `ssh` to log into the TrueNAS® 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 TrueNAS® 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 TrueNAS® 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`.

12.13.2 Troubleshooting SSH

Keywords listed in [sshd_config\(5\)](https://www.freebsd.org/cgi/man.cgi?query=sshd_config) (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 TrueNAS® 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 302) to read these messages as they occur:

```
tail -f /var/log/messages
```

Additional messages regarding authentication errors are found in `/var/log/auth.log`.

12.14 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 TrueNAS® 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 12.15 shows the TFTP configuration screen and Table 12.13 summarizes the available options.

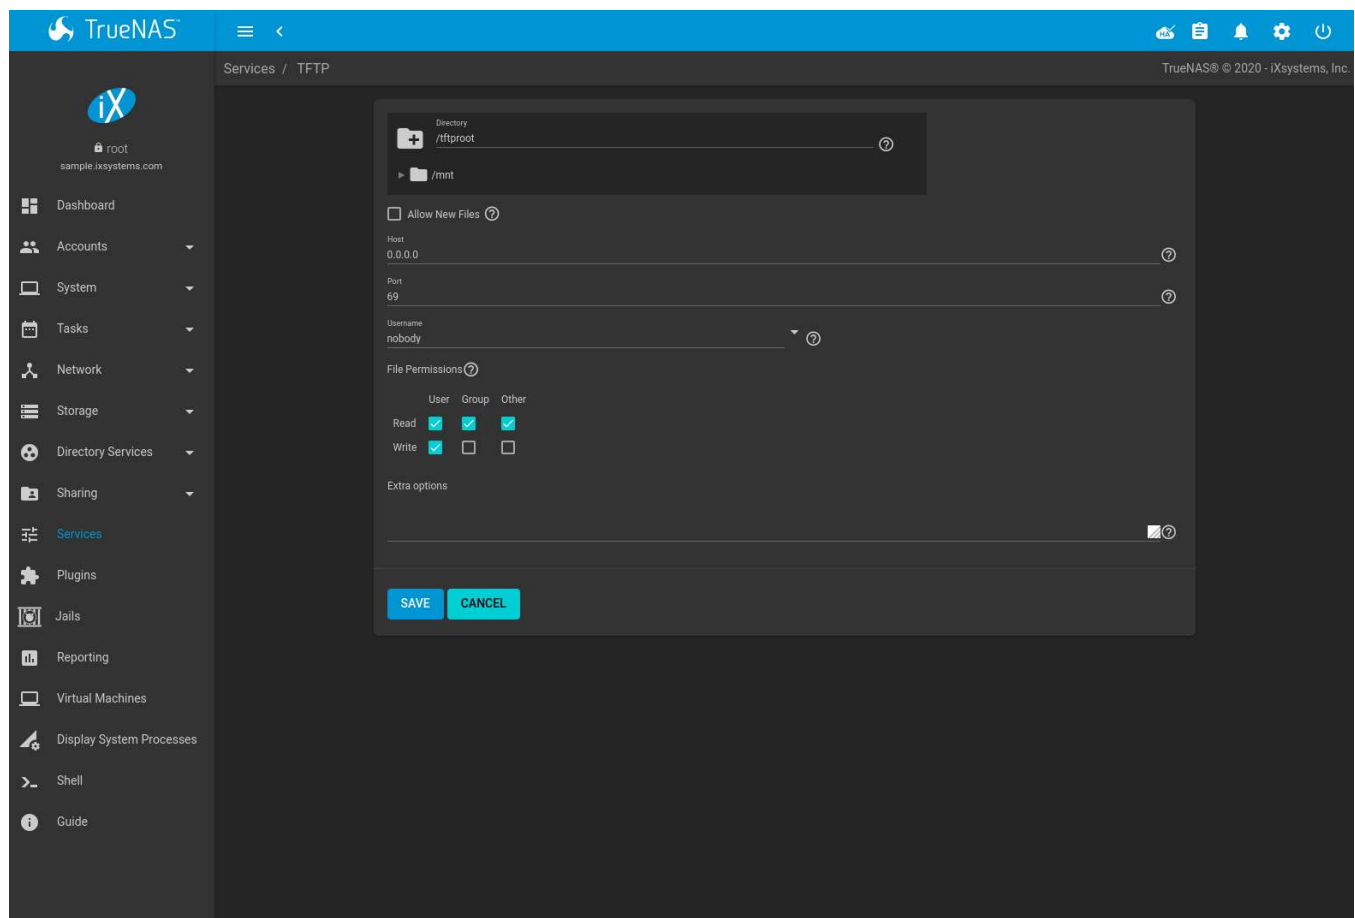


Fig. 12.15: TFTP Configuration

Table 12.13: 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.
Host	IP address	The default host to use for TFTP transfers. Enter an IP address. Example: <i>192.0.2.1</i> .
Port	integer	The UDP port number that listens for TFTP requests. Example: <i>8050</i> .
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.

12.15 UPS

TrueNAS® uses [NUT](https://networkupstools.org/) (Network UPS Tools) to provide UPS support. If the TrueNAS® system is connected to a UPS device, configure the UPS service in *Services* → *UPS* → *Configure*.

Figure 12.16 shows the UPS configuration screen:

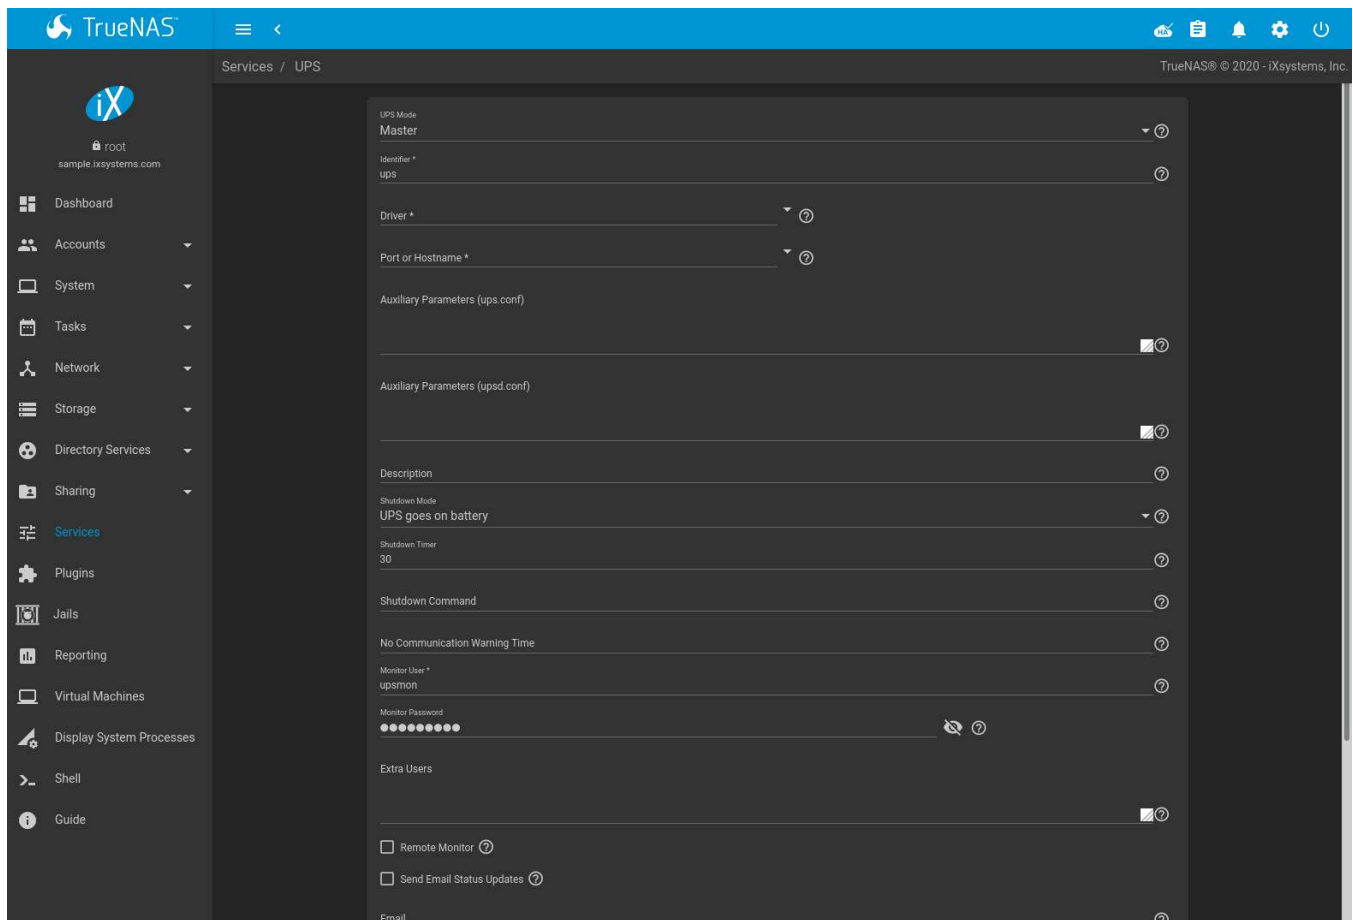


Fig. 12.16: UPS Configuration Screen

Table 12.14 summarizes the options in the UPS Configuration screen.

Table 12.14: 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	Required. Describe the UPS device. Can contain alphanumeric, period, comma, hyphen, and underscore characters.
Driver / Remote Host	combo-box	Required. For a list of supported devices, see the Network UPS Tools compatibility list (https://networkupstools.org/stable-hcl.html). The field suggests drivers based on the text entered. To search for a specific driver, begin typing the name of the driver. The search is case sensitive. The <i>Driver</i> field changes to <i>Remote Host</i> when <i>UPS Mode</i> is set to <i>Slave</i> . Enter the IP address of the system configured as the UPS <i>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.

Continued on next page

Table 12.14 – continued from previous page

Setting	Value	Description
Port or Hostname	drop-down menu	Serial or USB port connected to the UPS. To automatically detect and manage the USB port settings, open the drop-down menu and select <i>auto</i> . If the specific USB port must be chosen, see this note (page 250) about identifying the USB port used by the UPS. When an SNMP driver is selected, enter the IP address or hostname of the SNMP UPS device. <i>Port or Hostname</i> becomes <i>Remote Port</i> when the <i>UPS Mode</i> is set to <i>Slave</i> . Enter the open network port number of the UPS <i>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	Required. Enter a user to associate with this service. The recommended default user is <i>upsmon</i> .
Monitor Password	string	Required. Default is the known value <i>fixmepass</i> . Change this to enhance system security. Cannot contain a space or #.
Extra Users	string	Enter accounts that have 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 TrueNAS® 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 TrueNAS® system.
Host Sync	integer	Enter a time in seconds for UPSMON(8) (https://www.freebsd.org/cgi/man.cgi?query=upsmon) to wait in master mode for the slaves to disconnect during a shutdown.

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.

Some UPS models might be unresponsive with the default polling frequency. This can show in TrueNAS® logs as a recurring error like: `libusb_get_interrupt: Unknown error`.

If this error occurs, decrease 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 302) 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.

12.15.1 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>) and *NUT User Manual Pages* (https://networkupstools.org/docs/man/index.html#User_man).

12.16 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 207) 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 12.17* shows the WebDAV configuration screen. *Table 12.15* summarizes the available options.

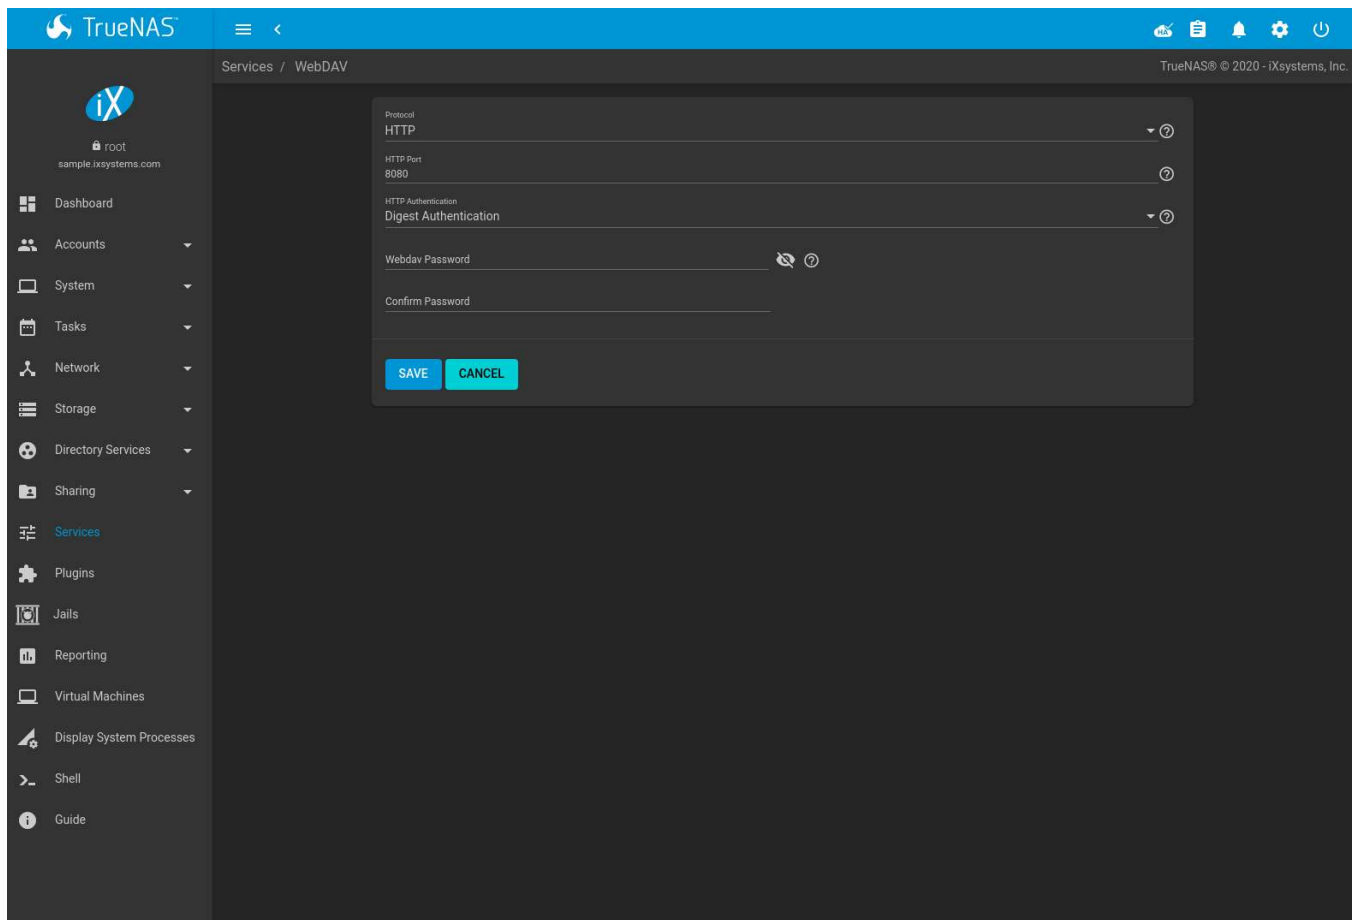


Fig. 12.17: WebDAV Configuration Screen

Table 12.15: 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.

Note: This is a TrueNAS® licensed feature only. For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 6:00AM to 6:00PM Pacific Standard Time: <ul style="list-style-type: none">• US-only toll-free: 855-473-7449 option 2• Local and international: 408-943-4100 option 2
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none">• US-only toll-free: 855-499-5131• International: 408-878-3140 (international calling rates will apply)

PLUGINS

TrueNAS® provides the ability to extend the built-in NAS services by providing two methods for installing additional software.

Plugins (page 254) 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) (https://en.wikipedia.org/wiki/Freebsd_jail) that cannot install additional software.

Jails (page 260) 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 254) and *Jails* (page 260) 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 260) must be configured before plugins are available on TrueNAS®. This means having a suitable *pool* (page 131) created to store plugins.

13.1 Installing Plugins

A plugin is a self-contained application installer designed to integrate into the TrueNAS® web interface. A plugin offers several advantages:

- the TrueNAS® web interface provides a browser for viewing the list of available plugins
- the TrueNAS® web interface provides buttons for installing, starting, managing, and uninstalling plugins
- if the plugin has configuration options, a management screen is added to the TrueNAS® web interface for these options to be configured

View available plugins by clicking *Plugins*.

Note: If the list of available plugins is not displayed, open *Shell* (page 302) and verify that the TrueNAS® system can *ping* an address on the Internet. If it cannot, add a default gateway address and DNS server address in *Network* → *Global Configuration*.

Click *REFRESH INDEX* to refresh the current list of plugins.

Click a plugin icon to see the description, whether it is an Official plugin, the version available, and the number of installed instances.

To install the selected plugin, click *INSTALL*.

Note: A warning will display when an unofficial plugin is selected for installation.

Enter a *Jail Name*. A unique name is required, since multiple installations of the same plugin are supported. Names can contain letters, numbers, periods (.), dashes (-), and underscores (_).

Most plugins default to *NAT*. This setting is recommended as it does not require manual configuration of multiple available IP addresses and prevents addressing conflicts on the network.

Some plugins default to *DHCP* as their management utility conflicts with *NAT*. Keep these plugins set to *DHCP* unless manually configuring an IP address is preferred.

If both *NAT* and *DHCP* are unset, an IPv4 or IPv6 address can be manually entered. If desired, an IPv4 or IPv6 interface can be selected. If no interface is selected the jail IP address uses the current active interface. The IPv4 or IPv6 address must be in the range of the local network.

Click *ADVANCED PLUGIN INSTALLATION* to show all options for the plugin jail. The options are described in *Advanced Jail Creation* (page 263).

To start the installation, click *SAVE*.

Depending on the size of the application, the installation can take several minutes to download and install. A confirmation message is shown when the installation completes, along with any post-installation notes.

Installed plugins appear on the *Plugins* page as shown in [Figure 13.1](#).

Note: Plugins are also added to *Jails* as a *pluginv2* jail. This type of jail is editable like a standard jail, but the *UUID* cannot be altered. See *Managing Jails* (page 271) for more details about modifying jails.

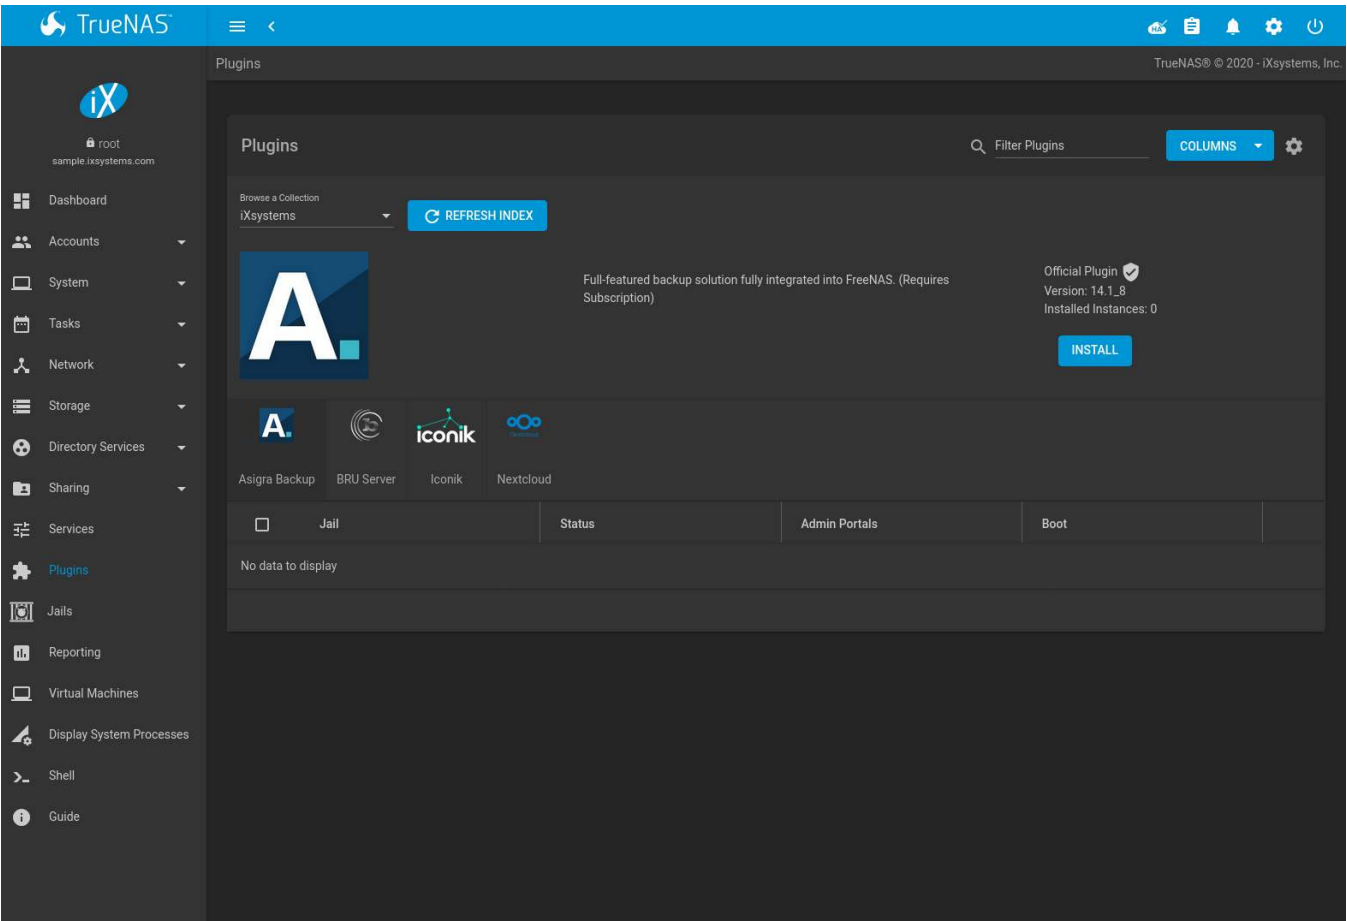


Fig. 13.1: Viewing Installed Plugins

Plugins are immediately started after installation. By default, all plugins are started when the system boots. Un-setting *Boot* means the plugin will not start when the system boots and must be started manually.

In addition to the *Jail* name, the *Columns* menu can be used to display more information about installed Plugins. More information such as *RELEASE* and *VERSION* is shown by clicking > (Expand). Options to *RESTART*, *STOP*, *UPDATE*, *MANAGE*, and *UNINSTALL* the plugin are also displayed. If an installed plugin has notes, the notes can be viewed by clicking *POST INSTALL NOTES*.

Plugins with additional documentation also have a *DOCUMENTATION* button which opens the README in the plugin repository.

The plugin must be started before the installed application is available. Click > (Expand) and *START*. The plugin *Status* changes to *up* when it starts successfully.

Stop and immediately start an *up* plugin by clicking > (Expand) and *RESTART*.

Click > (Expand) and *MANAGE* to open a management or configuration screen for the application. Plugins with a management interface show the IP address and port to that page in the *Admin Portal* column.

Note: Not all plugins have a functional management option. See [Managing Jails](#) (page 271) for more instructions about interacting with a plugin jail with the shell.

Some plugins have options that need to be set before their service will successfully start. Check the website of the application to see what documentation is available. If there are any difficulties using a plugin, refer to the official documentation for that application.

If the application requires access to the data stored on the TrueNAS® system, click the entry for the associated jail in the *Jails* page and add storage as described in [Additional Storage](#) (page 275).

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 -f /var/log/messages` to see if any errors were logged.

13.2 Updating Plugins

When a newer version of a plugin or release becomes available in the official repository, click > (Expand) and *UPDATE*. Updating a plugin updates the operating system and version of the plugin.




Fig. 13.2: Updating a Plugin

Updating a plugin also restarts that plugin. To update or upgrade the plugin jail operating system, see [Jail Updates and Upgrades](#) (page 273).

13.3 Uninstalling Plugins

Installing a plugin creates an associated jail. Uninstalling a plugin deletes the jail because it is no longer required. This means all **datasets or snapshots that are associated with the plugin are also deleted**. Make sure to back up any important data from the plugin **before** uninstalling it.

Figure 13.3 shows an example of uninstalling a plugin by expanding the plugin's entry and clicking *UNINSTALL*. A two-step dialog opens to confirm the action. **This is the only warning**. Enter the plugin name, set the *Confirm* checkbox, and click *DELETE* to remove the plugin and the associated jail, dataset, and snapshots.




Fig. 13.3: Uninstalling a Plugin and its Associated Jail and Dataset

13.4 Asigra Plugin

The Asigra plugin connects TrueNAS® to a third party service and is subject to licensing. Please read the [Asigra Software License Agreement](https://www.asigra.com/legal/software-license-agreement) (https://www.asigra.com/legal/software-license-agreement) before using this plugin.

To begin using Asigra services after installing the plugin, open the plugin options and click *Register*. A new browser tab opens to [register a user with Asigra](https://licenseportal.asigra.com/licenseportal/user-registration.do) (https://licenseportal.asigra.com/licenseportal/user-registration.do).

The TrueNAS® system must have a public static IP address for Asigra services to function.

Refer to the Asigra documentation for details about using the Asigra platform:

- [DS-Operator Management Guide](https://s3.amazonaws.com/asigra-documentation/Help/v14.1/DS-System%20Help/index.html) (https://s3.amazonaws.com/asigra-documentation/Help/v14.1/DS-System%20Help/index.html): Using the DS-Operator interface to manage the plugin DS-System service. Click *Management* in the plugin options to open the DS-Operator interface.
- [DS-Client Installation Guide](https://s3.amazonaws.com/asigra-documentation/Guides/Cloud%20Backup/v14.1/Client_Software_Installation_Guide.pdf) (https://s3.amazonaws.com/asigra-documentation/Guides/Cloud%20Backup/v14.1/Client_Software_Installation_Guide.pdf): How to install the DS-Client system. DS-Client aggregates backup content from endpoints and transmits it to the DS-System service.
- [DS-Client Management Guide](https://s3.amazonaws.com/asigra-documentation/Help/v14.1/DS-Client%20Help/index.html) (https://s3.amazonaws.com/asigra-documentation/Help/v14.1/DS-Client%20Help/index.html): Managing the DS-Client system after it has been successfully installed at one or more locations.

Note: This is a TrueNAS® licensed feature only. For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 6:00AM to 6:00PM Pacific Standard Time: <ul style="list-style-type: none">• US-only toll-free: 855-473-7449 option 2• Local and international: 408-943-4100 option 2
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none">• US-only toll-free: 855-499-5131• International: 408-878-3140 (international calling rates will apply)

JAILS

Jails are a lightweight, operating-system-level virtualization. One or multiple services can run in a jail, isolating those services from the host TrueNAS® system. TrueNAS® uses [iocage](https://github.com/iocage/iocage) (<https://github.com/iocage/iocage>) for jail and [plugin](#) (page 254) management. 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 TrueNAS® host operating system and any other jails running on that system.

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 131) must be created before using jails or [Plugins](#) (page 254). 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 TrueNAS® system.

If pools exist, but none have been chosen for use with jails or plugins, a dialog appears to choose a pool. Select a pool and click *CHOOSE*.

To select a different pool for jail and plugin storage, click ⚙ (Settings). A dialog shows the active pool. A different pool can be selected from the drop-down.

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 177).
- [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 TrueNAS® 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.

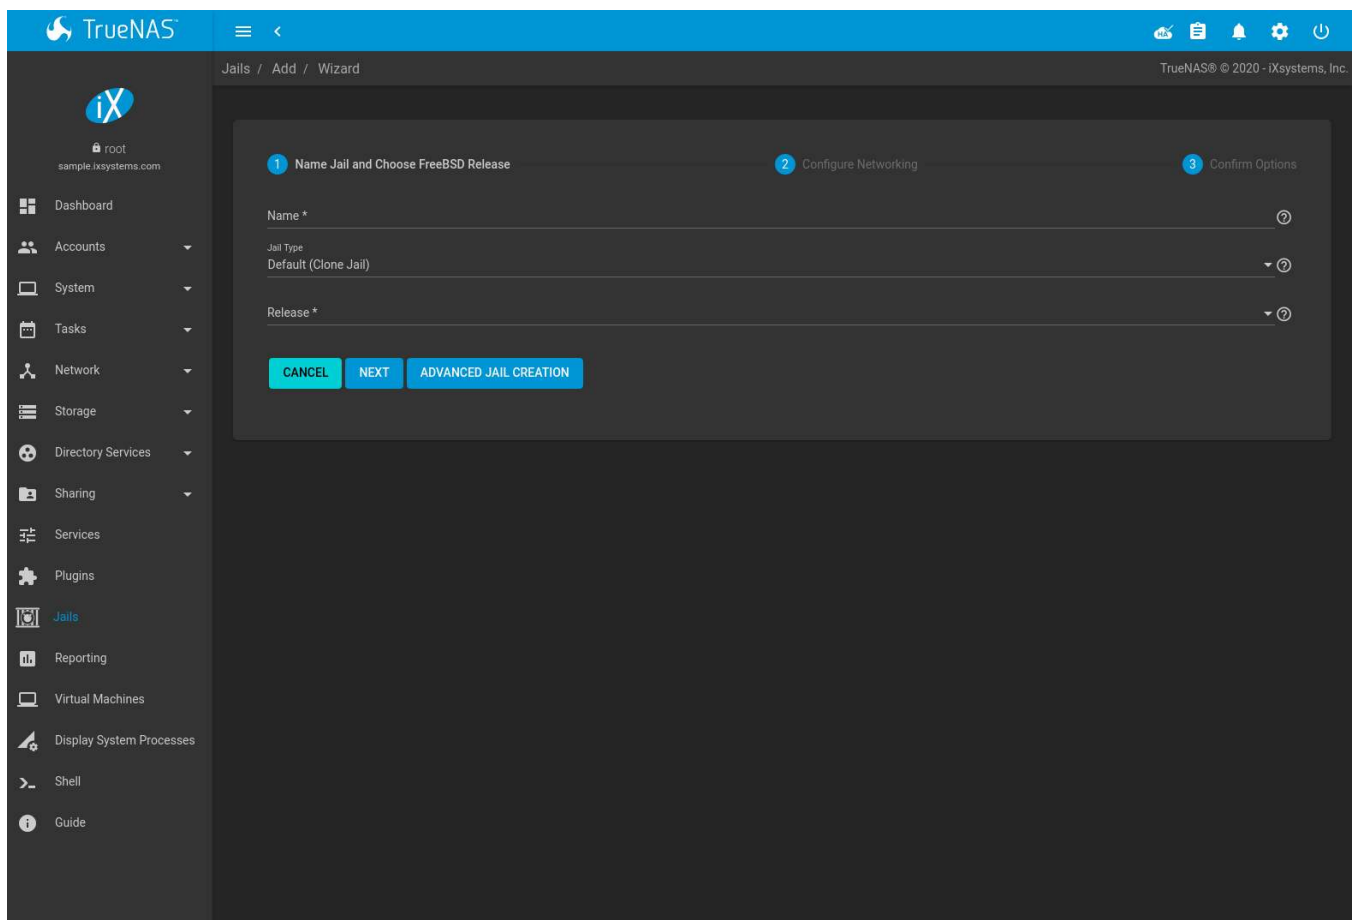
Note: `iocage` jail configs are stored in `/mnt/poolname/iocage/jails/jailname`. When `iocage` is updated, the `config.json` configuration file is backed up as `/mnt/poolname/iocage/jails/jailname/config_backup.json`. The backup file can be renamed to `config.json` to restore previous jail settings.

14.2 Creating Jails

TrueNAS® 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.



The screenshot shows the TrueNAS web interface with the 'Jails' menu item selected in the left sidebar. The main content area displays the 'Jails / Add / Wizard' screen. The wizard consists of three steps: 1. Name Jail and Choose FreeBSD Release, 2. Configure Networking, and 3. Confirm Options. Step 1 is currently active, showing fields for 'Name *', 'Jail Type' (set to 'Default (Clone Jail)'), and 'Release *'. Below these fields are three buttons: 'CANCEL', 'NEXT', and 'ADVANCED JAIL CREATION'. The top of the interface shows the TrueNAS logo and a navigation bar with icons for cloud, home, notifications, settings, and power.

Fig. 14.1: Jail Creation Wizard

The wizard provides the simplest process to create and configure a new jail.


Enter a *Jail Name*. Names can contain letters, numbers, periods (.), dashes (-), and underscores (_).

Choose a *Jail Type*: *Default (Clone Jail)* or *Basejail*. Clone jails are clones of the specified FreeBSD RELEASE. They are linked to that RELEASE, even if they are upgraded. Basejails mount the specified RELEASE directories as nullfs mounts over the jail directories. Basejails are not linked to the original RELEASE when upgraded.

Jails can run FreeBSD versions up to the same version as the host TrueNAS® system. Newer releases are not shown.

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.

Click *NEXT* to see a simplified list of networking options. Jails support several different networking solutions:

- *VNET* can be set to add a virtual network interface to the jail. This interface can be used to set NAT, DHCP, or static jail network configurations. Since *VNET* provides the jail with an independent networking stack, it can broadcast an IP address, which is required by some applications.
- The jail can use [Network Address Translation \(NAT\)](https://en.wikipedia.org/wiki/Network_address_translation) (https://en.wikipedia.org/wiki/Network_address_translation), which uses the TrueNAS® IP address and sets a unique port for the jail to use. *VNET* is required when *NAT* is selected.
- Configure the jail to receive its IP address from a DHCP server by setting *DHCP Autoconfigure IPv4*.
- Networking can be manually configured by entering values for the *IPv4 Address* or *IPv6 Address* fields. Any combination of these fields can be configured. Multiple interfaces are supported for IPv4 and IPv6 addresses. To add more interfaces and addresses, click *ADD*. Setting the *IPv4 Default Router* and *IPv6 Default Router* fields to *auto* automatically configures these values. *VNET* must be set to enable the *IPv4 Default Router* field. If no interface is selected when manually configuring IP addresses, TrueNAS® automatically assigns the given IP address of the jail to the current active interface of the host system.
- Leaving all checkboxes unset and fields empty initializes the jail without any networking abilities. Networking can be added to the jail after creation by going to *Jails* → > (Expand) →  *EDIT* → *Basic Properties*.

Setting a proxy in the TrueNAS® [network settings](#) (page 119) also configures new jails to use the proxy settings, except when performing DNS lookups. Make sure a firewall is properly configured to maximize system security.

When pairing the jail with a physical interface, edit the [network interface](#) (page 121) and set *Disable Hardware Offloading*. This prevents a network interface reset when the jail starts.

The screenshot displays the TrueNAS web interface for adding a new jail. The left sidebar contains a navigation menu with items like Dashboard, Accounts, System, Tasks, Network, Storage, Directory Services, Sharing, Services, Plugins, Jails, Reporting, Virtual Machines, Display System Processes, Shell, and Guide. The main content area shows a three-step wizard: 1. Name Jail and Choose FreeBSD Release, 2. Configure Networking (current step), and 3. Confirm Options. Under step 2, there are checkboxes for 'DHCP Autoconfigure IPv4', 'NAT', and 'VNET'. Below these are sections for IPv4 and IPv6 configuration. Each section has a dropdown for 'Interface', a text field for 'Address', a dropdown for 'Netmask' (IPv4) or 'Prefix' (IPv6), and a text field for 'Default Router'. At the bottom of the configuration area are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

Fig. 14.2: Configure Jail Networking

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.

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 14.3](#) is shown.

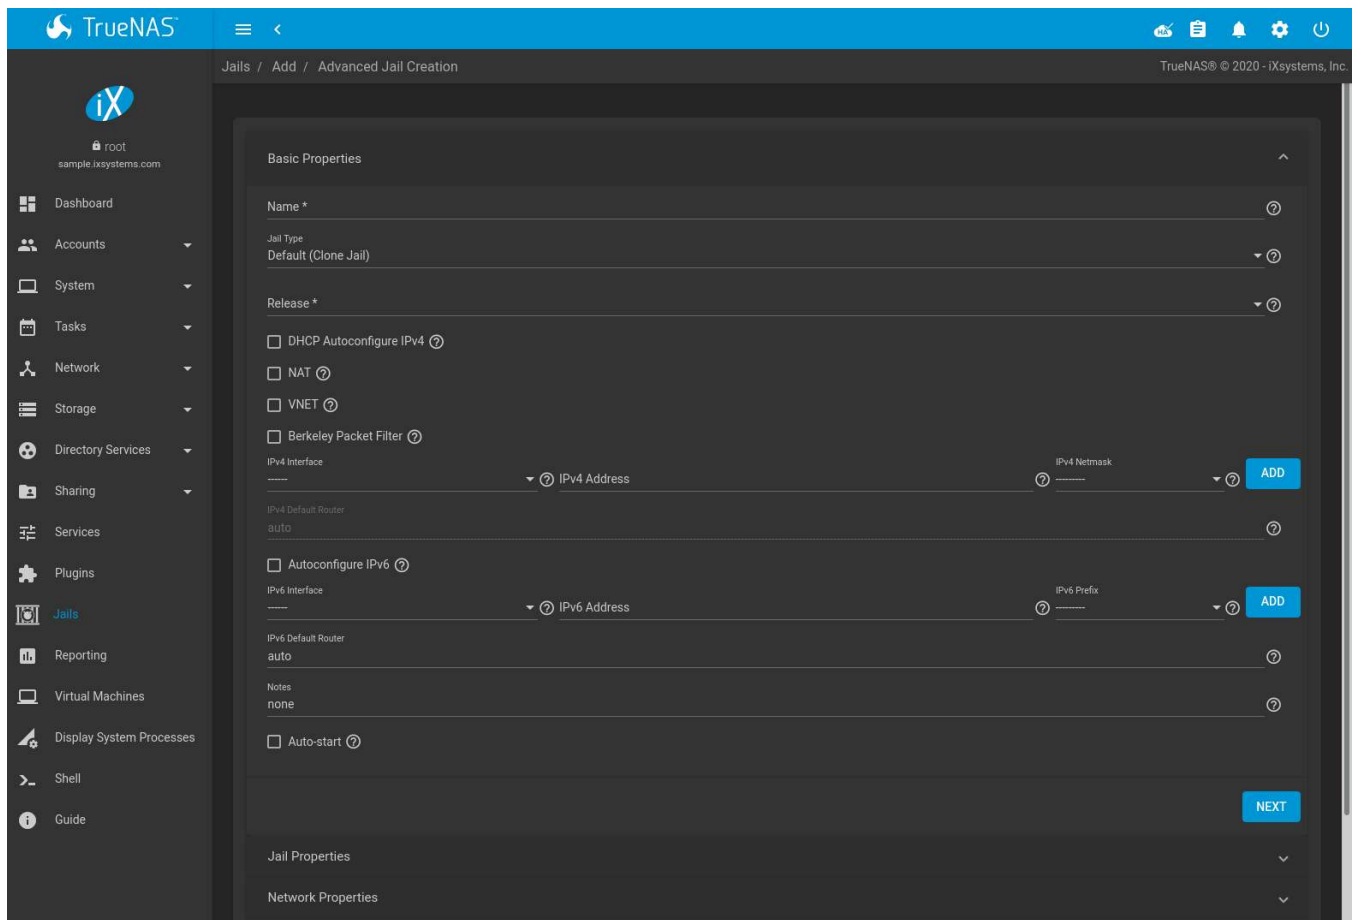


Fig. 14.3: Creating a Jail

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. Can contain letters, numbers, periods (.), dashes (-), and underscores (_).
Jail Type	drop-down	<i>Default (Clone Jail)</i> are clones of the specified RELEASE. They are linked to that RELEASE, even if they are upgraded. <i>Basejail</i> mount the specified RELEASE directories as nullfs mounts over the jail directories. Basejails are not linked to the original RELEASE when upgraded.
Release	drop-down menu	Required. Jails can run FreeBSD versions up to the same version as the host TrueNAS® system. Newer releases are not shown.
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.

Continued on next page

Table 14.1 – continued from previous page

Setting	Value	Description
NAT	checkbox	Network Address Translation (NAT). When set, the jail is given an internal IP address and connections are forwarded from the host to the jail. When NAT is set, <i>Berkeley Packet Filter</i> cannot be set. Adds the <i>NAT Port Forwarding</i> options to the jail Network Properties (page 269).
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. Cannot be set when <i>NAT</i> is set.
vnet_default_interface	drop-down	Set the default VNET interface. Only takes effect when <i>VNET</i> is set. Choose a specific interface, or set to <i>auto</i> to use the interface that has the default route. Choose <i>none</i> to not set a default VNET interface.
IPv4 Interface	drop-down menu	Choose a network interface to use for this IPv4 connection. See <i>note</i> (page ??) to add more.
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> . See <i>note</i> (page ??) to add more.
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 auto-configure IPv6 in the jail.
IPv6 Interface	drop-down menu	Choose a network interface to use for this IPv6 connection. See <i>note</i> (page ??) to add more.
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> . See <i>note</i> (page ??) to add more.
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.

Note: For static configurations not using DHCP or NAT, multiple IPv4 and IPv6 addresses and interfaces can be added to the jail by clicking *ADD*.

Similar to the [Jail Wizard](#) (page 261), 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_statfs</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_prestart</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 <code>SIGTERM</code> 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 <code>SIGTERM</code> 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 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 <code>passwd</code> .
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), fs-tatfs(2) (https://www.freebsd.org/cgi/man.cgi?query=statfs), 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 (<i>default</i>), 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! Use sysvmsg, sysvsem, and sysvshm instead.</i>

Continued on next page

Table 14.2 – continued from previous page

Setting	Value	Description
allow_raw_sockets	checkbox	Allow the jail to use raw sockets (https://en.wikipedia.org/wiki/Network_socket#Raw_socket). When set, the jail has access to lower-level network layers. This allows 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) to work in the jail, but has security implications and should only be used on jails running trusted software.
allow_chflags	checkbox	Treat jail users as privileged and allow the manipulation of system file flags. <i>securelevel</i> constraints are still enforced.
allow_mlock	checkbox	Allow jail to run services that use mlock(2) (https://www.freebsd.org/cgi/man.cgi?query=mlock) to lock physical pages in memory.
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_fusefs	checkbox	Allow privileged users inside the jail to mount and unmount fusefs. The jail must have FreeBSD 12.0 or newer installed. 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_vmm	checkbox	Grants the jail access to the Bhyve Virtual Machine Monitor (VMM). The jail must have FreeBSD 12.0 or newer installed with the vmm(4) (https://www.freebsd.org/cgi/man.cgi?query=vmm) kernel module loaded.
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.

Continued on next page

Table 14.2 – continued from previous page

Setting	Value	Description
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.
vnet_interfaces	string	Space-delimited list of network interfaces to attach to a VNET-enabled jail after it is created. Interfaces are automatically released when the jail is removed.

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.
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 NAME/UUID.
exec_fib	integer	Enter a number to define the routing table (FIB) to set when running commands inside the jail.
ip4.saddrsel	checkbox	Disables IPv4 source address selection for the jail in favor of the primary IPv4 address of the jail. Only available when the jail is not configured to use VNET.
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	Disable IPv6 source address selection for the jail in favor of the primary IPv6 address of the jail. Only available when the jail is not configured to use VNET.
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>
vnet0_mac	string	Leave this blank to generate random MAC addresses for the host and jail. To assign fixed MAC addresses, enter the host MAC address and the jail MAC address separated by a space.
vnet1_mac	string	Leave this blank to generate random MAC addresses for the host and jail. To assign fixed MAC addresses, enter the host MAC address and the jail MAC address separated by a space.
vnet2_mac	string	Leave this blank to generate random MAC addresses for the host and jail. To assign fixed MAC addresses, enter the host MAC address and the jail MAC address separated by a space.

Continued on next page

Table 14.3 – continued from previous page

Setting	Value	Description
vnet3_mac	string	Leave this blank to generate random MAC addresses for the host and jail. To assign fixed MAC addresses, enter the host MAC address and the jail MAC address separated by a space.

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> .
hostid_strict_check	checkbox	Check the jail <i>hostid</i> property. Prevents the jail from starting if the <i>hostid</i> does not match the host.
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 <i>/dev/proc</i> directory.
mount_linprocfs	checkbox	Allow mounting of a linprocfs(5) (https://www.freebsd.org/cgi/man.cgi?query=linprocfs) filesystem in the jail.
template	checkbox	Convert the jail into a template. Template jails can be used to quickly create jails with the same configuration.
host_time	checkbox	Synchronize the time between jail and host.
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>
allow_tun	checkbox	Expose host tun(4) (https://www.freebsd.org/cgi/man.cgi?query=tun) devices in the jail. Allow the jail to create tun devices.
Autoconfigure IPv6 with rtsold	checkbox	Use rtsold(8) (https://www.freebsd.org/cgi/man.cgi?query=rtsold) as part of IPv6 autoconfiguration. Send ICMPv6 Router Solicitation messages to interfaces to discover new routers.
ip_hostname	checkbox	Use DNS records during jail IP configuration to search the resolver and apply the first open IPv4 and IPv6 addresses. See jail(8) (https://www.freebsd.org/cgi/man.cgi?query=jail).
assign_localhost	checkbox	Add network interface <i>lo0</i> to the jail and assign it the first available localhost address, starting with <i>127.0.0.2</i> . <i>VNET</i> cannot be set. Jails using <i>VNET</i> configure a localhost as part of their virtualized network stack.

Click *SAVE* when the desired jail properties have been set. New jails are added to the primary list in the *Jails* menu.

14.2.2.1 Creating Template Jails

Template jails are basejails that can be used as a template to efficiently create jails with the same configuration. These steps create a template jail:

1. Go to *Jails* → *ADD* → *ADVANCED JAIL CREATION*.
2. Select *Basejail* as the *Jail Type*. Configure the jail with desired options.
3. Set *template* in the *Custom Properties* tab.
4. Click *Save*.
5. Click *ADD*.
6. Enter a name for the template jail. Leave *Jail Type* as *Default (Clone Jail)*. Set *Release* to *basejailname(template)*, where *basejailname* is the name of the base jail created earlier.
7. Complete the jail creation wizard.

14.3 Managing Jails

Clicking *Jails* shows a list of installed jails. An example is shown in Figure 14.4.

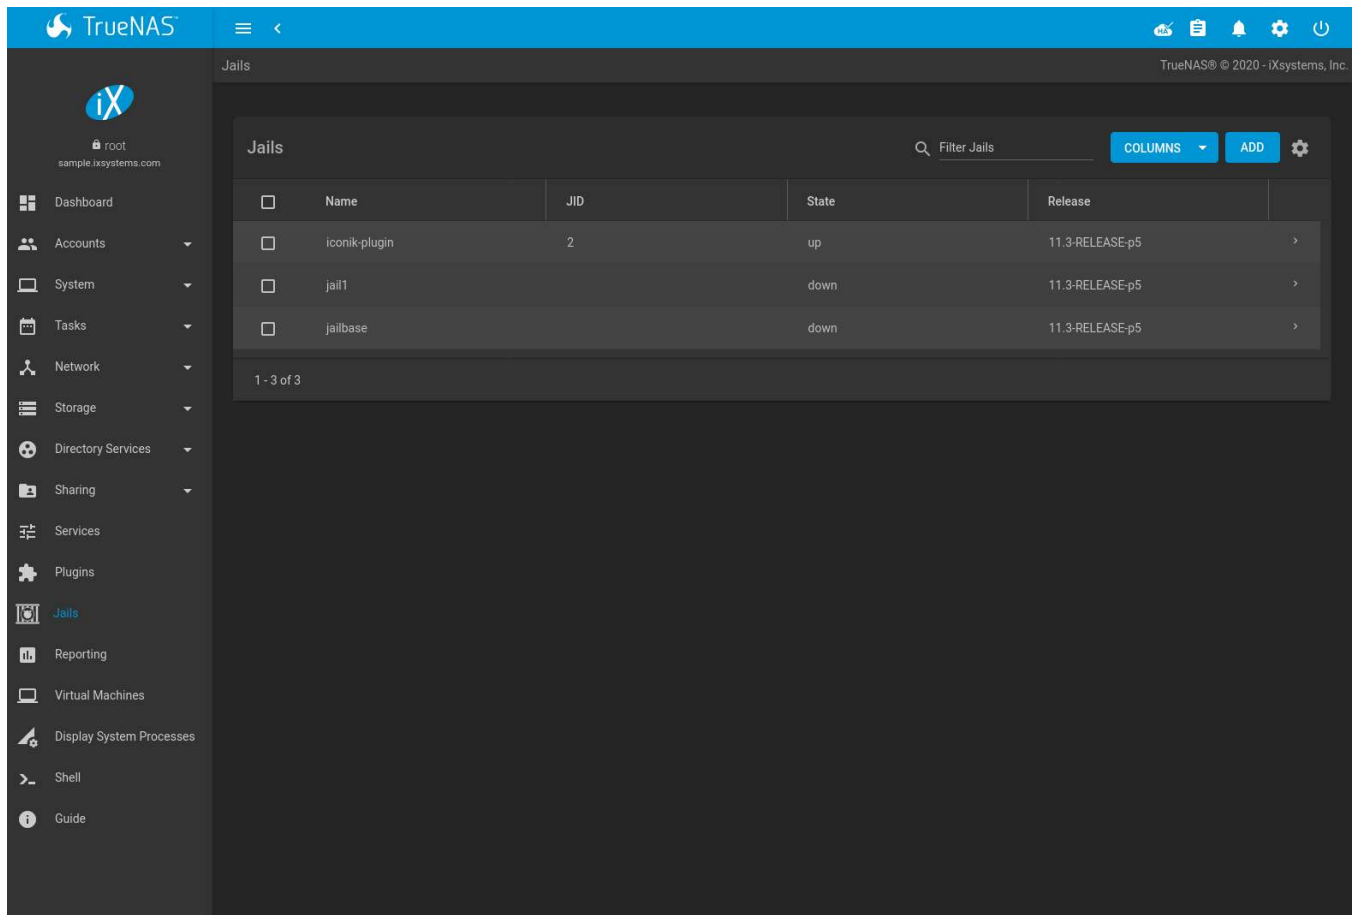


Fig. 14.4: Jail Overview Section

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.

More information such as *IPV4*, *IPV6*, *TYPE* of jail, and whether it is a *TEMPLATE* jail or *BASEJAIL* can be shown by clicking > (Expand). Additional options for that jail are also displayed. These are described in Table 14.5. Figure 14.5 shows the menu that appears.

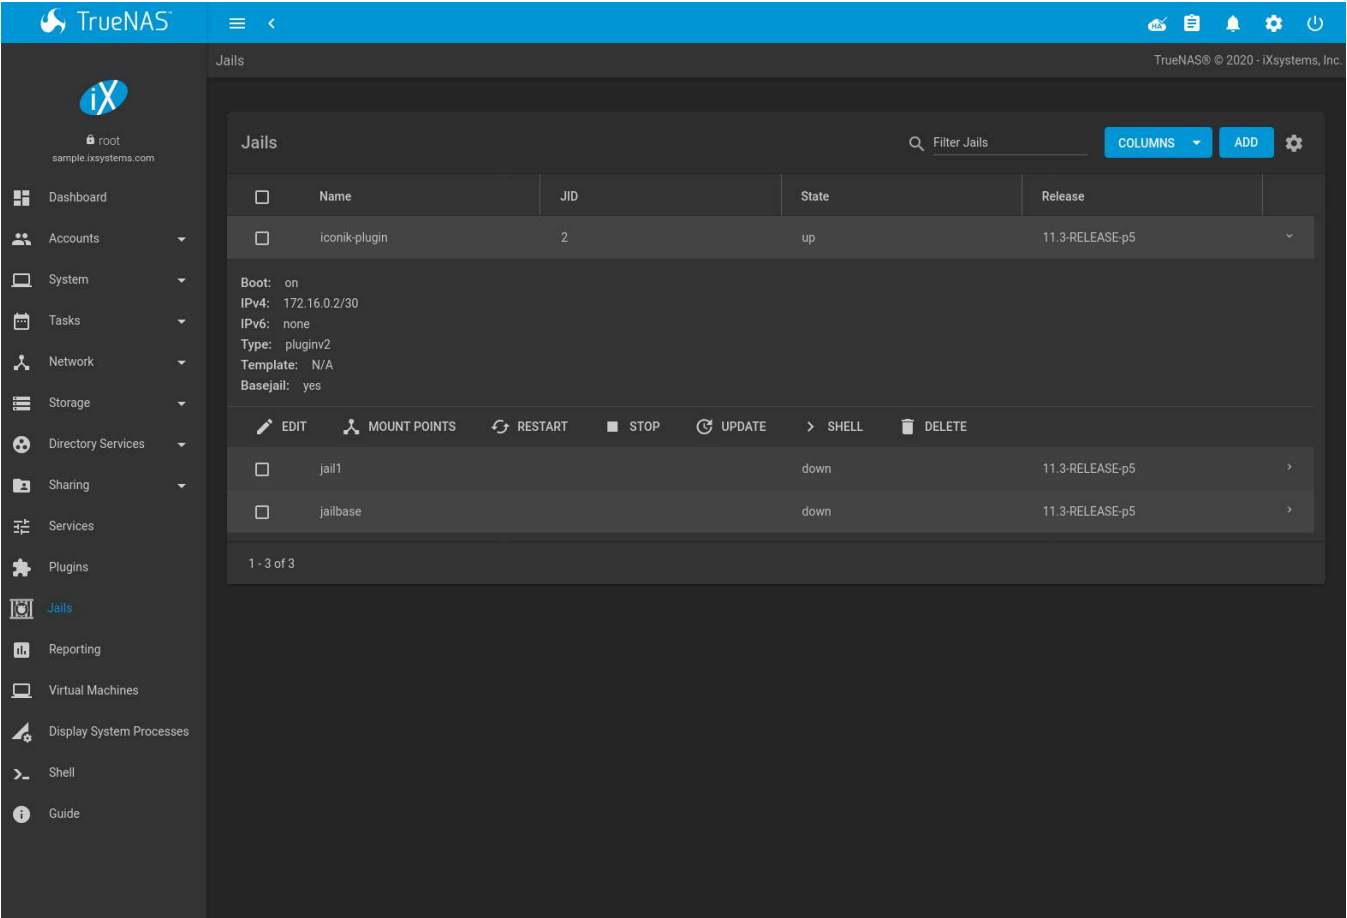


Fig. 14.5: Jail Options Menu

Warning: Modify the IP address information for a jail by clicking > (Expand) → *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 TrueNAS® reboot.

Table 14.5: Jail Option Menu Entry Descriptions

Option	Description
EDIT	Used to modify the settings described in <i>Advanced Jail Creation</i> (page 263). A jail cannot be edited while it is running. The settings can be viewed, but are read only.
MOUNT POINTS	Select an existing mount point to <i>EDIT</i> or click <i>ACTIONS</i> → <i>Add Mount Point</i> to create a mount point for the jail. 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 mount point. See <i>Additional Storage</i> (page 275) for more details.
RESTART	Stop and immediately start an <i>up</i> jail.
START	Start a jail that has a current <i>STATE</i> of <i>down</i> .
STOP	Stop a jail that has a current <i>STATE</i> of <i>up</i> .

Continued on next page

Table 14.5 – continued from previous page

Option	Description
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 <code>root</code> command prompt to interact with a jail directly from the command line. Type <code>exit</code> to leave the command prompt.
DELETE	Caution: deleting the jail also deletes all of the jail contents and all associated <i>snapshots</i> (page 152). Back up the jail 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.

Jail status messages and command output are stored in `/var/log/iocage.log`.

14.3.1 Jail Updates and Upgrades

Click **>** (Expand) → *Update* to update a jail to the most current patch level of the installed FreeBSD release. This does **not** change the release. For example, a jail installed with *FreeBSD 11.2-RELEASE* can update to *p15* or the latest patch of 11.2, but not an 11.3-RELEASE-*p#* version of FreeBSD.

A jail *upgrade* replaces the jail FreeBSD operating system with a new release of FreeBSD, such as taking a jail from FreeBSD 11.2-RELEASE to 11.3-RELEASE. Upgrade a jail by stopping it, opening the *Shell* (page 302) and entering `iocage upgrade name -r release`, where *name* is the plugin jail name and *release* is the desired release to upgrade to.

Tip: It is possible to *manually remove* (page 144) 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 *STATE* must be *up* before the *SHELL* option is available. If the jail is not up, start it by clicking *Jails* → **>** (Expand) → *START* for the desired jail. Click **>** (Expand) → *SHELL* to open a shell in 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 TrueNAS® 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 TrueNAS® system. It is possible to give a FreeBSD jail access to an area of storage on the TrueNAS® 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 TrueNAS® system. For example, Transmission is a plugin that stores data using BitTorrent. The TrueNAS® external storage is added using the `mount_nullfs(8)` (https://www.freebsd.org/cgi/man.cgi?query=mount_nullfs) mechanism, which links data that resides outside of the jail as a storage area within a jail.

› (Expand) → *MOUNT POINTS* shows any added storage and allows adding more storage.

Note: A jail must have a *STATE* of *down* before adding a new mount point. Click › (Expand) and *STOP* for a jail to change the jail *STATE* to *down*.

Storage can be added by clicking *Jails* → › (Expand) → *MOUNT POINTS* for the desired jail. The *MOUNT POINT* section is a list of all of the currently defined mount points.

Go to *MOUNT POINTS* → *ACTIONS* → *Add Mount Point* to add storage to a jail. This opens the screen shown in [Figure 14.6](#).

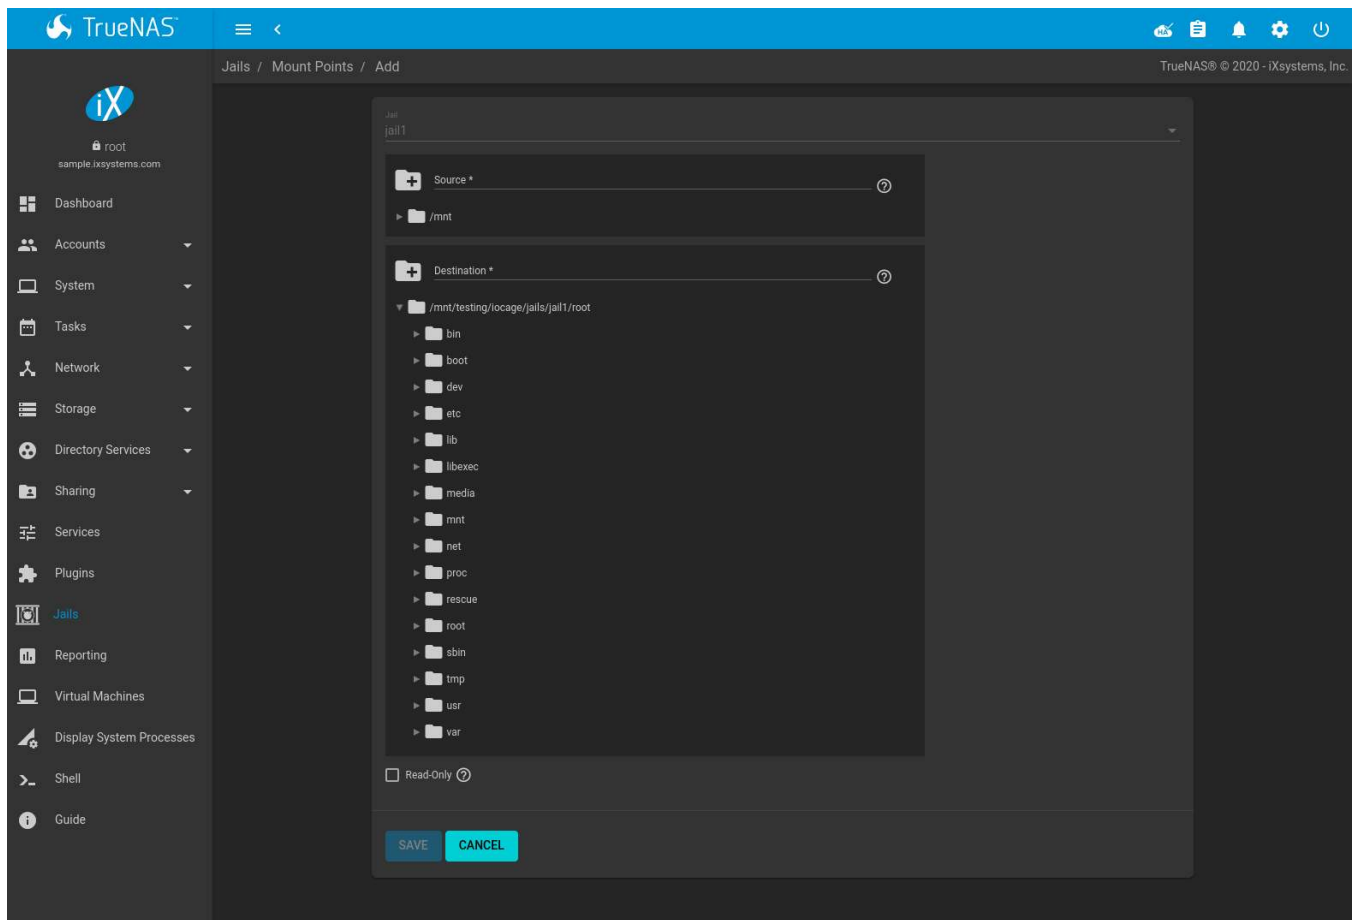


Fig. 14.6: Adding Storage to a Jail

Browse to the *Source* and *Destination*, where:

- *Source*: is the directory or dataset on the TrueNAS® system which will be accessed by the jail. TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 > (Expand) → *MOUNT POINTS* → *ACTIONS* → *Add Mount Point* to select 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 14.7](#), a dataset named `pool1/smb-backups` has been chosen as the *Source* as it contains the files stored on the TrueNAS® system. The user entered `/mnt/iocage/jails/jail1/root/mounted` as the directory to be mounted in the *Destination* field. To users inside the jail, this data appears in the `/root/mounted` directory.

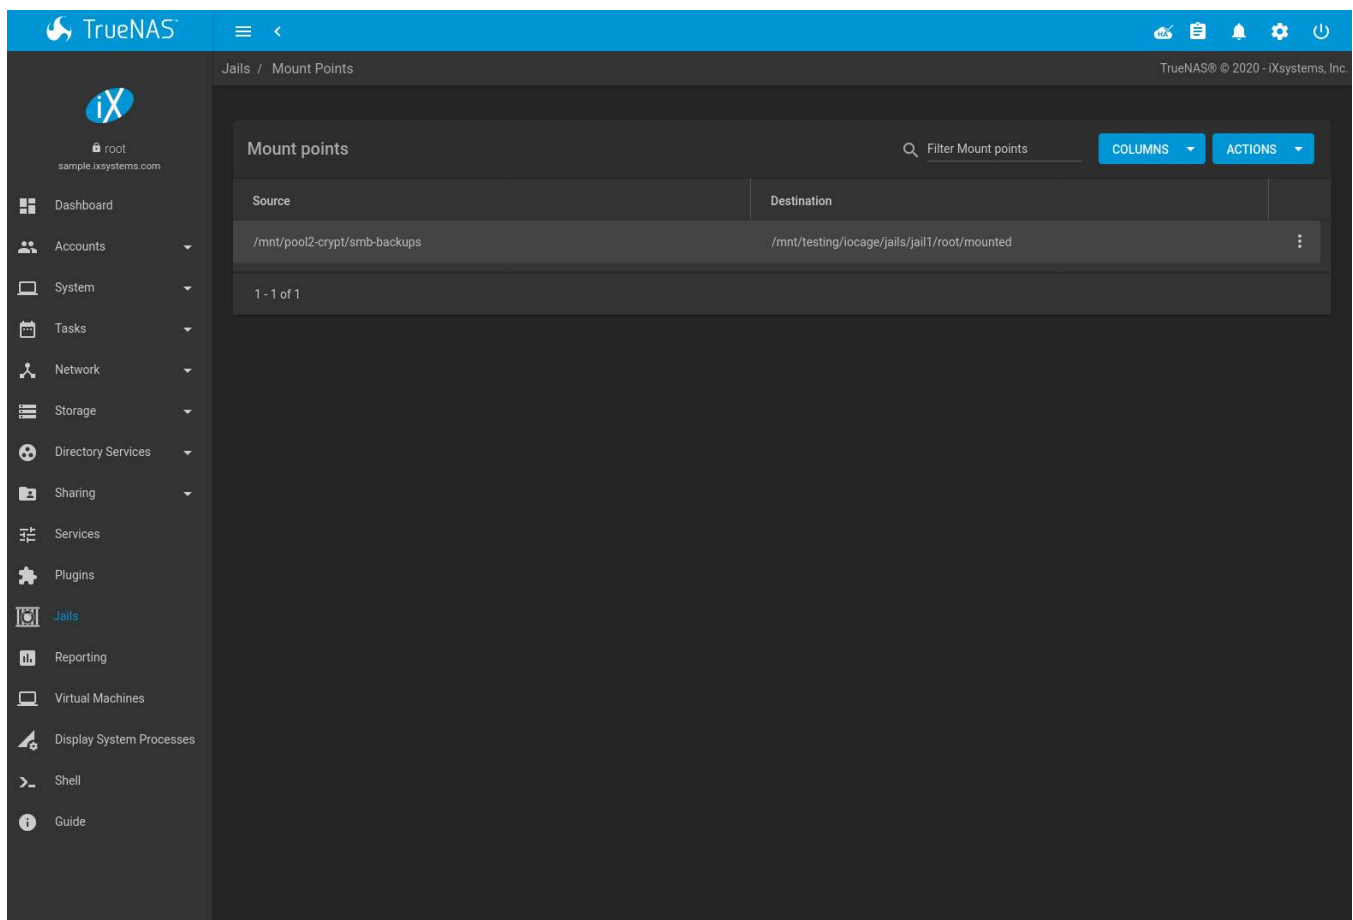


Fig. 14.7: Example Storage

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 TrueNAS® 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 TrueNAS® 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. To install more software, start the jail and click > SHELL.



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 precompiled 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 TrueNAS® 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 TrueNAS® 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 14.8](#) shows the *Configuration Options* for *audiotag*, a utility for renaming multiple audio files.



If you buy from Amazon USA, please support us by using [this link](#).

Follow us

[Blog](#)
[Twitter](#)
[Status page](#)

Port details

audiotag Command-line tool for mass tagging/renaming of audio files
0.19_1 [audio](#) [Z=1](#) [🔍](#) [🔗](#) [🔖](#)

There is no maintainer for this port.

Any concerns regarding this port should be directed to the FreeBSD Ports mailing list via ports@FreeBSD.org [🔗](#)

Port Added: 2008-04-15 13:43:37

Last Update: 2016-12-02 09:21:59

SVN Revision: 427548

Also Listed In: [multimedia](#)

License: GPLv2+

Audiotag is a command-line tool for mass tagging/renaming of audio files
it supports the vorbis comment, id3 tags, and MP4 tags.

WWW: <https://github.com/Daenyth/audiotag>

SVNWeb : [Homepage](#) : [PortsMon](#)

Pseudo-**pkg-plist** information, but much better, from make generate-plist
▼ [Expand this list](#) (4 items)

Dependency line: `audiotag>0:audio/audiotag`

To install the port: `cd /usr/ports/audio/audiotag/ && make install clean`

To add the package: `pkg install audiotag`

PKGNAME: audiotag

There is no flavor information for this port.

distinfo:

```
SHA256 (audiotag-0.19.tar.bz2) = 7b6a2de751058a95755f0842b83f2b1d8b94e5cd7634cbe71d67257208bf4646
SIZE (audiotag-0.19.tar.bz2) = 15816
```

NOTE: FreshPorts displays only information on required and default dependencies. Optional dependencies are not covered.

Runtime dependencies:

1. flac : [audio/flac](#)
2. id3tag : [audio/id3lib](#)
3. AtomicParsley : [multimedia/atomicparsley](#)
4. vorbiscomment : [audio/vorbis-tools](#)
5. perl5 >= 5.24 < 5.25 : [lang/perl5.24](#)

There are no ports dependent upon this port

Configuration Options

====> The following configuration options are available for audiotag-0.19.1:

```
DOC=on: Build and/or install documentation
FLAC=on: FLAC lossless audio codec support
ID3=on: ID3 tags support
MP4=on: MP4 media format support
VORBIS=on: Ogg Vorbis audio codec support
====> Use 'make config' to modify these settings
```

USES:

tar:bzip2 shebangfix perl5

Master Sites:

1. <https://cloud.github.com/downloads/Daenyth/audiotag/>

Login

[User Login](#)
[Create account](#)

Servers and bandwidth provided
by
[New York Internet](#), [SuperNews](#)
and [RootBSD](#)

This site

[What is FreshPorts?](#)
[About the authors](#)
[Issues](#)
[FAQ](#)
[How big is it?](#)
[The latest upgrade!](#)
[Privacy](#)
[Forums](#)
[Blog](#)
[Contact](#)

Search

Enter Keywords:

[more...](#)

Latest

Vulnerabilities

powerdns	Mar 19
rubygem-actionview4	Mar 18
rubygem-actionview5	Mar 18
rubygem-actionview50	Mar 18
putty	Mar 17
putty-gtk2	Mar 17
putty-nogtk	Mar 17
py-notebook	Mar 16
ruby-gems	Mar 15
ntp	Mar 07
openssl*	Mar 07
openssl111	Mar 07
rsh	Mar 06
rt42	Mar 06
rt44	Mar 06

13 vulnerabilities
affecting 42 ports have
been reported in the
past 14 days

* - modified, not new

[All vulnerabilities](#)

Last updated:
2019-03-19 14:51:44

Ports

[Home](#)
[Categories](#)
[Deleted ports](#)
[Sanity Test Failures](#)
[Newsfeeds](#)

Statistics

[Graphs](#)
[NEW Graphs](#)
(javascript)
[Traffic](#)

Calculated hourly:

Port count	36286
Broken	116
Deprecated	119
Ignore	337
Forbidden	5
Restricted	161
No CDROM	74
Vulnerable	37
Expired	6
Set to expire	93
Interactive	0
new 24 hours	4
new 48 hours	12
new 7 days	35
new fortnight	65
new month	138

Number of commits found: 22

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
29 May 2012 14:01:15 🔍 🔗 🔖 0.19_1	bapt 🔗	Convert to new options framework While here activate flac by default share descriptions of mp4, id3, flac and vorbis

Packages are built with 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 14.9 is displayed:




Fig. 14.9: Configuration Options for Audiotag Port

This port has several configurable options: *DOCS*, *FLAC*, *ID3*, *MP4*, and *VORBIS*. Selected options are shown with a *. Use the arrow keys to select an option and press `spacebar` to toggle the value. Press `Enter` when satisfied with the 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
```


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

TrueNAS® 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.

Note: Reporting graphs do not appear if there is no related data.

Graphs are grouped by category on the Reporting page:

- *CPU*
 - *CPU* (<https://collectd.org/wiki/index.php/Plugin:CPU>) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle. Graphs of short-, mid-, and long-term load are shown, along with CPU temperature graphs.
- *Disk*
 - *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.

Note: Temperature monitoring for the disk is disabled if *HDD Standby* is enabled in *Disks* (page 156).

- *Memory*
 - *Memory* (<https://collectd.org/wiki/index.php/Plugin:Memory>) displays memory usage.
 - *Swap* (<https://collectd.org/wiki/index.php/Plugin:Swap>) displays the amount of free and used swap space.
- *Network*
 - *Interface* (<https://collectd.org/wiki/index.php/Plugin:Interface>) shows received and transmitted traffic in megabytes per second for each configured interface.
- *NFS*
 - *NFS* (<https://collectd.org/wiki/index.php/Plugin:NFS>) shows information about the number of procedure calls for each procedure and whether the system is a server or client.
- *Partition*
 - *Disk space* (<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>) displays the number of processes. It is grouped by state.
- *Target*
 - Target shows bandwidth statistics for iSCSI ports.
- *UPS*
 - *UPS* (<https://collectd.org/wiki/index.php/Plugin:NUT>) displays statistics about an uninterruptible power supply (UPS) using *Network UPS tools* (<https://networkupstools.org/>). Statistics include voltages, currents, power, frequencies, load, and temperatures.
- *ZFS*
 - *ZFS* (https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) shows compressed physical ARC size, hit ratio, demand data, demand metadata, 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/`.

<p>Warning: Reporting data is frequently written and should not be stored on the boot pool or operating system device.</p>

Note: This is a TrueNAS® licensed feature only. For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 6:00AM to 6:00PM Pacific Standard Time: <ul style="list-style-type: none">• US-only toll-free: 855-473-7449 option 2• Local and international: 408-943-4100 option 2
Telephone	After Hours (24x7 Gold Level Support only): <ul style="list-style-type: none">• US-only toll-free: 855-499-5131• International: 408-878-3140 (international calling rates will apply)

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 260), although there is additional overhead. A portion of system RAM is assigned to each VM, and each VM uses a *zvol* (page 145) for storage. While a VM is running, these resources are not available to the host computer or other VMs.

TrueNAS® 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). VMs cannot be created unless the host system supports these features.

To verify that an Intel processor has the required features, use *Shell* (page 302) 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 302) to run `grep POPCNT /var/run/dmesg.boot`. If the output shows the POPCNT feature, this processor can be used with *bhyve*.

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.

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 CPU.

Virtual Machines shows a list of installed virtual machines and available memory. The available memory changes depending on what the system is doing, including which virtual machines are running.

A log file for each VM is written to `/var/log/vm/vmname`.




Fig. 16.1: Virtual Machines

Name, *State*, and *Autostart* are displayed on the *Virtual Machines* page. Click > (Expand) to view additional options for controlling and modifying VMs:

- *Start* boots a VM. VMs can also be started by clicking the slide toggle on the desired VM.
If there is insufficient memory to start the VM, a dialog will prompt to *Overcommit Memory*. Memory overcommitment allows the VM to launch even though there is insufficient free memory. Proceeding with the overcommitment option should be used with caution.
To start a VM when the host system boots, set *Autostart*. If *Autostart* is set and the VM is in an encrypted, locked pool, the VM starts when the pool is unlocked.
- *Edit* changes VM settings.
- *Delete* removes the VM. *Zvols* (page 145) used in *disk devices* (page 296) and image files used in *raw file* (page 297) devices are *not* removed when a VM is deleted. These resources can be removed manually in *Storage* → *Pools* after it is determined that the data in them has been backed up or is no longer needed.
- *Devices* is used to add, remove, or edit devices attached to a virtual machine.
- *Clone* copies the VM. A new name for the clone can be specified. If a custom name is not entered, the name assigned is `vmname_cloneN`, where *vmname* is the original VM name and *N* is the clone number. Each clones is given a new VNC port.

These additional options in > (Expand) are available when a VM is running:

- *Power off* immediately halts the VM. This is equivalent to unplugging the power cord from a computer.
- *Stop* shuts down the VM.
- *Restart* shuts down and immediately starts the VM.

- VMs with *Enable VNC* set show a *VNC* button. VNC connections permit remote graphical access to the VM.
- *SERIAL* opens a connection to a virtual serial port on the VM. `/dev/nmdm1B` is assigned to the first VM, `/dev/nmdm2B` is assigned to the second VM, and so on. These virtual serial ports allow connections to the VM console from the *Shell* (page 302).

Tip: The `nmdm` (<https://www.freebsd.org/cgi/man.cgi?query=nmdm>) device is dynamically created. The actual `nmdm XY` name varies on each VM.

To connect to the first VM, type `cu -l /dev/nmdm1B -s 9600` in the *Shell* (page 302). See `cu(1)` (<https://www.freebsd.org/cgi/man.cgi?query=cu>) for more information.

16.1 Creating VMs

Click **ADD** to open the wizard in Figure 16.2:




Fig. 16.2: Add VM


The configuration options for a Virtual Machine (VM) type are described in Table 16.1.

Table 16.1: VM Wizard Options

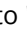

Screen #	Setting	Value	Description
1	Guest Operating System	drop-down menu	Choose the VM operating system type. Choices are: <i>Windows</i> , <i>Linux</i> , or <i>FreeBSD</i> . See this guide (https://github.com/FreeBSD-UPB/freebsd/wiki/How-to-launch-different-guest-OS) for detailed instructions about using a different guest OS.
1	Name	string	Name of the VM. Alphanumeric characters and <code>_</code> are allowed. The name must be unique.
1	Description	string	Description (optional).
1	System Clock	drop-down menu	Virtual Machine system time. Options are <i>Local</i> and <i>UTC</i> . <i>Local</i> is default.
1	Boot Method	drop-down menu	Choices are <i>UEFI</i> , <i>UEFI-CSM</i> , and <i>Grub</i> . Select <i>UEFI</i> for newer operating systems, or <i>UEFI-CSM</i> (Compatibility Support Mode) for older operating systems that only understand <i>BIOS booting</i> . <i>VNC connections are only available with *UEFI</i> . <i>Grub</i> is not supported by <i>Windows</i> guest operating systems.
1	Start on Boot	checkbox	Set to start the VM when the system boots.
1	Enable VNC	checkbox	Add a VNC remote connection. Requires <i>UEFI</i> booting.
1	Delay VM Boot Until VNC Connects	checkbox	Wait to start VM until VNC client connects. Only appears when <i>Enable VNC</i> is set.
1	Bind	drop-down menu	VNC network interface IP address. The primary interface IP address is the default. A different interface IP address can be chosen.
2	Virtual CPUs	integer	Number of virtual CPUs to allocate to the VM. The maximum is 16 unless limited by the host CPU. The VM operating system might also have operational or licensing restrictions on the number of CPUs.
2	Memory Size	integer	Set the amount of RAM for the VM. Allocating too much memory can slow the system or prevent VMs from running. This is a humanized field (page 14).
3	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> and choose an existing zvol from the <i>Select Existing zvol</i> drop-down.
3	Select Disk Type	drop-down menu	Select the disk type. Choices are <i>AHCI</i> and <i>VirtIO</i> . Refer to Disk Devices (page 296) for more information about these disk types.
3	Size (Examples: 500 KiB, 500M, 2TB)		Allocate the amount of storage for the zvol. This is a humanized field (page 14). Numbers without unit letters are interpreted as megabytes. For example, <code>500</code> sets the zvol size to 500 megabytes.
3	Zvol Location		When <i>Create new disk image</i> is chosen, select a pool or dataset for the new zvol.
3	Select existing zvol	drop-down menu	When <i>Use existing disk image</i> is chosen, select an existing zvol for the VM.
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 16.1 – continued from previous page

Screen #	Setting	Value	Description
4	Attach NIC	drop-down menu	Select the physical interface to associate with the VM.
5	Optional: Choose installation media image	browse button	Click  (Browse) to select an installer ISO or image file on the TrueNAS [®] system.
5	Upload ISO	checkbox and	Set to upload an installer ISO or image file to the TrueNAS [®] system.

The final screen of the Wizard displays the chosen options for the new Virtual Machine (VM) type. Click *SUBMIT* to create the VM or *BACK* to change any settings.

After the VM has been installed, remove the install media device. Go to *Virtual Machines* →  (Options) → *Devices*. Remove the *CDROM* device by clicking  (Options) → *Delete*. This prevents the virtual machine from trying to boot with the installation media after it has already been installed.

This example creates a FreeBSD VM:

1. *Guest Operating System* is set to *FreeBSD*. *Name* is set to *samplevm*. Other options are left at defaults.
2. *Virtual CPUs* is set to 2 and *Memory Size (MiB)* is set to 2048.
3. *Create new disk image* is selected. The zvol size is set to 20 GiB and stored on the pool named *pool1*.
4. Network settings are left at default values.
5. A FreeBSD ISO installation image has been selected and uploaded to the TrueNAS[®] system. The *Choose installation media image* field is populated when the upload completes.
6. After verifying the *VM Summary* is correct, *SUBMIT* is clicked.

Figure 16.3 shows the confirmation step and basic settings for the new virtual machine:

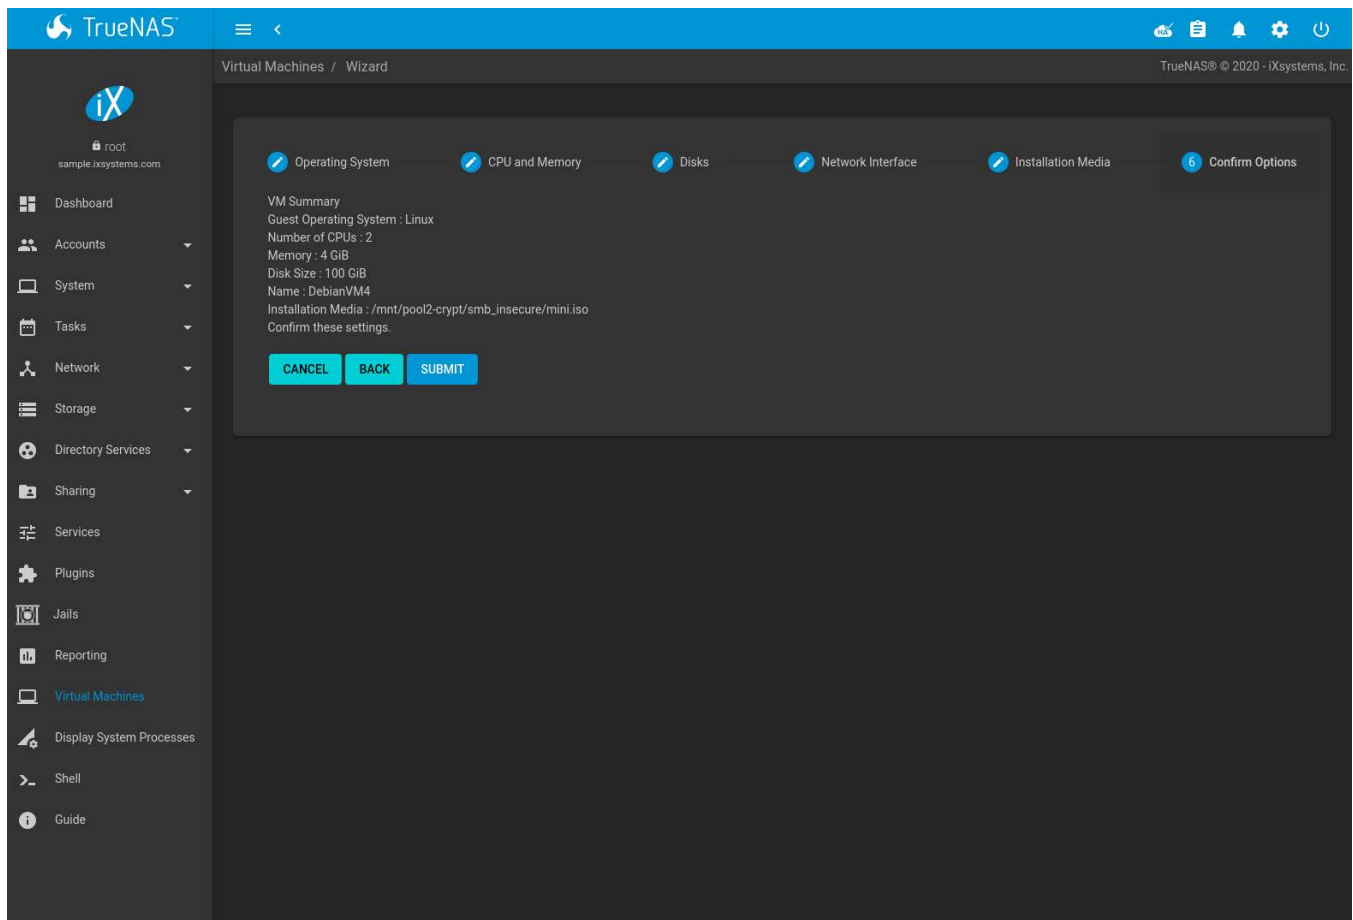


Fig. 16.3: Creating a Sample Virtual Machine

16.2 Installing Docker

Docker (<https://www.docker.com/>) can be used on TrueNAS® by installing it on a Linux virtual machine.

Choose a Linux distro and install it on TrueNAS® by following the steps in [Creating VMs](#) (page 290). Using **Ubuntu** (<https://ubuntu.com/>) is recommended.

After the Linux operating system has been installed, start the VM. Connect to it by clicking > (Expand) → VNC. Follow the [Docker documentation](https://docs.docker.com/) (<https://docs.docker.com/>) for Docker installation and usage.

16.3 Adding Devices to a VM

Go to *Virtual Machines*, ⋮ (Options) → *Devices*, and click **ADD** to add a new VM device.

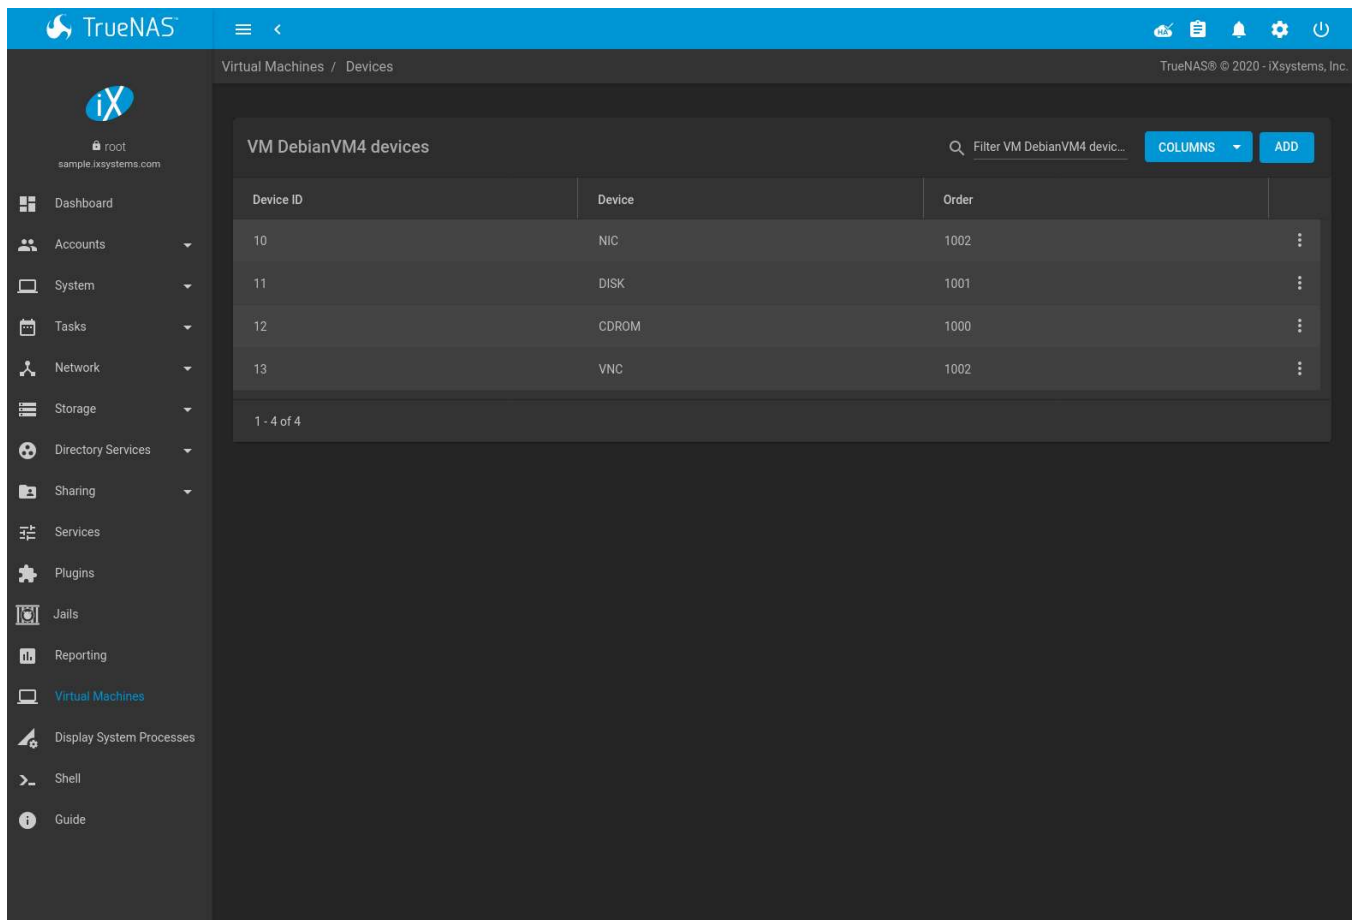


Fig. 16.4: VM Devices

Select the new device from the *Type* field. These devices are available:

- [CD-ROM](#) (page 294)
- [NIC \(Network Interface Card\)](#) (page 295)
- [Disk Device](#) (page 296)
- [Raw File](#) (page 297)
- [VNC Interface](#) (page 298) (only available on virtual machines with *Boot Loader Type* set to *UEFI*)

Virtual Machines → ⋮ (Options) → *Devices* is also used to edit or delete existing devices. Click ⋮ (Options) for a device to display *Edit*, *Delete*, *Change Device Order*, and *Details* options:

- *Edit* modifies a device.
- *Delete* removes the device from the VM.
- *Change Device Order* sets the priority number for booting this device. Smaller numbers are higher in boot priority.
- *Details* shows additional information about the specific device. This includes the physical interface and MAC address in a *NIC* device, the path to the zvol in a *DISK* device, and the path to an *.iso* or other file for a *CDROM* device.

16.3.1 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 TrueNAS® storage. In this example, a FreeBSD installation

image is shown:




Fig. 16.5: CD-ROM Device

Note: VMs from other virtual machine systems can be recreated for use in TrueNAS®. Back up the original VM, then create a new TrueNAS® VM with virtual hardware as close as possible to the original VM. Binary-copy the disk image data into the *zvol* (page 145) created for the TrueNAS® 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 TrueNAS® VM, and restore the data into the new VM.

16.3.2 NIC (Network Interfaces)

Figure 16.6 shows the fields that appear after going to *Virtual Machines* → ⋮ (Options) → *Devices*, clicking *ADD*, and selecting *NIC* as the *Type*.




Fig. 16.6: 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. To prevent a network interface reset when the VM starts, edit the [network interface](#) (page 121) and set *Disable Hardware Offloading*.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.

Tip: To check which interface is attached to a VM, start the VM and go to the [Shell](#) (page 302). Type `ifconfig` and find the `tap` (<https://en.wikipedia.org/wiki/TUN/TAP>) interface that shows the name of the VM in the description.

16.3.3 Disk Devices

Zvols (page 145) are typically used as virtual hard drives. After [creating a zvol](#) (page 145), associate it with the VM by clicking *Virtual Machines* → ⋮ (Options) → *Devices*, clicking *ADD*, and selecting *Disk* as the *Type*.

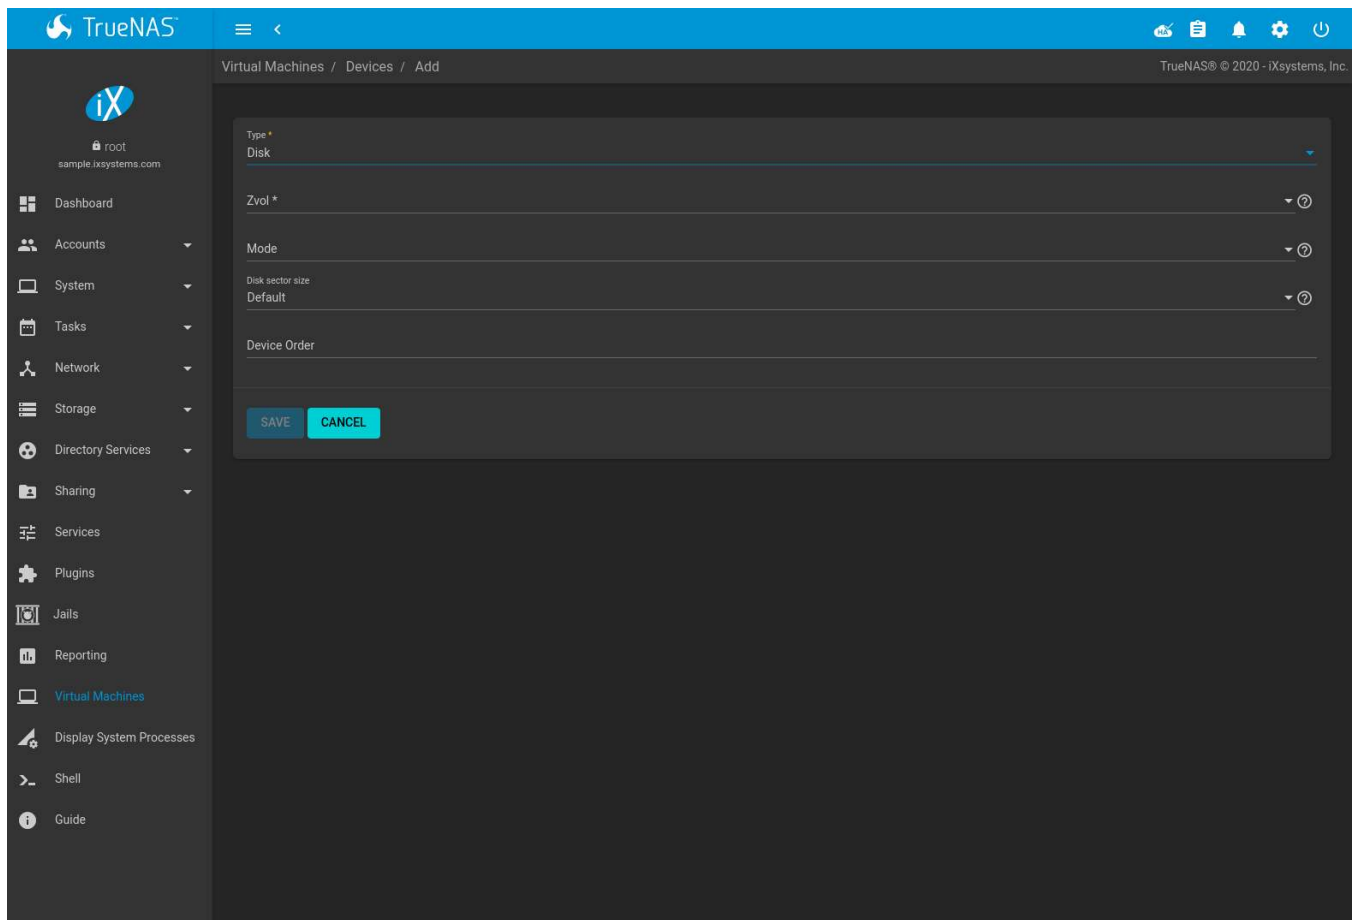


Fig. 16.7: Disk Device

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. This is recommended for Windows VMs.
- *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.

16.3.4 Raw Files

Raw Files are similar to *Zvol* (page 145) 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 TrueNAS® system, click *Virtual Machines* → ⋮ (Options) → *Devices*, click *ADD*, then set the *Type* to *Raw File*.


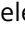


Fig. 16.8: Raw File Disk Device

Click  (Browse) to select the image file. If a specific sector size is required, choose it from *Disk sector size*. The *Default* value automatically selects a preferred sector size for the file.

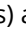
Setting disk *Mode* to *AHCI* emulates an AHCI hard disk for best software compatibility. *VirtIO* uses paravirtualized drivers and can provide better performance, but requires the operating system installed in the VM to support *VirtIO* disk devices.

Set a *Device Order* number to determine the boot order of this device. A lower number means a higher boot priority.



Set the size of the file in GiB.

16.3.5 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.

Each VM can have a single VNC device. An existing VNC interface can be changed by clicking  (Options) and *Edit*.

Note: Using a non-US keyboard with 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 16.9 shows the fields that appear after going to *Virtual Machines* →  (Options) → *Devices*, and clicking  (Options) → *Edit* for VNC.




Fig. 16.9: VNC Device

Setting *Port* to 0 automatically assigns a port when the VM is started. If a fixed, preferred port number is needed, enter it here.

Set *Delay VM Boot until VNC Connects* to wait to start the VM until a VNC client connects.

Resolution sets the default screen resolution used for the VNC session.

Use *Bind* to select the IP address for VNC connections.

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*, *ProtocolVersion*, 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.

VCENTER PLUGIN

vCenter Server (<https://www.vmware.com/products/vcenter-server.html>) is server management software that uses a single console to manage a virtual infrastructure across a hybrid cloud of physical and virtual machines. The TrueNAS[®] vCenter Plugin makes it possible to provision and use TrueNAS[®] storage from within vCenter Server.

For more information, please contact iXsystems Support at support@ixsystems.com or by phone:

- US-only toll-free: 855-473-7449 option 2
- Local and international: 408-943-4100 option 2

ADDITIONAL OPTIONS

This section covers the remaining miscellaneous options available from the TrueNAS® graphical administrative interface.

18.1 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 18.1.




Fig. 18.1: System Processes Running on TrueNAS®

The display automatically refreshes itself. The display is read-only.

18.2 Shell

The TrueNAS® web interface provides a web shell, making it convenient to run command line tools from the web browser as the *root* user.




Fig. 18.2: Web Shell

The prompt shows that the current user is *root*, the hostname is *freenas*, and the current working directory is *~*, the home directory of the logged-in user.

Note: The default shell for a new install of TrueNAS® is *zsh* (<https://www.freebsd.org/cgi/man.cgi?query=zsh>). TrueNAS® systems which have been upgraded from an earlier version will continue to use *cs*h as the default shell. The default shell can be changed in *Accounts* → *Users*. Click **⋮** (Options) and *Edit* for the *root* user. Choose the desired shell from the *Shell* drop-down and click *SAVE*.

The *Set font size* slider adjusts the size of text displayed in the Shell. Click *RESTORE DEFAULT* to reset the shell font and size.

A history of previous commands is available. Use the up and down arrow keys to scroll through previously entered commands. Edit the command if desired, then press *Enter* to re-enter the command.

Home, End, and Delete keys are supported. Tab completion is also available. Type a few letters and press Tab to complete a command name or filename in the current directory. Right-clicking in the terminal window displays a reminder about using *Command+c* and *Command+v* or *Ctrl+Insert* and *Shift+Insert* for copy and paste operations in the TrueNAS® shell.

Type *exit* to leave the session.

Clicking other web interface menus closes the shell session and stops commands running in the shell.

Note: Not all shell features render correctly in Chrome. Firefox is the recommended browser when using the shell.

Most FreeBSD command line utilities are available in the *Shell*.

18.3 Log Out, Restart, or Shut Down

The ⏻ (Power) button is used to log out of the web interface or restart or shut down the TrueNAS® system.

18.3.1 Log Out

To log out, click ⏻ (Power), then *Log Out*. After logging out, the login screen is displayed.

18.3.2 Restart

Click *Restart* shows the warning message in Figure 18.3.




Fig. 18.3: Restart Warning Message

If a scrub or resilver is in progress when a restart is requested, an additional warning asks if you wish to proceed. In this case, it is recommended to *Cancel* the restart request and to periodically run `zpool status` from *Shell*

(page 302) until it is verified that the scrub or resilver process is complete. Once complete, the restart request can be re-issued.

Click the *Cancel* button to cancel the reboot request. Otherwise, set *Confirm* and click *Reboot* to reboot the system. Rebooting the system disconnects all clients, including the web interface. Wait a few minutes for the system to boot. If the login screen does not appear, access the system using IPMI to determine if a problem is preventing the system from resuming normal operation.

18.3.3 Shut down

Clicking *Shut down* shows the warning message in Figure 18.4.




Fig. 18.4: Shutdown Warning Message

If a scrub or resilver is in progress when a shut down is requested, an additional warning will ask for confirmation to proceed. In this case, it is recommended to *Cancel* the shutdown request and to periodically run `zpool status` from *Shell* (page 302) until it is verified that the scrub or resilver process is complete. Once complete, the shut down request can be re-issued.

Click the *Cancel* button to cancel the shutdown request. Otherwise, set *Confirm* and click *SHUT DOWN* to halt the system. Shutting down the system will disconnect all clients, including the web interface, and will power off the TrueNAS® system. If the system has High Availability (HA) with *Failover* (page 81) enabled, the system failover to the standby TrueNAS controller.

18.4 Alert

The TrueNAS® 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 18.5, the system is warning that a pool is degraded.




Fig. 18.5: Example Alert Message

Table 18.1 shows the icons that indicate notification, warning, critical, and one-shot critical alerts. Critical messages are also emailed to the root account. One-shot critical alerts must be dismissed by the user.

Table 18.1: TrueNAS® Alert Icons

Alert Level	Icon
Notification	🔔
Warning	⚠️
Critical	🔴
One-shot Critical	🔔🔴

Close an alert message by clicking *Dismiss*. There is also an option to *Dismiss All Alerts*. Dismissing all alerts removes the notification badge from the alerts icon. Dismissed alerts can be re-opened by clicking *Re-Open*. 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 2.1) or the Web Shell (Figure

18.2) by running `midclt call alert.list`. Alert messages indicate which [High Availability \(HA\)](#) (page 81) TrueNAS controller generated the alert.

Notifications for specific alerts are adjusted in the [Alert Settings](#) (page 52) 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 [ZFS Feature Flags](#) (page 312) are available for the pool; this alert can be adjusted in [Alert Settings](#) (page 52) if a pool upgrade is not desired at present
- a new update is available
- hardware events detected by an attached [IPMI](#) (page 126) controller
- an error with the [Active Directory](#) (page 165) connection
- 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
- [syslog-ng\(8\)](#) (<https://www.freebsd.org/cgi/man.cgi?query=syslog-ng>) is not running
- a periodic snapshot or replication task fails
- a VMware login or a [VMware-Snapshots](#) (page 155) task fails
- a [Cloud Sync task](#) (page 113) 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
- an Active Directory domain goes offline; by default the winbindd connection manager will try to reconnect every 30 seconds and will clear the alert when the domain comes back online
- LDAP failed to bind to the domain
- any member interfaces of a lagg interface are not active
- a device is slowing pool I/O
- [Rsync task](#) (page 90) status
- a scrub has been paused for more than eight hours
- a connected Uninterruptible Power Supply (UPS) switches to battery power, switches to line power, communication with the UPS is lost or established, the battery is low, or the battery needs to be replaced
- a Fibre Channel (FC) Host Bus Adapter (HBA) configured as an iSCSI target is not detected
- the interface which is set as critical for failover is not found or is not configured
- attached SATADOM has 20% or less lifetime remaining
- NVDIMM problems
- HA is configured but the connection is not established
- one TrueNAS controller of an HA pair gets stuck applying its configuration journal as this condition could block future configuration changes from being applied to the standby TrueNAS controller
- TrueNAS controllers do not have the same number of connected disks
- the boot volume of the standby TrueNAS controller is not *HEALTHY*


-
- 30 days before the license expires, and when the license expires
 - the usage of a HA link goes above 10MB/s
 - an IPMI query to a standby TrueNAS controller fails, indicating the standby TrueNAS controller is down
 - *Proactive Support* (page 85) is enabled but any of the configuration fields are empty
 - ticket creation fails while Proactive Support is enabled
 - if VMware failed to log in (usually preceding a VMware snapshot)
 - if an unlicensed expansion shelf is connected
 - if a USB storage device has been attached which could prevent booting or failover
 - when the standby TrueNAS controller cannot be contacted
 - when it is 180, 90, 30, or 14 days before support contract expiration

Note: If *Proactive Support* (page 85) is enabled with Silver or Gold support coverage, and there is an internet connection, alerts which can indicate a hardware issue automatically create a support ticket with iXsystems Support. These alerts include a ZFS pool status change, a multipath failure, a failed S.M.A.R.T. test, and a failed re-key operation.

TASK MANAGER

The task manager shows a list of tasks performed by the TrueNAS® system starting with the most recent. Click a task name to display its start time, progress, finish time, and whether the task succeeded. If a task failed, the error status is shown.

Tasks with log file output have a *View Logs* button to show the log files.

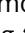
The task manager can be opened by clicking  (Task Manager). Close the task manager by clicking *CLOSE*, clicking anywhere outside the task manager dialog, or by pressing `ESC`.

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.

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. TrueNAS® automatically schedules a monthly scrub for each ZFS pool and the results of the scrub are displayed by selecting the [Pools](#) (page 130), 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 TrueNAS®, [Pools](#) (page 130) 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 TrueNAS® 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 TrueNAS®, 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 302) 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 302) 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. TrueNAS® 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 an adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of *random* reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 32 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 TrueNAS® 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 97) and [Replication Tasks](#) (page 107) 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 interruptible. 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%.
- 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 Part 1 \(video\)](https://www.youtube.com/watch?v=tPsV_8k-aVU) (https://www.youtube.com/watch?v=tPsV_8k-aVU) and [Becoming a ZFS Ninja Part 2 \(video\)](https://www.youtube.com/watch?v=wy6cJRVHiYU) (<https://www.youtube.com/watch?v=wy6cJRVHiYU>)
- [The Z File System \(ZFS\)](https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/zfs.html) (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/zfs.html)
- [ZFS: The Last Word in File Systems - Part 1 \(video\)](https://www.youtube.com/watch?v=aTXKxpL_0OI&list=PL5AD0E439599) (https://www.youtube.com/watch?v=aTXKxpL_0OI&list=PL5AD0E439599)
- [The Zettabyte Filesystem](https://www.youtube.com/watch?v=ptY6-K78McY) (<https://www.youtube.com/watch?v=ptY6-K78McY>)

20.1 ZFS Feature Flags

To differentiate itself from Oracle ZFS version numbers, OpenZFS uses feature flags. Feature flags are used to tag features with unique names 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. TrueNAS[®] uses OpenZFS and each new version of TrueNAS[®] keeps up-to-date with the latest feature flags and OpenZFS bug fixes.

See [zpool-features\(7\)](https://www.freebsd.org/cgi/man.cgi?query=zpool-features) (<https://www.freebsd.org/cgi/man.cgi?query=zpool-features>) for a complete listing of all OpenZFS feature flags available on FreeBSD.

VMWARE RECOMMENDATIONS

This section offers TrueNAS® configuration recommendations and troubleshooting tips when using TrueNAS® with a [VMware](https://www.vmware.com/) (https://www.vmware.com/) hypervisor.

21.1 TrueNAS® as a VMware Guest

This section has recommendations for configuring TrueNAS® when it is installed as a Virtual Machine (VM) in VMware.

Configure and use the [vmx\(4\)](https://www.freebsd.org/cgi/man.cgi?query=vmx) (https://www.freebsd.org/cgi/man.cgi?query=vmx) drivers for the TrueNAS® system.

Network connection errors for plugins or jails inside the TrueNAS® VM can be

caused by a misconfigured [virtual switch](https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.wssdk.pg.doc%2FPG_Networking.11.4.html) (https://pubs.vmware.com/vsphere-

51/index.jsp?topic=%2Fcom.vmware.wssdk.pg.doc%2FPG_Networking.11.4.html) or [VMware port group](https://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.server_configclassic.doc_40/esx_server)

(https://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.server_configclassic.doc_40/esx_server

Make sure MAC spoofing and promiscuous mode are enabled on the switch first, and then the port group the VM is using.

21.2 Hosting VMware Storage with TrueNAS®

This section has recommendations for configuring TrueNAS® when the system is being used as a VMware datastore.

Be sure to set up ALUA when using [iSCSI Sharing](#) (page 183) and VMware on a TrueNAS® High Availability (HA) system. This improves the resiliency of guest VMs during a [failover](#) (page 81) event.

Make sure guest VMs have the latest version of `vmware-tools` installed. VMware provides instructions to [install VMware Tools](#) (https://www.vmware.com/support/ws5/doc/new_guest_tools_ws.html) on different guest operating systems.

Increase the VM disk timeouts to better survive long disk operations. This also helps VMs deal with TrueNAS® High Availability (HA) [failovers](#) (page 81). Set the timeout to a minimum of *300 seconds*. See the guest operating system documentation for setting disk timeouts. VMware provides instructions for setting disk timeouts on some specific guest operating systems:

- Windows guest operating system: <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.storage.doc/GUID-EA1E1AAD-7130-457F-8894-70A63BD0623A.html>
- Linux guests running kernel version 2.6: <https://kb.vmware.com/s/article/1009465>

When TrueNAS® is used as a VMware datastore, [coordinated ZFS and VMware snapshots](#) (page 155) can be used.

21.3 VAAI for iSCSI

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.

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. This 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 TrueNAS®, this threshold is configured at the pool level when using zvols (see [Table 11.2](#)) or at the extent level (see [Table 11.7](#)) 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 TrueNAS® 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 TrueNAS® 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) (<https://tools.ietf.org/html/rfc2616>), such as GET, PUT, POST, or DELETE.

As shown in [Figure 22.1](#), an online version of the API is available at api.ixsystems.com/freenas (<https://api.ixsystems.com/freenas/>).




Fig. 22.1: API Documentation

The rest of this section shows code examples to illustrate the use of the API.

Note: A new API was released with TrueNAS® 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 TrueNAS® system at the `/api/docs/` URL. For example, if the TrueNAS® 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 TrueNAS® documentation will include relevant examples that make use of the new API.

22.1 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 TrueNAS® 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 user section at api.ixsystems.com/freenas (<https://api.ixsystems.com/freenas/>) describes this in more detail. Allowed parameters are listed in the JSON Pa-

rameters section of that resource. Since this resource creates a FreeBSD user, the values entered must be valid for a FreeBSD user account. [Table 22.1](#) summarizes acceptable values. This resource uses JSON, so the boolean values are *True* or *False*.

Table 22.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
<code>bsdusr_username</code>	string	Maximum 32 characters, though a maximum of 8 is recommended for interoperability. Can include numerals but cannot include a space.
<code>bsdusr_full_name</code>	string	May contain spaces and uppercase characters.
<code>bsdusr_password</code>	string	Can include a mix of upper and lowercase letters, characters, and numbers.
<code>bsdusr_uid</code>	integer	By convention, user accounts have an ID greater than 1000 with a maximum allowable value of 65,535.
<code>bsdusr_group</code>	integer	If <code>bsdusr_creategroup</code> is set to <i>False</i> , specify the numeric ID of the group to create.
<code>bsdusr_creategroup</code>	boolean	Set <i>True</i> to automatically create a primary group with the same numeric ID as <code>bsdusr_uid</code> .
<code>bsdusr_mode</code>	string	Sets default numeric UNIX permissions of a user home directory.
<code>bsdusr_shell</code>	string	Specify the full path to a UNIX shell that is installed on the system.
<code>bsdusr_password_disabled</code>	boolean	Set to <i>True</i> to disable user login.
<code>bsdusr_locked</code>	boolean	Set to <i>True</i> to disable user login.
<code>bsdusr_sudo</code>	boolean	Set to <i>True</i> to enable <code>sudo</code> for the user.
<code>bsdusr_sshpubkey</code>	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*.

22.2 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 *Ser-vices* resource.

```

1 def service_start(self, name):
2     self.request('services/services/%s' % name, method='PUT', data={
3         'srv_enable': True,
4     })
5

```

USER GUIDE

The TrueNAS® User Guide with complete configuration instructions is available either by clicking *Guide* in the TrueNAS® user interface or going to <https://www.ixsystems.com/documentation/truenas/>.

APPENDIX A: END-USER LICENSE AGREEMENT

TrueNAS® EULA:

Important - Please Read This EULA Carefully

PLEASE CAREFULLY READ THIS END USER LICENSE AGREEMENT (EULA) BEFORE CLICKING THE AGREE BUTTON. THIS AGREEMENT SERVES AS A LEGALLY BINDING DOCUMENT BETWEEN YOU AND IXSYSTEMS, INC. BY CLICKING THE AGREE BUTTON, DOWNLOADING, INSTALLING, OR OTHERWISE USING TRUENAS SOFTWARE, YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS AGREEMENT). IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS IN THIS AGREEMENT, DO NOT USE OR INSTALL TRUENAS SOFTWARE.

This agreement is provided in accordance with the Commercial Arbitration Rules of the American Arbitration Association (the “AAA Rules”) under confidential binding arbitration held in Santa Clara County, California. To the fullest extent permitted by applicable law, no arbitration under this EULA will be joined to an arbitration involving any other party subject to this EULA, whether through class arbitration proceedings or otherwise. Any litigation relating to this EULA shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California. All matters arising out of or relating to this agreement shall be governed by and construed in accordance with the internal laws of the State of California without giving effect to any choice or conflict of law provision or rule.

1.0 Definitions

1.1 “Company”, “iXsystems” and “iX” means iXsystems, Inc., on behalf of themselves, subsidiaries, and affiliates under common control.

1.2 “TrueNAS Software” means the TrueNAS storage management software.

1.3 “TrueNAS Device” means the TrueNAS hardware storage appliances and peripheral equipment.

1.4 “Product” means, individually and collectively, the TrueNAS Software and the TrueNAS Device.

1.5 “Open Source Software” means various open source software components licensed under the terms of applicable open source license agreements, each of which has its own copyright and its own applicable license terms.

1.6 “Licensee”, “You” and “Your” refers to the person, organization, or entity that has agreed to be bound by this EULA including any employees, affiliates, and third party contractors that provide services to You.

1.6 “Agreement” refers to this document, the TrueNAS End User License Agreement.

2.0 License

Subject to the terms set forth in this Agreement, iXsystems grants You a non-exclusive, non-transferable, revocable, limited license without the option to sublicense, to use TrueNAS Software on Your TrueNAS Device(s) in accordance with Your authorized purchase and use of a TrueNAS Device(s) for Your internal business purposes. This use includes but is not limited to using or viewing the instructions, specifications, and documentation provided with the Product.

3.0 License Restrictions

TrueNAS Software is only authorized for use with a TrueNAS Device identified by a specific serial number and manufactured by iXsystems. This license may be extended to a second TrueNAS Device if an additional TrueNAS Device was purchased for high availability data protection. The Product, including the TrueNAS Software, is protected by copyright laws and international treaties, as well as other intellectual property laws, statutes, and treaties. The Product is licensed, not sold to You the end user. You do not acquire any ownership interest in the

Product, including TrueNAS Software, or any other rights to such Product, other than to use such Product in accordance with the license granted under this Agreement, subject to all terms, conditions, and restrictions. iXsystems reserves and shall retain its entire right, title, and interest in and to the Product, and all intellectual property rights arising out of or relating to the Product, subject to the license expressly granted to You in this Agreement.

The Product, including TrueNAS Software, may contain iXsystems' trademarks, trade secrets, and proprietary collateral. iXsystems strictly prohibits the acts of decompiling, reverse engineering, or disassembly of the Product, including TrueNAS Software. You agree to use commercially reasonable efforts to safeguard the Product and iXsystems' intellectual property, trade secrets, or other proprietary information You may have access to, from infringement, misappropriation, theft, misuse, or unauthorized access. You will promptly notify iXsystems if You become aware of any infringement of the Product and cooperate with iXsystems in any legal action taken by iXsystems to enforce its intellectual property rights. By accepting this Agreement, You agree You will not disclose, copy, transfer, or publish benchmark results relating to the Product without the express written consent of iXsystems. You agree not to use, or permit others to use, the Product beyond the scope of the license granted under Section 2, unless otherwise permitted by iXsystems, or in violation of any law, regulation or rule, and you will not modify, adapt, or otherwise create derivative works or improvements of the Product. You are responsible and liable for all uses of the Product through access thereto provided by You, directly or indirectly.

4.0 General

4.1 Entire Agreement - This Agreement, together with any associated purchase order, service level agreement, and all other documents and policies referenced herein, constitutes the entire and only agreement between You and iXsystems for use of the TrueNAS Software and the TrueNAS Device and all other prior negotiations, representations, agreements, and understandings are superseded hereby. No agreements altering or supplementing the terms hereof may be made except by means of a written document signed by Your duly authorized representatives and those of iXsystems.

4.2 Waiver and Modification - No failure of either party to exercise or enforce any of its rights under this EULA will act as a waiver of those rights. This EULA may only be modified, or any rights under it waived, by a written document executed by the party against which it is asserted.

4.3 Severability - If any provision of this EULA is found illegal or unenforceable, it will be enforced to the maximum extent permissible, and the legality and enforceability of the other provisions of this EULA will not be affected.

4.4 United States Government End Users - For any Product licensed directly or indirectly on behalf of a unit or agency of the United States Government, this paragraph applies. Company's proprietary software embodied in the Product: (a) was developed at private expense and is in all respects Company's proprietary information; (b) was not developed with government funds; (c) is Company's trade secret for all purposes of the Freedom of Information Act; (d) is a commercial item and thus, pursuant to Section 12.212 of the Federal Acquisition Regulations (FAR) and DFAR Supplement Section 227.7202, Government's use, duplication or disclosure of such software is subject to the restrictions set forth by the Company and Licensee shall receive only those rights with respect to the Product as are granted to all other end users.

4.5 Foreign Corrupt Practices Act - You will comply with the requirements of the United States Foreign Corrupt Practices Act (the "FCPA") and will refrain from making, directly or indirectly, any payments to third parties which constitute a breach of the FCPA. You will notify Company immediately upon Your becoming aware that such a payment has been made. You will indemnify and hold harmless Company from any breach of this provision.

4.6. Title - iXsystems retains all rights, titles, and interest in TrueNAS Software and in TrueNAS Device(s) and all related copyrights, trade secrets, patents, trademarks, and any other intellectual and industrial property and proprietary rights, including registrations, applications, registration keys, renewals, and extensions of such rights.

4.7 Contact Information - If You have any questions about this Agreement, or if You want to contact iXsystems for any reason, please email legal@ixsystems.com.

4.8 Maintenance and Support - You may be entitled to support services from iXsystems after purchasing a TrueNAS Device or a support contract. iXsystems will provide these support services based on the length of time of the purchased support contract. This maintenance and support is only valid for the length of time that You have purchased with Your TrueNAS Device. iXsystems may from time to time and at their sole discretion vary the terms and conditions of the maintenance and support agreement based on different business environmental and personnel factors. For more information on our Maintenance and Support contract, refer to https://ixsystems.com/TrueNAS_SLA.

4.9 Force Majeure - iXsystems will not be deemed to be in default of any of the provisions of this Agreement or be liable for any delay or failure in performance due to Force Majeure, which shall include without limitation acts of God, earthquake, weather conditions, labor disputes, changes in law, regulation or government policy, riots, war, fire, epidemics, acts or omissions of vendors or suppliers, equipment failures, transportation difficulties, malicious or criminal acts of third parties, or other occurrences which are beyond iXsystems' reasonable control.

4.10 Termination - iXsystems may terminate or suspend Your license to use the Product and cease any and all support, services, or maintenance under this Agreement without prior notice, or liability, and for any reason whatsoever, without limitation, if any of the terms and conditions of this Agreement are breached. Upon termination, rights to use the Product will immediately cease. Other provisions of this Agreement will survive termination including, without limitation, ownership provisions, warranty disclaimers, indemnity, and limitations of liability.

4.11 Open Source Software Components - iXsystems uses Open Source Software components in the development of the Product. Open Source Software components that are used in the Product are composed of separate components each having their own trademarks, copyrights, and license conditions.

4.12 Assignment - Licensee shall not assign or otherwise transfer any of its rights, or delegate or otherwise transfer any of its obligations or performance, under this Agreement, in each case whether voluntarily, involuntarily, by operation of law, or otherwise, without iXsystems' prior written consent. No delegation or other transfer will relieve Licensee of any of its obligations or performance under this Agreement. Any purported assignment, delegation, or transfer in violation of this Section is void. iXsystems may freely assign or otherwise transfer all or any of its rights, or delegate or otherwise transfer all or any of its obligations or performance, under this Agreement without Licensee's consent. This Agreement is binding upon and inures to the benefit of the parties hereto and their respective permitted successors and assigns.

5.0 Export Control Regulations

The Product may be subject to US export control laws, including the US Export Administration Act and its associated regulations. You shall not, directly or indirectly, export, re-export, or release the Product to, or make the Product accessible from, any jurisdiction or country to which export, re-export, or release is prohibited by law, rule, or regulation. You shall comply with all applicable federal laws, regulations, and rules, and complete all required undertakings (including obtaining any necessary export license or other governmental approval), prior to exporting, re-exporting, releasing, or otherwise making the Product available outside the US.

6.0 Data Collection and Privacy

TrueNAS Software may collect information relating to Your use of the Product, including information that has been provided directly or indirectly through automated means. Usage of TrueNAS Software, geolocation information, user login credentials, and device and operating system identification are allowed according to iXsystems' [privacy policy](https://www.ixsystems.com/privacy-policy/) (<https://www.ixsystems.com/privacy-policy/>). By accepting this Agreement and continuing to use the Product, you agree that iXsystems may use any information provided through direct or indirect means in accordance with our privacy policy and as permitted by applicable law, for purposes relating to management, compliance, marketing, support, security, update delivery, and product improvement.

7.0 Limitation of Liability and Disclaimer of Warranty

THE PRODUCT IS PROVIDED "AS IS" AND WITH ALL FAULTS AND DEFECTS WITHOUT WARRANTY OF ANY KIND. TO THE MAXIMUM EXTENT PERMITTED UNDER APPLICABLE LAW, IXSYSTEMS, ON ITS OWN BEHALF AND ON BEHALF OF ITS AFFILIATES AND ITS AND THEIR RESPECTIVE LICENSORS AND SERVICE PROVIDERS, EXPRESSLY DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE, WITH RESPECT TO THE PRODUCT, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NON-INFRINGEMENT, AND WARRANTIES THAT MAY ARISE OUT OF COURSE OF DEALING, COURSE OF PERFORMANCE, USAGE, OR TRADE PRACTICE. WITHOUT LIMITATION TO THE FOREGOING, IXSYSTEMS PROVIDES NO WARRANTY OR UNDERTAKING, AND MAKES NO REPRESENTATION OF ANY KIND THAT THE PRODUCT WILL MEET THE LICENSEE'S REQUIREMENTS, ACHIEVE ANY INTENDED RESULTS, BE COMPATIBLE, OR WORK WITH ANY OTHER SOFTWARE, APPLICATIONS, SYSTEMS, OR SERVICES, OPERATE WITHOUT INTERRUPTION, MEET ANY PERFORMANCE OR RELIABILITY STANDARDS OR BE ERROR FREE, OR THAT ANY ERRORS OR DEFECTS CAN OR WILL BE CORRECTED.

TO THE FULLEST EXTENT PERMITTED UNDER APPLICABLE LAW: (A) IN NO EVENT WILL IXSYSTEMS OR ITS AFFILIATES, OR ANY OF ITS OR THEIR RESPECTIVE LICENSORS OR SERVICE PROVIDERS, BE LIABLE TO LICENSEE, LICENSEE'S AFFILIATES, OR ANY THIRD PARTY FOR ANY USE, INTERRUPTION, DELAY, OR INABILITY TO USE THE

PRODUCT; LOST REVENUES OR PROFITS; DELAYS, INTERRUPTION, OR LOSS OF SERVICES, BUSINESS, OR GOODWILL; LOSS OR CORRUPTION OF DATA; LOSS RESULTING FROM SYSTEM OR SYSTEM SERVICE FAILURE, MALFUNCTION, OR SHUTDOWN; FAILURE TO ACCURATELY TRANSFER, READ, OR TRANSMIT INFORMATION; FAILURE TO UPDATE OR PROVIDE CORRECT INFORMATION; SYSTEM INCOMPATIBILITY OR PROVISION OF INCORRECT COMPATIBILITY INFORMATION; OR BREACHES IN SYSTEM SECURITY; OR FOR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY, SPECIAL, OR PUNITIVE DAMAGES, WHETHER ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, REGARDLESS OF WHETHER SUCH DAMAGES WERE FORESEEABLE AND WHETHER OR NOT IXSYSTEMS WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES; (B) IN NO EVENT WILL IXSYSTEMS' AND ITS AFFILIATES', INCLUDING ANY OF ITS OR THEIR RESPECTIVE LICENSORS' AND SERVICE PROVIDERS', COLLECTIVE AGGREGATE LIABILITY UNDER OR IN CONNECTION WITH THIS AGREEMENT OR ITS SUBJECT MATTER, UNDER ANY LEGAL OR EQUITABLE THEORY, INCLUDING BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY, AND OTHERWISE, EXCEED THE TOTAL AMOUNT PAID TO IXSYSTEMS PURSUANT TO THIS AGREEMENT FOR THE PRODUCT THAT IS THE SUBJECT OF THE CLAIM; (C) THE LIMITATIONS SET FORTH IN THIS SECTION SHALL APPLY EVEN IF THE LICENSEE'S REMEDIES UNDER THIS AGREEMENT FAIL OF THEIR ESSENTIAL PURPOSE.

You hereby acknowledge that you have read and understand this Agreement and voluntarily accept the duties and obligations set forth herein by clicking accept on this Agreement.

APPENDIX B: TRUENAS® PRODUCT CATALOG


iXsystems offers many storage products that are designed to fully use the many features of TrueNAS®. These products have enterprise features like High Availability (HA), high-efficiency storage optimization, and fast networking speeds for data transfer. Scalability and modularity are also featured.

Each product makes full use of TrueNAS® and have a variety of features at different price points. To see hardware specifications or request a quote, visit the [iXsystems website](https://www.ixsystems.com/truenas) (<https://www.ixsystems.com/truenas>) or contact an iX representative at sales@ixsystems.com.

25.1 TrueNAS® Unified Storage Arrays

TrueNAS® storage products are divided into two families, the X-Series and M-Series. Each product supports High Availability (HA) hardware features and all software features included with TrueNAS®. The individual specifications of each unit are tuned to maximize efficiency, reliability, and affordability for different size IT environments. iXsystems provides [detailed specifications and comparisons](https://static.ixsystems.co/uploads/2019/07/StorageProductMatrix_Datasheet_WEB.pdf) (https://static.ixsystems.co/uploads/2019/07/StorageProductMatrix_Datasheet_WEB.pdf) of each TrueNAS® product.

25.1.1 X-Series



The [X-Series](https://static.ixsystems.co/uploads/2019/07/BSG-X-Series-1.3_screen.pdf) (https://static.ixsystems.co/uploads/2019/07/BSG-X-Series-1.3_screen.pdf) are 2U, 12-bay (front-loading), hybrid data storage arrays. They support single or dual TrueNAS controller configurations and can be connected with expansion shelves to increase storage capacity. Physical dimensions are 21"L x 19"W x 3.5"H (531 x 447 x 89 mm). They weigh 44 lbs (20 kg) and draw 138-200 W of power when fully loaded with drives.

25.1.2 M-Series



The **M-Series** (https://static.ixsystems.co/uploads/2019/07/BSG-M-Series_screen.pdf) are 4U, 24-bay (front-loading), hybrid data storage arrays. They support single or dual TrueNAS controller configurations and can be connected with expansion shelves to increase storage capacity to a maximum of 2-10.4 PB. Physical dimensions are 27"L x 19"W x 7"H (686 x 483 x 178 mm). They weigh 114 lbs (52 kg) when fully loaded with drives. Maximum power draw is determined by the total number of drives attached to the system. Without drives, M-Series units draw 404-905 W of power, depending on TrueNAS controller configuration.

25.2 Expansion Shelves

Expansion shelves are an easy way to expand TrueNAS® capacity. All TrueNAS® expansion shelves feature redundant, hot-swappable, high-efficiency power supplies, redundant cooling, and are **RoHS 6/6** (<https://www.rohsguide.com/rohs-faq.htm>) compliant.

25.2.1 ES12



The TrueNAS® **ES12** (https://static.ixsystems.co/uploads/2019/07/BSG-ES12-1.3_screen.pdf) is a 2U, 12-bay (front-loading), SAS3 (12 Gb/s) expansion shelf with dual expansion controllers and redundant power supplies. Physical dimensions are 21"L x 19"W x 3.5"H (531 x 447 x 89 mm). It weighs 44 lbs (20 kg) and draws 112-180 W of power when fully loaded with drives.

25.2.2 ES24



The TrueNAS® ES24 (https://static.ixsystems.co/uploads/2019/08/BSG-ES24-1.4_screen.pdf) is a 4U, 24-bay (front-loading), SAS3 (12 Gb/s) expansion shelf with dual expansion controllers and redundant power supplies. Physical dimensions are 20.5"L x 19"W x 7"H (521 x 483 x 178 mm). It weighs 76 lbs (34.5 kg) and draws 231-288 W of power when fully loaded with drives.

25.2.3 ES60



The TrueNAS® ES60 (https://static.ixsystems.co/uploads/2019/07/BSG-ES60_screen.pdf) is a 4U, 60-bay (top-loading), SAS3 (12 Gb/s) expansion shelf with dual expansion controllers and redundant power supplies. Physical dimensions are 33.38"L x 19"W x 6.9"H (848 x 483 x 176 mm). It weighs 175 lbs (80 kg) and draws 479-821 W of power when fully loaded with drives.