

TrueNAS® 11.1-U4 User Guide

March 2018 Edition

Copyright iXsystems 2011-2018

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

Welcome	1
Typographic Conventions	2
1 Introduction	3
1.1 Path and Name Lengths	3
2 Initial Setup	4
2.1 Out-of-Band Management	4
2.2 Console Setup Menu	14
2.3 Accessing the Administrative GUI	16
3 Account	19
3.1 Groups	19
3.2 Users	22
4 System	26
4.1 Information	26
4.2 General	27
4.3 Boot	30
4.4 Advanced	33
4.4.1 Autotune	34
4.5 Email	34
4.6 System Dataset	36
4.7 Tunables	36
4.8 Update	39
4.8.1 Preparing for Updates	39
4.8.2 Updates and Trains	40
4.8.3 Checking for Updates	41
4.8.4 Applying Updates	41
4.8.5 Manual Updates	41
4.8.6 Updating from the Shell	41
4.8.7 Updating an HA System	41
4.8.8 If Something Goes Wrong	42
4.8.9 Upgrading a ZFS Pool	42
4.9 Cloud Credentials	43
4.10 Alert Services	44
4.10.1 How it Works	45
4.11 CAs	45
4.12 Certificates	48
4.13 Support	51
4.14 Proactive Support	53
4.15 View Enclosure	54
4.16 Failover	55
5 Tasks	58

5.1	Cloud Sync	58
5.1.1	Cloud Sync Example	61
5.2	Cron Jobs	63
5.3	Init/Shutdown Scripts	65
5.4	Rsync Tasks	66
5.4.1	Rsync Module Mode	70
5.4.2	Rsync over SSH Mode	70
5.5	S.M.A.R.T. Tests	73
6	Network	76
6.1	Global Configuration	76
6.2	Interfaces	78
6.3	IPMI	80
6.4	Link Aggregations	82
6.4.1	LACP, MPIO, NFS, and ESXi	82
6.4.2	Creating a Link Aggregation	83
6.5	Network Summary	86
6.6	Static Routes	86
6.7	VLANs	87
7	Storage	89
7.1	Volumes	89
7.1.1	Volume Manager	89
	Encryption	91
	Encryption Performance	92
	Manual Setup	92
	Extending a ZFS Volume	93
7.1.2	Change Permissions	94
7.1.3	Create Dataset	96
	Compression	98
7.1.4	Create zvol	99
7.1.5	Import Disk	100
7.1.6	Import Volume	101
	Importing an Encrypted Pool	101
7.1.7	View Disks	102
7.1.8	Volumes	104
	Managing Encrypted Volumes	106
	Additional Controls for Encrypted Volumes	107
7.1.9	View Multipaths	110
7.1.10	Replacing a Failed Drive	110
	Replacing an Encrypted Drive	112
	Removing a Log or Cache Device	112
7.1.11	Replacing Drives to Grow a ZFS Pool	112
7.1.12	Hot Spares	113
7.2	Periodic Snapshot Tasks	113
7.3	Replication Tasks	115
7.3.1	Examples: Common Configuration	115
	Alpha (Source)	115
	Beta (Destination)	116
7.3.2	Example: TrueNAS® to TrueNAS® Semi-Automatic Setup	116
7.3.3	Example: TrueNAS® to TrueNAS® Dedicated User Replication	118
7.3.4	Example: TrueNAS® to TrueNAS® or Other Systems, Manual Setup	119
	Encryption Keys	119
7.3.5	Replication Options	122
7.3.6	Replication Encryption	123
7.3.7	Limiting Replication Times	123
7.3.8	Replication Topologies and Scenarios	123

Star Replication	123
Tiered Replication	124
N-way Replication	124
Disaster Recovery	124
7.3.9 Troubleshooting Replication	124
SSH	124
Compression	125
Manual Testing	125
7.4 Resilver Priority	125
7.5 Scrubs	126
7.6 Snapshots	129
7.7 VMware-Snapshot	131
8 Directory Services	132
8.1 Active Directory	132
8.1.1 Troubleshooting Tips	136
8.1.2 If the System Will not Join the Domain	137
8.2 LDAP	137
8.3 NIS	140
8.4 Kerberos Realms	141
8.5 Kerberos Keytabs	141
8.6 Kerberos Settings	142
9 Sharing	143
9.1 Apple (AFP) Shares	144
9.1.1 Creating AFP Guest Shares	146
9.1.2 Creating Authenticated and Time Machine Shares	147
9.2 Unix (NFS) Shares	151
9.2.1 Example Configuration	155
9.2.2 Connecting to the Share	155
From BSD or Linux	155
From Microsoft	156
From Mac OS X	156
9.2.3 Troubleshooting NFS	157
9.3 WebDAV Shares	158
9.4 Windows (SMB) Shares	159
9.4.1 Configuring Unauthenticated Access	164
9.4.2 Configuring Authenticated Access Without a Domain Controller	165
9.4.3 Configuring Shadow Copies	168
9.5 Block (iSCSI)	169
9.5.1 Target Global Configuration	171
9.5.2 Portals	172
9.5.3 Initiators	173
9.5.4 Authorized Accesses	174
9.5.5 Targets	176
9.5.6 Extents	177
9.5.7 Target/Extents	180
9.5.8 Fibre Channel Ports	181
9.5.9 Connecting to iSCSI	184
9.5.10 Growing LUNs	185
Zvol Based LUN	185
File Extent Based LUN	186
10 Services	187
10.1 Control Services	187
10.2 AFP	189
10.2.1 Troubleshooting AFP	191
10.3 Domain Controller	191

10.3.1 Samba Domain Controller Backup	193
10.4 Dynamic DNS	193
10.5 FTP	194
10.5.1 Anonymous FTP	197
10.5.2 FTP in chroot	198
10.5.3 Encrypting FTP	199
10.5.4 Troubleshooting FTP	199
10.6 iSCSI	199
10.7 LLDP	200
10.8 Netdata	200
10.9 NFS	201
10.10 Rsync	203
10.10.1 Configure Rsyncd	203
10.10.2 Rsync Modules	204
10.11 S3	205
10.12 S.M.A.R.T.	206
10.13 SMB	208
10.13.1 Troubleshooting SMB	211
10.14 SNMP	212
10.15 SSH	213
10.15.1 SCP Only	215
10.15.2 Troubleshooting SSH	215
10.16 TFTP	216
10.17 UPS	216
10.17.1 Multiple Computers with One UPS	220
10.18 WebDAV	220
11 vCenter Plugin	222
11.1 New Features in 2.1.0	222
11.2 Configuration	222
11.3 Secure Connections	224
11.4 Using the vCenter Plugin	224
11.5 Hosts	224
11.5.1 Host Information	225
11.5.2 Manage Host	225
11.6 Datastores	225
11.6.1 VMFS Datastores	225
Creating VMFS Datastores	225
Extending VMFS Datastores	228
11.6.2 NFS Datastores	229
Creating NFS Datastores	229
11.6.3 Cloning Datastores	231
11.6.4 Snapshotting a Datastore	232
11.7 Snapshots	232
11.7.1 Scheduling: Periodic Snapshots	233
11.8 Replication	234
11.9 Zvols	236
11.10 iSCSI Functions	237
11.10.1 Targets	237
11.10.2 Extents	238
11.10.3 Associated Targets	239
11.11 Role-Based Access Control (RBAC)	240
11.11.1 Add a Role to an Existing vCenter User	240
11.12 Copying iSCSI Configuration	240
12 OpenStack Cinder Driver	241
12.1 Requirements	241

12.2	Installation	241
12.3	Configuration	241
12.4	Using the Cinder Driver	243
12.4.1	Managing Volumes with Cinder by Project	243
	Project Volume Snapshots	244
12.4.2	Administering All Volumes with Cinder	245
	Managing Volumes	245
	Manage Volume Types	246
	Manage Volume Snapshots	246
13	Reporting	248
14	Wizard	250
15	Additional Options	257
15.1	Display System Processes	257
15.2	Shell	257
15.3	Log Out	259
15.4	Reboot	259
15.5	Shutdown	259
15.6	Support Icon	260
15.7	Guide	260
15.8	Alert	260
16	ZFS Primer	263
17	Hardware Setup	267
17.1	Contacting iXsystems	267
17.2	Z Series Unified Storage Array	267
17.2.1	Check the Contents of the Box	268
17.2.2	Become Familiar with the System	269
17.2.3	Mounting in an EIA Rack or Cabinet	270
17.3	X Series Unified Storage Array	277
17.3.1	Become Familiar With the X Series System	278
17.3.2	Rail Kit Assembly	279
17.3.3	Install Drive Trays	280
17.3.4	Connect Expansion Shelves	281
17.3.5	Connect Network Cables	281
17.3.6	Connect Power Cords	281
17.3.7	Install Bezel (Optional)	281
17.3.8	Discover Out-of-Band Management IP Address	282
	Preset	282
	DHCP	282
	Serial Cable	282
17.3.9	Connect to the X Series Console	284
	With IPMI	284
	With the Serial Cable	285
17.3.10	Using the X Series Console	286
17.3.11	Perform TrueNAS® Initial Software Configuration	286
17.4	ES12 Expansion Shelf	287
17.4.1	Become Familiar With the ES12 System	288
17.4.2	Rail Kit Assembly	289
17.4.3	Install Drive Trays	290
17.4.4	Connect Power Cords	291
17.4.5	Connect SAS Cables	292
17.4.6	Use with the ES60 Expansion Shelf (X20 Only)	292
17.4.7	Install Bezel (Optional)	293
17.5	ES24 Expansion Shelf	293

17.5.1	Become Familiar With the ES24	294
17.5.2	Rail Kit Assembly	295
	Remove Cabinet Rails from Rack Rails	295
	Mount Cabinet Rails	295
	Mount the Rack Rails	296
	Mount the Unit in the Rack	296
17.5.3	Install Drive Trays	297
17.5.4	Connect Power Cords	298
17.5.5	Connect SAS Cables	298
17.5.6	Install Bezel (Optional)	299
17.6	ES60 Expansion Shelf	300
17.6.1	Become Familiar With the ES60	301
17.6.2	Rail Kit Assembly	302
	Separate Cabinet Rails from Rack Rails	302
	Mount Cabinet Rails	302
	Mount Rack Rails	302
17.6.3	Mount Unit in the Rack	303
17.6.4	Drive Tray Installation	304
	Remove Top Cover	304
	Install Drive Trays	304
17.6.5	ES60 Cable Management Arm	305
	Install the Cable Management Arm	306
17.6.6	Connect Power Cords	307
17.6.7	Connect SAS Cables	308
	Connecting the ES60 to a TrueNAS® X20	308
	Connecting the ES60 to a TrueNAS® Z35	308
17.6.8	Install Bezel (Optional)	310
18	VAAI	311
18.1	VAAI for iSCSI	311
18.2	VAAI for NAS	311
19	Using the API	312
19.1	A Simple API Example	313
19.2	A More Complex Example	314
20	Appendix A: End-User License Agreement	316
	Index	319

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.

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

INTRODUCTION

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 Path and Name Lengths

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

Table 1.1: Path and Name Lengths

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

INITIAL SETUP

Before beginning software configuration, please see the [Hardware Setup](#) (page 267) section for specific racking and connection information.

Depending on the degree of pre-configuration requested from iXsystems, most of the initial TrueNAS® setup might already be complete.

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. They are located on the back of each chassis.

2.1 Out-of-Band Management

Before attempting to configure TrueNAS® for out-of-band management, ensure that the out-of-band management port is connected to an appropriate network. Refer to the guide included with the TrueNAS® Storage Array for detailed instructions on how to connect to a network.

Connect the out-of-band management port before powering on the TrueNAS® Storage Array.

In most cases, the out-of-band management interface will have been pre-configured by iXsystems. This section contains instructions for configuring it from the BIOS if needed. The same settings can be configured using the instructions in [IPMI](#) (page 80).

Press **F2** at the splash screen while the TrueNAS® Storage Array is booting to access the system BIOS. This opens the menu shown in [Figure 2.1](#).

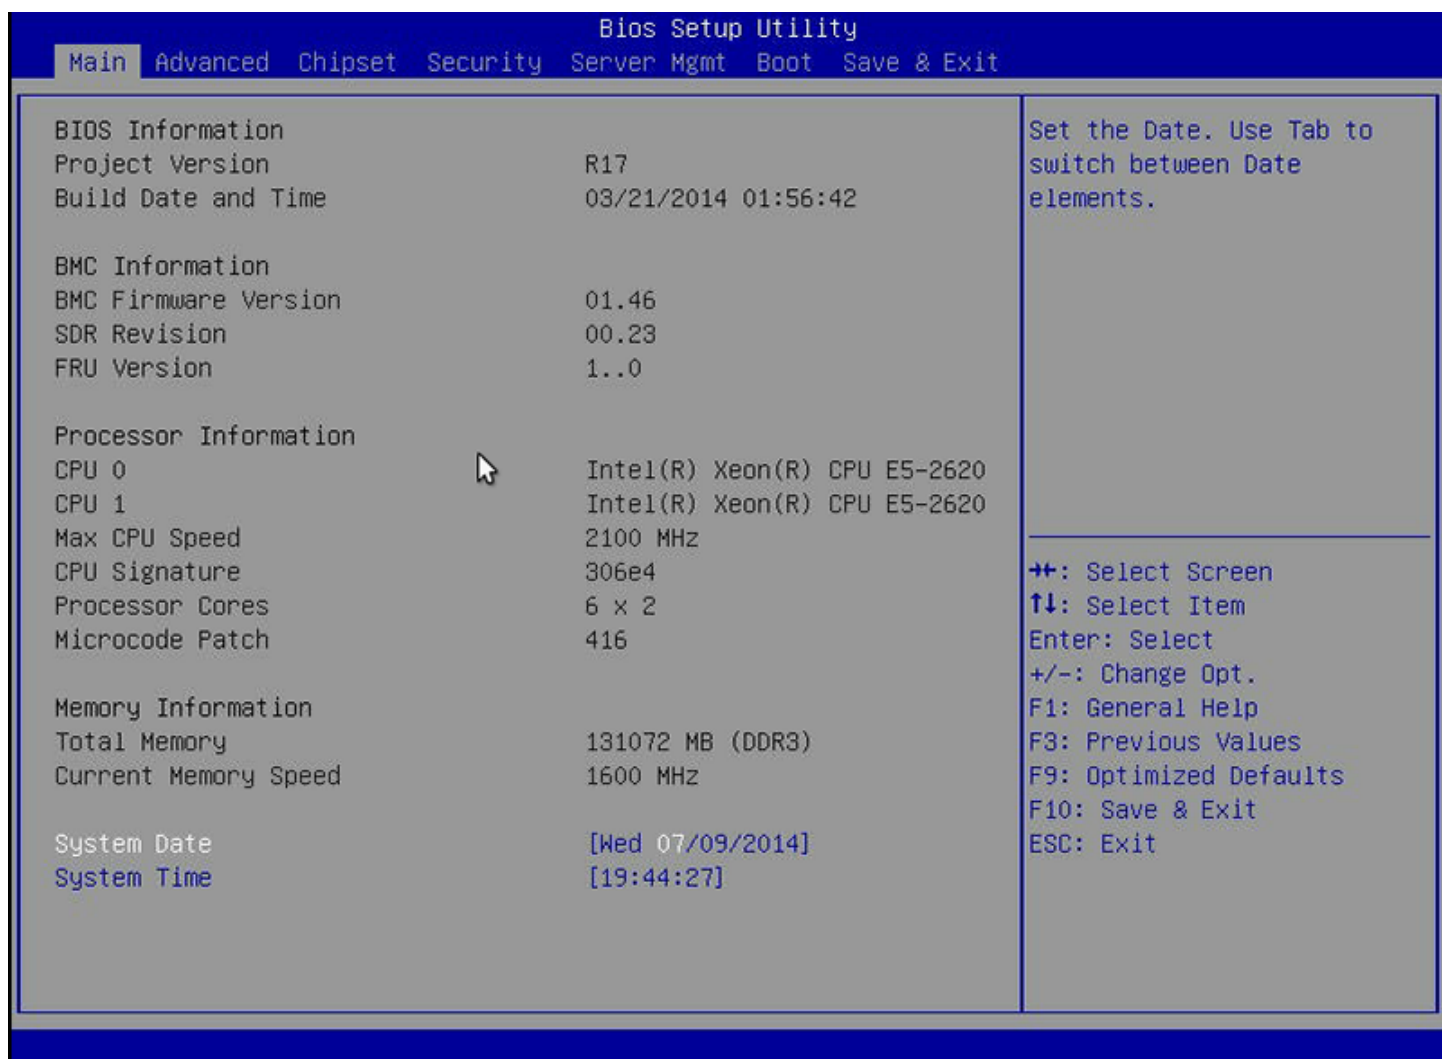


Fig. 2.1: Initial BIOS Screen

Navigate to the Server Mgmt menu and then BMC LAN Configuration, as shown in Figure 2.2.

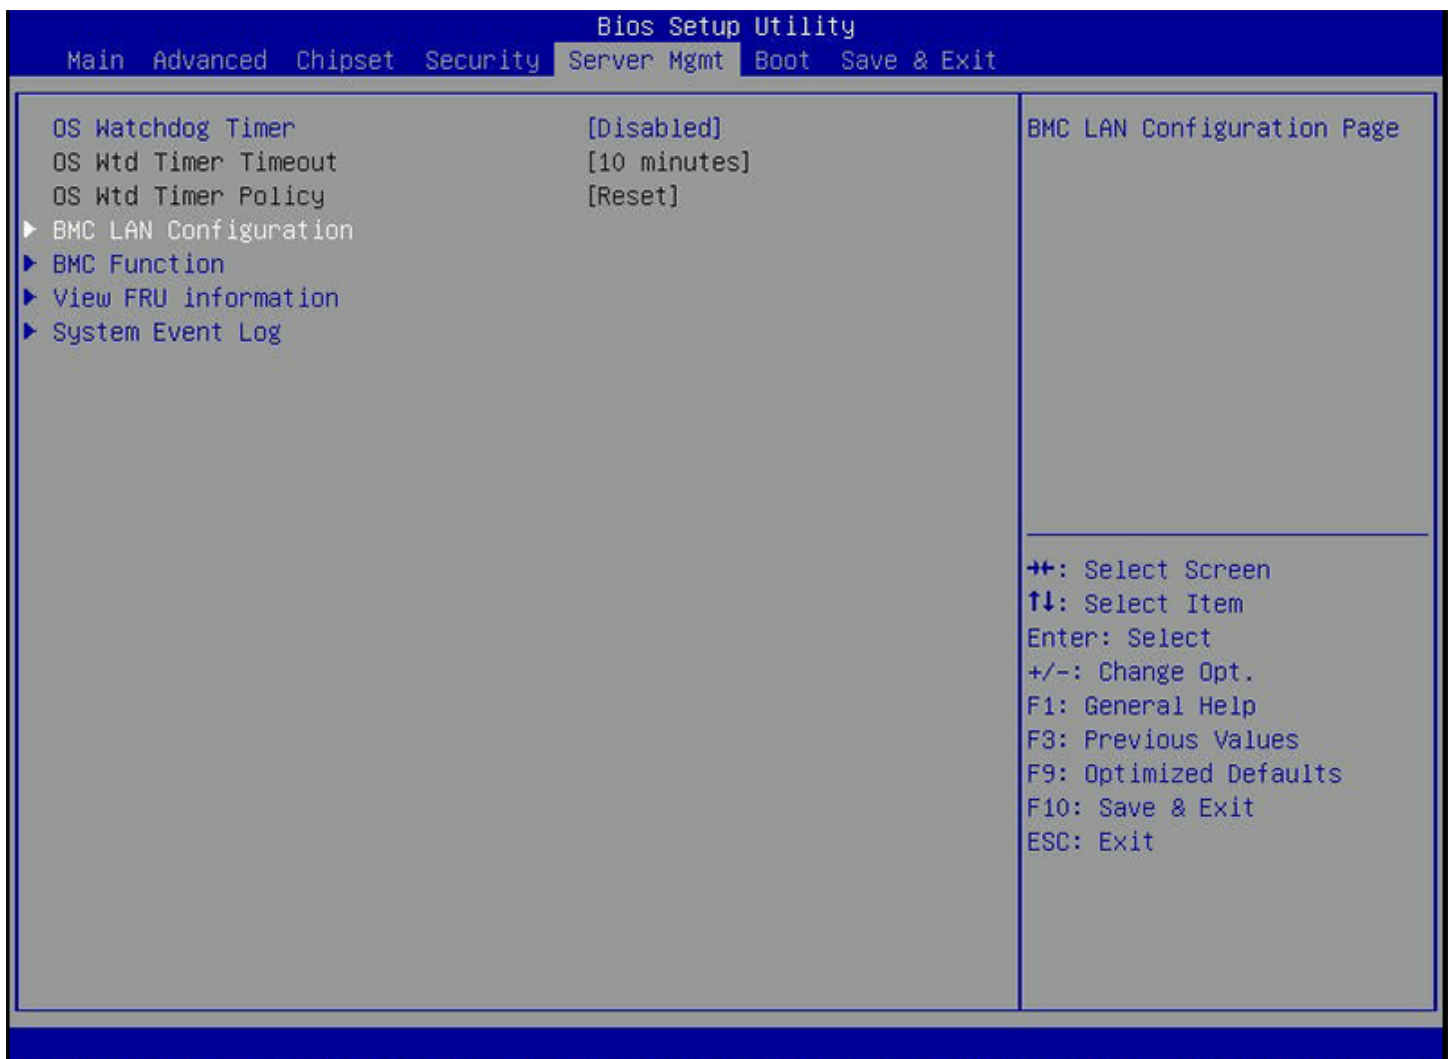


Fig. 2.2: Navigate to BMC LAN Configuration

When using DHCP to assign the out-of-band management IP address, leave the Configuration Source set to Dynamic in the screen shown in [Figure 2.3](#). If an IP has been assigned by DHCP, it is displayed.

Bios Setup Utility		
Server Mgmt		
Lan Channel 1		Select to configure LAN channel parameters statically or dynamically(DHCP). Do nothing option will not modify any BMC network parameters during BIOS phase
Configuration Source	[Dynamic]	
IP Address	010.005.000.149	
Subnet Mask	255.255.000.000	
Default Gateway Address	010.005.000.001	
BMC MAC Address	74-D4-35-43-44-2F	
Load BMC IP		
		↔: Select Screen ↑↓: Select Item Enter: Select +/-: Change Opt. F1: General Help F3: Previous Values F9: Optimized Defaults F10: Save & Exit ESC: Exit

Fig. 2.3: Configuring a Dynamic IP Address

To assign a static IP address for out-of-band management, set the Configuration Source to Static, as shown in Figure 2.4. Enter the desired IP Address into the IP Address setting, filling out all four octets completely.

Bios Setup Utility		
Server Mgmt		
Lan Channel 1		Enter station IP address
Configuration Source	[Static]	
IP Address	010.005.000.149	
Subnet Mask		
Default Gateway Address		
BMC MAC Address	74-D4-35-43-44-2F	
Load BMC IP		
<div>IP Address 010.005.000.149</div>		
		↔: Select Screen ↑↓: Select Item Enter: Select +/-: Change Opt. F1: General Help F8: Previous Values F9: Optimized Defaults F10: Save & Exit ESC: Exit

Fig. 2.4: Configuring a Static IP Address

Next, enter the Subnet Mask of the out-of-band management network subnet. An example is shown in [Figure 2.5](#).




Fig. 2.5: Entering the Subnet Mask

Finally, set the Default Gateway Address for the network to which the out-of-band management port is connected. An example is shown in [Figure 2.6](#).

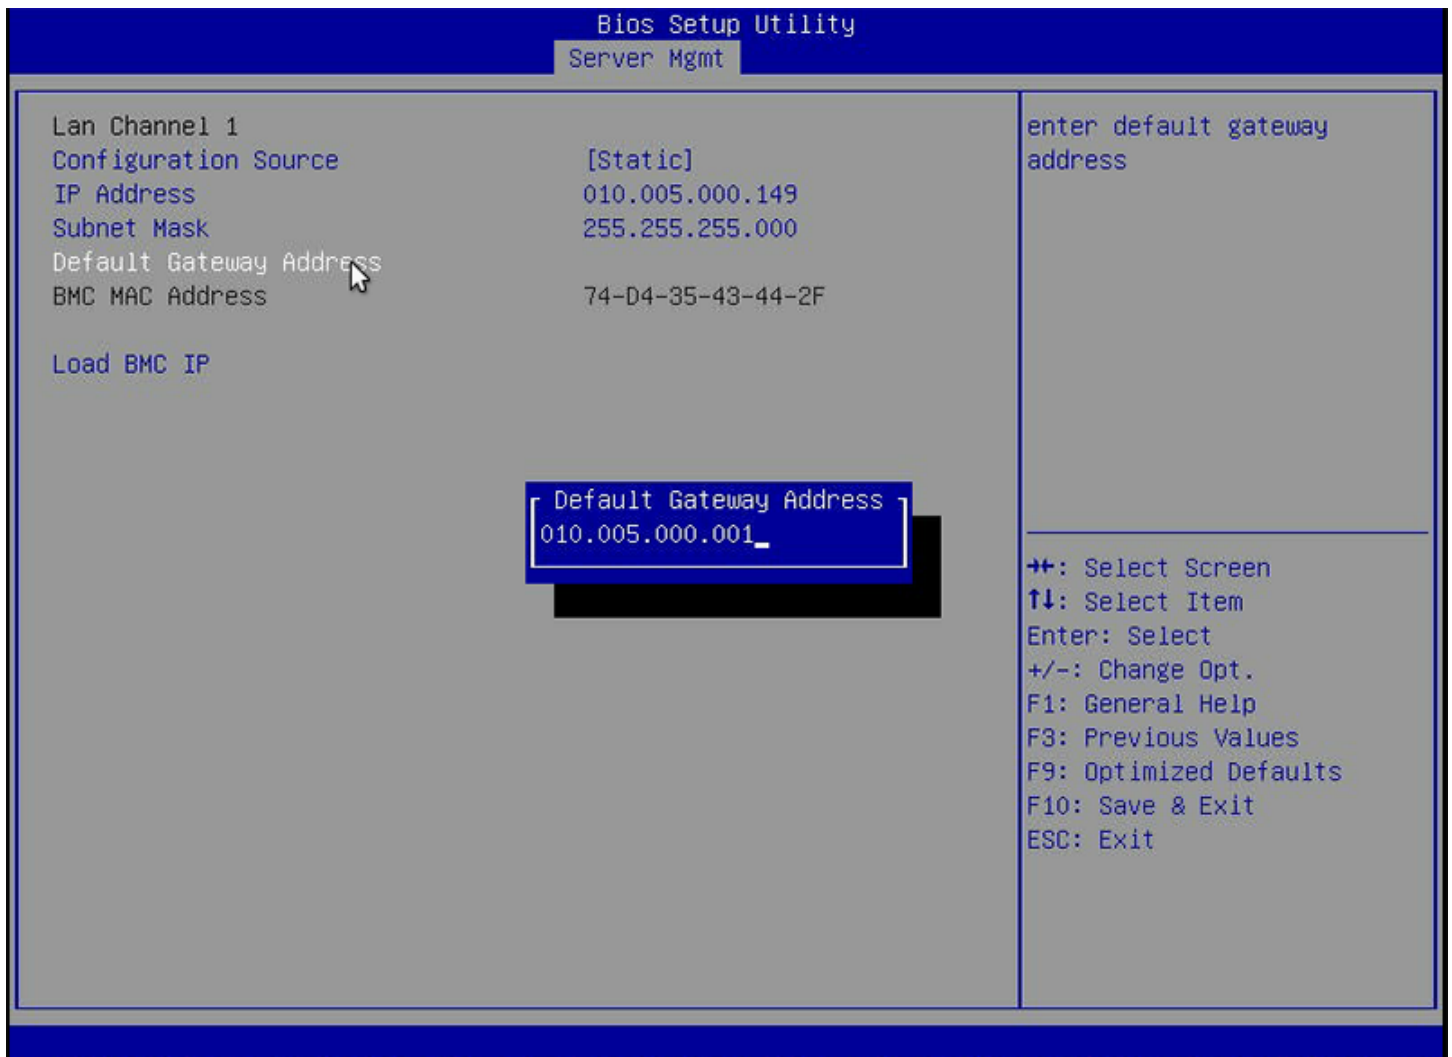


Fig. 2.6: Entering the Default Gateway Address

Save the changes, exit the BIOS, and allow the system to boot.

To connect to the TrueNAS® Storage Array's out-of-band management port, enter the IP address into a web browser from a computer that is either within the same network or which is directly wired to the array. As shown in [Figure 2.7](#), a login prompt appears.

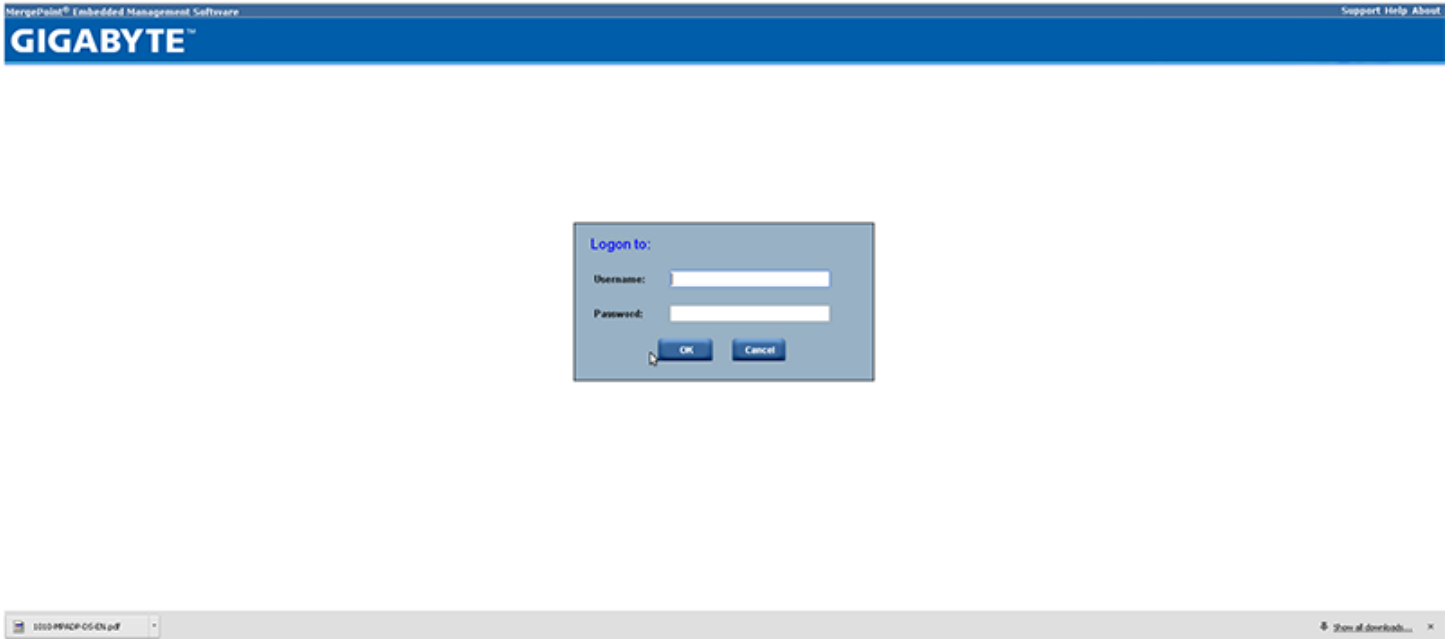


Fig. 2.7: Connecting to the IPMI Graphical Interface

Log in using the default Username of admin and the default Password of password.

The administrative password can be changed using the instructions in [IPMI](#) (page 80).

After logging in, click the vKVM and Media button at the top right to download the Java KVM Client. Run the client by clicking the Launch Java KVM Client button shown in [Figure 2.8](#).




Fig. 2.8: Launching the Java KVM Client

When prompted for a program to open the file with, select the Java Web Start Launcher shown in Figure 2.9.




Fig. 2.9: Configure the Launch Program

If asked to verify running a program from an unknown publisher, check the box indicating that you understand the risks and press Run. An example is shown in Figure 2.10.

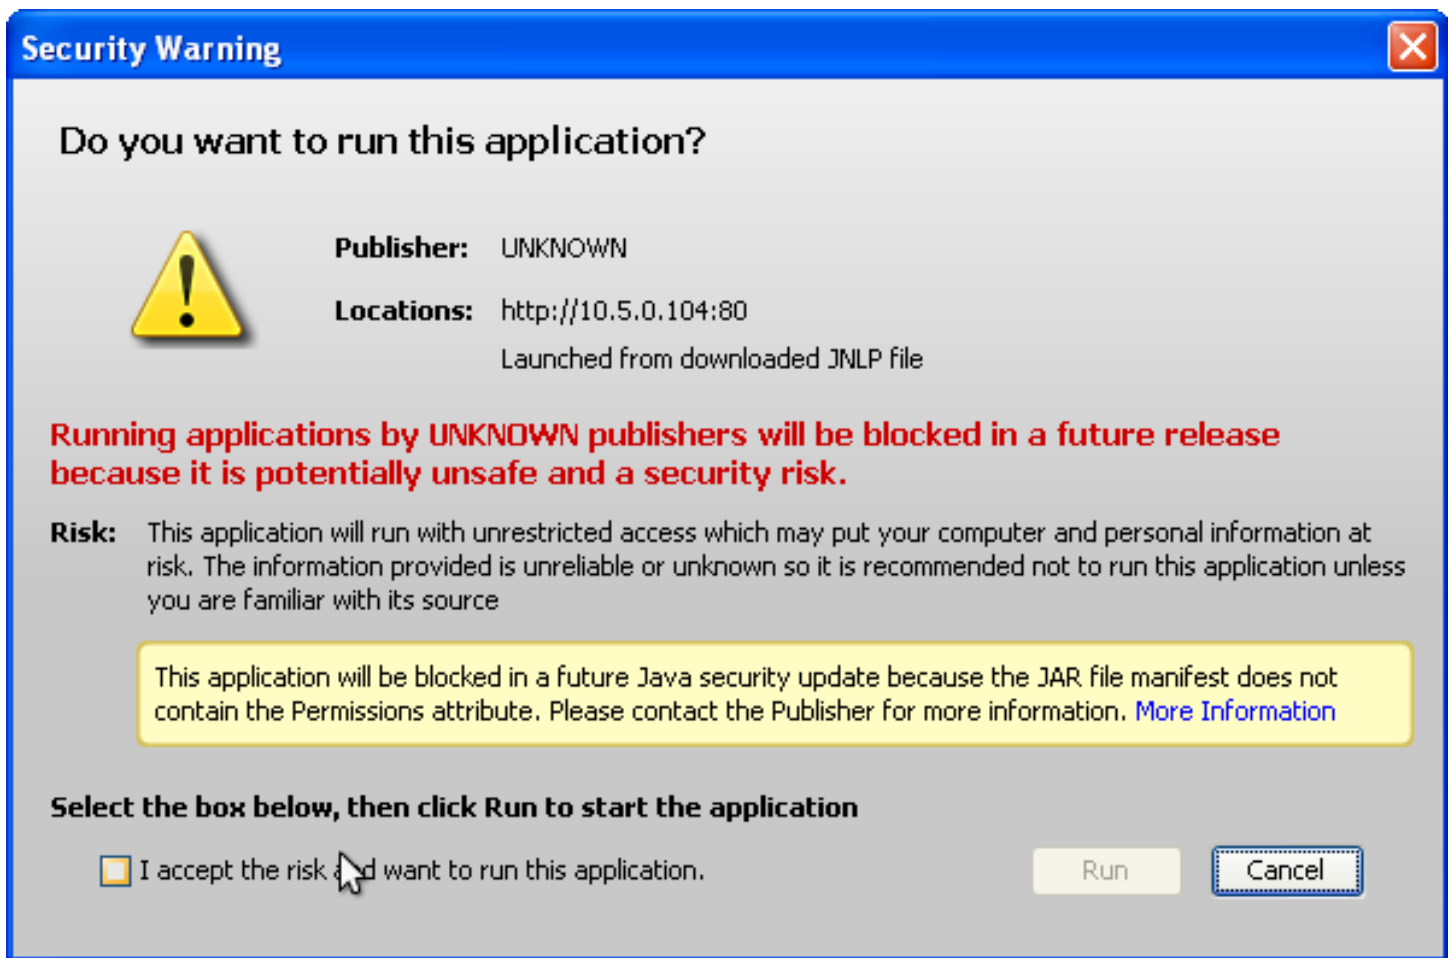


Fig. 2.10: Respond to Warning

When prompted that the connection is untrusted, as shown in [Figure 2.11](#), press Continue.

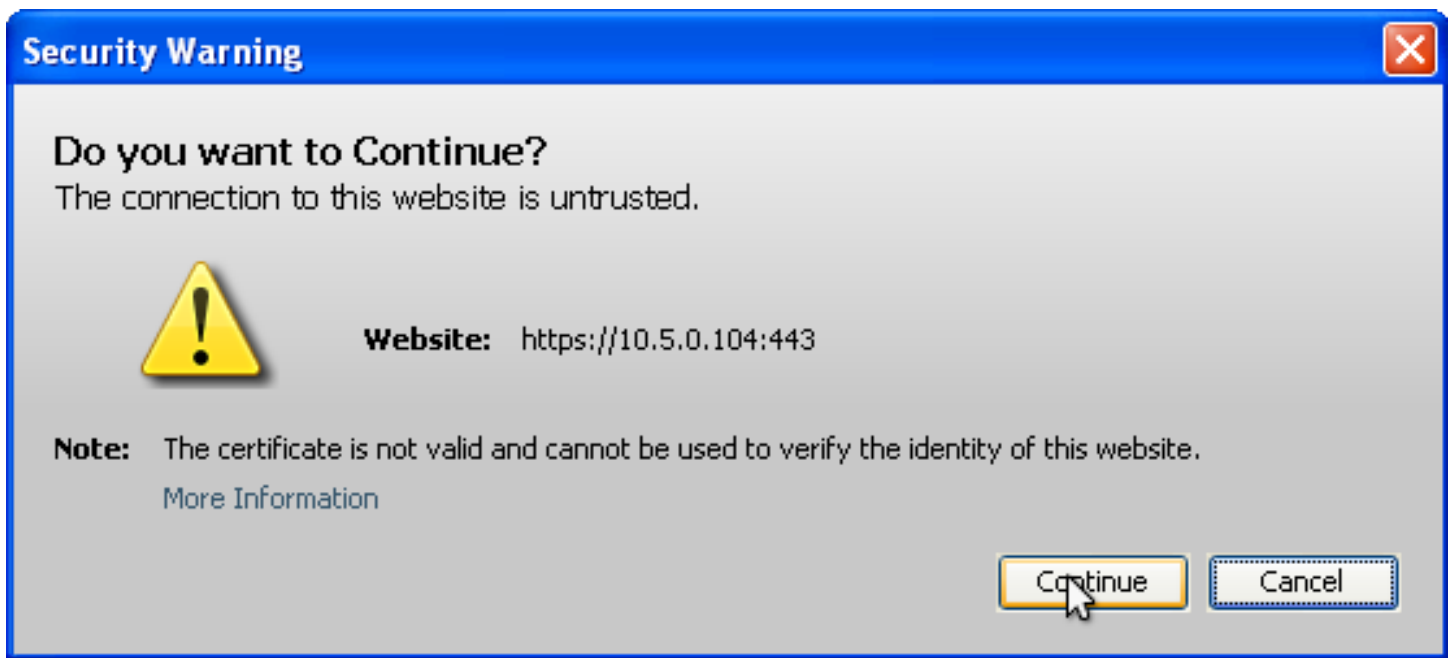


Fig. 2.11: Continue Through this Screen

With the out-of-band console open, the TrueNAS® Storage Array can be controlled as if using a directly-connected keyboard and monitor.

2.2 Console Setup Menu

The Console Setup menu, shown in [Figure 2.12](#), 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 257), the Console Setup menu is not shown by default. It can be started by the root user or another user with root permissions by typing `/etc/netcli`.

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

```
Console setup
-----

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

The web user interface is at:

http://10.0.0.102

Enter an option from 1-11: █
```

Fig. 2.12: Console Setup Menu

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 Node) and (Node B).
- 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 → Volumes → Import Volume` can be used to re-import volumes.
- 9) Shell starts a shell for running FreeBSD commands. To leave the shell, type `exit`.
- 10) Reboot reboots the system.
- 11) Shut Down shuts down the system.

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

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 2.12](#), 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
Reset network configuration (y/n) n
Configure interface for DHCP? (y/n) n
Configure IPv4? (y/n) y
Interface name:      (press enter, the name can be blank)
Several input formats are supported
Example 1 CIDR Notation:
    192.168.1.1/24
Example 2 IP and Netmask separate:
    IP: 192.168.1.1
    Netmask: 255.255.255.0, or /24 or 24
IPv4 Address: 192.168.1.108/24
Saving interface configuration: Ok
Configure IPv6? (y/n) n
Restarting network: ok

...

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

2.3 Accessing the Administrative GUI

After the system has an IP address, enter that address into a graphical web browser from a computer on the same network as the TrueNAS® system. A prompt appears to enter the password for the root user, as shown in [Figure 2.13](#).




Fig. 2.13: Enter the Root Password

Enter the default password of `abcd1234`.

Note: The default root password can be changed to a more secure value by going to `Account → Users → View Users`. Highlight the entry for `root`, click the `Modify User` button, enter the new password in the `Password` and `Password confirmation` fields, and click `OK` to save the new password to use on subsequent logins.

On the first login, the EULA found in [Appendix A: End-User License Agreement](#) (page 316) is displayed, along with a box where the license key for the TrueNAS® array can be pasted. Read the EULA, paste in the license key, then click OK. The administrative GUI appears, as shown in the example in [Figure 2.14](#).

Note: Entering the license key for a High Availability pair is not allowed unless both the active and standby computers are up. The key is entered on the active computer.

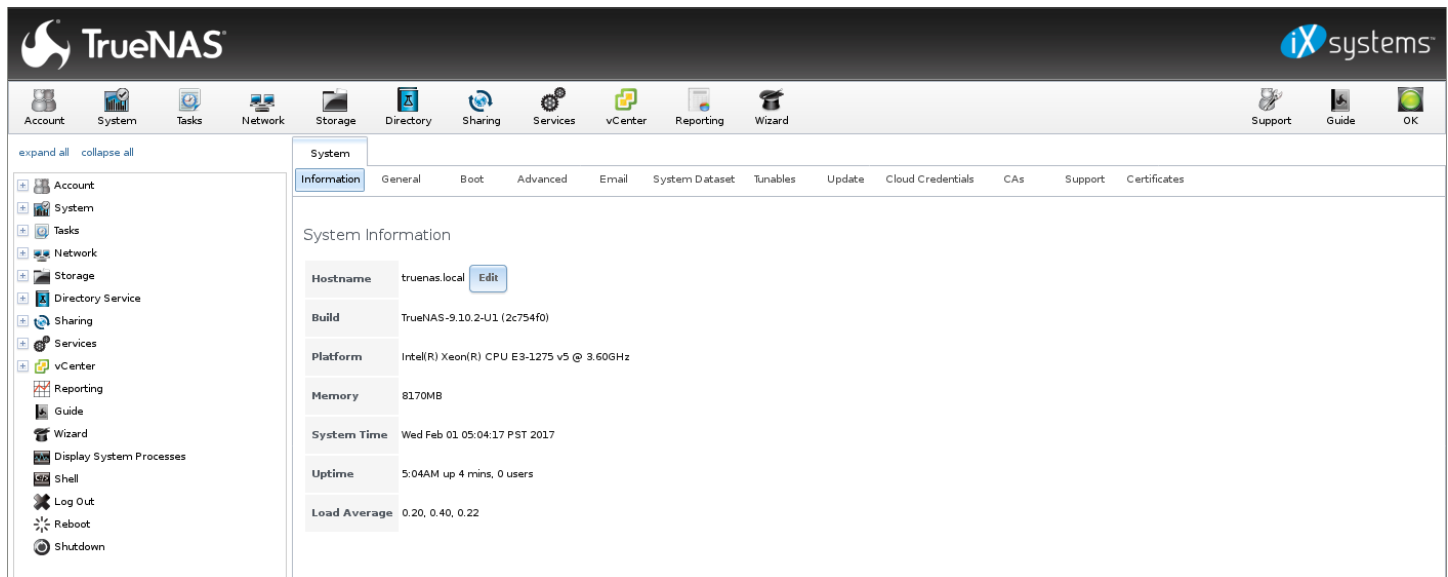


Fig. 2.14: TrueNAS® Graphical Configuration Menu

Note: If the storage devices have been encrypted, a prompt appears for the passphrase. It must be correctly entered for the data on the disks to be accessible. If the system has also been licensed for High Availability (HA), the passphrase will be remembered as long as either node in the HA unit remains up. If both nodes are powered off, the passphrase must be re-entered when the first node powers back up.

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

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

The rest of this Guide describes all of the configuration screens available within the TrueNAS® graphical administrative interface. The screens are listed in the order that they appear within the tree, or the left frame of the graphical interface.

Note: iXsystems recommends that you contact your iXsystems Support Representative for initial setup and configuration assistance.

Once the system has been configured and you are familiar with the configuration workflow, the rest of this document can be used as a reference guide to the features built into the TrueNAS® Storage Array.

Note: It is important to use the graphical interface (or the console setup menu) for all non-ZFS configuration changes. TrueNAS® uses a configuration database to store its settings. If changes are made at the command line, they will not be written to the configuration database. This means that these changes will not persist after a reboot and will be overwritten by the values in the configuration database during an upgrade.

ACCOUNT

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

- [Groups](#) (page 19): used to manage UNIX-style groups on the TrueNAS® system.
- [Users](#) (page 22): used to manage UNIX-style accounts on the TrueNAS® system.

Each entry is described in more detail in this section.

3.1 Groups

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

Note: If a directory service is running on the network, it is not necessary to recreate the network's users or groups. Instead, import the existing account information into TrueNAS®. Refer to [Directory Services](#) (page 132) for details.

This section describes how to create a group and assign user accounts to it. The next section, [Users](#) (page 22), describes creating user accounts.

Click `Groups` → `View Groups` to see a screen like [Figure 3.1](#).

Account

Groups

Users

Add Group

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

Members

Fig. 3.1: Group Management

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

The Add Group button opens the screen shown in Figure 3.2. Table 3.1 summarizes the available options when creating a group.

Add Group

Group ID:

1001

Group Name:

Permit Sudo:

☐

Allow repeated GIDs:

☐

OK

Cancel

Fig. 3.2: Creating a New Group

Table 3.1: Group Creation Options

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

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

In the example shown in Figure 3.3, the data1 group has been created and the user1 user account has been created with a primary group of user1. The Members button for the data1 group has been selected and user1 has been added as a member of the group.

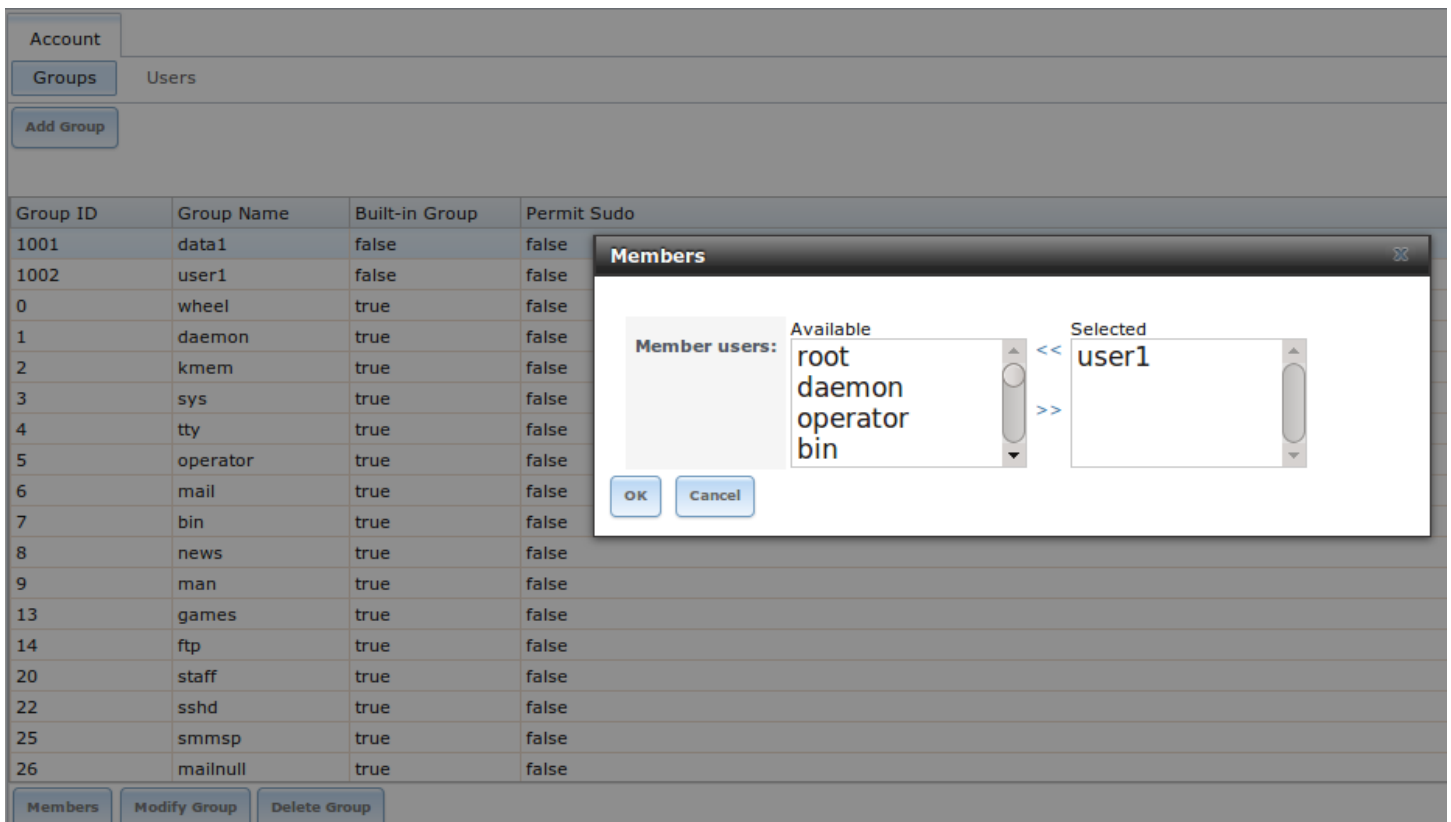


Fig. 3.3: Assigning a User to a Group

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

3.2 Users

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

1. Create a guest account that all users will use or create a user account for every user in the network where the name of each account is the same as a logon name used on a computer. For example, if a Windows system has a login name of bobsmith, create a user account with the name bobsmith on TrueNAS®. A common strategy is to create groups with different sets of permissions on shares, then assign users to those groups.
2. If your network uses a directory service, import the existing account information using the instructions in [Directory Services](#) (page 132).

Account → Users → View Users provides a listing of all of the system accounts that were installed with the TrueNAS® operating system, as shown in [Figure 3.4](#).

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

Fig. 3.4: Managing User Accounts

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

Click a user account to cause these buttons to appear:

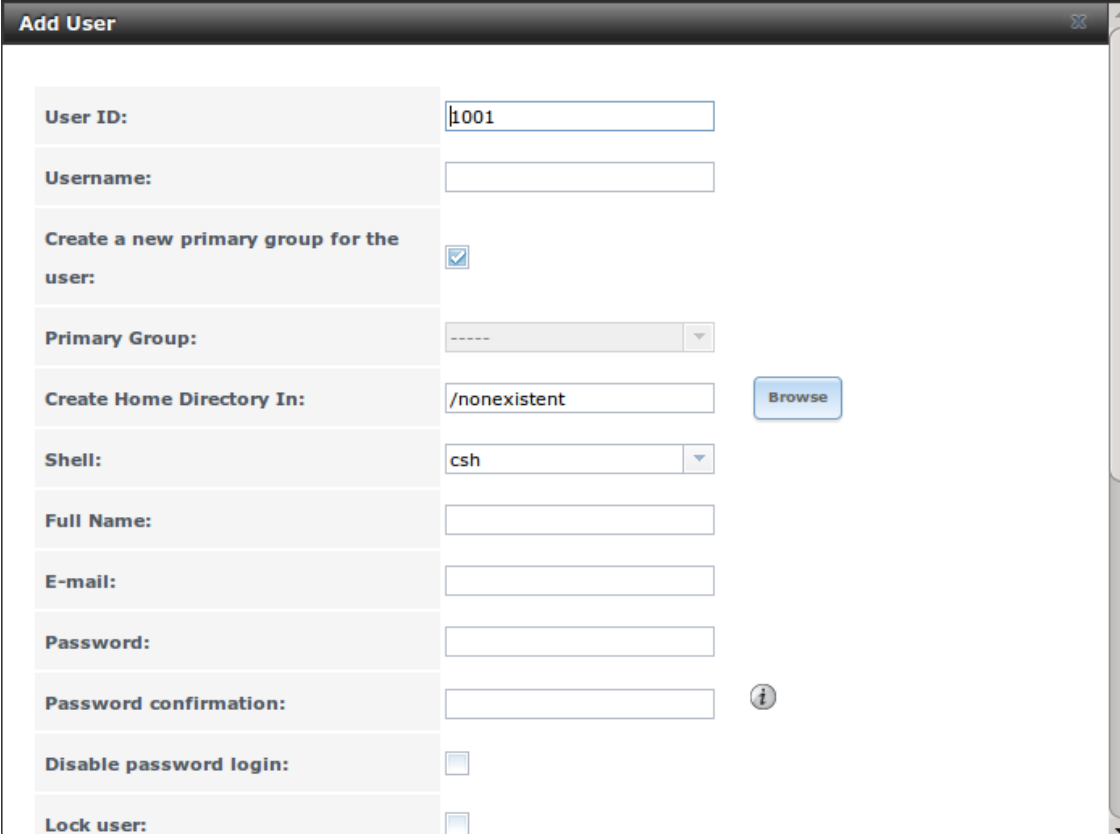
- **Modify User:** used to modify the account's settings, as listed in [Table 3.2](#).
- **Change E-mail:** used to change the email address associated with the account.

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

Except for the root user, the accounts that come with 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)` (<http://www.freebsd.org/cgi/man.cgi?query=nologin>). For security reasons, and to prevent breakage of system services, do not modify the system accounts.

The Add User button opens the screen shown in [Figure 3.5](#). Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in `System → Advanced`. [Table 3.2](#) summarizes the options which are available when user accounts are created or modified.

Warning: When using [Active Directory](#) (page 132), Windows user passwords must be set from within Windows.



The screenshot shows the 'Add User' form with the following fields and values:

Field	Value
User ID:	h001
Username:	
Create a new primary group for the user:	<input checked="" type="checkbox"/>
Primary Group:	-----
Create Home Directory In:	/nonexistent
Shell:	csh
Full Name:	
E-mail:	
Password:	
Password confirmation:	
Disable password login:	<input type="checkbox"/>
Lock user:	<input type="checkbox"/>

Fig. 3.5: Adding or Editing a User Account

Table 3.2: User Account Configuration

Setting	Value	Advanced Mode	Description
User ID	integer		grayed out if user already created; when creating an account, the next numeric ID will be suggested; by convention, user accounts have an ID greater than 1000 and system accounts have an ID equal to the default port number used by the service
Username	string		grayed out if user already created; maximum 16 characters though a maximum of 8 is recommended for interoperability; cannot begin with a hyphen, if a \$ is used it can only be the last character, and it cannot contain a space, tab, or the characters , : + & # % ^ & () ! @ ~ * ? < > =
Create a new primary group	checkbox		by default, a primary group with the same name as the user will be created; uncheck this box to select a different primary group name
Primary Group	drop-down menu		must uncheck Create a new primary group to access this menu; for security reasons, FreeBSD will not give a user <code>su</code> permissions if wheel is their primary group; to give a user <code>su</code> access, add them to the wheel group in Auxiliary groups
Create Home Directory In	browse button		browse to the name of an existing volume or dataset that the user will be assigned permission to access
Home Directory Mode	checkboxes	✓	sets default Unix permissions of user's home directory; read-only for built-in users
Shell	drop-down menu		select shell to use for local and SSH logins; see Table 3.3 for an overview of available shells
Full Name	string		mandatory, may contain spaces
E-mail	string		email address associated with the account
Password	string		mandatory unless check box Disable password login; cannot contain a ?
Password confirmation	string		must match the value of Password
Disable password login	checkbox		when checked, disables password logins and authentication to SMB shares; to undo this setting, set a password for the user using the Modify User button for the user in View Users; checking this box grays out Lock user and Permit Sudo, which are mutually exclusive
Lock user	checkbox		a checked box prevents user from logging in until the account is unlocked (box is unchecked); checking this box will gray out Disable password login which is mutually exclusive
Permit Sudo	checkbox		if checked, members of the group have permission to use <code>sudo</code> (http://www.sudo.ws/); when using sudo, a user will be prompted for their own password
Microsoft Account	checkbox		check this box if the user will be connecting from a Windows 8 or higher system
SSH Public Key	string		paste the user's public SSH key to be used for key-based authentication (do not paste the private key!)
Auxiliary groups	mouse selection		highlight the groups to which the user is to be added; click the >> button to add the user to the highlighted groups

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

Table 3.3: Available Shells

Shell	Description
netcli.sh	user is shown the Console Setup menu (Figure 2.12) on connection, even if it is disabled in <code>System → Advanced → Enable Console Menu</code> ; the user must be root or have root permissions (effective user ID 0, like toor)
csch	C shell (https://en.wikipedia.org/wiki/C_shell)
sh	Bourne shell (https://en.wikipedia.org/wiki/Bourne_shell)
tcsh	Enhanced C shell (https://en.wikipedia.org/wiki/Tcsh)
nologin	use when creating a system account or to create a user account that can authenticate with shares but which cannot login to the FreeNAS system using <code>ssh</code>
bash	Bourne Again shell (https://en.wikipedia.org/wiki/Bash_%28Unix_shell%29)
ksh93	Korn shell (http://www.kornshell.com/)
mksh	mirBSD Korn shell (https://www.mirbsd.org/mksh.htm)
rbash	Restricted bash (http://www.gnu.org/software/bash/manual/html_node/The-Restricted-Shell.html)
rzsh	Restricted zsh (http://www.csse.uwa.edu.au/programming/linux/zsh-doc/zsh_14.html)
scponly	select scponly (https://github.com/scponly/scponly/wiki) to restrict the user's SSH usage to only the <code>scp</code> and <code>sftp</code> commands
zsh	Z shell (http://www.zsh.org/)
git-shell	restricted git shell (http://git-scm.com/docs/git-shell)

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

SYSTEM

The System section of the administrative GUI contains these entries:

- [Information](#) (page 26) provides general TrueNAS® system information such as hostname, operating system version, platform, and uptime
- [General](#) (page 27) configures general settings such as HTTPS access, the language, and the timezone
- [Boot](#) (page 30) creates, renames, and deletes boot environments
- [Advanced](#) (page 33) configures advanced settings such as the serial console, swap space, and console messages
- [Email](#) (page 34) configures the email address to receive notifications
- [System Dataset](#) (page 36) configures the location where logs and reporting graphs are stored
- [Tunables](#) (page 36) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- [Update](#) (page 39) performs upgrades and checks for system updates
- [Cloud Credentials](#) (page 43) is used to enter connection credentials for remote cloud service providers
- [Alert Services](#) (page 44) configures services used to notify the administrator about system events.
- [CAs](#) (page 45): import or create internal or intermediate CAs (Certificate Authorities)
- [Certificates](#) (page 48): import existing certificates or create self-signed certificates
- [Support](#) (page 51): view licensing information or create a support ticket.
- [Proactive Support](#) (page 53): enable and configure automatic proactive support (Silver or Gold support coverage only).
- [View Enclosure](#) (page 54): view status of disk enclosures.
- [Failover](#) (page 55): manage High Availability.

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

4.1 Information

`System` → `Information` displays general information about the TrueNAS® system. An example is seen in [Figure 4.1](#).

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

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

System

InformationGeneralBootAdvancedEmailSystem DatasetTunablesUpdateCloud CredentialsAlert ServicesCAsCertificatesSupportProactive SupportView EnclosureFailover

System Information

Hostname	<div></div> <div>Edit</div>
Build	TrueNAS-11.1
Platform	Intel(R) Xeon(R) CPU D-1531 @ 2.20GHz
Memory	32641MB
System Time	Wed, 10 Jan 2018 10:29:07 -0500
Uptime	10:29AM up 1 day, 17:40, 7 users
Load Average	0.45, 0.37, 0.26
System Serial	<div></div>
System Product	TRUENAS-X10-MODEL-HA
License	Gold contract, expires at <div></div>

System Information (Standby Node)

Hostname	<div></div>
Build	TrueNAS-11.1
Platform	Intel(R) Xeon(R) CPU D-1531 @ 2.20GHz
Memory	32641MB
System Time	Wed, 10 Jan 2018 10:29:08 -0500
Uptime	10:29AM up 1 day, 17:19, 1 user
Load Average	0.30, 0.33, 0.21
System Serial	<div></div>
System Product	TRUENAS-X10-HA

Fig. 4.1: System Information Tab

4.2 General

System → General is shown in [Figure 4.2](#).

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
Alert Services
CAs
Certificates
Support
Proactive Support
View Enclosure
Failover

Protocol: HTTP
Certificate:
WebGUI IPv4 Address: 0.0.0.0
WebGUI IPv6 Address: ::
WebGUI HTTP Port: 80
WebGUI HTTPS Port: 443
WebGUI HTTP -> HTTPS Redirect:
Language (Require UI reload): English
Console Keyboard Map: United States of America
Timezone: America/New_York
Syslog level: Info
Syslog server:
Save
Reset Configuration to Defaults
Save Config
Upload Config
HTTP Servers

Fig. 4.2: General Screen

Table 4.1 summarizes the settings that can be configured using the General tab:

Table 4.1: General Configuration Settings

Setting	Value	Description
Protocol	drop-down menu	protocol to use when connecting to the administrative GUI from a browser; if modified from the default of HTTP to HTTPS or to HTTP+HTTPS, select the certificate to use in Certificate; if you do not have a certificate, first create a CA (in CAs (page 45)), then the certificate itself (in Certificates (page 48))
Certificate	drop-down menu	required for HTTPS; browse to the location of the certificate to use for encrypted connections
WebGUI IPv4 Address	drop-down menu	choose from a list of recent IP addresses to limit the one to use when accessing the administrative GUI; the built-in HTTP server will automatically bind to the wildcard address of 0.0.0.0 (any address) and will issue an alert if the specified address becomes unavailable
WebGUI IPv6 Address	drop-down menu	choose from a list of recent IPv6 addresses to limit the one to use when accessing the administrative GUI; the built-in HTTP server will automatically bind to any address and will issue an alert if the specified address becomes unavailable
WebGUI HTTP Port	integer	allows configuring a non-standard port for accessing the administrative GUI over HTTP; changing this setting might also require changing a Firefox configuration setting (http://www.redbrick.dcu.ie/%7Ed_fens/articles/Firefox:_This_Address_is_Restricted)
WebGUI HTTPS Port	integer	allows configuring a non-standard port for accessing the administrative GUI over HTTPS
WebGUI HTTP -> HTTPS Redirect	checkbox	when this box is checked, HTTP connections are automatically redirected to HTTPS if HTTPS is selected in Protocol, otherwise such connections will fail
Language	drop-down menu	select the localization from the drop-down menu and reload the browser; view the status of localization at pootle.freenas.org (http://pootle.freenas.org/)

Continued on next page

Table 4.1 – continued from previous page

Setting	Value	Description
Console Keyboard Map	drop-down menu	select the keyboard layout
Timezone	drop-down menu	select the timezone from the drop-down menu
Syslog level	drop-down menu	when Syslog server is defined, only logs matching this level are sent
Syslog server	string	IP address_or_hostname:optional_port_number of remote syslog server to send logs to; once set, log entries are written to both the console and the remote server

After making any changes, click the Save button.

This screen also contains these buttons:

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

Save Config: save a backup copy of the current configuration database in the format hostname-version-architecture to the computer accessing the administrative interface. Saving the configuration after making any configuration changes is highly recommended. 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 213) keys are not stored in the configuration database and must be backed up separately.

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

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

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

NTP Servers: The network time protocol (NTP) is used to synchronize the time on the computers in a network. Accurate time is necessary for the successful operation of time sensitive applications such as Active Directory or other directory services. By default, TrueNAS® is pre-configured to use three public NTP servers. If your 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 TrueNAS® system's physical location.

NTP servers are added by clicking on `NTP Servers → Add NTP Server` to open the screen shown in [Figure 4.3](#). [Table 4.2](#) summarizes the options available when adding an NTP server. [ntp.conf\(5\)](#) (<http://www.freebsd.org/cgi/man.cgi?query=ntp.conf>) explains these options in more detail.




Fig. 4.3: Add an NTP Server

Table 4.2: NTP Servers Configuration Options

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

4.3 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. If the update fails, reboot the system and select the previous boot environment from the boot menu to instruct the system to go back to that system state.

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

As seen in [Figure 4.4](#), two boot environments are created when TrueNAS® is installed. The system will boot into the default boot environment and users can make their changes and update from this version. The other boot environment, named Initial-Install can be booted into if the system needs to be returned to a pristine, non-configured version of the installation.

If the [Wizard](#) (page 250) was used, a third boot environment called `Wizard-date` is also created, indicating the date and time the [Wizard](#) (page 250) was run.

System			
Information	General	Boot	Advanced Email System Dataset Tunables Update Cloud Credentials Alert Services CAs Certificates Support Proactive Support View Enclosure Failover
Create	Scrub Boot	Status	Boot Volume Condition: HEALTHY Size: 103.0 GiB Last Scrub Run on: Sat Jan 6 03:46:21 2018 Used: 11.6 GiB (11%)
7 Automatic scrub interval (in days)			
Name	Active	Created	Keep
default		2017-08-18 13:33:00	No
Initial-Install		2017-08-18 13:42:00	No
11.1	On Reboot, Now	2018-01-08 13:16:00	No

Fig. 4.4: Viewing Boot Environments

Each boot environment entry contains this information:

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

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

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

The buttons above the boot entries can be used to:

- Create: a manual boot environment. A pop-up menu prompts for entry of a Name for the boot environment. When entering the name, only alphanumeric characters, underscores, and dashes are allowed.
- Scrub Boot: can be used to perform a manual scrub of the boot devices. By default, the boot device is scrubbed every 7 days. To change the default interval, change the number in the Automatic scrub interval (in days) field. The date and results of the last scrub are also listed in this screen. The condition of the boot device should be listed as HEALTHY.
- Status: click this button to see the status of the boot devices. [Figure 4.5](#), shows only one boot device, which is ONLINE.

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

Replace

Fig. 4.5: Viewing the Status of the Boot Device

If one of the boot devices has a Status of OFFLINE, click the device to replace, select the new replacement device, and click Replace Disk to rebuild the boot mirror.

Figure 4.6 shows a sample boot menu.




Fig. 4.6: Boot Environments in Boot Menu

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

4.4 Advanced

`System → Advanced` is shown in [Figure 4.7](#). The configurable settings are summarized in [Table 4.3](#).

System

InformationGeneralBootAdvancedEmailSystem DatasetTunablesUpdateCloud CredentialsAlert ServicesCAsCertificatesSupportProactive SupportView EnclosureFailover

Show Text Console without Password Prompt:☒

Use Serial Console:☒

Serial Port Address:

0x3f8

Serial Port Speed:

115200

Enable screen saver:☐

Enable powerd (Power Saving Daemon):☐

Show console messages in the footer:☒

Show tracebacks in case of fatal errors:☒

Show advanced fields by default:☐

Enable autotune:☒

Enable debug kernel:☐

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

MOTD banner:

Welcome to TrueNAS

Periodic Notification User:

root

Report CPU usage in percentage:☐

Remote Graphite Server Hostname:

Use FQDN for logging:☐

Save

Save Debug

Fig. 4.7: Advanced Screen

Table 4.3: Advanced Configuration Settings

Setting	Value	Description
Show Text Console Without Password Prompt	checkbox	unchecking this box replaces the console menu shown in Figure 2.12 with a login prompt
Use Serial Console	checkbox	do not check this box if the serial port is disabled
Serial Port Address	string	serial port address in hex
Serial Port Speed	drop-down menu	select the speed used by the serial port
Enable screen saver	checkbox	enable or disable the console screen saver
Enable powerd (Power Saving Daemon)	checkbox	powerd(8) (http://www.freebsd.org/cgi/man.cgi?query=powerd) monitors the system state and sets the CPU frequency accordingly
Continued on next page		

Table 4.3 – continued from previous page

Setting	Value	Description
Show console messages in the footer	checkbox	display console messages in real time at bottom of browser; click the console to bring up a scrollable screen; check the Stop refresh box in the scrollable screen to pause updating and uncheck the box to continue to watch the messages as they occur
Show tracebacks in case of fatal errors	checkbox	provides a pop-up of diagnostic information when a fatal error occurs
Show advanced fields by default	checkbox	several GUI menus provide an Advanced Mode button to access additional features; enabling this shows these features by default
Enable autotune	checkbox	enables Autotune (page 34) which attempts to optimize the system depending upon the hardware which is installed
Enable debug kernel	checkbox	when checked, next boot uses a debug version of the kernel
Enable automatic upload of kernel crash dumps and daily telemetry	checkbox	when checked, kernel crash dumps and telemetry (some system stats, collectedd RRDs, and select syslog messages) are automatically sent to the development team for diagnosis
MOTD banner	string	message to be shown when a user logs in with SSH
Periodic Notification User	drop-down menu	user to receive security output emails; this output runs nightly but only sends an email when the system reboots or encounters an error
Report CPU usage in percentage	checkbox	when checked, CPU usage is reported as percentages in Reporting (page 248)
Remote Graphite Server hostname	string	IP address or hostname of a remote server running Graphite (http://graphite.wikidot.com/)
Use FQDN for logging	checkbox	when checked, include the Fully-Qualified Domain Name in logs to precisely identify systems with similar hostnames

Click the Save button after making any changes.

This tab also contains this button:

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

4.4.1 Autotune

TrueNAS® provides an autotune script which optimizes the system. The Enable autotune checkbox in [System](#) → [Advanced](#) is checked by default, so this script runs automatically. It is recommended to leave autotune enabled 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 your support engineer beforehand as 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`.

4.5 Email

An automatic script sends a nightly email to the root user account containing important information such as the health of the disks. [Alert](#) (page 260) events are also emailed to the root user account. Problems with [Scrubs](#) (page 126) are reported separately in an email sent at 03:00AM.

Note: [S.M.A.R.T.](#) (page 206) 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. Select **Account** → **Users**, click on **root** to highlight that user, then click **Modify User**. In the **E-mail** field, enter the email address on the remote system where email is to be sent, like **admin@example.com**. Click **OK** to save the settings.

Additional configuration is performed with **System** → **Email**, shown in [Figure 4.8](#).

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

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

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

Fig. 4.8: Email Screen

Table 4.4: Email Configuration Settings

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

Click the **Send Test Mail** button to verify that the configured email settings are working. If the test email fails, double-check the destination email address by clicking the **Change E-mail** button for the **root** account in **Account** → **Users** → **View Users**. Test mail cannot be sent unless the root email address has been set.

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

4.6 System Dataset

System → System Dataset, shown in [Figure 4.9](#), is used to select the pool which will contain the persistent system dataset. The system dataset stores debugging core files and Samba4 metadata such as the user/group cache and share level permissions. If the TrueNAS® system is configured to be a Domain Controller, all of the domain controller state is stored there as well, including domain controller users and groups.

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

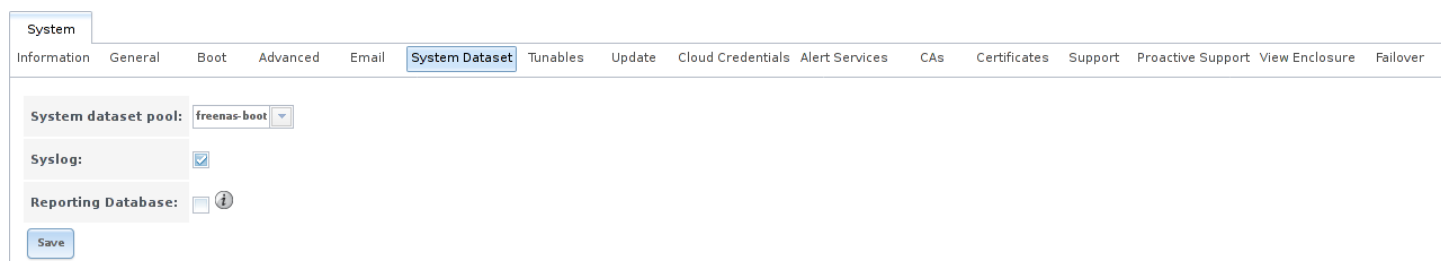


Fig. 4.9: System Dataset Screen

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

The system dataset can optionally be configured to also store the system log and [Reporting](#) (page 248) information. If there are lots of log entries or reporting information, moving these to the system dataset will prevent `/var/` on the device holding the operating system from filling up as `/var/` has limited space.

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

Note: It is recommended to store the system dataset on the `freenas-boot` pool. For this reason, a yellow system alert will be generated when the system dataset is configured to use another pool.

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

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

Click the Save button to save changes.

If the pool storing the system dataset is changed at a later time, 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.

4.7 Tunables

System → Tunables can be used to manage the following:

1. FreeBSD sysctls: a [sysctl\(8\)](http://www.freebsd.org/cgi/man.cgi?query=sysctl) (<http://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&manpath=FreeBSD+11.0-RELEASE) (<https://www.freebsd.org/cgi/man.cgi?query=rc.conf&manpath=FreeBSD+11.0-RELEASE>) is used to pass system configuration options to the system startup scripts as the system boots. Since 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 unless you understand and have tested the ramifications of that change.

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

To add a loader, sysctl, or rc.conf option, go to `System → Tunables → Add Tunable`, to access the screen shown in seen in [Figure 4.10](#).

Fig. 4.10: Adding a Tunable

[Table 4.5](#) summarizes the options when adding a tunable.

Table 4.5: Adding a Tunable

Setting	Value	Description
Variable	string	typically the name of the sysctl or driver to load, as indicated by its man page
Value	integer or string	value to associate with Variable; typically this is set to YES to enable the sysctl or driver specified by the “Variable”
Type	drop-down menu	choices are Loader, rc.conf, or Sysctl
Comment	string	optional, but a useful reminder for the reason behind adding this tunable
Continued on next page		

Table 4.5 – continued from previous page

Setting	Value	Description
Enabled	checkbox	uncheck if you would like to disable the tunable without deleting it

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

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

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

```
kern.metadelay=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 GUI does not display the loaders that are pre-set when TrueNAS® is installed. TrueNAS® 11.1 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.hptrr.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"
hint.ahciem.0.disabled="1"
hint.ahciem.1.disabled="1"
kern.msgbufsize="524288"
hw.cxgbe.toecaps_allowed=0
hw.cxgbe.rdmacaps_allowed=0
```

```
hw.cxgbe.iscsicaps_allowed=0
vfs.nfsd.fha.write=0
vfs.nfsd.fha.max_nfsds_per_fh=32
kern.ipc.nmbclusters="262144"
kern.hwpmc.nbuffers="4096"
kern.hwpmc.nsamples="4096"
hw.memtest.tests="0"
vfs.zfs.trim.enabled="0"
kern.cam.ctl.ha_mode=2
kern.geom.label.ufs.enable=0
kern.geom.label.ufsid.enable=0
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 as doing so might make the system unusable.

The ZFS version used in 11.1 deprecates these tunables:

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

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

4.8 Update

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

4.8.1 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 difficulties are encountered. 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, use [Boot](#) (page 30) to remove unneeded boot environments.

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

Available ZFS version upgrades are indicated by an [Alert](#) (page 260) in the graphical user 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®.

4.8.2 Updates and Trains

TrueNAS® is updated with signed update files. This provides flexibility in deciding when to upgrade the system with patches, new drivers, or new features. It also allows “test driving” an upcoming release. Combined with boot environments, new features or system patches can be tested while still being able to revert to a previous version of the operating system (see [If Something Goes Wrong](#) (page 42)). Digital signing of update files eliminates the need to manually download both an upgrade file and the associated checksum to verify file integrity.

Figure 4.11 shows an example of the `System` → `Update` screen.

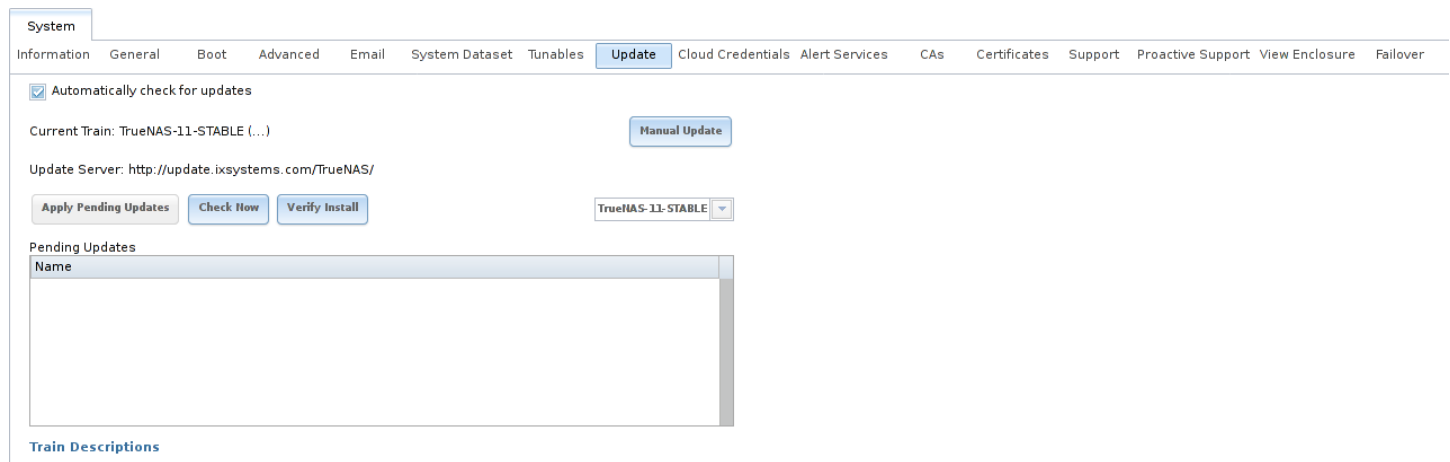


Fig. 4.11: Update Options

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

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

Several trains are available for updates.

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

These trains are available:

For Production Use

- TrueNAS-11-STABLE (Recommended)

After new fixes and features have been tested as production-ready, they are added to this train. Following this train and applying any pending updates from it is recommended.

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 iX support engineer.

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

4.8.3 Checking for Updates

To see if any updates are available, click the Check Now button. Any available updates are listed.

4.8.4 Applying Updates

Make sure the system is in a low-usage state as described above in [Preparing for Updates](#) (page 39).

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

Warning: Each update creates a boot environment. If the update process needs more space, it attempts to remove old boot environments. Boot environments marked with the Keep attribute as shown in [Boot](#) (page 30) will not be removed. If space for a new boot environment is not available, the upgrade fails. Space on the boot device can be manually freed using `System → Boot`. Review the boot environments and remove the Keep attribute or delete any boot environments that are no longer needed.

Updates can also be downloaded and applied later. To do so, uncheck the Apply updates after downloading box before pressing OK. In this case, this screen closes after updates are downloaded. Downloaded updates are listed in the Pending Updates section of the screen shown in [Figure 4.11](#). When ready to apply the previously downloaded updates, click the Apply Pending Updates button. Remember that the system might reboot after the updates are applied.

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

4.8.5 Manual Updates

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

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

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

4.8.6 Updating from the Shell

Updates can also be performed from the [Shell](#) (page 257) 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`.

4.8.7 Updating an HA System

If the TrueNAS® array has been configured for High Availability (HA), the update process must be started on the active node. Once the update is complete, the standby node will automatically reboot. Wait for it to come back up by monitoring the remote console or the graphical administrative interface of the standby node.

After the standby node has finished booting, it is important to perform a failover by rebooting the current active node. This action tells the standby node to import the current configuration and restart services.

Once the previously active node comes back up as a standby node, use `System` → `Update` to apply the update on the current active node (which was previously the passive node). Once complete, the now standby node will reboot a second time.

4.8.8 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.

4.8.9 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 feature flags.
- before performing any operation that may affect the data on a storage disk, always back up all data first and verify the integrity of the backup. While it is unlikely that the pool upgrade will affect the data, it is always better to be safe than sorry.
- upgrading a ZFS pool is optional. Do not upgrade the pool if the the possibility of reverting to an earlier version of 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. If a pool is upgraded to the latest feature flags, it will not be possible to import that pool into another operating system that does not yet support those feature flags.

To perform the ZFS pool upgrade, go to `Storage` → `Volumes` → `View Volumes` and highlight the volume (ZFS pool) to upgrade. Click the Upgrade button as shown in [Figure 4.12](#).

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




Fig. 4.12: Upgrading a ZFS Pool

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

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

4.9 Cloud Credentials

TrueNAS® can use cloud services for features like [Cloud Sync](#) (page 58). The credentials to provide secure connections with cloud services are entered here. Amazon S3, Azure Blob Storage, Backblaze B2, and Google Cloud Storage are supported.

Select **System** → **Cloud Credentials** → **Add Cloud Credential** to display the dialog shown in [Figure 4.13](#).

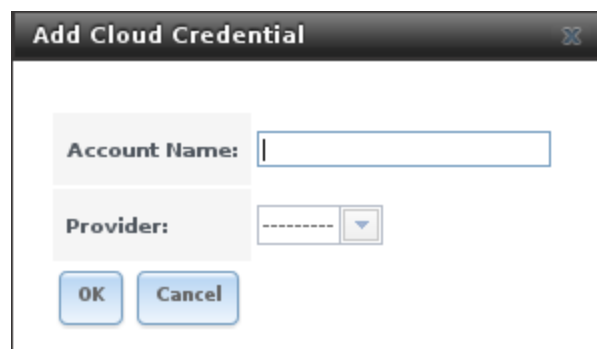


Fig. 4.13: Adding Cloud Credentials

Enter a descriptive name for the cloud credential in the Account Name field, then select a provider. The remaining options vary by provider, and are shown in [Table 4.6](#).

Table 4.6: Cloud Credential Options

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

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

4.10 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 260) marked WARN or CRITICAL.

Currently available alert services:

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

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

Select **System** → **Alert Services** to go to the Alert Services screen. Click **Add Service** to display the dialog shown in [Figure 4.14](#).

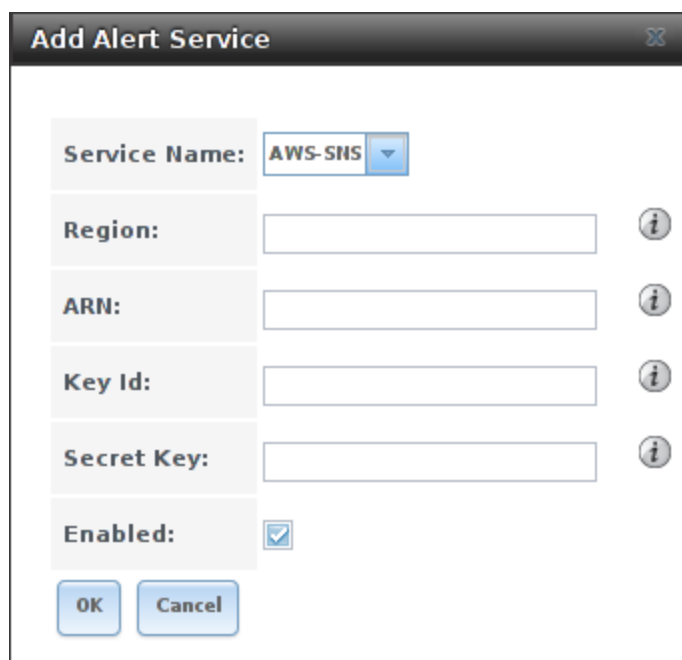


Fig. 4.14: Add Alert Service

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

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

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

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

4.10.1 How it Works

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

4.11 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 4.15 shows the screen after clicking `System → CAs`.

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
Alert Services
CAs
Certificates
Support
Proactive Support
View Enclosure
Failover

Import CA
Create Internal CA
Create Intermediate CA

Name	Internal	Issuer	Certificates	Distinguished Name	From	Until
No entry has been found						

Fig. 4.15: Initial CA Screen

If your organization already has a CA, the CA's certificate and key can be imported. Click the Import CA button to open the configuration screen shown in Figure 4.16. The configurable options are summarized in Table 4.7.

Import CA

Identifier:

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

Certificate:

i

Private Key:

i

Passphrase:

i

Confirm Passphrase:

Serial:

i

OK

Cancel

Fig. 4.16: Importing a CA

Table 4.7: Importing a CA Options

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

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

To create a CA for internal use only, click the Create Internal CA button which will open the screen shown in Figure 4.17.

Create Internal CA

Identifier:

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

Key length:

2048

Digest Algorithm:

SHA256

Lifetime:

3,650

Country:

United States

State:

Locality:

Organization:

Email Address:

Common Name:

Subject Alternate Names:

OK

Cancel

Fig. 4.17: Creating an Internal CA

The configurable options are described in [Table 4.8](#). When completing the fields for the certificate authority, supply the information for your organization.

Table 4.8: Internal CA Options

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

Continued on next page

Table 4.8 – continued from previous page

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

To instead create an intermediate CA which is part of a certificate chain, click the Create Intermediate CA button. This screen adds one more option to the screen shown in [Figure 4.17](#):

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

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

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

- **Sign CSR:** used to sign internal Certificate Signing Requests created using `System → Certificates → Create Certificate Signing Request`.
- **Export Certificate:** prompts to browse to the location to save a copy of the CA's X.509 certificate on the computer being used to access the 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.

4.12 Certificates

TrueNAS® can import existing certificates, create new certificates, and issue certificate signing requests so that created certificates can be signed by the CA which was previously imported or created in [CAs](#) (page 45).

[Figure 4.18](#) shows the initial screen after clicking `System → Certificates`.




Fig. 4.18: Initial Certificates Screen

To import an existing certificate, click the Import Certificate button to open the configuration screen shown in [Figure 4.19](#). 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 55) systems, the imported certificate must include the IP addresses or DNS hostnames of both nodes 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 4.9](#).

Import Certificate

Identifier:

Certificate:

Private Key:

Passphrase:

Confirm Passphrase:

OK Cancel

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

Fig. 4.19: Importing a Certificate

Table 4.9: Certificate Import Options

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

To instead create a new self-signed certificate, click the Create Internal Certificate button to see the screen shown in [Figure 4.20](#). The configurable options are summarized in [Table 4.10](#). When completing the fields for the certificate authority, use the information for your organization. Since this is a self-signed certificate, use the CA that was imported or created with CAs (page 45) as the signing authority.

Create Internal Certificate

Signing Certificate Authority:

Identifier:

Key length:

2048

Digest Algorithm:

SHA256

Lifetime:

3,650

Country:

United States

State:

Locality:

Organization:

Email Address:

Common Name:

Subject Alternate Names:

OK

Cancel

Fig. 4.20: Creating a New Certificate

Table 4.10: Certificate Creation Options

Setting	Value	Description
Signing Certificate Authority	drop-down menu	required; select the CA which was previously imported or created using CAs (page 45)
Identifier	string	required; enter a descriptive name for the certificate using only alphanumeric, underscore (<code>_</code>), and dash (<code>-</code>) characters
Key Length	drop-down menu	for security reasons, a minimum of 2048 is recommended
Digest Algorithm	drop-down menu	the default is acceptable unless your organization requires a different algorithm
Lifetime	integer	in days
Country	drop-down menu	select the country for the organization
State	string	required; enter the state or province for the organization
Locality	string	required; enter the location for the organization
Organization	string	required; enter the name of the company or organization

Continued on next page

Table 4.10 – continued from previous page

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

If you need to use a certificate that is signed by an external CA, such as Verisign, instead create a certificate signing request. To do so, click the Create Certificate Signing Request button. A screen like the one in [Figure 4.20](#) opens, but without the Signing Certificate Authority field.

Certificates that are imported, self-signed, or for which a certificate signing request is created are added as entries to `System` → `Certificates`. In the example shown in [Figure 4.21](#), a self-signed certificate and a certificate signing request have been created for the fictional organization My Company. The self-signed certificate was issued by the internal CA named My Company and the administrator has not yet sent the certificate signing request to Verisign so that it can be signed. Once that certificate is signed and returned by the external CA, it should be imported using the Import Certificate button so that is available as a configurable option for encrypting connections.

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

Fig. 4.21: Managing Certificates

Clicking an entry activates these configuration buttons:

- View: use this option to view or edit the contents of an existing certificate. These fields can be edited: Identifier (name), Certificate, and Private Key.
- Export Certificate saves a copy of the certificate or certificate signing request to the system being used to access the 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.

4.13 Support

The TrueNAS® Support tab, shown in [Figure 4.22](#), is used to view or update the system's license information. It also provides a built-in ticketing system for generating support requests.

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
Alert Services
CAs
Certificates
Support
Proactive Support
View Enclosure
Failover

License Information
Update License
Userguide (PDF)

ModelX10System SerialContract TypeContract Date
Customer NameFeaturesNoneAdditional HardwareNone

Name
E-mail
Phone
CategoryBug
EnvironmentProduction
CriticalityInquiry
Attach Debug Info
Subject
Description

Attachments
x Browse... No file selected.
+
Submit

Fig. 4.22: Support Tab

In this example, the system has a valid license which indicates 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 your iXsystems support engineer. Once you have the new license string, click the Update License button, paste in the new license, and click OK. The new details will be displayed.

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.
- E-mail is the email address of the person to contact.
- Phone is the phone number of the person to contact.
- Category is a drop-down menu to select whether the ticket is to report a software bug, report a hardware failure, ask for assistance in installing or configuring the system, or request assistance in diagnosing a performance bottleneck.
- Environment is a drop-down menu to indicate the role of the affected system. Choices are Production, Staging, Test, Prototyping, or Initial Deployment/Setup.
- Criticality is a drop-down menu to indicate the criticality level. Choices are Inquiry, Loss of Functionality, or Total Down.
- Attach Debug Info allows an overview of the system hardware and configuration to be automatically generated and included with the ticket. It is recommended to leave this box checked.
- 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.
- Attachments is an optional field where configuration files or screenshots of any errors or tracebacks can be included.

After completing the fields, click the Submit button to generate and send the support ticket to iXsystems. A pop-up menu provides a clickable URL to view the status of or add additional information to that support ticket. When not already logged into the [iXsystems Support page](https://support.ixsystems.com/) (<https://support.ixsystems.com/>), clicking this URL prompts for a login, or to register a new login.

4.14 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 for information on upgrading from other support levels.

Fig. 4.23: Proactive Support Tab

The Proactive Support fields are:

- Enable automatic support alerts to iXsystems allows enabling or disabling Proactive Support emails to iXsystems. It is recommended to enable this automatic reporting.
- Name of Primary Contact is the name of the first person to be contacted by iXsystems Support to assist with issues.
- Title is the title of the primary contact person.
- E-mail is the email address of the primary contact person.
- Phone is the phone number of the primary contact person.
- Name of Secondary Contact is the name of the person to be contacted when the primary contact person is not available.
- Secondary Title is the title of the secondary contact person.
- Secondary E-mail is the email address of the secondary contact person.
- Secondary Phone is the phone number of the secondary contact person.

To enable Proactive Support, complete the fields, make sure the Enable automatic support alerts to iXsystems box is checked, then click Save.

4.15 View Enclosure

Click Storage → Volumes → View Enclosure to display a status summary of the connected disks and hardware. An example is shown in Figure 4.24.

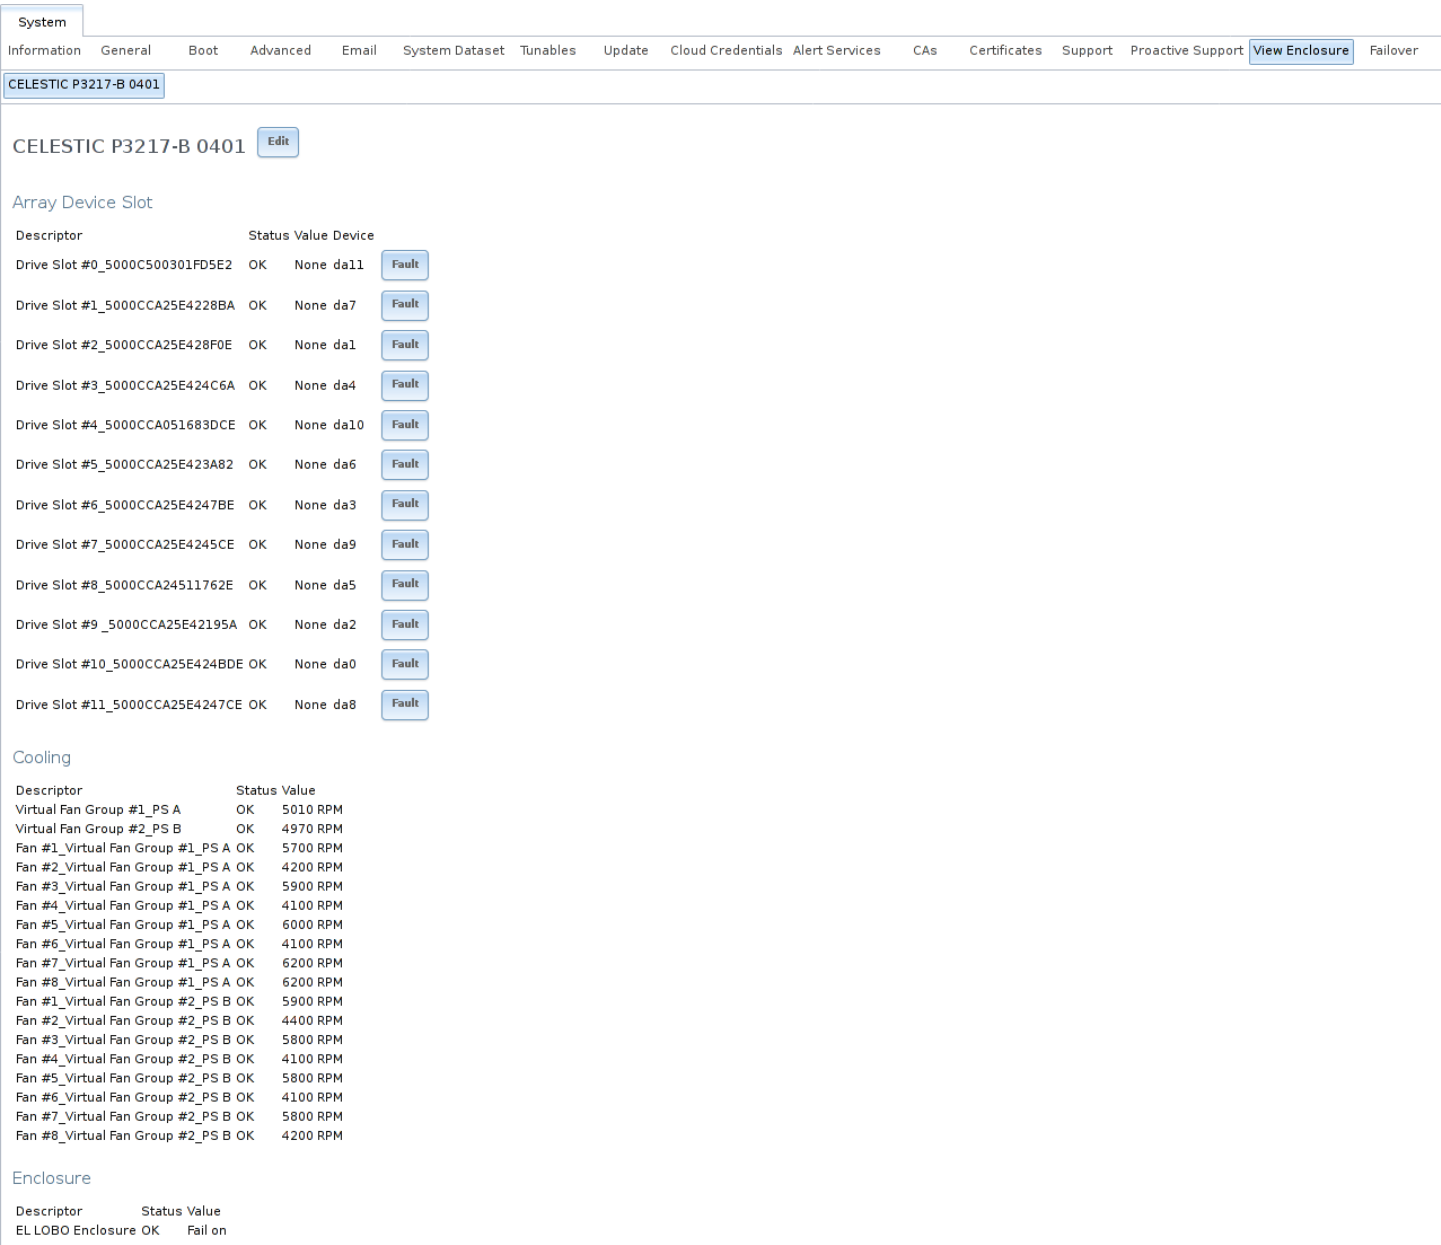


Fig. 4.24: View Enclosure

The screen is divided into these sections:

Array Device Slot: has an entry for each slot in the storage array, indicating the current disk status and FreeBSD device name. To blink the status light for that disk as a visual indicator, click the Identify button.

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

Enclosure: shows the status of the enclosure.

Power Supply: shows the status of each power supply.

SAS Expander: shows the status of the expander.

Temperature Sensor: shows the current temperature of each expander and the disk chassis.

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

4.16 Failover

If the TrueNAS® array has been licensed for High Availability (HA), a Failover tab is added to System.

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

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.

Warning: Seamless failover is only available with iSCSI or NFSv4. Other protocols will failover, but connections will be disrupted by the failover event.

To configure HA, turn on both units in the array. Use the instructions in the [Console Setup Menu](#) (page 14) to log into the graphical interface for one of the units (it does not matter which one). If this is the first login, the Upload License screen is automatically displayed. Otherwise, click `System` → `Support` → `Upload License`.

Paste the HA license received from iXsystems and press OK to activate it. The license contains the serial numbers for both units in the chassis. After the license is activated, the Failover tab is added to System and some fields are modified in Network so that the peer IP address, peer hostname, and virtual IP can be configured. An extra IPMI (Node A/B) tab will also be added so that [IPMI](#) (page 80) can be configured for the other unit.

Note: The modified fields refer to this node as This Node and the other node as either A or B. The node value is hard-coded into each unit and the value that appears is automatically generated. For example, on node A, the fields refer to node B, and vice versa.

To configure HA networking, go to `Network` → `Global Configuration`. The Hostname field is replaced by two fields:

- Hostname (Node A/B): enter the hostname to use for the other node.
- Hostname (This Node): enter the hostname to use for this node.

Next, go to `Network` → `Interfaces` → `Add Interface`. The HA license adds several fields to the usual [Interfaces](#) (page 78) screen:

- IPv4 Address (Node A/B): if the other node will use a static IP address, rather than DHCP, set it here.
- IPv4 Address (This Node): if this node will use a static IP address, rather than DHCP, set it here.
- Virtual IP: enter the IP address to use for administrative access to the array.
- Virtual Host ID: the Virtual Host ID (VHID) must be unique on the broadcast segment of the network. It can be any unused number between 1 and 255.
- Critical for Failover: check this box if a failover should occur when this interface becomes unavailable. How many seconds it takes for that failover to occur depends upon the value of the Timeout, as described in [Table 4.11](#). This checkbox is interface-specific, allowing you to have different settings for a management network and a data network.

Note that checking this box requires the Virtual IP to be set and that at least one interface needs to be set as Critical for Failover to configure failover.

- Group: this drop-down menu is grayed out unless the Critical for Failover checkbox is checked. This box allows grouping multiple, critical-for-failover interfaces. In this case, all of the interfaces in a group must go down before failover occurs. This can be a useful configuration in a multipath scenario.

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

The passive or standby node indicates the virtual IP address that is used for configuration management. The standby node also has a red Standby icon and no longer accepts logins as all configuration changes must occur on the active node.

Note: After the Virtual IP address is configured, all subsequent logins should use that address.

After HA has been configured, an HA Enabled icon is shown to the right of the Alert icon on the active node.

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

The icon is red when HA is starting up, disabled, or has encountered a problem. When HA is functioning normally, the icon turns green.

The options available in `System → Failover` are shown in [Figure 4.25](#): and described in [Table 4.11](#).

Fig. 4.25: Example Failover Screen

Table 4.11: Failover Options

Setting	Value	Description
Disabled	checkbox	when checked, administratively disable failover which changes the HA Enabled icon to HA Disabled and activates the Master field; an error message is generated if the standby node is not responding or failover has not been configured
Master	checkbox	grayed out unless Disabled is checked; in that case, this box is automatically checked on the master system, allowing the master to automatically take over when the Disabled box is unchecked
Timeout	integer	specify, in seconds, how quickly failover occurs after a network failure; the default of 0 indicates that failover either occurs immediately or, if the system is using a link aggregation, after 2 seconds
Sync to Peer	button	open a dialog window to force the TrueNAS® configuration to sync from the active node to the standby node; after the sync, the standby node must be rebooted (enabled by default) to load the new configuration; do not use this unless requested by an iX support engineer, the HA daemon normally handles configuration sync automatically

Continued on next page

Table 4.11 – continued from previous page

Setting	Value	Description
Sync From Peer	button	open a dialog window to force the TrueNAS® configuration to sync from the standby node to the active node; do not use this unless requested by an iX support engineer, the HA daemon normally handles configuration sync automatically

Warning: Booting an HA pair with failover disabled causes both nodes to come up in standby mode. The GUI shows an additional Force Takeover button which can be used to force that node to take control.

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

TASKS

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

- [Cloud Sync](#) (page 58) schedules data synchronization to cloud providers
- [Cron Jobs](#) (page 63) schedules a command or script to automatically execute at a specified time
- [Init/Shutdown Scripts](#) (page 65) configures a command or script to automatically execute during system startup or shutdown
- [Rsync Tasks](#) (page 66) schedules data synchronization to another system
- [S.M.A.R.T. Tests](#) (page 73) schedules disk tests

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

Note: By default, [Scrubs](#) (page 126) are run once a month by an automatically-created task. [S.M.A.R.T. Tests](#) (page 73) and [Periodic Snapshot Tasks](#) (page 113) must be set up manually.

5.1 Cloud Sync

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

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

Selecting [Tasks](#) → [Cloud Sync](#) shows the screen in [Figure 5.1](#). This screen shows a single cloud sync called backup-acctg that “pushes” a file to cloud storage. The last run finished with a status of SUCCESS.

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

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

Fig. 5.1: Cloud Sync Status

[Cloud Credentials](#) (page 43) must be defined before a cloud sync is created. One set of credentials can be used for more than one cloud sync. For example, a single set of credentials for Amazon S3 can be used for separate cloud syncs that push different sets of files or directories.

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

After the credentials and receiving bucket have been created, a cloud sync task is created with `Tasks → Cloud Sync → Add Cloud Sync`. The Add Cloud Sync dialog is shown in [Figure 5.2](#).

Add Cloud Sync

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

Description:

Direction:

Push

Provider:

Credential

Path:

Browse

Transfer Mode:

Sync

Minute:

Every N minute

Each selected minute

00

01

02

03

04

05

06

07

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

Fig. 5.2: Adding a Cloud Sync

Table 5.1 shows the configuration options for Cloud Syncs.

Table 5.1: Cloud Sync Options

Setting	Value Type	Description
Description	string	a descriptive name for this Cloud Sync
Direction	string	Push to send data to cloud storage, or Pull to pull data from the cloud storage
Provider	drop-down menu	select the cloud storage provider; the list of providers is defined by Cloud Credentials (page 43)
Amazon S3 Buckets	drop-down menu	only appears when an S3 credential is the Provider; select the bucket to use

Continued on next page

Table 5.1 – continued from previous page

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

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

Provider is the name of the cloud storage provider. These providers are defined by entering credentials in [Cloud Credentials](#) (page 43).

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

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

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

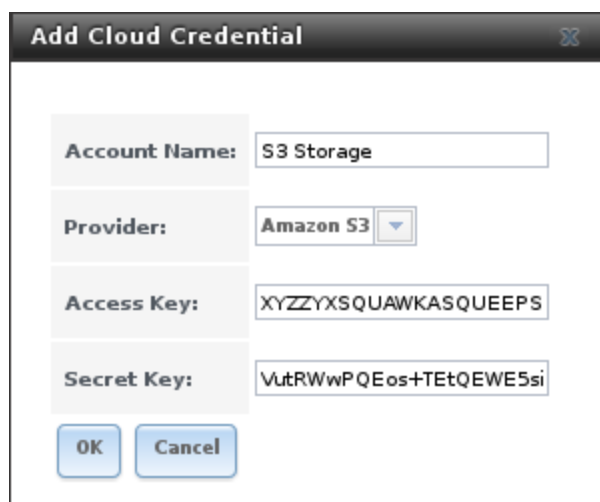
Finally, the Enabled field makes it possible temporarily disable a cloud sync job without deleting it.

5.1.1 Cloud Sync Example

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

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

`System` → `Cloud Credentials` → `Add Cloud Credential` is used to enter the credentials for storage on an Amazon AWS account. The credential is given the name `S3 Storage`, as shown in [Figure 5.3](#):



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

Fig. 5.3: Example: Adding Cloud Credentials

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

The Path to the data file is selected.

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

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

The completed dialog is shown in [Figure 5.4](#):

Add Cloud Sync

Description:

backup-acctg

Direction:

Push

Provider:

Credential

S3 Storage

Amazon S3 Buckets

cloudsync-bucket

Folder

Path:

/mnt/volume1/smb-storage/

Close

/

mnt

volume1

smb-storage

accounting-backup.bin

Minute:

Every N minute

Each selected minute

00

01

02

03

04

05

06

07

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

Hour:

Every N hour

Each selected hour

1

Day of month:

Every N day of month

Each selected day of month

1

Month:

January

February

March

April

May

June

July

August

September

October

November

December

Day of week:

Monday

Tuesday

Wednesday

Thursday

Friday

Saturday

Sunday

Enabled:

OK

Cancel

Fig. 5.4: Example: Adding a Cloud Sync

5.2 Cron Jobs

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

Figure 5.5 shows the screen that opens after clicking `Tasks` → `Cron Jobs` → `Add Cron Job`.

Add Cron Job

User:

The user to run the command

Command:

Short description:

Minute:

Every N minute

Each selected minute

00

01

02

03

04

05

06

07

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

Hour:

Every N hour

Each selected hour

1

Day of month:

Every N day of month

Each selected day of month

1

Month:

January

Fig. 5.5: Creating a Cron Job

Table 5.2 lists the configurable options for a cron job.

Table 5.2: Cron Job Options

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

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

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

5.3 Init/Shutdown Scripts

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

Figure 5.6 shows the screen that opens after clicking **Tasks** → **Init/Shutdown Scripts** → **Add Init/Shutdown Script**. Table 5.3 summarizes the options.

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

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




Fig. 5.6: Add an Init/Shutdown Script

Table 5.3: Options When Adding an Init/Shutdown Script

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

5.4 Rsync Tasks

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

Rsync is most effective when only a relatively small amount of the data has changed. There are also [some limitations when using Rsync with Windows files](https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/) (<https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/>). For large amounts of data, data that has many changes from the previous copy, or Windows files, [Replication Tasks](#) (page 115) are often the faster and better solution.

Rsync is single-threaded, so gains little from multiple processor cores. To see whether rsync is currently running, use `pgrep rsync` from the [Shell](#) (page 257).

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:

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

This section summarizes the options when creating an Rsync Task. It then provides a configuration example between two 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 5.7 shows the screen that appears after selecting `Tasks → Rsync Tasks → Add Rsync Task`. **Table 5.4** summarizes the options that can be configured when creating an rsync task.

Add Rsync Task

Path:
Browse

User:
i

Remote Host:
i

Rsync mode:
Rsync module

Remote Module Name:
i

Direction:
Push
i

Short description:

Minute:

Every N minute
Each selected minute

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

i

Hour:

Every N hour
Each selected hour

1

i

Day of month:

Every N day of month
Each selected day of month

1

Fig. 5.7: Adding an Rsync Task

Table 5.4: Rsync Configuration Options

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

Continued on next page

Table 5.4 – continued from previous page

Setting	Value	Description
Enabled	checkbox	uncheck to disable the rsync task without deleting it; note that when the Rsync (page 203) service is OFF, the rsync task will continue to look for the server unless this checkbox is unchecked

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

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

5.4.1 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 volume named `/mnt/remote`. It will be the rsync server, meaning that it will receive the contents of `/mnt/local/images`. An rsync module needs to be defined on this system and the rsyncd service needs to be started. It will be referred to as PULL.

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

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

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

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

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

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

5.4.2 Rsync over SSH Mode

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

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

- the SSH service must be running on PULL

To create the public/private key pair for the rsync user account, open [Shell](#) (page 257) 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.

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

```
su - user1
```

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

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

Go to PULL and paste (or append) the copied key into the SSH Public Key field of Account → Users → View Users → root → Modify User, or the username of the specified rsync user account. The paste for the above example is shown in [Figure 5.8](#). When pasting the key, ensure that it is pasted as one long line and, if necessary, remove any extra spaces representing line breaks.

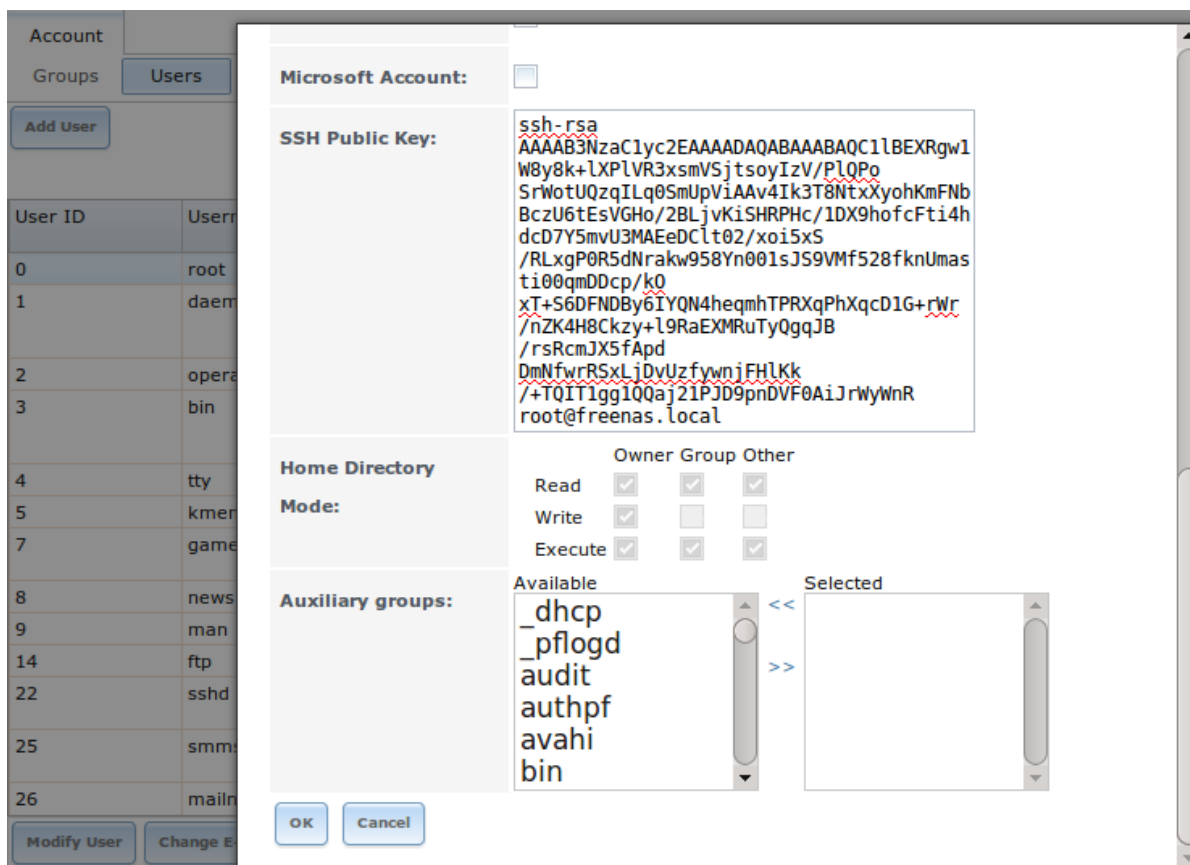


Fig. 5.8: Pasting the User's SSH Public Key

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

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

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

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

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

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

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

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

5.5 S.M.A.R.T. Tests

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

Figure 5.9 shows the configuration screen that appears after selecting `Tasks → S.M.A.R.T. Tests → Add S.M.A.R.T. Test`. Tests are listed under `View S.M.A.R.T. Tests`. After creating tests, check the configuration in `Services → S.M.A.R.T.`, then click the slider to ON for the S.M.A.R.T. service in `Services → Control Services`. The S.M.A.R.T. service will not start if there are no volumes.

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

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

Disks:

ada0
ada1
ada2
ada3

Type:

Short description:

Hour:

Every N hour

Each selected hour

1

Day of month:

Every N day of month

Each selected day of month

1

Month:

☒

January

☒

February

☒

March

☒

April

☒

May

☒

June

☒

July

☒

August

☒

September

☒

October

☒

November

☒

December

Fig. 5.9: Adding a S.M.A.R.T. Test

Table 5.5 summarizes the configurable options when creating a S.M.A.R.T. test.

Table 5.5: S.M.A.R.T. Test Options

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

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

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

Which tests will run and when can be verified by typing `smartd -q showtests` within [Shell](#) (page 257).

The results of a test can be checked from [Shell](#) (page 257) by specifying the name of the drive. For example, to see the results for disk `ada0`, type:

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

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

NETWORK

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

- [Global Configuration](#) (page 76): general network settings.
- [Interfaces](#) (page 78): settings for each network interface.
- [IPMI](#) (page 80): settings controlling connection to the appliance through the hardware side-band management interface if the graphical user interface becomes unavailable.
- [Link Aggregations](#) (page 82): settings for network link aggregation and link failover.
- [Network Summary](#) (page 86): display an overview of the current network settings.
- [Static Routes](#) (page 86): add static routes.
- [VLANs](#) (page 87): configure IEEE 802.1q tagging for virtual LANs.

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

6.1 Global Configuration

Network → Global Configuration, shown in [Figure 6.1](#), is for general network settings that are not unique to any particular network interface.

Network

Global Configuration

Interfaces

Link Aggregations

Network Summary

Static Routes

VLANs

Hostname (This Node):

192.168.1.1

Hostname (Node B):

192.168.1.2

Hostname (Virtual):

Domain:

example.com

Additional domains:

IPv4 Default Gateway:

192.168.1.1

IPv6 Default Gateway:

Nameserver 1:

192.168.1.1

Nameserver 2:

192.168.1.2

Nameserver 3:

HTTP Proxy:

Enable netwait feature:

☐

Netwait IP list:

Host name data base:

Save

Fig. 6.1: Global Network Configuration

Table 6.1 summarizes the settings on the Global Configuration tab. Hostname and domain fields are pre-filled as shown in Figure 6.1, but can be changed to meet requirements of the local network.

Table 6.1: Global Configuration Settings

Setting	Value	Description
Hostname (This Node)	string	host name of first storage controller
Continued on next page		

Table 6.1 – continued from previous page

Setting	Value	Description
Hostname (Node B)	string	host name of second storage controller
Hostname (Virtual)	string	virtual host name; when using a virtualhost, this is also used as the Kerberos principal name; must be the fully qualified hostname plus the domain name
Domain	string	system domain name
Additional domains	string	can enter up to 6 space delimited search domains; adding multiple domains may result in slower DNS lookups
IPv4 Default Gateway	IP address	typically not set (see Note below); if set, used instead of default gateway provided by DHCP
IPv6 Default Gateway	IP address	typically not set (see Note below)
Nameserver 1	IP address	primary DNS server (typically in Windows domain)
Nameserver 2	IP address	secondary DNS server
Nameserver 3	IP address	tertiary DNS server
HTTP Proxy	string	enter the proxy information for the network in the format <code>http://my.proxy.server:3128</code> or <code>http://user:password@my.proxy.server:3128</code>
Enable netwait feature	checkbox	if enabled, network services are not started at boot until the interface is able to ping the addresses listed in Netwait IP list
Netwait IP list	string	if Enable netwait feature is checked, list of IP addresses to ping; otherwise, ping the default gateway
Host name database	string	used to add one entry per line which will be appended to <code>/etc/hosts</code> ; use the format <code>IP_address space hostname</code> where multiple hostnames can be used if separated by a space

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

If the network does not have a DNS server, or NFS, SSH, or FTP users are receiving “reverse DNS” or timeout errors, add an entry for the IP address of the 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 86) to be able to reach external DNS, NTP, and mail servers which are configured with static IP addresses. When a gateway to the Internet is added, make sure that the TrueNAS® system is protected by a properly configured firewall.

6.2 Interfaces

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

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

[Figure 6.2](#) shows the screen that opens on clicking `Interfaces` → `Add Interface`. [Table 6.2](#) summarizes the configuration options shown when adding an interface or editing an already configured interface. Note that if any changes to this

screen require a network restart, the screen will turn red when the OK button is clicked and a pop-up message will point out that network connectivity to the TrueNAS® system will be interrupted while the changes are applied.

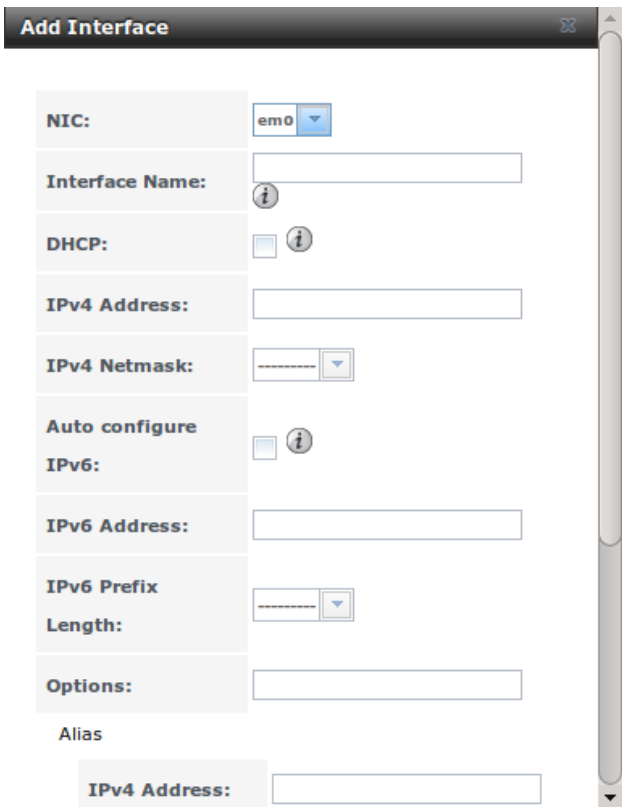


Fig. 6.2: Adding or Editing an Interface

Table 6.2: Interface Configuration Settings

Setting	Value	Description
NIC	drop-down menu	the FreeBSD device name of the interface; a read-only field when editing an interface
Interface Name	string	description of interface
DHCP	checkbox	requires static IPv4 or IPv6 configuration if unchecked; only one interface can be configured for DHCP
IPv4 Address	IP address	enter a static IP address if DHCP is unchecked
IPv4 Netmask	drop-down menu	enter a netmask if DHCP is unchecked
Auto configure IPv6	checkbox	only one interface can be configured for this option; if unchecked, manual configuration is required to use IPv6
IPv6 Address	IPv6 address	must be unique on network
IPv6 Prefix Length	drop-down menu	match the prefix used on network
Options	string	additional parameters from <code>ifconfig(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ifconfig), separate multiple parameters with a space; for example: <code>mtu 9000</code> increases the MTU for interfaces which support jumbo frames (but see this note (page 86) about MTU and lagg interfaces)

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

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

Note: The ability to delete interfaces is disabled if [Failover](#) (page 55) has been configured and enabled.

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

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

6.3 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 in order 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/) (<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 6.3](#), provides a shortcut to the most basic IPMI configuration. Those already familiar with IPMI management tools can use them instead. [Table 6.3](#) summarizes the options available when configuring IPMI with the TrueNAS® GUI.

Network

Global Configuration

Interfaces

IPMI

Link Aggregations

Network Summary

Static Routes

VLANs

Channel:

1

Password:

Password confirmation:

DHCP:

☒

IPv4 Address:

IPv4 Netmask:

/16 (255.255.0.0)

IPv4 Default Gateway:

VLAN ID:

OK

Cancel

Identify Light

Fig. 6.3: IPMI Configuration

Table 6.3: IPMI Options

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

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

After configuration, the IPMI interface is accessed using a web browser and the IP address specified in the configuration. The management interface prompts for a username (the default is admin) and the configured password.

After logging in to the management interface, the administrative username can be changed and additional users can be created.

Refer to [Figure 2.8](#) through [Figure 2.11](#) in [Out-of-Band Management](#) (page 4) for additional instructions on how to configure the Java KVM Client used by the IPMI management interface.

A command-line utility called `ipmitool` is available to control many features of the IPMI interface. See [How To: Change IPMI](#)

Sensor Thresholds using `ipmitool` (<https://forums.freenas.org/index.php?resources/how-to-change-ipmi-sensor-thresholds-using-ipmitool.35/>) for some examples.

6.4 Link Aggregations

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

Aggregation works best on switches supporting LACP, which distributes traffic bi-directionally while responding to failure of individual links. 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 after that are used as failover devices. By default, received traffic is only accepted when received through the active port. This constraint can be relaxed, which is useful for certain bridged network setups, by creating a tunable with a Variable of `net.link.lagg.failover_rx_all`, a Value of a non-zero integer, and a Type of `Sysctl` in `System` → `Tunables` → `Add Tunable`.

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 cases there will only be one LAG which contains all ports. In the event of changes in physical connectivity, link aggregation will quickly converge to a new configuration. LACP must be configured on the switch, and LACP does not support mixing interfaces of different speeds. Only interfaces that use the same driver, like two `igb` ports, are recommended for LACP. Using LACP for iSCSI is not recommended, as iSCSI has built-in multipath features which are more efficient.

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

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

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

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

6.4.1 LACP, MPIO, NFS, and ESXi

LACP bonds Ethernet connections to improve bandwidth. For example, four physical interfaces can be used to create one mega interface. However, it cannot increase the bandwidth for a single conversation. It is designed to increase bandwidth

when multiple clients are simultaneously accessing the same system. It also assumes that quality Ethernet hardware is used and it will not make much difference when using inferior Ethernet chipsets such as a Realtek.

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

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

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

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

6.4.2 Creating a Link Aggregation

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

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

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

Figure 6.4 shows the configuration options when adding a lagg interface using `Network → Link Aggregations → Create Link Aggregation`.

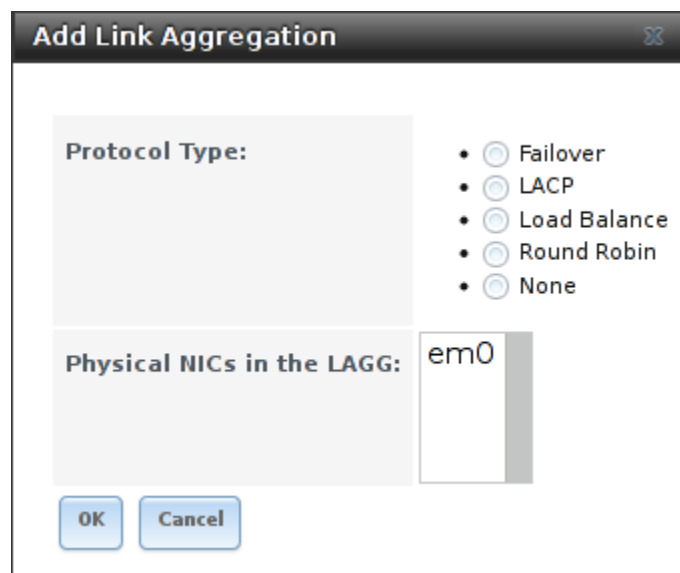


Fig. 6.4: Creating a lagg Interface

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

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

Clicking the Edit button for a lagg opens the configuration screen shown in Figure 6.5. Table 6.4 describes the options in this screen.

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

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




Fig. 6.5: Editing a lagg

Table 6.4: Configurable Options for a lagg

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

Continued on next page

Table 6.4 – continued from previous page

Setting	Value	Description
Auto configure IPv6	checkbox	check only if DHCP server available to provide IPv6 address info
IPv6 Address	string	optional
IPv6 Prefix Length	drop-down menu	required if an IPv6 address is entered
Options	string	additional <code>ifconfig(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ifconfig) options

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

Click the Edit Members button, click the entry for a member, then click its Edit button to see the configuration screen shown in Figure 6.6. The configurable options are summarized in Table 6.5.




Fig. 6.6: Editing a Member Interface

Table 6.5: Configuring a Member Interface

Setting	Value	Description
LAGG Interface group	drop-down menu	select the member interface to configure
LAGG Priority Number	integer	order of selected interface within the lagg; configure a failover to set the master interface to 0 and the other interfaces to 1, 2, etc.
LAGG Physical NIC	drop-down menu	physical interface of the selected member

Continued on next page

Table 6.5 – continued from previous page

Setting	Value	Description
Options	string	additional parameters from ifconfig(8) (http://www.freebsd.org/cgi/man.cgi?query=ifconfig)

Options can be set at the lagg level using the Edit button, or at the individual parent interface level using the Edit Members button. Changes are typically made at the lagg level ([Figure 6.5](#)) as each interface member will inherit from the lagg. To configure at the interface level ([Figure 6.6](#)) instead, the configuration must be repeated for each interface within the lagg. Some options can only be set on the parent interfaces and are inherited by the lagg interface. For example, to set the MTU on a lagg, use Edit Members to set the MTU for each parent interface.

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

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

To see if the link aggregation is properly load balancing, run this command from [Shell](#) (page 257):

```
systat -ifstat
```

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

6.5 Network Summary

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

6.6 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 with [Network](#) → [Static Routes](#) → [Add Static Route](#), shown in [Figure 6.7](#).

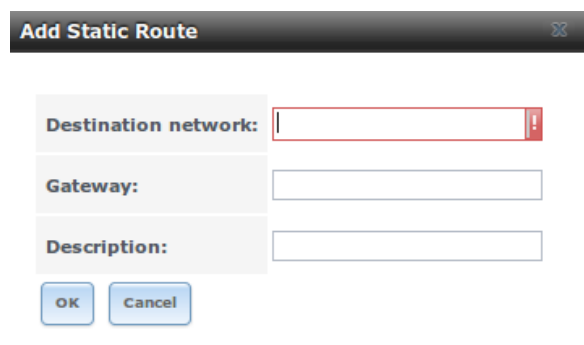


Fig. 6.7: Adding a Static Route

The available options are summarized in [Table 6.6](#).

Table 6.6: Static Route Options

Setting	Value	Description
Destination network	integer	use the format A.B.C.D/E where E is the CIDR mask

Continued on next page

Table 6.6 – continued from previous page

Setting	Value	Description
Gateway	integer	enter the IP address of the gateway
Description	string	optional

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

6.7 VLANs

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

Note: VLAN tagging is the only 802.1q feature that is implemented.

Click **Network** → **VLANs** → **Add VLAN**, to see the screen shown in [Figure 6.8](#).

Fig. 6.8: Adding a VLAN

[Table 6.7](#) summarizes the configurable fields.

Table 6.7: Adding a VLAN

Setting	Value	Description
Virtual Interface	string	use the format vlanX where X is a number representing a vlan interface not currently being used as a parent
Parent Interface	drop-down menu	usually an Ethernet card connected to a properly configured switch port; note that newly created Link Aggregations (page 82) will not appear in the drop-down until the system is rebooted
VLAN Tag	integer	number between 1 and 4095 which matches a numeric tag set up in the switched network

Continued on next page

Table 6.7 – continued from previous page

Setting	Value	Description
Priority Code Point	drop-down menu	available 802.1p Class of Service ranges from Best Effort (default) to Network Control (highest)
Description	string	optional

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

Warning: Creating a vlan will cause network connectivity to be interrupted and, if [Failover](#) (page 55) has been configured, a failover event. Accordingly, the GUI will provide a warning and an opportunity to cancel the vlan creation.

STORAGE

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

- [Volumes](#) (page 89) creates and manages storage volumes.
- [Periodic Snapshot Tasks](#) (page 113) schedules automatic creation of filesystem snapshots.
- [Replication Tasks](#) (page 115) automate the replication of snapshots to a remote system.
- [Resilver Priority](#) (page 125) controls the priority of resilvers.
- [Scrubs](#) (page 126) schedules scrubs as part of ongoing disk maintenance.
- [Snapshots](#) (page 129) manages local snapshots.
- [VMware-Snapshot](#) (page 131) coordinates ZFS snapshots with a VMware datastore.

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

7.1 Volumes

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

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

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

7.1.1 Volume Manager

The Volume Manager is used to add disks to a ZFS pool. Any old data on added disks is overwritten, so save it elsewhere before reusing a disk. Please see the [ZFS Primer](#) (page 263) for information on ZFS redundancy with multiple disks before using Volume Manager. It is important to realize that different layouts of virtual devices (vdevs) affect which operations can be performed on that volume later. For example, drives can be added to a mirror to increase redundancy, but that is not possible with RAIDZ arrays.

Selecting `Storage → Volumes → Volume Manager` opens a screen like the example shown in [Figure 7.1](#).

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

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

Fig. 7.1: Creating a ZFS Pool Using Volume Manager

Table 7.1 summarizes the configuration options of this screen.

Table 7.1: ZFS Volume Creation Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html); choosing a name that will stick out in the logs (e.g. not a generic term like <code>data</code> or <code>freenas</code>) is recommended
Volume to extend	drop-down menu	extend an existing ZFS pool; see Extending a ZFS Volume (page 93) for more details
Encryption	checkbox	see the warnings in Encryption (page 91) before enabling encryption
Available disks	display	display the number and size of available disks; hover over show to list the available device names; click the + to add all of the disks to the pool
Volume layout	drag and drop	click and drag the icon to select the desired number of disks for a vdev; when at least one disk is selected, the layouts supported by the selected number of disks are added to the drop-down menu
Add Extra Device	button	configure multiple vdevs or add log or cache devices during pool creation
Manual setup	button	create a pool manually (not recommended); see Manual Setup (page 92) for details

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

Note: For performance and capacity reasons, this screen does not allow creating a volume from disks of differing sizes. While it is not recommended, it is possible to create a volume of differently-sized disks with the Manual setup button. Follow the instructions in [Manual Setup](#) (page 92).

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

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

When more than five disks are used, consideration must be given to the optimal layout for the best performance and scalability. An overview of the recommended disk group sizes as well as more information about log and cache devices can be found in the [ZFS Primer](#) (page 263).

The Add Volume button warns that existing data will be cleared. In other words, creating a new volume reformats the selected disks. To preserve existing data, click the Cancel button and refer to [Import Disk](#) (page 100) and [Import Volume](#) (page 101) to see if the existing format is supported. If so, perform that action instead. If the current storage format is not supported, it is necessary to back up the data to external media, format the disks, then restore the data to the new volume.

Depending on the size and number of disks, the type of controller, and whether encryption is selected, creating the volume may take some time. After the volume is created, the screen refreshes and the new volume is listed in the tree under `Storage` → `Volumes`. Click the + next to the volume name to access [Change Permissions](#) (page 94), [Create Dataset](#) (page 96), and [Create zvol](#) (page 99) options for that volume.

Encryption

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

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

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

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

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

- At present, there is no one-step way to encrypt an existing, unencrypted volume. Instead, the data must be backed up, the existing pool destroyed, a new encrypted volume created, and the backup restored to the new volume.
- Hybrid pools are not supported. Added vdevs must match the existing encryption scheme. The [Volume Manager](#) (page 89) automatically encrypts a new vdev being added to an existing encrypted pool.

To create an encrypted volume, check the Encryption box shown in [Figure 7.1](#). A pop-up message shows a reminder that it is extremely important to make a backup of the key. Without the key, the data on the disks is inaccessible. See [Managing Encrypted Volumes](#) (page 106) for instructions.

Encryption Performance

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

Manual Setup

The Manual Setup button shown in [Figure 7.1](#) can be used to create a ZFS volume manually. While this is not recommended, it can, for example, be used to create a non-optimal volume containing disks of different sizes.

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

[Figure 7.2](#) shows the Manual Setup screen. [Table 7.2](#) shows the available options.

Fig. 7.2: Manually Creating a ZFS Volume

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

Table 7.2: Manual Setup Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (http://docs.oracle.com/cd/E19082-01/817-2271/gbcpt/index.html); choose a name that will stand out in the logs (e.g. not <code>data</code> or <code>freenas</code>)
Encryption	checkbox	see the warnings in Encryption (page 91) before using encryption
Member disks	list	highlight desired number of disks from list of available disks
Deduplication	drop-down menu	do not change this setting unless instructed to do so by an iXsystems support engineer
ZFS Extra	bullet selection	specify disk usage: storage (None), a log device, a cache device, or a spare

Extending a ZFS Volume

The Volume to extend drop-down menu in `Storage → Volumes → Volume Manager`, shown in [Figure 7.1](#), is used to add disks to an existing ZFS volume to increase capacity. This menu is empty if there are no ZFS volumes yet.

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

Note: If the existing volume is encrypted, a warning message shows a reminder that extending a volume resets the passphrase and recovery key. After extending the volume, immediately recreate both using the instructions in [Managing Encrypted Volumes](#) (page 106).

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

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

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

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

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

Adding L2ARC or SLOG Devices

Storage → Volumes → Volume Manager (see [Figure 7.1](#)) is also used to add L2ARC or SLOG SSDs to improve specific types of volume performance. This is described in more detail in the [ZFS Primer](#) (page 263).

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

7.1.2 Change Permissions

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

The chapter on [Sharing](#) (page 143) contains configuration examples for several types of permission scenarios. This section provides an overview of the screen that is used to set permissions.

Note: For users and groups to be available, they must either be first created using the instructions in [Account](#) (page 19) or imported from a directory service using the instructions in [Directory Services](#) (page 132). If more than 50 users or groups are available, the drop-down menus described in this section will automatically truncate their display to 50 for performance reasons. In this case, start to type in the desired user or group name so that the display narrows its search to matching results.

After a volume or dataset is created, it is listed by its mount point name in Storage → Volumes. Clicking the Change Permissions icon for a specific volume/dataset displays the screen shown in [Figure 7.3](#). [Table 7.3](#) summarizes the options in this screen.

Change Permissions

Change permission

Change permission on /mnt/volume1 to:

Apply Owner (user): ☒

Owner (user): root

Apply Owner (group): ☒

Owner (group): wheel

Apply Mode: ☒

Mode:

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

Permission Type:

- ☒ Unix
- ☐ Mac
- ☐ Windows

Set permission recursively: ☐

Fig. 7.3: Changing Permissions on a Volume or Dataset

Table 7.3: Options When Changing Permissions

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

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

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

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

The Mac Permission Type is used with AFP shares.

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

7.1.3 Create Dataset

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

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

Selecting an existing ZFS volume in the tree and clicking Create Dataset shows the screen in [Figure 7.1.3](#).

Create Dataset

Create ZFS dataset in volume1

Dataset Name:

Comments:

Sync:

Inherit (standard)

Compression level:

Inherit (lz4)

Share type:

UNIX

Enable atime:

☒ Inherit (on)

☐ On

☐ Off

ZFS Deduplication:

Inherit (off)

Case Sensitivity:

Sensitive

Add Dataset

Cancel

Advanced Mode

Dedup feature not activated. Contact [TrueNAS Support](#) for assistance.

Fig. 7.4: Creating a ZFS Dataset

Table 7.4 shows the options available when creating a dataset. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button, or configure the system to always display advanced settings by checking the box Show advanced fields by default in `System → Advanced`. Most attributes, except for the Dataset Name, Case Sensitivity, and Record Size, can be changed after dataset creation by highlighting the dataset name and clicking the Edit Options button in `Storage → Volumes`.

Table 7.4: ZFS Dataset Options

Setting	Value	Description
Dataset Name	string	mandatory; enter a unique name for the dataset
Comments	string	short comments or user notes about this dataset
Sync	drop-down menu	data write synchronization: Inherit inherits the sync settings from the parent dataset; Standard uses the sync settings that have been requested by the client software; Always always waits for data writes to complete; Disabled never waits for writes to complete
Compression Level	drop-down menu	see the section on Compression (page 98) for a description of the available algorithms

Continued on next page

Table 7.4 – continued from previous page

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

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

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

Compression

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

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

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

7.1.4 Create zvol

A zvol is a feature of ZFS that creates a raw block device over ZFS. This allows you to use a zvol as an [iSCSI](#) (page 199) device extent.

To create a zvol, select an existing ZFS volume or dataset from the tree then click Create zvol to open the screen shown in [Figure 7.5](#).

Create zvol

Create zvol on volume1

zvol name:

Comments:

Size for this zvol: ⓘ

Force size: ☐ ⓘ

Sync: Inherit (standard) ▼

Compression level: Inherit (lz4) ▼

ZFS Deduplication: Enabling dedup can drastically reduce performance and affect the ability to access data. Compression usually offers similar space savings with much lower performance impact and overhead.
Inherit (off) ▼

Sparse volume: ☐ ⓘ

Add zvol Cancel Advanced Mode

Fig. 7.5: Creating a Zvol

The configuration options are described in [Table 7.5](#). Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking Show advanced fields by default in `System → Advanced`.

Table 7.5: zvol Configuration Options

Setting	Value	Description
zvol Name	string	mandatory; enter a name for the zvol; note that there is a 63-character limit on device path names in devfs, so using long zvol names can prevent accessing zvols as devices; for example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent
Comments	string	short comments or user notes about this zvol
Size for this zvol	integer	specify size and value such as 10Gib; if the size is more than 80% of the available capacity, the creation will fail with an “out of space” error unless Force size is checked
Force size	checkbox	by default, the system will not let you create a zvol if that operation will bring the pool to over 80% capacity; while NOT recommended, checking this box will force the creation of the zvol in this situation
Compression level	drop-down menu	see the section on Compression (page 98) for a description of the available algorithms
Sparse volume	checkbox	used to provide thin provisioning; use with caution for when this option is selected, writes will fail when the pool is low on space
Block size	drop-down menu	only available in Advanced Mode and by default is based on the number of disks in pool; can be set to match the block size of the filesystem which will be formatted onto the iSCSI target

7.1.5 Import Disk

The **Volume → Import Disk** screen, shown in [Figure 7.6](#), is used to import a single disk that has been formatted with the UFS, NTFS, MSDOS, or EXT2 filesystem. The import is meant to be a temporary measure to copy the data from a disk to an existing ZFS dataset. Only one disk can be imported at a time.

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

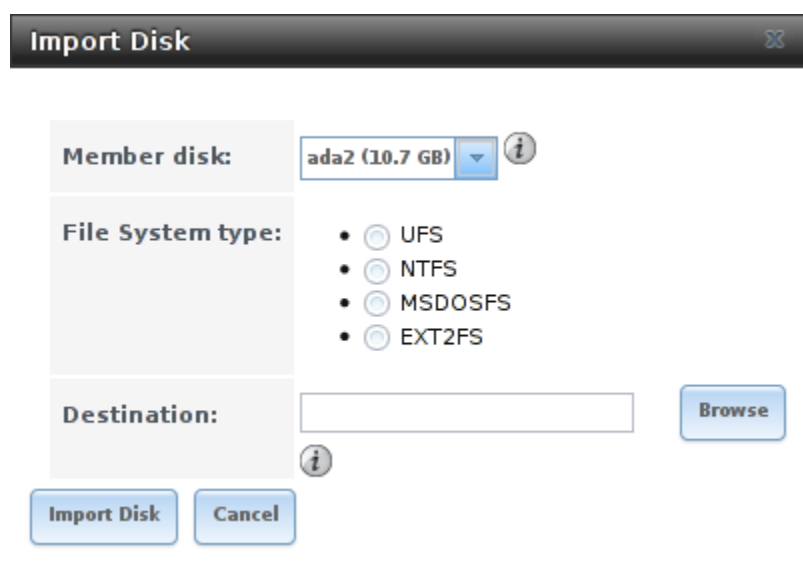


Fig. 7.6: Importing a Disk

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

7.1.6 Import Volume

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

Figure 7.7 shows the initial pop-up window that appears when you import a volume.

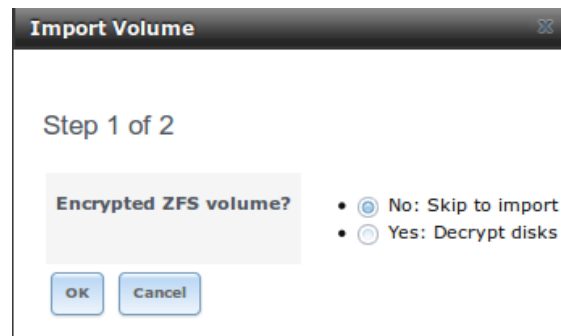


Fig. 7.7: Initial Import Volume Screen

If you are importing an unencrypted ZFS pool, select No: Skip to import to open the screen shown in Figure 7.8.



Fig. 7.8: Importing a Non-Encrypted Volume

Existing volumes should be available for selection from the drop-down menu. In the example shown in Figure 7.8, the TrueNAS® system has an existing, unencrypted ZFS pool. Once the volume is selected, click the OK button to import the volume.

If an existing ZFS pool does not show in the drop-down menu, run `zpool import` from [Shell](#) (page 257) to import the pool.

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

Importing an Encrypted Pool

Disks in existing GELI-encrypted ZFS pools must be decrypted before importing the pool. In the Import Volume dialog shown in Figure 7.7, select Yes: Decrypt disks. The screen shown in Figure 7.9 is then displayed.




Fig. 7.9: Decrypting Disks Before Importing a ZFS Pool

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

Note: The encryption key is required to decrypt the pool. If the pool cannot be decrypted, it cannot be re-imported after a failed upgrade or lost configuration. This means that it is very important to save a copy of the key and to remember the passphrase that was configured for the key. Refer to [Managing Encrypted Volumes](#) (page 106) for instructions on how to manage the keys for encrypted volumes.

After the pool is decrypted, it appears in the drop-down menu of [Figure 7.8](#). Click the OK button to finish the volume import.

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

7.1.7 View Disks

Storage → Volumes → View Disks shows all of the disks recognized by the TrueNAS® system. An example is shown in [Figure 7.10](#).

View Disks										
Name	Serial	Disk Size	Description	Transfer Mode	HDD Standby	Advanced Power Management	Acoustic Level	Enable S.M.A.R.T.	S.M.A.R.T. extra options	Enclosure Slot
da0	STM000199402	8.0 GB		Auto	Always On	Disabled	Disabled	true		1
da1	STM000190111	800.2 GB		Auto	Always On	Disabled	Disabled	true		2
da2		6.0 TB		Auto	Always On	Disabled	Disabled	true		3
da3		6.0 TB		Auto	Always On	Disabled	Disabled	true		4
da4		6.0 TB		Auto	Always On	Disabled	Disabled	true		5
da5		6.0 TB		Auto	Always On	Disabled	Disabled	true		6
da6		6.0 TB		Auto	Always On	Disabled	Disabled	true		7
da7		6.0 TB		Auto	Always On	Disabled	Disabled	true		8
da8		6.0 TB		Auto	Always On	Disabled	Disabled	true		9
da9		6.0 TB		Auto	Always On	Disabled	Disabled	true		10
da10		6.0 TB		Auto	Always On	Disabled	Disabled	true		11
da11		6.0 TB		Auto	Always On	Disabled	Disabled	true		12
da12		6.0 TB		Auto	Always On	Disabled	Disabled	true		13
da13		6.0 TB		Auto	Always On	Disabled	Disabled	true		14
da14		6.0 TB		Auto	Always On	Disabled	Disabled	true		15
da15		6.0 TB		Auto	Always On	Disabled	Disabled	true		16

Fig. 7.10: Viewing Disks

The current configuration of each device is displayed. Click a disk entry and the Edit button to change its configuration. The configurable options are described in [Table 7.6](#).

Table 7.6: Disk Options

Setting	Value	Description
Name	string	read-only value showing FreeBSD device name for disk
Serial	string	read-only value showing the disk's serial number
Description	string	optional
HDD Standby	drop-down menu	indicates the time of inactivity (in minutes) before the drive enters standby mode in order to conserve energy; this forum post (https://forums.freenas.org/index.php?threads/how-to-find-out-if-a-drive-is-spinning-down-properly.2068/) demonstrates how to determine if a drive has spun down
Advanced Power Management	drop-down menu	default is Disabled, can select a power management profile from the menu
Acoustic Level	drop-down menu	default is Disabled; can be modified for disks that understand AAM (https://en.wikipedia.org/wiki/Automatic_acoustic_management)
Enable S.M.A.R.T.	checkbox	enabled by default if the disk supports S.M.A.R.T.; unchecking this box will disable any configured S.M.A.R.T. Tests (page 73) for the disk
S.M.A.R.T. extra options	string	additional smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools) options

Note: If a disk's serial number is not displayed in this screen, use the `smartctl` command from [Shell](#) (page 257). For example, to determine the serial number of disk `ada0`, type `smartctl -a /dev/ada0 | grep Serial`.

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

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

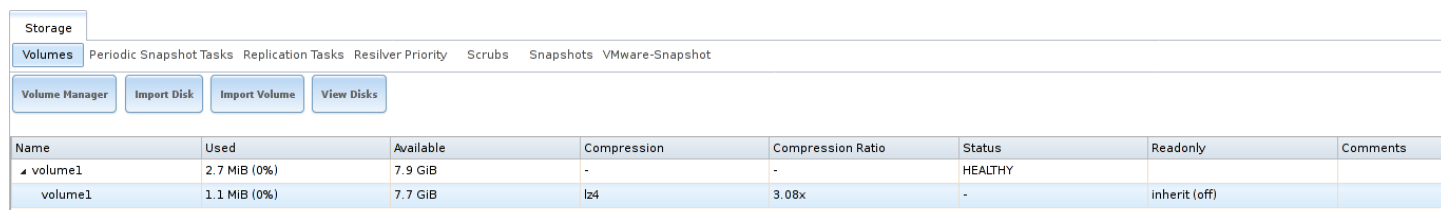
Quick wipes take only a few seconds. A Full with zeros wipe of a large disk can take several hours, and a Full with random data takes longer. A progress bar is displayed during the wipe to track status.

7.1.8 Volumes

Storage → Volumes is used to view and further configure existing ZFS pools, datasets, and zvols. The example shown in Figure 7.11 shows one ZFS pool (volume1) with two datasets (the one automatically created with the pool, volume1, and dataset1) and one zvol (zvol1).

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

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



The screenshot shows the 'Storage' section with the 'Volumes' tab selected. Below the tab are buttons for 'Volume Manager', 'Import Disk', 'Import Volume', and 'View Disks'. A table lists the following volumes:

Name	Used	Available	Compression	Compression Ratio	Status	Readonly	Comments
volume1	2.7 MiB (0%)	7.9 GiB	-	-	HEALTHY		
volume1	1.1 MiB (0%)	7.7 GiB	lz4	3.08x	-	inherit (off)	

Fig. 7.11: Viewing Volumes

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

Note: When the system has [High Availability \(HA\)](#) (page 55) active, volumes cannot be exported or destroyed.

Detach Volume: allows exporting the pool or deleting the contents of the pool, depending upon the choice made in the screen shown in Figure 7.12. The Detach Volume screen displays the current used space and indicates whether there are any shares, provides checkboxes to Mark the disks as new (destroy data) and to Also delete the share's configuration, and asks if you are sure about doing this. The browser window turns red to indicate that some choices will make the data inaccessible. When the box to mark the disks as new is left unchecked, the volume is exported. The data is not destroyed and the volume can be re-imported at a later time. When moving a ZFS pool from one system to another, perform this export action first as it flushes any unwritten data to disk, writes data to the disk indicating that the export was done, and removes all knowledge of the pool from the system.

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




Fig. 7.12: Detach or Delete a Volume

Scrub Volume: scrubs and scheduling them are described in more detail in [Scrubs](#) (page 126). This button allows manually initiating a scrub. Scrubs are I/O intensive and can negatively impact performance. Avoid initiating a scrub when the system is busy.

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

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

Volume Status: as shown in the example in [Figure 7.13](#), this screen shows the device name and status of each disk in the ZFS pool as well as any read, write, or checksum errors. It also indicates the status of the latest ZFS scrub. Clicking the entry for a device causes buttons to appear to edit the device's options (shown in [Figure 7.14](#)), offline or online the device, or replace the device (as described in [Replacing a Failed Drive](#) (page 110)).

Upgrade: used to upgrade the pool to the latest ZFS features, as described in [Upgrading a ZFS Pool](#) (page 42). This button does not appear if the pool is running the latest version of feature flags.

Volume Status				
Scrub				
Status: Completed				
Errors: 0 Repaired: 0 Date: Sun Jan 25 00:00:02 2015				
Name	Read	Write	Checksum	Status
▲ tank	0	0	0	ONLINE
▲ mirror-6	0	0	0	ONLINE
da15p1	0	0	0	ONLINE
da14p1	0	0	0	ONLINE
▲ mirror-5	0	0	0	ONLINE
da13p1	0	0	0	ONLINE
da12p1	0	0	0	ONLINE
▲ mirror-4	0	0	0	ONLINE
da11p1	0	0	0	ONLINE
da10p1	0	0	0	ONLINE
▲ mirror-3	0	0	0	ONLINE
da9p1	0	0	0	ONLINE

Fig. 7.13: Volume Status

Selecting a disk in Volume Status and clicking its Edit Disk button shows the screen in [Figure 7.14](#). [Table 7.6](#) summarizes the configurable options.

Edit

Name:

ada0

Serial:

JP2940HZ3SNPDC

Description:

HDD Standby:

Always On

Advanced Power Management:

Disabled

Acoustic Level:

Disabled

Enable S.M.A.R.T.

☒

S.M.A.R.T. extra options:

OK

Cancel

Fig. 7.14: Editing a Disk

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

Change Permissions: edit the dataset's permissions as described in [Change Permissions](#) (page 94).

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use [Periodic Snapshot Tasks](#) (page 113).

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

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

Edit Options: edit the volume's properties described in [Table 7.1.3](#). Note that it will not allow changing the dataset's name.

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

Create zvol: create a child zvol within this dataset.

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

Managing Encrypted Volumes

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

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

- **Key stored locally, no passphrase:** the encrypted volume is decrypted and accessible when the system running. Protects "data at rest" only.

- Key stored locally, with passphrase: the encrypted volume is not accessible until the passphrase is entered by the TrueNAS® administrator.
- Key not stored locally: the encrypted volume is not accessible until the TrueNAS® administrator provides the key. If a passphrase is set on the key, it must also be entered before the encrypted volume can be accessed ([two factor authentication](https://en.wikipedia.org/wiki/Multi-factor_authentication) (https://en.wikipedia.org/wiki/Multi-factor_authentication)).

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

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

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

Additional Controls for Encrypted Volumes

If the Encryption box is checked during the creation of a pool, additional buttons appear in the entry for the volume in Storage → Volumes. An example is shown in [Figure 7.15](#).

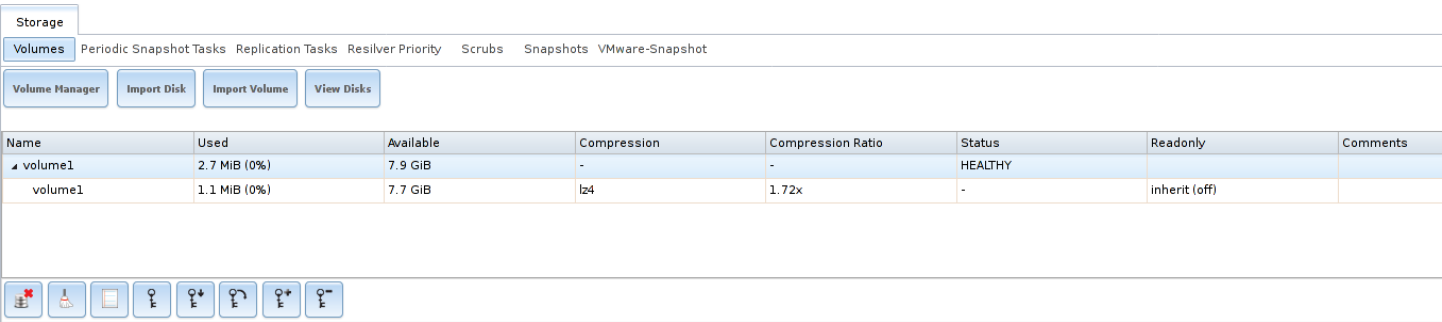


Fig. 7.15: Encryption Icons Associated with an Encrypted Volume

These additional encryption buttons are used to:

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



Fig. 7.16: Add or Change a Passphrase to an Encrypted Volume

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

Encrypted volumes with a passphrase display an additional lock button:



Fig. 7.17: Lock Button

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



Fig. 7.18: Unlock Button

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




Fig. 7.19: Unlock Locked Volume

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

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

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

Note: A re-key is not allowed if [Failover](#) (page 55) (High Availability) has been enabled and the standby node is down.

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

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

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

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

7.1.9 View Multipaths

TrueNAS® uses [`gmlinux\(8\)`](http://www.freebsd.org/cgi/man.cgi?query=gmlinux(8)) (<http://www.freebsd.org/cgi/man.cgi?query=gmlinux>) 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 hardware that is capable of multipath. An example would be a dual SAS expander backplane in the chassis or an external JBOD.

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

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

7.1.10 Replacing a Failed Drive

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

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

Note: If the volume is encrypted with GELI, refer to [Replacing an Encrypted Drive](#) (page 112) before proceeding.

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

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

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

2. After the disk has been replaced and is showing as OFFLINE, click the disk again and then click its Replace button. Select the replacement disk from the drop-down menu and click the Replace Disk button. After clicking the Replace Disk button, the ZFS pool begins resilvering.
3. After the drive replacement process is complete, re-add the replaced disk in the [S.M.A.R.T. Tests](#) (page 73) screen.

In the example shown in [Figure 7.20](#), a failed disk is being replaced by disk `ada5` in the volume named `volume1`.

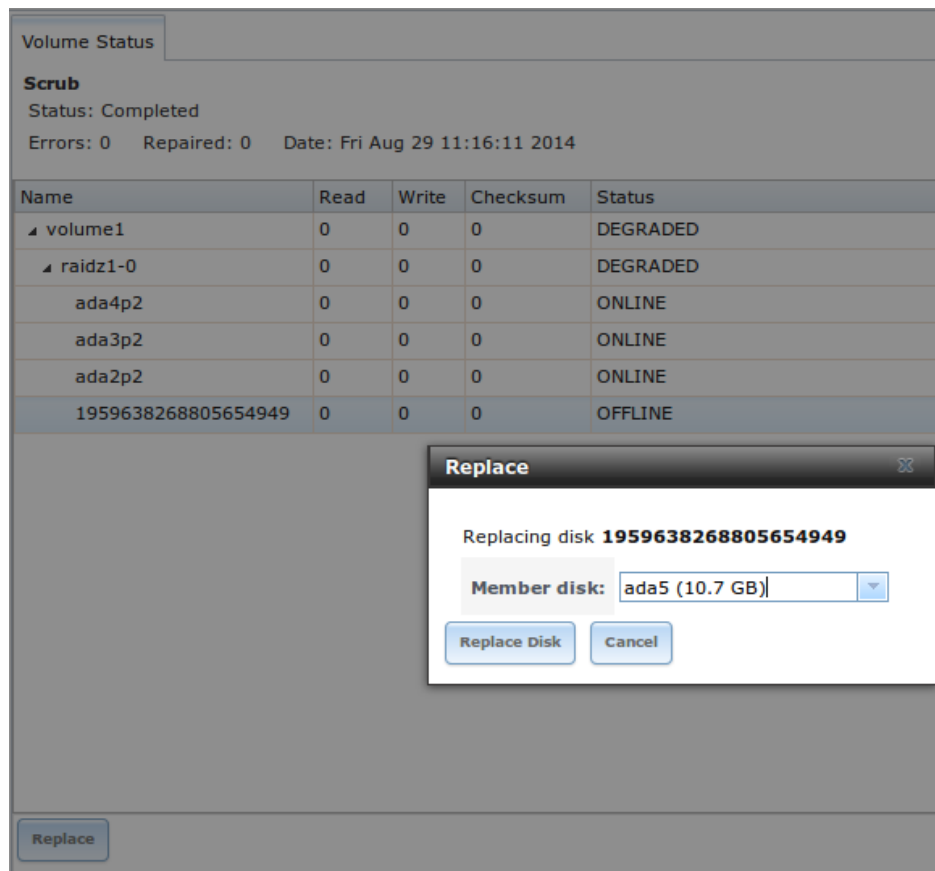


Fig. 7.20: Replacing a Failed Disk

After the resilver is complete, Volume Status shows a Completed resilver status and indicates any errors. [Figure 7.21](#) indicates that the disk replacement was successful in this example.

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

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

Fig. 7.21: Disk Replacement is Complete

Replacing an Encrypted Drive

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

First, make sure that a passphrase has been set using the instructions in [Encryption](#) (page 91) before attempting to replace the failed drive. Then, follow the steps 1 and 2 as described above. During step 3, you will be prompted to input and confirm the passphrase for the pool. Enter this information then click the Replace Disk button. Wait until the resilvering is complete.

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

1. Highlight the pool that contains the disk that was just replaced and click the Encryption Re-key button in the GUI. Entry of the root password will be required.

Note: A re-key is not allowed if [Failover](#) (page 55) (High Availability) has been enabled and the standby node is down.

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

Removing a Log or Cache Device

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

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

7.1.11 Replacing Drives to Grow a ZFS Pool

The recommended method for expanding the size of a ZFS pool is to pre-plan the number of disks in a vdev and to stripe additional vdevs using [Volume Manager](#) (page 89) as additional capacity is needed.

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

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

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

If a spare drive port is not available, a drive can be replaced with a larger one using the instructions in [Replacing a Failed Drive](#) (page 110). This process is slow and places the system in a degraded state. Since a failure at this point could be disastrous, do not attempt this method unless the system has a reliable backup. Replace one drive at a time and wait for the resilver process to complete on the replaced drive before replacing the next drive. After all the drives are replaced and the final resilver completes, the added space will appear in the pool.

7.1.12 Hot Spares

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

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

7.2 Periodic Snapshot Tasks

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

Snapshots provide a clever way of keeping a history of files, providing a way to recover an older copy or even a deleted file. For this reason, many administrators take snapshots often (perhaps every fifteen minutes), store them for a period of time (possibly a month), and store them on another system (typically using [Replication Tasks](#) (page 115)). Such a strategy allows the administrator to roll the system back to a specific point in time. If there is a catastrophic loss, an off-site snapshot can be used to restore the system up to the time of the last snapshot.

An existing ZFS volume is required before creating a snapshot. Creating a volume is described in [Volume Manager](#) (page 89).

To create a periodic snapshot task, click *Storage* → *Periodic Snapshot Tasks* → *Add Periodic Snapshot* which opens the screen shown in [Figure 7.22](#). [Table 7.7](#) summarizes the fields in this screen.

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

Fig. 7.22: Creating a Periodic Snapshot

Table 7.7: Options When Creating a Periodic Snapshot

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

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

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

display Edit and Delete buttons for it.

7.3 Replication Tasks

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

The basic configuration requires a source system with the original data and a destination system where the data will be replicated. The destination system is prepared to receive replicated data, a [periodic snapshot](#) (page 113) of the data on the source system is created, and then a replication task is created. As snapshots are automatically created on the source computer, they are automatically replicated to the destination computer.

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

Note: The target dataset on the receiving system is automatically created in read-only mode to protect the data. To mount or browse the data on the receiving system, create a clone of the snapshot and use the clone. Clones are created in read/write mode, making it possible to browse or mount them. See [Snapshots](#) (page 129) for more information on creating clones.

7.3.1 Examples: Common Configuration

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

Alpha (Source)

Alpha is the source computer with the data to be replicated. It is at IP address 10.0.0.102. A [volume](#) (page 89) named alphavol has already been created, and a [dataset](#) (page 96) named alphadata has been created on that volume. This dataset contains the files which will be snapshotted and replicated onto Beta.

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

Create a periodic snapshot of the source dataset by selecting `Storage → Periodic Snapshot Tasks`. Click the alphavol/alphadata dataset to highlight it. Create a [periodic snapshot](#) (page 113) of it by clicking Periodic Snapshot Tasks, then Add Periodic Snapshot as shown in [Figure 7.23](#).

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

Fig. 7.23: Create a Periodic Snapshot for Replication

Beta (Destination)

Beta is the destination computer where the replicated data will be copied. It is at IP address 10.0.0.118. A [volume](#) (page 89) named betavol has already been created.

Snapshots are transferred with [SSH](#) (page 213). To allow incoming connections, this service is enabled on Beta. The service is not required for outgoing connections, and so does not need to be enabled on Alpha.

7.3.2 Example: TrueNAS® to TrueNAS® Semi-Automatic Setup

TrueNAS® offers a special semi-automatic setup mode that simplifies setting up replication. Create the replication task on Alpha by clicking Replication Tasks and Add Replication. alphavol/alphadata is selected as the dataset to replicate. betavol is the destination volume where alphadata snapshots are replicated. The Setup mode dropdown is set to Semi-automatic as shown in [Figure 7.24](#). The IP address of Beta is entered in the Remote hostname field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP → HTTPS Redirect has been enabled in `System → General` on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

Add Replication

Volume/Dataset:

alphavol/alphadata

Remote ZFS Volume/Dataset:

betavol

Recursively replicate child dataset's snapshots:

☐

Delete stale snapshots on remote system:

☐

Replication Stream Compression:

lz4 (fastest)

Limit (kB/s):

0

Begin:

00:00:00

End:

23:59:00

Enabled:

☒

Setup mode:

Semi-automatic

Remote hostname:

10.0.0.118

Remote HTTP/HTTPS Port:

80

Remote HTTPS:

☐

Remote Auth Token:

Dedicated User Enabled:

☐

Dedicated User:

Encryption Cipher:

Standard

OK

Cancel

This method only works with remote version greater or equal than 9.10.2

On the remote host go to Storage -> Replication Tasks, click the Temporary Auth Token button and paste the resulting value in to this field.

Fig. 7.24: Add Replication Dialog, Semi-Automatic

The Remote Auth Token field expects a special token from the Beta computer. On Beta, choose `Storage` → `Replication Tasks`, then click `Temporary Auth Token`. A dialog showing the temporary authorization token is shown as in [Figure 7.25](#).

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



Fig. 7.25: Temporary Authentication Token on Destination

On the Alpha system, paste the copied temporary authorization token string into the Remote Auth Token field as shown in Figure 7.26.




Fig. 7.26: Temporary Authentication Token Pasted to Source

Finally, click the OK button to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See [Limiting Replication Times](#) (page 123) for information about restricting when replication is allowed to run.

Note: The temporary authorization token is only valid for a few minutes. If a Token is invalid message is shown, get a new temporary authorization token from the destination system, clear the Remote Auth Token field, and paste in the new one.

7.3.3 Example: TrueNAS® to TrueNAS® Dedicated User Replication

A dedicated user can be used for replications rather than the root user. This example shows the process using the semi-automatic replication setup between two TrueNAS® systems with a dedicated user named repluser. SSH key authentication is used to allow the user to log in remotely without a password.

In this example, the periodic snapshot task has not been created yet. If the periodic snapshot shown in the [example configuration](#) (page 115) has already been created, go to **Storage → Periodic Snapshot Tasks**, click on the task to select it, and click Delete to remove it before continuing.

On Alpha, select **Account → Users**. Click the Add User. Enter repluser for Username, enter /mnt/alphavol/repluser in the Create Home Directory In field, enter Replication Dedicated User for the Full Name, and set the Disable password login checkbox. Leave the other fields at their default values, but note the User ID number. Click OK to create the user.

On Beta, the same dedicated user must be created as was created on the sending computer. Select **Account → Users**. Click the Add User. Enter the User ID number from Alpha, repluser for Username, enter /mnt/betavol/repluser in the Create Home Directory In field, enter Replication Dedicated User for the Full Name, and set the Disable password login checkbox. Leave the other fields at their default values. Click OK to create the user.

A dataset with the same name as the original must be created on the destination computer, Beta. Select **Storage → Volumes**, click on betavol, then click the Create Dataset icon at the bottom. Enter alphadata as the Dataset Name, then click Add Dataset.

The replication user must be given permissions to the destination dataset. Still on Beta, open a [Shell](#) (page 257) and enter this command:

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

The destination dataset must also be set to read-only. Enter this command in the [Shell](#) (page 257):

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

Close the [Shell](#) (page 257) by typing `exit` and pressing `Enter`.

The replication user must also be able to mount datasets. Still on Beta, go to `System` → `Tunables`. Click `Add Tunable`. Enter `vfs.usermount` for the Variable, 1 for the Value, and choose `Sysctl` from the Type drop-down. Click `OK` to save the tunable settings.

Back on Alpha, create a periodic snapshot of the source dataset by selecting `Storage` → `Periodic Snapshot Tasks`. Click the `alphavol/alphadata` dataset to highlight it. Create a [periodic snapshot](#) (page 113) of it by clicking `Periodic Snapshot Tasks`, then `Add Periodic Snapshot` as shown in [Figure 7.23](#).

Still on Alpha, create the replication task by clicking `Replication Tasks` and `Add Replication`. `alphavol/alphadata` is selected as the dataset to replicate. `betavol/alphadata` is the destination volume and dataset where `alphadata` snapshots are replicated.

The Setup mode dropdown is set to `Semi-automatic` as shown in [Figure 7.24](#). The IP address of Beta is entered in the Remote hostname field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP → HTTPS Redirect has been enabled in `System` → `General` on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

The Remote Auth Token field expects a special token from the Beta computer. On Beta, choose `Storage` → `Replication Tasks`, then click `Temporary Auth Token`. A dialog showing the temporary authorization token is shown as in [Figure 7.25](#).

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

On the Alpha system, paste the copied temporary authorization token string into the Remote Auth Token field as shown in [Figure 7.26](#).

Set the `Dedicated User` checkbox. Choose `repluser` in the `Dedicated User` drop-down.

Click the `OK` button to create the replication task.

Note: The temporary authorization token is only valid for a few minutes. If a Token is invalid message is shown, get a new temporary authorization token from the destination system, clear the Remote Auth Token field, and paste in the new one.

Replication will begin when the periodic snapshot task runs.

Additional replications can use the same dedicated user that has already been set up. The permissions and read only settings made through the [Shell](#) (page 257) must be set on each new destination dataset.

7.3.4 Example: TrueNAS® to TrueNAS® or Other Systems, Manual Setup

This example uses the same basic configuration of source and destination computers shown above, but the destination computer is not required to be a TrueNAS® system. Other operating systems can receive the replication if they support SSH, ZFS, and the same features that are in use on the source system. The details of creating volumes and datasets, enabling SSH, and copying encryption keys will vary when the destination computer is not a TrueNAS® system.

Encryption Keys

A public encryption key must be copied from Alpha to Beta to allow a secure connection without a password prompt. On Alpha, select `Storage` → `Replication Tasks` → `View Public Key`, producing the window shown in [Figure 7.27](#). Use the mouse to highlight the key data shown in the window, then copy it.

Add Replication



Volume/Dataset:

alphavol/alphadata



Remote ZFS Volume/Dataset:

betavol



Recursively replicate child dataset's snapshots:

☐

Delete stale snapshots on remote system:

☐

Replication Stream Compression:

lz4 (fastest)



Limit (kB/s):

0



Begin:

00:00:00



End:

23:59:00



Enabled:

☒


Setup mode:

Manual



Remote hostname:

10.0.0.118

Remote port:

22

Dedicated User Enabled:

☐


Dedicated User:

Encryption Cipher:

Standard



Remote hostkey:

```
10.0.0.118 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQCA4WnS+kfJa
CDL1SnPWEqHwuVjE0k8pl+kU8JlS8yyfOALP1/aB
c82DdZoNGwtJjn14xTyxA1XJKXio1YYkTnTiLj7M
R+S905HLt+vwSUhkfs3EdD8/oOCFmeiw
/00dzjT9oiCrqqnHiL+dySqBjAE0yfoQyTGfzbsy
FYG9BZ6aLSzA+oEd7i+aJlE++n6oRCENUCopeFGF
m9gADtWwETiHxJkY292JRqhY02k7JrhyzYPSLZvL
Yy3mw0bSG1Xjf8D2xGgxs7qdi ai3r6aKl+TRA4Bi
/d8GxVAKwzJPgv
/K/aWiibmaUcVBavUbM60yaRFg9uuhn43HYMHbJa
4fE/r1
10.0.0.118 ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlz
dHAyNTYAAABBBBANGLOmMyTZl/Fp1aScYX
/8S/b3nvXibX
/levDCDwJecuD1ASWY5Xx+Wp8YkraJzLv9bonf1w
yc2fCL4gzFs0Ag=
10.0.0.118 ssh-ed25519
AAAAC3NzaC1lZDI1NTE5AAAAIOZtUTtc59hv90WH
7nDoD4li3GdRKaZR/V70gzT8t7GE
```

OK

Cancel

SSH Key Scan

7.3.5 Replication Options

Table 7.8 describes the options in the replication task dialog.

Table 7.8: Replication Task Options

Setting	Value	Description
Volume/Dataset	drop-down menu	ZFS volume or dataset on the source computer containing the snapshots to be replicated; the drop-down menu is empty if a snapshot does not already exist
Remote ZFS Volume/Dataset	string	ZFS volume on the remote or destination computer which will store the snapshots; if the destination dataset is not present, it will be created; <code>/mnt/</code> is assumed, do not include it in the path
Recursively replicate child dataset's snapshots	checkbox	when checked, also replicate snapshots of datasets that are children of the main dataset
Delete stale snapshots	checkbox	when checked, delete previous snapshots on the remote or destination computer which are no longer present on the source computer
Replication Stream Compression	drop-down menu	choices are lz4 (fastest), pigz (all rounder), plzip (best compression), or Off (no compression); selecting a compression algorithm can reduce the size of the data being replicated
Limit (kB/s)	integer	limit replication speed to the specified value in kilobytes/second; default of 0 is unlimited
Begin	drop-down menu	replication is not allowed to start before this time; times entered in the Begin and End fields set when replication can occur
End	drop-down menu	replication must start by this time; once started, replication will continue until it is finished
Enabled	checkbox	uncheck to disable the scheduled replication task without deleting it
Setup mode	drop-down menu	Manual or Semi-automatic
Remote hostname	string	IP address or DNS name of remote computer where replication is sent
Remote port	string	the port used by the SSH server on the remote or destination computer
Dedicated User Enabled	checkbox	allow a user account other than root to be used for replication
Dedicated User	drop-down menu	only available if Dedicated User Enabled is checked; select the user account to be used for replication
Encryption Cipher	drop-down menu	Standard, Fast, or Disabled
Remote hostkey	string	use the SSH Key Scan button to retrieve the public host key of the remote or destination computer and populate this field with that key

The replication task runs after a new periodic snapshot is created. The periodic snapshot and any new manual snapshots of the same dataset are replicated onto the destination computer.

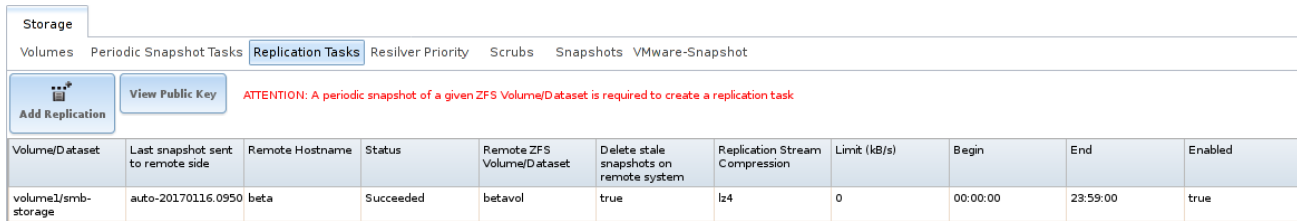
When multiple replications have been created, replication tasks run serially, one after another. Completion time depends on the number and size of snapshots and the bandwidth available between the source and destination computers.

The first time a replication runs, it must duplicate data structures from the source to the destination computer. This can take much longer to complete than subsequent replications, which only send differences in data.

Warning: Snapshots record incremental changes in data. If the receiving system does not have at least one snapshot that can be used as a basis for the incremental changes in the snapshots from the sending system, there is no way to identify only the data that has changed. In this situation, the snapshots in the receiving system target dataset are removed so a complete initial copy of the new replicated data can be created.

Selecting **Storage** → **Replication Tasks** displays [Figure 7.30](#), the list of replication tasks. The Last snapshot sent to remote side column shows the name of the last snapshot that was successfully replicated, and Status shows the current

status of each replication task. The display is updated every five seconds, always showing the latest status.



Volume/Dataset	Last snapshot sent to remote side	Remote Hostname	Status	Remote ZFS Volume/Dataset	Delete stale snapshots on remote system	Replication Stream Compression	Limit (kB/s)	Begin	End	Enabled
volume1/smb-storage	auto-20170116.0950	beta	Succeeded	betavol	true	lz4	0	00:00:00	23:59:00	true

Fig. 7.30: Replication Task List

Note: The encryption key that was copied from the source computer (Alpha) to the destination computer (Beta) is an RSA public key located in the `/data/ssh/replication.pub` file on the source computer. The host public key used to identify the destination computer (Beta) is from the `/etc/ssh/ssh_host_rsa_key.pub` file on the destination computer.

7.3.6 Replication Encryption

The default Encryption Cipher Standard setting provides good security. Fast is less secure than Standard but can give reasonable transfer rates for devices with limited cryptographic speed. For networks where the entire path between source and destination computers is trusted, the Disabled option can be chosen to send replicated data without encryption.

7.3.7 Limiting Replication Times

The Begin and End times in a replication task make it possible to restrict when replication is allowed. These times can be set to only allow replication after business hours, or at other times when disk or network activity will not slow down other operations like snapshots or [Scrubs](#) (page 126). The default settings allow replication to occur at any time.

These times control when replication task are allowed to start, but will not stop a replication task that is already running. Once a replication task has begun, it will run until finished.

7.3.8 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.

Star Replication

In a star topology, a single TrueNAS® computer replicates data to multiple destination computers. This can provide 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.

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.

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.

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 can be 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's 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.

7.3.9 Troubleshooting Replication

Replication depends on SSH, disks, network, compression, and encryption to work. A failure or misconfiguration of any of these can prevent successful replication.

SSH

SSH (page 213) must be able to connect from the source system to the destination system with an encryption key. This can be tested from [Shell](#) (page 257) by making an [SSH](#) (page 213) connection from the source system to the destination system. From the previous example, this is a connection from Alpha to Beta at 10.0.0.118. Start the [Shell](#) (page 257) on the source machine (Alpha), then enter this command:

```
ssh -vv -i /data/ssh/replication 10.0.0.118
```

On the first connection, the system might say

```
No matching host key fingerprint found in DNS.  
Are you sure you want to continue connecting (yes/no)?
```

Verify that this is the correct destination computer from the preceding information on the screen and type `yes`. At this point, an [SSH](#) (page 213) shell connection is open to the destination system, Beta.

If a password is requested, SSH authentication is not working. See [Figure 7.27](#) above. This key value must be present in the `/root/.ssh/authorized_keys` file on Beta, the destination computer. The `/var/log/auth.log` file can show diagnostic errors for login problems on the destination computer also.

Compression

Matching compression and decompression programs must be available on both the source and destination computers. This is not a problem when both computers are running 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.

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 257) and manually send a single snapshot to the destination computer, Beta. The snapshot used in this example is named `auto-20161206.1110-2w`. As before, it is located in the `alphavol/alphadata` dataset. A `@` symbol separates the name of the dataset from the name of the snapshot in the command.

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

If a snapshot of that name already exists on the destination computer, the system will refuse to overwrite it with the new snapshot. The existing snapshot on the destination computer can be deleted by opening a [Shell](#) (page 257) on Beta and running this command:

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

Then send the snapshot manually again. Snapshots on the destination system, Beta, can be listed from the [Shell](#) (page 257) with `zfs list -t snapshot` or by going to `Storage` → `Snapshots`.

Error messages here can indicate any remaining problems.

7.4 Resilver Priority

Resilvering, or the process of copying data to a replacement disk, is best completed as quickly as possible. Increasing the priority of resilvers can help them to complete more quickly. The Resilver Priority tab makes it possible to increase the priority of resilvering at times where the additional I/O or CPU usage will not affect normal usage. Select `Storage` → `Resilver Priority` to display the screen shown in [Figure 7.31](#). [Table 7.9](#) describes the fields on this screen.

Storage

Volumes

Periodic Snapshot Tasks

Replication Tasks

Resilver Priority

Scrubs

Snapshots

VMware-Snapshot

Enabled:

Begin higher priority resilvering at this time:

6:00 PM

End higher priority resilvering at this time:

9:00 AM

Weekday:

Monday

Tuesday

Wednesday

Thursday

Friday

Saturday

Sunday

Save

Fig. 7.31: Resilver Priority

Table 7.9: Resilver Priority Options

Setting	Value	Description
Enabled	checkbox	check to enable higher-priority resilvering
Begin higher priority resilvering at this time	drop-down	start time to begin higher-priority resilvering
End higher priority resilvering at this time	drop-down	end time to begin higher-priority resilvering
Weekday	checkboxes	use higher-priority resilvering on these days of the week

7.5 Scrubs

A scrub is the process of ZFS scanning through the data on a volume. Scrubs help to identify data integrity problems, detect silent data corruptions caused by transient hardware issues, and provide early alerts of impending disk failures. TrueNAS® makes it easy to schedule periodic automatic scrubs.

Each volume should be scrubbed at least once a month. Bit errors in critical data can be detected by ZFS, but only when that data is read. Scheduled scrubs can find bit errors in rarely-read data. The amount of time needed for a scrub is proportional to the quantity of data on the volume. Typical scrubs take several hours or longer.

The scrub process is I/O intensive and can negatively impact performance. Schedule scrubs for evenings or weekends to minimize impact to users. Make certain that scrubs and other disk-intensive activity like [S.M.A.R.T. Tests](#) (page 73) 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 73) of Type Long Self-Test to run once or twice a month.

Scrubs are scheduled and managed with `Storage → Scrubs`.

When a volume is created, a ZFS scrub is automatically scheduled. An entry with the same volume name is added to `Storage` → `Scrubs`. A summary of this entry can be viewed with `Storage` → `Scrubs` → `View Scrubs`. [Figure 7.32](#) displays the default settings for the volume named `volume1`. In this example, the entry has been highlighted and the `Edit` button clicked to display the `Edit` screen. [Table 7.10](#) summarizes the options in this screen.

Table 7.10: ZFS Scrub Options

Setting	Value	Description
Volume	drop-down menu	volume to be scrubbed
Threshold days	integer	prevent scrub from running for this number of days after a scrub has completed, regardless of the calendar schedule; the default is a multiple of 7 to ensure that the scrub always occurs on the same day of the week
Description	string	optional text description of scrub
Minute	slider or minute selections	if the slider is used, a scrub occurs every N minutes; if specific minutes are chosen, a scrub runs only at the selected minute values
Hour	slider or hour selections	if the slider is used, a scrub occurs every N hours; if specific hours are chosen, a scrub runs only at the selected hour values
Day of Month	slider or month selections	if the slider is used, a scrub occurs every N days; if specific days of the month are chosen, a scrub runs only on the selected days of the selected months
Month	checkboxes	a scrub occurs on the selected months
Day of week	checkboxes	a scrub occurs on the selected days; the default is Sunday to least impact users; note that this field and the Day of Month field are ORed together: setting Day of Month to 01,15 and Day of week to Thursday will cause scrubs to run on the 1st and 15th days of the month, but also on any Thursday
Enabled	checkbox	uncheck to disable the scheduled scrub without deleting it

Review the default selections and, if necessary, modify them to meet the needs of the environment. Note that the Threshold field is used to prevent scrubs from running too often, and overrides the schedule chosen in the other fields.

Scheduled crubs can be deleted with the Delete button, but this is not recommended. Scrubs can provide an early indication of disk issues before a disk failure. If a scrub is too intensive for the hardware, consider temporarily unchecking the Enabled button for the scrub until the hardware can be upgraded.

7.6 Snapshots

The Snapshots tab is used to review the listing of available snapshots. An example is shown in [Figure 7.33](#).

Note: If snapshots do not appear, check that the current time configured in [Periodic Snapshot Tasks](#) (page 113) 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 257).

Storage

VolumesPeriodic Snapshot TasksReplication TasksResilver PriorityScrubsSnapshotsVMware-Snapshot















	Volume/Dataset	Snapshot Name	Used	Refer	Available Actions
 ...	No filter applied				
	volume1	auto-20171018.0840-2w	0	88.0 KiB	 
	volume1	auto-20171018.0850-2w	0	88.0 KiB	 
	volume1	auto-20171018.0900-2w	0	88.0 KiB	 
	volume1	auto-20171018.0910-2w	0	88.0 KiB	  

Fig. 7.33: Viewing Available Snapshots

The listing includes the name of the volume or dataset, the name of each snapshot, and the amount of used and referenced data.

Used is the amount of space consumed by this dataset and all of its descendants. This value is checked against the dataset's quota and reservation. The space used does not include the dataset's reservation, but does take into account the reservations of any descendent datasets. The amount of space that a dataset consumes from its parent, as well as the amount of space that are freed if this dataset is recursively destroyed, is the greater of its space used and its reservation. When a snapshot is created, the space is initially shared between the snapshot and the filesystem, and possibly with previous snapshots. As the filesystem changes, space that was previously shared becomes unique to the snapshot, and is counted in the snapshot's space used. Additionally, deleting snapshots can increase the amount of space unique to (and used by) other snapshots. The amount of space used, available, or referenced does not take into account pending changes. While pending changes are generally accounted for within a few seconds, disk changes do not necessarily guarantee that the space usage information is updated immediately.

Tip: Space used by individual snapshots can be seen by running `zfs list -t snapshot` from [Shell](#) (page 257).

Refer indicates the amount of data accessible by this dataset, which may or may not be shared with other datasets in the pool. When a snapshot or clone is created, it initially references the same amount of space as the file system or snapshot it was created from, since its contents are identical.

Snapshots have icons on the right side for several actions.

Clone Snapshot prompts for the name of the clone to create. A clone is a writable copy of the snapshot. Since a clone is actually a dataset which can be mounted, it appears in the Volumes tab rather than the Snapshots tab. By default, `-clone` is added to the name of a snapshot when a clone is created.

Destroy Snapshot a pop-up message asks for confirmation. Child clones must be destroyed before their parent snapshot can be destroyed. While creating a snapshot is instantaneous, deleting a snapshot can be I/O intensive and can take a long time, especially when deduplication is enabled. In order to delete a block in a snapshot, ZFS has to walk all the allocated blocks to see if that block is used anywhere else; if it is not, it can be freed.

The most recent snapshot also has a Rollback Snapshot icon. Clicking the icon asks for confirmation before rolling back to this snapshot state. Confirming by clicking Yes causes any files that have changed since the snapshot was taken to be reverted back to their state at the time of the snapshot.

Note: Rollback is a potentially dangerous operation and causes any configured replication tasks to fail as the replication system uses the existing snapshot when doing an incremental backup. To restore the data within a snapshot, the recommended steps are:

1. Clone the desired snapshot.
2. Share the clone with the share type or service running on the TrueNAS® system.
3. After users have recovered the needed data, destroy the clone in the Active Volumes tab.

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

A range of snapshots can be selected with the mouse. Click on the checkbox in the left column of the first snapshot, then press and hold `Shift` and click on the checkbox for the end snapshot. This can be used to select a range of obsolete snapshots to be deleted with the Destroy icon at the bottom. Be cautious and careful when deleting ranges of snapshots.

Periodic snapshots can be configured to appear as shadow copies in newer versions of Windows Explorer, as described in [Configuring Shadow Copies](#) (page 168). Users can access the files in the shadow copy using Explorer without requiring any interaction with the TrueNAS® graphical administrative interface.

The ZFS Snapshots screen allows the creation of filters to view snapshots by selected criteria. To create a filter, click the Define filter icon (near the text No filter applied). When creating a filter:

- Select the column or leave the default of Any Column.
- Select the condition. Possible conditions are: contains (default), is, starts with, ends with, does not contain, is not, does not start with, does not end with, and is empty.
- Enter a value that meets your view criteria.

- Click the Filter button to save the filter and exit the define filter screen. Alternately, click the + button to add another filter.

When creating multiple filters, select the filter to use before leaving the define filter screen. After a filter is selected, the No filter applied text changes to Clear filter. Clicking Clear filter produces a pop-up message indicates that this removes the filter and all available snapshots are listed.

7.7 VMware-Snapshot

Storage → VMware-Snapshot allows you to coordinate ZFS snapshots when using TrueNAS® as a VMware datastore. Once this type of snapshot is created, TrueNAS® will automatically snapshot any running VMware virtual machines before taking a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore. The temporary VMware snapshots are then deleted on the VMware side but still exist in the ZFS snapshot and can be used as stable resurrection points in that snapshot. These coordinated snapshots will be listed in [Snapshots](#) (page 129).

Figure 7.34 shows the menu for adding a VMware snapshot and Table 7.11 summarizes the available options.

Fig. 7.34: Adding a VMware Snapshot

Table 7.11: VMware Snapshot Options

Setting	Value	Description
Hostname	string	IP address or hostname of VMware host; when clustering, this is the vCenter server for the cluster
Username	string	user on VMware host with enough permission to snapshot virtual machines
Password	string	password associated with Username
ZFS Filesystem	drop-down menu	the filesystem to snapshot
Datastore	drop-down menu	after entering the Hostname, Username, and Password, click Fetch Datastores to populate the menu and select the datastore with which to synchronize

DIRECTORY SERVICES

TrueNAS® supports integration with these directory services:

- [Active Directory](#) (page 132) (for Windows 2000 and higher networks)
- [LDAP](#) (page 137)
- [NIS](#) (page 140)

It also supports [Kerberos Realms](#) (page 141), [Kerberos Keytabs](#) (page 141), and the ability to add additional parameters to [Kerberos Settings](#) (page 142).

This section summarizes each of these services and their available configurations within the TrueNAS® GUI.

8.1 Active Directory

Active Directory (AD) is a service for sharing resources in a Windows network. AD can be configured on a Windows server that is running Windows Server 2000 or higher or on a Unix-like operating system that is running [Samba version 4](#) (https://wiki.samba.org/index.php/Samba4/HOWTO#Provisioning_The_Samba_Active_Directory). Since AD provides authentication and authorization services for the users in a network, it is not necessary to recreate these user accounts on the TrueNAS® system. Instead, configure the Active Directory service so that it can import the account information and imported users can be authorized to access the SMB shares on the 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® 11.1 before attempting Active Directory integration.

Before configuring the Active Directory service, ensure name resolution is properly configured by `ping`ing the domain name of the Active Directory domain controller from [Shell](#) (page 257) on the TrueNAS® system. If the `ping` fails, check the DNS server and default gateway settings in `Network → Global Configuration` on the TrueNAS® system.

Next, add a DNS record for the TrueNAS® system on the Windows server and verify that the hostname of the TrueNAS® system can be pinged from the domain controller.

Active Directory relies on Kerberos, which is a time sensitive protocol. The time on both the TrueNAS® system and the Active Directory Domain Controller cannot be out of sync by more than a few minutes. The best way to ensure that the same time is running on both systems is to configure both systems to:

- use the same NTP server (set in `System → NTP Servers` on the TrueNAS® system)
- have the same timezone
- be set to either localtime or universal time at the BIOS level

[Figure 8.1](#) shows the screen that appears when `Directory Service → Active Directory` is chosen. [Table 8.1](#) describes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box `Show advanced fields` by default in `System → Advanced`.

Directory Service
Active Directory
LDAP
NIS
Kerberos Realms
Kerberos Keytabs
Kerberos Settings

Domain Name (DNS/Realm-Name):
Domain Account Name:
Domain Account Password:
AD check connectivity frequency (seconds):
How many recovery attempts:
Enable Monitoring:
Enable:

Save
Advanced Mode
Rebuild Directory Service Cache

Fig. 8.1: Configuring Active Directory

Table 8.1: Active Directory Configuration Options

Setting	Value	Advanced Mode	Description
Domain Name (DNS/Realm-Name)	string		name of Active Directory domain (example.com) or child domain (sales.example.com); this setting is mandatory and the GUI will refuse to save the settings if the domain controller for the specified domain cannot be found
Domain Account Name	string		name of the Active Directory administrator account; this setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this account name
Domain Account Password	string		password for the Active Directory administrator account; this setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this password
AD check connectivity frequency (seconds)	integer		how often to verify that Active Directory services are active
How many recovery attempts	integer		number of times to attempt reconnecting to the Active Directory server; tries forever when set to 0
Enable Monitoring	checkbox		restart Active Directory automatically if the service is disconnected
Encryption Mode	drop-down menu	✓	choices are Off, SSL, or TLS
Certificate	drop-down menu	✓	select the certificate of the Active Directory server if SSL connections are used; if a certificate does not exist yet, create a CA (page 45), then create a certificate on the Active Directory server and import it to the TrueNAS® system with Certificates (page 48)
Verbose logging	checkbox	✓	when checked, logs attempts to join the domain to <code>/var/log/messages</code>
Continued on next page			

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
UNIX extensions	checkbox	✓	only check this box if the AD server has been explicitly configured to map permissions for UNIX users; checking this box provides persistent UIDs and GIDs, otherwise, users/groups are mapped to the UID/GUID range configured in Samba
Allow Trusted Domains	checkbox	✓	should only be enabled if network has active do-main/forest trusts (https://technet.microsoft.com/en-us/library/cc757352(WS.10).aspx) and you need to manage files on multiple domains; use with caution as it will generate more winbindd traffic, slowing down the ability to filter through user/group information
Use Default Domain	checkbox	✓	when unchecked, the domain name is prepended to the username; if Allow Trusted Domains is checked and multiple domains use the same usernames, uncheck this box to prevent name collisions
Allow DNS updates	checkbox	✓	when unchecked, disables Samba from doing DNS updates when joining a domain
Disable Active Directory user/group cache	checkbox	✓	when checked, disables caching AD users and groups; useful if you cannot bind to a domain with a large number of users or groups
User Base	string	✓	distinguished name (DN) of the user container in Active Directory
Group Base	string	✓	distinguished name (DN) of the group container in Active Directory
Site Name	string	✓	the relative distinguished name of the site object in Active Directory
Domain Controller	string	✓	will automatically be added to the SRV record for the domain and, when multiple controllers are specified, TrueNAS® selects the closest DC which responds; use the short form of the FQDN (Example: sampleserver)
Global Catalog Server	string	✓	if the hostname of the global catalog server to use is specified, make sure it is resolvable
Kerberos Realm	drop-down menu	✓	select the realm created using the instructions in Kerberos Realms (page 141)
Kerberos Principal	drop-down menu	✓	browse to the location of the keytab created using the instructions in Kerberos Keytabs (page 141)
AD timeout	integer	✓	in seconds, increase if the AD service does not start after connecting to the domain
DNS timeout	integer	✓	in seconds, increase if AD DNS queries timeout
Idmap backend	drop-down menu and Edit	✓	select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs; see Table 8.2 for a summary of the available backends; click the Edit link to configure that backend's editable options
Windbind NSS Info	drop-down menu	✓	defines the schema to use when querying AD for user/group info; rfc2307 uses the RFC2307 schema support included in Windows 2003 R2, sfu20 is for Services For Unix 3.0 or 3.5, and sfu is for Services For Unix 2.0
SASL wrapping	drop-down menu	✓	defines how LDAP traffic is transmitted; choices are plain (plain text), sign (signed only), or seal (signed and encrypted); Windows 2000 SP3 and higher can be configured to enforce signed LDAP connections
Enable	checkbox		Enable the Active Directory service

Continued on next page

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
NetBIOS Name (This Node)	string	✓	limited to 15 characters; automatically populated with the system's original hostname; it must be different from the Workgroup name
NetBIOS Name (Node B)	string	✓	limited to 15 characters; when using Failover (page 55), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	✓	limited to 15 characters; when using Failover (page 55), this is the NetBIOS name that resolves to either node

[Table 8.2](#) summarizes the backends which are available in the Idmap backend drop-down menu. Each backend has its own [man page](http://samba.org.ru/samba/docs/man/manpages/) (<http://samba.org.ru/samba/docs/man/manpages/>) which gives implementation details. Since selecting the wrong backend will break Active Directory integration, a pop-up menu will appear whenever changes are made to this setting.

Table 8.2: ID Mapping Backends

Value	Description
ad	AD server uses RFC2307 or Services For Unix schema extensions; mappings must be provided in advance by adding the uidNumber attributes for users and gidNumber attributes for groups in the AD
autorid	similar to rid, but automatically configures the range to be used for each domain, so there is no need to specify a specific range for each domain in the forest; the only needed configuration is the range of UID/GIDs to use for user/group mappings and an optional size for the ranges
fruit	generate IDs the way Apple Mac OS X does, so UID and GID can be identical on all TrueNAS® servers on the network; for use in LDAP (page 137) environments where Apple's Open Directory is the authoritative LDAP server
ldap	stores and retrieves mapping tables in an LDAP directory service; default for LDAP directory service
nss	provides a simple means of ensuring that the SID for a Unix user is reported as the one assigned to the corresponding domain user
rfc2307	an AD server is required to provide the mapping between the name and SID and an LDAP server is required to provide the mapping between the name and the UID/GID
rid	default for AD; requires an explicit idmap configuration for each domain, using disjoint ranges where a writeable default idmap range should be defined, using a backend like tdb or ldap
script	stores mapping tables for clustered environments in the winbind_cache tdb
tdb	default backend used by winbindd for storing mapping tables
tdb2	substitute for tdb used by winbindd in clustered environments

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

Note: Active Directory places restrictions on which characters are allowed in Domain and NetBIOS names, a limits the length of those names to 15 characters. If there are problems connecting to the realm, [verify](https://support.microsoft.com/en-us/kb/909264) (<https://support.microsoft.com/en-us/kb/909264>) that your settings do not include any disallowed characters. Also, the Administrator account password cannot contain the \$ character. If a \$ exists in the domain administrator's password, `kinit` will report a "Password Incorrect" error and `ldap_bind` will report an "Invalid credentials (49)" error.

It can take a few minutes after configuring the Active Directory service for the AD information to be populated to the TrueNAS® system. Once populated, the AD users and groups will be available in the drop-down menus of the Permissions screen of a volume/dataset. For performance reasons, every available user may not show in the listing. However, it will

autocomplete all applicable users when typing in a username.

The Active Directory users and groups that have been imported to the TrueNAS® system can be shown by using these commands from the TrueNAS® [Shell](#) (page 257). To view users:

```
wbinfo -u
```

To view groups:

```
wbinfo -g
```

In addition, `wbinfo -t` will test the connection and, if successful, will show a message similar to:

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

To manually check that a specified user can authenticate:

```
net ads join -S dcname -U username
```

If no users or groups are listed in the output, these commands can provide more troubleshooting information:

```
getent passwd
```

```
getent group
```

If the `wbinfo` commands display the network users, but they do not show up in the drop-down menu of a Permissions screen, it may be because it is taking longer than the default ten seconds for the TrueNAS® system to join Active Directory. Try bumping up the value of AD timeout to 60 seconds.

Tip: To change a certificate, set the Encryption Mode to Off, then disable AD by unchecking Enable. Click the Save button. Select the new Certificate, set the Encryption Mode as desired, set the Enable checkbox to re-enable AD, and click the Save button to restart AD.

8.1.1 Troubleshooting Tips

When running AD in a 2003/2008 mixed domain, [refer to](https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) (https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) for instructions on how to prevent the secure channel key from becoming corrupt.

Active Directory uses DNS to determine the location of the domain controllers and global catalog servers in the network. Use the `host -t srv _ldap._tcp.domainname.com` command to determine the network's SRV records and, if necessary, change the weight and/or priority of the SRV record to reflect the fastest server. More information about SRV records can be found in the Technet article [How DNS Support for Active Directory Works](https://technet.microsoft.com/en-us/library/cc759550(WS.10).aspx) (https://technet.microsoft.com/en-us/library/cc759550(WS.10).aspx).

The realm that is used depends upon the priority in the SRV DNS record, meaning that DNS can override your Active Directory settings. When unable to connect to the correct realm, check the SRV records on the DNS server. [This article](http://www.informit.com/guides/content.aspx?g=security&seqNum=37&rll=1) (http://www.informit.com/guides/content.aspx?g=security&seqNum=37&rll=1) describes how to configure KDC discovery over DNS and provides some examples of records with differing priorities.

If the cache becomes out of sync due to an AD server being taken off and back online, resync the cache using `Directory Service → Active Directory → Rebuild Directory Service Cache`.

An expired password for the administrator account will cause kinit to fail, so ensure that the password is still valid. Also, double-check that the password on the AD account being used does not include any spaces or special symbols, and is not unusually long.

If the Windows server version is lower than 2008 R2, try creating a Computer entry on the Windows server's OU. When creating this entry, enter the TrueNAS® hostname in the name field. Make sure that it is under 15 characters and that it is the same name as the one set in the Hostname field in `Network → Global Configuration` and the NetBIOS Name

in Directory Service → Active Directory settings. Make sure the hostname of the domain controller is set in the Domain Controller field of Directory Service → Active Directory.

8.1.2 If the System Will not Join the Domain

If the system will not join the Active Directory domain, run these commands in the order listed. If any of the commands fail or result in a traceback, create a bug report at <https://redmine.ixsystems.com/projects/freenas/issues> that includes the commands in the order in which they were run and the exact wording of the error message or traceback.

Start with these commands, where the `echo` commands should return a value of 0 and the `klist` command should show a Kerberos ticket:

```
sqlite3 /data/freenas-v1.db "update directoryservice_activedirectory set ad_enable=1;"
echo $?
service ix-kerberos start
service ix-nsswitch start
service ix-kinit start
service ix-kinit status
echo $?
klist
```

Next, only run these two commands if the Unix extensions box is checked in Advanced Mode and a keytab has been uploaded using [Kerberos Keytabs](#) (page 141):

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

Finally, run these commands. Again, the `echo` command should return a 0:

```
python /usr/local/www/freenasUI/middleware/notifier.py start cifs
service ix-activedirectory start
service ix-activedirectory status
echo $?
python /usr/local/www/freenasUI/middleware/notifier.py restart cifs
service ix-pam start
service ix-cache start &
```

8.2 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 Microsoft Server (2000 and newer), Mac OS X Server, Novell eDirectory, and OpenLDAP running on a BSD or Linux system. If an LDAP server is running on your network, configure the 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](http://download.gna.org/smbldap-tools/) (<http://download.gna.org/smbldap-tools/>) and instructions for using it can be found at [The Linux Samba-OpenLDAP Howto](http://download.gna.org/smbldap-tools/docs/samba-ldap-howto/#htoc29) (<http://download.gna.org/smbldap-tools/docs/samba-ldap-howto/#htoc29>). In addition, the LDAP server must support SSL/TLS and the certificate for the LDAP server CA must be imported with System → Certificates → Import Certificate. Note that non-CA certificates are not supported at this time.

Tip: Apple's [Open Directory](https://manuals.info.apple.com/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) (https://manuals.info.apple.com/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) is an LDAP-compatible directory service into which TrueNAS® can be integrated. See [FreeNAS with Open Directory in](#)

Figure 8.2 shows the LDAP Configuration screen that is seen after clicking `Directory Service` → `LDAP`.

Directory Service

Active DirectoryLDAPNISKerberos RealmsKerberos KeytabsKerberos Settings

Hostname:

i

Base DN:

i

Bind DN:

i

Bind password:

i

Enable:

☐

Save

Advanced Mode

Rebuild Directory Service Cache

Fig. 8.2: Configuring LDAP

Table 8.3 summarizes the available configuration options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in `System` → `Advanced`.

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

Table 8.3: LDAP Configuration Options

Setting	Value	Advanced Mode	Description
Hostname	string		hostname or IP address of LDAP server
Base DN	string		top level of the LDAP directory tree to be used when searching for resources (e.g. dc=test,dc=org)
Bind DN	string		name of administrative account on LDAP server (e.g. cn=Manager,dc=test,dc=org)
Bind password	string		password for Root bind DN
Allow Anonymous Binding	checkbox	✓	instructs LDAP server to not provide authentication and to allow read and write access to any client
User Suffix	string	✓	optional; can be added to name when user account added to LDAP directory (e.g. dept. or company name)
Group Suffix	string	✓	optional; can be added to name when group added to LDAP directory (e.g. dept. or company name)

Continued on next page

Table 8.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Password Suffix	string	✓	optional; can be added to password when password added to LDAP directory
Machine Suffix	string	✓	optional; can be added to name when system added to LDAP directory (e.g. server, accounting)
SUDO Suffix	string	✓	use if LDAP-based users need superuser access
Kerberos Realm	drop-down menu	✓	select the realm created using the instructions in Kerberos Realms (page 141)
Kerberos Principal	drop-down menu	✓	browse to the location of the principal in the keytab created as described in Kerberos Keytabs (page 141)
Encryption Mode	drop-down menu	✓	choices are Off, SSL, or TLS; note that either SSL or TLS and a Certificate must be selected in order for authentication to work
Certificate	drop-down menu	✓	select the certificate of the LDAP CA (required if authentication is used); the certificate for the LDAP server CA must first be imported with <code>System → Certificates → Import Certificate</code>
LDAP timeout	integer	✓	increase this value (in seconds) if obtaining a Kerberos ticket times out
DNS timeout	integer	✓	increase this value (in seconds) if DNS queries timeout
Idmap backend	drop-down menu and Edit	✓	select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs; see Table 8.2 for a summary of the available backends; click the Edit link to configure the backend's editable options
Samba Schema	checkbox	✓	only check this box if you need LDAP authentication for SMB shares and have already configured the LDAP server with Samba attributes
Auxiliary Parameters	string	✓	additional options for <code>sssd.conf(5)</code> (https://jhrozek.fedorapeople.org/sss/1.11.6/man/sss.conf.5.html)
Schema	drop-down menu	✓	if Samba Schema is checked, select the schema to use; choices are <code>rfc2307</code> and <code>rfc2307bis</code>
Enable	checkbox		uncheck to disable the configuration without deleting it
NetBIOS Name (This Node)	string	✓	limited to 15 characters; automatically populated with the system's original hostname; it must be different from the Workgroup name
NetBIOS Name (Node B)	string	✓	limited to 15 characters; when using Failover (page 55), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	✓	limited to 15 characters; when using Failover (page 55), this is the NetBIOS name that resolves to either node

Click the Rebuild Directory Service Cache button after adding a user to LDAP who needs immediate access to TrueNAS®. Otherwise this occurs automatically once a day as a cron job.

Note: TrueNAS® automatically appends the root DN. This means that the scope and root DN should not be included when configuring the user, group, password, and machine suffixes.

LDAP users and groups appear in the drop-down menus of the Permissions screen of a volume/dataset after configuring the LDAP service. Type `getent passwd` from [Shell](#) (page 257) to verify that the users have been imported. Type `getent group` to verify that the groups have been imported.

If the users and groups are not listed, refer to [Common errors encountered when using OpenLDAP Software](#) (<http://www.openldap.org/doc/admin24/appendix-common-errors.html>) for common errors and how to fix them. When troubleshooting LDAP, open [Shell](#) (page 257) and look for error messages in `/var/log/auth.log`.

8.3 NIS

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

Note: In Windows Server 2016, Microsoft removed the Identity Management for Unix (IDMU) and NIS Server Role. See [Clarification regarding the status of Identity Management for Unix \(IDMU\) & NIS Server Role in Windows Server 2016 Technical Preview and beyond](https://blogs.technet.microsoft.com/activedirectoryua/2016/02/09/identity-management-for-unix-idmu-is-deprecated-in-windows-server/) (<https://blogs.technet.microsoft.com/activedirectoryua/2016/02/09/identity-management-for-unix-idmu-is-deprecated-in-windows-server/>).

Figure 8.3 shows the configuration screen which opens when you click `Directory Service` → `NIS`. Table 8.4 summarizes the configuration options.

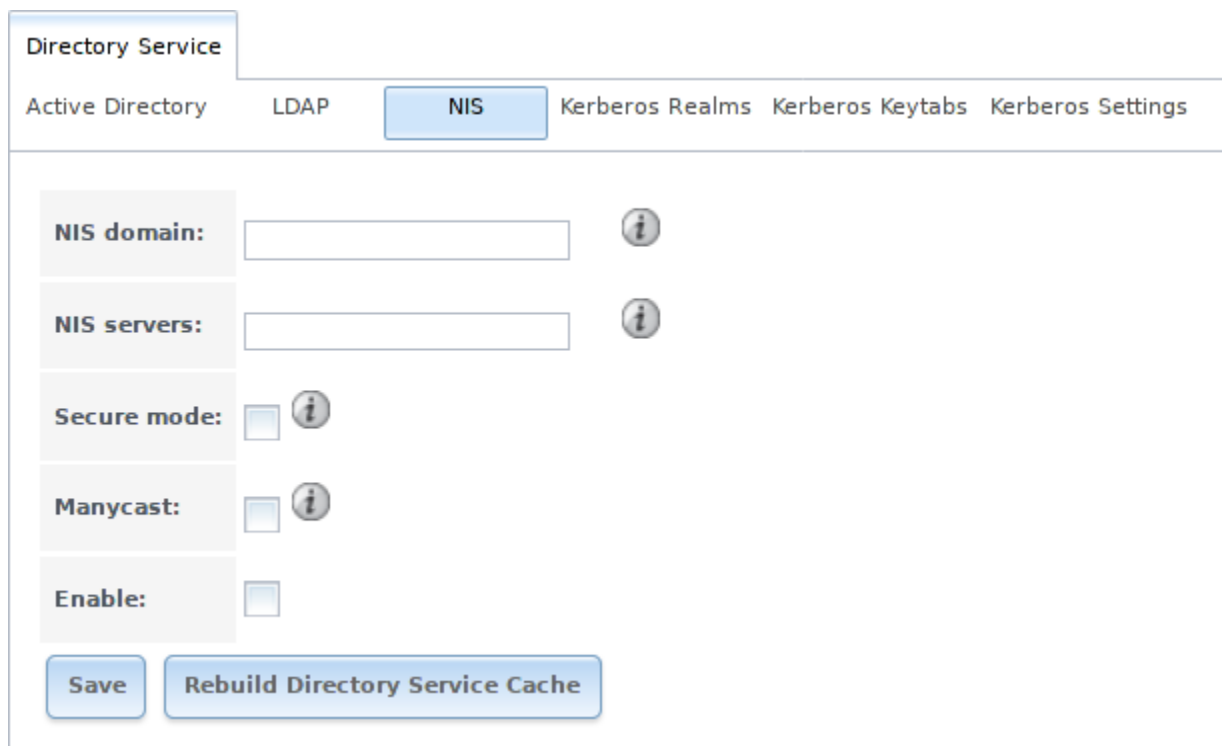


Fig. 8.3: NIS Configuration

Table 8.4: NIS Configuration Options

Setting	Value	Description
NIS domain	string	name of NIS domain
NIS servers	string	comma delimited list of hostnames or IP addresses
Secure mode	checkbox	if checked, <code>ypbind(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=ypbind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, <code>ypbind</code> will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet
Enable	checkbox	uncheck to disable the configuration without deleting it

Click the `Rebuild Directory Service Cache` button after adding a user to NIS who needs immediate access to TrueNAS®.

Otherwise this occurs automatically once a day as a cron job.

8.4 Kerberos Realms

A default Kerberos realm is created for the local system in TrueNAS®. `Directory Service → Kerberos Realms` can be used to view and add Kerberos realms. If the network contains a KDC, click the `Add kerberos realm` button to add the Kerberos realm. This configuration screen is shown in [Figure 8.4](#).

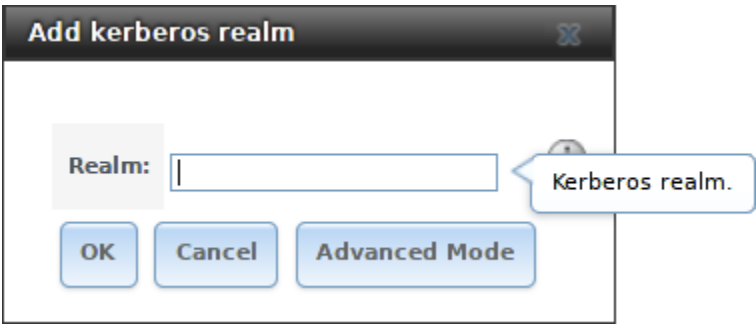


Fig. 8.4: Adding a Kerberos Realm

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

Table 8.5: Kerberos Realm Options

Setting	Value	Advanced Mode	Description
Realm	string		mandatory; name of the realm
KDC	string	✓	name of the Key Distribution Center
Admin Server	string	✓	server where all changes to the database are performed
Password Server	string	✓	server where all password changes are performed

8.5 Kerberos Keytabs

Kerberos keytabs are used to do Active Directory or LDAP joins without a password. This means that the password for the Active Directory or LDAP administrator account does not need to be saved into the 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 these commands:

```
ktpass.exe -out hostname.keytab host/ hostname@DOMAINNAME -ptype KRB5_NT_PRINCIPAL -mapuser_
↪DOMAIN\username -pass userpass

setspn -A host/ hostname@DOMAINNAME DOMAIN\username
```

where:

- `hostname` is the fully qualified hostname of the domain controller
- `DOMAINNAME` is the domain name in all caps

- DOMAIN is the pre-Windows 2000 short name for the domain
- username is the privileged account name
- userpass is the password associated with username

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

After the keytab is generated, use `Directory Service → Kerberos Keytabs → Add kerberos keytab` to add it to the TrueNAS® system.

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

To instruct LDAP to use a principal from the keytab, select the principal from the drop-down Kerberos Principal menu in `Directory Service → LDAP`.

8.6 Kerberos Settings

To configure additional Kerberos parameters, use `Directory Service → Kerberos Settings`. Figure 8.5 shows the fields available:

- Appdefaults auxiliary parameters: contains settings used by some Kerberos applications. The available settings and their syntax are listed in the [\[appdefaults\] section of krb.conf\(5\)](http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#appdefaults) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#appdefaults).
- Libdefaults auxiliary parameters: contains settings used by the Kerberos library. The available settings and their syntax are listed in the [\[libdefaults\] section of krb.conf\(5\)](http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#libdefaults) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#libdefaults).

Fig. 8.5: Additional Kerberos Settings

SHARING

Shares are created to make part or all of a volume accessible to other computers on the network. The type of share to create depends on factors like which operating systems are being used by computers on the network, security requirements, and expectations for network transfer speeds.

TrueNAS® provides a [Wizard](#) (page 250) for creating shares. The [Wizard](#) (page 250) automatically creates the correct type of dataset and permissions for the type of share, sets the default permissions for the share type, and starts the service needed by the share. It is recommended to use the Wizard to create shares, fine-tune the share settings using the instructions in the rest of this chapter if needed, then fine-tune the default permissions from the client operating system to meet the requirements of the network.

Note: Shares are created to provide and control access to an area of storage. Before creating shares, it is recommended to make a list of the users that need access to storage data, which operating systems these users are using, whether all users should have the same permissions to the stored data, and whether these users should authenticate before accessing the data. This information can help determine which type of shares are needed, whether multiple datasets are needed to divide the storage into areas with different access and permissions, and how complex it will be to set up those permission requirements. Note that shares are used to provide access to data. When a share is deleted, it removes access to data but does not delete the data itself.

These types of shares and services are available:

- [AFP](#) (page 144): Apple File Protocol shares are often used when the client computers all run Mac OS X. Apple has slowly shifted to preferring [SMB](#) (page 159) for modern networks, although Time Machine still requires AFP.
- [Unix \(NFS\)](#) (page 151): Network File System shares are accessible from Mac OS X, Linux, BSD, and the professional and enterprise versions (but not the home editions) of Windows. This can be a good choice when the client computers do not all run the same operating system but NFS client software is available for all of them.
- [WebDAV](#) (page 158): WebDAV shares are accessible using an authenticated web browser (read-only) or [WebDAV client](#) (<https://en.wikipedia.org/wiki/WebDAV#Clients>) running on any operating system.
- [SMB](#) (page 159): Server Message Block shares, also known as Common Internet File System (CIFS) shares, are accessible by Windows, Mac OS X, Linux, and BSD computers. Access is slower than an NFS share due to the single-threaded design of Samba. SMB provides more configuration options than NFS and is a good choice on a network for Windows systems. However, it is a poor choice if the CPU on the TrueNAS® system is limited; if the CPU is maxed out, upgrade the CPU or consider another type of share.
- [Block \(iSCSI\)](#) (page 169): 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 194) service instead of a share and using a cross-platform FTP file manager application such as [Filezilla](#) (<https://filezilla-project.org/>). Secure FTP can be configured if the data needs to be encrypted.

When data security is a concern and the network users are familiar with SSH command line utilities or [WinSCP](#) (<http://winscp.net/eng/index.php>), consider using the [SSH](#) (page 213) service instead of a share. It is slower than unencrypted FTP due to the encryption overhead, but the data passing through the network is encrypted.

Note: It is generally a mistake to share a volume or dataset with more than one share type or access method. Different types of shares and services use different file locking methods. For example, if the same volume is configured to use both NFS and FTP, NFS will lock a file for editing by an NFS user, but an FTP user can simultaneously edit or delete that file. This results in lost edits and confused users. Another example: if a volume is configured for both AFP and SMB, Windows users can be confused by the “extra” filenames used by Mac files and delete them. This corrupts the files on the AFP share. Pick the one type of share or service that makes the most sense for the types of clients accessing that volume, and use that single type of share or service. To support multiple types of shares, divide the volume into datasets and use one dataset per share.

This section demonstrates configuration and fine-tuning of AFP, NFS, SMB, WebDAV, and iSCSI shares. FTP and SSH configurations are described in [Services](#) (page 187).

9.1 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 created using the [Wizard](#) (page 250). It then provides configuration examples for using the [Wizard](#) (page 250) to create a guest share, configuring Time Machine to back up to a dataset on the TrueNAS® system, and for connecting to the share from a Mac OS X client.

To view the AFP share created by the Wizard, click [Sharing](#) → [Apple \(AFP\)](#) and highlight the name of the share. Click its [Edit](#) button to see the configuration options shown in [Figure 9.1](#). The values showing for these options will vary, depending upon the information given when the share was created.

The screenshot shows the TrueNAS web interface. At the top, there's a 'Sharing' tab. Below it, there are tabs for 'Apple (AFP)', 'UNIX (NFS)', 'WebDAV', 'Windows (SMB)', and 'Block (iSCSI)'. The 'Apple (AFP)' tab is selected. Below the tabs is a button labeled 'Add Apple (AFP) Share'. Below that is a table with columns 'Path', 'Name', and 'Share Comment'. The table is empty, with the text 'No entry has been found' below it. A modal dialog box titled 'Add Apple (AFP) Share' is open in the foreground. The dialog has the following fields and controls:

- Path:** A text field containing '/mnt/volume1/afp1' and a 'Browse' button.
- Name:** A text field containing 'afp1' and an information icon.
- Time Machine:** A checkbox that is unchecked and an information icon.
- Auxiliary Parameters:** A large text area and an information icon.
- At the bottom are three buttons: 'OK', 'Cancel', and 'Advanced Mode'.

Fig. 9.1: Creating an AFP Share

Note: [Table 9.1](#) summarizes the options available to fine-tune an AFP share. These options should usually be left at the default settings. Changing them might cause unexpected behavior. Most settings are only available with Advanced Mode.

Do not change an advanced option without fully understanding the function of that option. Refer to [Setting up Netatalk](http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html) (<http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html>) for a more detailed explanation of these options.

Table 9.1: AFP Share Configuration Options

Setting	Value	Advanced Mode	Description
Path	browse button		browse to the volume/dataset to share; do not nest additional volumes, datasets, or symbolic links beneath this path because Netatalk does not fully support that
Name	string		volume name which appears in the Mac computer's connect to server dialog; limited to 27 characters and cannot contain a period
Share Comment	string	✓	optional comment
Allow List	string	✓	comma-delimited list of allowed users and/or groups where groupname begins with a @; note that adding an entry will deny any user/group that is not specified
Deny List	string	✓	comma-delimited list of denied users and/or groups where groupname begins with a @; note that adding an entry will allow all users/groups that are not specified
Read-only Access	string	✓	comma-delimited list of users and/or groups who only have read access where groupname begins with a @
Read-write Access	string	✓	comma-delimited list of users and/or groups who have read and write access where groupname begins with a @
Time Machine	checkbox		when checked, TrueNAS® advertises itself as a Time Machine disk so it can be found by Macs; due to a limitation in how the Mac deals with low-diskspace issues when multiple Macs share the same volume, checking Time Machine on multiple shares could result in intermittent failed backups
Time Machine Quota	checkbox		only appears when Time Machine is checked; when checked, each time machine backup on the share has its own quota
Zero Device Numbers	checkbox	✓	enable when the device number is not constant across a reboot
No Stat	checkbox	✓	if checked, AFP does not stat the volume path when enumerating the volumes list; useful for automounting or volumes created by a preexec script
AFP3 UNIX Privs	checkbox	✓	enable Unix privileges supported by OSX 10.5 and higher; do not enable this if the network contains Mac OS X 10.4 clients or lower as they do not support this feature
Default file permission	checkboxes	✓	only works with Unix ACLs; new files created on the share are set with the selected permissions
Default directory permission	checkboxes	✓	only works with Unix ACLs; new directories created on the share are set with the selected permissions
Default umask	integer	✓	umask used for newly created files, default is 000 (anyone can read, write, and execute)
Hosts Allow	string	✓	comma-, space-, or tab-delimited list of allowed hostnames or IP addresses
Hosts Deny	string	✓	comma-, space-, or tab-delimited list of denied hostnames or IP addresses
Auxiliary Parameters	string		additional afp.conf (http://netatalk.sourceforge.net/3.1/htmldocs/afp.conf.5.html) parameters not covered by other option fields

9.1.1 Creating AFP Guest Shares

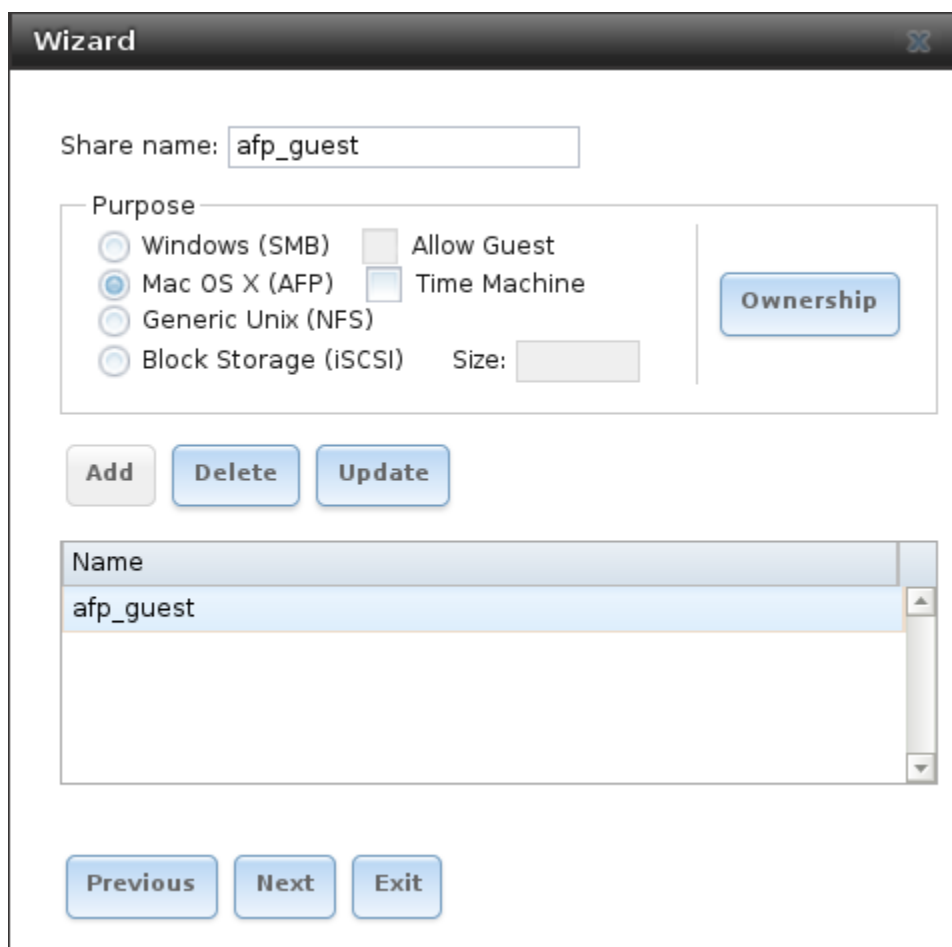
AFP supports guest logins, meaning that Mac OS X users can access the AFP share without requiring their user accounts to first be created on or imported into the TrueNAS® system.

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 make sure that the Guest Access box is checked.

To create the AFP guest share, click *Wizard*, then click the *Next* button twice to display the screen shown in [Figure 9.2](#). Complete these fields in this screen:

1. Share name: enter a name for the share that is identifiable but less than 27 characters long. This name cannot contain a period. In this example, the share is named `afp_guest`.
2. Click the button for Mac OS X (AFP).
3. Click the *Ownership* button. Click the drop-down *User* menu and select *nobody*. Click the *Return* button to return to the previous screen.
4. Click the *Add* button. The share is not created until the button is clicked. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.



The screenshot shows the 'Wizard' window for creating a share. At the top, the title bar says 'Wizard'. Below it, the 'Share name' field contains 'afp_guest'. Under the 'Purpose' section, there are four radio buttons: 'Windows (SMB)', 'Mac OS X (AFP)' (which is selected), 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' and 'Time Machine'. Below these is a 'Size' field. To the right of the 'Purpose' section is a button labeled 'Ownership'. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. Below these is a list box labeled 'Name' containing the entry 'afp_guest'. At the bottom of the window are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 9.2: Creating a Guest AFP Share

Click the Next button twice, then the Confirm button to create the share. The Wizard automatically creates a dataset for the share that contains the correct default permissions and starts the AFP service so the share is immediately available. The new share is also added as an entry to `Sharing → Apple (AFP)`.

Mac OS X users can connect to the guest AFP share by clicking `Go → Connect to Server`. In the example shown in [Figure 9.3](#), the user has 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. 9.3: Connect to Server Dialogue

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

9.1.2 Creating Authenticated and Time Machine Shares

Mac OS X includes the Time Machine application which can be used to schedule automatic backups. In this configuration example, a Time Machine user will be configured to backup to an AFP share on a TrueNAS® system. It is recommended to create a separate Time Machine share for each user that will be using Time Machine to backup their Mac OS X system to TrueNAS®. The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

To use the Wizard to create an authenticated or Time Machine share, enter the following information, as seen in the example in [Figure 9.4](#).

1. Share name: enter a name for the share that is identifiable but less than 27 characters long. The name cannot contain a period. In this example, the share is named `backup_user1`.
2. Click the button for Mac OS X (AFP) and check the box for Time Machine.
3. Click the Ownership button. If the user already exists on the TrueNAS® system, click the drop-down User menu to select their user account. If the user does not yet exist on the TrueNAS® system, type their name into the User field and check the Create User checkbox. If the user will be a member of a group that already exists on the TrueNAS® system, click the drop-down Group menu to select the group name. To create a new group to be used by Time Machine users, enter the name in the Group field and check the Create Group checkbox. Otherwise, enter the same name as the user. In

the example shown in [Figure 9.5](#), both a new user1 user and a new tm_backups group will be created. Since a new user is being created, this screen prompts for the user password to be used when accessing the share. It also provides an opportunity to change the default permissions on the share. When finished, click Return to return to the screen shown in [Figure 9.4](#).

4. Click the Add button. Remember to do this or the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

To configure multiple authenticated or Time Machine shares, repeat for each user, giving each user their own Share name and Ownership. When finished, click the Next button twice, then the Confirm button to create the shares. The Wizard automatically creates a dataset for each share with the correct ownership and starts the AFP service so the shares are immediately available. The new shares are also added to *Sharing* → *Apple* (AFP).

The screenshot shows a 'Wizard' window with a dark title bar. Inside, the 'Share name' field is set to 'backup_user1'. Below this, the 'Purpose' section has four radio buttons: 'Windows (SMB)', 'Mac OS X (AFP)' (selected), 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' (unchecked) and 'Time Machine' (checked). An 'Ownership' button is to the right of the checkboxes. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. Below these is a list box with a header 'Name' and one entry 'backup_user1'. At the bottom of the window are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 9.4: Creating a Time Machine Share

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

Buttons: Return, Cancel

Fig. 9.5: Creating an Authenticated User

At this point, it may be desirable to configure a quota for each Time Machine share, to restrict backups from using all of the available space on the TrueNAS® system. The first time Time Machine makes a backup, it will create a full backup after waiting two minutes. It will then create a one hour incremental backup for the next 24 hours, and then one backup each day, each week and each month. Since the oldest backups are deleted when a Time Machine share becomes full, make sure that the quota size is sufficient to hold the desired number of backups. Note that a default installation of Mac OS X is ~21 GB in size.

To configure a quota, go to **Storage** → **Volumes** and highlight the entry for the share. In the example shown in [Figure 9.6](#), the Time Machine share name is `backup_user1`. Click the **Edit Options** button for the share, then **Advanced Mode**. Enter a value in the **Quota** for this dataset field, then click **Edit Dataset** to save the change. In this example, the Time Machine share is restricted to 200 GB.




Fig. 9.6: Setting a Quota

Note: An alternative is to create a global quota using the instructions in [Set up Time Machine for multiple machines with OSX Server-Style Quotas](https://forums.freenas.org/index.php?threads/how-to-set-up-time-machine-for-multiple-machines-with-osx-server-style-quotas.47173/) (<https://forums.freenas.org/index.php?threads/how-to-set-up-time-machine-for-multiple-machines-with-osx-server-style-quotas.47173/>).

To configure Time Machine on the Mac OS X client, go to `System Preferences → Time Machine` which opens the screen shown in [Figure 9.7](#). Click ON and a pop-up menu shows the TrueNAS® system as a backup option. In our example, it is listed as backup_user1 on "freenas". Highlight the TrueNAS® system and click Use Backup Disk. A connection bar opens and prompts for the user account's password—in this example, the password that was set for the user1 account.



Fig. 9.7: Configuring Time Machine on Mac OS X Lion

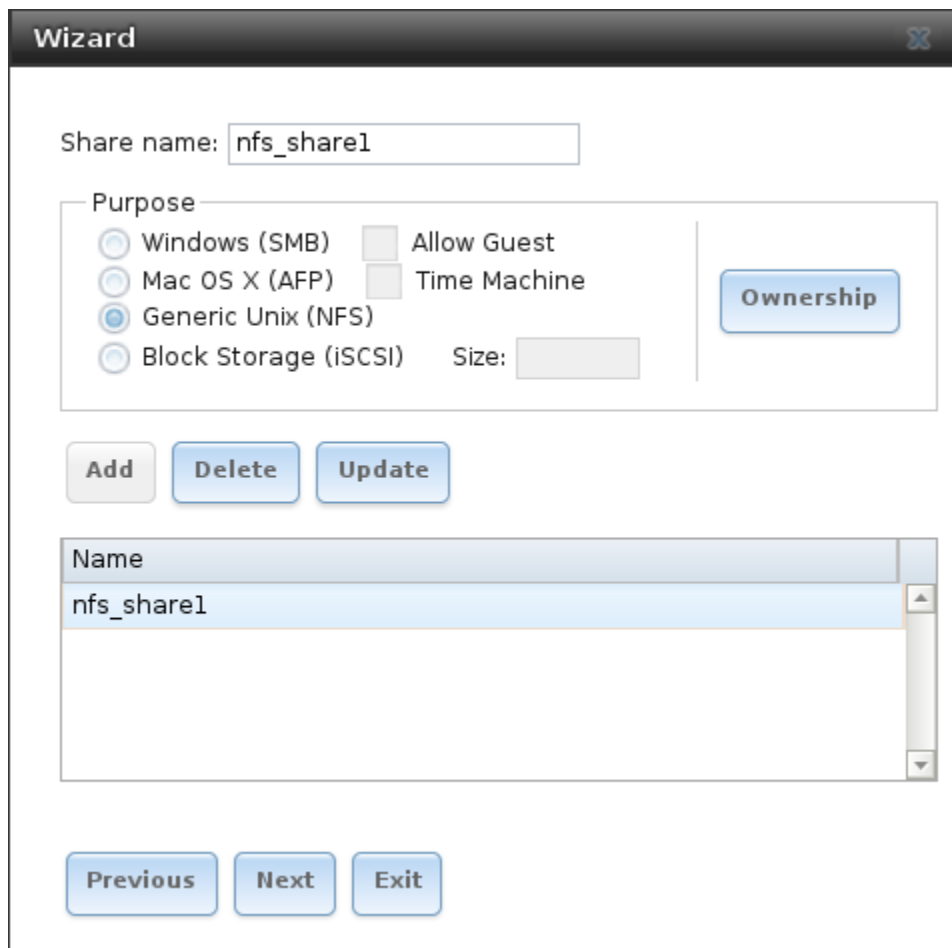
If Time Machine could not complete the backup. The backup disk image could not be created (error 45) is shown when backing up to the TrueNAS[®] system, a sparsebundle image must be created using [these instructions](http://forum1.netgear.com/showthread.php?t=49482) (<http://forum1.netgear.com/showthread.php?t=49482>).

If Time Machine completed a verification of your backups. To improve reliability, Time Machine must create a new backup for you. is shown, follow the instructions in [this post](http://www.garth.org/archives/2011,08,27,169,fix-time-machine-sparsebundle-nas-based-backup-errors.html) (<http://www.garth.org/archives/2011,08,27,169,fix-time-machine-sparsebundle-nas-based-backup-errors.html>) to avoid making another backup or losing past backups.

9.2 Unix (NFS) Shares

TrueNAS[®] supports sharing over the Network File System (NFS). Clients use the `mount` command to mount the share. Once mounted, the NFS share appears as just another directory on the client system. Some Linux distros require the installation of additional software in order to mount an NFS share. On Windows systems, enable Services for NFS in the Ultimate or Enterprise editions or install an NFS client application.

To create an NFS share using the [Wizard](#) (page 250), click the Next button twice to display the screen shown in [Figure 9.8](#). Enter a Share name. Spaces are not allowed in these names. Click the button for Generic Unix (NFS), then click Add so the share name appears in the Name frame. When finished, click the Next button twice, then the Confirm button to create the share. Creating an NFS share using the wizard automatically creates a new dataset for the share, starts the services required for NFS, and adds an entry in `Sharing → Unix (NFS) Shares`. Depending on your requirements, the IP addresses that are allowed to access the NFS share can be restricted, or the permissions adjusted.



The image shows a 'Wizard' dialog box for configuring NFS shares. At the top, the title bar says 'Wizard'. Below it, there's a text field for 'Share name:' containing 'nfs_share1'. Under the 'Purpose' section, there are four radio buttons: 'Windows (SMB)', 'Mac OS X (AFP)', 'Generic Unix (NFS)' (which is selected), and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' and 'Time Machine'. A 'Size:' label is followed by an empty text field. A blue button labeled 'Ownership' is positioned to the right of the checkboxes. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. A list box below these buttons has a header 'Name' and contains one entry, 'nfs_share1', which is highlighted. At the bottom of the dialog are three buttons: 'Previous', 'Next', and 'Exit'.

Wizard

Share name:

Purpose

☐ Windows (SMB) ☐ Allow Guest

☐ Mac OS X (AFP) ☐ Time Machine

☒ Generic Unix (NFS)

☐ Block Storage (iSCSI) Size:

Ownership

Add **Delete** **Update**

Name
nfs_share1

Previous **Next** **Exit**

Fig. 9.8: NFS Share Wizard

NFS shares are edited by clicking *Sharing* → *Unix (NFS)*, highlighting the entry for the share, and clicking its Edit button. In the example shown in [Figure 9.9](#), the configuration screen is open for the `nfs_share1` share.




Fig. 9.9: NFS Share Settings

Table 9.2 summarizes the available configuration options in this screen. Some settings are only available by clicking the Advanced Mode button.

Table 9.2: NFS Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		browse to the volume or dataset to be shared; click Add extra path to select multiple paths
Comment	string		set the share name; if left empty, share name is the list of selected Path entries
Authorized networks	string	✓	list of allowed networks in network/mask CIDR notation, like 1.2.3.0/24, space-delimited; leave empty to allow all
Authorized IP addresses or hosts	string	✓	list of allowed IP addresses or hostnames, space-delimited; leave empty to allow all
All directories	checkbox		when checked, allow the client to mount any subdirectory within the Path
Read only	checkbox		prohibit writing to the share
Quiet	checkbox	✓	inhibit otherwise-useful syslog diagnostics to avoid some annoying error messages; see exports(5) (http://www.freebsd.org/cgi/man.cgi?query=exports) for examples
Maproot User	drop-down menu	✓	when a user is selected, the root user is limited to that user's permissions

Continued on next page

Table 9.2 – continued from previous page

Setting	Value	Advanced Mode	Description
Maproot Group	drop-down menu	✓	when a group is selected, the root user is also limited to that group's permissions
Mapall User	drop-down menu	✓	the specified user's permissions are used by all clients
Mapall Group	drop-down menu	✓	the specified group's permissions are used by all clients
Security	selection	✓	only appears if Enable NFSv4 is checked in <i>Services</i> → <i>NFS</i> ; choices are <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

When creating NFS shares, keep these points in mind:

1. Clients will specify the Path when mounting the share.
2. The Maproot and Mapall options are exclusive, meaning only one can be used—the GUI does not allow both. The Mapall options supersede the Maproot options. To restrict only the root user's permissions, set the Maproot option. To restrict permissions of all users, set the Mapall options.
3. Each volume or dataset is considered to be its own filesystem and NFS is not able to cross filesystem boundaries.
4. The network or host must be unique per share and per filesystem or directory.
5. The All directories option can only be used once per share per filesystem.

To better understand these restrictions, consider a scenario where there are:

- two networks, 10.0.0.0/8 and 20.0.0.0/8
- a ZFS volume named `volume1` with 2 datasets named `dataset1` and `dataset2`
- `dataset1` contains a directory named `directory1`

Because of restriction #3, an error is shown when trying to create one NFS share like this:

- Authorized networks set to 10.0.0.0/8 20.0.0.0/8
- Path set to `/mnt/volume1/dataset1` and `/mnt/volume1/dataset1/directory1`

Instead, set a Path of `/mnt/volume1/dataset1` and check the All directories box.

That directory could also be restricted to one of the networks by creating two shares instead:

First NFS share:

- Authorized networks set to 10.0.0.0/8
- Path set to `/mnt/volume1/dataset1`

Second NFS share:

- Authorized networks set to 20.0.0.0/8
- Path set to `/mnt/volume1/dataset1/directory1`

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

9.2.1 Example Configuration

By default, the Mapall fields are not set. This means that when a user connects to the NFS share, the user has the permissions associated with their user account. This is a security risk if a user is able to connect as root as they will have complete access to the share.

A better option is to do this:

1. Specify the built-in nobody account to be used for NFS access.
2. In the Change Permissions screen of the volume/dataset that is being shared, change the owner and group to nobody and set the permissions according to your requirements.
3. Select nobody in the Mapall User and Mapall Group drop-down menus for the share in *Sharing* → *Unix (NFS) Shares*.

With this configuration, it does not matter which user account connects to the NFS share, as it will be mapped to the nobody user account and will only have the permissions that were specified on the volume/dataset. For example, even if the root user is able to connect, it will not gain root access to the share.

9.2.2 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/volume1/nfs_share1` is created and the permissions set to the nobody user account and the nobody group.
3. An NFS share is created with these attributes:
 - Path: `/mnt/volume1/nfs_share1`
 - Authorized Networks: 192.168.2.0/24
 - All Directories checkbox is checked
 - MapAll User is set to nobody
 - MapAll Group is set to nobody

From BSD or Linux

NFS shares are mounted on BSD or Linux clients with this command executed as the superuser (root) or with `sudo`:

```
mount -t nfs 192.168.2.2:/mnt/volume1/nfs_share1 /mnt
```

- `-t nfs` specifies the filesystem type of the share
- 192.168.2.2 is the IP address of the TrueNAS® system
- `/mnt/volume/nfs_share1` is the name of the directory to be shared, a dataset in this case
- `/mnt` is the mountpoint on the client system. This must be an existing, empty directory. The data in the NFS share appears in this directory on the client computer.

A successful mounting of the share returns to the command prompt without any status or error messages.

Note: If this command fails on a Linux system, make sure that the [nfs-utils](http://sourceforge.net/projects/nfs/files/nfs-utils/) (<http://sourceforge.net/projects/nfs/files/nfs-utils/>) package is installed.

This configuration allows users on the client system to copy files to and from `/mnt` (the mount point). All files are owned by `nobody:nobody`. Changes to any files or directories in `/mnt` are written to the TrueNAS[®] system's `/mnt/volume1/nfs_share1` dataset.

Settings cannot be changed on the NFS share if it is mounted on any client computers. The `umount` command is used to unmount the share on BSD and Linux clients. Run it as the superuser or with `sudo` on each client computer:

```
umount /mnt
```

From Microsoft

Windows NFS client support varies with versions and releases. For best results, use [Windows \(SMB\) Shares](#) (page 159).

From Mac OS X

To mount the NFS volume from a Mac OS X client, click on `Go → Connect to Server`. In the Server Address field, enter `nfs://` followed by the IP address of the TrueNAS[®] system and the name of the volume/dataset being shared by NFS. The example shown in [Figure 9.10](#) continues with our example of `192.168.2.2:/mnt/volume1/nfs_share1`.

Finder opens automatically after connecting. The IP address of the TrueNAS[®] system is displayed in the SHARED section in the left frame and the contents of the share are displayed in the right frame. In the example shown in [Figure 9.11](#), `/mnt/data` has one folder named `images`. The user can now copy files to and from the share.




Fig. 9.10: Mounting the NFS Share from Mac OS X




Fig. 9.11: Viewing the NFS Share in Finder

9.2.3 Troubleshooting NFS

Some NFS clients do not support the NLM (Network Lock Manager) protocol used by NFS. This is the case if the client receives an error that all or part of the file may be locked when a file transfer is attempted. To resolve this error, add the option `-o nolock` when running the `mount` command on the client to allow write access to the NFS share.

If a “time out giving up” error is shown when trying to mount the share from a Linux system, make sure that the portmapper service is running on the Linux client. If portmapper is running and timeouts are still shown, force the use of TCP by including `-o tcp` in the `mount` command.

If a “RPC: Program not registered” error is shown, upgrade to the latest version of 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 data base field in `Network → Global Configuration`.

Some older versions of NFS clients default to UDP instead of TCP and do not auto-negotiate for TCP. By default, TrueNAS®

uses TCP. To support UDP connections, go to *Services* → *NFS* and check the box *Serve UDP NFS clients*.
The `nfsstat -c` or `nfsstat -s` commands can be helpful to detect problems from the *Shell* (page 257). A high proportion of retries and timeouts compared to reads usually indicates network problems.

9.3 WebDAV Shares

In TrueNAS®, WebDAV shares can be created so that authenticated users can browse the contents of the specified volume, dataset, or directory from a web browser.

Configuring WebDAV shares is a two step process. First, create the WebDAV shares to specify which data can be accessed. Then, configure the WebDAV service by specifying the port, authentication type, and authentication password. Once the configuration is complete, the share can be accessed using a URL in the format:

```
protocol://IP_address:port_number/share_name
```

where:

- `protocol`: is either `http` or `https`, depending upon the Protocol configured in *Services* → *WebDAV*.
- `IP address`: is the IP address or hostname of the TrueNAS® system. Take care when configuring a public IP address to ensure that the network's firewall only allows access to authorized systems.
- `port_number`: is configured in *Services* → *WebDAV*. If the TrueNAS® system is to be accessed using a public IP address, consider changing the default port number and ensure that the network's firewall only allows access to authorized systems.
- `share_name`: is configured in *Sharing* → *WebDAV Shares*.

Entering the URL in a web browser brings up an authentication pop-up message. Enter a username of `webdav` and the password configured in *Services* → *WebDAV*.

Warning: At this time, only the `webdav` user is supported. For this reason, it is important to set a good password for this account and to only give the password to users which should have access to the WebDAV share.

To create a WebDAV share, click *Sharing* → *WebDAV Shares* → *Add WebDAV Share* which will open the screen shown in Figure 9.12.




Fig. 9.12: Adding a WebDAV Share

Table 9.3 summarizes the available options.

Table 9.3: WebDAV Share Options

Setting	Value	Description
Share Path Name	string	input a name for the share
Comment	string	optional
Path	browse button	browse to the volume/dataset to share
Read Only	checkbox	if checked, users cannot write to the share
Change User & Group Ownership	checkbox	if checked, automatically sets the share's contents to the webdav user and group

After clicking OK, a pop-up asks about enabling the service. Once the service starts, review the settings in *Services* → *WebDAV* as they are used to determine which URL is used to access the WebDAV share and whether or not authentication is required to access the share. These settings are described in *WebDAV* (page 220).

9.4 Windows (SMB) Shares

TrueNAS® uses *Samba* (<https://www.samba.org/>) to share volumes using Microsoft's SMB protocol. SMB is built into the Windows and Mac OS X operating systems and most Linux and BSD systems pre-install the Samba client in order to provide support for SMB. If your distro did not, install the Samba client using the distro's software repository.

The SMB protocol supports many different types of configuration scenarios, ranging from the very simple to quite complex. The complexity of the scenario depends upon the types and versions of the client operating systems that will connect to the share, whether the network has a Windows server, and whether Active Directory is being used. Depending on the authentication requirements, it might be necessary to create or import users and groups.

Samba supports server-side copy of files on the same share with clients from Windows 8 and higher. Copying between two different shares is not server-side. Windows 7 clients support server-side copying with *Robocopy* (<https://technet.microsoft.com/en-us/library/cc733145>).

This chapter starts by summarizing the available configuration options. It demonstrates some common configuration scenarios as well as offering some troubleshooting tips. It is recommended to first read through this entire chapter before creating any SMB shares to get a better idea of the configuration scenario that best meets your network's needs.

Tip: *SMB Tips and Tricks* (<https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/>) shows helpful hints for configuring and managing SMB networking. The *FreeNAS and Samba (CIFS) permissions* (<https://www.youtube.com/watch?v=RxggaE935PM>) and *Advanced Samba (CIFS) permissions on FreeNAS* (<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/>).

Tip: Run `smbstatus` from the *Shell* (page 257) for a list of active connections and users.

Figure 9.13 shows the configuration screen that appears after clicking *Sharing* → *Windows (SMB Shares)* → *Add Windows (SMB) Share*.

Add Windows (SMB) Share

Path: **Browse**

Use as home share: ☐

Name:

Apply Default Permissions: ☒ ⓘ

Allow Guest Access: ☐ ⓘ

OK **Cancel** **Advanced Mode**

Fig. 9.13: Adding an SMB Share

Table 9.4 summarizes the options when creating a SMB share. Some settings are only available after clicking the Advanced Mode button. For simple sharing scenarios, Advanced Mode options are not needed. For more complex sharing scenarios, only change an Advanced Mode option after fully understanding the function of that option. [smb.conf\(5\)](https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>) provides more details for each configurable option.

Table 9.4: Options for a SMB Share

Setting	Value	Advanced Mode	Description
Path	browse button		select volume/dataset/directory to share
Use as home share	checkbox		check this box if the share is meant to hold user home directories; only one share can be the homes share
Name	string		mandatory; name of share
Comment	string	✓	optional description
Apply Default Permissions	checkbox		sets the ACLs to allow read/write for owner/group and read-only for others; should only be unchecked when creating a share on a system that already has custom ACLs set
Export Read Only	checkbox	✓	prohibits write access to the share
Browsable to Network Clients	checkbox	✓	when checked, users see the contents of /homes (including other home directories of other users) and when unchecked, users see only their own home directory
Export Recycle Bin	checkbox	✓	deleted files are moved to a hidden <code>.recycle</code> in the root folder of the share; the <code>.recycle</code> directory can be deleted to reclaim space and is automatically recreated when a file is deleted
Show Hidden Files	checkbox	✓	if enabled, the Windows hidden attribute is not set when file-names that begin with a dot (a Unix hidden file) are created; existing files are not affected

Continued on next page

Table 9.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Allow Guest Access	checkbox		if checked, a password is not required to connect to the share; connections with a bad password are rejected unless the user account does not exist, in which case it is mapped to the guest account and granted the permissions of the guest user defined in the SMB (page 208) service
Only Allow Guest Access	checkbox	✓	requires Allow guest access to also be checked; forces guest access for all connections
Access Based Share Enumeration	checkbox	✓	when checked, users can only see the shares they have permission to access; to change the default that grants Everyone access, use the computer management MMC on Windows or the sharesec command-line utility
Hosts Allow	string	✓	comma-, space-, or tab-delimited list of allowed hostnames or IP addresses
Hosts Deny	string	✓	comma-, space-, or tab-delimited list of denied hostnames or IP addresses; allowed hosts take precedence so can use ALL in this field and specify allowed hosts in Hosts Allow
VFS Objects	selection	✓	adds virtual file system modules to enhance functionality; Table 9.5 summarizes the available modules
Periodic Snapshot Task	drop-down menu	✓	used to configure directory shadow copies on a per-share basis; select the pre-configured periodic snapshot task to use for the share's shadow copies; periodic snapshot must be recursive
Auxiliary Parameters	string	✓	additional <code>smb4.conf</code> parameters not covered by other option fields

Note the following regarding some of the Advanced Mode settings:

- Hostname lookups add some time to accessing the SMB share. If you only use IP addresses, uncheck the Hostnames lookups box in `Services` → `SMB`.
- When the Browseable to Network Clients box is checked (the default), the share is visible through Windows File Explorer or through `net view`. When the Use as a home share box is checked, unchecking the Browseable to Network Clients box hides the share named homes so that only the dynamically generated share containing the authenticated user's home directory will be visible. By default, the homes share and the user's home directory are both visible. Users are not automatically granted read or write permissions on browsable shares. This option provides no real security because shares that are not visible in Windows File Explorer can still be accessed with a UNC path.
- If some files on a shared volume should be hidden and inaccessible to users, put a `veto files=` line in the Auxiliary Parameters field. The syntax for the veto files option and some examples can be found in the [smb.conf manual page](#) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>).

Samba disables NTLMv1 authentication by default for security. Standard configurations of Windows XP and some configurations of later clients like Windows 7 will not be able to connect with NTLMv1 disabled. [Security guidance for NTLMv1 and LM network authentication](#) (<https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-network-authentication>) has information about the security implications and ways to enable NTLMv2 on those clients. If changing the client configuration is not possible, NTLMv1 authentication can be enabled by checking the box NTLMv1 auth in `Services` → `SMB`.

[Table 9.5](#) provides an overview of the available VFS modules. Be sure to research each module before adding or deleting it from the Selected column of the VFS Objects field of the share. Some modules need additional configuration after they are added. Refer to [Stackable VFS modules](#) (<https://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/VFS.html>) and the [vfs_* man pages](#) (<https://www.samba.org/samba/docs/man/manpages/>) for more details.

Table 9.5: Available VFS Modules

Value	Description
acl_tdb	stores NTFS ACLs in a tdb file to enable full mapping of Windows ACLs
acl_xattr	stores NTFS ACLs in Extended Attributes (EAs) to enable the full mapping of Windows ACLs
aio_fork	enables async I/O
aio_pthread	implements async I/O in Samba vfs using a pthread pool instead of the internal Posix AIO interface
audit	logs share access, connects/disconnects, directory opens/creates/removes, and file opens/closes/renames/unlinks/chmods to syslog
cacheprime	primes the kernel file data cache
cap	translates filenames to and from the CAP encoding format, commonly used in Japanese language environments
catia	improves Mac interoperability by translating characters that are unsupported by Windows
commit	tracks the amount of data written to a file and synchronizes it to disk when a specified amount accumulates
crossrename	allows server side rename operations even if source and target are on different physical devices
default_quota	stores the default quotas that are reported to a windows client in the quota record of a user
dfs_samba4	distributed file system for providing an alternative name space, load balancing, and automatic failover
dirsort	sorts directory entries alphabetically before sending them to the client
expand_msdfs	enables support for Microsoft Distributed File System (DFS)
extd_audit	sends audit logs to both syslog and the Samba log files
fake_acls	stores file ownership and ACLs as extended attributes
fake_perms	allows roaming profile files and directories to be set as read-only
fruit	enhances OS X support by providing the SMB2 AAPL extension and Netatalk interoperability; automatically loads catia and streams_xattr but read the caveat in NOTE below table
full_audit	record selected client operations to the system log; if selected, a warning will indicate that Windows 10 clients may experience issues when transferring files to the NAS system when this module is enabled
linux_xfs_sgid	used to work around an old Linux XFS bug
media_harmony	allows Avid editorial workstations to share a network drive
netatalk	eases the co-existence of SMB and AFP shares
offline	marks all files in the share with the DOS offline attribute; this can prevent Windows Explorer from reading files just to make thumbnail images
posix_eadb	provides Extended Attributes (EAs) support so they can be used on filesystems which do not provide native support for EAs

Continued on next page

Table 9.5 – continued from previous page

Value	Description
preopen	useful for video streaming applications that want to read one file per frame
readahead	useful for Windows Vista clients reading data using Windows Explorer
readonly	marks a share as read-only for all clients connecting within the configured time period
shadow_copy	allows Microsoft shadow copy clients to browse shadow copies on Windows shares
shadow_copy_test	shadow copy testing
shell_snap	provides shell-script callouts for snapshot creation and deletion operations issued by remote clients using the File Server Remote VSS Protocol (FSRVP)
skel_opaque	implements dummy versions of all VFS modules (useful to VFS module developers)
skel_transparent	implements dummy passthrough functions of all VFS modules (useful to VFS module developers)
snapper	provides the ability for remote SMB clients to access shadow copies of FSRVP snapshots using Windows Explorer
streams_depot	experimental module to store alternate data streams in a central directory; the association with the primary file can be lost due to inode numbers changing when a directory is copied to a new location (see http://marc.info/?l=samba&m=132542069802160&w=2)
streams_xattr	enables storing of NTFS alternate data streams in the file system
syncops	ensures metadata operations are performed synchronously
time_audit	logs system calls that take longer than the number of defined milliseconds
unityed_media	allows multiple Avid clients to share a network drive
winmsa	emulate Microsoft's MoveSecurityAttributes=0 registry option, setting the ACL for file and directory hierarchies to inherit from the parent directory into which they are moved
worm	controls the writability of files and folders depending on their change time and an adjustable grace period
xattr_tdb	stores Extended Attributes (EAs) in a tdb file so they can be used on filesystems which do not provide support for EAs
zfs_space	correctly calculates ZFS space used by the share, including space used by ZFS snapshots, quotas, and reservations; enabled by default
zfsacl	provide ACL extensions for proper integration with ZFS; enabled by default

Note: Be careful when using multiple SMB shares, some with and some without fruit. OS X clients negotiate SMB2 AAPL protocol extensions on the first connection to the server, so mixing shares with and without fruit will globally disable AAPL if the first connection occurs without fruit. To resolve this, all OS X clients need to disconnect from all SMB shares and the first reconnection to the server has to be to a fruit-enabled share.

These VFS objects do not appear in the selection box:

-
- recycle: moves deleted files to the recycle directory instead of deleting them. Controlled by Export Recycle Bin in the [SMB share options](#) (page 160).
 - shadow_copy2: a more recent implementation of shadow_copy with some additional features. shadow_copy2 and the associated parameters are automatically added to the `smb4.conf` when a Periodic Snapshot Task is selected.

9.4.1 Configuring Unauthenticated Access

SMB supports guest logins, meaning that users can access the SMB share without needing to provide a username or password. This type of share is convenient as it is easy to configure, easy to access, and does not require any users to be configured on the 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.

To configure an unauthenticated SMB share, click Wizard, then click the Next button twice to display the screen shown in [Figure 9.14](#). Complete the following fields in this screen:

1. Share name: enter a name for the share that is useful to you. In this example, the share is named `smb_insecure`.
2. Click the button for Windows (SMB) and check the box for Allow Guest.
3. Click the Ownership button. Click the drop-down User menu and select nobody. Click the Return button to return to the previous screen.
4. Click the Add button. If you forget to do this, the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

Wizard

Share name:

Purpose

☒ Windows (SMB) ☒ Allow Guest
☐ Mac OS X (AFP) ☐ Time Machine
☐ Generic Unix (NFS)
☐ Block Storage (iSCSI) Size:

Ownership

Add **Delete** **Update**

Name
smb_insecure

Previous **Next** **Exit**

Fig. 9.14: Creating an Unauthenticated SMB Share

Click the Next button twice, then the Confirm button to create the share. The Wizard automatically creates a dataset for the share and starts the SMB service so the share is immediately available. The new share is also be added to *Sharing* → *Windows (SMB)*.

Users can now access the share from any SMB client and will not be prompted for their username or password. For example, to access the share from a Windows system, open Explorer and click on Network. For this configuration example, a system named FREENAS appears with a share named insecure_smb. The user can copy data to and from the unauthenticated SMB share.

9.4.2 Configuring Authenticated Access Without a Domain Controller

Most configuration scenarios require each user to have their own user account and to authenticate before accessing the share. This allows the administrator to control access to data, provide appropriate permissions to that data, and to determine who accesses and modifies stored data. A Windows domain controller is not needed for authenticated SMB shares, which means that additional licensing costs are not required. However, since there is no domain controller to provide authentication for the network, each user account needs to be created on the TrueNAS® system. This type of configuration scenario is often used in home and small networks as it does not scale well if many users accounts are needed.

Before configuring this scenario, determine which users will need authenticated access. While not required for the configuration, it eases troubleshooting if the username and password that will be created on the TrueNAS® system matches that information on the client system. Next, determine if each user should have their own share to store their own data or if several users will be using the same share. The simpler configuration is to make one share per user as it does not require the creation of groups, adding the correct users to the groups, and ensuring that group permissions are set correctly.

To use the Wizard to create an authenticated SMB share, enter the following information, as shown in the example in [Figure 9.15](#).

1. Share name: enter a name for the share that is useful to you. In this example, the share is named smb_user1.
2. Click the button for Windows (SMB).
3. Click the Ownership button. To create the user account on the TrueNAS® system, type their name into the User field and check the Create User checkbox. The user's password is then entered and confirmed. If the user will not be sharing this share with other users, type their name into the Group field and click Create Group. If, however, the share will be used by several users, instead type in a group name and check the Create Group box. In the example shown in [Figure 9.16](#), user1 has been used for both the user and group name, meaning that this share will only be used by user1. When finished, click Return to return to the screen shown in [Figure 9.15](#).
4. Click the Add button. If you forget to do this, the share will not be created. Clicking the Add button adds an entry to the Name frame with the name that was entered in Share name.

If you wish to configure multiple authenticated shares, repeat for each user, giving each user their own Share name and Ownership. When finished, click Next twice, then Confirm to create the shares. The Wizard automatically creates a dataset with the correct ownership for each share and starts the SMB service so the shares are available immediately. The new shares are also added to `Sharing → Windows (SMB)`.

The screenshot shows the 'Wizard' window with a dark title bar. Inside, the 'Share name' field contains 'smb_user1'. Below it, the 'Purpose' section has four radio buttons: 'Windows (SMB)' (selected), 'Mac OS X (AFP)', 'Generic Unix (NFS)', and 'Block Storage (iSCSI)'. To the right of these are two checkboxes: 'Allow Guest' (checked) and 'Time Machine' (unchecked). A 'Size' field is also present. An 'Ownership' button is located to the right of the 'Purpose' section. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. A table with one row and one column is shown, with the header 'Name' and the value 'smb_user1'. At the bottom are three buttons: 'Previous', 'Next', and 'Exit'.

Fig. 9.15: Creating an Authenticated SMB Share

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

Fig. 9.16: Creating the User and Group

The authenticated share can now be tested from any SMB client. For example, to test an authenticated share from a Windows system, open Explorer and click on Network. For this configuration example, a system named FREENAS appears with a share named smb_user1. If you click on smb_user1, a Windows Security pop-up screen prompts for that user's username and password. Enter the values that were configured for that share, in this case user user1. After authentication, the user can copy data to and from the SMB share.

To prevent Windows Explorer from hanging when accessing the share, map the share as a network drive. To do this, right-click the share and select Map network drive.... Choose a drive letter from the drop-down menu and click the Finish button.

Note that Windows systems cache a user's credentials. This can cause issues when testing or accessing multiple authenticated shares as only one authentication is allowed at a time. If you are having problems authenticating to a share and are sure that you are entering the correct username and password, type cmd in the Search programs and files box and use the following command to see if you have already authenticated to a share. In this example, the user has already authenticated to the smb_user1 share:

```
net use
New connections will be remembered.
```

Status	Local	Remote	Network
OK		\\FREENAS\smb_user1	Microsoft Windows Network

```
The command completed successfully.
```

To clear the cache:

```
net use * /DELETE
You have these remote connections:
    \\FREENAS\smb_user1
Continuing will cancel the connections.

Do you want to continue this operation? <Y/N> [N]: y
```

An additional warning is shown if the share is currently open in Explorer:

```
There are open files and/or incomplete directory searches pending on the connection
to \\FREENAS\smb_user1.

Is it OK to continue disconnecting and force them closed? <Y/N> [N]: y
The command completed successfully.
```

The next time a share is accessed with Explorer, you will be prompted to authenticate.

9.4.3 Configuring Shadow Copies

Shadow Copies (https://en.wikipedia.org/wiki/Shadow_copy), also known as the Volume Shadow Copy Service (VSS) or Previous Versions, is a Microsoft service for creating volume snapshots. Shadow copies allow you to easily restore previous versions of files from within Windows Explorer. Shadow Copy support is built into Vista and Windows 7. Windows XP or 2000 users need to install the **Shadow Copy client** (<http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=16220>).

When you create a periodic snapshot task on a ZFS volume that is configured as a SMB share in 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 you are unable to see any previous versions of files to restore, use Windows Update to make sure that the system is fully up-to-date.
- Shadow copy support only works for ZFS pools or datasets. This means that the SMB share must be configured on a volume or dataset, not on a directory.
- Datasets are filesystems and shadow copies cannot traverse filesystems. If you want to be able to see the shadow copies in your child datasets, create separate shares for them.
- Shadow copies will not work with a manual snapshot, you must create a periodic snapshot task for the pool or dataset being shared by SMB or a recursive task for a parent dataset.
- The periodic snapshot task should be created and at least one snapshot should exist before creating the SMB share. If the SMB share was created first, restart the SMB service in `Services → Control Services`.
- Appropriate permissions must be configured on the volume/dataset being shared by SMB.
- Users cannot delete shadow copies on the Windows system due to the way Samba works. Instead, the administrator can remove snapshots from the TrueNAS® administrative GUI. The only way to disable shadow copies completely is to remove the periodic snapshot task and delete all snapshots associated with the SMB share.

To configure shadow copy support, use the instructions in [Configuring Authenticated Access Without a Domain Controller](#) (page 165) to create the desired number of shares. In this configuration example, a Windows 7 computer has two users: user1 and user2. For this example, two authenticated shares are created so that each user account has their own share. The first share is named user1 and the second share is named user2. Then:

1. Use `Storage → Periodic Snapshot Tasks → Add Periodic Snapshot` to create at least one periodic snapshot task. You can either create a snapshot task for each user's dataset, in this example the datasets `/mnt/volume1/user1` and `/mnt/volume1/user2`, or you can create one periodic snapshot task for the entire volume, in this case `/mnt/volume1`. Before continuing to the next step, confirm that at least one snapshot for each defined task is displayed in the `Storage → Snapshots` tab. When creating the schedule for the periodic snapshot tasks, keep in mind how often your users need to access modified files and during which days and time of day they are likely to make changes.
2. Go to `Sharing → Windows (SMB) Shares`. Highlight a share and click `Edit`, then `Advanced Mode`. Click the `Periodic Snapshot Task` drop-down menu and select the periodic snapshot task to use for that share. Repeat for each share being configured as a shadow copy. For this example, the share named `/mnt/volume1/user1` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/volume1/user1` dataset and the share named `/mnt/volume1/user2` is configured to use a periodic snapshot task that was configured to take snapshots of the `/mnt/volume1/user2` dataset.
3. Verify that the SMB service is set to ON in `Services → Control Services`.

Figure 9.17 provides an example of using shadow copies while logged in as user1 on the Windows system. In this example, the user right-clicked modified file and selected `Restore previous versions` from the menu. This particular file has three versions: the current version, plus two previous versions stored on the TrueNAS® system. The user can choose to open one of the previous versions, copy a previous version to the current folder, or restore one of the previous versions, overwriting the existing file on the Windows system.

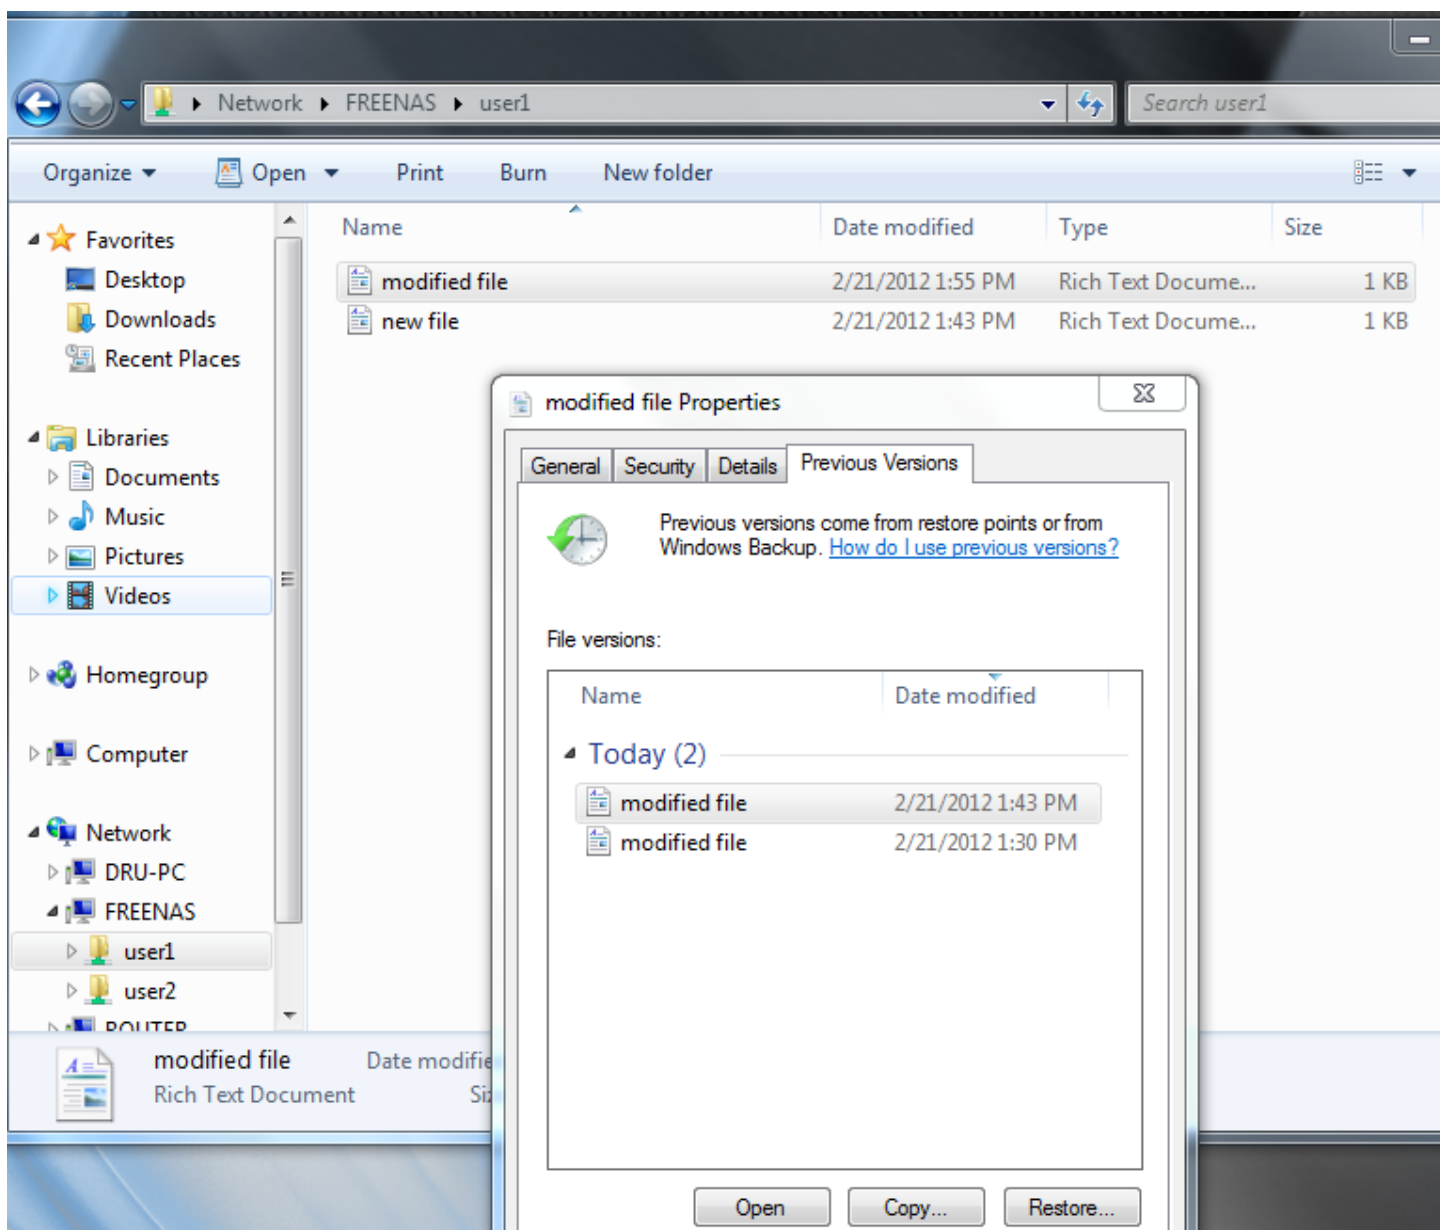


Fig. 9.17: Viewing Previous Versions within Explorer

9.5 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://technet.microsoft.com/en-us/library/hh831628) (<https://technet.microsoft.com/en-us/library/hh831628>), meaning that file copies happen locally, rather than over the network. It also supports the [VAAI](#) (page 311) (vStorage APIs for Array Integration) primitives for efficient operation of storage tasks directly on the NAS. To take advantage of the VAAI primitives, create a zvol using the instructions in [Create zvol](#) (page 99) and use it to create a device extent, as described in [Extents](#) (page 177).

To configure iSCSI:

1. Review the target global configuration parameters.
2. Create at least one portal.
3. Determine which hosts are allowed to connect using iSCSI and create an initiator.
4. Decide if authentication will be used, and if so, whether it will be CHAP or mutual CHAP. If using authentication, create an authorized access.
5. Create a target.
6. Create either a device or a file extent to be used as storage.
7. Associate a target with an extent.
8. Start the iSCSI service in `Services` → `Control Services`.

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

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 181) for details.

9.5.1 Target Global Configuration

Sharing → Block (iSCSI) → Target Global Configuration, shown in [Figure 9.18](#), contains settings that apply to all iSCSI shares. [Table 9.6](#) summarizes the settings that can be configured in the Target Global Configuration screen.

Some built-in values affect iSNS usage. Fetching of allowed initiators from iSNS is not implemented, so target ACLs must be configured manually. To make iSNS registration useful, iSCSI targets should have explicitly configured port IP addresses. This avoids initiators attempting to discover unconfigured target portal addresses like 0.0.0.0.

The iSNS registration period is 900 seconds. Registered Network Entities not updated during this period are unregistered. The timeout for iSNS requests is 5 seconds.

Sharing

Apple (AFP)UNIX (NFS)WebDAVWindows (SMB)Block (iSCSI)

Target Global ConfigurationPortalsInitiatorsAuthorized AccessTargetsExtentsAssociated Targets

Base Name:

iqn.2005-10.org.freenas.ctl

ISNS Servers:

Pool Available Space Threshold (%):

Enable iSCSI ALUA:

☐

Save

Fig. 9.18: iSCSI Target Global Configuration Variables

Table 9.6: Target Global Configuration Settings

Setting	Value	Description
Base Name	string	see the “Constructing iSCSI names using the iqn. format” section of RFC 3721 (https://tools.ietf.org/html/rfc3721.html) if unfamiliar with this format
ISNS Servers	string	space delimited list of hostnames or IP addresses of ISNS servers with which to register the system’s iSCSI targets and portals
Pool Available Space Threshold	integer	enter the percentage of free space that should remain in the pool; when this percentage is reached, the system issues an alert, but only if zvols are used; see VAAI (page 311) Threshold Warning
Continued on next page		

Table 9.6 – continued from previous page

Setting	Value	Description
Enable iSCSI ALUA	checkbox	enable ALUA for automatic best path discovery when supported by clients; this option is only available on HA systems

9.5.2 Portals

A portal specifies the IP address and port number to be used for iSCSI connections. [Sharing](#) → [Block \(iSCSI\)](#) → [Portals](#) → [Add Portal](#) brings up the screen shown in [Figure 9.19](#).

[Table 9.19](#) summarizes the settings that can be configured when adding a portal. If you need to assign additional IP addresses to the portal, click the link [Add extra Portal IP](#).

The screenshot shows the 'Add Portal' dialog box in the iSCSI configuration interface. The dialog is titled 'Add Portal' and contains the following fields and controls:

- Comment:** A text input field with an information icon.
- Discovery Auth Method:** A dropdown menu currently set to 'None'.
- Discovery Auth Group:** A dropdown menu currently set to 'None'.
- Portal IP:** A section header for the IP address and port configuration.
- IP Address:** A dropdown menu currently set to '0.0.0.0'.
- Port:** A text input field currently set to '3260'.
- Add extra Portal IP:** A link to add additional IP addresses.
- OK** and **Cancel** buttons at the bottom.

Fig. 9.19: Adding an iSCSI Portal

Table 9.7: Portal Configuration Settings

Setting	Value	Description
Comment	string	optional description; portals are automatically assigned a numeric group ID
Discovery Auth Method	drop-down menu	configures the authentication level required by the target for discovery of valid devices, where None will allow anonymous discovery while CHAP and Mutual CHAP require authentication
Discovery Auth Group	drop-down menu	select a user created in Authorized Access if the Discovery Auth Method is set to CHAP or Mutual CHAP
IP address	drop-down menu	select the IP address associated with an interface or the wildcard address of 0.0.0.0 (any interface)
Port	integer	TCP port used to access the iSCSI target; default is 3260

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 the following addresses:

192.168.1.1/24

192.168.2.1/24

192.168.3.1/24

192.168.4.1/24

You could create a portal containing the first two IP addresses (group ID 1) and a portal containing the remaining two IP addresses (group ID 2). You could then create a target named A with a Portal Group ID of 1 and a second target named B with a Portal Group ID of 2. In this scenario, the iSCSI service would listen on all four interfaces, but connections to target A would be limited to the first two networks and connections to target B would be limited to the last two networks.

Another scenario would be to create a portal which includes every IP address except for the one used by a management interface. This would prevent iSCSI connections to the management interface.

9.5.3 Initiators

The next step is to configure authorized initiators, or the systems which are allowed to connect to the iSCSI targets on the TrueNAS® system. To configure which systems can connect, use **Sharing → Block (iSCSI) → Initiators → Add Initiator**, shown in Figure 9.20.

The screenshot shows a window titled "Add Initiator" with a close button in the top right corner. Inside the window, there are three labeled input fields on the left and their corresponding values in text boxes on the right. The first field is "Initiators" with the value "ALL". The second field is "Authorized network" with the value "ALL". The third field is "Comment" with an empty text box. To the right of each text box is a small circular icon with an "i" inside, representing an information tooltip. At the bottom left of the window are two buttons: "OK" and "Cancel".

Fig. 9.20: Adding an iSCSI Initiator

Table 9.8 summarizes the settings that can be configured when adding an initiator.

Table 9.8: Initiator Configuration Settings

Setting	Value	Description
Initiators	string	use ALL keyword or a list of initiator hostnames separated by spaces
Authorized network	string	use ALL keyword or a network address with CIDR mask such as 192.168.2.0/24
Comment	string	optional description

In the example shown in Figure 9.21, two groups have been created. Group 1 allows connections from any initiator on any network. Group 2 allows connections from any initiator on the 10.10.1.0/24 network. Click an initiator's entry to display its Edit and Delete buttons.

Note: Attempting to delete an initiator causes a warning that indicates if any targets or target/extent mappings depend upon the initiator. Confirming the delete causes these to be deleted also.

Sharing

Apple (AFP)UNIX (NFS)WebDAVWindows (SMB)Block (iSCSI)

Target Global ConfigurationPortalsInitiatorsAuthorized AccessTargetsExtentsAssociated Targets

Add Initiator

Group ID	Initiators	Authorized network	Comment
1	ALL	ALL	
2	ALL	10.10.1.0/24	

Fig. 9.21: Sample iSCSI Initiator Configuration

9.5.4 Authorized Accesses

If you will be using CHAP or mutual CHAP to provide authentication, you must create an authorized access in Sharing → Block (iSCSI) → Authorized Accesses → Add Authorized Access. This screen is shown in Figure 9.22.

Note: This screen sets login authentication. This is different from discovery authentication which is set in Target Global Configuration (page 171).

Add Authorized Access

Group ID: 1

User: ⓘ

Secret: ⓘ

Secret (Confirm): ⓘ

Peer User: ⓘ

Peer Secret: ⓘ

Peer Secret (Confirm): ⓘ

OK Cancel

Fig. 9.22: Adding an iSCSI Authorized Access

Table 9.9 summarizes the settings that can be configured when adding an authorized access:

Table 9.9: Authorized Access Configuration Settings

Setting	Value	Description
Group ID	integer	allows different groups to be configured with different authentication profiles; for instance, all users with a Group ID of 1 will inherit the authentication profile associated with Group 1
User	string	name of user account to create for CHAP authentication with the user on the remote system; many initiators default to using the initiator name as the user
Secret	string	password to be associated with User; the iSCSI standard requires that this be between 12 and 16 characters
Peer User	string	only input when configuring mutual CHAP; in most cases it will need to be the same value as User
Peer Secret	string	the mutual secret password which must be different than the Secret; required if Peer User is set

Note: CHAP does not work with GlobalSAN initiators on Mac OS X.

As authorized accesses are added, they will be listed under View Authorized Accesses. In the example shown in Figure 9.23, three users (test1, test2, and test3) and two groups (1 and 2) have been created, with group 1 consisting of one CHAP user and group 2 consisting of one mutual CHAP user and one CHAP user. Click an authorized access entry to display its Edit and Delete buttons.

Sharing

Apple (AFP)

UNIX (NFS)

WebDAV

Windows (SMB)

Block (iSCSI)

Target Global Configuration

Portals

Initiators

Authorized Access

Targets

Extents

Associated Targets

Add Authorized Access

Group ID	User	Peer User
1	test1	
2	test2	test2
2	test3	

Fig. 9.23: Viewing Authorized Accesses

9.5.5 Targets

Next, create a Target using `Sharing → Block (iSCSI) → Targets → Add Target`, as shown in [Figure 9.24](#). A target combines a portal ID, allowed initiator ID, and an authentication method. [Table 9.10](#) summarizes the settings that can be configured when creating a Target.

Note: An iSCSI target creates a block device that may be accessible to multiple initiators. A clustered filesystem is required on the block device, such as VMFS used by VMware ESX/ESXi, in order for multiple initiators to mount the block device read/write. If a traditional filesystem such as EXT, XFS, FAT, NTFS, UFS, or ZFS is placed on the block device, care must be taken that only one initiator at a time has read/write access or the result will be filesystem corruption. If multiple clients need access to the same data on a non-clustered filesystem, use SMB or NFS instead of iSCSI, or create multiple iSCSI targets (one per client).

Add Target

Target Name:

Target Alias:

iSCSI Group

Portal Group ID:

Initiator Group ID:

Auth Method: **None**

Authentication Group number: **None**

[Add extra iSCSI Group](#)

OK **Cancel**

Base Name will be appended automatically when starting without 'iqn.', 'eui.' or 'naa.'

Fig. 9.24: Adding an iSCSI Target

Table 9.10: Target Settings

Setting	Value	Description
Target Name	string	required value; base name will be appended automatically if it does not start with iqn
Target Alias	string	optional user-friendly name
Portal Group ID	drop-down menu	leave empty or select number of existing portal to use
Initiator Group ID	drop-down menu	select which existing initiator group has access to the target
Auth Method	drop-down menu	choices are None, Auto, CHAP, or Mutual CHAP
Authentication Group number	drop-down menu	None or integer representing number of existing authorized access

9.5.6 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, a hardware RAID volume, or a [HAST device](http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html) (http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html).

File extents provide virtual storage access to an individual file.

Tip: For typical use as storage for virtual machines where the virtualization software is the iSCSI initiator, device extents with zvols provide the best performance and most features. For other applications, device extents sharing a raw device can be appropriate. File extents do not have the performance or features of device extents, but do allow creating multiple extents

on a single filesystem.

Virtualized zvols support all the TrueNAS® [VAAI](#) (page 311) primitives and are recommended for use with virtualization software as the iSCSI initiator.

The ATS, WRITE SAME, XCOPY and STUN, primitives are supported by both file and device extents. The UNMAP primitive is supported by zvols and raw SSDs. The threshold warnings primitive is fully supported by zvols and partially supported by file extents.

Virtualizing a raw device like a single disk or hardware RAID volume limits performance to the abilities of the device. Because this bypasses ZFS, such devices do not benefit from ZFS caching or provide features like block checksums or snapshots.

Virtualizing a zvol adds the benefits of ZFS, such as read and write cache. Even if the client formats a device extent with a different filesystem, the data still resides on a ZFS volume and benefits from ZFS features like block checksums and snapshots.

Warning: For performance reasons and to avoid excessive fragmentation, keep the used space of the pool below 50% when using iSCSI. The capacity of an existing extent can be increased as shown in [Growing LUNs](#) (page 185).

To add an extent, go to `Sharing → Block (iSCSI) → Extents → Add Extent`. In the example shown in [Figure 9.25](#), the device extent is using the `export` zvol that was previously created from the `/mnt/volume1` volume.

[Table 9.11](#) summarizes the settings that can be configured when creating an extent. Note that file extent creation will fail if you do not append the name of the file to be created to the volume/dataset name.

Add Extent

Extent Name:

Extent Type: Device

Device: ada1 (10.0 GiB)

Serial:

Logical Block Size: 512

Disable Physical Block Size Reporting: ☐

Comment:

Enable TPC: ☒

Xen initiator compat mode: ☐

LUN RPM: SSD

Read-only: ☐

String identifier of the extent.

Fig. 9.25: Adding an iSCSI Extent

Table 9.11: Extent Configuration Settings

Setting	Value	Description
Extent Name	string	name of extent; if the Extent size is not 0, it cannot be an existing file within the volume/dataset
Extent Type	drop-down menu	select from File or Device
Device	drop-down menu	only appears if Device is selected; select the unformatted disk, controller, zvol, zvol snapshot, or HAST device
Serial	string	unique LUN ID; the default is generated from the system's MAC address
Path to the extent	browse button	only appears if File is selected; browse to an existing file and use 0 as the Extent size, or browse to the volume or dataset, click Close, append the Extent Name to the path, and specify a value in Extent size; extents cannot be created inside the jail root directory
Extent size	integer	only appears if File is selected; if the size is specified as 0, the file must already exist and the actual file size will be used; otherwise, specify the size of the file to create
Logical Block Size	drop-down menu	only override the default if the initiator requires a different block size

Continued on next page

Table 9.11 – continued from previous page

Setting	Value	Description
Disable Physical Block Size Reporting	checkbox	if the initiator does not support physical block size values over 4K (MS SQL), check this box
Available Space Threshold	string	only appears if File or a zvol is selected; when the specified percentage of free space is reached, the system issues an alert; see VAAI (page 311) Threshold Warning
Comment	string	optional
Enable TPC	checkbox	if checked, an initiator can bypass normal access control and access any scannable target; this allows <code>xcopy</code> operations otherwise blocked by access control
Xen initiator compat mode	checkbox	check this box when using Xen as the iSCSI initiator
LUN RPM	drop-down menu	do NOT change this setting when using Windows as the initiator; only needs to be changed in large environments where the number of systems using a specific RPM is needed for accurate reporting statistics
Read-only	checkbox	check this box to prevent the initiator from initializing this LUN

9.5.7 Target/Extents

The last step is associating an extent to a target within `Sharing → Block (iSCSI) → Associated Targets → Add Target/Extent`. This screen is shown in [Figure 9.26](#). Use the drop-down menus to select the existing target and extent. Click OK to add an entry for the LUN.

Fig. 9.26: Associating a Target With an Extent

[Table 9.12](#) summarizes the settings that can be configured when associating targets and extents.

Table 9.12: Target/Extents Configuration Settings

Setting	Value	Description
Target	drop-down menu	select the pre-created target
LUN ID	drop-down menu	select the value to use or type in a value between 1 and 1023; note that some initiators expect a value below 256
Extent	drop-down menu	select the pre-created extent

It is recommended to always associate extents to targets in a one-to-one manner, even though the GUI will allow multiple extents to be associated with the same target.

Note: Each LUN entry has Edit and Delete buttons for modifying the settings or deleting the LUN entirely. A verification popup appears when the Delete button is clicked. If an initiator has an active connection to the LUN, it is indicated in red text. It is recommended to clear initiator connections to a LUN before deleting it.

After iSCSI has been configured, remember to start it in `Services → Control Services`. Click the red OFF button next to iSCSI. After a second or so, it will change to a blue ON, indicating that the service has started.

9.5.8 Fibre Channel Ports

If the TrueNAS® system has Fibre Channel ports, `Sharing → Block (iSCSI)` will appear as `Sharing → Block (iSCSI/FC)` and an extra Fibre Channel Ports tab is added. An example is shown in Figure 9.27.

Sharing

Apple (AFP)UNIX (NFS)WebDAVWindows (CIFS)Block (iSCSI/FC)

Target Global ConfigurationPortals (iSCSI)Initiators (iSCSI)Authorized Access (iSCSI)TargetsExtentsAssociated TargetsFibre Channel Ports

Base Name:

iqn.2005-10.org.freenas.ctl

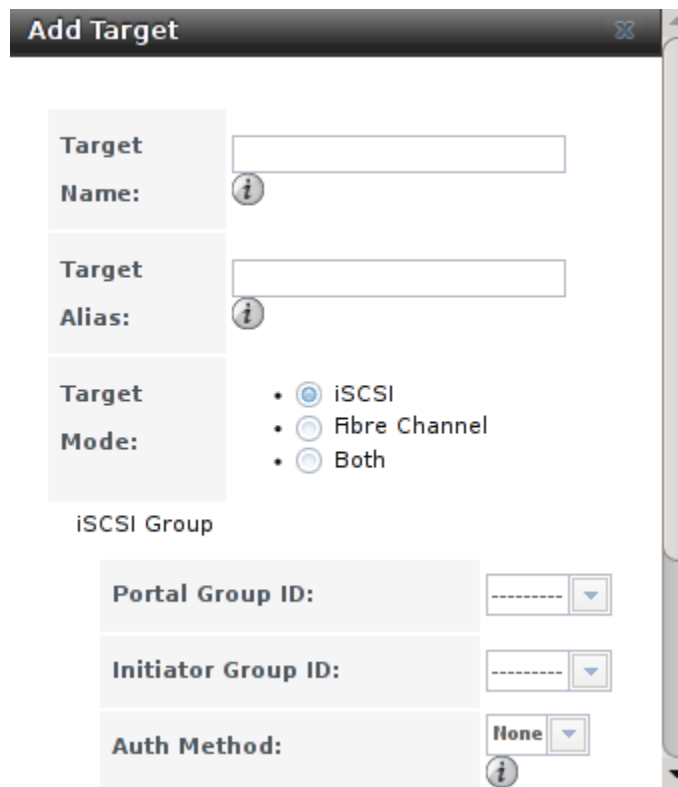
ISNS Servers:

Pool Available Space Threshold (%):

Save

Fig. 9.27: Block (iSCSI) Screen

Otherwise, the Target Global Configuration screen is the same as described in [Target Global Configuration](#) (page 171). 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 seen in [Figure 9.28](#), the `Targets → Add Target` screen has an extra Target Mode option for indicating whether the target to create is iSCSI, Fibre Channel, or both.



Add Target

Target Name: ⓘ

Target Alias: ⓘ

Target Mode:

- ☒ iSCSI
- ☐ Fibre Channel
- ☐ Both

iSCSI Group

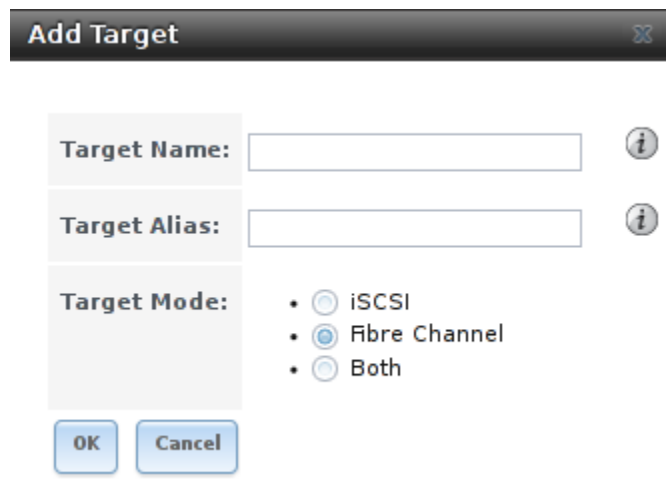
Portal Group ID: ▾

Initiator Group ID: ▾

Auth Method: None ▾ ⓘ

Fig. 9.28: Add Target Screen

If you select Fibre Channel, this screen will change so only the Target Name and Target Alias fields remain, as those are the only applicable fields for a Fibre Channel connection. An example is shown in [Figure 9.29](#).



Add Target

Target Name: ⓘ

Target Alias: ⓘ

Target Mode:

- ☐ iSCSI
- ☒ Fibre Channel
- ☐ Both

OK **Cancel**

Fig. 9.29: Configuring a Fibre Channel Target

The screens for adding an extent and associating a target are the same as described in [Extents](#) (page 177) and [Target/Extents](#) (page 180).

An example of the Fibre Channel Ports screen is shown in [Figure 9.30](#).

Sharing

Apple (AFP)

UNIX (NFS)

WebDAV

Windows (SMB)

Block (iSCSI/FC)

Target Global Configuration

Portals (iSCSI)

Initiators (iSCSI)

Authorized Access (iSCSI)

Targets

Extents

Associated Targets

Fibre Channel Ports

<div>isp0 - Ready (8 Gbps)</div> <div>WWPN: naa.21000024ff4ce7ea</div> <div> <input type="radio"/> Initiator <input checked="" type="radio"/> Target <div>fc-target</div> <input type="radio"/> Disabled </div>	<div>Connected Initiators</div> <div>- naa.21000024ff5105c1</div> <div>- naa.21000024ff5105c0 (Node B)</div>
<div>isp0/1 - Ready (8 Gbps)</div> <div>WWPN: naa.22000024ff4ce7ea</div> <div> <input type="radio"/> Initiator <input checked="" type="radio"/> Target <div>fc-target2</div> <input type="radio"/> Disabled </div>	<div>Connected Initiators</div> <div>- naa.21000024ff5105c1</div> <div>- naa.21000024ff5105c0 (Node B)</div>
<div>isp0/2 - No Link</div> <div>WWPN: naa.23000024ff4ce7ea</div> <div> <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled </div>	
<div>isp0/3 - No Link</div> <div>WWPN: naa.24000024ff4ce7ea</div> <div> <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled </div>	
<div>isp0/4 - No Link</div> <div>WWPN: naa.25000024ff4ce7ea</div> <div> <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled </div>	

Fig. 9.30: Configuring a Fibre Channel Port

This screen shows the status of each attached fibre channel port, where:

- Initiator: indicates that the port is acting as a client and has access to any physically attached storage.
- Target: indicates that clients are connecting to the specified target through this port.
- Disabled: indicates that this fibre channel port is not in use.

Note: The Target tab of [Reporting](#) (page 248) provides Fibre Channel port bandwidth graphs.

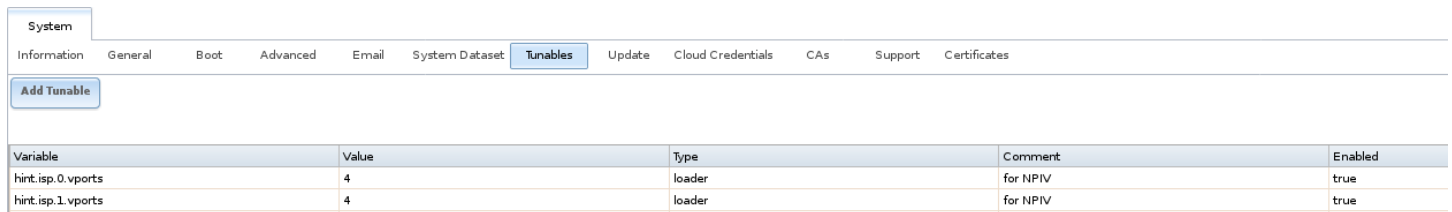
This example has also been configured for NPIV (N_Port ID Virtualization). Note that the physical interface isp0 has two virtual ports (isp0/1 and isp0/2) displayed in [Figure 9.30](#). NPIV allows the administrator to use switch zoning to configure each virtual port as if it was a physical port in order 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 your 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 → Add Tunable` and enter the following:

- 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 9.31](#), 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 248) where its bandwidth usage can be viewed.



Variable	Value	Type	Comment	Enabled
hint.isp.0.vports	4	loader	for NPV	true
hint.isp.1.vports	4	loader	for NPV	true

Fig. 9.31: Adding Virtual Ports

9.5.9 Connecting to iSCSI

To access the iSCSI target, clients must use iSCSI initiator software.

An iSCSI Initiator client is pre-installed with Windows 7. A detailed how-to for this client can be found [here](http://www.windowsnetworking.com/articles-tutorials/windows-7/Connecting-Windows-7-iSCSI-SAN.html) (<http://www.windowsnetworking.com/articles-tutorials/windows-7/Connecting-Windows-7-iSCSI-SAN.html>). A client for Windows 2000, XP, and 2003 can be found [here](http://www.microsoft.com/en-us/download/details.aspx?id=18986) (<http://www.microsoft.com/en-us/download/details.aspx?id=18986>). This [how-to](http://blog.pluralsight.com/freenas-8-iscsi-target-windows-7) (<http://blog.pluralsight.com/freenas-8-iscsi-target-windows-7>) shows how to create an iSCSI target for a Windows 7 system.

Mac OS X does not include an initiator. [globalSAN](http://www.studionetworksolutions.com/globalsan-iscsi-initiator/) (<http://www.studionetworksolutions.com/globalsan-iscsi-initiator/>) is a commercial, easy-to-use Mac initiator.

BSD systems provide command line initiators: [iscontrol\(8\)](http://www.freebsd.org/cgi/man.cgi?query=iscontrol) (<http://www.freebsd.org/cgi/man.cgi?query=iscontrol>) comes with FreeBSD versions 9.x and lower, [iscsictl\(8\)](https://www.freebsd.org/cgi/man.cgi?query=iscsictl) (<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://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/.man8/iscsid.8?query=iscsid) (<http://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/.man8/iscsid.8?query=iscsid>) comes with OpenBSD.

Some Linux distros provide the command line utility `iscsiadm` from [Open-iSCSI](http://www.open-iscsi.com/) (<http://www.open-iscsi.com/>). Use a web search to see if a package exists for your distribution should the command not exist on your Linux system.

If a LUN is added while `iscsiadm` is already connected, it will not see the new LUN until rescanned with `iscsiadm -m node -R`. Alternately, use `iscsiadm -m discovery -t st -p portal_IP` to find the new LUN and `iscsiadm -m node -T LUN_Name -l` to log into the LUN.

Instructions for connecting from a VMware ESXi Server can be found at [How to configure FreeNAS 8 for iSCSI and connect to ESXi\(i\)](http://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/) (<http://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/>). Note that the requirements for booting vSphere 4.x off iSCSI differ between ESX and ESXi. ESX requires a hardware iSCSI adapter while ESXi requires specific iSCSI boot firmware support. The magic is on the booting host side, meaning that there is no difference to the TrueNAS® configuration. See the [iSCSI SAN Configuration Guide](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) (http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) for details.

The VMware firewall only allows iSCSI connections on port 3260 by default. If a different port has been selected, outgoing connections to that port must be manually added to the firewall before those connections will work.

If the target can be seen but does not connect, check the Discovery Auth settings in Target Global Configuration.

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

9.5.10 Growing LUNs

The method used to grow the size of an existing iSCSI LUN depends on whether the LUN is backed by a file extent or a zvol. Both methods are described in this section.

Enlarging a LUN with one of the methods below gives it more unallocated space, but does not automatically resize filesystems or other data on the LUN. This is the same as binary-copying a smaller disk onto a larger one. More space is available on the new disk, but the partitions and filesystems on it must be expanded to use this new space. Resizing virtual disk images is usually done from virtual machine management software. Application software to resize filesystems is dependent on the type of filesystem and client, but is often run from within the virtual machine. For instance, consider a Windows VM with the last partition on the disk holding an NTFS filesystem. The LUN is expanded and the partition table edited to add the new space to the last partition. The Windows disk manager must still be used to resize the NTFS filesystem on that last partition to use the new space.

Zvol Based LUN

To grow a zvol based LUN, go to **Storage** → **Volumes** → **View Volumes**, highlight the zvol to be grown, and click **Edit zvol**. In the example shown in [Figure 9.32](#), the current size of the zvol named zvol1 is 4GB.

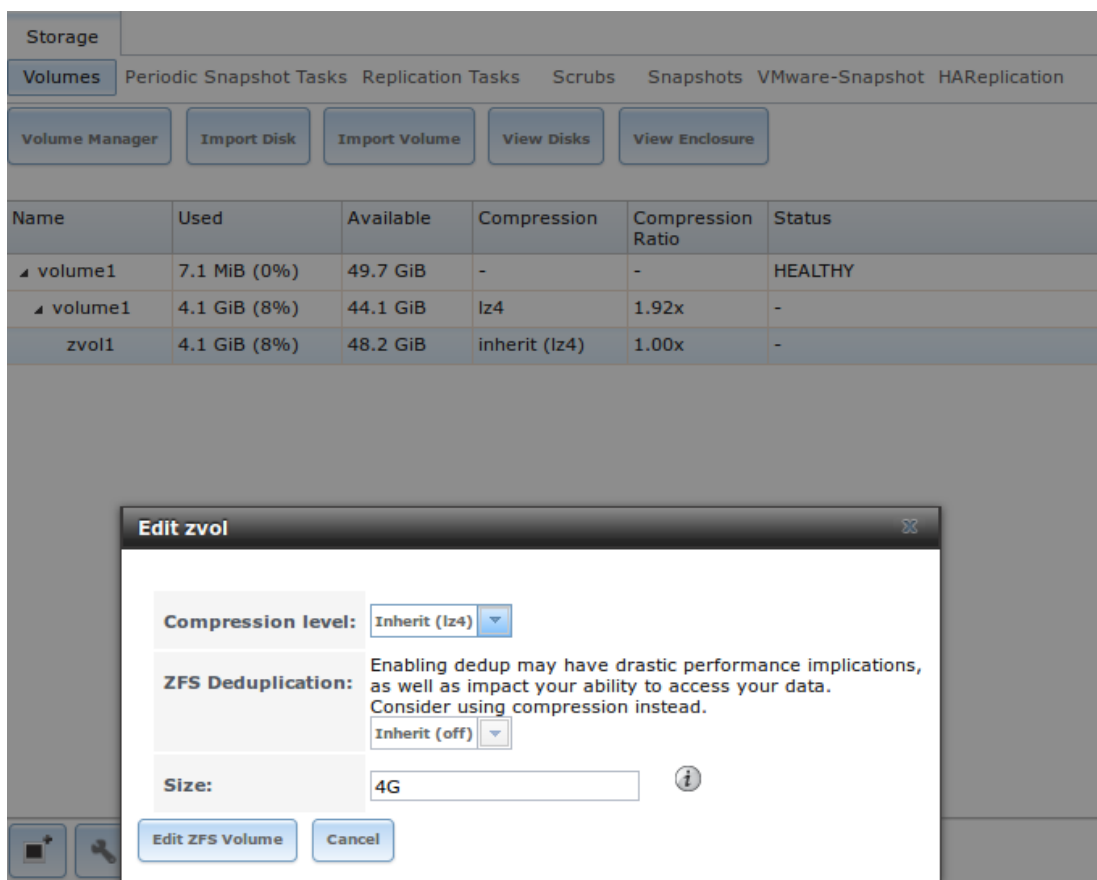


Fig. 9.32: Editing an Existing Zvol

Enter the new size for the zvol in the **Size** field and click **Edit ZFS Volume**. This menu closes and the new size for the zvol is immediately shown in the **Used** column of the **View Volumes** screen.

Note: The GUI does not allow reducing (shrinking) the size of the zvol, as doing so could result in loss of data. It also does not allow increasing the size of the zvol past 80% of the volume size.

File Extent Based LUN

To grow a file extent based LUN, go to `Services → iSCSI → File Extents → View File Extents` to determine the path of the file extent to grow. Open Shell to grow the extent. This example grows `/mnt/volume1/data` by 2 G:

```
truncate -s +2g /mnt/volume1/data
```

Go back to `Services → iSCSI → File Extents → View File Extents` and click the Edit button for the file extent. Set the size to 0 as this causes the iSCSI target to use the new size of the file.

SERVICES

The Services section of the GUI is where various services that ship with the TrueNAS® system are configured, started, or stopped. TrueNAS® includes these built-in services:

- [AFP](#) (page 189)
- [Domain Controller](#) (page 191)
- [Dynamic DNS](#) (page 193)
- [FTP](#) (page 194)
- [iSCSI](#) (page 199)
- [LLDP](#) (page 200)
- [Netdata](#) (page 200)
- [NFS](#) (page 201)
- [Rsync](#) (page 203)
- [S3](#) (page 205)
- [S.M.A.R.T.](#) (page 206)
- [SMB](#) (page 208)
- [SNMP](#) (page 212)
- [SSH](#) (page 213)
- [TFTP](#) (page 216)
- [UPS](#) (page 216)
- [WebDAV](#) (page 220)

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

10.1 Control Services

`Services` → `Control Services`, shown in [Figure 10.1](#), shows which services are currently running and can start, stop, or configure them. The S.M.A.R.T. service is enabled by default, but only runs if the storage devices support [S.M.A.R.T. data](#) (<http://en.wikipedia.org/wiki/S.M.A.R.T.>) Other services default to off until started.

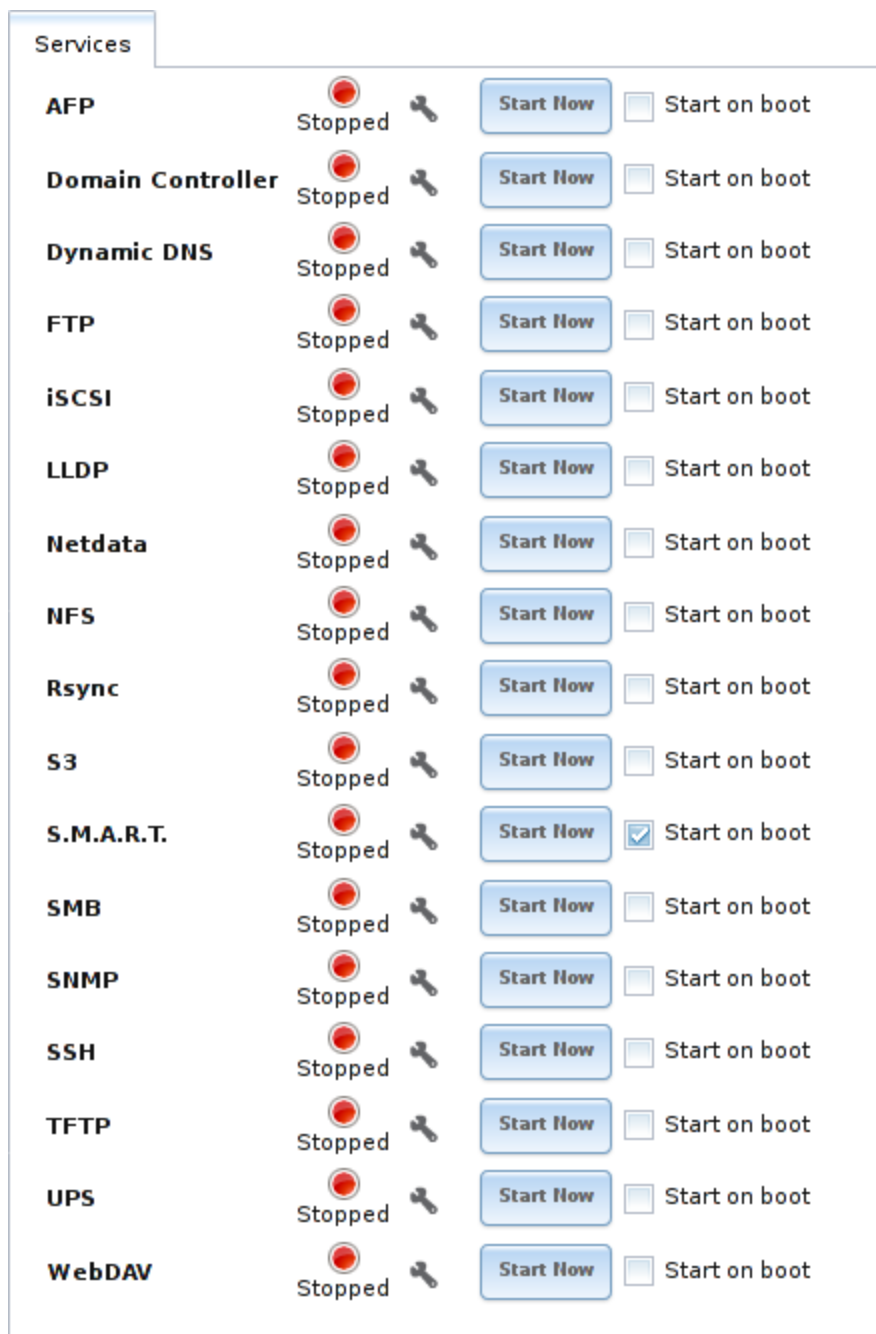


Fig. 10.1: Control Services

Stopped services show a red stop symbol and a Start Now button. Running services show a green light with a Stop Now button.

Tip: Using a proxy server can prevent the list of services from being displayed. If a proxy server is used, configure it to not proxy local network connections or websocket connections. VPN software can also cause problems. If the list of services is displayed when connecting on the local network but not when connecting through the VPN, check the VPN software configuration.

Services are configured by clicking the wrench icon or the name of the service in the Services section of the tree menu.

If a service does not start, go to `System` → `Advanced` and check the box `Show console messages` in the footer. Console

messages appear at the bottom of the browser. Clicking the console message area makes it into a pop-up window, allowing scrolling through or copying the messages. Watch these messages for errors when stopping or starting the problematic service.

To read the system logs for more information about a service failure, open [Shell](#) (page 257) and type `more /var/log/messages`.

10.2 AFP

The settings that are configured when creating AFP Shares in `Sharing → Apple (AFP) Shares → Add Apple (AFP) Share` are specific to each configured AFP Share. In contrast, global settings which apply to all AFP shares are configured in `Services → AFP`.

[Figure 10.2](#) shows the available global AFP configuration options which are described in [Table 10.1](#).

Settings

Guest Access:

☐

Guest account:

nobody

Max. Connections:

50

Enable home directories:

☐

Home directories:

Browse

Home share name:

Home Share Time Machine:

☐

Database Path:

Browse

Global auxiliary parameters:

Map ACLs:

Rights

Chmod Request:

Preserve

Bind IP Addresses:

☐ 10.0.0.102

OK

Cancel

Fig. 10.2: Global AFP Configuration

Table 10.1: Global AFP Configuration Options

Setting	Value	Description
Guest Access	checkbox	if checked, clients will not be prompted to authenticate before accessing AFP shares
Guest account	drop-down menu	select account to use for guest access; the selected account must have permissions to the volume or dataset being shared
Max Connections	integer	maximum number of simultaneous connections
Continued on next page		

Table 10.1 – continued from previous page

Setting	Value	Description
Enable home directories	checkbox	if checked, any user home directories located under Home directories will be available over the share
Home directories	browse button	select the volume or dataset which contains user home directories
Home share name	string	overrides default home folder name with the specified value
Home Share Time Machine	checkbox	when checked, enables Time Machine Lock Stealing; Apple recommends that shares designated for Time Machine backups be used exclusively for Time Machine backups
Database Path	browse button	select the path to store the CNID databases used by AFP (default is the root of the volume); the path must be writable
Global auxiliary parameters	string	additional afp.conf(5) (http://netatalk.sourceforge.net/3.0/html/docs/afp.conf.5.html) parameters not covered elsewhere in this screen
Map ACLs	drop-down menu	choose mapping of effective permissions for authenticated users; Rights (default, Unix-style permissions), Mode (ACLs), or None
Chmod Request	drop-down menu	indicates how to handle ACLs; choices are Ignore, Preserve, or Simple
Bind IP Addresses	selection	specify the IP addresses to listen for FTP connections; highlight the desired IP addresses in the Available list and use the >> button to add to the Selected list

When configuring home directories, it is recommended to create a dataset to hold the home directories which contains a child dataset for each user. As an example, create a dataset named `volume1/homedirs` and browse to this dataset when configuring the Home directories field of the AFP service. Then, as you create each user, first create a child dataset for that user. For example, create a dataset named `volume1/homedirs/user1`. When you create the user1 user, browse to the `volume1/homedirs/user1` dataset in the Home Directory field of the Add New User screen.

10.2.1 Troubleshooting AFP

You can determine which users are connected to an AFP share by typing `afpusers`.

If Something wrong with the volume's CNID DB is shown, run this command from [Shell](#) (page 257), replacing the path to the problematic AFP share:

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

This command may take a while, depending upon the size of the volume or dataset being shared. This command will wipe the CNID database and rebuild it from the CNIDs stored in the AppleDouble files.

10.3 Domain Controller

TrueNAS® can be configured to act either as the domain controller for a network or to join an existing [Active Directory](#) (page 132) network as a domain controller.

Note: This section demonstrates how to configure the TrueNAS® system to act as a domain controller. If the goal is to integrate with an existing [Active Directory](#) (page 132) network to access its authentication and authorization services, configure [Active Directory](#) (page 132) instead.

Be aware that configuring a domain controller is a complex process that requires a good understanding of how [Active Directory](#) (page 132) works. While `Services → Domain Controller` makes it easy to enter the needed settings into the administrative graphical interface, it is important to understand what those settings should be. Before beginning configuration, read through the [Samba AD DC HOWTO](https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO) (https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO). After TrueNAS®

is configured, use the RSAT utility from a Windows system to manage the domain controller. The Samba AD DC HOWTO includes instructions for installing and configuring RSAT.

Figure 10.3 shows the configuration screen for creating a domain controller and Table 10.2 summarizes the available options.

Settings

Realm:

Domain:

Server Role:

active directory domain controller

DNS Forwarder:

Domain Forest Level:

2003

Administrator Password:

Confirm Administrator Password:

Kerberos Realm:

OK

Cancel

Delete

Fig. 10.3: Domain Controller Settings

Table 10.2: Domain Controller Configuration Options

Setting	Value	Description
Realm	string	capitalized DNS realm name
Domain	string	capitalized domain name
Server Role	drop-down menu	at this time, the only supported role is as the domain controller for a new domain
DNS Forwarder	string	IP address of DNS forwarder; required for recursive queries when SAMBA_INTERNAL is selected
Domain Forest Level	drop-down menu	choices are 2000, 2003, 2008, 2008_R2, 2012, or 2012_R2; refer to Understanding Active Directory Domain Services (AD DS) Functional Levels (https://technet.microsoft.com/en-us/library/understanding-active-directory-functional-levels(Ws.10).aspx) for details
Administrator password	string	password to be used for the Active Directory (page 132) administrator account
Kerberos Realm	drop-down menu	auto-populates with information from the Realm when the settings in this screen are saved

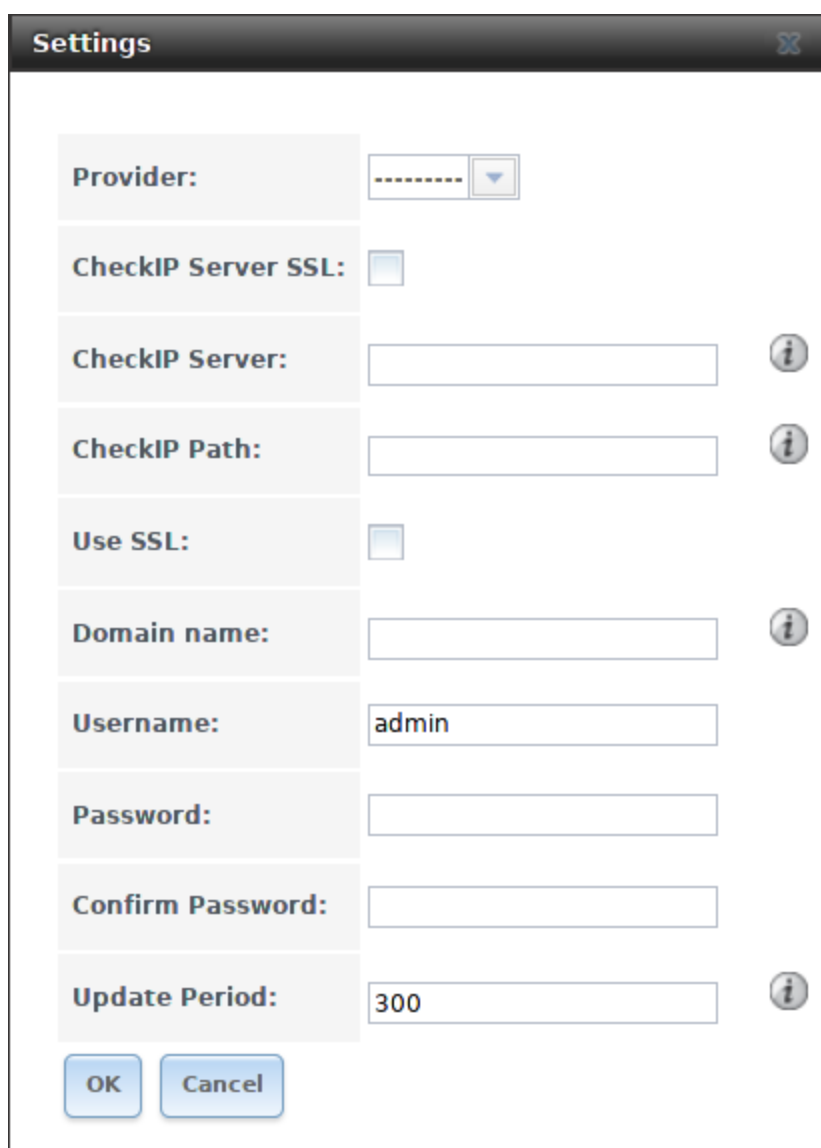
10.3.1 Samba Domain Controller Backup

A `samba_backup` script is available to back up Samba4 domain controller settings is available. From the [Shell](#) (page 257), run `/usr/local/bin/samba_backup --usage` to show the input options.

10.4 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 you to access the TrueNAS[®] system even if the IP address changes. DDNS requires you to register with a DDNS service such as [DynDNS](http://dyn.com/dns/) (<http://dyn.com/dns/>).

[Figure 10.4](#) shows the DDNS configuration screen and [Table 10.3](#) summarizes the configuration options. The values for these fields are provided by the DDNS provider. After configuring DDNS, remember to start the DDNS service in [Services](#) → [Control Services](#).



The screenshot shows a 'Settings' window for DDNS configuration. It contains the following fields and controls:

- Provider:** A dropdown menu with a dashed line and a downward arrow.
- CheckIP Server SSL:** A checkbox.
- CheckIP Server:** A text input field with an information icon (i) to its right.
- CheckIP Path:** A text input field with an information icon (i) to its right.
- Use SSL:** A checkbox.
- Domain name:** A text input field with an information icon (i) to its right.
- Username:** A text input field containing the value 'admin'.
- Password:** A text input field.
- Confirm Password:** A text input field.
- Update Period:** A text input field containing the value '300' with an information icon (i) to its right.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom left.

Fig. 10.4: Configuring DDNS

Table 10.3: DDNS Configuration Options

Setting	Value	Description
Provider	drop-down menu	several providers are supported; if your provider is not listed, select Custom Provider and enter the information in the Custom Server and Custom Path fields
CheckIP Server SSL	string	when checked, HTTPS is used for the connection to the CheckIP Server
CheckIP Server	string	enter the name and port of the server that reports the external IP address, in the format server.name.org:port
CheckIP Path	string	enter the path that is requested by the CheckIP Server to determine the user's IP address
Use SSL		when checked, HTTPS is used for the connection to the server that updates the DNS record
Domain name	string	fully qualified domain name (e.g. yourname.dyndns.org)
Username	string	username used to logon to the provider and update the record
Password	string	password used to logon to the provider and update the record
Update period	integer	how often the IP is checked in seconds

When using `he.net`, enter the domain name for Username and enter the DDNS key generated for that domain's A entry at the `he.net` website for Password.

10.5 FTP

TrueNAS® uses the `proftpd` (<http://www.proftpd.org/>) FTP server to provide FTP services. Once the FTP service is configured and started, clients can browse and download data using a web browser or FTP client software. The advantage of FTP is that easy-to-use cross-platform utilities are available to manage uploads to and downloads from the 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 you are concerned about sensitive data, see [Encrypting FTP](#).

This section provides an overview of the FTP configuration options. It then provides examples for configuring anonymous FTP, specified user access within a chroot environment, encrypting FTP connections, and troubleshooting tips.

[Figure 10.5](#) shows the configuration screen for `Services` → `FTP`. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box `Show advanced fields by default` in `System` → `Advanced`.

The screenshot shows a window titled "FTP Settings" with a dark header bar. Below the header, there are several configuration fields and checkboxes. The fields are: "Port:" with a text box containing "21", "Clients:" with a text box containing "5", "Connections:" with a text box containing "2", "Login Attempts:" with a text box containing "1", and "Timeout:" with a text box containing "600". Each of these fields has a small circular icon with an "i" to its right. Below these are three checkboxes: "Allow Root Login:", "Allow Anonymous Login:", and "Allow Local User Login:", all of which are currently unchecked. The "Path:" field is followed by a text box and a "Browse" button. Below that is the "Display Login:" field, which is a larger text box. At the bottom is the "Allow Transfer Resumption:" checkbox, which is also unchecked. A small circular icon with an "i" is located below the "Display Login:" field.

Fig. 10.5: Configuring FTP

Table 10.4 summarizes the available options when configuring the FTP server.

Table 10.4: FTP Configuration Options

Setting	Value	Advanced Mode	Description
Port	integer		port the FTP service listens on
Clients	integer		maximum number of simultaneous clients
Connections	integer		maximum number of connections per IP address where 0 means unlimited
Login Attempts	integer		maximum number of attempts before client is disconnected; increase this if users are prone to typos
Timeout	integer		maximum client idle time in seconds before client is disconnected
Allow Root Login	checkbox		discouraged as increases security risk
Allow Anonymous Login	checkbox		enables anonymous FTP logins with access to the directory specified in Path
Path	browse button		root directory for anonymous FTP connections
Allow Local User Login	checkbox		required if Anonymous Login is disabled
Display Login	string		message displayed to local login users after authentication; not displayed to anonymous login users
File Permission	checkboxes	✓	sets default permissions for newly created files
Directory Permission	checkboxes	✓	sets default permissions for newly created directories

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Enable FXP (https://en.wikipedia.org/wiki/File_eXchange_Protocol)	checkbox	✓	enables File eXchange Protocol which is discouraged as it makes the server vulnerable to FTP bounce attacks
Allow Transfer Resumption	checkbox		allows FTP clients to resume interrupted transfers
Always Chroot	checkbox		a local user is only allowed access to their home directory unless the user is a member of group wheel
Require IDENT Authentication	checkbox	✓	will result in timeouts if <code>identd</code> is not running on the client
Perform Reverse DNS Lookups	checkbox		perform reverse DNS lookups on client IPs; can cause long delays if reverse DNS is not configured
Masquerade address	string		public IP address or hostname; set if FTP clients cannot connect through a NAT device
Minimum passive port	integer	✓	used by clients in PASV mode, default of 0 means any port above 1023
Maximum passive port	integer	✓	used by clients in PASV mode, default of 0 means any port above 1023
Local user upload bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Local user download bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Anonymous user upload bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Anonymous user download bandwidth	integer	✓	in KB/s, default of 0 means unlimited
Enable TLS	checkbox	✓	enables encrypted connections and requires a certificate to be created or imported using Certificates (page 48)
TLS policy	drop-down menu	✓	the selected policy defines whether the control channel, data channel, both channels, or neither channel of an FTP session must occur over SSL/TLS; the policies are described here (http://www.proftpd.org/docs/directives/linked/config_ref_TLSRequired.h)
TLS allow client renegotiations	checkbox	✓	checking this box is not recommended as it breaks several security measures; for this and the rest of the TLS fields, refer to mod_tls (http://www.proftpd.org/docs/contrib/mod_tls.html) for more details
TLS allow dot login	checkbox	✓	if checked, the user's home directory is checked for a <code>.tlslogin</code> file which contains one or more PEM-encoded certificates; if not found, the user is prompted for password authentication
TLS allow per user	checkbox	✓	if checked, the user's password may be sent unencrypted
TLS common name required	checkbox	✓	if checked, the common name in the certificate must match the FQDN of the host
TLS enable diagnostics	checkbox	✓	if checked when troubleshooting a connection, logs more verbosely
TLS export certificate data	checkbox	✓	if checked, exports the certificate environment variables
TLS no certificate request	checkbox	✓	try checking this box if the client cannot connect and it is suspected that the client software is not properly handling the server's certificate request
TLS no empty fragments	checkbox	✓	checking this box is not recommended as it bypasses a security mechanism
TLS no session reuse required	checkbox	✓	checking this box reduces the security of the connection, so only use it if the client does not understand reused SSL sessions

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
TLS export standard vars	checkbox	✓	if checked, sets several environment variables
TLS DNS name required	checkbox	✓	if checked, the client's DNS name must resolve to its IP address and the cert must contain the same DNS name
TLS IP address required	checkbox	✓	if checked, the client's certificate must contain the IP address that matches the IP address of the client
Certificate	drop-down menu		the SSL certificate to be used for TLS FTP connections; to create a certificate, use System → Certificates
Auxiliary parameters	string	✓	used to add proftpd(8) (http://linux.die.net/man/8/proftpd) parameters not covered elsewhere in this screen

This example demonstrates the auxiliary parameters that prevent all users from performing the FTP DELETE command:

```
<Limit DELE>
DenyAll
</Limit>
```

10.5.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the TrueNAS® system is not accessible from the Internet and everyone in your internal network needs easy access to the stored data. Anonymous FTP does not require you to create a user account for every user. In addition, passwords are not required so it is not necessary to manage changed passwords on the TrueNAS® system.

To configure anonymous FTP:

1. Give the built-in ftp user account permissions to the volume/dataset to be shared in [Storage → Volumes](#) as follows:
 - Owner(user): select the built-in ftp user from the drop-down menu
 - Owner(group): select the built-in ftp group from the drop-down menu
 - Mode: review that the permissions are appropriate for the share

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means that you always use Unix ACLs, even if Windows clients will be accessing TrueNAS® via FTP.

2. Configure anonymous FTP in [Services → FTP](#) by setting the following attributes:
 - check the box Allow Anonymous Login
 - Path: browse to the volume/dataset/directory to be shared
3. Start the FTP service in [Services → Control Services](#). Click the Start Now button next to FTP. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to Stop Now.
4. Test the connection from a client using a utility such as [Filezilla](https://filezilla-project.org/) (<https://filezilla-project.org/>).

In the example shown in [Figure 10.6](#), the user has enter the following 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. 10.6: Connecting Using Filezilla

The messages within the client indicate that the FTP connection is successful. The user can now navigate the contents of the root folder on the remote site—this is the volume/dataset that was specified in the FTP service configuration. The user can also transfer files between the local site (their system) and the remote site (the TrueNAS® system).

10.5.2 FTP in chroot

If you require your users to authenticate before accessing the data on the TrueNAS® system, you will need to either create a user account for each user or import existing user accounts using [Active Directory](#) (page 132) or LDAP. If you then create a ZFS dataset for each user, you can chroot each user so that they are limited to the contents of their own home directory. Datasets provide the added benefit of configuring a quota so that the size of the user's home directory is limited to the size of the quota.

To configure this scenario:

1. Create a ZFS dataset for each user in **Storage** → **Volumes**. Click an existing ZFS volume → **Create ZFS Dataset** and set an appropriate quota for each dataset. Repeat this process to create a dataset for every user that needs access to the FTP service.
2. If you are not using AD or LDAP, create a user account for each user in **Account** → **Users** → **Add User**. For each user, browse to the dataset created for that user in the **Home Directory** field. Repeat this process to create a user account for every user that needs access to the FTP service, making sure to assign each user their own dataset.
3. Set the permissions for each dataset in **Storage** → **Volumes**. Click the **Change Permissions** button for a dataset to assign a user account as Owner of that dataset and to set the desired permissions for that user. Repeat for each dataset.

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means that you always use Unix ACLs, even if Windows clients will be accessing TrueNAS® via FTP.

4. Configure FTP in **Services** → **FTP** with these attributes:
 - Path: browse to the parent volume containing the datasets
 - make sure the boxes for **Allow Anonymous Login** and **Allow Root Login** are unchecked
 - check the box **Allow Local User Login**
 - check the box **Always Chroot**

5. Start the FTP service in *Services* → *Control Services*. Click the Start Now button next to FTP. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to Stop Now.
6. Test the connection from a client using a utility such as Filezilla.

To test this configuration in Filezilla, use the IP address of the TrueNAS® system, the Username of a user that has been associated with a dataset, and the Password for that user. The messages should indicate that the authorization and the FTP connection are successful. The user can now navigate the contents of the root folder on the remote site—this time it is not the entire volume but the dataset that was created for that user. The user should be able to transfer files between the local site (their system) and the remote site (their dataset on the TrueNAS® system).

10.5.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

1. Import or create a certificate authority using the instructions in [CAs](#) (page 45). Then, import or create the certificate to use for encrypted connections using the instructions in [Certificates](#) (page 48).
2. In *Services* → *FTP*, check the box Enable TLS and select the certificate in the Certificate drop-down menu.
3. Specify secure FTP when accessing the 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 OK to accept the certificate and negotiate an encrypted connection.
4. To force encrypted connections, select on for the TLS Policy.

10.5.4 Troubleshooting FTP

The FTP service will not start if it cannot resolve the system's hostname to an IP address using DNS. To see if the FTP service is running, open [Shell](#) (page 257) 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*, check the box Show console messages in the footer and click Save. Next, go to *Services* → *Control Services* and switch the FTP service off, then back on. Watch the console messages at the bottom of the browser for errors.

If the error refers to DNS, either create an entry in the local DNS server with the TrueNAS® system's hostname and IP address or add an entry for the IP address of the TrueNAS® system in the Host name database field of *Network* → *Global Configuration*.

10.6 iSCSI

Refer to [Block \(iSCSI\)](#) (page 169) for instructions on configuring iSCSI. To start the iSCSI service, click its entry in *Services*.

Note: A warning message is shown if you stop the iSCSI service when initiators are connected. Type `ctladm islist` to determine the names of the connected initiators.

10.7 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) (<https://github.com/sspan/ladvd>) LLDP implementation. If your network contains managed switches, configuring and starting the LLDP service will tell the TrueNAS® system to advertise itself on the network.

Figure 10.7 shows the LLDP configuration screen and Table 10.5 summarizes the configuration options for the LLDP service.




Fig. 10.7: Configuring LLDP

Table 10.5: LLDP Configuration Options

Setting	Value	Description
Interface De- scription	checkbox	when checked, receive mode is enabled and received peer information is saved in interface descriptions
Country Code	string	required for LLDP location support; enter a two-letter ISO 3166 country code
Location	string	optional; specify the physical location of the host

10.8 Netdata

Netdata is a real-time performance and monitoring system. It displays data as web dashboards.

Start the Netdata service from the [Services](#) (page 187) screen. Click the wrench icon to display the Netdata settings dialog shown in Figure 10.8.




Fig. 10.8: Netdata Settings Dialog

Click the Take me to the Netdata UI button to view the web dashboard as shown in Figure 10.9.




Fig. 10.9: Netdata Web Dashboard

More information on configuring and using Netdata is available at the [Netdata website \(https://my-netdata.io/\)](https://my-netdata.io/).

10.9 NFS

The settings that are configured when creating NFS Shares in **Sharing → Unix (NFS) Shares → Add Unix (NFS) Share** are specific to each configured NFS Share. In contrast, global settings which apply to all NFS shares are configured in **Services → NFS**.

VAAI for NAS is supported through the NFS service. See [VAAI for NAS](#) (page 311) for more details.

Figure 10.10 shows the configuration screen and Table 10.6 summarizes the configuration options for the NFS service.

Settings

Number of servers:

4

Serve UDP NFS clients:

Bind IP Addresses:

10.0.0.142

Allow non-root mount:

Enable NFSv4:

NFSv3 ownership model for NFSv4:

Require Kerberos for NFSv4:

mountd(8) bind port:

rpc.statd(8) bind port:

rpc.lockd(8) bind port:

Support >16 groups:

Log mountd(8) requests:

Log rpc.statd(8) and rpc.lockd(8):

OK

Cancel

Fig. 10.10: Configuring NFS

Table 10.6: NFS Configuration Options

Setting	Value	Description
Number of servers	integer	the number of servers can be increased if NFS client responses are slow; to limit CPU context switching, keep this number less than or equal to the number of CPUs reported by <code>sysctl -n kern.smp.cpus</code> .
Serve UDP NFS clients	checkbox	check if NFS clients need to use UDP
Bind IP Addresses	checkboxes	IP addresses to listen on for NFS requests; when unchecked, NFS listens on all available addresses

Continued on next page

Table 10.6 – continued from previous page

Setting	Value	Description
Allow non-root mount	checkbox	check this box only if the NFS client requires it
Enable NFSv4	checkbox	NFSv3 is the default, check this box to switch to NFSv4
NFSv3 owner-ship model for NFSv4	checkbox	grayed out unless Enable NFSv4 is checked and, in turn, grays out Support>16 groups which is incompatible; check this box if NFSv4 ACL support is needed without requiring the client and the server to sync users and groups
Require Kerberos for NFSv4	checkbox	when checked, NFS shares will fail if the Kerberos ticket is unavailable
mountd(8) bind port	integer	optional; specify port that mountd(8) (http://www.freebsd.org/cgi/man.cgi?query=mountd) binds to
rpc.statd(8) bind port	integer	optional; specify port that rpc.statd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.statd) binds to
rpc.lockd(8) bind port	integer	optional; specify port that rpc.lockd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) binds to
Support>16 groups	checkbox	check this box if any users are members of more than 16 groups (useful in AD environments); note that this assumes that group membership has been configured correctly on the NFS server
Log mountd(8) requests	checkbox	enable logging of mountd(8) (http://www.freebsd.org/cgi/man.cgi?query=mountd) requests by syslog
Log rpc.statd(8) and rpc.lockd(8)	checkbox	enable logging of rpc.statd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.statd) and rpc.lockd(8) (http://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) requests by syslog

Note: NFSv4 sets all ownership to nobody:nobody if user and group do not match on client and server.

10.10 Rsync

`Services` → `Rsync` is used to configure an rsync server when using rsync module mode. Refer to [Rsync Module Mode](#) (page 70) for a configuration example.

This section describes the configurable options for the `rsyncd` service and rsync modules.

10.10.1 Configure Rsyncd

[Figure 10.11](#) shows the rsyncd configuration screen which is accessed from `Services` → `Rsync` → `Configure Rsyncd`.

Configure Rsyncd

TCP Port

873

Auxiliary parameters

OK

Cancel

Fig. 10.11: Rsyncd Configuration

Table 10.7 summarizes the options that can be configured for the rsync daemon:

Table 10.7: Rsyncd Configuration Options

Setting	Value	Description
TCP Port	integer	port for <code>rsyncd</code> to listen on, default is 873
Auxiliary parameters	string	additional parameters from rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html)

10.10.2 Rsync Modules

Figure 10.12 shows the configuration screen that appears after clicking `Services → Rsync → Rsync Modules → Add Rsync Module`.

Table 10.8 summarizes the options that can be configured when creating a rsync module.

Fig. 10.12: Adding an Rsync Module

Table 10.8: Rsync Module Configuration Options

Setting	Value	Description
Module name	string	mandatory; needs to match the setting on the rsync client
Comment	string	optional description
Path	browse button	volume/dataset to hold received data
Access Mode	drop-down menu	choices are Read and Write, Read-only, or Write-only
Maximum connections	integer	0 is unlimited
User	drop-down menu	select user that file transfers to and from that module should take place as
Group	drop-down menu	select group that file transfers to and from that module should take place as
Hosts allow	string	see rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html) for allowed formats
Hosts deny	string	see rsyncd.conf(5) for allowed formats
Auxiliary parameters	string	additional parameters from rsyncd.conf(5)

10.11 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 10.13 shows the S3 service configuration screen and Table 10.9 summarizes the configuration options. After configuring the S3 service, start it in *Services* → *Control Services*.

Settings

IP Address:

0.0.0.0

i

Port:

9000

i

Access key of 5 to 20 characters in length:

i

Secret key of 8 to 40 characters in length:

i

Confirm S3 Key:

Disks:

Browse

i

Certificate:

Enable Browser:

i

OK

Cancel

Fig. 10.13: Configuring S3

Table 10.9: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	the IP address on which to run the S3 service; 0.0.0.0 sets the server to listen on all addresses
Port	string	TCP port on which to provide the S3 service (default 9000)
Access Key	string	the S3 user name
Secret Key	string	the password to be used by connecting S3 systems; must be at least 8 but no more than 40 characters long
Confirm S3 Key	string	re-enter the S3 password to confirm
Disks	string	S3 filesystem directory
Certificate	drop-down menu	the SSL certificate to be used for secure S3 connections; to create a certificate, use <i>System</i> → <i>Certificates</i>
Enable Browser	checkbox	Enable the web user interface for the S3 service

10.12 S.M.A.R.T.

S.M.A.R.T., or *Self-Monitoring, Analysis, and Reporting Technology* (<http://en.wikipedia.org/wiki/S.M.A.R.T.>), is an industry standard for disk monitoring and testing. Drives can be monitored for status and problems, and several types of self-tests can be run to check the drive health.

Tests run internally on the drive. Most tests can run at the same time as normal disk usage. However, a running test can greatly reduce drive performance, so they should be scheduled at times when the system is not busy or in normal use. It is very important to avoid scheduling disk-intensive tests at the same time. For example, do not schedule S.M.A.R.T. tests to run at the same time, or preferably, even on the same days as [Scrubs](#) (page 126).

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\(8\)](http://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in) (<http://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in>) service to monitor S.M.A.R.T. information. A complete configuration consists of:

- 1. Scheduling when S.M.A.R.T. tests are run in [Tasks](#) → [S.M.A.R.T. Tests](#) → [Add S.M.A.R.T. Test](#).
- 2. Enabling or disabling S.M.A.R.T. for each disk member of a volume in [Volumes](#) → [View Disks](#). This setting is enabled by default for disks that support S.M.A.R.T.
- 3. Checking the configuration of the S.M.A.R.T. service as described in this section.
- 4. Starting the S.M.A.R.T. service with [Services](#) → [Control Services](#).

[Figure 10.14](#) shows the configuration screen that appears after clicking [Services](#) → [S.M.A.R.T.](#)

S.M.A.R.T. Settings

Check interval:

30

i

Power mode:

Never - Check the device

Difference:

0

i

Informational:

0

i

Critical:

0

i

Email to report:

i

OK

Cancel

Fig. 10.14: S.M.A.R.T Configuration Options

Note: `smartd` wakes up at the configured Check Interval. It checks the times configured in [Tasks](#) → [S.M.A.R.T. Tests](#) to see whether tests should be run. Since the smallest time increment for a test is an hour (60 minutes), it does not make sense to set a Check Interval value higher than 60 minutes. For example, if the Check Interval is set to 120 minutes and the smart test to every hour, the test will only be run every two hours because `smartd` only wakes up every two hours.

[Table 10.10](#) summarizes the options in the S.M.A.R.T configuration screen.

Table 10.10: S.M.A.R.T Configuration Options

Setting	Value	Description
Check interval	integer	in minutes, how often <code>smartd</code> wakes up to check if any tests have been configured to run
Power mode	drop-down menu	tests are not performed if the system enters the specified power mode; choices are: Never, Sleep, Standby, or Idle

Continued on next page

Table 10.10 – continued from previous page

Setting	Value	Description
Difference	integer in degrees Celsius	default of 0 disables this check, otherwise reports if the temperature of a drive has changed by N degrees Celsius since last report
Informational	integer in degrees Celsius	default of 0 disables this check, otherwise will message with a log level of LOG_INFO if the temperature is higher than specified degrees in Celsius
Critical	integer in degrees Celsius	default of 0 disables this check, otherwise will message with a log level of LOG_CRIT and send an email if the temperature is higher than specified degrees in Celsius
Email to report	string	email address of person or alias to receive S.M.A.R.T. alerts

10.13 SMB

The settings that are configured when creating SMB Shares in *Sharing* → *Windows (SMB) Shares* → *Add Windows (SMB) Share* are specific to each configured SMB Share. In contrast, global settings which apply to all SMB shares are configured in *Services* → *SMB*.

Note: After starting the SMB service, it can take several minutes for the *master browser election* (<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/NetworkBrowsing.html#id2581357>) to occur and for the TrueNAS® system to become available in Windows Explorer.

Figure 10.15 shows the global SMB configuration options which are described in Table 10.11. This configuration screen is really a front-end to *smb4.conf* (<https://www.freebsd.org/cgi/man.cgi?query=smb4.conf&manpath=FreeBSD+11.0-RELEASE+and+Ports>).

SMB

NetBIOS name:

truenas

NetBIOS alias:

Workgroup:

WORKGROUP

i

Description:

TrueNAS Server

i

DOS charset:

CP437

UNIX charset:

UTF-8

Log level:

Minimum

Use syslog only:

☐

Local Master:

☒

Domain logons:

☐

Time Server for Domain:

☒

Guest account:

nobody

i

Fig. 10.15: Global SMB Configuration

Table 10.11: Global SMB Configuration Options

Setting	Value	Description
NetBIOS Name (This Node)	string	automatically populated with the system's original hostname; limited to 15 characters; it must be different from the Workgroup name
NetBIOS Name (Node B)	string	limited to 15 characters; when using Failover (page 55), set a unique Net-BIOS name for the standby node
NetBIOS Alias	string	limited to 15 characters; when using Failover (page 55), this is the NetBIOS name that resolves to either node
Workgroup	string	must match Windows workgroup name; this setting is ignored if the Active Directory (page 132) or LDAP (page 137) service is running
Description	string	optional
Continued on next page		

Table 10.11 – continued from previous page

Setting	Value	Description
DOS charset	drop-down menu	the character set Samba uses when communicating with DOS and Windows 9x/ME clients; default is CP437
UNIX charset	drop-down menu	default is UTF-8 which supports all characters in all languages
Log level	drop-down menu	choices are Minimum, Normal, or Debug
Use syslog only	checkbox	when checked, authentication failures are logged to <code>/var/log/messages</code> instead of the default of <code>/var/log/samba4/log.smbd</code>
Local Master	checkbox	determines whether or not the system participates in a browser election; should be disabled when network contains an AD or LDAP server and is not necessary if Vista or Windows 7 machines are present
Domain logons	checkbox	only check if need to provide the netlogin service for older Windows clients
Time Server for Domain	checkbox	determines whether or not the system advertises itself as a time server to Windows clients; should be disabled when network contains an AD or LDAP server
Guest Account	drop-down menu	account to be used for guest access; default is nobody; account must have permission to access the shared volume/dataset; if Guest Account user is deleted, resets to nobody
File mask	integer	overrides default file creation mask of 0666 which creates files with read and write access for everybody
Directory mask	integer	overrides default directory creation mask of 0777 which grants directory read, write and execute access for everybody
Allow Empty Password	checkbox	if checked, users can just press <code>Enter</code> when prompted for a password; requires that the username/password be the same as the Windows user account
Auxiliary parameters	string	<code>smb.conf</code> options not covered elsewhere in this screen; see the Samba Guide (http://www.oreilly.com/openbook/samba/book/appb_02.html) for additional settings
Unix Extensions	checkbox	allows non-Windows SMB clients to access symbolic links and hard links, has no effect on Windows clients
Zeroconf share discovery	checkbox	enable if Mac clients will be connecting to the SMB share
Hostname lookups	checkbox	allows using hostnames rather than IP addresses in the Hosts Allow or Hosts Deny fields of a SMB share; uncheck if IP addresses are used to avoid the delay of a host lookup
Allow execute always	checkbox	if checked, Samba will allow the user to execute a file, even if that user's permissions are not set to execute
Obey pam restrictions	checkbox	uncheck this box to allow cross-domain authentication, to allow users and groups to be managed on another forest, or to allow permissions to be delegated from Active Directory (page 132) users and groups to domain admins on another forest
NTLMv1 auth	checkbox	when checked, allow NTLMv1 authentication, required by Windows XP clients and sometimes by clients in later versions of Windows
Bind IP Addresses	checkboxes	check the IP addresses on which SMB should listen
Idmap Range Low	integer	the beginning UID/GID for which this system is authoritative; any UID/GID lower than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs
Idmap Range High	integer	the ending UID/GID for which this system is authoritative; any UID/GID higher than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs

Changes to SMB settings take effect immediately. Changes to share settings only take effect after the client and server negotiate a new session.

Note: Do not set the directory name cache size as an Auxiliary parameter. Due to differences in how Linux and BSD handle file descriptors, directory name caching is disabled on BSD systems to improve performance.

Note: [SMB](#) (page 208) cannot be disabled while [Active Directory](#) (page 132) is enabled.

10.13.1 Troubleshooting SMB

Windows automatically caches file sharing information. If changes are made to an SMB share or to the permissions of a volume/dataset being shared by SMB and the share becomes inaccessible, try logging out and back into the Windows system. Alternately, users can type `net use /delete` from the command line to clear their SMB sessions.

Windows also automatically caches login information. To require users to log in every time access is required, reduce the cache settings on the client computers.

Where possible, avoid using a mix of case in filenames as this can cause confusion for Windows users. [Representing and resolving filenames with Samba](#) (http://www.oreilly.com/openbook/samba/book/ch05_04.html) explains in more detail.

If a particular user cannot connect to a SMB share, make sure that their password does not contain the `?` character. If it does, have the user change the password and try again.

If permissions work for Windows users but not for OS X users, try disabling Unix Extensions and restarting the SMB service.

If the SMB service will not start, run this command from [Shell](#) (page 257) to see if there is an error in the configuration:

```
testparm /usr/local/etc/smb4.conf
```

If clients have problems connecting to the SMB share, go to `Services` → `SMB` and verify that Server maximum protocol is set to SMB2.

It is recommended to use a dataset for SMB sharing. When creating the dataset, make sure that the Share type is set to Windows.

Do not use `chmod` to attempt to fix the permissions on a SMB share as it destroys the Windows ACLs. The correct way to manage permissions on a SMB share is to manage the share security from a Windows system as either the owner of the share or a member of the group that owns the share. To do so, right-click on the share, click Properties and navigate to the Security tab. If you already destroyed the ACLs using `chmod`, `winacl` can be used to fix them. Type `winacl` from [Shell](#) (page 257) for usage instructions.

The [Common Errors](#) (<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/domain-member.html#id2573692>) section of the Samba documentation contains additional troubleshooting tips.

The Samba [Performance Tuning](#) (https://wiki.samba.org/index.php/Performance_Tuning) page describes options to improve performance.

Directory listing speed in folders with a large number of files is sometimes a problem. A few specific changes can help improve the performance. However, changing these settings can affect other usage. In general, the defaults are adequate. Do not change these settings unless there is a specific need.

- Hostname Lookups and Log Level can also have a performance penalty. When not needed, they can be disabled or reduced in the [global SMB service options](#) (page 209).
- Make Samba datasets case insensitive by setting Case Sensitivity to Insensitive when creating them. This ZFS property is only available when creating a dataset. It cannot be changed on an existing dataset. To convert such datasets, back up the data, create a new case-insensitive dataset, create an SMB share on it, set the share level auxiliary parameter `case sensitive = true`, then copy the data from the old one onto it. After the data has been checked and verified on the new share, the old one can be deleted.
- If present, remove options for extended attributes and DOS attributes in the share's [Auxiliary Parameters](#) (page 160).
- Disable as many VFS Objects as possible in the [share settings](#) (page 160). Many have performance overhead.

10.14 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/) (<http://net-snmp.sourceforge.net/>) to provide SNMP. When you start the SNMP service, the following 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 10.16 shows the SNMP configuration screen. Table 10.12 summarizes the configuration options.

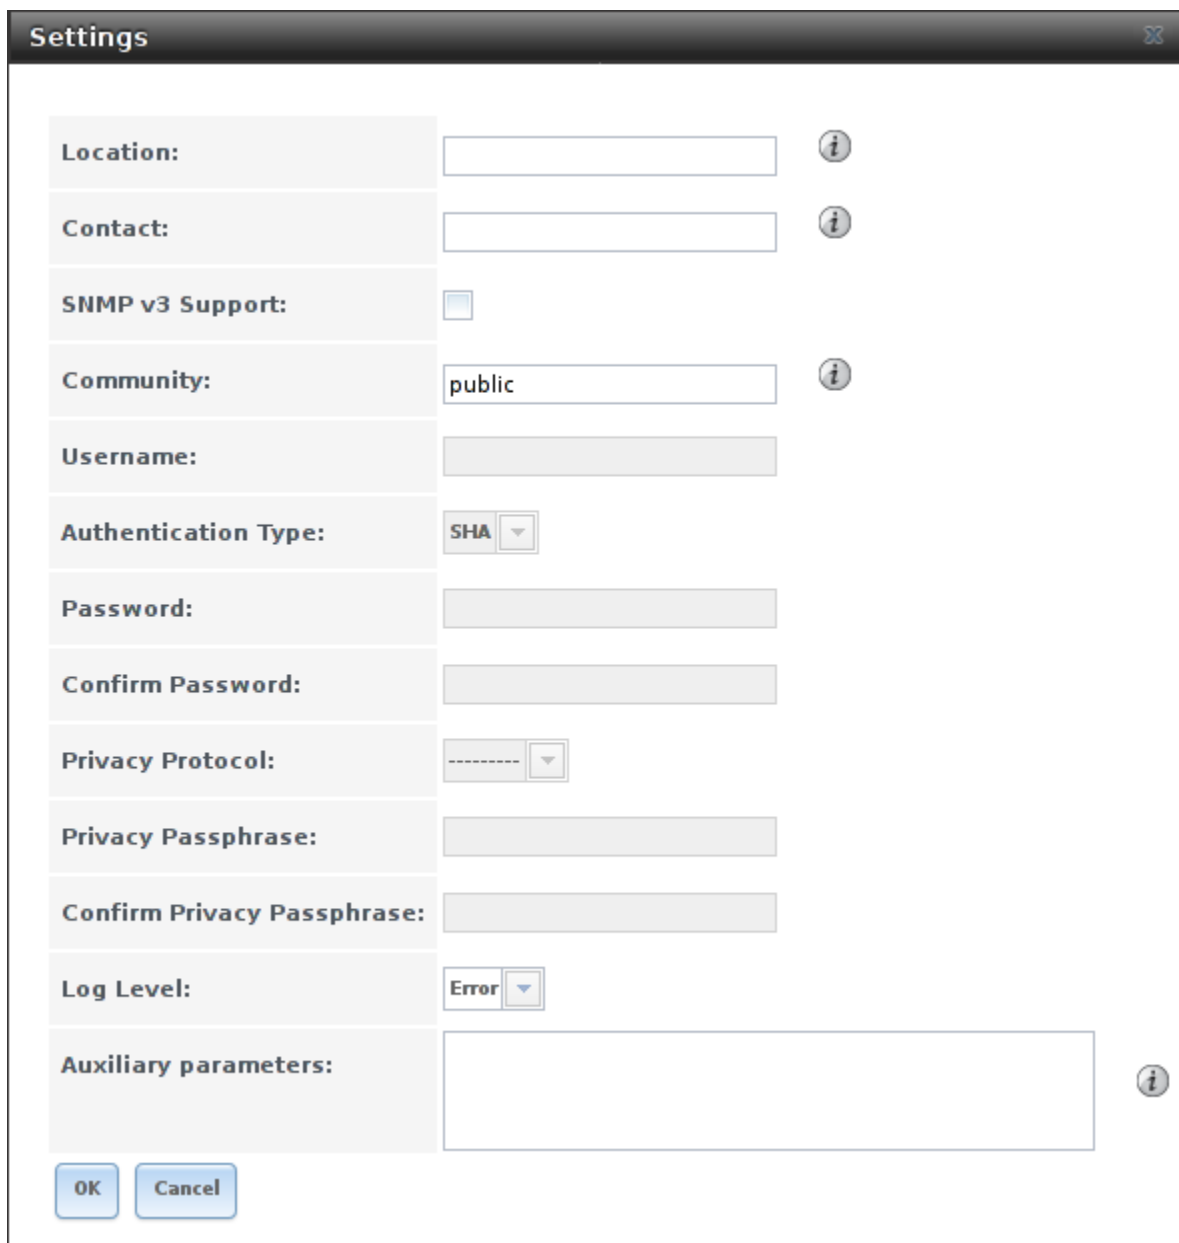
A screenshot of the 'Settings' window for SNMP configuration. The window has a title bar with 'Settings' and a close button. The main area contains several configuration fields: 'Location:' (text box), 'Contact:' (text box), 'SNMP v3 Support:' (checkbox), 'Community:' (text box with 'public' entered), 'Username:' (text box), 'Authentication Type:' (dropdown menu with 'SHA' selected), 'Password:' (text box), 'Confirm Password:' (text box), 'Privacy Protocol:' (dropdown menu with a dashed line selected), 'Privacy Passphrase:' (text box), 'Confirm Privacy Passphrase:' (text box), 'Log Level:' (dropdown menu with 'Error' selected), and 'Auxiliary parameters:' (large text box). Information icons (i) are present next to 'Location:', 'Contact:', 'Community:', and 'Auxiliary parameters:'. At the bottom left are 'OK' and 'Cancel' buttons.

Fig. 10.16: Configuring SNMP

Table 10.12: SNMP Configuration Options

Setting	Value	Description
Location	string	optional description of system's location
Contact	string	optional email address of administrator
SNMP v3 Support	checkbox	check this box to enable support for SNMP version 3
Community	string	default is public and should be changed for security reasons; can only contain alphanumeric characters, underscores, dashes, periods, and spaces; this value can be empty for SNMPv3 networks
Username	string	only applies if SNMP v3 Support is checked; specify the username to register with this service; refer to snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) for more information regarding the configuration of this setting as well as the Authentication Type, Password, Privacy Protocol, and "Privacy Passphrase" fields
Authentication Type	drop-down menu	only applies if SNMP v3 Support is checked; choices are MD5 or SHA
Password	string	only applies if SNMP v3 Support is checked; specify and confirm a password of at least eight characters
Privacy Protocol	drop-down menu	only applies if SNMP v3 Support is checked; choices are AES or DES
Privacy Passphrase	string	if not specified, Password is used
Log Level	drop-down menu	choices range from the least log entries (Emergency) to the most (Debug)
Auxiliary Parameters	string	additional snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) options not covered in this screen, one per line

10.15 SSH

Secure Shell (SSH) allows for files to be transferred securely over an encrypted network. If you configure your TrueNAS® system as an SSH server, the users in your network will need to use [SSH client software](https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) (https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) to transfer files with SSH.

This section shows the TrueNAS® SSH configuration options, demonstrates an example configuration that restricts users to their home directory, and provides some troubleshooting tips.

[Figure 10.17](#) shows the `Services → SSH` configuration screen. After configuring SSH, remember to start it in `Services → Control Services`.

The image shows a window titled "SSH" with a close button in the top right corner. Inside the window, there are five configuration options, each with a label and a control element:

- TCP Port:** A text input field containing the number "22" and an information icon to its right.
- Login as Root with password:** A checkbox that is currently unchecked, with an information icon to its right.
- Allow Password Authentication:** A checkbox that is currently checked.
- Allow TCP Port Forwarding:** A checkbox that is currently unchecked.
- Compress Connections:** A checkbox that is currently unchecked.

At the bottom of the window, there are three buttons: "OK", "Cancel", and "Advanced Mode".

Fig. 10.17: SSH Configuration

Table 10.13 summarizes the configuration options. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button, or configure the system to always display these settings by checking the box Show advanced fields by default in *System* → *Advanced*.

Table 10.13: SSH Configuration Options

Setting	Value	Advanced Mode	Description
Bind Interfaces	selection	✓	by default, SSH listens on all interfaces unless specific interfaces are highlighted in the Available field and added to the Selected field
TCP Port	integer		port to open for SSH connection requests; 22 by default
Login as Root with password	checkbox		for security reasons, root logins are discouraged and disabled by default if enabled, password must be set for root user in View Users
Allow Password Authentication	checkbox		if unchecked, key-based authentication for all users is required; requires additional setup (http://the.earth.li/%7Esgtatham/putty/0.55/html/doc/Chapter8.html) on both the SSH client and server
Allow Kerberos Authentication	checkbox		before checking this box, ensure that Kerberos Realms (page 141) and Kerberos Keytabs (page 141) have been configured and that the TrueNAS® system can communicate with the KDC
Allow TCP Port Forwarding	checkbox		allows users to bypass firewall restrictions using SSH's port forwarding feature (http://www.symantec.com/connect/articles/ssh-port-forwarding)
Compress Connections	checkbox		may reduce latency over slow networks
SFTP Log Level	drop-down menu	✓	select the syslog(3) (http://www.freebsd.org/cgi/man.cgi?query=syslog) level of the SFTP server
SFTP Log Facility	drop-down menu	✓	select the syslog(3) (http://www.freebsd.org/cgi/man.cgi?query=syslog) facility of the SFTP server

Continued on next page

Table 10.13 – continued from previous page

Setting	Value	Advanced Mode	Description
Extra Options	string	✓	additional sshd_config(5) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options not covered in this screen, one per line; these options are case-sensitive and misspellings may prevent the SSH service from starting

A few [sshd_config\(5\)](http://www.freebsd.org/cgi/man.cgi?query=sshd_config) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options that are useful to enter in the Extra Options field include:

- increase the `ClientAliveInterval` if SSH connections tend to drop
- `ClientMaxStartup` defaults to 10; increase this value if you need more concurrent SSH connections

10.15.1 SCP Only

When you configure SSH, authenticated users with a user account created using `Account → Users → Add User` can use the `ssh` command to login to the TrueNAS® system over the network. A user's home directory will be the volume/dataset specified in the Home Directory field of their TrueNAS® user account. While the SSH login will default to the user's home directory, users are able to navigate outside of their home directory, which can pose a security risk.

It is possible to allow users to use the `scp` and `sftp` commands to transfer files between their local computer and their home directory on the TrueNAS® system, while restricting them from logging into the system using `ssh`. To configure this scenario, go to `Account → Users → View Users`, select the user and click `Modify User`, and change the user's Shell to `scponly`. Repeat for each user that needs restricted SSH access.

Test the configuration from another system by running the `sftp`, `ssh`, and `scp` commands as the user. The `sftp` and `scp` commands should work but the `ssh` should fail.

Note: Some utilities such as WinSCP and Filezilla can bypass the `scponly` shell. This section assumes that users are accessing the system using the command line versions of `scp` and `sftp`.

10.15.2 Troubleshooting SSH

When adding any Extra Options, be aware that the keywords listed in [sshd_config\(5\)](http://www.freebsd.org/cgi/man.cgi?query=sshd_config) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) are case sensitive. This means that your configuration will fail to do what you intended if you do not match the upper and lowercase letters of the keyword.

If your clients are receiving “reverse DNS” or timeout errors, add an entry for the IP address of the TrueNAS® system in the Host name database field of `Network → Global Configuration`.

When configuring SSH, always test your configuration as an SSH user account to ensure that the user is limited to what you have configured and that they have permission to transfer files within the intended directories. If the user account is experiencing problems, the SSH error messages are usually pretty specific to what the problem is. Type the following command within [Shell](#) (page 257) to read these messages as they occur:

```
tail -f /var/log/messages
```

Additional messages regarding authentication errors may be found in `/var/log/auth.log`.

10.16 TFTP

Trivial File Transfer Protocol (TFTP) is a light-weight version of FTP usually used to transfer configuration or boot files between machines, such as routers, in a local environment. TFTP provides an extremely limited set of commands and provides no authentication.

If the TrueNAS® system will be used to store images and configuration files for the network's devices, configure and start the TFTP service. Starting the TFTP service will open UDP port 69.

Figure 10.18 shows the TFTP configuration screen and Table 10.14 summarizes the available options:

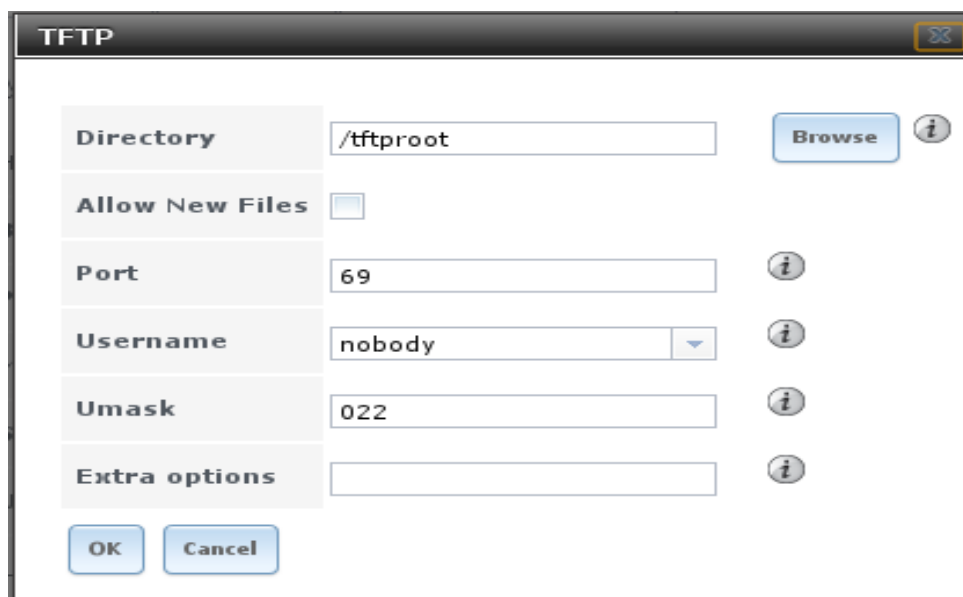


Fig. 10.18: TFTP Configuration

Table 10.14: TFTP Configuration Options

Setting	Value	Description
Directory	browse button	browse to an existing directory to be used for storage; some devices require a specific directory name, refer to the device's documentation for details
Allow New Files	checkbox	enable if network devices need to send files to the system (for example, to back up their configuration)
Port	integer	UDP port to listen for TFTP requests, 69 by default
Username	drop-down menu	account used for tftp requests; must have permission to the Directory
Umask	integer	umask for newly created files, default is 022 (everyone can read, nobody can write); some devices require a less strict umask
Extra options	string	additional <code>tftpd(8)</code> (http://www.freebsd.org/cgi/man.cgi?query=tftpd) options not shown in this screen, one per line

10.17 UPS

TrueNAS® uses [NUT](http://www.networkupstools.org/) (Network UPS Tools) to provide UPS support. If the TrueNAS® system is connected to a UPS device, configure the UPS service then start it in `Services → Control Services`.

Figure 10.19 shows the UPS configuration screen:

Settings

UPS Mode:

Master

Identifier:

ups

Driver:

Port:

Auxiliary parameters (ups.conf):

Auxiliary parameters (upsd.conf):

Description:

Shutdown mode:

UPS goes on battery

Shutdown timer:

30

Shutdown Command:

/sbin/shutdown -p now

No Communication Warning Time:

Monitor User:

upsmon

Monitor Password:

fixmepass

Extra users (upsd.users):

Remote Monitor:

☐

Send Email Status Updates:

☐

To email:

Email Subject:

UPS report generated by %h

Power Off UPS:

☐

OK

Cancel

Fig. 10.19: UPS Configuration Screen

Table 10.15 summarizes the options in the UPS Configuration screen.

Table 10.15: UPS Configuration Options

Setting	Value	Description
UPS Mode	drop-down menu	select from Master or Slave
Identifier	string	can contain alphanumeric, period, comma, hyphen, and underscore characters
Driver	drop-down menu	supported UPS devices are listed at http://www.networkupstools.org/stable-hcl.html
Port	drop-down menu	select the serial or USB port the UPS is plugged into (see NOTE below)
Auxiliary Parameters (ups.conf)	string	additional options from ups.conf(5) (http://www.networkupstools.org/docs/man/ups.conf.html)
Auxiliary Parameters (upsd.conf)	string	additional options from upsd.conf(5) (http://www.networkupstools.org/docs/man/upsd.conf.html)
Description	string	optional
Shutdown mode	drop-down menu	choices are UPS goes on battery and UPS reaches low battery
Shutdown timer	integer	in seconds; will initiate shutdown after this many seconds after UPS enters UPS goes on battery, unless power is restored
Shutdown Command	string	the command to run to shut down the computer when battery power is low or shutdown timer runs out
No Communication Warning Time	string	the frequency, in seconds, of email notifications during the loss of UPS communications; the default is 300
Monitor User	string	default is upsmon
Monitor Password	string	default is known value fixmepass and should be changed; cannot contain a space or #
Extra users	string	defines the accounts that have administrative access; see upsd.users(5) (http://www.networkupstools.org/docs/man/upsd.users.html) for examples
Remote monitor	checkbox	if enabled, be aware that the default is to listen on all interfaces and to use the known values user upsmon and password fixmepass
Send Email Status Updates	checkbox	if checked, activates the To email field
To email	email address	if Send Email box checked, email address to receive status updates; separate multiple email addresses with a semicolon
Email Subject	string	subject line to be used in the email
Power Off UPS	checkbox	if checked, the UPS will also power off after shutting down the FreeNAS system

Note: For USB devices, the easiest way to determine the correct device name is to check the box Show console messages in System → Advanced. Plug in the USB device and look for a /dev/ugen or /dev/uhid device name in the console messages.

[upsc\(8\)](http://www.networkupstools.org/docs/man/upsc.html) (<http://www.networkupstools.org/docs/man/upsc.html>) can be used to get status variables from the UPS daemon such as the current charge and input voltage. It can be run from Shell using the following syntax. The man page gives some other usage examples.

```
upsc ups@localhost
```

[upscmd\(8\)](http://www.networkupstools.org/docs/man/upscmd.html) (<http://www.networkupstools.org/docs/man/upscmd.html>) can be used to send commands directly to the UPS, assuming that the hardware supports the command being sent. Only users with administrative rights can use this command. These users are created in the Extra users field.

10.17.1 Multiple Computers with One UPS

A UPS with adequate capacity can be used to power multiple computers. One computer is connected to the UPS data port with a serial or USB cable. This master makes UPS status available on the network for other computers. These slave computers are powered by the UPS, but receive UPS status data from the master computer. See the [NUT User Manual](http://networkupstools.org/docs/user-manual.chunked/index.html) (<http://networkupstools.org/docs/user-manual.chunked/index.html>) and [NUT User Manual Pages](http://networkupstools.org/docs/man/index.html#User_man) (http://networkupstools.org/docs/man/index.html#User_man).

10.18 WebDAV

The WebDAV service can be configured to provide a file browser over a web connection. Before starting this service, you must create at least one WebDAV share using `Sharing → WebDAV Shares → Add WebDAV Share`. Refer to [WebDAV Shares](#) (page 158) for instructions on how to create a share and then how to connect to it once the service is configured and started.

The settings in the WebDAV service apply to all WebDAV shares. [Figure 10.20](#) shows the WebDAV configuration screen. [Table 10.16](#) summarizes the available options.




Fig. 10.20: WebDAV Configuration Screen

Table 10.16: WebDAV Configuration Options

Setting	Value	Description
Protocol	drop-down menu	choices are HTTP (connection always unencrypted), HTTPS (connection always encrypted), or HTTP+HTTPS (both types of connections allowed)
HTTP Port	string	only appears if the selected Protocol is HTTP or HTTP+HTTPS and is used to specify the port to be used for unencrypted connections; the default of 8080 should work, if you change it, do not use a port number already being used by another service
HTTPS Port	string	only appears if the selected Protocol is HTTPS or HTTP+HTTPS and is used to specify the port to be used for encrypted connections; the default of 8081 should work, if you change it, do not use a port number already being used by another service
Webdav SSL Certificate	drop-down menu	only appears if the selected Protocol is HTTPS or HTTP+HTTPS; select the SSL certificate to be used for encrypted connections; to create a certificate, use <code>System → Certificates</code>
HTTP Authentication	drop-down menu	choices are No Authentication, Basic Authentication (unencrypted), or Digest Authentication (encrypted)

Continued on next page

Table 10.16 – continued from previous page

Setting	Value	Description
Webdav Password	string	default is davtest; this should be changed as it is a known value

VCENTER PLUGIN

The TrueNAS® vCenter Plugin integrates the control and operation of TrueNAS® into VMware vCenter Server (<https://www.vmware.com/products/vcenter-server>).

11.1 New Features in 2.1.0

- Advanced provisioning
- Datastore snapshot management
- LUN management
- ZFS pool management
- ZFS pool health status
- Role base access control

Note: The current vCenter plugin works with versions of vCenter up to vCenter 6.

11.2 Configuration

Click vCenter to open the plugin configuration screen shown in [Figure 11.1](#).

vCenter Plugin Configuration

Plugin Name:
TrueNAS vCenter Plugin

Available Plugin Version:
2.2.0

Installed Plugin Version:
2.2.0

TrueNAS Management IP Address:
--Select--

vCenter Hostname/IP Address:

vCenter Port:
443

vCenter Username:

vCenter Password:

Install
Uninstall
Upgrade
Repair

Fig. 11.1: Configuring the vCenter Plugin

Table 11.1 lists the options on this screen.

Table 11.1: vCenter Plugin Options

Setting	Value	Description
TrueNAS Management IP Address	drop-down menu	select the CARP address of the TrueNAS® system
vCenter Hostname/IP Address	string	the IP address or resolveable hostname of the vCenter Server
vCenter Port	integer	the port number on which the vCenter Server is listening
vCenter Username	string	the username for the vCenter Server
vCenter Password	string	the password associated with vCenter Username

Enter the information, then click the Install button to add the TrueNAS® system as an object in vCenter Server. From vCenter Server, click the object to create its datastores.

In addition to the Install button, these buttons are available:

Uninstall: click to remove the TrueNAS® object from vCenter Server.

Upgrade: as more features are added to the vCenter plugin, the Available Plugin Version number is incremented. Click this button to upgrade to the newer version.

Repair: click this button at the request of an iXsystems support engineer. This reinstalls the TrueNAS® object to repair a corrupted object.

Note: In an HA system, the Upgrade button can only be used on the system that originally installed the plugin. The Upgrade button is disabled on the other HA node.

11.3 Secure Connections

The vCenter plugin can be configured for a secure connection by clicking `vCenter` → `vCenter Auxiliary Settings` in the left tree. In the secure connection screen shown in Figure 11.2, check the `Enable vCenter Plugin over https` box.




Fig. 11.2: Securing the vCenter Plugin Connection

11.4 Using the vCenter Plugin

Log in to the vCenter server by entering the IP address in a web browser. From the main vSphere window, click vCenter Inventory Lists, TrueNAS Hosts, and TrueNAS to see the list of defined TrueNAS® hosts.

11.5 Hosts

The TrueNAS® hosts are listed on this screen. New hosts can be added by clicking the + button. Enter the new TrueNAS® host IP address, username, and password of that host to add it to the list of hosts managed by vCenter. Figure 11.3 shows the Add New Host window.

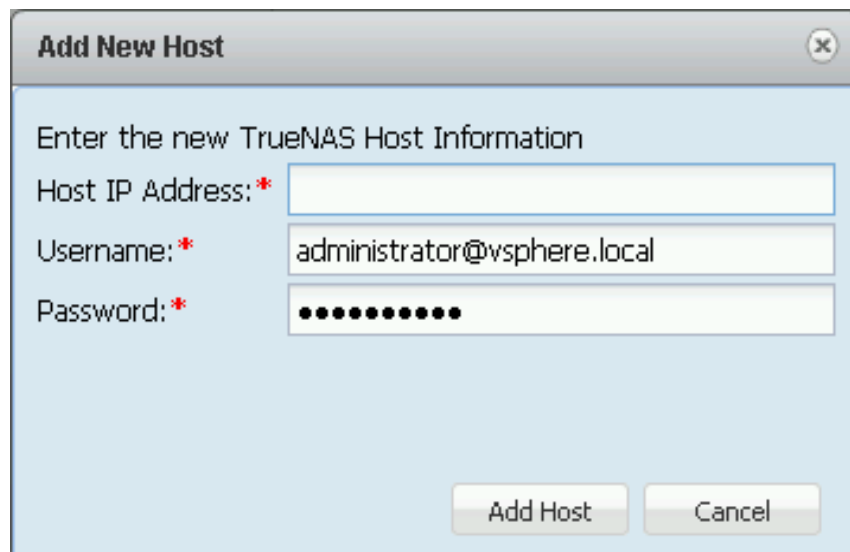


Fig. 11.3: Adding a New TrueNAS® Host

After a TrueNAS® host has been added, click its entry in the sidebar to view Summary and Manage options.

11.5.1 Host Information

Click the Summary tab to view basic information about the connected TrueNAS® host. IP address, Version, Storage Sizes, and status of additional features like NFS and HA are provided.

11.5.2 Manage Host

Click the Manage tab to view a number of options to modify a TrueNAS® host. Each option displays related modification buttons and, except for the TrueNAS Configuration option, a table that shows existing configurations. Highlight a table header and click the drop-down menu to sort the table or choose which columns to display.

11.6 Datastores

The vCenter plugin can create VMFS datastores on TrueNAS® hosts for iSCSI block-level access, or NFS datastores for file-level access.

11.6.1 VMFS Datastores

VMFS datastores provide block-level (iSCSI) storage for virtual machines.

Creating VMFS Datastores

Select the Datastore tab, then click the + button to create a datastore. Select the ESXi host on which the datastore is to be available, as shown in [Figure 11.4](#), then click Next.

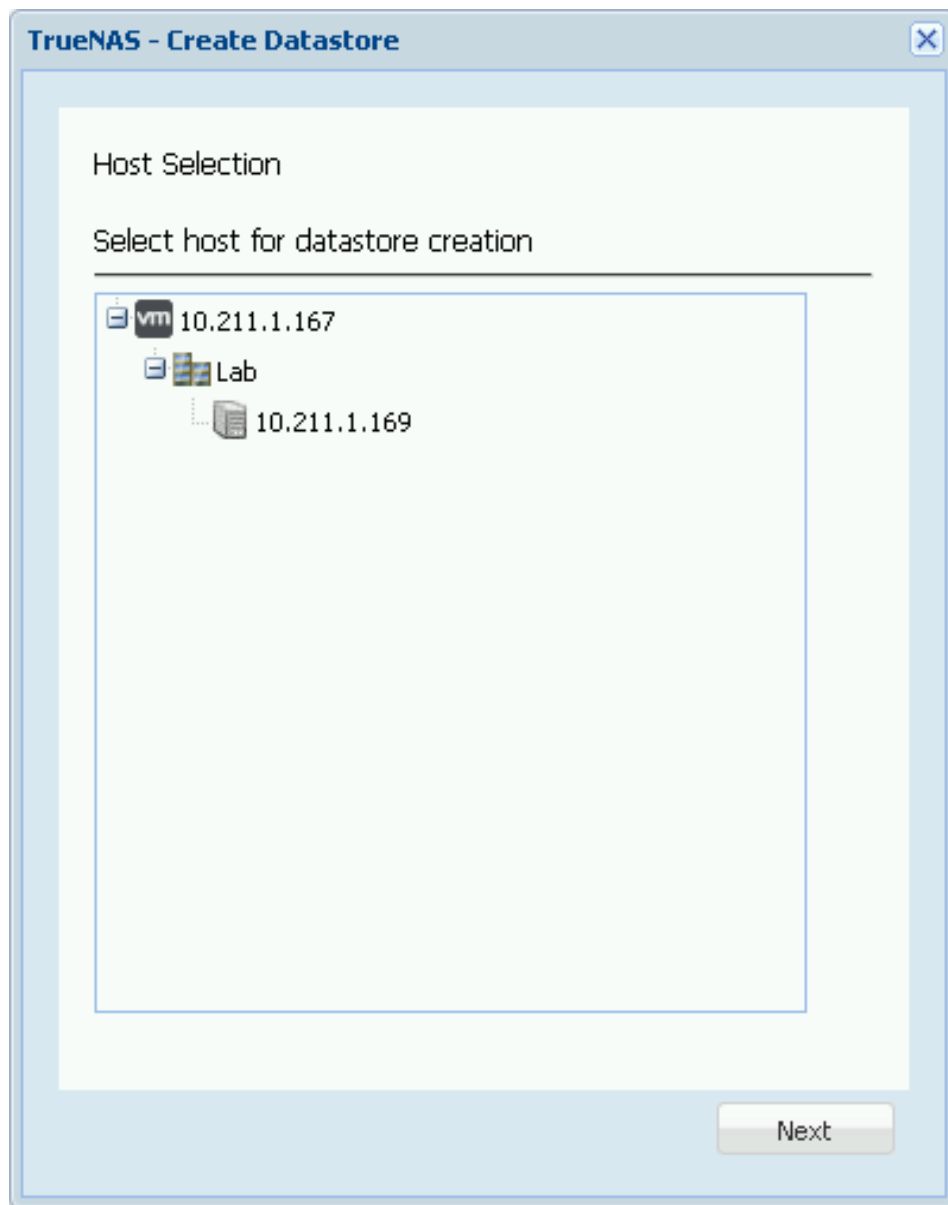
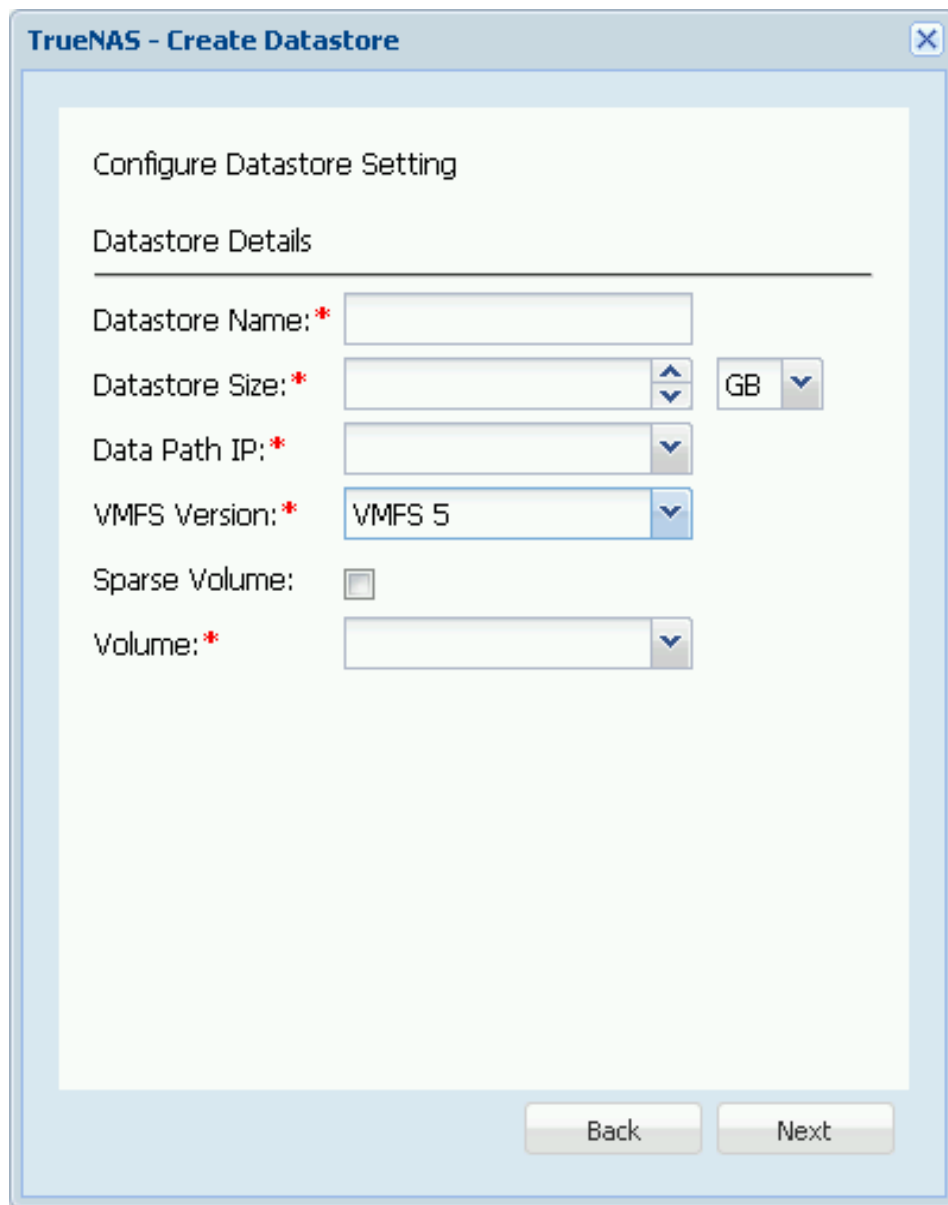


Fig. 11.4: Choosing an ESXi Host

Select VMFS as the datastore type and click Next. [Figure 11.5](#) shows the configuration screen for the new datastore:



The image shows a 'TrueNAS - Create Datastore' dialog box. It has a title bar with a close button. The main area is titled 'Configure Datastore Setting' and contains a section 'Datastore Details'. This section includes several input fields: 'Datastore Name' with a red asterisk, 'Datastore Size' with a red asterisk, 'Data Path IP' with a red asterisk, 'VMFS Version' with a red asterisk and a dropdown menu showing 'VMFS 5', 'Sparse Volume' with an unchecked checkbox, and 'Volume' with a red asterisk and a dropdown menu. To the right of the 'Datastore Size' field is a unit selector showing 'GB'. At the bottom of the dialog are 'Back' and 'Next' buttons.

TrueNAS - Create Datastore

Configure Datastore Setting

Datastore Details

Datastore Name: *

Datastore Size: * GB

Data Path IP: *

VMFS Version: * VMFS 5

Sparse Volume: ☐

Volume: *

Back Next

Fig. 11.5: Configuring a VMFS Datastore

Enter a name for the new datastore and the desired size. Datastore names can only contain letters, numbers, and – and . characters. Enter the IP address of the TrueNAS® system in the Data Path IP box. Select the desired VMFS Version from the drop-down. If a Sparse Volume is desired, check the box.

Select the TrueNAS® volume on which to create the datastore, then click Next. [Figure 11.6](#) shows the final summary screen:

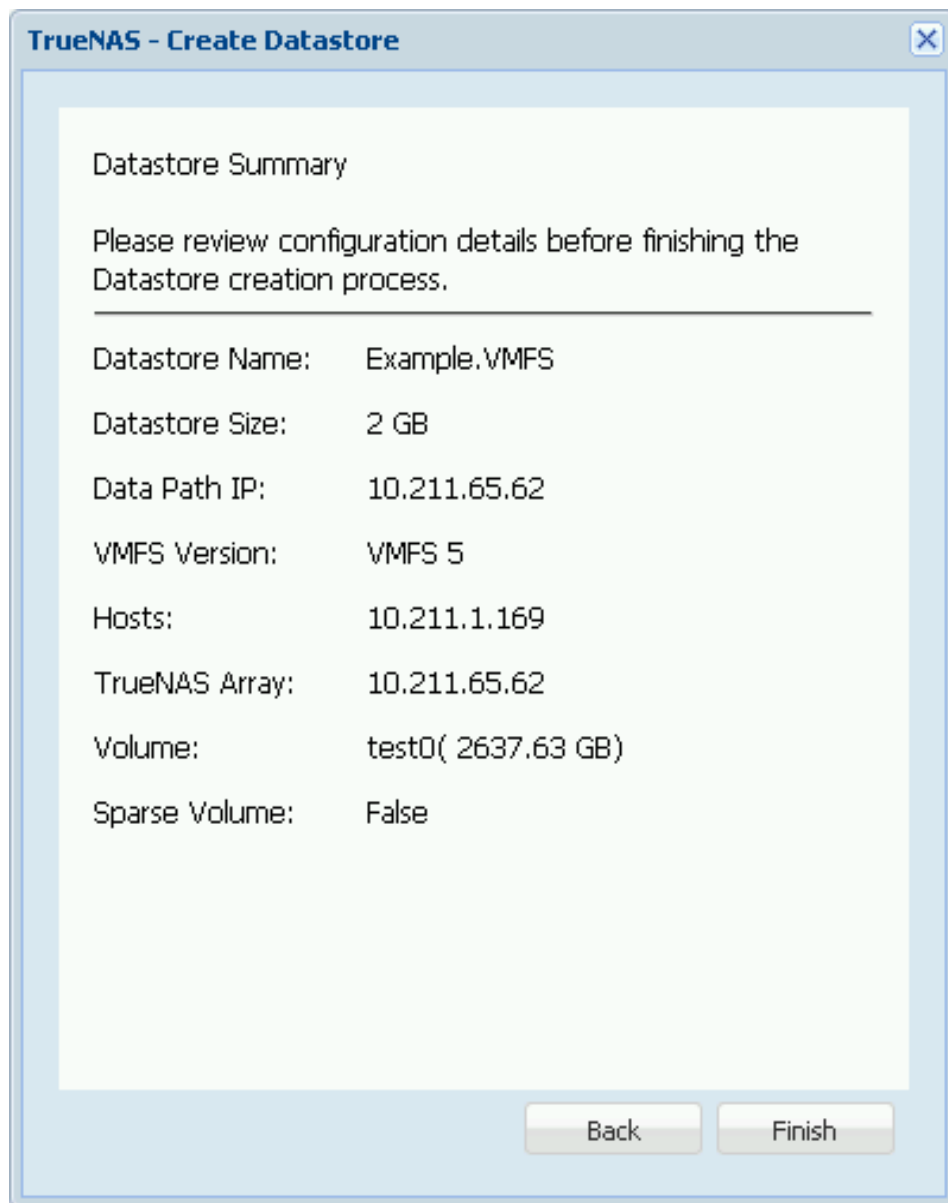


Fig. 11.6: Finish Creating the Datastore

Click Finish to create the datastore. The datastore may not be immediately visible in the list of datastores. The Refresh button refresh the list contents to show new datastores.

Extending VMFS Datastores

Highlight a VMFS datastore from the list, then click the Edit button to begin extending the datastore. [Figure 11.7](#) shows the Extend Datastore window.




Fig. 11.7: Extending a Datastore

Choose a new size for the datastore, then click Extend Datastore. A pop-up window appears, showing the system is extending the datastore.

Note: A datastore uses some of the space internally, resulting in the available capacity being slightly less than the amount chosen.

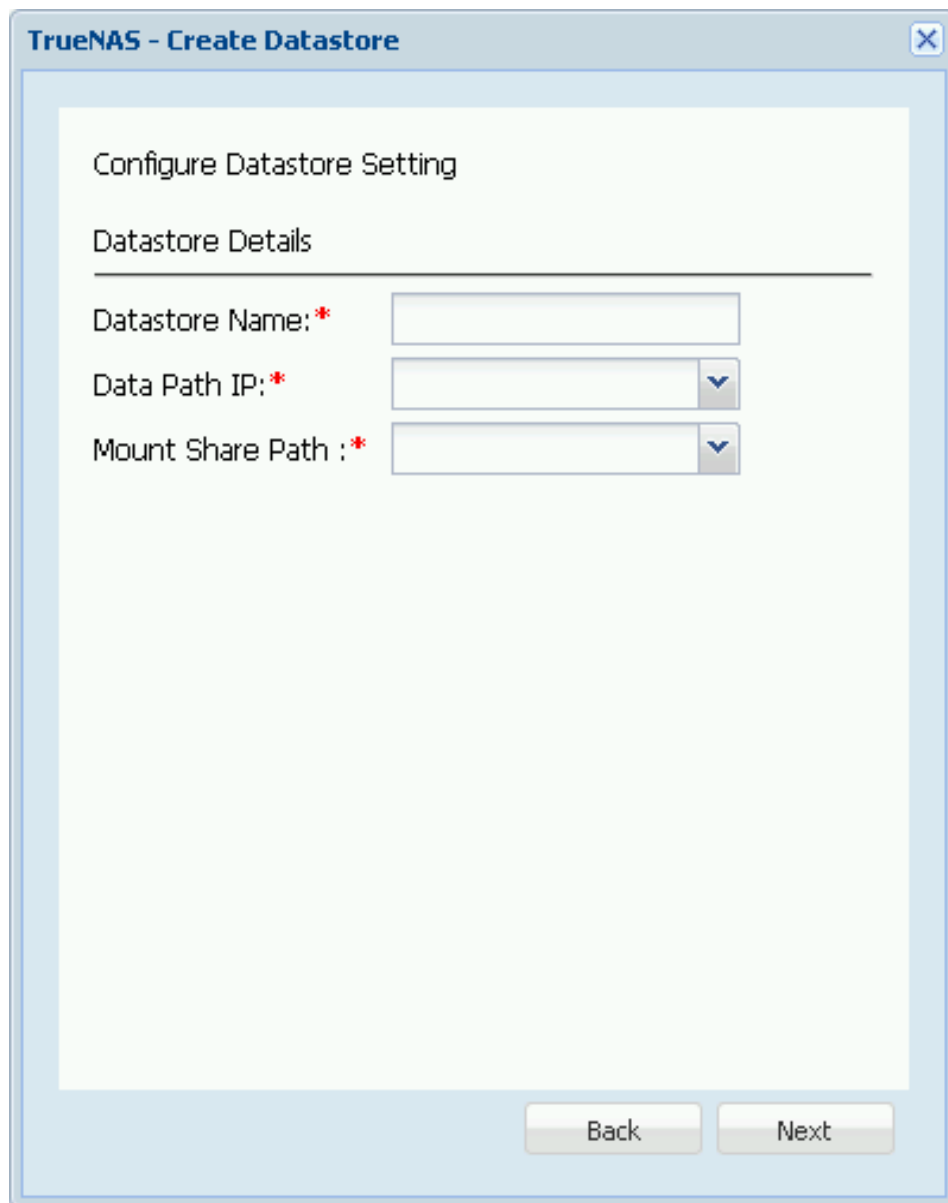
11.6.2 NFS Datastores

NFS datastores provide file-level storage access.

Creating NFS Datastores

Select the Datastore tab, then click the + button to create a datastore. Select the [ESXi host](#) (page 226) on which the datastore is to be available. Click Next.

Select NFS as the datastore type, then click Next. [Figure 11.8](#) shows the configuration window for the new NFS datastore.



The image shows a window titled "TrueNAS - Create Datastore" with a close button in the top right corner. Inside the window, the text "Configure Datastore Setting" is displayed. Below this, a section titled "Datastore Details" is separated by a horizontal line. This section contains three input fields, each with a red asterisk indicating it is required: "Datastore Name:" followed by a text box; "Data Path IP:" followed by a text box and a dropdown arrow; and "Mount Share Path :" followed by a text box and a dropdown arrow. At the bottom right of the window, there are two buttons: "Back" and "Next".

Fig. 11.8: Configuring an NFS Datastore

Enter a name for the new datastore. Enter the IP address of the TrueNAS® system in the Data Path IP box. The path to the NFS share on the TrueNAS® system is entered in the Mount Share Path box. Click Next.

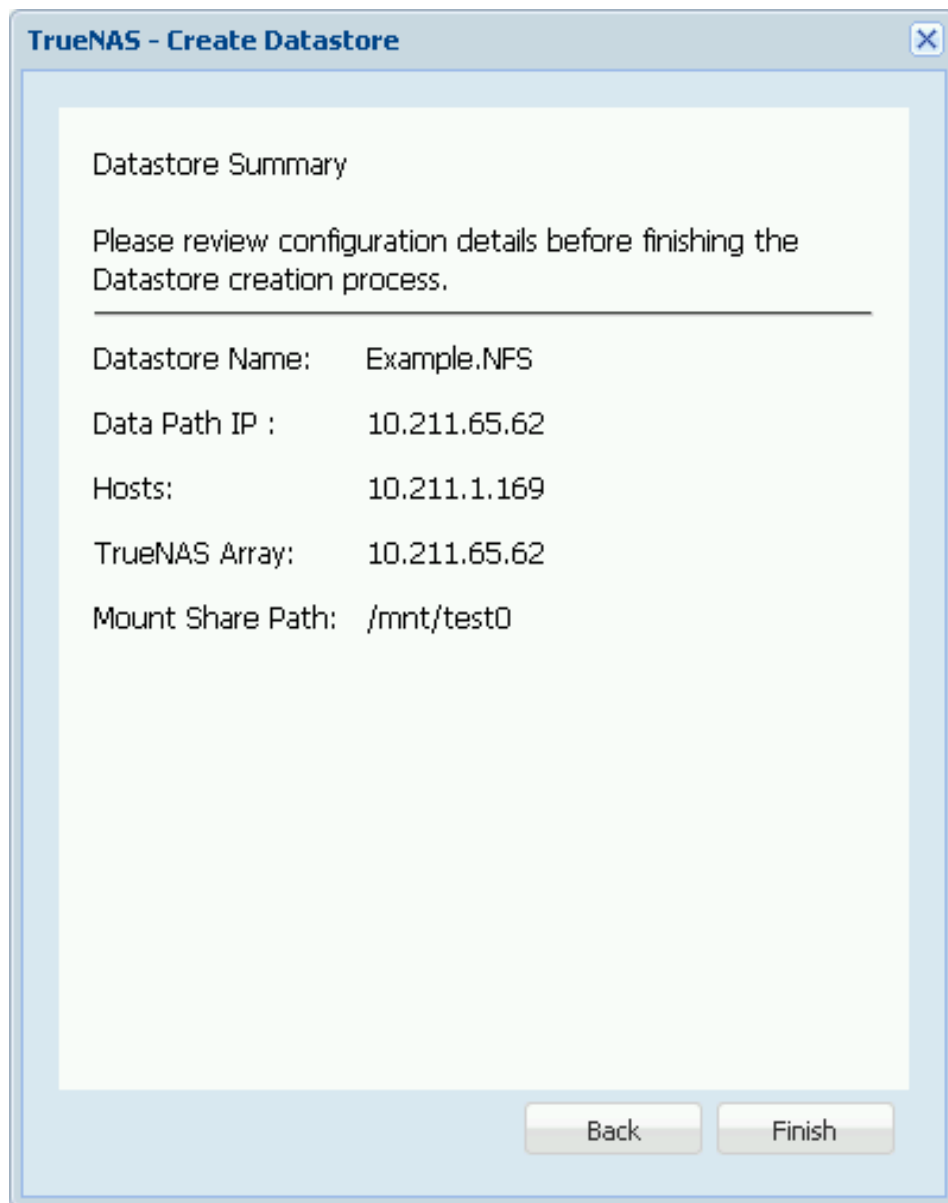


Fig. 11.9: New NFS Datastore Summary

The summary screen in [Figure 11.9](#) is shown. Click Finish to create the datastore. The datastore might not be immediately visible in the list of datastores. Click the Refresh button to refresh the list contents.

11.6.3 Cloning Datastores

A datastore of either type can be cloned by selecting it in the list of datastores, then clicking the Clone button. Select the ESXi host on which the datastore is to be available. Click Next to see the window shown in [Figure 11.10](#).




Fig. 11.10: Creating a Datastore Clone

Type a name for the clone datastore and click Clone Datastore. A pop-up displays when cloning begins. The new clone datastore may not be immediately visible. Click the Refresh button to refresh the list contents.

11.6.4 Snapshotting a Datastore

Highlight a datastore from the list and click the Create Snapshot button to snapshot that datastore. Type a name for the snapshot and click Create to create the new snapshot. The new snapshot is added to the Snapshot tab, separate from the Datastore tab.

11.7 Snapshots

See the [Snapshots](#) (page 129) section for more details about creating and managing snapshots directly on the TrueNAS® system.

Available snapshots are listed in this tab. Figure 11.11 shows an example.




Fig. 11.11: Managing Snapshots

Buttons to Revert and Delete snapshots are available.

Revert returns a datastore to the state saved by the selected snapshot. When choosing a snapshot, be sure it is associated with the proper datastore. Click Revert and confirm this is correct by clicking Yes in the window that appears. The system then rolls the datastore back to the state saved in the snapshot.

11.7.1 Scheduling: Periodic Snapshots

See the [Periodic Snapshot Tasks](#) (page 113) section for more details about native TrueNAS® periodic snapshots.

The Scheduling tab is used to Create, Edit, and Delete periodic snapshot schedules.

Click the + button to begin creating a new snapshot schedule, as shown in [Figure 11.12](#).

Periodic Snapshot

Datastore: * ▼

Recursive: ☐

Lifetime: * Hour(s) ▼

Begin: * ▼

End: * ▼

Interval: * ▼

Weekdays: *

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

Enable: ☐

OK Cancel

Fig. 11.12: Creating a New Periodic Snapshot Task

The Datastore, Lifetime, Begin, End, Interval, and Weekdays options must be configured before the schedule can be saved.

Note: For detailed descriptions of each option, see the [Periodic Snapshot Options](#) (page 114) table.

When satisfied with the configuration, click OK to save the schedule and add it to the list. If the new schedule is not immediately visible, click the Refresh button to repopulate the list of saved schedules.

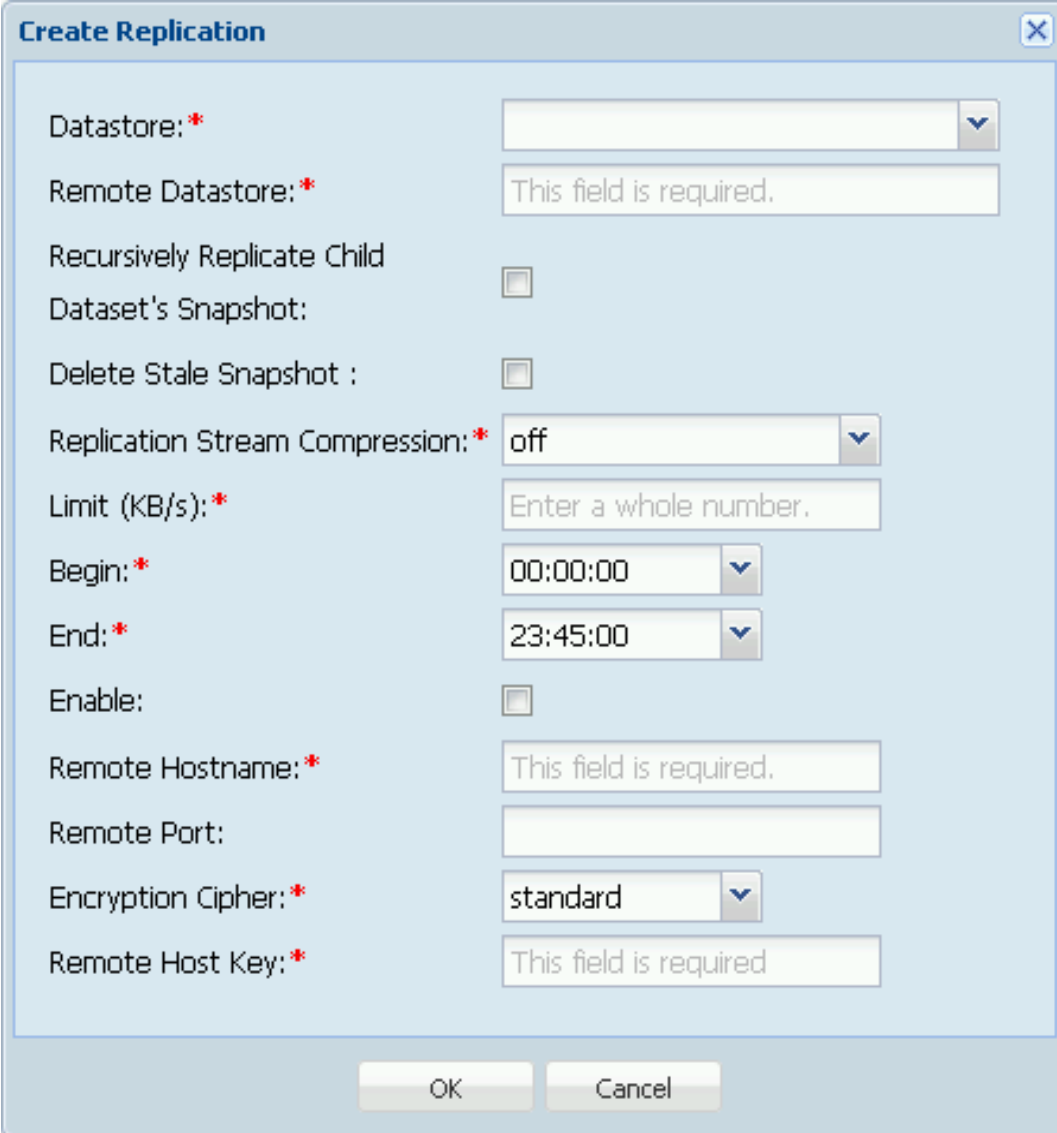
Select a schedule from the list and click Edit to open the configuration window for that schedule. Make any needed adjustments, then click Save to save the updated schedule.

11.8 Replication

The [Replication Tasks](#) (page 115) section shows examples of different replication configurations. Refer to this section for more general details about creating and modifying replications.

Note: A periodic snapshot task must be available before creating a new replication task.

Click the + button to open the Create Replication window, as shown in [Figure 11.13](#).

The image shows a 'Create Replication' dialog box with a light blue background and a standard Windows-style title bar with a close button. The dialog contains several configuration fields. 'Datastore:' is a dropdown menu. 'Remote Datastore:' is a text field with a red asterisk and the message 'This field is required.'. 'Recursively Replicate Child' is a checkbox. 'Dataset's Snapshot:' is a text field. 'Delete Stale Snapshot :' is a checkbox. 'Replication Stream Compression: *' is a dropdown menu set to 'off'. 'Limit (KB/s): *' is a text field with the message 'Enter a whole number.'. 'Begin: *' and 'End: *' are time selection fields set to '00:00:00' and '23:45:00' respectively. 'Enable:' is a checkbox. 'Remote Hostname: *' is a text field with the message 'This field is required.'. 'Remote Port:' is a text field. 'Encryption Cipher: *' is a dropdown menu set to 'standard'. 'Remote Host Key: *' is a text field with the message 'This field is required'. At the bottom are 'OK' and 'Cancel' buttons.

Create Replication

Datastore: *

Remote Datastore: * This field is required.

Recursively Replicate Child ☐

Dataset's Snapshot:

Delete Stale Snapshot : ☐

Replication Stream Compression: * off

Limit (KB/s): * Enter a whole number.

Begin: * 00:00:00

End: * 23:45:00

Enable: ☐

Remote Hostname: * This field is required.

Remote Port:

Encryption Cipher: * standard

Remote Host Key: * This field is required

OK Cancel

Fig. 11.13: Create a new replication schedule

The Datastore, Remote Datastore, Replication Stream Compression, Limit (KB/s), Begin, End, Remote Hostname, Encryption Cipher, and Remote Host Key options must be set to create a replication.

Note: For detailed descriptions of each option, see [Replication Task Options](#) (page 122).

Instructions for obtaining the Remote Host Key are shown in [Encryption Keys](#) (page 120).

Click Ok to save the new replication schedule and add it to the list. If the task does not appear in the list, click the Refresh button.

Select an existing replication from the list and click Edit to modify the replication. Editing a replication provides the same options as [creating a replication](#) (page 235). Make any adjustments to the replication configuration and click OK to save the

new settings.

11.9 Zvols

A *zvol* (page 99) is a raw block device over ZFS.

In the vCenter ZVol tab, click the + button to begin creating a new zvol. Figure 11.14 shows the Create ZVol window.

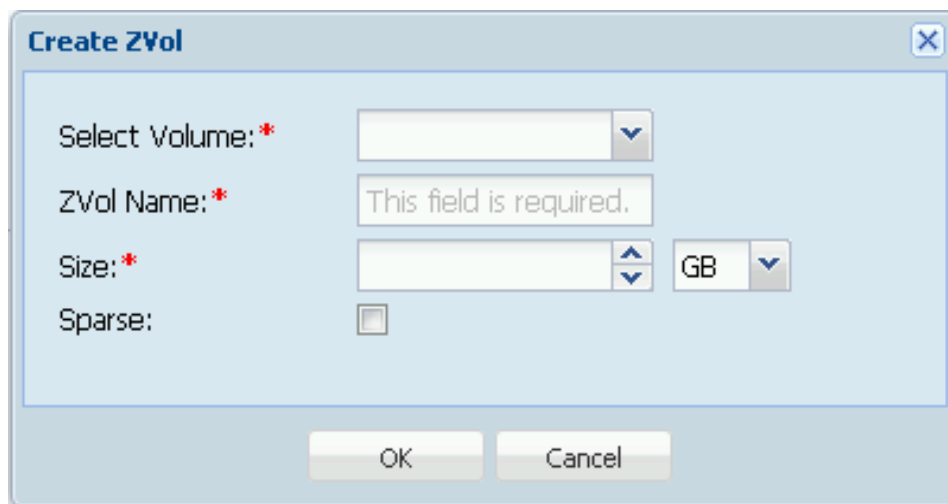
The 'Create ZVol' dialog box is shown. It has a title bar with a close button. Inside, there are four fields: 'Select Volume: *' with a dropdown menu, 'ZVol Name: *' with a text box containing 'This field is required.', 'Size: *' with a numeric input field and a 'GB' unit dropdown, and 'Sparse:' with an unchecked checkbox. At the bottom are 'OK' and 'Cancel' buttons.

Fig. 11.14: Create a new zvol

Select a volume from the drop-down menu. Type a name for the new zvol. Adjust the size of the zvol with the arrow buttons and Size drop menu. If desired, check Sparse to activate thin provisioning.

Warning: Be careful when using Sparse. Thin provisioning can cause writes to fail when the pool is low on space.

Click Ok to save the configuration and begin creating the new zvol. Click the Refresh button to repopulate the list and see the new zvol.

Select a zvol from the list and click the Edit button to change the size of an existing zvol. As Figure 11.15 shows, the only parameter that can be modified for an existing zvol is size.

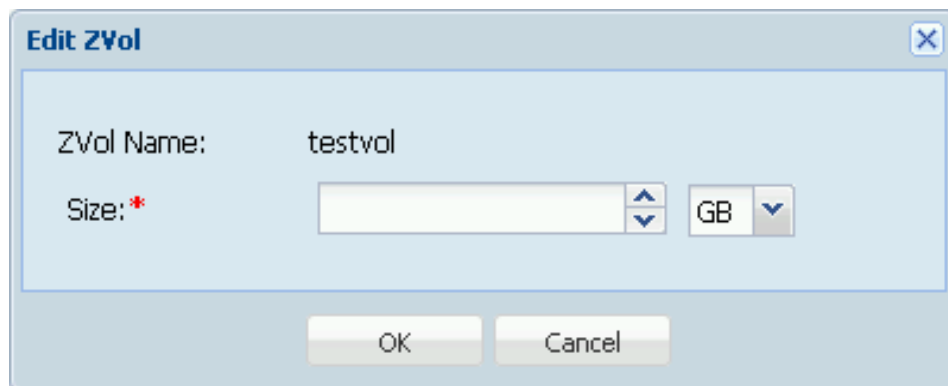
The 'Edit ZVol' dialog box is shown. It has a title bar with a close button. Inside, there are two fields: 'ZVol Name:' with a text box containing 'testvol', and 'Size: *' with a numeric input field and a 'GB' unit dropdown. At the bottom are 'OK' and 'Cancel' buttons.

Fig. 11.15: Edit a Zvol

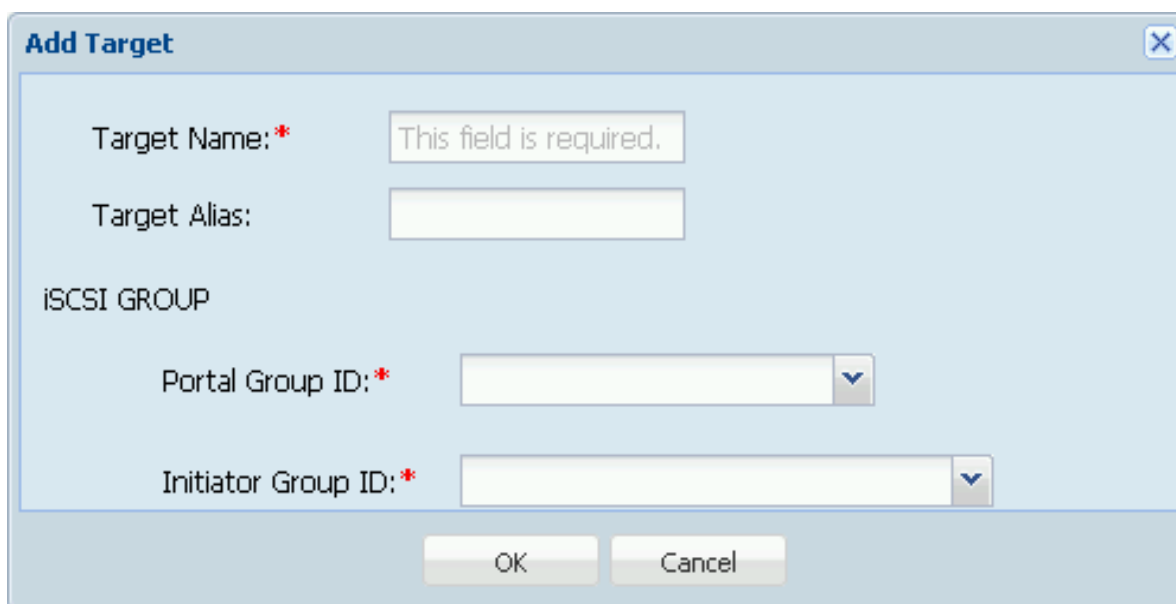
Adjust the size of the ZVol and click OK to save any changes.

11.10 iSCSI Functions

iSCSI is a protocol standard for the consolidation of storage data. Refer to [Block \(iSCSI\)](#) (page 169) for more details about this standard. This section describes configuring [Targets](#) (page 176), [Extents](#) (page 177), and [Associated Targets](#) (page 180) from the vCenter plugin.

11.10.1 Targets

Click the + button to begin creating a new [target](#) (page 176). The Add Target window opens, as shown in [Figure 11.16](#):

The image shows a dialog box titled "Add Target" with a close button (X) in the top right corner. The dialog contains several input fields: "Target Name:" with a red asterisk and a text box containing "This field is required."; "Target Alias:" with a text box; "iSCSI GROUP" as a section header; "Portal Group ID:" with a red asterisk and a dropdown menu; and "Initiator Group ID:" with a red asterisk and a dropdown menu. At the bottom are "OK" and "Cancel" buttons.

Add Target

Target Name:* This field is required.

Target Alias:

iSCSI GROUP

Portal Group ID:*

Initiator Group ID:*

OK Cancel

Fig. 11.16: Creating a Target

Type a target name. If desired, an alias for the target can also be provided. Choose a Portal Group ID and Initiator Group ID from the pulldown menus. Click OK to save and add the new target to the list. If not immediately visible, click the Refresh button to repopulate the list of targets.

Highlight an existing target and click the Edit button to open the Edit Target window, shown in [Figure 11.17](#).

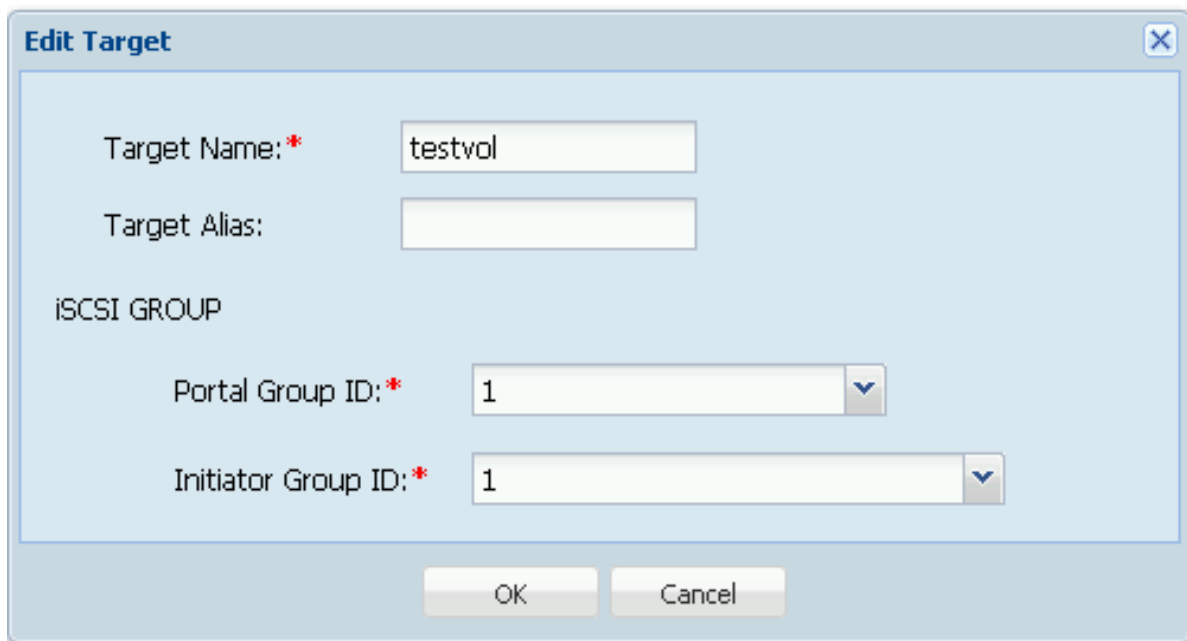
The image shows a dialog box titled "Edit Target" with a close button (X) in the top right corner. Inside the dialog, there are four input fields. The first is "Target Name:" with a red asterisk, containing the text "testvol". The second is "Target Alias:" which is empty. Below these is the section header "iSCSI GROUP". The third field is "Portal Group ID:" with a red asterisk, containing the value "1" and a dropdown arrow. The fourth field is "Initiator Group ID:" with a red asterisk, also containing the value "1" and a dropdown arrow. At the bottom of the dialog are two buttons: "OK" and "Cancel".

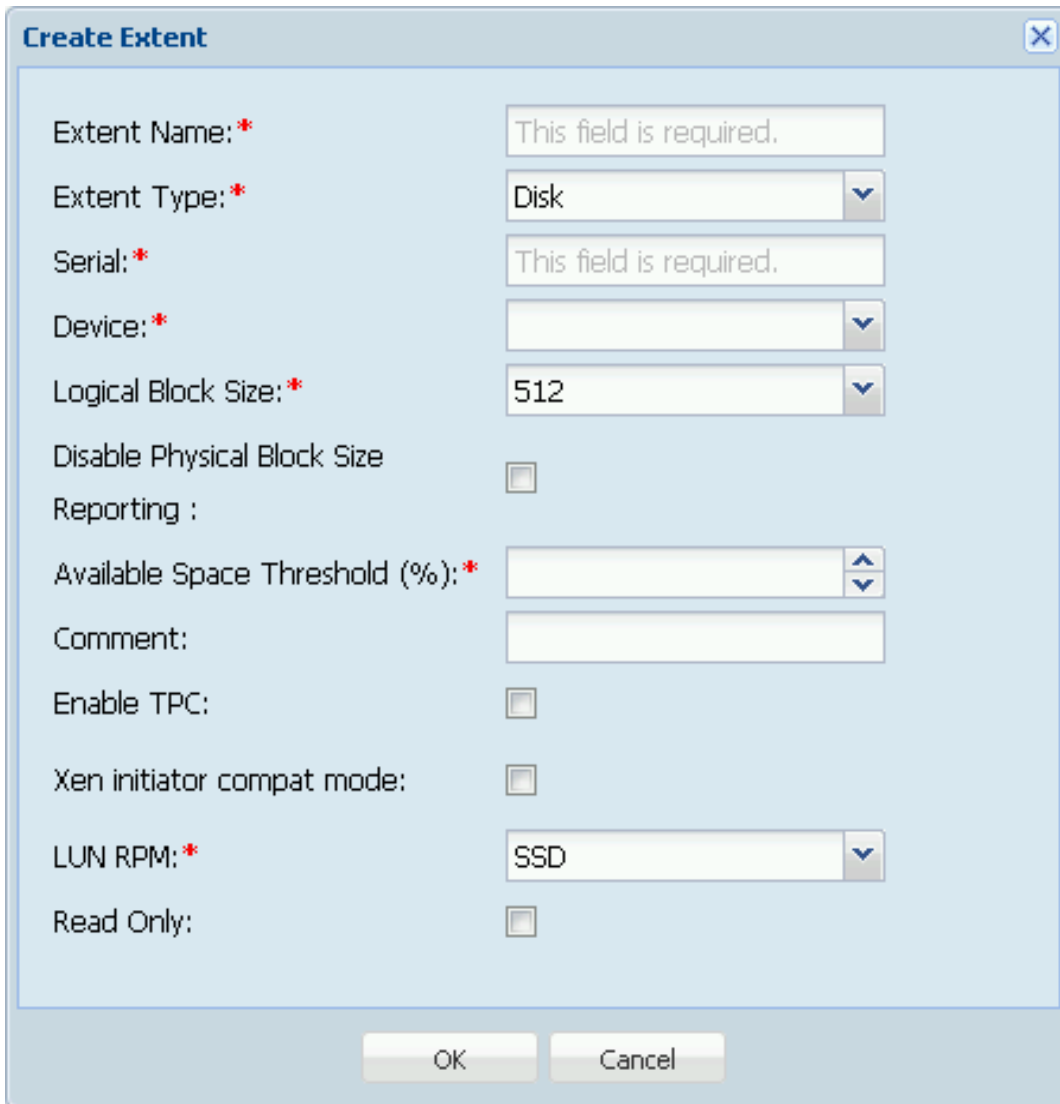
Fig. 11.17: Editing a Target

The current settings are displayed. Modify the values as needed, then click OK to save the new settings.

11.10.2 Extents

Note: The vCenter plugin only supports creating device/disk extents.

Click the + button to begin creating a new [extent](#) (page 177). [Figure 11.18](#) shows the Create Extent window.

A screenshot of the 'Create Extent' dialog box. It has a title bar with 'Create Extent' and a close button. The dialog contains several fields and checkboxes. 'Extent Name: *' is a text field with 'This field is required.' inside. 'Extent Type: *' is a dropdown menu with 'Disk' selected. 'Serial: *' is a text field with 'This field is required.' inside. 'Device: *' is a dropdown menu. 'Logical Block Size: *' is a dropdown menu with '512' selected. 'Disable Physical Block Size Reporting' is a checkbox. 'Available Space Threshold (%): *' is a text field with up and down arrows. 'Comment:' is a text field. 'Enable TPC:' is a checkbox. 'Xen initiator compat mode:' is a checkbox. 'LUN RPM: *' is a dropdown menu with 'SSD' selected. 'Read Only:' is a checkbox. At the bottom are 'OK' and 'Cancel' buttons.

Extent Name: *	This field is required.
Extent Type: *	Disk
Serial: *	This field is required.
Device: *	
Logical Block Size: *	512
Disable Physical Block Size Reporting	<input type="checkbox"/>
Available Space Threshold (%): *	
Comment:	
Enable TPC:	<input type="checkbox"/>
Xen initiator compat mode:	<input type="checkbox"/>
LUN RPM: *	SSD
Read Only:	<input type="checkbox"/>

Fig. 11.18: Creating a new extent

The Extent Name, Extent Type, Serial, Device, Logical Block Size, Available Space Threshold, and LUN RPM options must be set before the extent can be created.

Note: Detailed explanations of each option are available in [Extent Configuration Settings](#) (page 179).

Click OK to save the settings and begin creating the extent. Enter the data, then click the Refresh button to repopulate the list of extents.

Highlight an extent in the list and click Edit. The Edit Extent window opens, which is the same as [Create Extent](#) (page 239), but populated with the current extent settings. Make any changes to the options, then click OK to save the new configuration.

11.10.3 Associated Targets

The [Target/Extents](#) (page 180) section details the process of associating an extent with a target. vCenter requires previously created targets and extents before a new associated target can be created.

Click the + button to open the Add Target/Extent window. Using the drop-down menus, choose a target and extent to associate. It is recommended to leave the LUN ID set to Auto, but values from 0-24 are available.

Click OK to save the new associated target. The system may take some time to create the association. Click the Refresh button after a few moments to verify the new associated target.

Highlight an entry in the list and click the Edit button. The Target/Extent window appears, which allows modification of the Target, Extent, and LUN ID. Click OK to save any changes.

11.11 Role-Based Access Control (RBAC)

An administrator can grant vCenter users specific role-based access to a TrueNAS® system. The roles are listed in

Table 11.2: Roles

Role Name	User is allowed to:
Discovery	Add storage systems
Create Clones	Clone virtual machines
Create Storage	Create and provision storage
Modify Storage	Modify storage
Destroy Storage	Destroy storage

Each role gives the user the ability to perform the functions in not just that role, but all of the roles that precede it in the list. For example, a user with a Create Storage role can create storage, but can also do Discovery and Create Clones. The vCenter administrator can perform all of these operations by default.

Note: The vCenter plugin RBAC tab cannot create new vCenter users. That must be done from vCenter Administration. Hover over the House + lines icon on the top section of the window to view a drop-down menu of links. Click Home. From the Home menu, navigate from Administration → Single Sign-On and click Users and Groups. In the Users tab, click the + button to open the New User window. Enter a user name and password (twice) in the related fields. Optional fields can further describe the user, including first and last name and an email address for the account. Click OK to create the new user.

11.11.1 Add a Role to an Existing vCenter User

Click the + button to open the Add Role Based Access Control window. Type a user name in the form DOMAIN.NAME\username. Click the Assign Role drop-down menu and choose a role for the user. Click Add to add the associated user and role to the list. If the entry does not appear immediately, click the Refresh button to repopulate the list.

Highlight an entry in the list and click the Edit button. The Edit window opens, and the entry can be edited. Click the Save button to save any changes. Changes may take a moment to redraw. Click the Refresh button to repopulate the list.

11.12 Copying iSCSI Configuration

The TrueNAS Configuration tab is used to duplicate some of the current TrueNAS® system's iSCSI settings to another TrueNAS® system.

Note: Currently, only iSCSI targets are duplicated on the destination system.

Enter the IP address of the destination TrueNAS® system where the current TrueNAS® system's iSCSI configuration is to be copied. Provide a valid user name and password from the destination system. Click Submit to begin duplicating the TrueNAS® configuration to the destination system. This process may take some time. A popup window indicates when the task is complete.

OPENSTACK CINDER DRIVER

The TrueNAS® Cinder driver for OpenStack provides an iSCSI (Block) Storage service for OpenStack. Detection and management of connected TrueNAS® iSCSI resources is available from within the OpenStack web interface.

12.1 Requirements

1. A TrueNAS® 11.0 or newer system to provide storage.
2. A computer running the [Newton release](https://docs.openstack.org/newton/) (<https://docs.openstack.org/newton/>) of OpenStack or newer, with these minimum hardware requirements:
 - 4 GiB of RAM
 - 4-core CPU
 - 40 GiB storage disk

12.2 Installation

The Cinder driver is packaged as a .zip file for simple download and extraction. To obtain the driver, please contact iXsystems Customer Support at truenas-support@ixsystems.com.

Log in to a console on the OpenStack system. Switch user to stack with `su - stack`. Extract the contents of the Cinder driver .zip file to `/opt/stack/cinder/cinder/volume/drivers/`.

Note: Be sure the stack user has ownership or permissions for all the extracted Cinder files.

Log in to the TrueNAS® system and [create a new dataset](#) (page 96) named `cinder`. Note the TrueNAS® IP address, root username, and password. Locate the base name of the iSCSI target by navigating to `Sharing → Block(iSCSI) → Target → Global Configuration → Base Name`. This information is needed to configure the Cinder driver before it can be used in the OpenStack interface.

12.3 Configuration

On the OpenStack computer, edit `/etc/cinder/cinder.conf`. Some sections of this file must be edited and some user-defined parameters must be added.

Modify the existing lines in `cinder.conf` to these values:

```
default_volume_type = ixsystems-iscsi
enabled_backends = ixsystems-iscsi, lvmdriver-1
```

Table 12.1 explains the added parameters and their possible values:

Table 12.1: Configuration Parameters for `cinder.conf`

Parameter	Description	Possible Values
<code>iscsi_helper</code>	iSCSI target user-land tool; the default is recommended	<code>tgtadm</code> (default), <code>lioadm</code> , <code>scstadmin</code> , <code>iscsictl</code> , <code>ietadm</code> , <code>fake</code> (testing only)
<code>volume_dd_blocksize</code>	Default block size used when copying or clearing volumes	Block size string value: 512, 1M
<code>volume_driver</code>	Driver used in volume creation; provide the path to the driver	Default path for the iX driver: <code>cinder.volume.drivers.ixsystems.iscsi.FreeNASISCSIDriver</code>
<code>ixsystems_login</code>		Username of the connected host
<code>ixsystems_password</code>		Password of the connected host
<code>ixsystems_server_hostname</code>		IP address of the host
<code>ixsystems_volume_backend_name</code>	Backend name for a given driver implementation	The default is <code>ixsystems_FREENAS_Storage</code>
<code>ixsystems_iqn_prefix</code>	Base name of the iSCSI target	User defined. Found in the web interface by navigating to <code>Sharing -> Block(iscsi) -> Target Global Configuration -> Base Name</code>
<code>ixsystems_datastore_pool</code>	Name of the dataset on the connected system	If the cinder dataset is created, the value is <code>cinder-tank</code>
<code>ixsystems_vendor_name</code>		<code>iXsystems</code>
<code>ixsystems_storage_protocol</code>		<code>iscsi</code>

Here is an example of typical additional parameter settings for a TrueNAS® system at IP address 10.3.1.81:

```
[ixsystems-iscsi]
iscsi_helper = tgtadm
volume_dd_blocksize = 512
volume_driver = cinder.volume.drivers.ixsystems.iscsi.FreeNASISCSIDriver
ixsystems_login = root
ixsystems_password = thisisdummyspassword
ixsystems_server_hostname = 10.3.1.81
ixsystems_volume_backend_name = ixsystems_FREENAS_Storage
ixsystems_iqn_prefix = iqn.2005-10.org.freenas.ctl
ixsystems_datastore_pool = cinder-tank
ixsystems_vendor_name = iXsystems
ixsystems_storage_protocol = iscsi
```

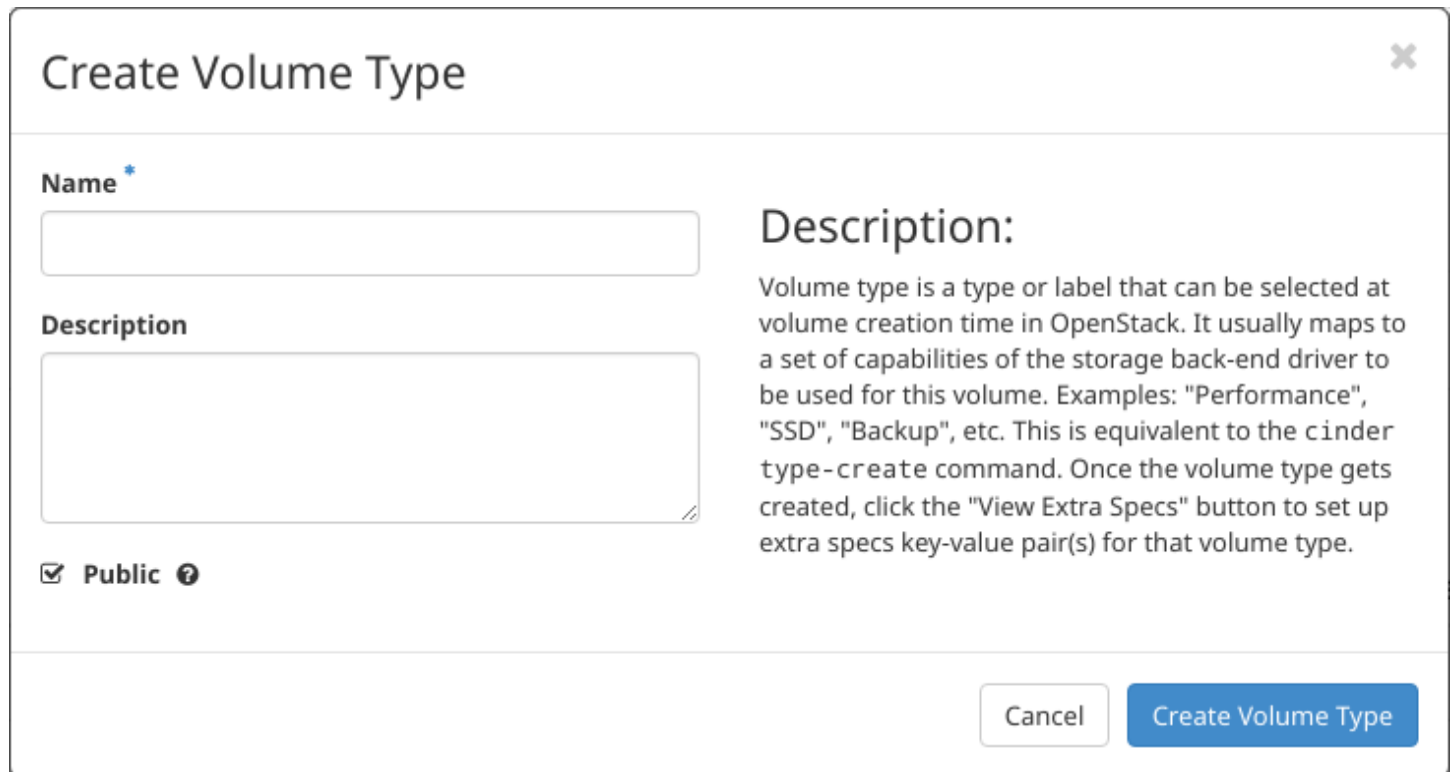
Tip: The [OpenStack documentation](https://docs.openstack.org/mitaka/config-reference/block-storage/block-storage-sample-configuration-files.html) (https://docs.openstack.org/mitaka/config-reference/block-storage/block-storage-sample-configuration-files.html) also provides examples of `cinder.conf` configurations.

Restart the OpenStack computer to reload `cinder.conf` and initialize the Cinder service properly. After this reboot, the Cinder service can be restarted manually:

```
# /usr/local/bin/cider-volume --config-file /etc/cinder/cinder.conf & echo $!
>/opt/stack/status/stack/c-vol.pid; fg || echo "c-vol failed to start" | tee
"/opt/stack/status/stack/c-vol.failure".
```

The Cinder driver is now available in the OpenStack web interface.

Open a web browser, go to the IP address of the OpenStack system, and log in with the OpenStack user name and password. Click Connect. Navigate from Admin → System → Volumes → Volume Types and click Create Volume Type. The window shown in Figure 12.1 appears.



Create Volume Type [X]

Name *

Description

☒ **Public** ?

Description:

Volume type is a type or label that can be selected at volume creation time in OpenStack. It usually maps to a set of capabilities of the storage back-end driver to be used for this volume. Examples: "Performance", "SSD", "Backup", etc. This is equivalent to the cinder type-create command. Once the volume type gets created, click the "View Extra Specs" button to set up extra specs key-value pair(s) for that volume type.

Cancel Create Volume Type

Fig. 12.1: Creating the ixsystems-iscsi Volume Type

Enter `ixsystems-iscsi` in the Name field and ensure the Public box is checked. An optional Description field is also available. Click Create Volume Type to create the volume type. The list is automatically refreshed when the task is finished.

12.4 Using the Cinder Driver

Cinder is used to manage volumes in the Project and Admin sections of the OpenStack interface.

12.4.1 Managing Volumes with Cinder by Project

The Project → Compute → Volumes section manages volumes on a per-project basis. To change projects, click the project name dropdown menu displayed in the upper left area of the interface, next to the openstack logo.

Figure 12.2 shows an example of the Project/Compute/Volumes interface with the Volumes tab visible:

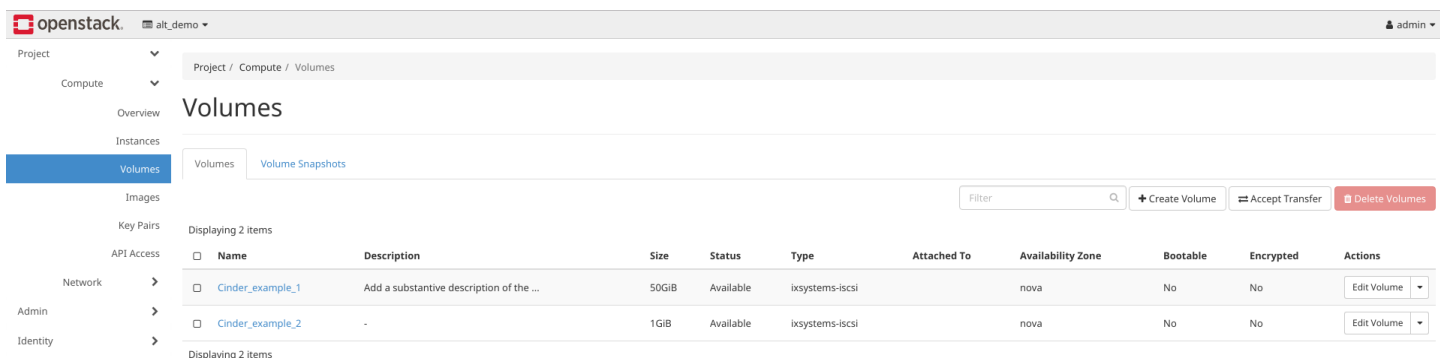


Fig. 12.2: Interface for Managing Project Volumes

Every volume in the project displays in a reference list. The list is sortable by any of the column titles. Users can also type custom filters to locate specific items. Options to Create, Accept Transfer of a volume from another project, and Delete Volumes are available.

An existing volume can be modified by checking the checkbox next to the volume Name, then opening the Actions pulldown menu. Table 12.2 lists the action, the equivalent `cinder` CLI command, and a description of the configurable settings:

Table 12.2: OpenStack Single Volume Actions

Action	Configurables
Edit Volume	Adjust name and description of a volume, or make a volume Bootable
Extend Volume	Enter a new size in GiB
Manage Attachments	View and adjust volume attachment to instances
Create Snapshot	Enter a snapshot name and description; snapshot limits are shown
Change Volume Type	Choose a new type and migration policy
Upload to Image	Enter an image name and choose a disk format: QCOW2, Raw, VDI, or VMDK
Create Transfer	Enter recipient project name
Update Metadata	Move items into the existing metadata column or create a custom metadata key

Project Volume Snapshots

Any snapshot of a volume displays in the Volume Snapshots tab, seen in Figure 12.3:

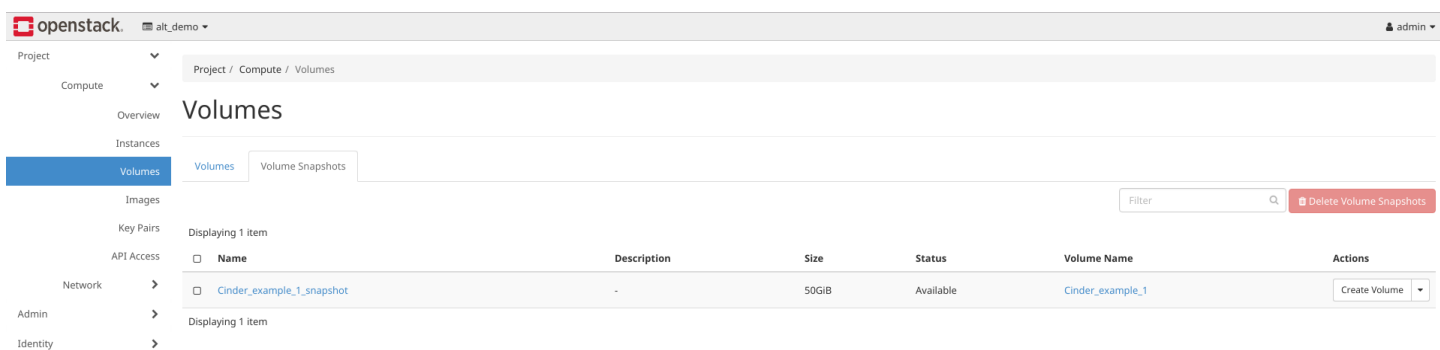


Fig. 12.3: Project Volume Snapshots Tab

As with the other interface elements, the list can be sorted and filtered. The Actions drop-down menu has options to create a volume from an existing snapshot, edit a snapshot name and description, delete snapshots, and update snapshot metadata.

12.4.2 Administering All Volumes with Cinder

Cinder functions within the Admin section manage volumes, volume types, and volume snapshots for all projects on the system. Navigate from Admin → System → Volumes to see these management tabs. Every created item is displayed in the master list regardless of project. This list can be sorted or filtered as needed.

Managing Volumes

Figure 12.4 shows the Volumes tab:

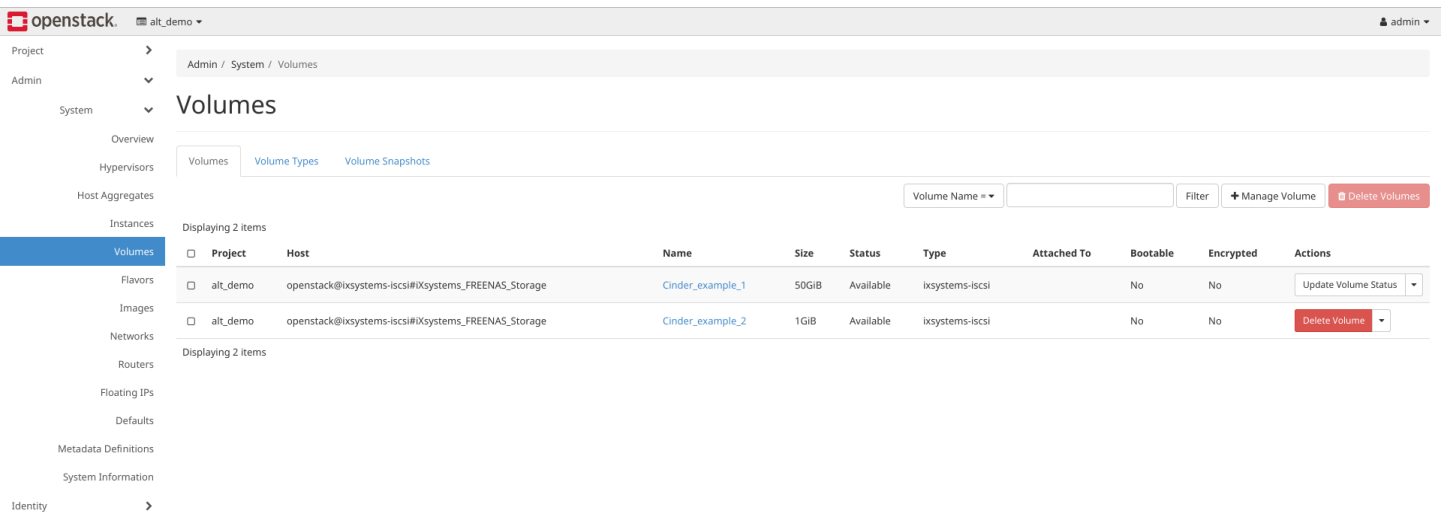


Fig. 12.4: OpenStack Volumes Tab

Table 12.3 shows the different options, the `cinder` CLI equivalents, and a description of the configurable elements.

Tip: Volumes can be managed or deleted in this section, but not created. Volume creation is handled in the Project → Compute → Volumes tab.

Table 12.3: Cinder Volume Management Options

Action	Configurables
Manage Volumes	Type the volume identifier, host, and volume name to make it visible in OpenStack; other configurable elements are Volume Name, Description, Metadata, Volume Type, Availability Zone, and Bootable
Delete Volumes	Warning: volumes deleted here cannot be recovered
Update Volume Status	Select a new status from the dropdown menu; volume status is normally set automatically
Unmanage Volume	Unmanaging a volume makes it invisible in OpenStack, but does not delete it
Migrate Volume	Choose a new Destination Host from the dropdown menu; there is also an option to Force Host Copy, which bypasses driver optimizations
Update Metadata	Choose items to move to the Existing Metadata column; custom keys can also be added

Manage Volume Types

Figure 12.5 shows how all created volume types are listed in this tab. If the `cinder.conf` file was edited as shown above in the Cinder driver installation, two types are visible: `ixsystems-iscsi` and `lvmdriver-1`.

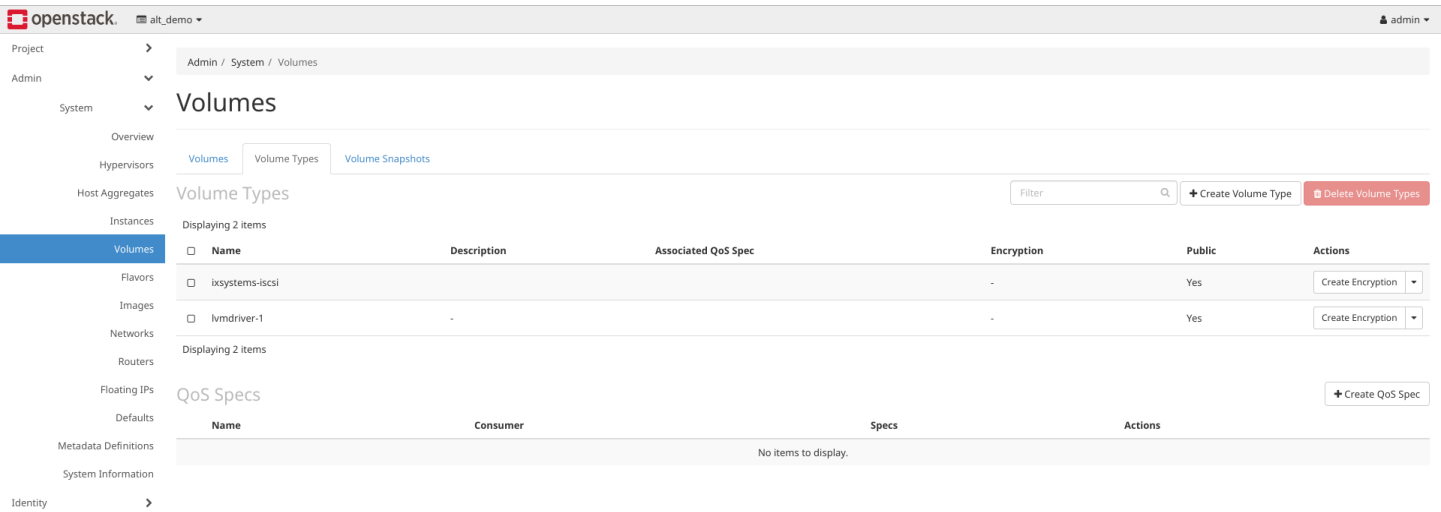


Fig. 12.5: OpenStack Volume Types Tab

A number of actions can be applied to either volume type, as Table 12.4 shows:

Table 12.4: Volume Type Actions

Action	Configurables
Create Volume Type	Name, description, and visible to “Public”
Delete Volume Types	Warning: action is permanent
Create QoS Spec	Provide a name and consumer for the new Quality of Service spec; choose back-end to associate the policy with Cinder
Create Encryption	Type a provider and control location; custom cipher and key size can also be designated
View/Create Extra Specs	Click Create to enter and save a new key/value pair
Manage QoS Spec Association	Select a QoS spec from the drop-down menu
Edit Volume Type	Modify the volume type name, description, and public fields
Update Metadata	Add new metadata items to the volume type; custom keys can be created

Manage Volume Snapshots

This tab is very similar to managing snapshots of a single project. The only difference is that snapshots from all projects are visible in the list. Figure 12.6 shows an example:

openstack

alt_demo

admin

Project

Admin

System

Overview

Hypervisors

Host Aggregates

Instances

Volumes

Flavors

Images

Networks

Routers

Floating IPs

Defaults

Metadata Definitions

System Information

Identity

Admin / System / Volumes

Volumes

Volumes

Volume Types

Volume Snapshots

Filter

Delete Volume Snapshots

Displaying 1 item

Project	Host	Name	Description	Size	Status	Volume Name	Actions
alt_demo	openstack@ixsystems-iscsi@ixsystems_FREENAS_Storage	Cinder_example_1_snapshot	-	50GiB	Available	Cinder_example_1	Delete Volume Snapshot

Displaying 1 item

Fig. 12.6: Manage All Projects Volume Snapshots

Options to Delete, Update Status, and Update Metadata for each listed snapshot are available.

REPORTING

Reporting displays several graphs, as seen in the example in Figure 13.1. Click the tab for a device type to see its graphs.

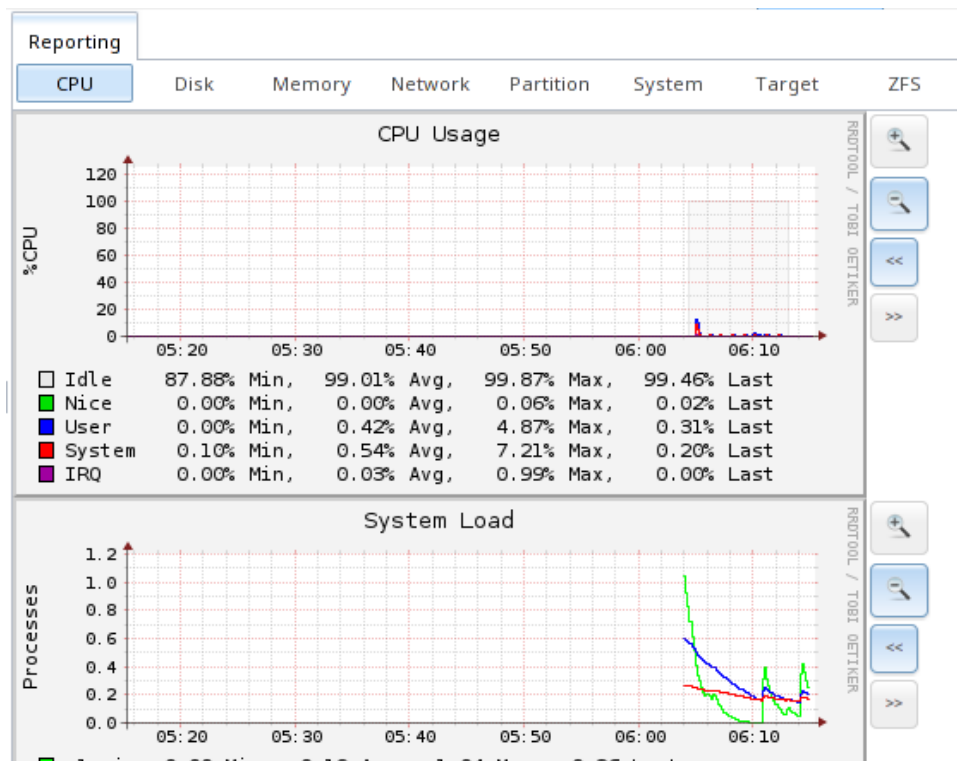


Fig. 13.1: Reporting Graphs

TrueNAS® uses [collectd](https://collectd.org/) (<https://collectd.org/>) to provide reporting statistics. The resulting graphs are grouped into several tabs on the Reporting page:

- CPU
 - CPU (<https://collectd.org/wiki/index.php/Plugin:CPU>) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle.
- Disk
 - Disk (<https://collectd.org/wiki/index.php/Plugin:Disk>) shows statistics on I/O, percent busy, latency, operations per second, and pending I/O requests.
- Memory
 - Memory (<https://collectd.org/wiki/index.php/Plugin:Memory>) displays memory usage.
 - Swap (<https://collectd.org/wiki/index.php/Plugin:Swap>) displays the amount of free and used swap space.

- Network
 - [Interface](https://collectd.org/wiki/index.php/Plugin:Interface) (<https://collectd.org/wiki/index.php/Plugin:Interface>) shows received and transmitted traffic in bits per second for each configured interface.
- Partition
 - [Disk space](https://collectd.org/wiki/index.php/Plugin:DF) (<https://collectd.org/wiki/index.php/Plugin:DF>) displays free and used space for each volume and dataset. However, the disk space used by an individual zvol is not displayed as it is a block device.
- System
 - [Processes and Uptime](https://collectd.org/wiki/index.php/Plugin:Processes) (<https://collectd.org/wiki/index.php/Plugin:Processes>) displays the number of processes, grouped by state.
 - [Uptime](https://collectd.org/wiki/index.php/Plugin:Uptime) (<https://collectd.org/wiki/index.php/Plugin:Uptime>) keeps track of the system uptime, the average running time, and the maximum reached uptime.
- Target
 - Target shows bandwidth statistics for iSCSI ports.
- ZFS
 - [ZFS](https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) (https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) shows ARC size, hit ratio, and requests.

Reporting data is saved to permit viewing and monitoring usage trends over time. This data is preserved across system upgrades and restarts.

Data files are saved in `/var/db/collectd/rrd/`.

The reporting data file recording method is controlled by the `System → System Dataset Reporting database` checkbox. When unchecked, data files are recorded in a temporary filesystem and copied hourly to on-disk files.

When `System → System Dataset Reporting database` is checked, data files are written directly to the [System Dataset](#) (page 36).

Warning: Reporting data is frequently written and should not be stored on the boot pool or boot device.

Use the magnifier buttons next to each graph to increase or decrease the displayed time increment from 10 minutes, hourly, daily, weekly, or monthly. The << and >> buttons can be used to scroll through the output.

[Update on using Graphite with FreeNAS](http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html) (<http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html>) contains instructions for sending the collected information to a [Graphite](http://graphite.wikidot.com/) (<http://graphite.wikidot.com/>) server.

WIZARD

TrueNAS® provides a wizard which helps complete the steps needed to quickly configure TrueNAS® for serving data over a network. The wizard can be run at any time by clicking the Wizard icon.

Figure 14.1 shows the first wizard configuration screen.




Fig. 14.1: Configuration Wizard

Note: You can exit the wizard at any time by clicking the Exit button. However, exiting the wizard will not save any selections. The wizard can always be run again by clicking the Wizard icon. Alternately, the TrueNAS® GUI can be used to configure the system, as described in the rest of this Guide.

This first screen can be used to change the default language, keyboard map, and timezone. After making your selections, click Next.

Note: Typically, a TrueNAS® system ships with pre-configured volumes. The screens shown in Figure 14.2 and Figure 14.3 will only appear if unformatted disks are available or the system has been reinstalled.

Figure 14.2 shows the configuration screen that appears if the storage disks have not yet been formatted.

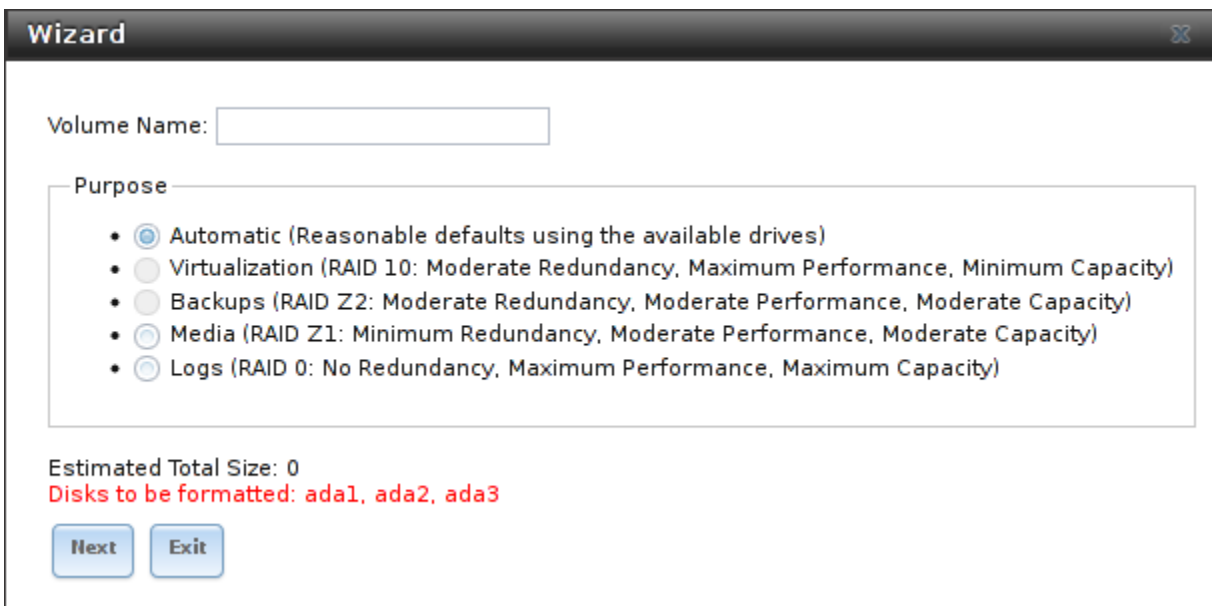


Fig. 14.2: Volume Creation Wizard

Note: The wizard will not recognize an encrypted ZFS pool. If your ZFS pool is GELI-encrypted, cancel the wizard and use the instructions in [Importing an Encrypted Pool](#) (page 101) to import the encrypted volume. You can then rerun the wizard afterwards, if you wish to use it for post-configuration, and it will recognize that the volume has been imported and will not prompt to reformat the disks.

Enter a name for the ZFS pool that conforms to these [naming conventions](http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html) (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html). It is recommended to choose a name that will stick out in the logs (e.g. not `data` or `truenas`).

Decide if the pool should provide disk redundancy, and if so, which type. The [ZFS Primer](#) (page 263) discusses RAIDZ redundancy in more detail. If you prefer to make a more complex configuration, click the Exit button to close the wizard and instead use [Volume Manager](#) (page 89).

These redundancy types are available:

- Automatic: automatically creates a mirrored, RAIDZ1, or RAIDZ2 pool, depending upon the number of disks. If you prefer to control the type of redundancy, select one of the other options.
- RAID 10: creates a striped mirror and requires a minimum of 4 disks.
- RAIDZ2: requires a minimum of 4 disks. Up to 2 disks can fail without data loss.
- RAIDZ1: requires a minimum of 3 disks. Up to 1 disk can fail without data loss.
- Stripe: requires a minimum of 1 disk. Provides no redundancy, meaning if any of the disks in the stripe fails, all data in the stripe is lost.

Once you have made your selection, click Next to continue.

If the system has been reinstalled and the disks are formatted as an unencrypted ZFS pool, a screen to import the volume will appear. This screen is shown in [Figure 14.3](#).



Fig. 14.3: Volume Import Screen

Select the existing volume from the drop-down menu and click Next to continue.

The next screen in the wizard is shown in [Figure 14.4](#).

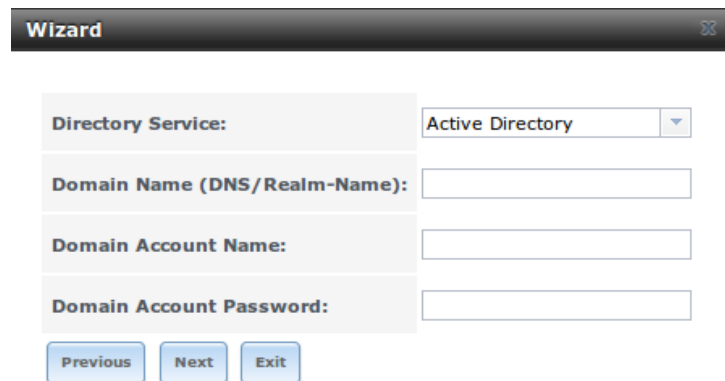


Fig. 14.4: Directory Service Selection

If the TrueNAS® system is on a network that does not contain an Active Directory, LDAP, or NIS server, click Next to skip to the next screen.

However, if the TrueNAS® system is on a network containing an Active Directory, LDAP, or NIS server and you wish to import the users and groups from that server, select the type of directory service in the Directory Service drop-down menu. The rest of the fields in this screen will vary, depending upon which directory service is selected. Available configuration options for each directory service are summarized in [Tables 14.1 through 14.3](#).

Note: Additional configuration options are available for each directory service. The wizard can be used to set the initial values required to connect to that directory service. You can then review the other available options in [Directory Services](#) (page 132) to determine if additional configuration is required.

Table 14.1: Active Directory Options

Setting	Value	Description
Domain Name	string	name of Active Directory domain (e.g. example.com) or child domain (e.g. sales.example.com)
Domain Account Name	string	name of the Active Directory administrator account
Domain Account Password	string	password for the Active Directory administrator account

Table 14.2: LDAP Options

Setting	Value	Description
Hostname	string	hostname or IP address of LDAP server
Continued on next page		

Table 14.2 – continued from previous page

Setting	Value	Description
Base DN	string	top level of the LDAP directory tree to be used when searching for resources (e.g. dc=test,dc=org)
Bind DN	string	name of administrative account on LDAP server (e.g. cn=Manager,dc=test,dc=org)
Base password	string	password for

Table 14.3: NIS Options

Setting	Value	Description
NIS domain	string	name of NIS domain
NIS servers	string	comma delimited list of hostnames or IP addresses
Secure mode	checkbox	if checked, ypbind(8) (http://www.freebsd.org/cgi/man.cgi?query=ypbind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, ypbind will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet

The next configuration screen, shown in [Figure 14.5](#), is used to create network shares.

Fig. 14.5: Network Shares

TrueNAS® supports several types of shares for providing storage data to the clients in a network. The initial wizard can be

used to quickly make shares using default permissions which should “just work” for common scenarios. For more complex scenarios, refer to the section on [Sharing](#) (page 143).

To create a share using the wizard, enter a name for the share, then select the Purpose of the share:

- Windows (SMB): this type of share can be accessed by any operating system using a SMB client. Check the box for Allow Guest to allow users to access the share without a password. SMB shares created with the wizard can be fine-tuned afterward with [Windows \(SMB\) Shares](#) (page 159).
- Mac OS X (AFP): this type of share can be accessed by Mac OS X users. Check the box for Time Machine if Mac users will be using the TrueNAS® system as a backup device. AFP shares created with the wizard can be fine-tuned afterward with [Apple \(AFP\) Shares](#) (page 144).
- Generic Unix (NFS): this type of share can be accessed by any operating system using a NFS client. NFS shares created using the wizard can be fine-tuned afterward with [Unix \(NFS\) Shares](#) (page 151).
- Block Storage (iSCSI): this type of share can be accessed by any operating system using iSCSI initiator software. Enter the size of the block storage to create in the format 20G (for 20 GB). iSCSI shares created with the wizard can be fine-tuned afterward with [iSCSI](#) (page 199).

After selecting the Purpose, click the Ownership button to see the screen shown in [Figure 14.6](#).

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

Fig. 14.6: Share Permissions

The default permissions for the share are displayed. To create a user or group, enter the desired name, then check the Create User box to create that user and the Create Group box to create the group. Check or uncheck the boxes in the Mode section to set the initial access permissions for the share. When finished, click the Return button to return to the share creation screen. Click the Add button to finish creating that share, which will then appear in the Name frame.

The Delete button can be used to remove the share highlighted in the Name frame. To edit a share, highlight it, make the change, then press the Update button.

When finished making shares, click the Next button to advance to the screen shown in [Figure 14.7](#).

Console messages:	<input type="checkbox"/>	i
Root E-mail:	<input type="text"/>	i
From email:	<input type="text" value="root@freenas.local"/>	i
Outgoing mail server:	<input type="text"/>	i
Port to connect to:	<input type="text" value="25"/>	i
TLS/SSL:	<input type="text" value="Plain"/>	i
Use SMTP Authentication:	<input type="checkbox"/>	
Username:	<input type="text"/>	i
Password:	<input type="password"/>	
Password confirmation:	<input type="password"/>	i

Previous Send Test Mail Next Exit

Fig. 14.7: Miscellaneous Settings

This screen can be used to configure these settings:

- **Console messages:** check this box if you would like to view system messages at the bottom of the graphical administrative interface. This can be handy when troubleshooting a service that will not start. When using the console message view, if you click the console messages area, it will pop-up as a window, allowing you to scroll through the output and to copy its contents.
- **Root E-mail:** TrueNAS® provides an “Alert” icon in the upper right corner to provide a visual indication of events that warrant administrative attention. The alert system automatically emails the root user account whenever an alert is issued. It is important to enter the email address of the person to receive these alerts and other administrative emails. The rest of the email settings in this screen should also be reviewed and edited as necessary. Before leaving this screen, click the “Send Test Mail” button to ensure that email notifications are working correctly.
- **From email:** the from email address to use when sending email notifications.
- **Outgoing mail server:** hostname or IP address of SMTP server.
- **Port to connect to:** port number used by the SMTP server.
- **TLS/SSL:** encryption type used by the SMTP server.
- **Use SMTP Authentication:** check this box if the SMTP server requires authentication.
- **Username:** enter the username if the SMTP server requires authentication.
- **Password:** enter the password if the SMTP server requires authentication.

When finished, click Next. A message will indicate that the wizard is ready to perform all of the saved actions. To make changes, click the Return to Wizard button to review your edits. If you click the Exit without saving button, none of your selections will be saved. To save your edits, click the Confirm button. A status bar will indicate when the wizard has completed applying the new settings.

In addition to the settings that you specify, the wizard will automatically enable [S.M.A.R.T. Tests](#) (page 73), create a boot environment, and add the new boot environment to the boot menu. If you also wish to save a backup of the configuration database to the system being used to access the administrative graphical interface, go to *System* → *General*, click the

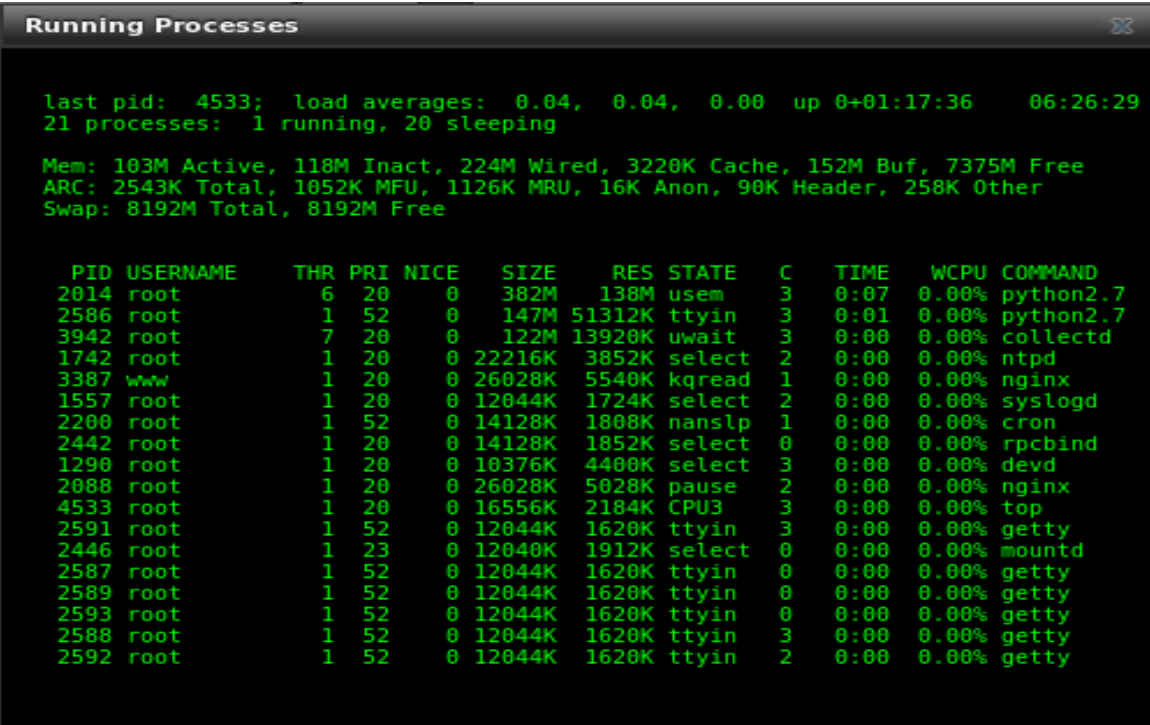
Save Config button, and browse to the directory where the configuration will be saved. Always back up your configuration after making any configuration changes.

ADDITIONAL OPTIONS

This section covers the remaining miscellaneous options available from the TrueNAS® graphical administrative interface.

15.1 Display System Processes

Clicking Display System Processes opens a screen showing the output of `top(1)` (<http://www.freebsd.org/cgi/man.cgi?query=top>). An example is shown in Figure 15.1.



```

Running Processes

last pid: 4533; load averages: 0.04, 0.04, 0.00 up 0+01:17:36 06:26:29
21 processes: 1 running, 20 sleeping

Mem: 103M Active, 118M Inact, 224M Wired, 3220K Cache, 152M Buf, 7375M Free
ARC: 2543K Total, 1052K MFU, 1126K MRU, 16K Anon, 90K Header, 258K Other
Swap: 8192M Total, 8192M Free

  PID USERNAME   THR PRI NICE   SIZE    RES STATE  C  TIME  WCPU COMMAND
  2014 root         6  20   0    382M    138M usem   3   0:07  0.00% python2.7
  2586 root         1  52   0    147M   51312K ttyin  3   0:01  0.00% python2.7
  3942 root         7  20   0    122M   13920K uwait   3   0:00  0.00% collectd
  1742 root         1  20   0   22216K   3852K select  2   0:00  0.00% ntpd
  3387 www         1  20   0   26028K   5540K kqread  1   0:00  0.00% nginx
  1557 root         1  20   0   12044K   1724K select  2   0:00  0.00% syslogd
  2200 root         1  52   0   14128K   1808K nanslp  1   0:00  0.00% cron
  2442 root         1  20   0   14128K   1852K select  0   0:00  0.00% rpcbind
  1290 root         1  20   0   10376K   4400K select  3   0:00  0.00% devd
  2088 root         1  20   0   26028K   5028K pause   2   0:00  0.00% nginx
  4533 root         1  20   0   16556K   2184K CPU3    3   0:00  0.00% top
  2591 root         1  52   0   12044K   1620K ttyin  3   0:00  0.00% getty
  2446 root         1  23   0   12040K   1912K select  0   0:00  0.00% mountd
  2587 root         1  52   0   12044K   1620K ttyin  0   0:00  0.00% getty
  2589 root         1  52   0   12044K   1620K ttyin  0   0:00  0.00% getty
  2593 root         1  52   0   12044K   1620K ttyin  0   0:00  0.00% getty
  2588 root         1  52   0   12044K   1620K ttyin  3   0:00  0.00% getty
  2592 root         1  52   0   12044K   1620K ttyin  2   0:00  0.00% getty
  
```

Fig. 15.1: System Processes Running on TrueNAS®

The display automatically refreshes itself. Click the X in the upper right corner to close the display when finished. This display is read-only, so it is not possible to give a `kill` command in it.

15.2 Shell

The TrueNAS® GUI provides a web shell, making it convenient to run command line tools from the web browser as the root user. The link to Shell is the fourth entry from the bottom of the menu tree. In Figure 15.2, the link has been clicked and

Shell is open.




Fig. 15.2: Web Shell

The prompt indicates that the current user is root, the hostname is truenas, and the current working directory is ~ (root's home directory).

To change the size of the shell, click the 80x25 drop-down menu and select a different size.

To copy text from shell, highlight the text, right-click, and select Copy from the right-click menu. To paste into the shell, click the Paste button, paste the text into the box that opens, and click the OK button to complete the paste operation.

While you are in Shell, you will not have access to any of the other GUI menus. If you need to have access to a prompt while using the GUI menus, use `tmux` instead as it supports multiple shell sessions and the detachment and reattachment of sessions.

Shell provides history (use your up arrow to see previously entered commands and press `Enter` to repeat the currently displayed command) and tab completion (type a few letters and press `tab` to complete a command name or filename in the current directory). When you are finished using Shell, type `exit` to leave the session.

Note: Not all of Shell's features render correctly in Chrome. Firefox is the recommended browser for using Shell.

Most FreeBSD command line utilities are available in Shell.

15.3 Log Out

Click the Log Out entry in the tree to log out of the TrueNAS® GUI. This causes an immediate logout. A message is displayed with a link to log back in.

15.4 Reboot

Click Reboot shows the warning message in Figure 15.3. The browser window background color changes to red to indicate that this option can negatively impact users of the TrueNAS® system.

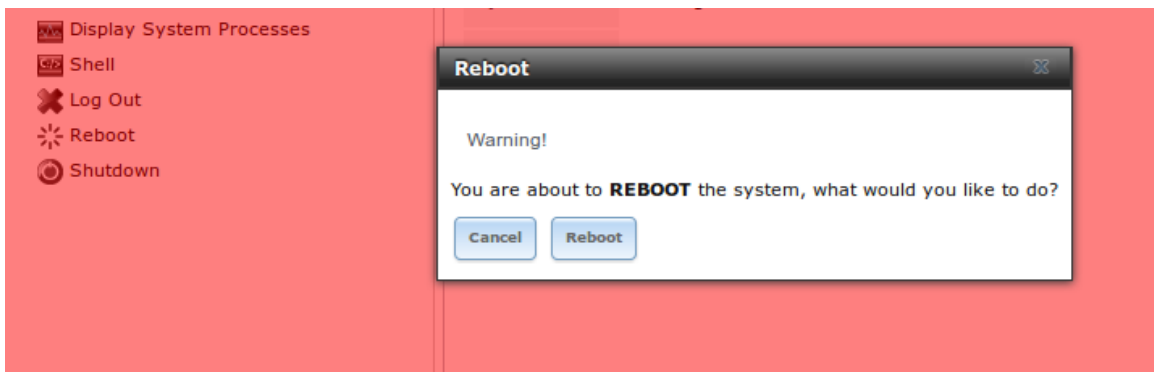


Fig. 15.3: Reboot Warning Message

If a scrub or resilver is in progress when a reboot is requested, an additional warning asks if you wish to proceed. In this case, it is recommended to Cancel the reboot request and to periodically run `zpool status` from [Shell](#) (page 257) until it is verified that the scrub or resilver process is complete. Once complete, the reboot request can be re-issued.

Click the Cancel button to cancel the reboot request. Otherwise, click the Reboot button to reboot the system. Rebooting the system disconnects all clients, including the web administration GUI. The URL in the web browser changes, adding `/system/reboot/` to the end of the IP address. Wait a few minutes for the system to boot, then use the browser's Back button to return to the TrueNAS® system's IP address and display the GUI login screen. 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.

15.5 Shutdown

Clicking Shutdown shows the warning message in Figure 15.4. The browser window background color changes to red to indicate that this is an option that will negatively impact users of the TrueNAS® system.



Fig. 15.4: Shutdown Warning Message

If a scrub or resilver is in progress when a shutdown 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 257) until it is verified that the scrub or resilver process is complete. Once complete, the shutdown request can be re-issued.

On High Availability (HA) systems with [Failover](#) (page 55), an additional checkbox is provided to shut down the standby node.

Click the Cancel button to cancel the shutdown request. Otherwise, click the Shutdown button to halt the system. Shutting down the system will disconnect all clients, including the web administration GUI, and will power off the TrueNAS® system.

15.6 Support Icon

The Support icon, the first icon on the right side of the menubar, provides a shortcut to `System` → `Support`. This screen can be used to verify the system license or to create a support ticket. Refer to [Support](#) (page 51) for detailed usage instructions.

15.7 Guide

The Guide icon, the second icon in the top menubar, provides links to both HTML (web) and PDF versions of the TrueNAS® User Guide (this documentation).

Both links are to local copies of the User Guide. No network connection is required.




Fig. 15.5: User Guide Menu

15.8 Alert

TrueNAS® provides an alert system to provide a visual warning of any conditions that require administrative attention. The Alert button in the far right corner flashes red when there is an outstanding alert. In the example alert shown in [Figure 15.6](#), the system is warning that the S.M.A.R.T. service is not running.




Fig. 15.6: Example Alert Message

Informational messages have a green OK, warning messages flash yellow, and messages requiring attention are listed as a red CRITICAL. CRITICAL messages are also emailed to the root user account. If you are aware of a critical condition but wish to remove the flashing alert until it can be dealt with, uncheck the box next to that message.

Behind the scenes, an alert daemon checks for various alert conditions, such as volume and disk status, and writes the current conditions to `/var/tmp/alert`. The daemon retrieves the current alert status every minute and changes the solid green alert icon to flashing red when a new alert is detected.

Current alerts can also be viewed from the Shell option of the Console Setup Menu ([Figure 2.12](#)) or from the Web Shell ([Figure 15.2](#)) by running `alertcli.py`. This can be useful when the alert originates from the standby node of a [High Availability \(HA\)](#) (page 55) system.

Some of the conditions that trigger an alert include:

- used space on a volume, dataset, or zvol goes over 80%; the alert goes red at 95%
- new OpenZFS feature flags are available for the pool; this alert can be unchecked if a pool upgrade is not desired at present
- a new update is available
- the system reboots itself
- non-optimal multipath states are detected
- ZFS pool status changes from HEALTHY
- a S.M.A.R.T. error occurs
- the system dataset does not reside on the boot pool
- the system is unable to bind to the WebGUI IPv4 Address set in `System → General`
- the system can not find an IP address configured on an iSCSI portal
- a periodic snapshot or replication task fails
- a VMware login or a [VMware-Snapshot](#) (page 131) task fails
- deleting a VMware snapshot fails
- a Certificate Authority or certificate is invalid or malformed
- an update failed, or the system needs to reboot to complete a successful update
- a re-key operation fails on an encrypted pool
- LDAP failed to bind to the domain
- any member interfaces of a lagg interface are not active
- the interface which is set as critical for failover is not found or is not configured
- HA is configured but the connection is not established
- one node of an HA pair gets stuck applying its configuration journal as this condition could block future configuration changes from being applied to the standby node
- 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 node fails, indicating the standby node is down
- [Proactive Support](#) (page 53) is enabled but any of the configuration fields are empty
- 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 passive node cannot be contacted
- when it is 180, 90, 30, or 14 days before support contract expiration

Note: If [Proactive Support](#) (page 53) 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.

ZFS PRIMER

ZFS is an advanced, modern filesystem that was specifically designed to provide features not available in traditional UNIX filesystems. It was originally developed at Sun with the intent to open source the filesystem so that it could be ported to other operating systems. After the Oracle acquisition of Sun, some of the original ZFS engineers founded [OpenZFS](http://open-zfs.org/wiki/Main_Page) (http://open-zfs.org/wiki/Main_Page) to provide continued, collaborative development of the open source version. To differentiate itself from Oracle ZFS version numbers, OpenZFS uses feature flags. Feature flags are used to tag features with unique names in order to provide portability between OpenZFS implementations running on different platforms, as long as all of the feature flags enabled on the ZFS pool are supported by both platforms. TrueNAS® uses OpenZFS and each new version of TrueNAS® keeps up-to-date with the latest feature flags and OpenZFS bug fixes.

Here is an overview of the features provided by ZFS:

ZFS is a transactional, Copy-On-Write (COW) (https://en.wikipedia.org/wiki/ZFS#Copy-on-write_transactional_model) filesystem. For each write request, a copy is made of the associated disk blocks and all changes are made to the copy rather than to the original blocks. When the write is complete, all block pointers are changed to point to the new copy. This means that ZFS always writes to free space, most writes are sequential, and old versions of files are not unlinked until a complete new version has been written successfully. ZFS has direct access to disks and bundles multiple read and write requests into transactions. Most filesystems cannot do this, as they only have access to disk blocks. A transaction either completes or fails, meaning there will never be a [write-hole](https://blogs.oracle.com/bonwick/entry/raid_z) (https://blogs.oracle.com/bonwick/entry/raid_z) and a filesystem checker utility is not necessary. Because of the transactional design, as additional storage capacity is added, it becomes immediately available for writes. To rebalance the data, one can copy it to re-write the existing data across all available disks. As a 128-bit filesystem, the maximum filesystem or file size is 16 exabytes.

ZFS was designed to be a self-healing filesystem. As ZFS writes data, it creates a checksum for each disk block it writes. As ZFS reads data, it validates the checksum for each disk block it reads. Media errors or “bit rot” can cause data to change, and the checksum no longer matches. When ZFS identifies a disk block checksum error on a pool that is mirrored or uses RAIDZ, it replaces the corrupted data with the correct data. Since some disk blocks are rarely read, regular scrubs should be scheduled so that ZFS can read all of the data blocks to validate their checksums and correct any corrupted blocks. While multiple disks are required in order to provide redundancy and data correction, ZFS will still provide data corruption detection to a system with one disk. TrueNAS® automatically schedules a monthly scrub for each ZFS pool and the results of the scrub are displayed by selecting the [Volume](#) (page 89) then clicking the Volume Status button. Checking scrub results can provide an early indication of potential disk problems.

Unlike traditional UNIX filesystems, it is not necessary to define partition sizes when filesystems are created. Instead, a group of disks, known as a vdev, are built into a ZFS pool. Filesystems are created from the pool as needed. As more capacity is needed, identical vdevs can be striped into the pool. In TrueNAS®, [Volume Manager](#) (page 89) can be used to create or extend ZFS pools. After a pool is created, it can be divided into dynamically-sized datasets or fixed-size zvols as needed. Datasets can be used to optimize storage for the type of data being stored as permissions and properties such as quotas and compression can be set on a per-dataset level. A zvol is essentially a raw, virtual block device which can be used for applications that need raw-device semantics such as iSCSI device extents.

ZFS supports real-time data compression. Compression happens when a block is written to disk, but only if the written data will benefit from compression. When a compressed block is accessed, it is automatically decompressed. Since compression happens at the block level, not the file level, it is transparent to any applications accessing the compressed data. By default, ZFS pools made using TrueNAS® version 9.2.1 or later will use the recommended LZ4 compression algorithm.

ZFS provides low-cost, instantaneous snapshots of the specified pool, dataset, or zvol. Due to COW, snapshots initially take no additional space. The size of a snapshot increases over time as changes to the files in the snapshot are written to disk.

Snapshots can be used to provide a copy of data at the point in time the snapshot was created. When a file is deleted, its disk blocks are added to the free list; however, the blocks for that file in any existing snapshots are not added to the free list until all referencing snapshots are removed. This makes snapshots a clever way to keep a history of files, useful for recovering an older copy of a file or a deleted file. For this reason, many administrators take snapshots often (e.g., every 15 minutes), store them for a period of time (e.g., for a month), and store them on another system. Such a strategy allows the administrator to roll the system back to a specific time. If there is a catastrophic loss, an off-site snapshot can restore the system up to the last snapshot interval (e.g., within 15 minutes of the data loss). Snapshots are stored locally but can also be replicated to a remote ZFS pool. During replication, ZFS does not do a byte-for-byte copy but instead converts a snapshot into a stream of data. This design means that the ZFS pool on the receiving end does not need to be identical and can use a different RAIDZ level, volume size, or compression settings.

ZFS boot environments provide a method for recovering from a failed upgrade. In 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 (https://blogs.oracle.com/realneel/entry/the_zfs_intent_log)). The ZIL is a storage area that temporarily holds *synchronous* writes until they are written to the ZFS pool (<https://pthree.org/2013/04/19/zfs-administration-appendix-a-visualizing-the-zfs-intent-log/>). Adding a fast (low-latency), power-protected SSD as a SLOG (Separate Log) device permits much higher performance. This is a necessity for NFS over ESXi, and highly recommended for database servers or other applications that depend on synchronous writes. More detail on SLOG benefits and usage is available in these blog and forum posts:

- [The ZFS ZIL and SLOG Demystified](http://www.freenas.org/blog/zfs-zil-and-slog-demystified/) (<http://www.freenas.org/blog/zfs-zil-and-slog-demystified/>)
- [Some insights into SLOG/ZIL with ZFS on FreeNAS®](https://forums.freenas.org/index.php?threads/some-insights-into-slog-zil-with-zfs-on-freenas.13633/) (<https://forums.freenas.org/index.php?threads/some-insights-into-slog-zil-with-zfs-on-freenas.13633/>)
- [ZFS Intent Log](http://nex7.blogspot.com/2013/04/zfs-intent-log.html) (<http://nex7.blogspot.com/2013/04/zfs-intent-log.html>)

Synchronous writes are relatively rare with SMB, AFP, and iSCSI, and adding a SLOG to improve performance of these protocols only makes sense in special cases. The `zilstat` utility can be run from `Shell` (page 257) to determine if the system will benefit from a SLOG. See [this website](http://www.richardelling.com/Home/scripts-and-programs-1/zilstat) (<http://www.richardelling.com/Home/scripts-and-programs-1/zilstat>) for usage information.

ZFS currently uses 16 GB of space for SLOG. Larger SSDs can be installed, but the extra space will not be used. SLOG devices cannot be shared between pools. Each pool requires a separate SLOG device. Bandwidth and throughput limitations require that a SLOG device must only be used for this single purpose. Do not attempt to add other caching functions on the same SSD, or performance will suffer.

In mission-critical systems, a mirrored SLOG device is highly recommended. Mirrored SLOG devices are required for ZFS pools at ZFS version 19 or earlier. ZFS pool version can be checked from the `Shell` (page 257) 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)` (<http://www.freebsd.org/cgi/man.cgi?query=top>) and includes the `arc_summary.py` and `arcstat.py` tools for monitoring the efficiency of the ARC. If an SSD is dedicated as a cache device, it is known as an L2ARC (<https://blogs.oracle.com/brendan/entry/test>). Additional read data is cached here, which can increase random read performance. L2ARC does not reduce the need for sufficient RAM. In fact, L2ARC needs RAM to function. If there is not enough RAM for an adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of random reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 32 GB of RAM, and the size of an L2ARC should not exceed ten times the amount of RAM. In some cases, it may be more efficient to have two separate pools: one on SSDs for active data, and another on hard drives for rarely used content. After adding an L2ARC device, monitor its effectiveness using tools such as `arcstat`. To increase the size of an existing L2ARC, stripe another cache device with it. The GUI will always stripe L2ARC, not mirror it, as the contents of L2ARC

are recreated at boot. Failure of an individual SSD from an L2ARC pool will not affect the integrity of the pool, but may have an impact on read performance, depending on the workload and the ratio of dataset size to cache size. Note that dedicated L2ARC devices cannot be shared between ZFS pools.

ZFS was designed to provide redundancy while addressing some of the inherent limitations of hardware RAID such as the write-hole and corrupt data written over time before the hardware controller provides an alert. ZFS provides three levels of redundancy, known as RAIDZ, where the number after the RAIDZ indicates how many disks per vdev can be lost without losing data. ZFS also supports mirrors, with no restrictions on the number of disks in the mirror. ZFS was designed for commodity disks so no RAID controller is needed. While ZFS can also be used with a RAID controller, it is recommended that the controller be put into JBOD mode so that ZFS has full control of the disks.

When determining the type of ZFS redundancy to use, consider whether the goal is to maximize disk space or performance:

- RAIDZ1 maximizes disk space and generally performs well when data is written and read in large chunks (128K or more).
- RAIDZ2 offers better data availability and significantly better mean time to data loss (MTTDL) than RAIDZ1.
- A mirror consumes more disk space but generally performs better with small random reads. For better performance, a mirror is strongly favored over any RAIDZ, particularly for large, uncacheable, random read loads.
- Using more than 12 disks per vdev is not recommended. The recommended number of disks per vdev is between 3 and 9. With more disks, use multiple vdevs.
- Some older ZFS documentation recommends that a certain number of disks is needed for each type of RAIDZ in order to achieve optimal performance. On systems using LZ4 compression, which is the default for TrueNAS® 9.2.1 and higher, this is no longer true.

These resources can also help determine the RAID configuration best suited to your storage needs:

- [Getting the Most out of ZFS Pools](https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/) (<https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/>)
- [A Closer Look at ZFS, Vdevs and Performance](http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance) (<http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance>)

Warning: RAID AND DISK REDUNDANCY ARE NOT A SUBSTITUTE FOR A RELIABLE BACKUP STRATEGY. BAD THINGS HAPPEN AND A GOOD BACKUP STRATEGY IS STILL REQUIRED TO PROTECT VALUABLE DATA. See [Periodic Snapshot Tasks](#) (page 113) and [Replication Tasks](#) (page 115) to use replicated ZFS snapshots as part of a backup strategy.

ZFS manages devices. When an individual drive in a mirror or RAIDZ fails and is replaced by the user, ZFS adds the replacement device to the vdev and copies redundant data to it in a process called resilvering. Hardware RAID controllers usually have no way of knowing which blocks were in use and must copy every block to the new device. ZFS only copies blocks that are in use, reducing the time it takes to rebuild the vdev. Resilvering is also interruptable. After an interruption, resilvering resumes where it left off rather than starting from the beginning.

While ZFS provides many benefits, there are some caveats:

- At 90% capacity, ZFS switches from performance- to space-based optimization, which has massive performance implications. For maximum write performance and to prevent problems with drive replacement, add more capacity before a pool reaches 80%. If you are using iSCSI, it is recommended to not let the pool go over 50% capacity to prevent fragmentation issues.
- When considering the number of disks to use per vdev, consider the size of the disks and the amount of time required for resilvering, which is the process of rebuilding the vdev. The larger the size of the vdev, the longer the resilvering time. When replacing a disk in a RAIDZ, it is possible that another disk will fail before the resilvering process completes. If the number of failed disks exceeds the number allowed per vdev for the type of RAIDZ, the data in the pool will be lost. For this reason, RAIDZ1 is not recommended for drives over 1 TB in size.
- It is recommended to use drives of equal sizes when creating a vdev. While ZFS can create a vdev using disks of differing sizes, its capacity will be limited by the size of the smallest disk.

For those new to ZFS, the [Wikipedia entry on ZFS](https://en.wikipedia.org/wiki/Zfs) (<https://en.wikipedia.org/wiki/Zfs>) provides an excellent starting point to learn more about its features. These resources are also useful for reference:

- [FreeBSD ZFS Tuning Guide](https://wiki.FreeBSD.org/ZFSTuningGuide) (<https://wiki.FreeBSD.org/ZFSTuningGuide>)
- [ZFS Administration Guide](http://docs.oracle.com/cd/E19253-01/819-5461/index.html) (<http://docs.oracle.com/cd/E19253-01/819-5461/index.html>)
- [Becoming a ZFS Ninja \(video\)](https://blogs.oracle.com/video/entry/becoming_a_zfs_ninja) (https://blogs.oracle.com/video/entry/becoming_a_zfs_ninja)
- [Slideshow explaining VDev, zpool, ZIL and L2ARC and other newbie mistakes!](https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-for-noobs.7775/) (<https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-for-noobs.7775/>)
- [A Crash Course on ZFS](http://www.bsdnow.tv/tutorials/zfs) (<http://www.bsdnow.tv/tutorials/zfs>)
- [ZFS: The Last Word in File Systems - Part 1 \(video\)](https://www.youtube.com/watch?v=uT2i2ryhCio) (<https://www.youtube.com/watch?v=uT2i2ryhCio>)
- [The Zettabyte Filesystem](https://www.youtube.com/watch?v=ptY6-K78McY) (<https://www.youtube.com/watch?v=ptY6-K78McY>)

HARDWARE SETUP

The TrueNAS® Z series (Z30 and Z35) consists of one or two main Unified Storage Array units. Optional E16, E24, or ES60 Expansion Shelves can be added to expand storage capacity.

The TrueNAS® X series (X10 and X20) consists of one or two 2U units. Optional ES12 Expansion Shelves can be added to expand storage capacity.

Specifications and data sheets for the Z series, X series, and expansion shelves are available at <https://www.ixsystems.com/truenas/>.

Racking and connection information is described below.

Note: Always perform the initial TrueNAS® setup in consultation with an iXsystems Support Representative.

17.1 Contacting iXsystems

For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@ixsystems.com
Telephone	Monday - Friday, 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)

17.2 Z Series Unified Storage Array











The TrueNAS® Unified Storage Array is an enterprise-grade 3U, 16-bay hybrid NAS control system that can be used alone, paired with a second Unified Storage Array to provide High Availability, or augmented with optional Expansion Shelves for additional storage.

Note: TrueNAS® units are carefully packed and shipped with trusted carriers to arrive in perfect condition. If there is any shipping damage or any parts are missing, please take photos and contact iXsystems support immediately at support@ixsystems.com or 1-855-GREP4-iX (1-855-473-7449) or 1-408-943-4100.

Please locate and record the hardware serial numbers on the back or side of each chassis for easy reference.

17.2.1 Check the Contents of the Box

The TrueNAS® Storage Array shipment comes with these components:

Quantity	Description	Image
1	TrueNAS® Unified Storage Array	
up to 16	Populated 3.5" drive trays	
1 pair	Outer rails, left and right	
8	#32 Thumbscrews	
2	Short screws	
2	Long screws	
2	Power cables	
1	Serial to 3.5mm cable	
1	Faceplate	
1	Printed setup guide	

Note: Network cables are highly configuration-dependent. Please contact your iXsystems Sales Representative for any questions about the included cables.

17.2.2 Become Familiar with the System




Fig. 17.1: Front View

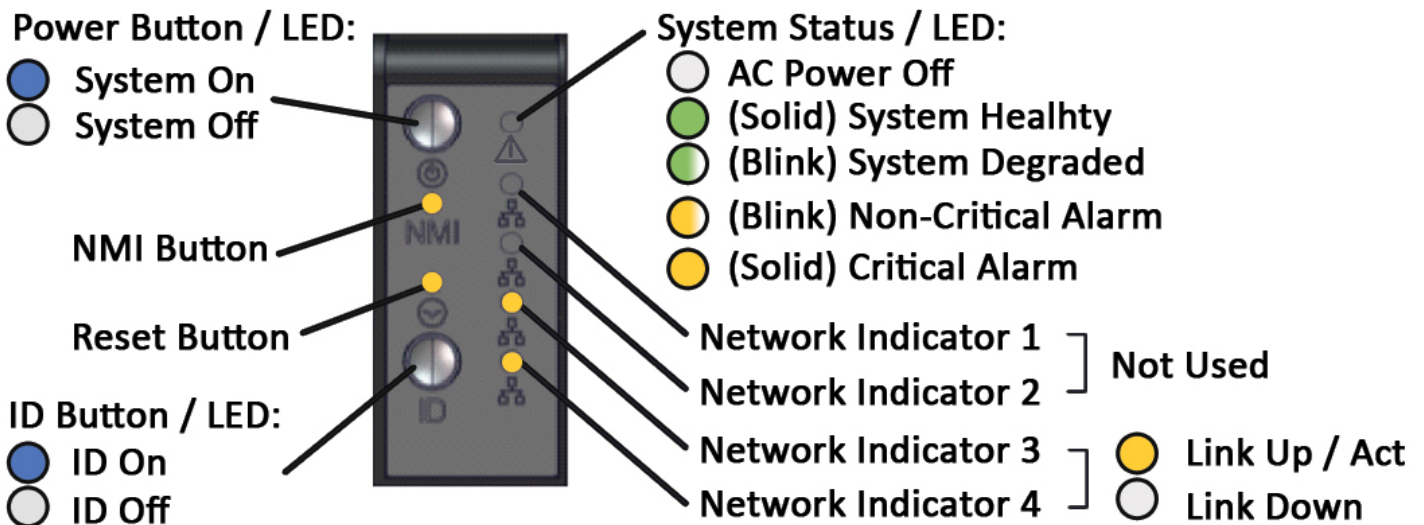


Fig. 17.2: Front Panel Buttons and Indicators




Fig. 17.3: Rear View




Fig. 17.4: Drive Tray

17.2.3 Mounting in an EIA Rack or Cabinet

EIA racks and cabinets have either round or square holes on the mounting frame.



Fig. 17.5: Racks with Square or Round Holes

Match the rail mounting bracket to the rack or cabinet holes. The outer rails shipped with a TrueNAS® Array are configured for a rack or cabinet with square holes by default.




Fig. 17.6: Outer Rail Configured for Square Holes (Default)




Fig. 17.7: Outer Rail Configured for Round Holes

Caution: The outer rails have a specific rear and front end and must be installed in the correct orientation. The front of the outer rail has a black plastic rail guide to assist guiding the chassis inner rail into it. If the outer rails are not attached to the rack in the correct orientation, the TrueNAS® Array cannot be mounted to them.

Press the button and rotate the bracket on the outer rail to match the type of holes in the rack or cabinet as shown:




Fig. 17.8: Rotate Rackmount Bracket

Make sure the rails are long enough for the rack or cabinet being used. Double-check the outer rail orientation, making sure the front of the outer rail is matched with the front of the rack or cabinet.




Fig. 17.9: Outer Rail Front




Fig. 17.10: Outer Rail Rear

For racks with square holes, snap the mounting brackets in the outer rail into the front and back of the rack frame. The brackets use a spring-loaded locking mechanism and do not require mounting screws. holes at either end of the rail into the mouting holes.





Fig. 17.11: Installing Rails in Racks with Square Holes

For racks with round holes, position the rails at the desired location in the rack and secure them with 8 #32 thumbscrews.




Fig. 17.12: Installing Rails in Racks with Round Holes

The TrueNAS® Storage Array can now be installed into the rack.

Caution: Two people are required to safely install or remove the TrueNAS® Storage Array in a rack or cabinet.

Carefully align the TrueNAS® Array inner rails with the guides in the outer rails attached to the rack. When aligned, slide the TrueNAS® Array into the rack until it locks. Press the lock releases, then push the unit the rest of the way in until the front is flush with the rack or cabinet. Finally, secure the system to the rack with 8 #32 screws. It might be necessary to adjust the position of the TrueNAS® Array to align the screw holes for securing the unit.




Fig. 17.13: Push Array into Rack and Release Pin-lock Latches

Install all of the populated drive trays into the front of the array. Refer to [Figure 17.14](#) for a detailed view.

Note: To avoid personal injury, do not install drives into the TrueNAS® Storage Array before racking.




Fig. 17.14: Drive Installation Instructions

Connect both network and storage cabling before turning on the TrueNAS® Storage Array for the first time.

Network cabling is highly dependent on the exact TrueNAS® model and environment. Please contact your iXsystems Support Representative if assistance is needed to connect the TrueNAS® Storage Array to the network.

The [Out-of-Band Management](https://support.ixsystems.com/truenasguide/truenas.html#out-of-band-management) (<https://support.ixsystems.com/truenasguide/truenas.html#out-of-band-management>) feature requires connection and configuration of the out-of-band management port before turning on the TrueNAS® Storage Array. Refer to [Figure 17.15](#) or the sticker on the storage controller handle for the location of the out-of-band management port.



Fig. 17.15: Back Panel Layout

Storage cabling instructions are shown in the [E16/E16F Expansion Shelf](https://support.ixsystems.com/truenasguide/tn_hardware.html#e16f-expansion-shelf) (https://support.ixsystems.com/truenasguide/tn_hardware.html#e16f-expansion-shelf) and [E24 Expansion Shelf](https://support.ixsystems.com/truenasguide/tn_hardware.html#e24-expansion-shelf) (https://support.ixsystems.com/truenasguide/tn_hardware.html#e24-expansion-shelf) setup instructions.

If the optional faceplate was included, attach it to the TrueNAS® Storage Array by inserting the two tabs on the right side of the faceplate into the holes in the right side handle section. Push the left side of the faceplate down until it clicks into place.

After all of the previous hardware setup steps are complete, plug the power cords into the AC receptacles on the back of the TrueNAS® Storage Array and secure them in place with the wire locks.

Note: Be sure to power on all TrueNAS® storage expansion shelves before powering on the TrueNAS® Storage Array.

Power on the TrueNAS® Storage Array by pressing the top left button on the control panel. Wait thirty seconds after turning on the first storage controller before powering on the second storage controller. This determines which controller is the active controller in High Availability configurations.

After the TrueNAS® Storage Array is fully operational, the TrueNAS® logo acts as a global fault light. By default, it is backlit in white. If there are any issues that need to be addressed, the light turns red. See the [Alert](https://support.ixsystems.com/truenasguide/tn_options.html#alert) (https://support.ixsystems.com/truenasguide/tn_options.html#alert) section of the TrueNAS® administrative graphical interface for more details about the error condition.

Copyright © 2018 iXsystems, Inc. All rights reserved.

All trademarks are the property of their respective owners.




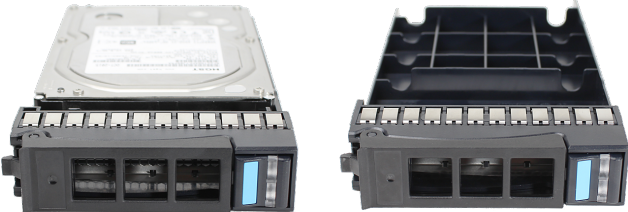


17.3 X Series Unified Storage Array

The TrueNAS® X Series Unified Storage Array is a 2U, 12-bay, hybrid unified data storage array.

Note: TrueNAS® units are carefully packed and shipped with trusted carriers to arrive in perfect condition. If there is any shipping damage or any parts are missing, please take photos and contact iXsystems support immediately at support@ixsystems.com or 1-855-GREP4-iX (1-855-473-7449) or 1-408-943-4100.


Please locate and record the hardware serial numbers on the back or side of each chassis for easy reference.

Carefully unpack the shipping boxes and locate these components:

	
X Series Unified Storage Array	X Series Bezel
	
Set of rackmount rails. The rails have a specific front end, identified by a label visible on the left above. The front ends of the rails must be installed facing the front of the rack.	A total of 12 populated or empty “air baffle” drive trays. Trays must be installed in all bays to maintain proper airflow for cooling. Up to ten drive trays are packed in a cardboard tray. Additional drive trays are packed with the accessory kit.
	
Accessory kit with 2 IEC C13 to NEMA 5-15P power cords, 2 IEC C14 to C14 cords, and velcro cable ties	Black USB to 3.5mm, 3.3V serial cable and rail extenders for racks over 30” deep

17.3.1 Become Familiar With the X Series System

The X Series has front panel indicators for power, locate ID, and fault. The fault indicator is on during the initial power-on self-test (POST) and turns off during normal operation. It turns on if the TrueNAS® software issues an **alert** (https://support.ixsystems.com/truenasguide/tn_options.html#alert).



The X Series contains one or two storage controllers in a side-by-side configuration. The connectors and features on each storage controller are:

1,2: Gigabit Ethernet connectors	7,8: HDmini SAS3 connectors 1 and 2
3: USB device (reserved, do not remove)	9: PCIe x8 expansion port
4: USB 2.0 connector	10: System console port (reserved)
5: Out-of-Band (OOB) serial port (3.5mm)	11: MAC address label
6: Out-of-Band Management Ethernet connector	12, 13: Redundant power supplies




Fig. 17.16: Back Panel


17.3.2 Rail Kit Assembly

On racks that are 30 inches deep or less, proceed to rail spring installation below.

Racks from 31 to 36 inches deep require installation of the included rail extenders. For these deeper racks, install cage nuts on the outside rear of the rack. The tabs on the cage nuts must be horizontal as shown. Using the included bolts, install the rail extender inside the rear of the rack. Repeat the process for the second extender, which is a mirror image of the first.




Rear of Rack
Viewed from Inside





Installed Extender
Viewed from Inside

Install a spring on the silver posts in the side of each rail.



Open the clamp latches on the ends of each rail. Place the rail in the rack with the front end toward the front of the rack, aligning the pins on both ends of the rail with the mounting holes in the rack. Swing the clamp latch closed to hold the rail in place. Use two of the supplied screws to secure the back end of the rail in place. Repeat the process for the second rail.



Caution: Two people are required to safely lift the chassis for rack installation or removal. Do not install drives until after the chassis has been installed in the rack, and remove all drives before removing the chassis from the rack.

Carefully place the chassis onto the rails mounted in the rack. Push the chassis in until the ears are flush with the front of the rack. Use two of the supplied screws to secure each ear to the rack.

17.3.3 Install Drive Trays

Drive trays are used to mount drives in the chassis. Each drive tray has a status LED which is blue when active or amber if a fault has occurred.

A tray must be placed in each drive bay to maintain proper airflow for cooling. If fewer than twelve drives are connected, empty “air baffle” trays must be placed in the empty bays.

A standard drive tray installation order simplifies support and is strongly recommended:


- SSD drives for SLOG, if present
- SSD drives for L2ARC, if present
- Hard drives or SSD drives for data storage
- Air baffle filler trays to fill any remaining empty bays

Install the first drive tray in the top left drive bay. Install the next drive tray to the right of the first. Install remaining drive trays to the right across the row. After a row is filled with drives, move down to the next row and start again with the left bay.

This example shows the proper order for a SLOG SSD, an L2ARC SSD, eight hard drives, and two empty air baffle trays.



To load an individual drive tray into a bay, press the blue button to open the latch. Carefully slide the tray into a drive bay until the left side of the latch touches the metal front edge of the chassis, then gently swing the latch closed until it clicks into place.



17.3.4 Connect Expansion Shelves

Refer to the installation instructions included with expansion shelves for details on connecting them.

17.3.5 Connect Network Cables


Note: Network cables vary by configuration and are not included. Please contact [iX Support](#) (page 267) with any questions.

Connect network cables to the Ethernet ports and Out-of-Band (OOB) management port before attempting to power on and configure the X series for the first time. The Out-of-Band (OOB) management port on the X series must be connected to a shielded Ethernet cable.

17.3.6 Connect Power Cords

If any TrueNAS® expansion shelves are connected to the X series, power them on first, then wait at least two minutes before connecting power cables to the X series.

Do not plug the power cords into a power outlet yet. Connect a power cord to the back of one power supply. Place the cord in the plastic clamp and press the tab into the latch to lock it in place. Repeat the process for the second power supply and cord.



After both power cords have been connected to the X series, they can be plugged into power outlets. The system is configured to automatically power on when connected to a power outlet. This design ensures that the X series comes back on when power is restored after a power failure.

17.3.7 Install Bezel (Optional)

The included bezel is not required for operation. If desired, install the bezel by aligning it with the pins on the bezel ears and pressing it into place.

17.3.8 Discover Out-of-Band Management IP Address

Several methods are available to determine the IP address currently assigned to the X series Out-of-Band management interface.

Preset

If the system was preconfigured by iXsystems, the Out-of-Band management interfaces have already been configured with the IP addresses requested by the user.

Otherwise, the Out-of-Band management IP addresses are set by default to static addresses:

Storage controller 1: 192.168.100.100, subnet mask 255.255.255.0

Storage controller 2 (if present): 192.168.100.101, subnet mask 255.255.255.0

DHCP

If the Out-of-Band management IP address has been configured to be assigned by DHCP, the IP address assigned by the DHCP server can be determined by checking the local DHCP server logs for the MAC addresses on the back panel of each X series storage controller, #11 on [Figure 17.16](#).

The local DHCP server can also be configured to provide a fixed IP address for the X series Out-of-Band management by using the MAC address.

Serial Cable

The Out-of-Band management IP address can be identified or changed by temporarily connecting the black USB serial cable to the Out-of-Band serial port, #5 on [Figure 17.16](#). Connect the USB end of the black cable to a laptop or desktop computer running serial terminal software.

Do not use the serial port for any purpose except checking the initial X series Out-of-Band management IP address or setting that address to be obtained by a different method. After use, disconnect the black USB serial cable from the X series.

Warning: The black USB serial cable is only for use with the Out-of-Band serial port on the X Series. Do not attempt to use it with any other equipment.

Out-of-Band Serial Terminal Communication Settings

Serial Port Device Names

The name of the serial port varies with operating systems. These are some typical examples: Windows: COM4, macOS: /dev/tty.usbserialxynnn, FreeBSD: /dev/cuaU0, Linux: /dev/ttyUSB0.

Serial Port Communication Parameters

Set the serial terminal program to use the appropriate port with these parameters: 38400 baud, 8 data bits, 1 stop bit, no parity, no flow control. Log in to the serial console with:

Username: sysadmin Password: superuser

The current Out-of-Band management IP address is displayed with:

```
ifconfig eth0 | grep 'inet addr'
    inet addr:10.20.1.227 Bcast:10.20.1.255 Mask:255.255.254.0
```

The current Out-of-Band network configuration settings are displayed with:

```
ipmitool -H 127.0.0.1 -U admin -P admin lan print
```

The Out-of-Band management system can be set to obtain an IP address from DHCP with:

```
ipmitool -H 127.0.0.1 -U admin -P admin lan set 1 ipsrc dhcp
```

The Out-of-Band management system can be set to use a static IP address and netmask. This example shows setting the IP address to 192.168.100.100 with a netmask of 255.255.255.0, and a default gateway of 192.168.100.1:

```
ipmitool -H 127.0.0.1 -U admin -P admin lan set 1 ipsrc static
ipmitool -H 127.0.0.1 -U admin -P admin lan set 1 ipaddr 192.168.100.10
ipmitool -H 127.0.0.1 -U admin -P admin lan set 1 netmask 255.255.255.0
ipmitool -H 127.0.0.1 -U admin -P admin lan set 1 defgw ipaddr 192.168.100.1
```

Log out of the Out-of-Band management system by typing `exit` and pressing `Enter`. After use, always disconnect the black USB serial cable from the X series system.

17.3.9 Connect to the X Series Console

With IPMI

Note: The IPMITool remote management utility must be installed on the laptop or desktop computer used to manage the X series remotely, and that computer must have access to the same network as the X series. FreeBSD, macOS, and Linux have package systems which can be used to install [IPMITool](https://sourceforge.net/projects/ipmitool/) (<https://sourceforge.net/projects/ipmitool/>). For Windows, a simple option is to install IPMITool through [Cygwin](https://www.cygwin.com/) (<https://www.cygwin.com/>).

Warning: Only use IPMITool for remote IPMI management on the X series. Other IPMI utilities may not work correctly or even damage the X series system.

When the Out-of-Band management IP address has been determined, the X series console is accessible through IPMI. In this example, 192.168.100.100 is the IP address assigned to the Out-of-Band management interface:

```
ipmitool -I lanplus -H 192.168.100.100 -U admin -a sol activate
```

Enter admin for the password, and the X series console is connected.

Tip: When a Serial Over LAN connection is already in use, SOL on another session is displayed when a laptop or desktop computer attempts to connect. The Serial Over LAN system can be reset from the remote laptop or desktop computer with:

```
ipmitool -H 192.168.100.100 -U admin bmc reset cold
```

Enter admin for the password, and the Serial Over LAN system is reset. Repeat the `sol activate` command above to connect.

The Serial Over LAN system can also be reset with the Out-of-Band serial port by attaching the black USB serial cable, connecting with a serial terminal program, and logging in as shown in [Serial Cable](#) (page 282). Then use

```
ifconfig eth0
```

to view the IP address of the eth0 network interface. Use the IP address, shown as eth0ipaddress in this example, in the reset command:

```
ipmitool -H eth0ipaddress -U admin bmc reset cold
```

Enter admin for the password, and the Serial Over LAN system is reset. Log out of the system with `exit` and disconnect the black USB serial cable from the X series system.

Tip: The Out-of-Band console password can be changed by attaching the black USB serial cable, connecting with a serial terminal program, and logging in as shown in [Serial Cable](#) (page 282). Then give this command to set the new password, shown as newpassword in this example:

```
ipmitool -H 127.0.0.1 -U admin -P admin user set password 2 newpassword
```

Log out of the system with `exit` and disconnect the black USB serial cable from the X series system.

Proceed to [Using the X Series Console](#) (page 286).

With the Serial Cable

The X series console can be directly connected to a serial terminal program by temporarily disconnecting the gray serial cable from the system console serial port, #10 on [Figure 17.16](#), and temporarily connecting the black USB serial cable to that port.

Connect the USB end of the black USB serial cable to a laptop or desktop computer running serial terminal software. See [Out-of-Band Serial Terminal Communication Settings](#) (page 282) for the serial device name. Set the terminal software to:

115200 baud, 8 data bits, 1 stop bit, no parity, no flow control

Wait two minutes after the X series has been connected to power, then press `Enter` to display the console menu. Find the message starting with `The web user interface is at:` and write down the IP address shown. After viewing the X series console, disconnect the black USB serial cable and reconnect the gray System Management cable to the system serial console port, #10 on [Figure 17.16](#).

17.3.10 Using the X Series Console

The X series console has two modes: SES (SCSI Enclosure Services) mode, and the standard x86 console mode.

If `ESM A =>` is displayed, the X series is in SES mode. Switch to the X86 console mode by typing these characters:

```
$%^0
```

Press `Enter` twice after typing the characters. The normal x86 console is displayed.

To switch back to the SES console, type these characters:

```
$%^2
```

17.3.11 Perform TrueNAS® Initial Software Configuration

The console displays the IP address of the TrueNAS® X series graphical web interface, 192.168.100.231 in this example:

```
The web user interface is at:
```

```
http://192.168.100.231
```

Enter the IP address into a browser on a computer on the same network to access the web user interface.




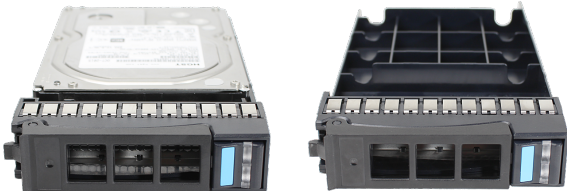


17.4 ES12 Expansion Shelf

The TrueNAS® ES12 is 2U, 12-bay, SAS3 (12 Gb/s) expansion shelf with dual expansion modules and redundant power supplies.

Note: TrueNAS® units are carefully packed and shipped with trusted carriers to arrive in perfect condition. If there is any shipping damage or any parts are missing, please take photos and contact iXsystems support immediately at support@ixsystems.com or 1-855-GREP4-iX (1-855-473-7449) or 1-408-943-4100.


Please locate and record the hardware serial numbers on the back or side of each chassis for easy reference.

Carefully unpack the shipping boxes and locate these components:

 <p>ES12 Expansion Shelf</p>	 <p>ES12 Bezel</p>
 <p>Set of rackmount rails. The rails have a specific front end, identified by a label visible on the left above. The front ends of the rails must be installed facing the front of the rack.</p>	 <p>A total of 12 populated or empty “air baffle” drive trays. Trays must be installed in all bays to maintain proper airflow for cooling. Up to ten drive trays are packed in a cardboard tray. Additional drive trays are packed with the accessory kit.</p>
 <p>Accessory kit with 2 IEC C13 to NEMA 5-15P power cords, 2 IEC C14 to C14 cords, and a set of velcro cable ties</p>	 <p>Two 3-meter Mini SAS HD to Mini SAS HD cables</p>


17.4.1 Become Familiar With the ES12 System

The ES12 has front panel indicators for power, locate ID, and fault. The fault indicator is on during the initial power-on self-test (POST) and turns off during normal operation. It turns on if the TrueNAS® software issues an **alert** (https://support.ixsystems.com/truenasguide/tn_options.html#alert).



The ES12 contains two expansion modules in a side-by-side configuration. The connectors and features on each module are:

1-3: HD Mini SAS3 connectors (T1-T3)
4: Debug port (TrueNAS® internal use only)
5: Redundant power supplies




17.4.2 Rail Kit Assembly

On racks that are 30 inches deep or less, proceed to rail spring installation below.

Racks from 31 to 36 inches deep require installation of the included rail extenders. For these deeper racks, install cage nuts on the outside rear of the rack. The tabs on the cage nuts must be horizontal as shown. Using the included bolts, install the rail extender inside the rear of the rack. Repeat the process for the second extender, which is a mirror image of the first.





Rear of Rack
Viewed from Inside





Installed Extender
Viewed from Inside

Install a spring on the silver posts in the side of each rail.



Open the clamp latches on the ends of each rail. Place the rail in the rack with the front end toward the front of the rack, aligning the pins on both ends of the rail with the mounting holes in the rack. Swing the clamp latch closed to hold the rail in place. Use two of the supplied screws to secure the back end of the rail in place. Repeat the process for the second rail.



Caution: Two people are required to safely lift the chassis for rack installation or removal. Do not install drives until after the chassis has been installed in the rack, and remove all drives before removing the chassis from the rack.

Carefully place the chassis onto the rails mounted in the rack. Push the chassis in until the ears are flush with the front of the rack. Use two of the supplied screws to secure each ear to the rack.

17.4.3 Install Drive Trays

Drive trays are used to mount drives in the chassis. Each drive tray has a status LED which is blue when active or amber if a fault has occurred.

A tray must be placed in each drive bay to maintain proper airflow for cooling. If fewer than twelve drives are connected, empty “air baffle” trays must be placed in the empty bays.

A standard drive tray installation order simplifies support and is strongly recommended:


- SSD drives for SLOG, if present
- SSD drives for L2ARC, if present
- Hard drives or SSD drives for data storage
- Air baffle filler trays to fill any remaining empty bays

Install the first drive tray in the top left drive bay. Install the next drive tray to the right of the first. Install remaining drive trays to the right across the row. After a row is filled with drives, move down to the next row and start again with the left bay.

This example shows the proper order for a SLOG SSD, an L2ARC SSD, eight hard drives, and two empty air baffle trays.




To load an individual drive tray into a bay, press the blue button to open the latch. Carefully slide the tray into a drive bay until the left side of the latch touches the metal front edge of the chassis, then gently swing the latch closed until it clicks into place.



17.4.4 Connect Power Cords

Do not plug the power cords into a power outlet yet. Connect a power cord to the back of one power supply. Place the cord in the plastic clamp and press the tab into the latch to lock it in place. Repeat the process for the second power supply and cord.




17.4.5 Connect SAS Cables

Shut down and power off the TrueNAS® X series. Remove the X series power cables from the power outlets.


Plug the ES12 power cords into power outlets. Wait two minutes for the drives to start.

Connect an SAS cable from the ES12 expander module 1 T1 connector to the X series node 1 SAS 1 connector (cable #1 in photo).

If the X series has a second node for HA (High Availability), connect a second SAS cable from the ES12 expander module 2 T1 connector to the X series node 2 SAS 1 connector (cable #2 in photo).




A second ES12 can also be connected to the X series. Connect the first ES12 as shown above. Then connect the second ES12, running SAS cables from the ES12 T1 connectors to the SAS 2 connectors on the X series.



17.4.6 Use with the ES60 Expansion Shelf (X20 Only)

The X20 supports using an ES60 Expansion Shelf in combination with the ES12. Connect the ES60 to the X20 SAS 1 ports, and connect the ES12 to the X20 SAS 2 ports.



After all SAS cables have been connected, plug the X series power cords into power outlets.

17.4.7 Install Bezel (Optional)

The included bezel is not required for operation. If desired, install the bezel by aligning it with the pins on the bezel ears and pressing it into place.

Copyright © 2018 iXsystems, Inc. All rights reserved.

All trademarks are the property of their respective owners.







17.5 ES24 Expansion Shelf

The TrueNAS® ES24 is a 4U, 24-bay, SAS3 (12 Gb/s) expansion shelf with dual expansion modules and redundant power supplies.

Note: TrueNAS® units are carefully packed and shipped with trusted carriers to arrive in perfect condition. If there is any shipping damage or any parts are missing, please take photos and contact iXsystems support immediately at support@ixsystems.com or 1-855-GREP4-iX (1-855-473-7449) or 1-408-943-4100.


Please locate and record the hardware serial numbers on the back or side of each chassis for easy reference.

Carefully unpack the shipping boxes and locate these components:

 <p>ES24 Expansion Shelf</p>	 <p>ES24 Bezel</p>
 <p>Set of rackmount rails</p>	 <p>Up to 24 drive trays populated with drives</p>
 <p>Two 3-meter Mini SAS HD to Mini SAS HD cables and one serial cable</p>	 <p>Two 2 IEC C13 to NEMA 5-15P power cords</p>

17.5.1 Become Familiar With the ES24

The ES24 has front panel buttons for power and alarm mute, and indicators for power, locate ID, and fault. The fault indicator is on during the initial power-on self-test (POST) or when the TrueNAS® software has issued an [alert](https://support.ixsystems.com/truenasguide/tn_options.html#alert) (https://support.ixsystems.com/truenasguide/tn_options.html#alert).






Fig. 17.17: Back Panel


The ES24 contains two expansion modules in an over/under configuration. The connectors and features on each module are:

1-2: Redundant power supplies
3-4: Fans
5-7: HD Mini SAS3 connectors
8: Serial port

17.5.2 Rail Kit Assembly


Remove Cabinet Rails from Rack Rails

Extend the cabinet rail until it stops. To remove the cabinet rail, press the white release tab to the right while pulling the cabinet rail.



Mount Cabinet Rails

The cabinet rails are mounted on both sides of the system. Align the cabinet rail keyholes with the three posts on the side of the chassis and pull the rail until the post is wedged in the smaller hole.




Align the rail holes with the screw holes and secure the rail with one of the supplied cabinet screws included with the rail kit. Repeat this process on the other side.

Mount the Rack Rails

Place the rail in the rack with the front end toward the front of the rack, aligning the pins with the mounting holes in the front of the rack. Push the pins into the holes until the latch clicks.

Place one square cage nut with the tabs vertical into the front of the rack, directly below the top rail pin. Repeat this process on the other front rail. The cage nuts will later support the screws used to secure the enclosure to the rack.


For the rear end of the rail, align the pins with the mounting holes on the rear rack. Pull the white latch toward the rear until the pins click in place. Repeat this process for the second rear rail.



Mount the Unit in the Rack

Caution: Two people are required to safely lift the chassis for rack installation or removal. Do not install drives until after the chassis has been installed in the rack, and remove all drives before removing the chassis from the rack.

Pull the front rack rail forward until it stops. Align the cabinet rail with the inside of the front rack rail and slide the cabinet rail forward until it is fully seated inside the rack rail. Repeat the process for the second rail.



When both cabinet rails are secured inside of the rack rails, gently push the chassis until it stops halfway in. Slide the blue release tabs on both cabinet rails toward the front of the system while pushing the unit in until it is flush with the front of the rack.



Anchor the unit in the rack on both sides with the screws included in the rail kit.

17.5.3 Install Drive Trays


Drive trays are used to mount drives in the array. Each drive tray has a status LED which is blue when active or amber if a fault has occurred.

Press the silver button to open the latch. Carefully slide the tray into a drive bay until the right side of the latch touches the metal front edge of the chassis, then gently swing the latch closed until it clicks into place.



17.5.4 Connect Power Cords

Do not plug the power cords into a power outlet yet. Connect a power cord to the back of one power supply. Place the cord into the plastic clamp and press the tab into the latch to lock it in place. Repeat the process for the second power supply and cord.



17.5.5 Connect SAS Cables

Plug the ES24 power cords into power outlets. Wait two minutes for the drives to start.

Note: The TrueNAS® Unified Array can remain on while the expansion shelf is connected.

Connect an SAS cable from ES24 expansion module 1 SAS 0 connector to the farthest left Z series node 1 SAS connector (cable #1 in photo).

If the Z series has a second node for HA (High Availability), connect a second SAS cable from ES24 expansion module 2 SAS 0 connector to the farthest left Z series node 2 SAS connector (cable #2 in photo).




Fig. 17.18: SAS Cable Connection for One ES24

A second ES24 can also be connected to the Z series. Connect the first ES24 as shown above. Then connect the second ES24 SAS 0 ports to the next unused SAS ports on the Z series as shown:




Fig. 17.19: SAS Cable Connection for Two ES24s

After all SAS cables have been connected, plug the Z series power cords into power outlets.

17.5.6 Install Bezel (Optional)

The included bezel is not required for operation. If desired, install the bezel by aligning it with the pins on the bezel ears and pressing it into place.


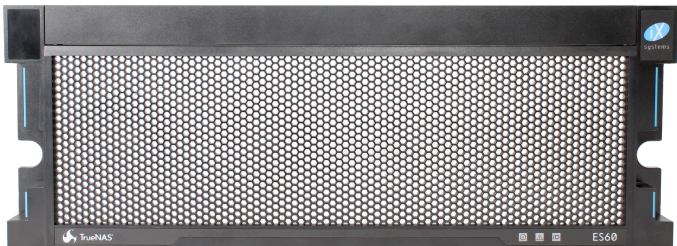




17.6 ES60 Expansion Shelf

The TrueNAS® ES60 Expansion Shelf is a 4U, 60-bay storage expansion unit designed specifically for use with the TrueNAS® Unified Storage Array.

Note: TrueNAS® units are carefully packed and shipped with trusted carriers to arrive in perfect condition. If there is any shipping damage or any parts are missing, please take photos and contact iXsystems support immediately at support@ixsystems.com or 1-855-GREP4-iX (1-855-473-7449) or 1-408-943-4100.

Please locate and record the hardware serial numbers on the back or side of each chassis for easy reference.

Carefully unpack the shipping boxes and locate these components:

 <p>ES60 Expansion Shelf</p>	 <p>ES60 Bezel</p>
 <p>Rail kit with mounting hardware</p>	 <p>Up to 60 drive trays with installed hard drives (shipped separately from the ES60)</p>
 <p>Two IEC C13 to NEMA 5-15P power cords, two IEC C14 to C14 power cords</p>	 <p>Two 3-meter Mini SAS HD to Mini SAS HD cables</p>

17.6.1 Become Familiar With the ES60

The ES60 has indicators on the front panel for power, fault, and locate ID. The fault indicator is on during the initial power-on self-test (POST) or when the TrueNAS® software has issued an **alert** (https://support.ixsystems.com/truenasguide/tn_options.html#alert).

Front view of the ES60 Expansion Shelf:



Front panel indicators:



Rear view of the ES60 Expansion Shelf:




The ES60 contains two expansion modules in a side-by-side configuration. The connectors and features on each module are:

1: Power supply	4: Locate ID
2: Power indicator	5: Management port (not used)
3: Alarm indicator	6,7: HD Mini SAS3 connectors

17.6.2 Rail Kit Assembly

Separate Cabinet Rails from Rack Rails

Each rack rail includes an inner cabinet rail that must be removed. Extend the cabinet rail until the white release tab is exposed.



Press the white release tab to the right while pulling the cabinet rail to remove it. Repeat this process for the second rail.


Mount Cabinet Rails

The cabinet rails are mounted on each side of the system. Align the cabinet rail keyholes with the posts on the side of the chassis. Slide the rail toward the rear of the system until the metal tab clicks and secures the rail in place. Repeat this process on the other side.



Mount Rack Rails

Install four cage nuts in the rack, one for each bracket. The lock screw for the bracket is already installed in the cage nut. Place the rail in the rack with the front end toward the front of the rack and rear toward the back of the rack, aligning the pins on both ends of the rail with the mounting holes in the rack. Push the pins into the rack holes until they lock in place.



17.6.3 Mount Unit in the Rack

Caution: Two people are required to safely lift the chassis for rack installation or removal. Do not install drives until after the chassis has been installed in the rack, and remove all drives before removing the chassis from the rack.

Slide the front rack rail forward to meet the rear of the cabinet rails. Align the cabinet rail with the inside of the front rack rail and slide the cabinet rail forward until it is fully seated inside the rack rail. Repeat the process for the second rail.



When both cabinet rails are secured inside the rack rails, carefully slide the chassis until the ears are flush with the front of the rack. Turn both blue screws on the ears of the chassis to the right to secure the unit in the rack.



17.6.4 Drive Tray Installation

Note: Do not install the drives until the chassis has been installed in the rack.

Remove Top Cover

Slide the unit out on the rails. Turn the blue screws counterclockwise to unlock the top cover. Slide the top cover toward the front of the system, then lift to remove it.



Install Drive Trays

Drive trays are used to mount drives in the array.

A standard drive tray installation order simplifies support and is strongly recommended:

- SSD drives for SLOG, if present
- SSD drives for L2ARC, if present
- Hard drives or SSD drives for data storage

Install the first drive tray in the front left drive bay. Install the next drive tray to the right of the first. Install remaining drive trays to the right across the row. After a row is filled with drives, move back to the next row and start again with the left bay. A label on the front left of the lid identifies the preferred order of drive connection.

Press the blue button to open the latch. Lower the drive tray into a drive bay until the latch begins to move into place. Push the latch the rest of the way until it locks into place.



17.6.5 ES60 Cable Management Arm

The included cable management arm (CMA) is not required for operation. If desired, the CMA can be used to help organize the ES60 power and data cables.




The tabs along the side of the flex housing can be unclipped from the top, the bottom, or removed entirely.



Install the Cable Management Arm

Locate the two posts on the left rear side of the E560. Align the holes on the CMA chassis bracket with the posts on the chassis. Slide the cable management arm forward and pull the lever on the latch upward to lock the bracket into place.




Locate the end of the flex housing with exposed pins. Unclip and open the two tabs closest to the end, allowing the flex housing to compress enough to fit into the bracket holes. Press the flex housing firmly into the bracket until the pins seat in the holes.




Remove the two screws already attached to the side of the CMA rail bracket. Align the screw holes with the holes in the rear of the left cabinet rail and attach the bracket to the rail with the screws.



Locate the end of the flex housing with exposed holes. Unclip and open the two tabs closest to the end, allowing the flex housing to expand enough to fit over the bracket pins. Press the flex housing firmly into the bracket until the holes seat on the pins.




Completed Cable Management Arm assembly:



Power and data cables are routed through the flex housing. The tabs can be opened or removed to allow access or space for cable ends. Remember to leave some slack in the cables at both ends to allow for movement of the arm and chassis.

17.6.6 Connect Power Cords

Do not plug the power cords into a power outlet yet. Connect a power cord to the back of one power supply, pressing it into the plastic clamp and pressing on the tab to lock it in place. Repeat the process for the second power supply and cord. This turns on the ES60.



Note: Service and management ports are not used during normal operation and should not be connected.

When setting up an ES60 for the first time, wait two minutes after turning on the ES60 before powering on the TrueNAS® Unified Array. If the TrueNAS® Unified Array is already in operation, the Expansion Shelf can be powered on at any time.


17.6.7 Connect SAS Cables

Plug the ES60 power cords into power outlets. Wait two minutes for the drives to start.

Note: The TrueNAS® Unified Array can remain on while the expansion shelf is connected.


Connecting the ES60 to a TrueNAS® X20

Connect an SAS cable from the ES60 expander module 1 connector 1 to the X20 1 SAS 1 connector (cable #1 in photo). If the X20 has a second node for HA (High Availability), connect a second SAS cable from the ES60 expander module 2 connector 1 to the X20 node 2 SAS 1 connector (cable #2 in photo).




Connecting the ES60 to a TrueNAS® Z35

Connect an SAS cable from the ES60 expander module 1 connector 1 to the Z35 Controller 1 SAS 1 connector (cable #1 in photo). If the Z35 has a second node for HA (High Availability), connect a second SAS cable from the ES60 expander module 2 connector 1 to the Z35 Controller 2 SAS 1 connector (cable #2 in photo).



A second ES60 can also be connected to the Z35. Connect the first ES60 as shown above. Then connect the second ES60 1 ports to the next unused SAS ports on the Z35 as shown:



Up to eight ES60 expansion shelves are supported on the TrueNAS® Z35.

17.6.8 Install Bezel (Optional)

The included bezel is not required for operation.

Line up the screw holes on the back of the bezel with the screw holes on the ears of the ES60. Install one upper screw from the back side of the left ES60 ear, then install a lower screw from the back of the right ES60 ear. Install the remaining two screws following the same diagonal pattern.

VAAI

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.

18.1 VAAI for iSCSI

VAAI for iSCSI supports these operations:

- Atomic Test and Set (ATS) allows multiple initiators to synchronize LUN access in a fine-grained manner rather than locking the whole LUN and preventing other hosts from accessing the same LUN simultaneously.
- Clone Blocks (XCOPY) copies disk blocks on the NAS. Copies occur locally rather than over the network. The operation is similar to [Microsoft ODX](https://technet.microsoft.com/en-us/library/hh831628) (<https://technet.microsoft.com/en-us/library/hh831628>).
- LUN Reporting allows a hypervisor to query the NAS to determine whether a LUN is using thin provisioning.
- Stun pauses running virtual machines when a volume runs out of space. The space issue can then be fixed and the virtual machines can continue rather than reporting write errors.
- Threshold Warning the system reports a warning when a configurable capacity is reached. In TrueNAS®, this threshold can be configured at the pool level when using zvols (see [Table 9.6](#)) or at the extent level (see [Table 9.11](#)) for both file- and device-based extents. Typically, the warning is set at the pool level, unless file extents are used, in which case it must be set at the extent level.
- Unmap informs 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.

18.2 VAAI for NAS

VAAI for NAS (<https://code.vmware.com/programs/vaai-nas>) is automatically enabled on TrueNAS® when the [NFS](#) (page 201) service is running. These operations are supported:

- Extended Statistics provides extended statistics on NFS shares.
- Full File Clone efficiently clones a file on the NAS without copying the data over the network.

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.html) (<https://tools.ietf.org/html/rfc2616.html>), such as GET, PUT, POST, or DELETE.

As shown in [Figure 19.1](#), an online version of the API is available at api.freenas.org (<http://api.freenas.org>).



Fig. 19.1: API Documentation

The rest of this section shows code examples to illustrate the use of the API.

Note: Beginning with TrueNAS® 9.10.2, a new API has been added. The old API is still present for compatibility. Documentation for the new API is available on the 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.

19.1 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 volume or dataset for the user's Home Directory. After creating the user, start the SSH service using **Services → Control Services**. That user will now be able to `ssh` to the IP address of the 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. The text in black should not be changed. After saving changes, run the script by typing `python scriptname.py`. If all goes well, the new user account will appear in **Account → Users → View Users** in the TrueNAS® GUI.

Here is the example script with an explanation of the line numbers below it.

```
1 import json
2 import requests
3 r = requests.post(
4     'https://freenas.mydomain/api/v1.0/account/users/',
5     auth=('root', 'freenas'),
6     headers={'Content-Type': 'application/json'},
7     verify=False,
8     data=json.dumps({
9         'bsdusr_uid': '1100',
10        'bsdusr_username': 'myuser',
11        'bsdusr_mode': '755',
12        'bsdusr_creategroup': 'True',
13        'bsdusr_password': '12345',
14        'bsdusr_shell': '/usr/local/bin/bash',
15        'bsdusr_full_name': 'Full Name',
16        'bsdusr_email': 'name@provider.com',
17    })
18 )
19 print r.text
```

Where:

Lines 1-2: import the Python modules used to make HTTP requests and handle data in JSON format.

Line 4: replace `freenas.mydomain` with the Hostname value in **System → System Information**. Note that the script will fail if the machine running it is not able to resolve that hostname. Change `https` to `http` to use HTTP rather than HTTPS to access the TrueNAS® system.

Line 5: replace `freenas` with the password used to access the TrueNAS® system.

Line 7: if you are using HTTPS and want to force validation of the SSL certificate, change `False` to `True`.

Lines 8-16: set the values for the user being created. The [Users resource](http://api.freenas.org/resources/account.html#users) (<http://api.freenas.org/resources/account.html#users>) describes this in more detail. Allowed parameters are listed in the JSON Parameters section of that resource. Since this resource creates a FreeBSD user, the values entered must be valid for a FreeBSD user account. [Table 19.1](#) summarizes acceptable values. This resource uses JSON, so the boolean values are `True` or `False`.

Table 19.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
<code>bsdusr_username</code>	string	maximum 32 characters, though a maximum of 8 is recommended for interoperability; can include numerals but cannot include a space
Continued on next page		

Table 19.1 – continued from previous page

JSON Parameter	Type	Description
bsdusr_full_name	string	may contain spaces and uppercase characters
bsdusr_password	string	can include a mix of upper and lowercase letters, characters, and numbers
bsdusr_uid	integer	by convention, user accounts have an ID greater than 1000 with a maximum allowable value of 65,535
bsdusr_group	integer	if bsdusr_creategroup is set to False, specify the numeric ID of the group to create
bsdusr_creategroup	boolean	if set to True, a primary group with the same numeric ID as bsdusr_uid will be created automatically
bsdusr_mode	string	sets default numeric UNIX permissions of user's home directory
bsdusr_shell	string	specify full path to a UNIX shell that is installed on the system
bsdusr_password_disabled	boolean	if set to True, user is not allowed to log in
bsdusr_locked	boolean	if set to True, user is not allowed to log in
bsdusr_sudo	boolean	if set to True, <code>sudo</code> is enabled for the user
bsdusr_sshpubkey	string	contents of SSH authorized keys file

Note: When using boolean values, JSON returns raw lowercase values but Python uses uppercase values. So use True or False in Python scripts even though the example JSON responses in the API documentation are displayed as true or false.

19.2 A More Complex Example

This section provides a walk-through of a more complex example found in the `startup.py` script. Use the searchbar within the API documentation to quickly locate the JSON parameters used here. This example defines a class and several methods to create a ZFS volume, create a ZFS dataset, share the dataset over CIFS, and enable the CIFS service. Responses from some methods are used as parameters in other methods. In addition to the import lines seen in the previous example, two additional Python modules are imported to provide parsing functions for command line arguments:

```
import argparse
import sys
```

It then creates a Startup class which is started with the hostname, username, and password provided by the user via the command line:

```
1 class Startup(object):
2     def __init__(self, hostname, user, secret):
3         self._hostname = hostname
4         self._user = user
5         self._secret = secret
6         self._ep = 'http://%s/api/v1.0' % hostname
7     def request(self, resource, method='GET', data=None):
8         if data is None:
9             data = ''
10        r = requests.request(
11            method,
12            '%s/%s/' % (self._ep, resource),
13            data=json.dumps(data),
14            headers={'Content-Type': 'application/json'},
15            auth=(self._user, self._secret),
16        )
17    if r.ok:
18        try:
19            return r.json()
20        except:
```



```

21         return r.text
22     raise ValueError(r)

```

A `get_disks` method is defined to get all the disks in the system as a `disk_name` response. The `create_pool` method uses this information to create a ZFS pool named `tank` which is created as a `stripe`. The `volume_name` and `layout` JSON parameters are described in the “Storage Volume” resource of the API documentation.

```

1  def _get_disks(self):
2      disks = self.request('storage/disk')
3      return [disk['disk_name'] for disk in disks]
4
5  def create_pool(self):
6      disks = self._get_disks()
7      self.request('storage/volume', method='POST', data={
8          'volume_name': 'tank',
9          'layout': [
10             {'vdevtype': 'stripe', 'disks': disks},
11         ],
12     })

```

The `create_dataset` method is defined which creates a dataset named `MyShare`:

```

1  def create_dataset(self):
2      self.request('storage/volume/tank/datasets', method='POST', data={
3          'name': 'MyShare',
4      })

```

The `create_cifs_share` method is used to share `/mnt/tank/MyShare` with guest-only access enabled. The `cifs_name`, `cifs_path`, `cifs_guestonly` JSON parameters, as well as the other allowable parameters, are described in the “Sharing CIFS” resource of the API documentation.

```

1  def create_cifs_share(self):
2      self.request('sharing/cifs', method='POST', data={
3          'cifs_name': 'My Test Share',
4          'cifs_path': '/mnt/tank/MyShare',
5          'cifs_guestonly': True
6      })

```

Finally, the `service_start` method enables the CIFS service. The `srv_enable` JSON parameter is described in the Services resource.

```

1  def service_start(self, name):
2      self.request('services/services/%s' % name, method='PUT', data={
3          'srv_enable': True,
4      })
5

```

APPENDIX A: END-USER LICENSE AGREEMENT

TrueNAS® EULA:

BY PURCHASING, DOWNLOADING, INSTALLING, OR OTHERWISE USING THE SOFTWARE, YOU AGREE TO BE BOUND BY THE TERMS OF THIS END-USER LICENSE AGREEMENT (EULA). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MAY NOT INSTALL OR USE THE SOFTWARE.

1. DEFINITIONS

“Company” means iXsystems, Inc.

“Product” means iXsystems Storage Array software (TrueNAS®).

“EULA” means this End User License Agreement

“You” means the natural person or the entity that is agreeing to be bound by this EULA, their employees and third party contractors that provide services to you.

“Open Source Software” means various open source software components licensed under the terms of applicable open source license agreements included in the materials relating to such software. Open Source Software is composed of individual software components, each of which has its own copyright and its own applicable license conditions.

“FreeNAS®” means a complete open source operating system available at <http://www.iXsystems.org>

“Site” means iXsystems, Inc. website: <http://www.iXsystems.com>

2. TERMS AND CONDITIONS

2.1. Company grants You a non-exclusive, non-sublicensable, non-transferable license to use the Product on a single computer, subject to the terms and conditions of this EULA and in accordance with the instructions, specifications and documentation provided with the Product (collectively, the “Documentation”). This license of Product may not be shared or used concurrently on different computers.

2.2. Product Warranty Disclaimer. THE PRODUCT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE. Company BEARS NO LIABILITY FOR ANY DAMAGES RESULTING FROM USE (OR ATTEMPTED USE) OF THE PRODUCT.

2.3. You agree that You will NOT without the express written authorization of Company: (a) copy, sell, sublicense, or otherwise transfer the Product to any third party; (b) remove any titles, trademarks or trade names, copyright notices, legends, or other proprietary markings on the software in the Product; (c) except to the extent expressly permitted by applicable law, and to the extent that the Company is not permitted by that applicable law to exclude or limit the following rights, You will not decompile, disassemble, reverse engineer, or otherwise attempt to derive source code from the Product, in whole or in part.

2.4. FreeNAS® software. The Product contains part of FreeNAS® software, which in turn contains a variety of Open Source Software components. You can redistribute and/or modify the Open Source Software under the terms and conditions of the corresponding open source licenses. You may obtain a copy of the source code corresponding to the binaries for the Open Source Software from the FreeNAS® home page at <http://www.FreeNAS.org>. You agree to comply with the applicable licenses and additional terms and notices of such Open Source Software components. Company makes no warranties or representations of any kind to You regarding Open Source Software components, or that the corresponding open source licenses may not change or be altered at any time.

2.5. Third party software. The Product may contain Third Party software that must be separately licensed. Any separately licensed software is licensed exclusively by that license and the terms of this License Agreement do not apply.

2.6. Software Modifications. Modifications of the Product software will not be supported by the Company unless indicated otherwise by express written authorization. Company will not be liable for any modifications to the Product software or any errors or damages resulting from such modifications.

2.7. Company may update or discontinue the Product or revise the Documentation at any time without prior notice to You, and the Product and/or the Documentation may become unavailable to You even after an order is placed. All prices mentioned on the Company Site are subject to change without notice.

2.8. Product Descriptions; Pricing; Errors. Company attempts to be as accurate as possible and eliminate errors in the Product and on the Site. However, Company does not warrant that the Product, its descriptions, photographs, pricing or other content of the Site is accurate, complete, reliable, stable, defect free, current, or error-free. In the event of an error, whether on the Site or otherwise, Company reserves the right to correct such error at any time, and Your sole remedy in the event of such error is stop using the Product.

3. TERMINATION

3.1. Termination. This License Agreement shall commence as of the date on which the submitted trial registration request has been received by Company and, unless terminated earlier in accordance with this License Agreement shall continue in perpetuity.

3.2. Company may terminate this EULA immediately and without notice if You fail to comply with any term of this EULA.

4. LIMITATION OF LIABILITY

4.1. Company PROVIDES THE PRODUCT WITHOUT ANY WARRANTIES OF ANY KIND, EXPRESS, IMPLIED, STATUTORY, OR IN ANY OTHER PROVISION OF THIS EULA OR COMMUNICATION WITH You. Company SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON- INFRINGEMENT.

4.2. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT WILL Company BE LIABLE FOR ANY LOST PROFITS OR BUSINESS OPPORTUNITIES, LOSS OF USE, BUSINESS INTERRUPTION, LOSS OF DATA, OR ANY OTHER INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES UNDER ANY THEORY OF LIABILITY, WHETHER BASED IN CONTRACT, TORT, NEGLIGENCE, PRODUCT LIABILITY, OR OTHERWISE.

5. GENERAL

5.1. Governing Law. This License Agreement shall be governed, construed and enforced in accordance with the laws of the United States of America and of the State of California.

5.2. Entire Agreement. This Agreement constitutes the entire and only agreement between the parties for Product 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 the duly authorized representatives of the parties.

5.3. 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.

5.4. 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.

5.5. United States Government End Users. For any Software 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.

5.6. 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.

5.7. Export Restrictions. You may not export or re-export the Product except in compliance with the United States Export Administration Act and the related rules and regulations and similar non-U.S. government restrictions, if applicable. The Product and accompanying documentation are deemed to be “commercial computer software” and “commercial computer software documentation” respectively, pursuant to DFAR Section 227.7202 and FAR Section 12.212(b), as applicable.

5.8. All disputes arising out of or relating to this EULA will be exclusively resolved 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.

5.9. Title. Company retains all right, title, and interest in and to the Software and the Software License Key and in all related copyrights, trade secrets, patents, trademarks, and any other intellectual and industrial property and proprietary rights, including registrations, applications, renewals, and extensions of such rights.

5.10. Contact Information. If You have any questions about this Agreement, or if You want to contact Company for any reason, please email sales@iXsystems.com.

Symbols

802.1Q, 87

A

Add Group, 20

Add User, 23

AFP, 144, 189

Alert, 260

Alert Services, 44

API, 311

Apple Filing Protocol, 144, 189

Attach the TrueNAS\ :sup:'®' Faceplate, 276

Autotune, 34

B

Become Familiar with the ES12 System, 288

Become Familiar with the ES24, 294

Become Familiar with the ES60, 301

Become Familiar with the X Series System, 278

Boot Environments, 30

C

CA, 45

Certificate Authority, 45

Certificates, 48

CIFS, 159, 208

Cloud Credentials, 43

Cloud Sync, 58

Compression, 98

Console Setup Menu, 14

Contacting iXsystems, 267

Create Dataset, 96

Create Group, 20

Create User, 23

Cron Jobs, 63

D

DC, 191

DDNS, 193

Delete Group, 21

Delete User, 25

Domain Controller, 191

Dynamic DNS, 193

E

Email, 34

Encryption, 91

ES12 Expansion Shelf, 287

ES12 Expansion Shelf Contents, 287

ES24 Expansion Shelf, 293

ES24 Expansion Shelf Contents, 293

ES60 Expansion Shelf, 300

ES60 Expansion Shelf Contents, 300

EtherChannel, 82

EULA, 315

F

Failover, 55

File Transfer Protocol, 194

FTP, 194

G

Groups, 19

GUI Access, 16

Guide, 260

H

Hot Spares, 113

I

Install Drive Trays into a TrueNAS Array, 275

Internet Small Computer System Interface, 169

iSCSI, 169

L

LACP, 82

LAGG, 82

Link Aggregation, 82

Link Layer Discovery Protocol, 199

LLDP, 199

Log Out, 258

M

Minio, 205

Multiple Boot Environments, 30

N

Netdata, 200

Network File System, 151, 201

Network Settings, 75

New Group, 20

New User, 23

NFS, 151, 201

O

OpenStack Cinder Driver, 240

Out-of-Band Management, 4

P

Path and Name Lengths, 3

Periodic Snapshot, 113

Plug in and Power on your TrueNAS\ :sup:® array, 276

Proactive Support, 53

Processes, 257

R

Rail Kit Assembly, 295

Reboot, 259

Remove Group, 21

Remove User, 25

Replace Failed Drive, 110

Replication, 115

Reporting, 247

RFC

RFC 2616, 312

RFC 3721, 171

Route, 86

Rsync, 203

Rsync Tasks, 66

S

S.M.A.R.T., 206

S.M.A.R.T. Tests, 73

S3, 205

Samba, 159, 208

SCP, 215

Scrub, 126

Secure Copy, 215

Secure Shell, 213

Services, 186

Shadow Copies, 168

Shell, 257

Shutdown, 259

Simple Network Management Protocol, 211

SMB, 159, 208

Snapshot, 113

Snapshots, 129

SNMP, 211

Spares, 113

SSH, 213

Start Service, 187

Static Route, 86

Stop Service, 187

Support, 51, 260

System Dataset, 35

T

Tasks, 57

TFTP, 215

Time Machine, 147

Trivial File Transfer Protocol, 215

Trunking, 87

Tunables, 36

U

Uninterruptible Power Supply, 216

Upgrade ZFS Pool, 42

UPS, 216

Users, 21

V

VAAI, 310

VAAI for iSCSI, 311

VAAI for NAS, 311

vCenter Copying iSCSI Configuration, 240

vCenter Datastores, 225

vCenter Hosts, 224

vCenter iSCSI Functions, 236

vCenter Periodic Snapshots, 233

vCenter Plugin, 221

vCenter Replication, 234

vCenter Role-Based Access Control (RBAC), 240

vCenter Snapshots, 232

vCenter Zvols, 236

VLAN, 87

VMware Snapshot, 131

Volumes, 89

W

WebDAV, 158, 220

Windows File Share, 208

Windows Shares, 159

Wizard, 249

X

X Series Unified Storage Array, 277

X Series Unified Storage Array Contents, 277

Z

Z Series Unified Storage Array, 267

ZVOL, 99