

TrueNAS® User Guide

April 2017 Edition

Copyright iXsystems 2011-2017

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

CONTENTS

Welcome	1
Typographic Conventions	2
1 Introduction	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	18
3.1 Groups	18
3.2 Users	21
4 System	25
4.1 Information	25
4.2 General	26
4.3 Boot	29
4.4 Advanced	32
4.4.1 Autotune	34
4.5 Email	34
4.6 System Dataset	36
4.7 Tunables	37
4.8 Update	40
4.8.1 Preparing for Updates	40
4.8.2 HA Updates	40
4.8.3 Updates and Trains	40
4.8.4 Checking for Updates	41
4.8.5 Applying Updates	41
4.8.6 Manual Updates	42
4.8.7 Updating from the Shell	42
4.8.8 Updating an HA System	42
4.8.9 If Something Goes Wrong	42
4.8.10 Upgrading a ZFS Pool	43
4.9 Cloud Credentials	44
4.10 CAs	44
4.11 Certificates	47
4.12 Support	51
4.13 Proactive Support	52
4.14 Failover	53
4.14.1 Failover Management	55

5	Tasks	57
5.1	Cloud Sync	57
5.1.1	Cloud Sync Example	60
5.2	Cron Jobs	62
5.3	Init/Shutdown Scripts	64
5.4	Rsync Tasks	65
5.4.1	Rsync Module Mode	67
5.4.2	Rsync over SSH Mode	68
5.5	S.M.A.R.T. Tests	71
6	Network	74
6.1	Global Configuration	74
6.2	Interfaces	75
6.3	IPMI	77
6.4	Link Aggregations	79
6.4.1	LACP, MPIO, NFS, and ESXi	80
6.4.2	Creating a Link Aggregation	80
6.5	Network Summary	84
6.6	Static Routes	84
6.7	VLANs	85
7	Storage	87
7.1	Volumes	87
7.1.1	Volume Manager	87
	Encryption	89
	Manual Setup	90
	Extending a ZFS Volume	91
7.1.2	Change Permissions	92
7.1.3	Create Dataset	94
	Compression	96
7.1.4	Create zvol	97
7.1.5	Import Disk	98
7.1.6	Import Volume	99
	Importing an Encrypted Pool	100
7.1.7	View Disks	100
7.1.8	View Enclosure	102
7.1.9	View Volumes	103
	Managing Encrypted Volumes	106
7.1.10	View Multipaths	109
7.1.11	Replacing a Failed Drive	109
	Replacing an Encrypted Drive	111
	Removing a Log or Cache Device	111
7.1.12	Replacing Drives to Grow a ZFS Pool	112
7.2	Periodic Snapshot Tasks	112
7.3	Replication Tasks	114
7.3.1	Examples: Common Configuration	114
	Alpha (Source)	114
	Beta (Destination)	115
7.3.2	Example: TrueNAS® to TrueNAS® Semi-Automatic Setup	115
7.3.3	Example: TrueNAS® to TrueNAS® or Other Systems, Manual Setup	117
	Encryption Keys	117
7.3.4	Replication Options	121
7.3.5	Replication Encryption	122
7.3.6	Limiting Replication Times	122

7.3.7	Replication Topologies and Scenarios	122
	Star Replication	123
	Tiered Replication	123
	N-way Replication	123
	Disaster Recovery	123
7.3.8	Troubleshooting Replication	124
	SSH	124
	Compression	124
	Manual Testing	124
7.4	Scrubs	125
7.5	Snapshots	126
7.6	VMware-Snapshot	128
8	Directory Services	130
8.1	Active Directory	130
	8.1.1 Troubleshooting Tips	135
	8.1.2 If the System Will not Join the Domain	135
8.2	LDAP	136
8.3	NIS	139
8.4	NT4	140
8.5	Kerberos Realms	141
8.6	Kerberos Keytabs	142
8.7	Kerberos Settings	142
9	Sharing	144
9.1	Apple (AFP) Shares	145
	9.1.1 Creating AFP Guest Shares	147
	9.1.2 Creating Authenticated and Time Machine Shares	149
9.2	Unix (NFS) Shares	153
	9.2.1 Example Configuration	157
	9.2.2 Connecting to the Share	157
	From BSD or Linux	157
	From Microsoft	158
	From Mac OS X	158
	9.2.3 Troubleshooting NFS	160
9.3	WebDAV Shares	161
9.4	Windows (SMB) Shares	162
	9.4.1 Configuring Unauthenticated Access	168
	9.4.2 Configuring Authenticated Access Without a Domain Controller	169
	9.4.3 Configuring Shadow Copies	171
9.5	Block (iSCSI)	173
	9.5.1 Target Global Configuration	174
	9.5.2 Portals	175
	9.5.3 Initiators	177
	9.5.4 Authorized Accesses	178
	9.5.5 Targets	180
	9.5.6 Extents	181
	9.5.7 Target/Extents	183
	9.5.8 Fibre Channel Ports	184
	9.5.9 Connecting to iSCSI	188
	9.5.10 Growing LUNs	189
	Zvol Based LUN	189
	File Extent Based LUN	190

10 Services	191
10.1 Control Services	191
10.2 AFP	193
10.2.1 Troubleshooting AFP	195
10.3 Domain Controller	195
10.3.1 Samba Domain Controller Backup	196
10.4 Dynamic DNS	197
10.5 FTP	198
10.5.1 Anonymous FTP	201
10.5.2 FTP in chroot	202
10.5.3 Encrypting FTP	203
10.5.4 Troubleshooting FTP	204
10.6 iSCSI	204
10.7 LLDP	204
10.8 NFS	205
10.9 Rsync	206
10.9.1 Configure Rsyncd	206
10.9.2 Rsync Modules	207
10.10 S.M.A.R.T.	208
10.11 SMB	210
10.11.1 Troubleshooting SMB	214
10.12 SNMP	215
10.13 SSH	217
10.13.1 SCP Only	218
10.13.2 Troubleshooting SSH	219
10.14 TFTP	219
10.15 UPS	220
10.15.1 Multiple Computers with One UPS	223
10.16 WebDAV	223
11 vCenter	225
12 Reporting	227
13 Wizard	229
14 Additional Options	236
14.1 Display System Processes	236
14.2 Shell	237
14.3 Log Out	238
14.4 Reboot	238
14.5 Shutdown	238
14.6 Support Icon	239
14.7 Guide	239
14.8 Alert	239
15 ZFS Primer	242
16 Hardware Setup	246
16.1 TrueNAS® Unified Storage Array	246
16.2 Hardware Installation	250
16.3 E16/E16F Expansion Shelf	255
16.4 E24 Expansion Shelf	264
17 VAAI	273

17.1 VAAI for iSCSI	273
17.2 VAAI for NAS	273
18 Using the API	275
18.1 A Simple API Example	276
18.2 A More Complex Example	277
19 Appendix A	279
Index	282

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.

Avago is a trademark of Avago Technologies.

Chelsio® is a registered trademark of Chelsio Communications.

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

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

Linux® is a registered trademark of Linus Torvalds.

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

UNIX® is a registered trademark of The Open Group.

VMware® is a registered trademark of VMware, Inc.

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

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

Typographic Conventions

Typographic Conventions

The TrueNAS® Administrator Guide uses these typographic conventions:

Table 1: Text Format Examples

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

INITIAL SETUP

Before beginning software configuration, please see the [Hardware Setup](#) (page 246) section for specific rack-ing 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 77).

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](#).

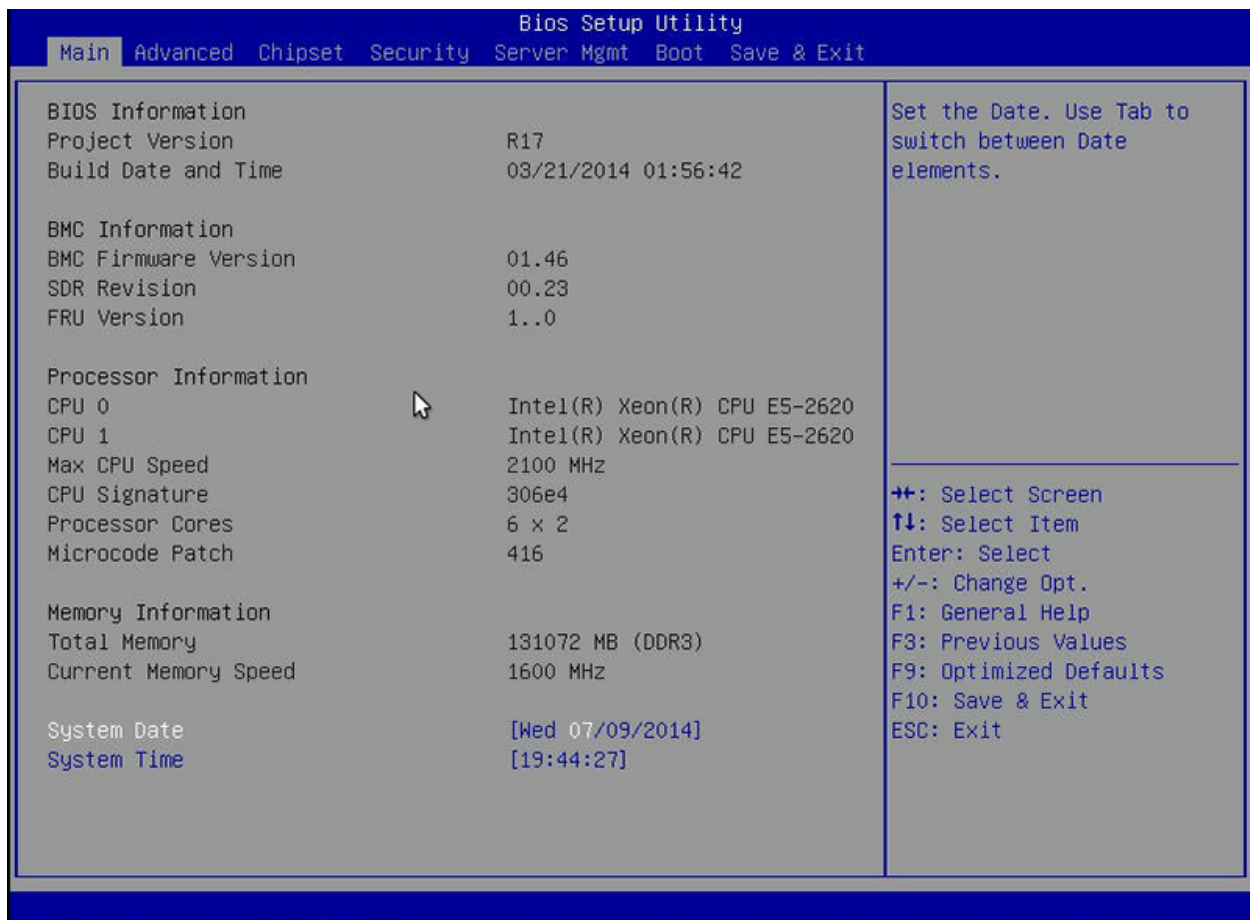


Fig. 2.1: Initial BIOS Screen

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

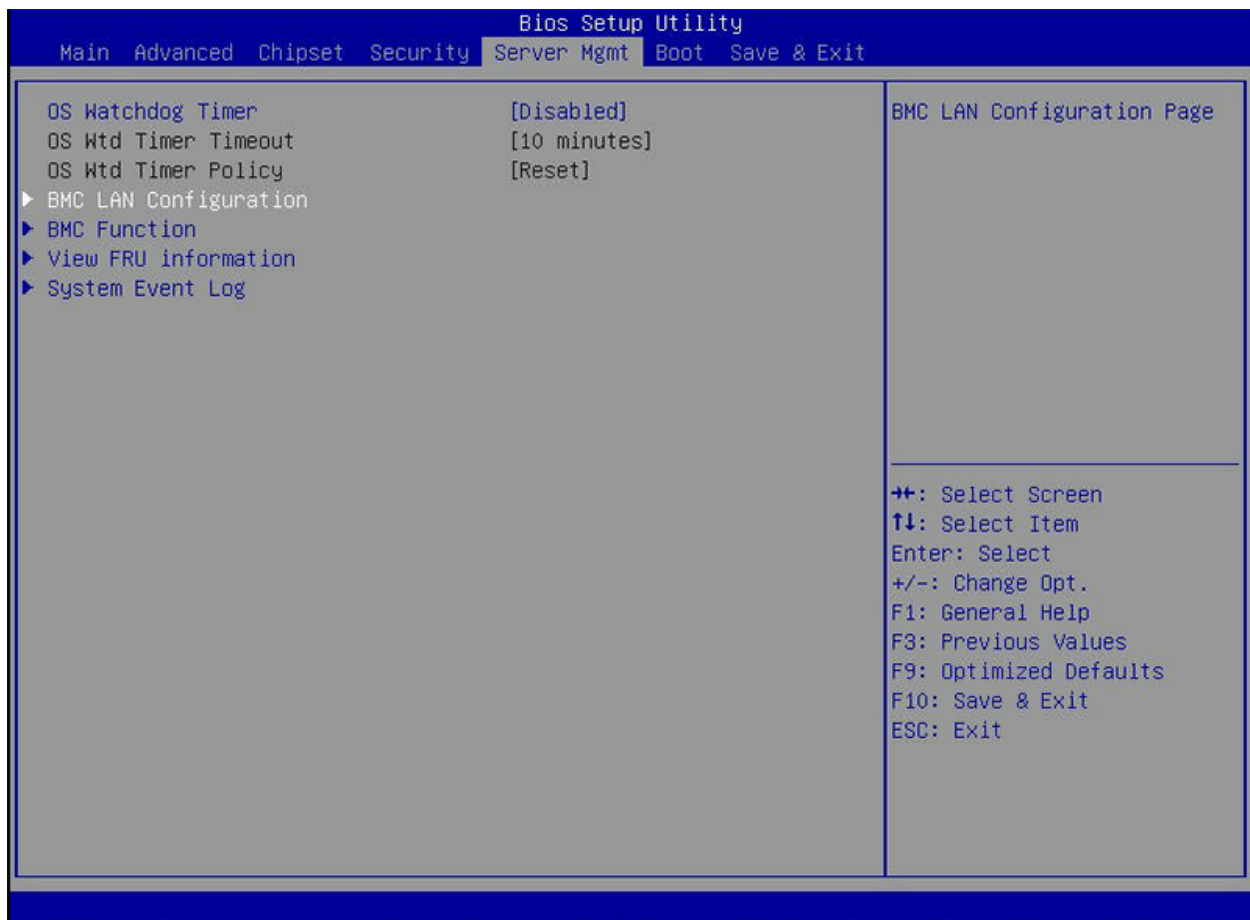


Fig. 2.2: Navigate to BMC LAN Configuration

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.

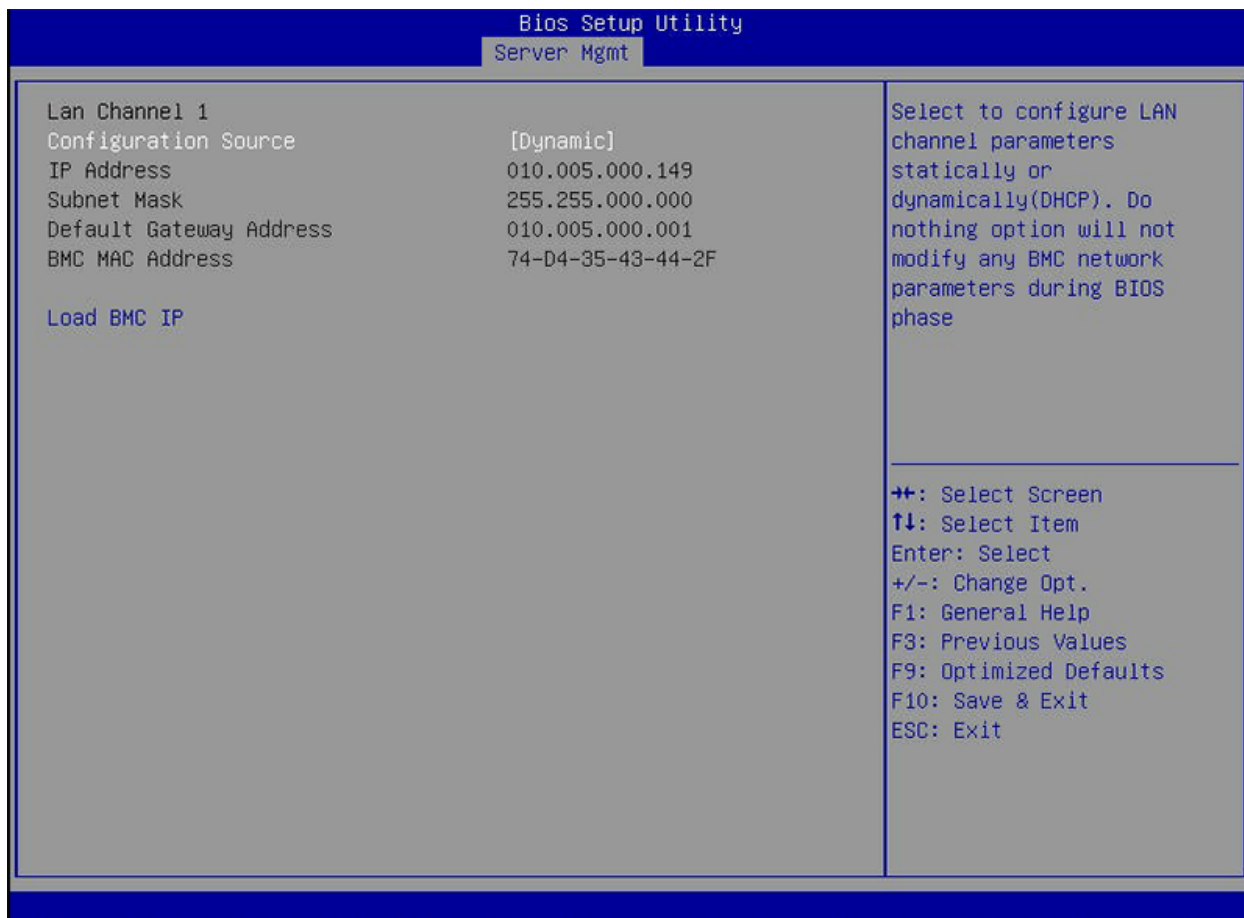


Fig. 2.3: Configuring a Dynamic IP Address

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.

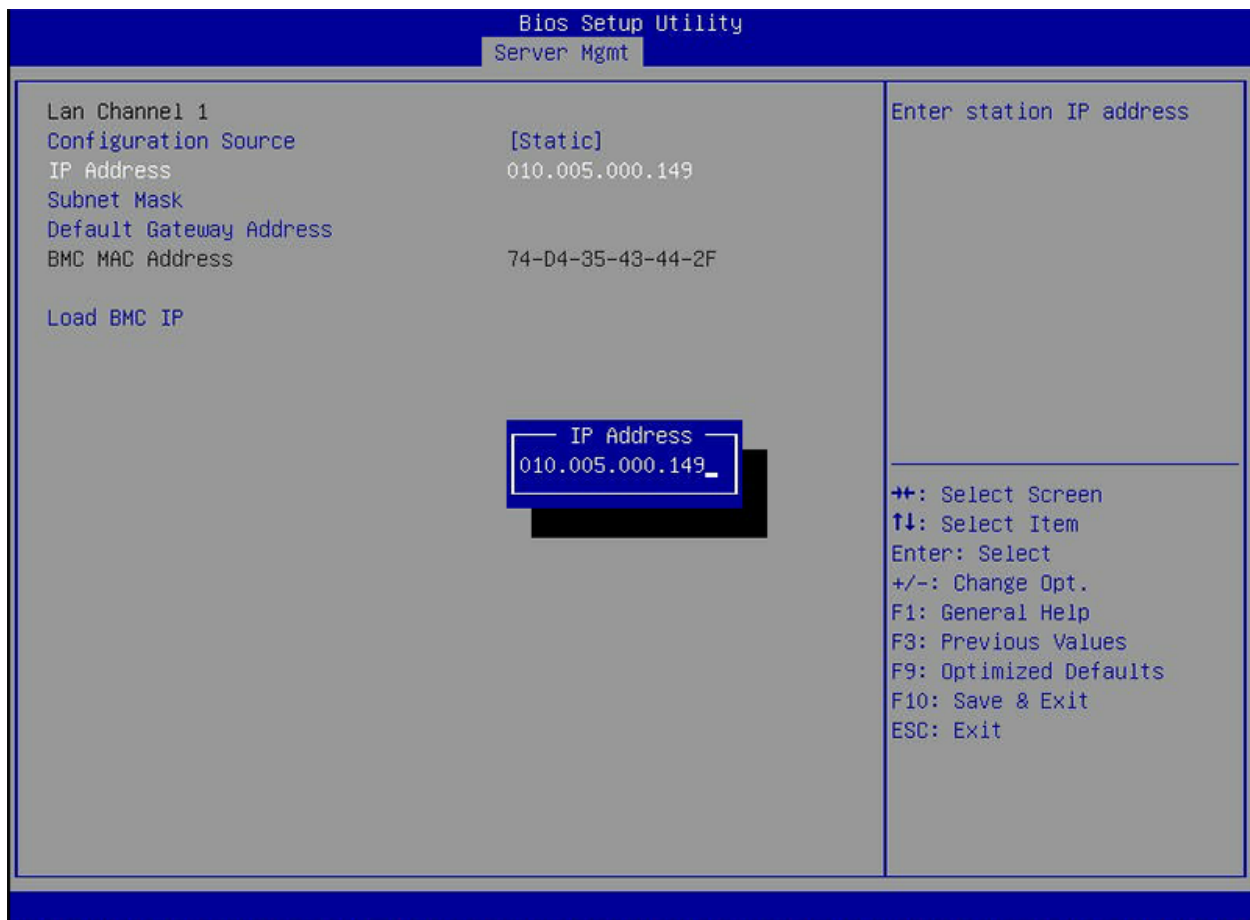


Fig. 2.4: Configuring a Static IP Address

Next, enter the *Subnet Mask* of the out-of-band management network subnet. An example is shown in Figure 2.5.

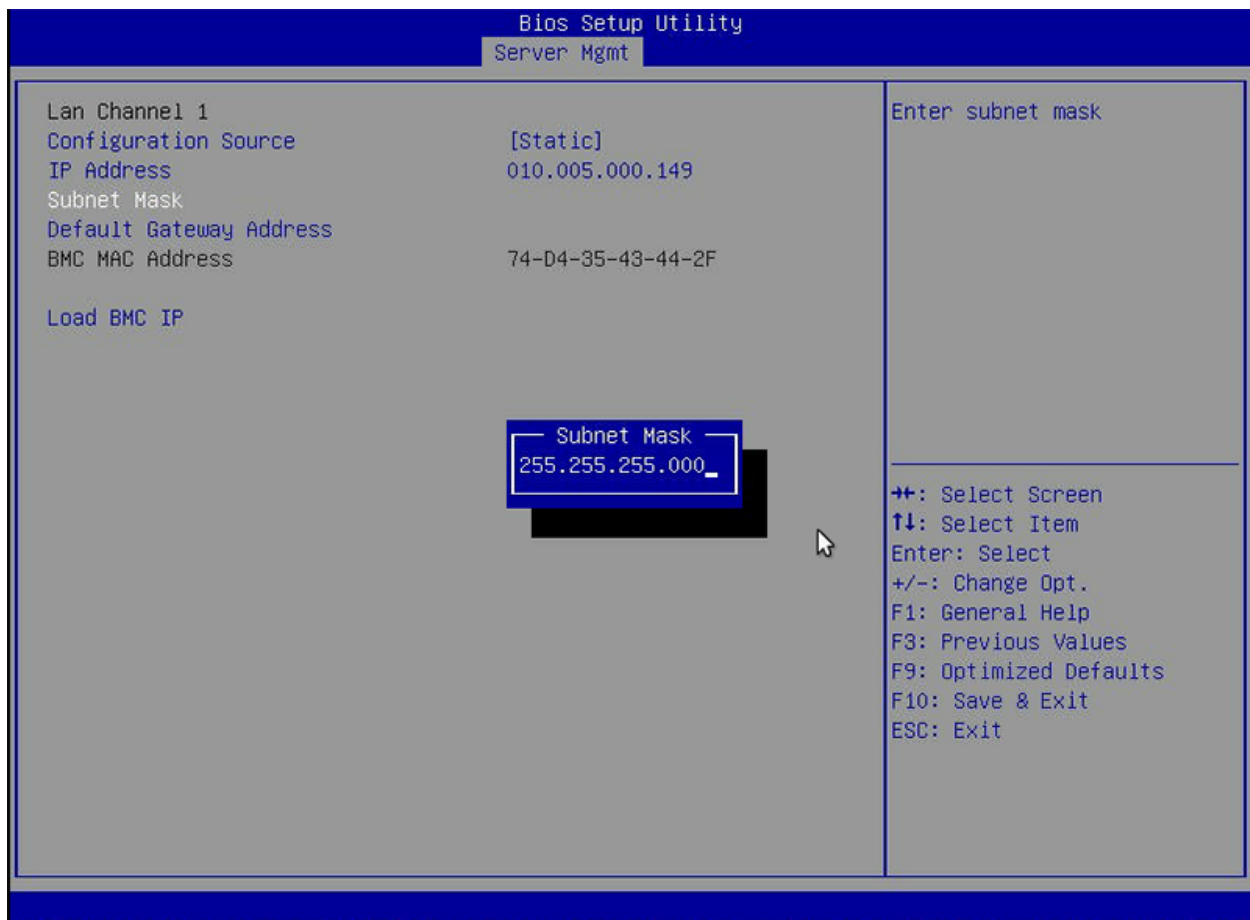


Fig. 2.5: Entering the Subnet Mask

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](#).

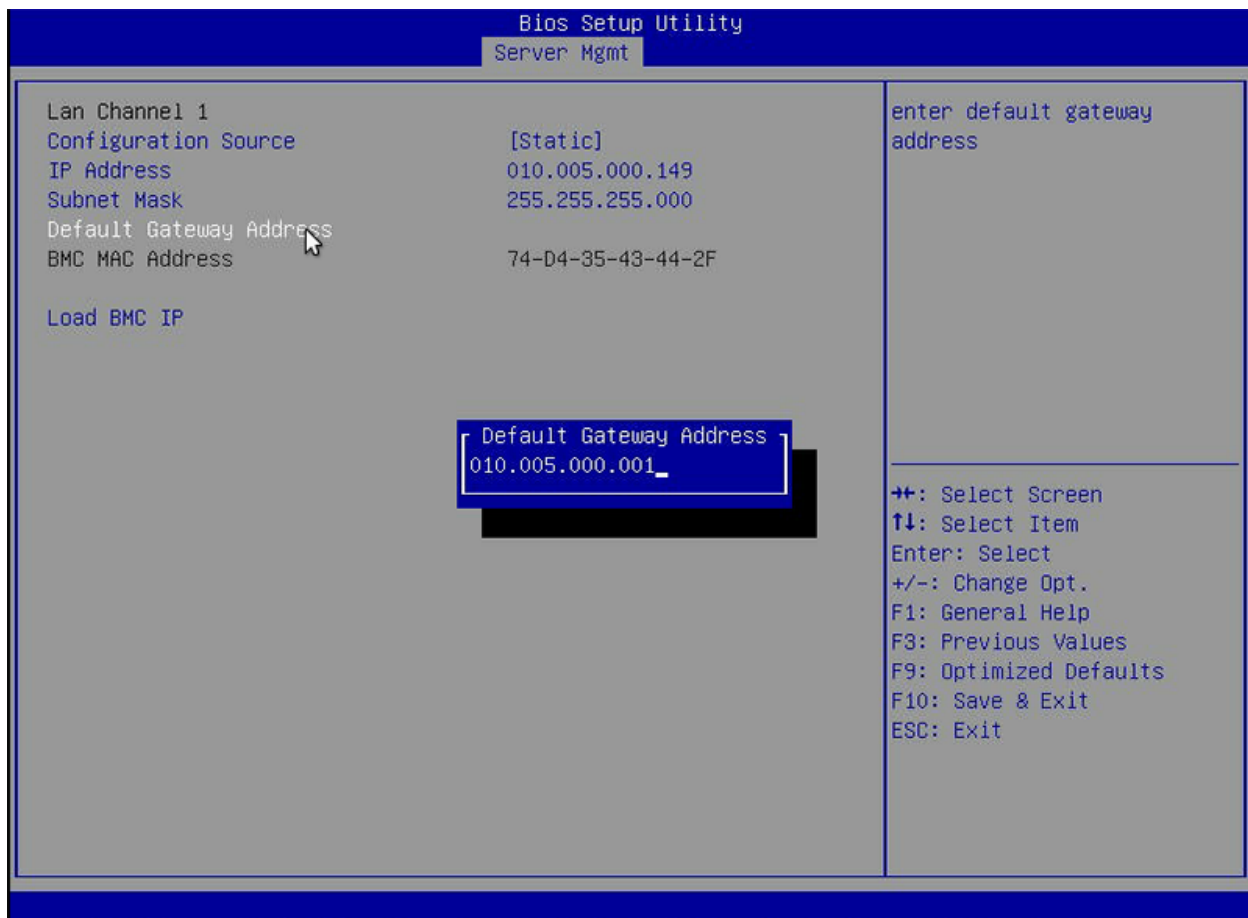


Fig. 2.6: Entering the Default Gateway Address

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.

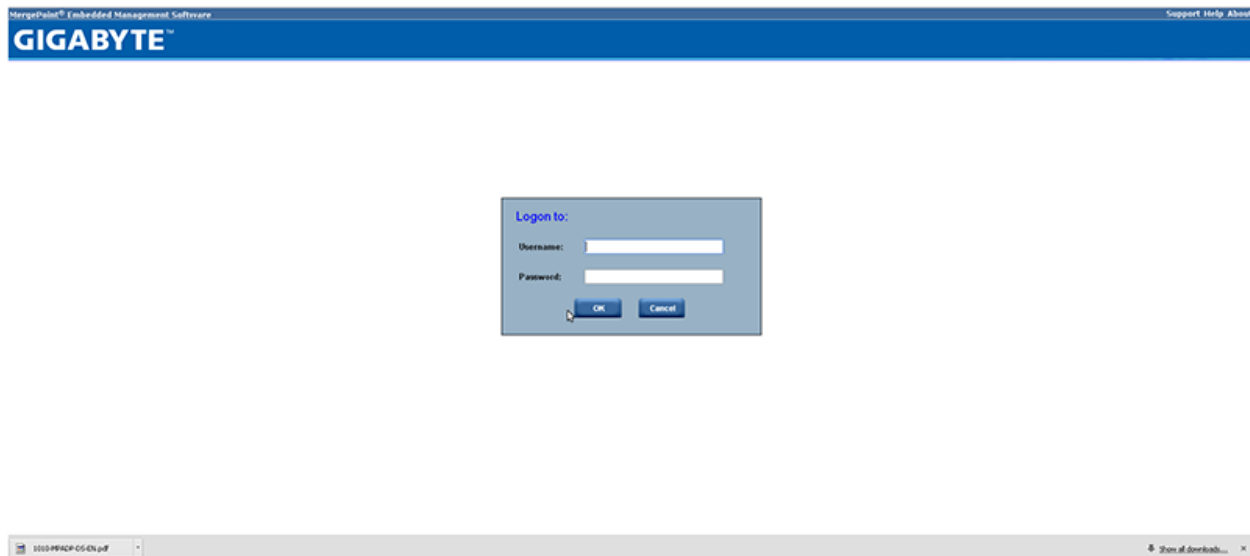


Fig. 2.7: Connecting to the IPMI Graphical Interface

Login using the default *Username* of *admin* and the default *Password* of *password*.

The administrative password can be changed using the instructions in [IPMI](#) (page 77).

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](#).

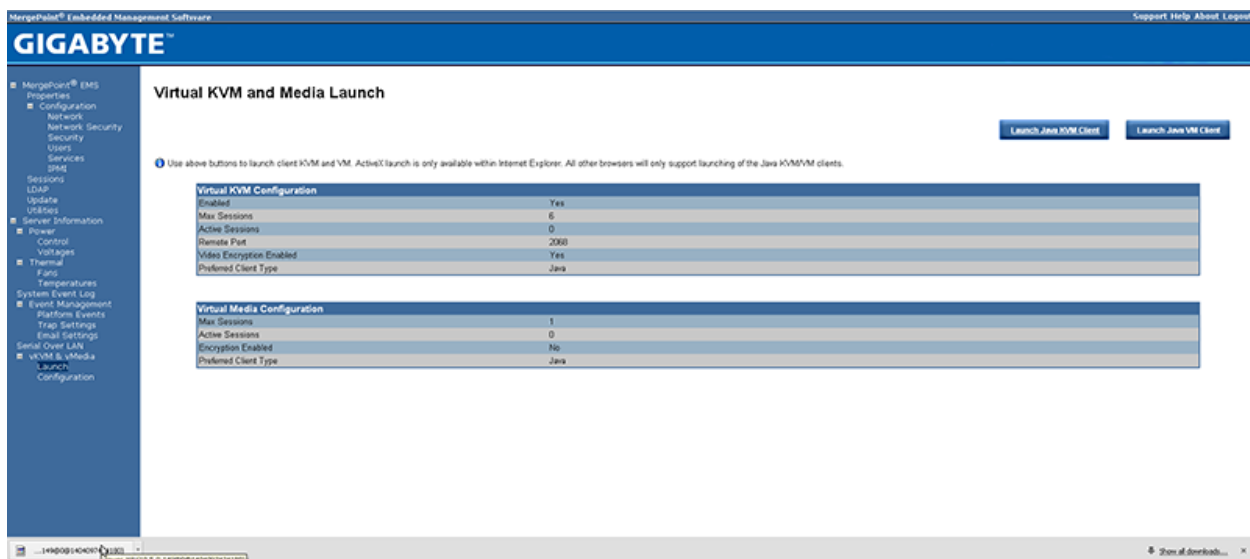


Fig. 2.8: Launching the Java KVM Client

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

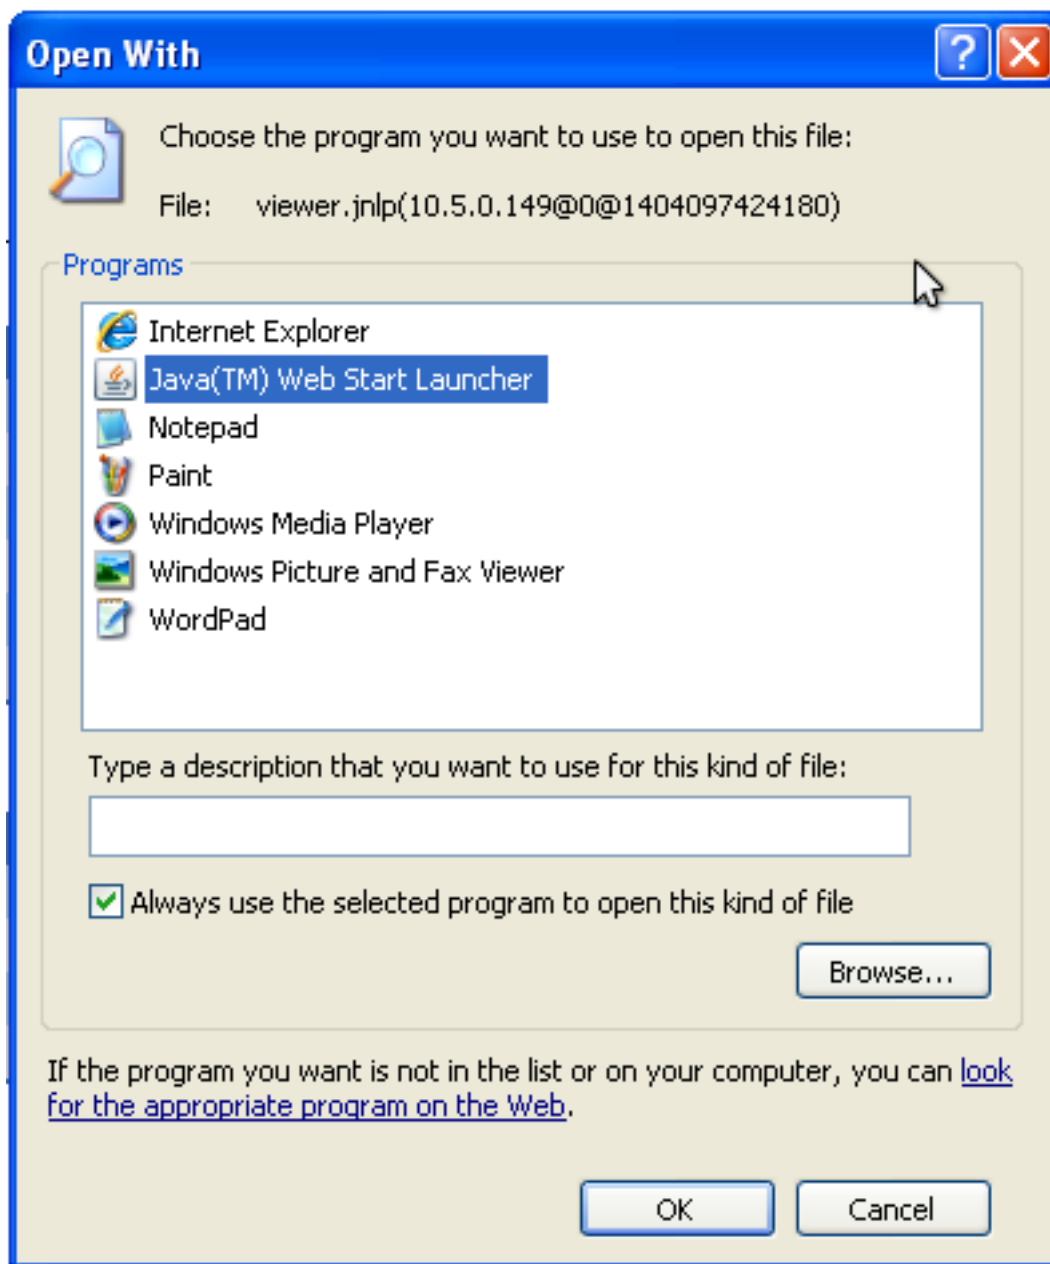


Fig. 2.9: Configure the Launch Program

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](#).

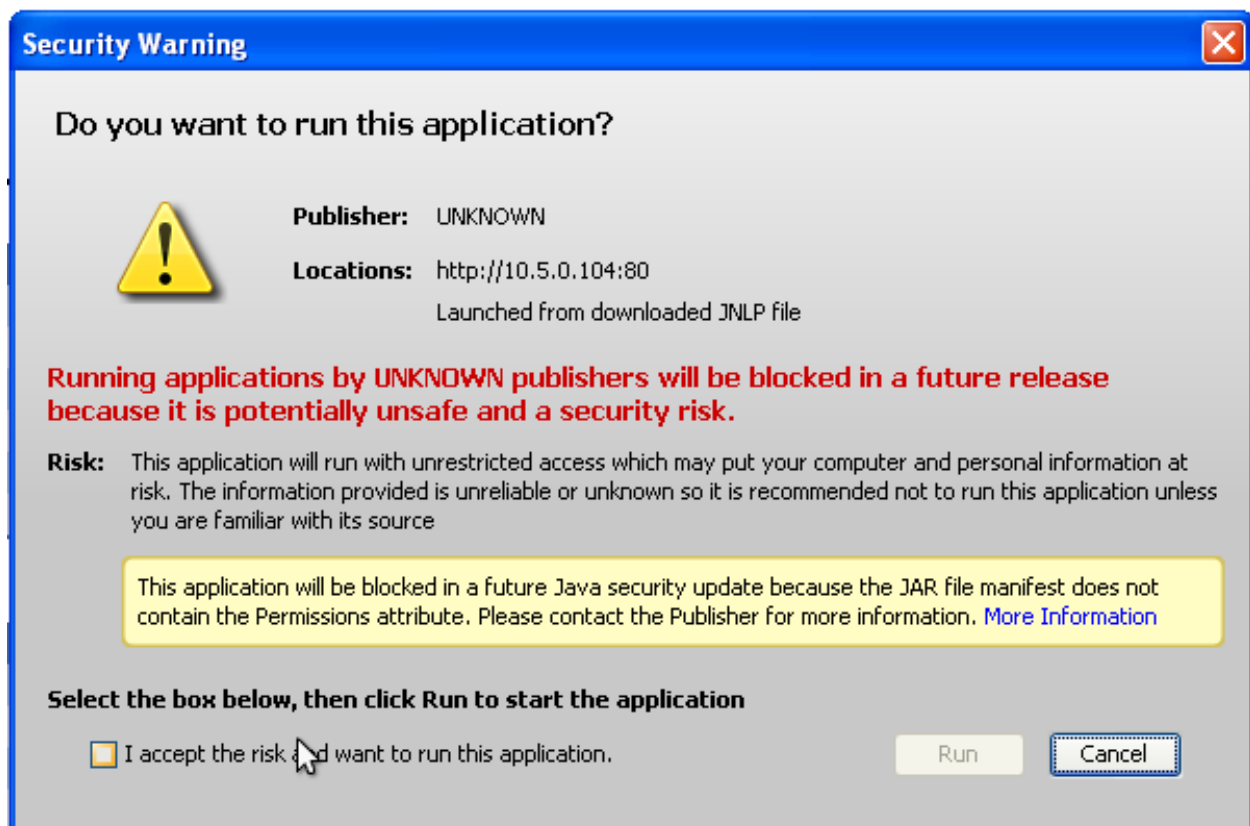


Fig. 2.10: Respond to Warning

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

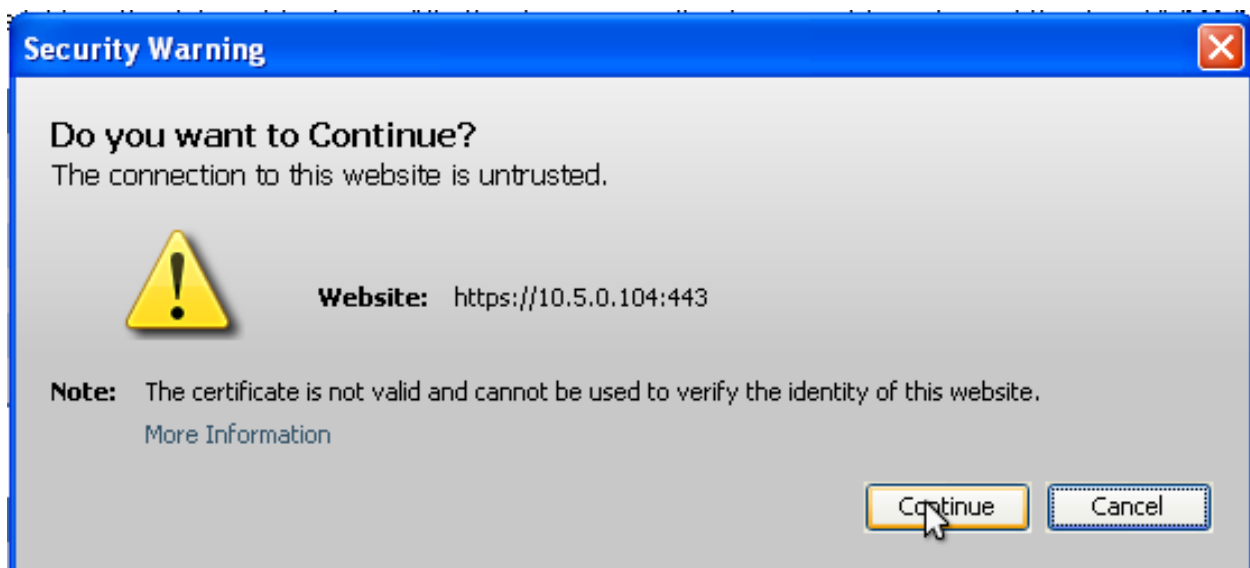


Fig. 2.11: Continue Through this Screen

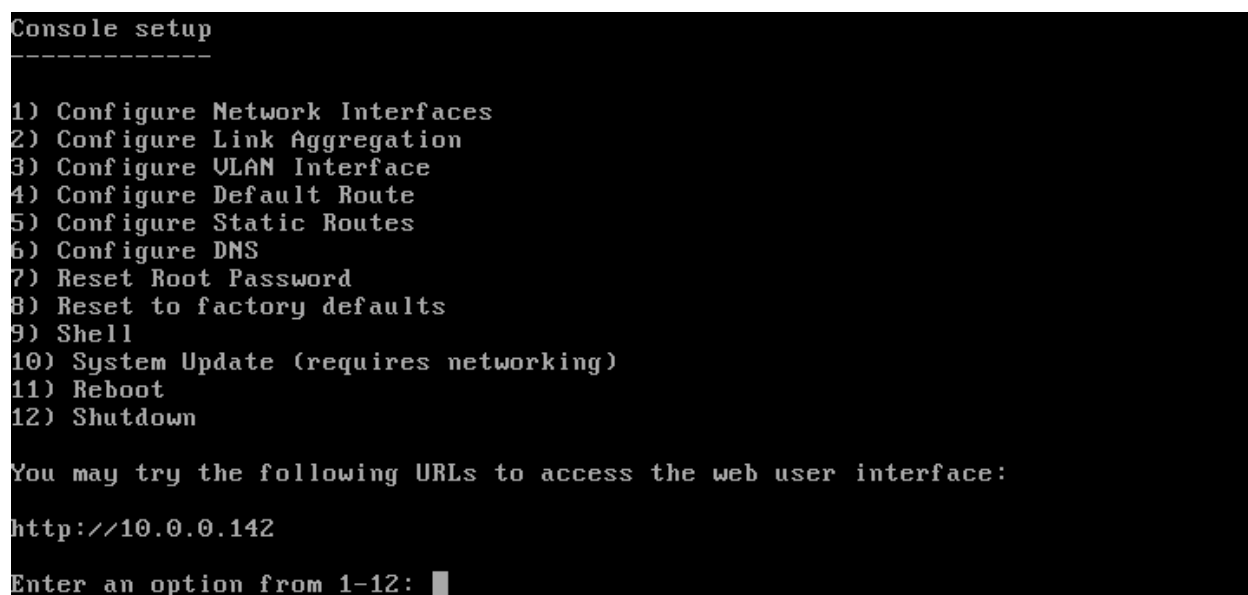
With the out-of-band console open, the TrueNAS® Storage Array can be controlled as if using a directly-

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 access to the TrueNAS® system's keyboard and monitor is available, this Console Setup menu can be used to administer the system even if the administrative GUI is not accessible.

Note: The Console Setup menu can be accessed from within the TrueNAS® GUI by typing `/etc/netcli` from [Shell](#) (page 237). The Console Setup menu can be disabled by unchecking the *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 to factory defaults
9) Shell
10) System Update (requires networking)
11) Reboot
12) Shutdown

You may try the following URLs to access the web user interface:
http://10.0.0.142

Enter an option from 1-12: █
```

Fig. 2.12: Console Setup Menu

This menu provides these options:

- 1) Configure Network Interfaces:** provides a configuration wizard to configure the system's network interfaces. If the system has been licensed for for High Availability (HA), the wizard will prompt to set the IP address for both (*This Node*) and (*Node B*).
- 2) Configure Link Aggregation:** allows creating a new link aggregation or deleting an existing link aggregation. If the system has been licensed for for High Availability (HA), this option prompts for the VHID when creating the link aggregation.
- 3) Configure VLAN Interface:** used to create or delete a VLAN interface.
- 4) Configure Default Route:** used to set the IPv4 or IPv6 default gateway. When prompted, input the IP address of the default gateway.
- 5) Configure Static Routes:** prompts for the destination network and the gateway IP address. Re-enter this option for each route to be added.
- 6) Configure DNS:** will prompt for the name of the DNS domain then the IP address of the first DNS server. To enter multiple DNS servers, press `Enter` to input the next one. When finished, press `Enter` twice to

leave this option.

7) Reset Root Password: if logging in to the graphical administrative interface fails, select this option and follow the prompts to set the *root* password.

8) Reset to factory defaults: to delete **all** of the configuration changes made in the administrative GUI, select this option. Once the configuration is reset, the system will reboot. It will be necessary to go to *Storage* → *Volumes* → *Import Volume* to re-import volumes.

9) Shell: starts a shell to run FreeBSD commands. To leave the shell, type **exit**.

10) System Update: if any system updates are available, they will automatically be downloaded and applied. The functionality is the same as described in *Update* (page 40), except that the updates will be applied immediately and access to the GUI is not required.

11) Reboot: reboot the system.

12) Shutdown: shut 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://10.0.0.142>.

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, the console network menu can be used to manually configure the interface as seen in *Example: Manually Setting an IP Address from the Console Menu* (page ??). In this example, the TrueNAS® system has one network interface, *em0*.

Manually Setting an IP Address from the Console Menu

```
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 as 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
You may try the following URLs to access the web user interface:
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.

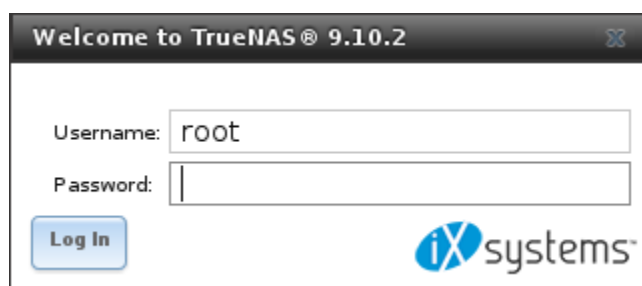


Fig. 2.13: Enter the Root Password

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](#) (page 279) 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.

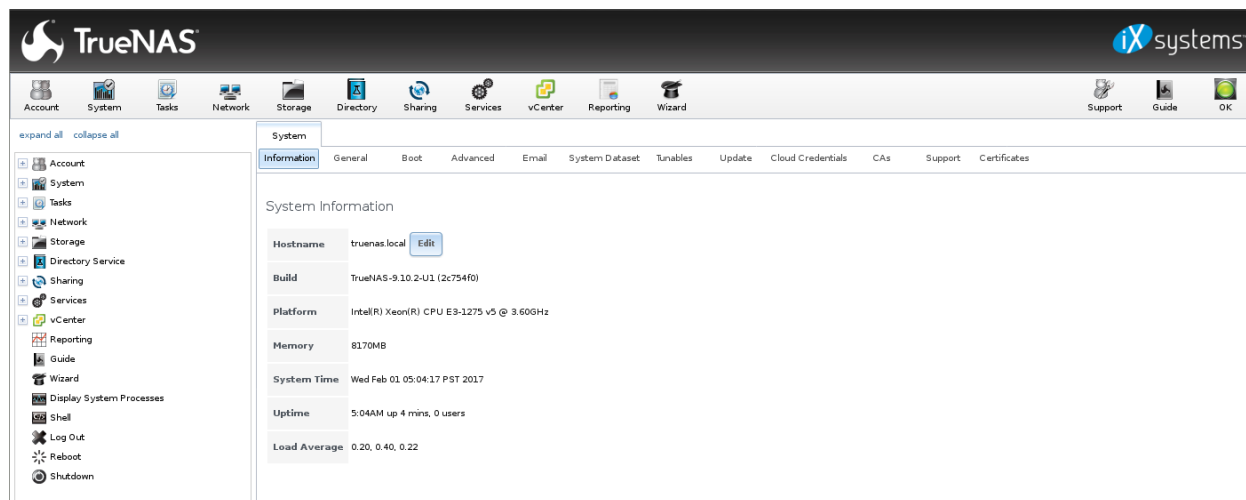


Fig. 2.14: TrueNAS® Graphical Configuration Menu

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 18): used to manage UNIX-style groups on the TrueNAS® system.
- [Users](#) (page 21): 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 130) for details.

This section describes how to create a group and assign user accounts to it. The next section, [Users](#) (page 21), 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 sudo (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.

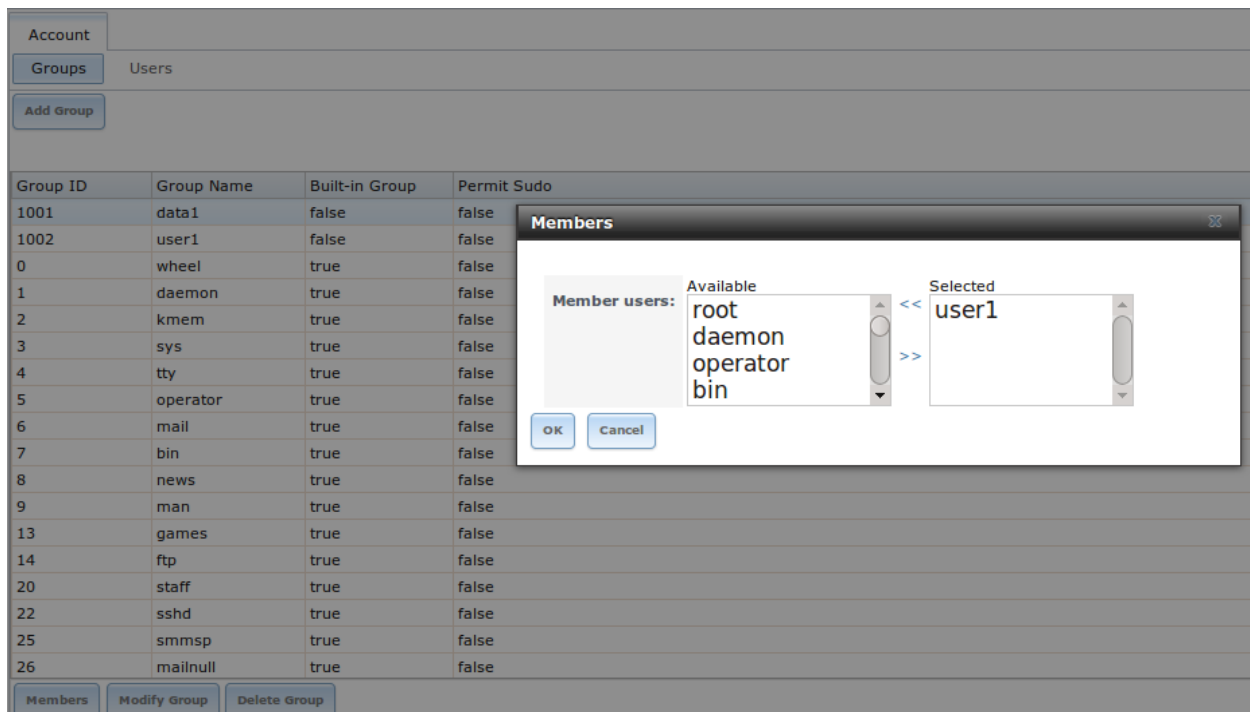


Fig. 3.3: Assigning a User to a Group

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 130).

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](#).

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	/usr/sbin/nologin	Sendmail Default	true		false	false	false	false

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 130), Windows user passwords must be set from within Windows.

The screenshot shows a window titled "Add User" with a close button in the top right corner. The window contains the following fields and controls:

- User ID:** A text input field containing the value "1001".
- Username:** An empty text input field.
- Create a new primary group for the user:** A checkbox that is checked.
- Primary Group:** A dropdown menu showing "-----".
- Create Home Directory In:** A text input field containing "/nonexistent" and a "Browse" button to its right.
- Shell:** A dropdown menu showing "csh".
- Full Name:** An empty text input field.
- E-mail:** An empty text input field.
- Password:** An empty text input field.
- Password confirmation:** An empty text input field with an information icon (i) to its right.
- Disable password login:** An unchecked checkbox.
- Lock user:** An unchecked 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
Continued on next page			

Table 3.2 – continued from previous page

Setting	Value	Advanced Mode	Description
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 <i>Create a new primary group</i> to access this menu; for security reasons, FreeBSD will not give a user su permissions if <i>wheel</i> is their primary group; to give a user su access, add them to the <i>wheel</i> group in <i>Auxiliary groups</i>
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 <i>Disable password login</i> ; cannot contain a ?
Password confirmation	string		must match the value of <i>Password</i>
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 <i>Modify User</i> button for the user in <i>View Users</i> ; checking this box grays out <i>Lock user</i> and <i>Permit Sudo</i> , 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 <i>Disable password login</i> which is mutually exclusive
Permit Sudo	checkbox		if checked, members of the group have permission to use sudo (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 can access the Console Setup menu shown in Figure 2.12 , even if it is disabled in System → Advanced → Enable Console Menu
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 ssh
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 scp and sftp 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 25) provides general TrueNAS® system information such as hostname, operating system version, platform, and uptime
- *General* (page 26) configures general settings such as HTTPS access, the language, and the timezone
- *Boot* (page 29) creates, renames, and deletes boot environments
- *Advanced* (page 32) 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 37) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- *Update* (page 40) performs upgrades and checks for system updates
- *Cloud Credentials* (page 44) is used to enter connection credentials for remote cloud service providers
- *CAs* (page 44): import or create internal or intermediate CAs (Certificate Authorities)
- *Support* (page 51): view licensing information or create a support ticket.
- *Proactive Support* (page 52): enable and configure automatic proactive support (Silver or Gold support coverage only).
- *Certificates* (page 47): import existing certificates or create self-signed certificates
- *Failover* (page 53): 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, and the current load average.

To change the system's hostname, click its *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* to the end of the hostname.

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
CAs
Support
Proactive Support
Certificates

System Information

Hostname	truenas.local	Edit
Build	TrueNAS-9.10.2-U2 (000c4d5c7)	
Platform	Intel(R) Xeon(R) CPU E3-1275 v5 @ 3.60GHz	
Memory	8170MB	
System Time	Wed Mar 29 12:18:40 PDT 2017	
Uptime	12:18PM up 1:24, 0 users	
Load Average	0.33, 0.31, 0.35	

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
CAs
Support
Proactive Support
Certificates

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:
Timezone: America/Los_Angeles
Syslog level: Info
Syslog server:

Save
Factory Restore
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 <i>HTTP</i> to <i>HTTPS</i> or to <i>HTTP+HTTPS</i> , select the certificate to use in <i>Certificate</i> ; if you do not have a certificate, first create a CA (in CAs (page 44)), then the certificate itself (in Certificates (page 47))
Certificate	drop-down menu	required for <i>HTTPS</i> ; 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 <i>0.0.0.0</i> (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, <i>HTTP</i> connections are automatically redirected to <i>HTTPS</i> if <i>HTTPS</i> is selected in <i>Protocol</i> , otherwise such connections will fail
Language	drop-down menu	select the localization from the drop-down menu and reload the browser; view the status of localization at pootle.freenas.org (http://pootle.freenas.org/)
Console Keyboard Map	drop-down menu	select the keyboard layout
Timezone	drop-down menu	select the timezone from the drop-down menu
Syslog level	drop-down menu	when <i>Syslog server</i> is defined, only logs matching this level are sent
Syslog server	string	<i>IP address_or_hostname:optional_port_number</i> 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:

Factory Restore: resets the configuration database to the default base version. However, it does not delete user SSH keys or any other data stored in a user's home directory. Since any configuration changes stored in the configuration database will be erased, this option is handy if a mistake has been made or to return a test system to the original configuration.

Save Config: saves a backup copy of the current configuration database in the format *hostname-version-architecture* to the system being used to access the administrative interface. It is recommended to always save the configuration after making any configuration changes. TrueNAS® automatically backs up the configuration database to the system dataset every morning at 3:45. However, this backup will 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 not be available. The location of the system dataset can be viewed or set using `System → System Dataset`.

Warning: Passwords are backed up with the system configuration. 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, must be stored in an encrypted form to prevent them from being visible as plain text in the saved system configuration. The key for this encryption is stored on the boot device. If TrueNAS® is installed on a new boot device and a backup system configuration is moved to that new boot device, the key is not present and these other passwords must be re-entered.

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.

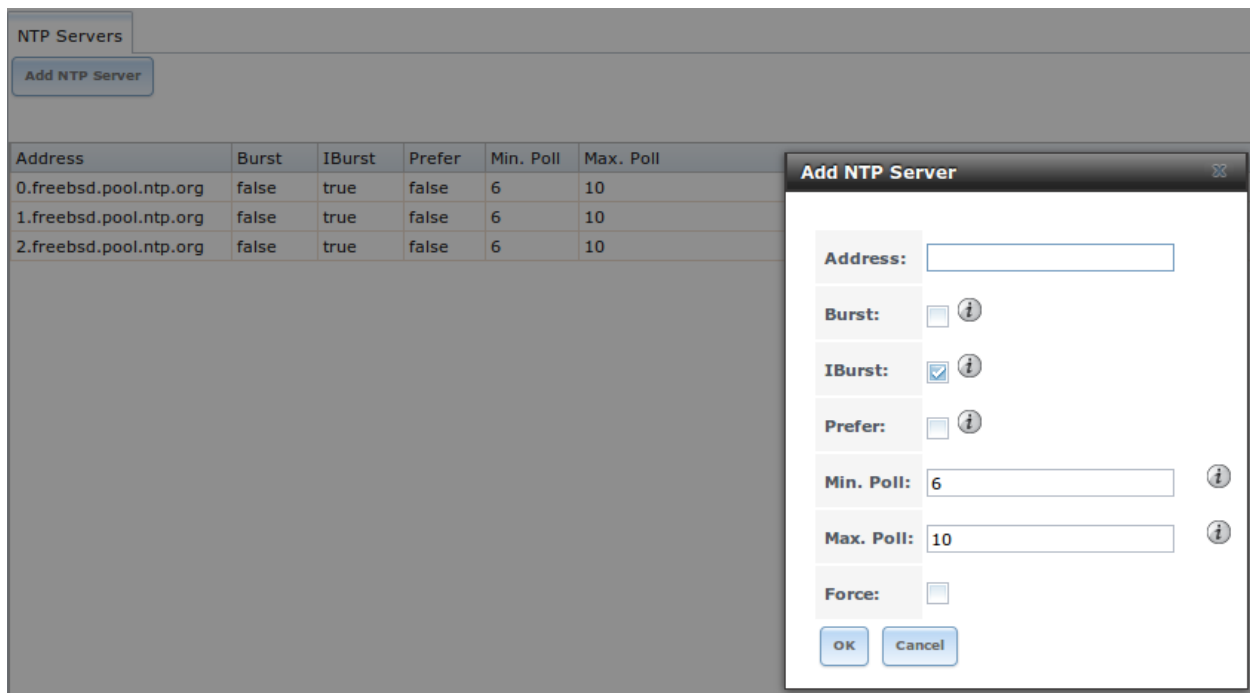


Fig. 4.3: Add an NTP Server

Table 4.2: NTP Servers Configuration Options

Setting	Value	Description
Address	string	name of NTP server
Burst	checkbox	recommended when <i>Max. Poll</i> 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 <i>Max. Poll</i>
Max. Poll	integer	power of 2 in seconds; cannot be higher than 17 or lower than <i>Min. Poll</i>
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 snap-

shot 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 229) was used, a third boot environment called *Wizard-date* is also created, indicating the date and time the *Wizard* (page 229) was run.

System			
Information	General	Boot	Advanced
Email	System Dataset	Tunables	Update
Cloud Credentials	CAs	Support	Proactive Support
Certificates			
Create	Scrub Boot	Status	
Boot Volume Condition: HEALTHY Size: 15.9 GiB			
Used: 712.2 MiB (4%)			
35 Automatic scrub interval (in days)			
Name	Active	Created	Keep
default	On Reboot, Now	2017-03-29 04:49:00	No
Initial-Install		2017-03-29 10:54:00	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 its configuration buttons. 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 will prompt you to input 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 35 days. To change the default interval, input a different 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. In the example shown in [Figure 4.5](#), there is only one boot device and it 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



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.

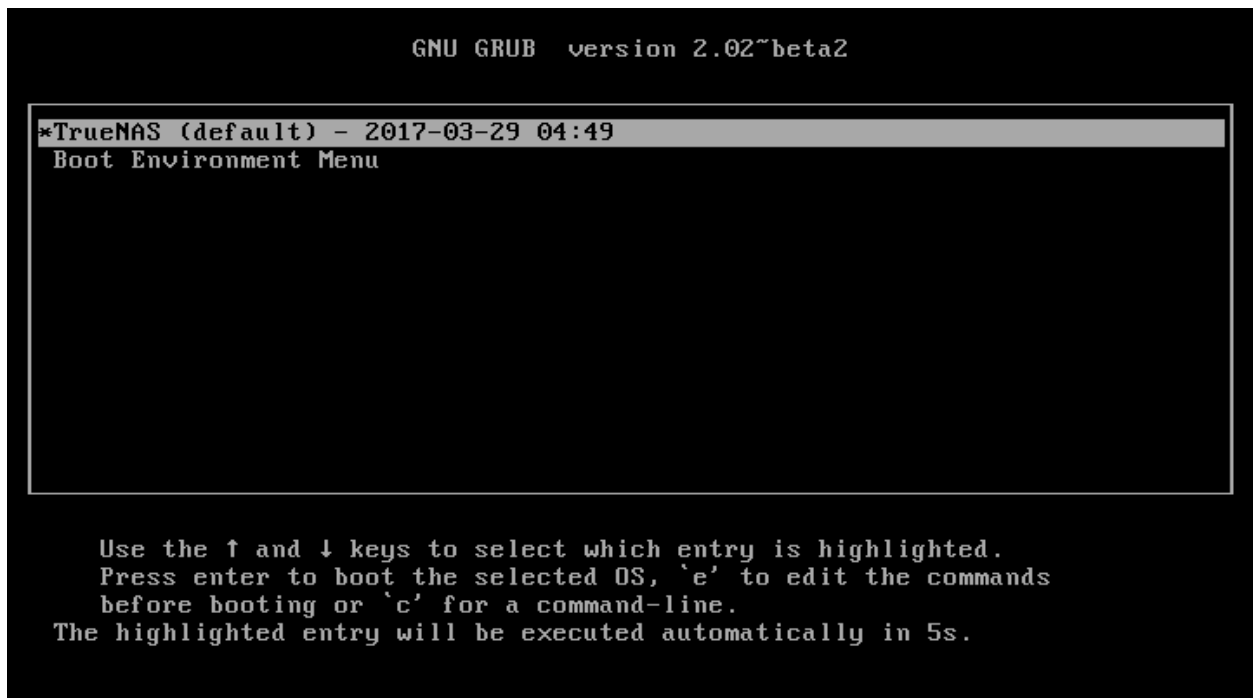


Fig. 4.6: Boot Environments in Boot Menu

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
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
CAs
Support
Proactive Support
Certificates

Enable Console Menu:

☒

Use Serial Console:

☐

Serial Port Address:

0x2f8 ⓘ

Serial Port Speed:

9600 ⓘ

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 ⓘ

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
Enable Console Menu	checkbox	unchecking this box removes the console menu shown in Figure 2.12
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
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 <i>Stop refresh</i> box in the scrollable screen to pause updating and uncheck the box to continue to watch the messages as they occur
Continued on next page		

Table 4.3 – continued from previous page

Setting	Value	Description
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 <i>Advanced Mode</i> 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, collectd 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
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 239) events are also emailed to the *root* user account. Problems with [Scrubs](#) (page 125) are reported separately in an email sent at 03:00AM.

Note: [S.M.A.R.T.](#) (page 208) 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 **Users** → **View Users**, click on *root* to highlight that user, then click *Change E-mail*. Enter the email address on the remote system where email is to be sent, like *admin@example.com*.

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

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 <i>Plain</i> , <i>SSL</i> , or <i>TLS</i>
Use SMTP Authentication	checkbox	enable/disable SMTP AUTH (http://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN SASL; if checked, enter the required <i>Username</i> and <i>Password</i>
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/) (<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.

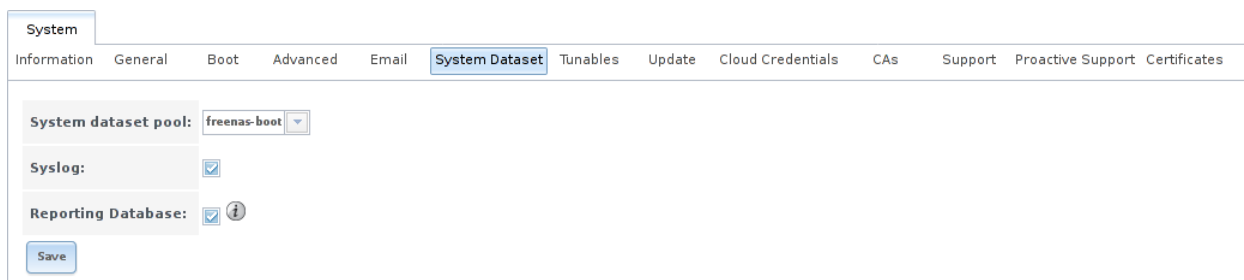


Fig. 4.9: System Dataset Screen

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 227) 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.

If you make any changes, click the *Save* button to save them.

If you change the pool storing the system dataset at a later time, TrueNAS® will automatically migrate 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>) 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&apropos=0&sektion=0&man=5&arch=default&format=html>) 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 <i>Variable</i> ; typically this is set to <i>YES</i> to enable the sysctl or driver specified by the “Variable”
Type	drop-down menu	choices are <i>Loader</i> , <i>rc.conf</i> , or <i>Sysctl</i>
Comment	string	optional, but a useful reminder for the reason behind adding this tunable
Enabled	checkbox	uncheck if you would like to disable the tunable without deleting it

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

Any tunables that you add will be listed in *System* → *Tunables*. To change the value of an existing tunable, click its *Edit* button. To remove a tunable, click its *Delete* button.

Some sysctls are read-only, meaning that they require a reboot in order to enable their setting. You can determine if a sysctl is read-only by first attempting to change it from *Shell* (page 237). For example, to change the value of *net.inet.tcp.delay_ack* to 1, use the command **sysctl net.inet.tcp.delay_ack=1**. If the sysctl value is read-only, an error message indicates that the setting is read-only. If no error is shown, the setting is now applied. For the setting to be persistent across reboots, the sysctl must still be added in *System* → *Tunables*.

The GUI does not display the sysctls that are pre-set when TrueNAS® is installed. TrueNAS® 9.10.2 ships with the following sysctls set:

```
kern.metadelat=3
kern.dirdelay=4
```

```
kern.filedelay=5
kern.coredump=0
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® 9.10.2 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"
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"
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_nvdim: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 9.10.2 deprecates these tunables:

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

After upgrading from an earlier version of 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 29) 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 239) 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 HA Updates

In HA (High Availability) systems, online upgrades usually cause a single failover event in each node. As the master node is updated, it fails over to the secondary node. Then the secondary node is updated, causing a failover back to the original master. These failovers cause short disruptions, usually less than 30 seconds for each.

4.8.3 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.

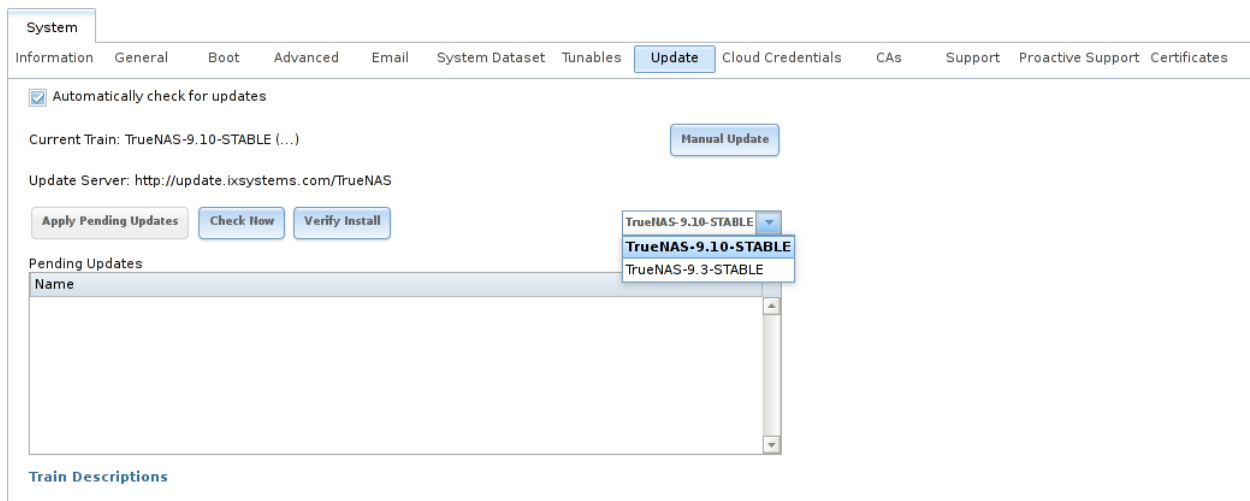


Fig. 4.11: Update Options

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 indicates which software branch, or *train*, is being tracked for updates. These trains are available:

- **TrueNAS-9.10-STABLE:** this is the **recommended train for production use**. Once new fixes and features have been tested as production-ready, they are added to this train. It is recommended to follow this train and to apply any of its pending updates.
- **TrueNAS-9.3-STABLE:** this is the maintenance-only mode for an older version of TrueNAS®. Unless an iX support engineer indicates otherwise, it is recommended to upgrade to *TrueNAS-9.10-STABLE*, by selecting that train, to ensure that the system receives bug fixes and new features.

The *Verify Install* button goes through the operating system files in the current installation, looking for any inconsistencies. When finished, a pop-up menu lists any files with checksum mismatches or permission errors.

4.8.4 Checking for Updates

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

4.8.5 Applying Updates

Make sure the system is in a low-usage state as described above in *Preparing for Updates* (page 40).

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 29) 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.

4.8.6 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.

Update files end with a `.tar` suffix.

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

4.8.7 Updating from the Shell

Updates can also be performed from the [Shell](#) (page 237) 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`.

4.8.8 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.

At this point, the active node may issue an alert indicating that there is a firmware version mismatch. This is expected when an update also updates the HBA version.

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.

4.8.9 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.

4.8.10 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 your 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**. It is not necessary to upgrade the pool if you do not need newer feature flags or if you want to keep the possibility of reverting to an earlier version of TrueNAS® or repurposing the disks in another operating system that supports ZFS. If you decide to upgrade the pool 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.

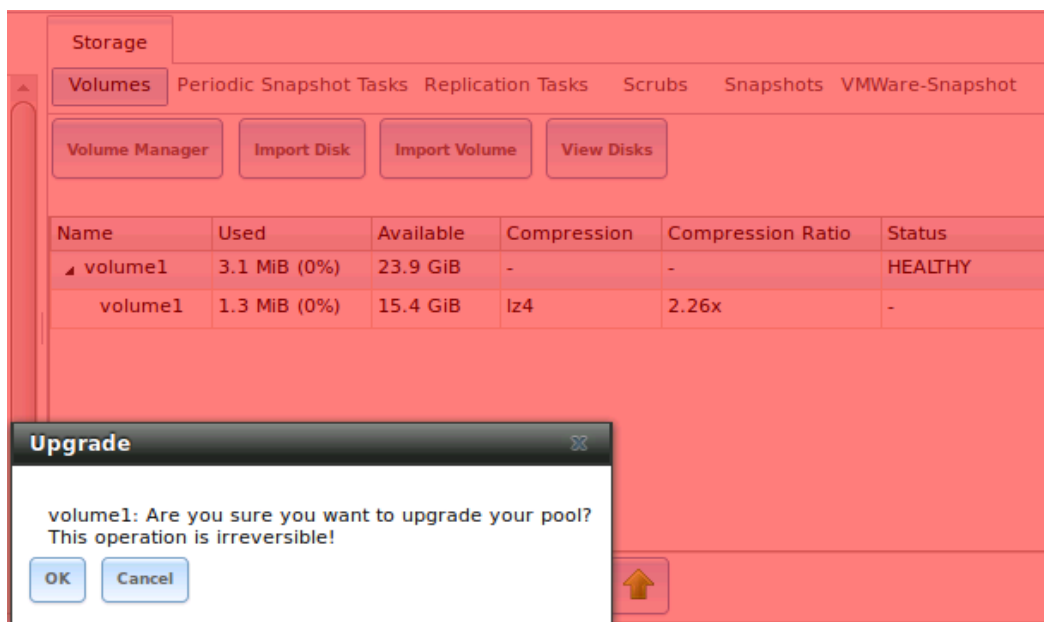


Fig. 4.12: Upgrading a ZFS Pool

The warning reminds you that a pool upgrade is irreversible. 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 57). The credentials to provide secure connections with cloud services are entered here. Select *System* → *Cloud Credentials* → *Add Cloud Credential* to display the dialog shown in [Figure 4.13](#).

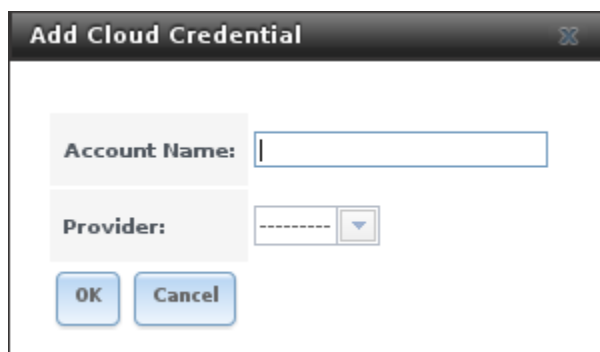


Fig. 4.13: Adding Cloud Credentials

The options are shown in [Table 4.6](#).

Table 4.6: Cloud Credential Options

Setting	Value	Description
Account Name	string	required; enter a descriptive name for the cloud credential
Provider	drop-down menu	required; select a cloud service provider
Access Key	string	shown when Amazon S3 is the <i>Provider</i> ; paste the Amazon account Access Key
Private Key	string	shown when Amazon S3 is the <i>Provider</i> ; paste the Amazon account Secret Key

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 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 4.14 shows the screen after clicking *System* → *CAs*.

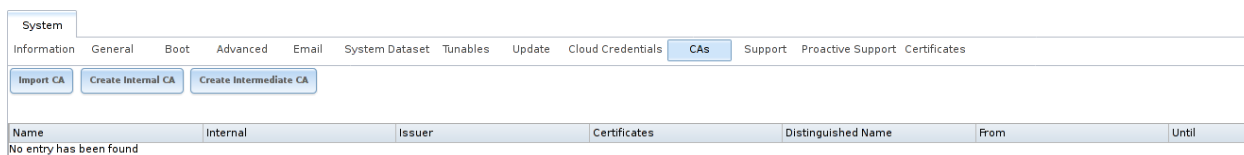


Fig. 4.14: Initial CA Screen

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 4.15. The configurable options are summarized in Table 4.7.

Fig. 4.15: Importing a CA

Table 4.7: Importing a CA Options

Setting	Value	Description
Name	string	mandatory; enter a descriptive name for the CA
Certificate	string	mandatory; paste in the certificate for the CA
Private Key	string	if there is a private key associated with the <i>Certificate</i> , paste it here
Passphrase	string	if the <i>Private Key</i> is protected by a passphrase, enter it here and repeat it in the "Confirm Passphrase" field
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 4.16.

Fig. 4.16: 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
Name	string	required; enter a descriptive name for the CA
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
Common Name	string	required; enter the fully-qualified hostname (FQDN) of the TrueNAS® system

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 4.16:

- **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:

- **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.

4.11 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 44).

Figure 4.17 shows the initial screen if you click `System → Certificates`.

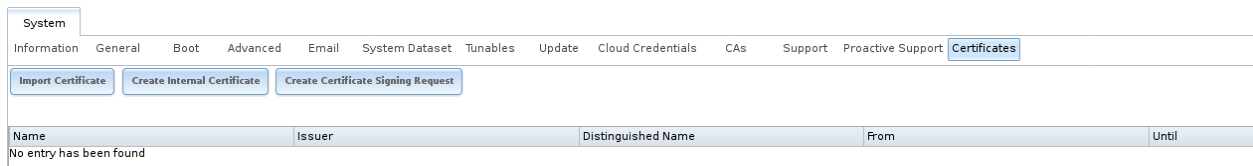
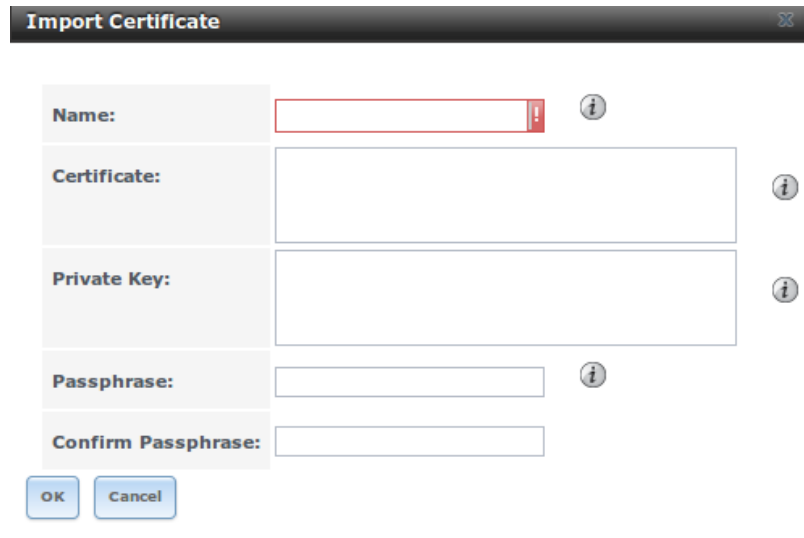


Fig. 4.17: Initial Certificates Screen

To import an existing certificate, click the *Import Certificate* button to open the configuration screen shown in [Figure 4.18](#). When importing a certificate chain, paste the primary certificate, followed by any intermediate certificates, followed by the root CA certificate.

The configurable options are summarized in [Table 4.9](#).



The image shows a dialog box titled "Import Certificate" with a close button (X) in the top right corner. It contains five input fields, each with an information icon (i) to its right:

- Name:** A text input field with a red border and a red exclamation mark icon on the right.
- Certificate:** A large text area for pasting the certificate content.
- Private Key:** A large text area for pasting the private key.
- Passphrase:** A text input field.
- Confirm Passphrase:** A text input field.

At the bottom left of the dialog are two buttons: "OK" and "Cancel".

Fig. 4.18: Importing a Certificate

Table 4.9: Certificate Import Options

Setting	Value	Description
Name	string	required; enter a descriptive name for the certificate; cannot contain the " (quote) character
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 <i>Confirm Passphrase</i> field

To instead create a new self-signed certificate, click the *Create Internal Certificate* button to see the screen shown in [Figure 4.19](#). 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 you imported or created using [CAs](#) (page 44) as the signing authority.

Fig. 4.19: 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 44)
Name	string	required; enter a descriptive name for the certificate; cannot contain the " (quote) character
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
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 TrueNAS® system

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 4.19](#) opens, but without the *Signing Certificate Authority* field.

All certificates that you import, self-sign, or make a certificate signing request for will be added as entries to `System → Certificates`. In the example shown in Figure 4.20, 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 it is available as a configurable option for encrypting connections.

System

Information

General

Boot

Advanced

Email

System Dataset

Tunables

Update

CAs

Certificates

Import Certificate

Create Internal Certificate

Create Certificate Signing Request

Name	Issuer	Distinguished Name	From	Until
self-signed cert	myCA	/C=US/ST=CA/L=San Jose/O=My Company/CN=Cert Wrangler/emailAddress=	Thu Nov 20 20:30:39 2014	Sun Nov 17 20:30:39 2024
verisign cert	external - signature pending	/C=US/ST=CA/L=San Jose/O=My Company/CN=Cert Wrangler/emailAddress=		

View

Export Certificate

Export Private Key

Delete

4.12 Support

The TrueNAS® *Support* tab, shown in Figure 4.21, 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 CAs **Support** Proactive Support Certificates

License Information [Update License](#) [Userguide \(PDF\)](#)

Model Z20 System Serial Contract Type Gold Contract Date

Customer Name Features None Additional Hardware None

Name

E-mail

Phone

Category Bug

Environment Production

Criticality Inquiry

Attach Debug Info ☒

Subject

Description

Attachments

X Browse... No file selected.

+

Submit

Fig. 4.21: 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.

4.13 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.

System
Information
General
Boot
Advanced
Email
System Dataset
Tunables
Update
Cloud Credentials
CAs
Support
Proactive Support
Certificates

Call iXsystems to upgrade to Silver/Gold support.

Enable automatic support alerts to iXsystems (Silver/Gold support only):
☐

Name of Primary Contact:

Title:

E-mail:

Phone:

Name of Secondary Contact:

Secondary Title:

Secondary E-mail:

Secondary Phone:

Upgrade to Silver/Gold support for this feature

Fig. 4.22: 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*.

4.14 Failover

If the TrueNAS® array has been licensed for High Availability (HA), a *Failover* tab is added to *System*. HA-licensed arrays use the Common Address Redundancy Protocol (CARP (<http://www.openbsd.org/faq/pf/carp.html>)) 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. TrueNAS® uses a two-unit active/standby model and provides an HA synchronization daemon to automatically monitor the status of the active node, synchronize any configuration changes between the active and the standby node, and failover to the standby node should the active node become unavailable.

Warning: Seamless failover is only available with iSCSI or NFS. 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 77) 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 75) 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:** input 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. A *HA Enabled* icon is added after the *Alert* icon on the active 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.

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 options available in `System → Failover` are shown in [Figure 4.23](#): and described in [Table 4.11](#).

The screenshot shows the 'Failover' configuration page under the 'System' menu. At the top, there's a navigation bar with tabs: Information, General, Boot, Advanced, Email, System Dataset, Tunables, Update, Cloud Credentials, CAs, Support, Proactive Support, Certificates, and Failover (which is active). Below the navigation bar, the 'Failover' section contains three settings: 'Disabled' with an unchecked checkbox, 'Master' with a checked checkbox, and 'Timeout' with a text input field containing '0'. At the bottom of the settings are three buttons: 'Save', 'Sync To Peer', and 'Sync From Peer'.

Fig. 4.23: Example Failover Screen

Table 4.11: Failover Options

Setting	Value	Description
Disabled	checkbox	when checked, administratively disable failover which changes the <i>HA Enabled</i> icon to <i>HA Disabled</i> and activates the <i>Master</i> field; an error message is generated if the standby node is not responding or failover has not been configured
Master	checkbox	grayed out unless <i>Disabled</i> is checked; in that case, this box is automatically checked on the master system, allowing the master to automatically take over when the <i>Disabled</i> 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	force configuration sync from the active node to the standby node; the standby node must be rebooted after the sync; the HA daemon does this automatically, do not use this unless requested by an iX support engineer
Sync From Peer	button	force configuration sync from the standby node to the active node; the HA daemon does this automatically, do not use this unless requested by an iX support engineer

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.

4.14.1 Failover Management

The `hact1` command line utility is included for managing existing failovers. Once a failover has been configured, it is recommended to use `hact1` instead of the GUI as any changes made using `System → Failover` will restart networking.

When this command is given without options, it indicates the failover status. This example was run on an active node:

```
hactl
Node status: Active
Failover status: Enabled
```

And this example is from a system that has not been configured for failover:

```
hactl
Node status: Not an HA node
```

Table 4.12 summarizes the options for this command.

Table 4.12: hactl Options

Option	Description
enable	administratively enable failover
disable	administratively disable failover
status	node type indicator: active, passive, or non-HA
takeover	can only be run from the passive node; gives a warning message that the current active node will reboot
giveback	cannot be run from the active node; gives a warning message that this node will reboot
-h or --help	show the help message (options) for this command
-q	prevent status display if this is a non-HA node

TASKS

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

- *Cloud Sync* (page 57) schedules data synchronization to cloud providers
- *Cron Jobs* (page 62) schedules a command or script to automatically execute at a specified time
- *Init/Shutdown Scripts* (page 64) configures a command or script to automatically execute during system startup or shutdown
- *Rsync Tasks* (page 65) schedules data synchronization to another system
- *S.M.A.R.T. Tests* (page 71) schedules disk tests

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

Note: By default, *Scrubs* (page 125) are run once a month by an automatically-created task. *S.M.A.R.T. Tests* (page 71) and *Periodic Snapshot Tasks* (page 112) 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 Raync Tasks S.M.A.R.T. Tests										
Add Cloud Sync										
Description	Direction	Path	Status	Minute	Hour	Day of month	Month	Day of week	Credential	Enabled
backup-acctg	PUSH	/mnt/volume1 /smb-storage /accounting-backup.bin	SUCCESS	00	Every hour	Everyday	Every month	Everyday	S3 Storage	true
<div> <div>Edit</div> <div>Delete</div> <div>Run Now</div> </div>										

Fig. 5.1: Cloud Sync Status

Cloud Credentials (page 44) 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

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

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	<i>Push</i> to send data to cloud storage, or <i>Pull</i> 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 44)
Path	browse button	select the directories or files to be sent for <i>Push</i> syncs or the destinations for <i>Pull</i> syncs
Minute	slider or minute selections	select <i>Every N minutes</i> and use the slider to choose a value, or select <i>Each selected minute</i> and choose specific minutes
Hour	slider or hour selections	select <i>Every N hours</i> and use the slider to choose a value, or select <i>Each selected hour</i> and choose specific hours
Days of month	slider or day of month selections	select <i>Every N days of month</i> and use the slider to choose a value, or select <i>Each selected day of month</i> and choose specific days
Months	checkboxes	months when the Cloud Sync runs
Days 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 44).

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.

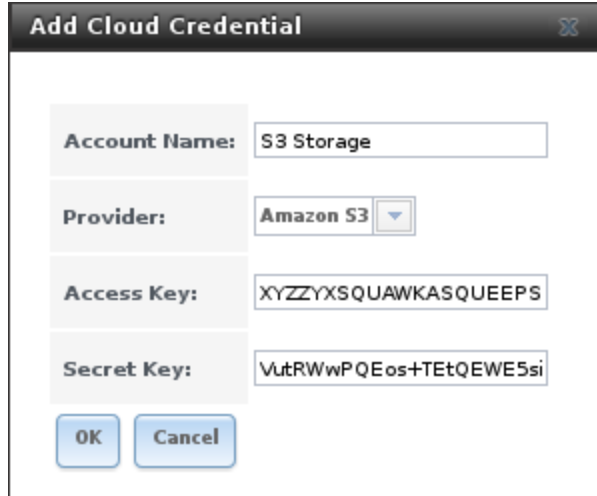
Finally, the *Enabled* field makes it possible temporarily disable a cloud sync job without deleting it.

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](#):

A screenshot of a dialog box titled "Add Cloud Credential". It contains four input fields: "Account Name" with the value "S3 Storage", "Provider" with a dropdown menu showing "Amazon S3", "Access Key" with the value "XYZZYXSQUAWKASQUEEPS", and "Secret Key" with the value "VutRWwPQEos+TetQEWE5si". At the bottom are "OK" and "Cancel" buttons.

Account Name:	S3 Storage
Provider:	Amazon S3
Access Key:	XYZZYXSQUAWKASQUEEPS
Secret Key:	VutRWwPQEos+TetQEWE5si

OK 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/s

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) (<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**.

62

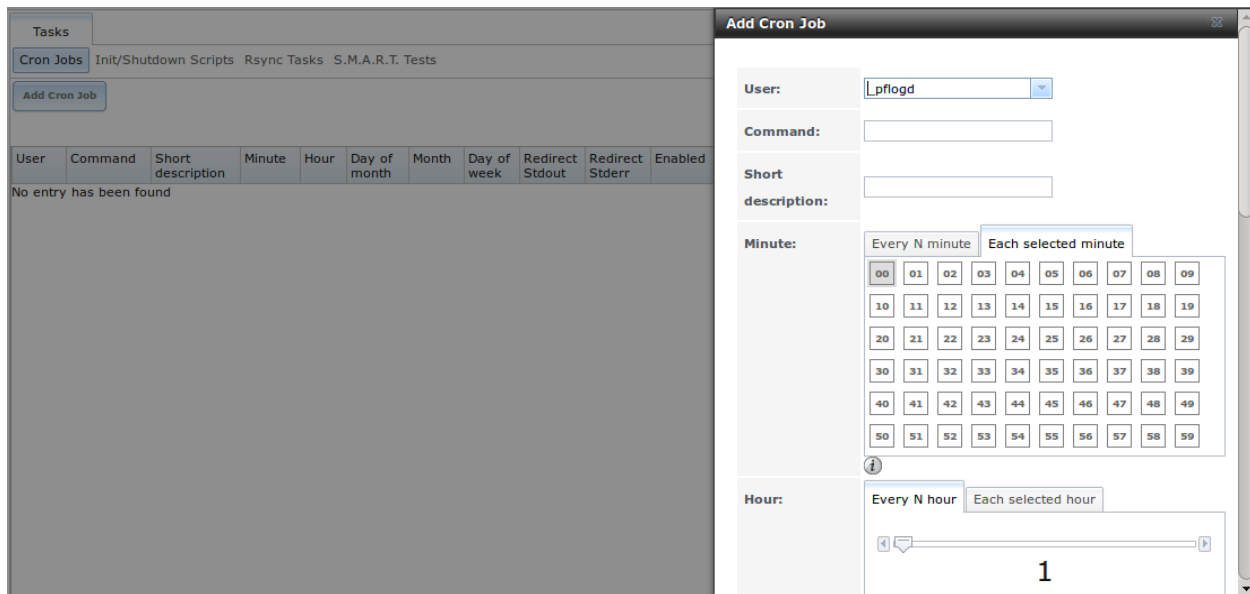


Fig. 5.5: Creating a Cron Job

Table 5.2 summarizes the configurable options when creating 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 <i>root</i> user account
Redirect Stderr	checkbox	disables emailing errors to the <i>root</i> user account
Enabled	checkbox	uncheck disable the cron job without deleting it

Created cron jobs will be listed 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.

When scheduling a command, make sure that the command is in the path or give the full path to the command. One way to test the path is to type **which command_name**. If the command is not found, it is not in your path.

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

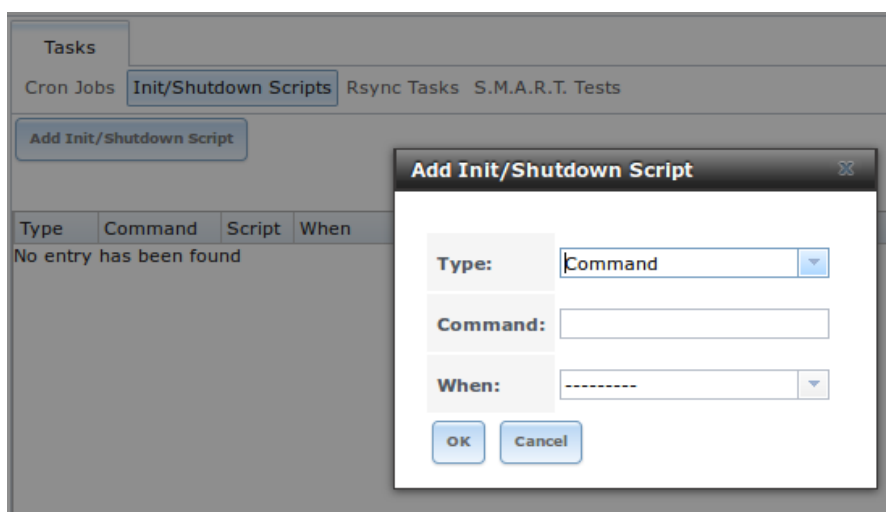


Fig. 5.6: Add an Init/Shutdown Script

Table 5.3: Options When Adding an Init/Shutdown Script

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

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 114) 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 237).

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) (<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.

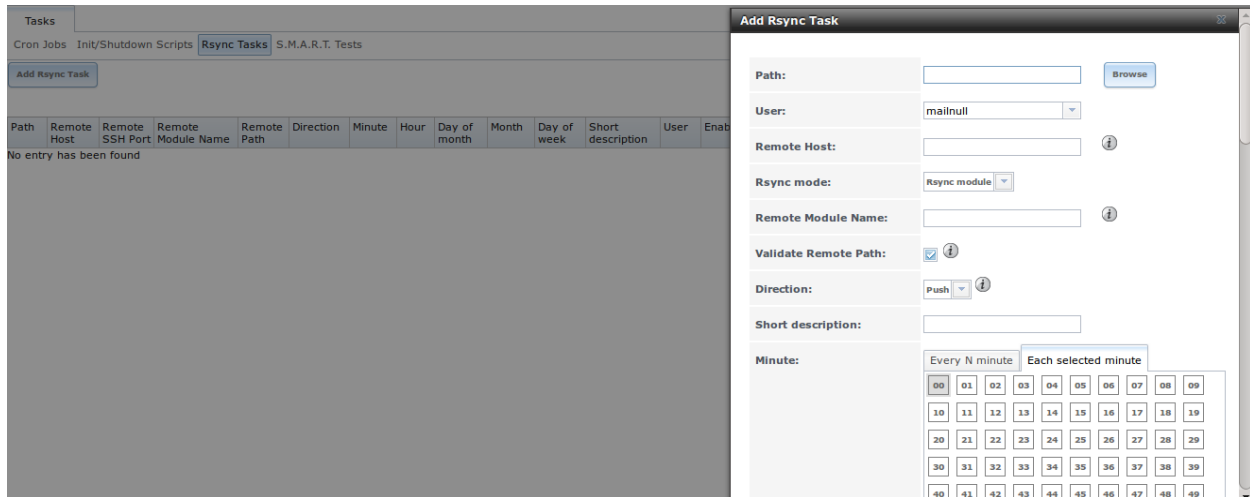


Fig. 5.7: Adding an Rsync Task

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 <i>username@remote_host</i> if the username differs on the remote host
Remote SSH Port	integer	only available in <i>Rsync over SSH</i> mode; allows specifying an SSH port other than the default of 22
Rsync mode	drop-down menu	choices are <i>Rsync module</i> or <i>Rsync over SSH</i>
Remote Module Name	string	only appears when using <i>Rsync module</i> mode, at least one module must be defined in rsyncd.conf(5) (http://www.samba.org/ftp/rsync/rsyncd.conf.html) of rsync server or in the <i>Rsync Modules</i> of another system
Remote Path	string	only appears when using <i>Rsync over SSH</i> mode, enter the existing path on the remote host to sync with (e.g. <i>/mnt/volume</i>); note that maximum path length is 255 characters
Validate Remote Path	checkbox	if the <i>Remote Path</i> does not yet exist, check this box to have it automatically created
Direction	drop-down menu	choices are <i>Push</i> or <i>Pull</i> ; 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

Continued on next page

Table 5.4 – continued from previous page

Setting	Value	Description
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>-r1ptgoD</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 <i>root</i>
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	rsync(1) (http://rsync.samba.org/ftp/rsync/rsync.html) options not covered by the GUI; if the * character is used, it must be escaped with a backslash (<code>*.txt</code>) or used inside single quotes (<code>'*.txt'</code>)
Enabled	checkbox	uncheck to disable the rsync task without deleting it; note that when the <i>Rsync</i> (page 206) 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, input `-password-file=/PATHTO/FILENAME` in the *Extra options* box, replacing `/PATHTO/FILENAME` with the appropriate path to the file containing the value of 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/*.

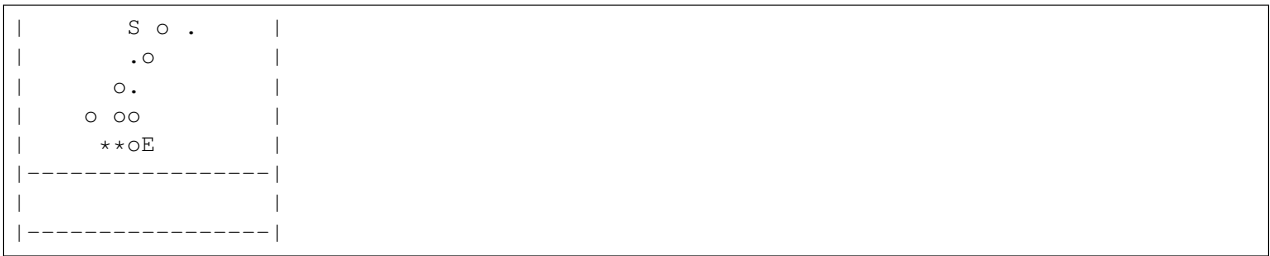
5.4.2 Rsync over SSH Mode

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 237) 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      |
```



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 AAAAB3NzaC1yc2EAAAADAQABAAQAC11BEXRgw1W8y8k+1XP1VR3xsmVSjtsoyIzV/PlQPoSrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4hdcD7Y5mvU3MAEeDClt02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmasti00qmDDcp/kOxT+S6DFNDBY6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+l9RaEXMRuTyQgqJB/rsRcmJX5fApdDmNfwrRSxLjDvUzfywnjFHlKk/+TQITlgg1QQaj21PJD9pnDVF0AiJrWyWnR 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.

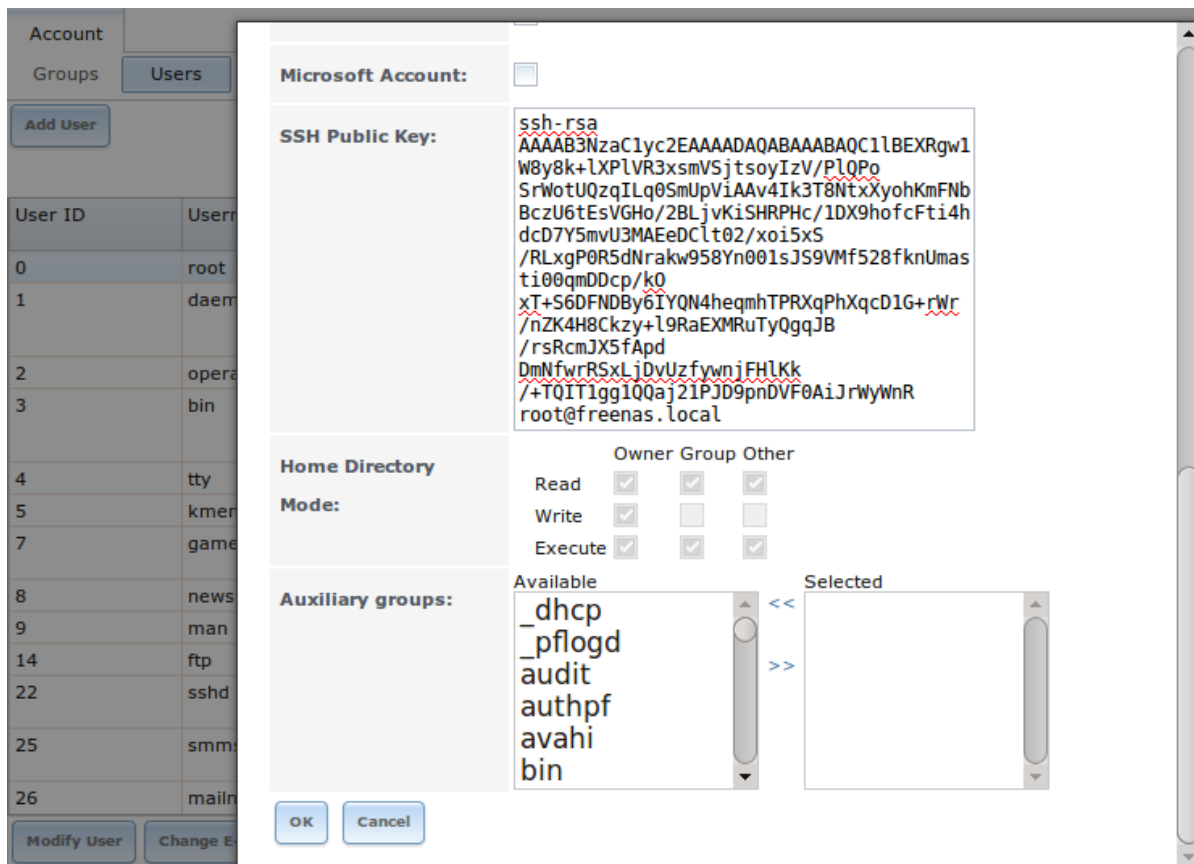


Fig. 5.8: Pasting the User's SSH Public Key

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.

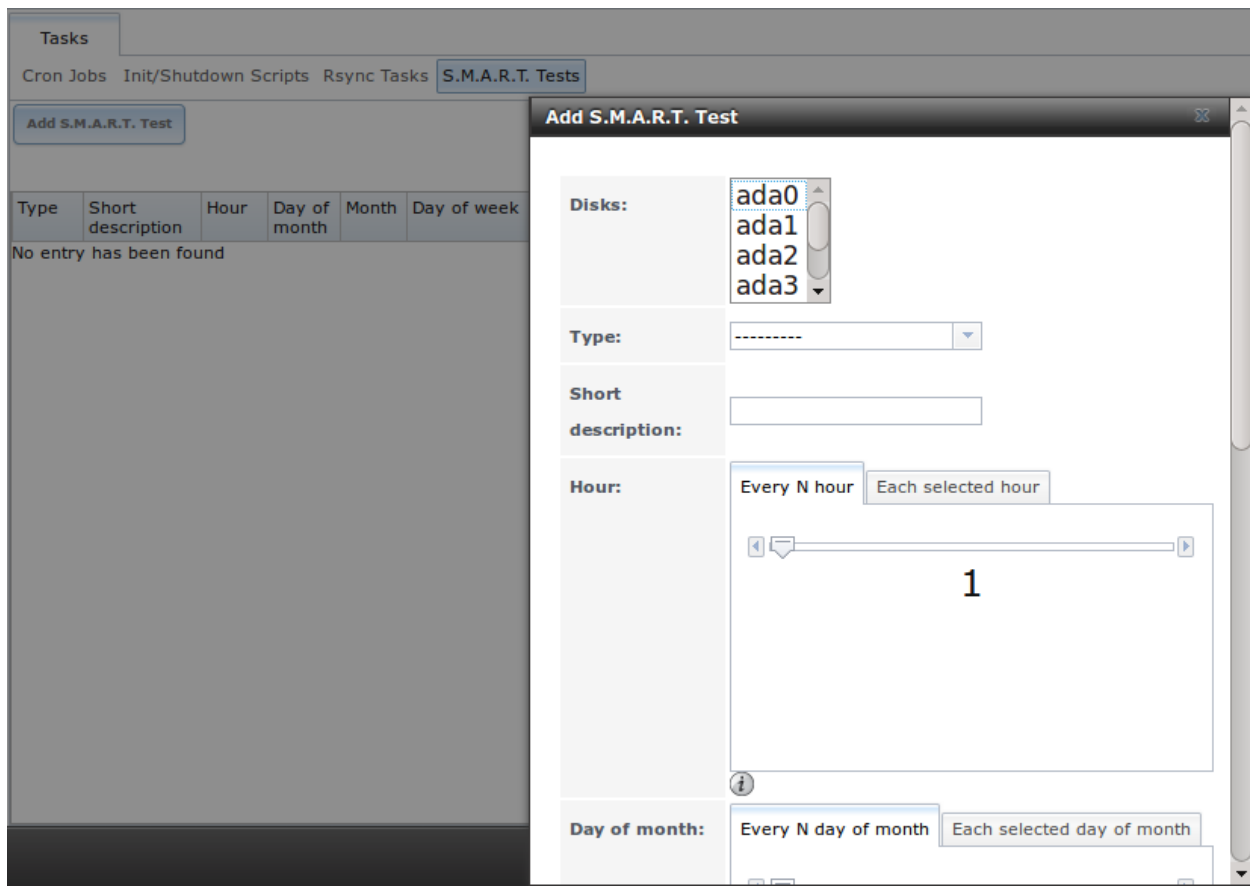


Fig. 5.9: Adding a S.M.A.R.T. Test

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 237).

The results of a test can be checked from *Shell* (page 237) 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 74): general network settings.
- *Interfaces* (page 75): settings for each network interface.
- *IPMI* (page 77): settings controlling connection to the appliance through the hardware side-band management interface if the graphical user interface becomes unavailable.
- *Link Aggregations* (page 79): settings for network link aggregation and link failover.
- *Network Summary* (page 84): display an overview of the current network settings.
- *Static Routes* (page 84): add static routes.
- *VLANs* (page 85): 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.

The screenshot shows the 'Global Configuration' tab selected under the 'Network' section. The following fields are visible:

Hostname:	freenas
Domain:	local
IPv4 Default Gateway:	
IPv6 Default Gateway:	
Nameserver 1:	
Nameserver 2:	
Nameserver 3:	
HTTP Proxy:	

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	string	system host name
Domain	string	system domain name
IPv4 Default Gateway	IP address	typically not set (see NOTE below)
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 <i>http://my.proxy.server:3128</i> or <i>http://user@password:my.proxy.server:3128</i>
Enable netwait feature	checkbox	if enabled, network services are not started at boot until the interface is able to ping the addresses listed in <i>Netwait IP list</i>
Netwait IP list	string	if <i>Enable netwait feature</i> 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 <i>/etc/hosts</i> ; use the format <i>IP_address space hostname</i> where multiple hostnames can be used if separated by a space

When 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 84) 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 but-

ton 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 <i>DHCP</i> is unchecked
IPv4 Netmask	drop-down menu	enter a netmask if <i>DHCP</i> 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

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 53) 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.

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
Network Summary
Link Aggregations
Static Routes
VLANs

Channel:

1

Password:

Password confirmation:

i

DHCP:

☒

IPv4 Address:

10.5.65.21

IPv4 Netmask:

/16 (255.255.0.0)

IPv4 Default Gateway:

10.5.0.1

VLAN ID:

OK

Cancel

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

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.

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, FEC, 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`.

FEC: supports Cisco EtherChannel on older Cisco switches. This is a static setup and does not negotiate aggregation with the peer or exchange frames to monitor the link.

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, as 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. This allows an iSCSI initiator to recognize multiple links to a target, utilizing 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**.

Note: 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`.

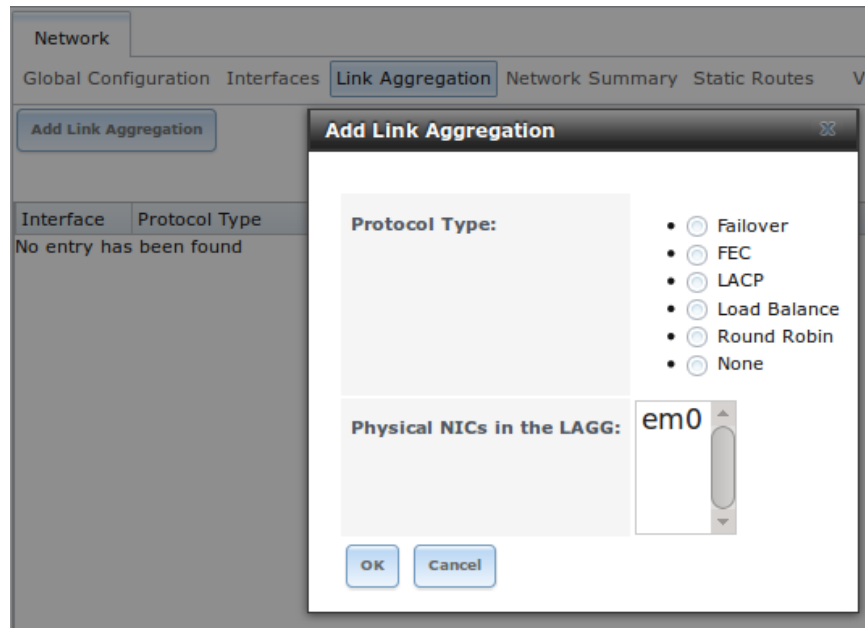


Fig. 6.4: Creating a lagg Interface

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.

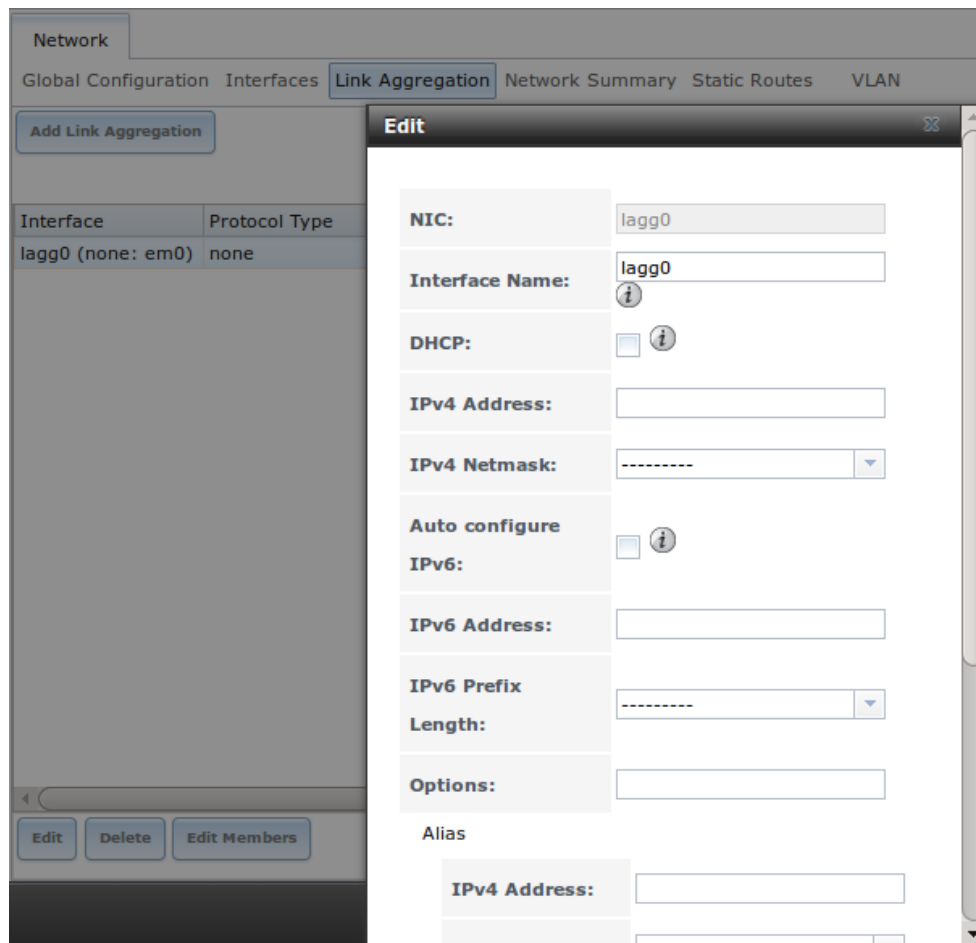


Fig. 6.5: Editing a lagg

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 <i>DHCP</i> is left unchecked
IPv4 Netmask	drop-down menu	enter a netmask if <i>DHCP</i> is left unchecked
Auto configure IPv6	checkbox	check only if DHCP server available to provide IPv6 address info
IPv6 Address	string	optional
IPv6 Prefix Length	drop-down menu	required if an IPv6 address is entered
Options	string	additional ifconfig(8) (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.

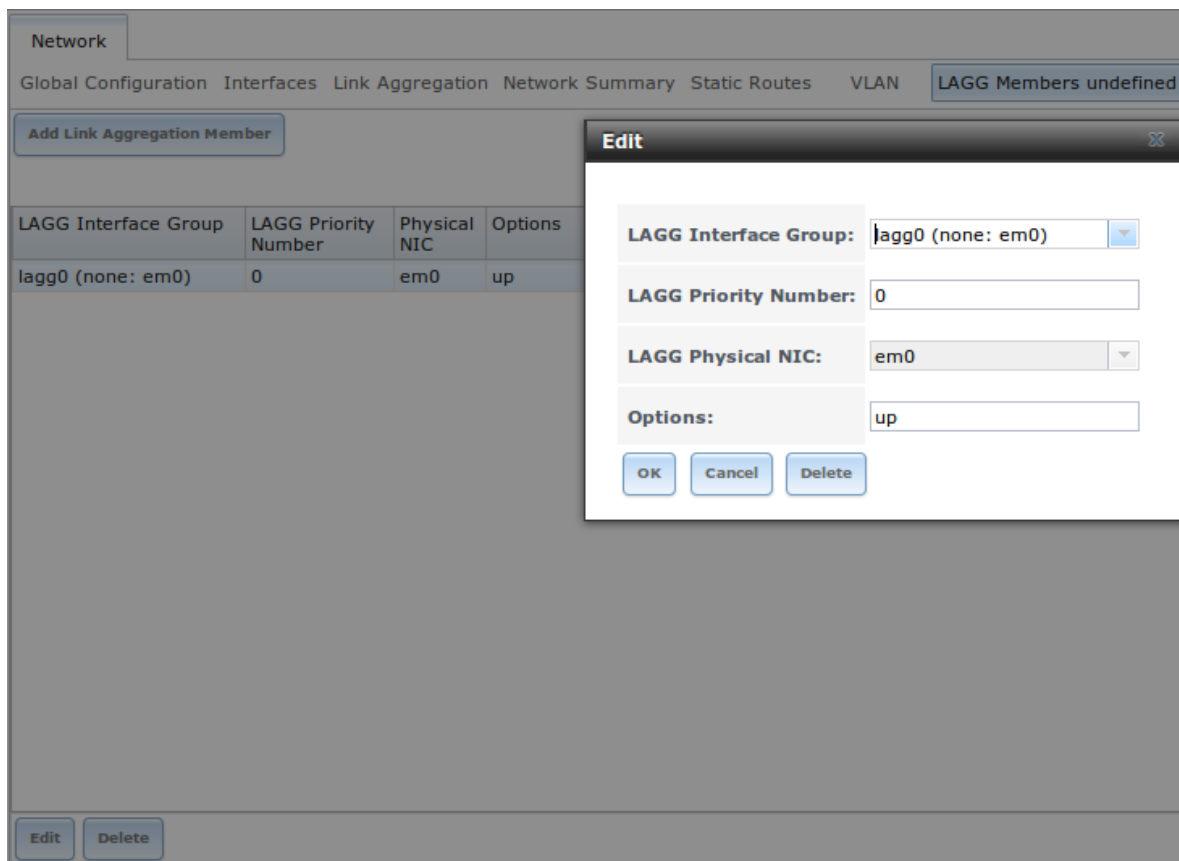


Fig. 6.6: Editing a Member Interface

Table 6.5: Configuring a Member Interface

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

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 237):

```
systat -ifstat
```

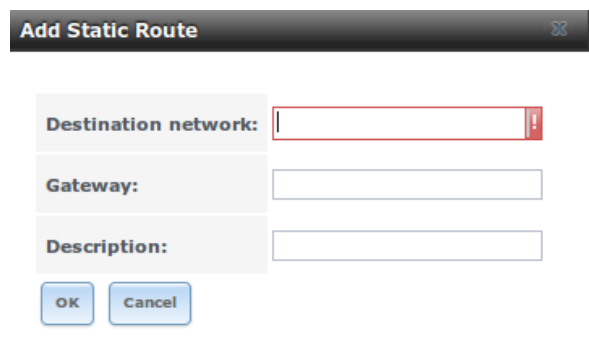
More information about this command can be found at [systat\(1\)](http://www.freebsd.org/cgi/man.cgi?query=systat) (<http://www.freebsd.org/cgi/man.cgi?query=systat>).

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](#).



The screenshot shows a dialog box titled "Add Static Route". It has three input fields: "Destination network:" (with a red border), "Gateway:", and "Description:". Below these fields are two buttons: "OK" and "Cancel".

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 <i>A.B.C.D/E</i> where <i>E</i> is the CIDR mask
Gateway	integer	enter the IP address of the gateway
Description	string	optional

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](#).

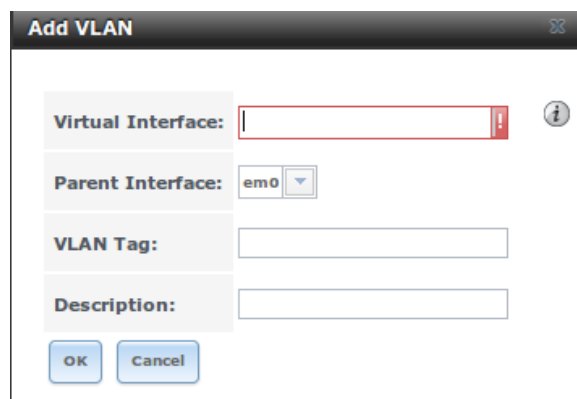


Fig. 6.8: Adding a VLAN

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

Table 6.7: Adding a VLAN

Setting	Value	Description
Virtual Inter- face	string	use the format <i>vlanX</i> where <i>X</i> is a number representing a vlan interface not currently being used as a parent
Parent Inter- face	drop-down menu	usually an Ethernet card connected to a properly configured switch port; note that newly created Link Aggregations (page 79) 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
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 53) 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 87) creates and manages storage volumes.
- [Periodic Snapshot Tasks](#) (page 112) schedules automatic creation of filesystem snapshots.
- [Replication Tasks](#) (page 114) automates the replication of snapshots to a remote system.
- [Scrubs](#) (page 125) schedules scrubs as part of ongoing disk maintenance.
- [Snapshots](#) (page 126) manages local snapshots.
- [VMware-Snapshot](#) (page 128) coordinates ZFS snapshots with a VMware datastore.

Note: If the TrueNAS® system has been configured as the passive node in a failover configuration, the screens shown in this chapter will be replaced by a message indicating that this node is passive. 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 ZFS pools, import a disk to copy its data into an existing pool, or import an existing ZFS pool. It can also be used to create ZFS datasets and zvols and to manage their permissions.

Note: In ZFS terminology, the storage that is managed by ZFS is 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. All of the features are described to help make it clear which will be the most benefit for your uses, and caveats or caveats or hardware restrictions which could limit their use.**

7.1.1 Volume Manager

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 242) for information on ZFS redundancy with multiple disks before using *Volume Manager*.

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

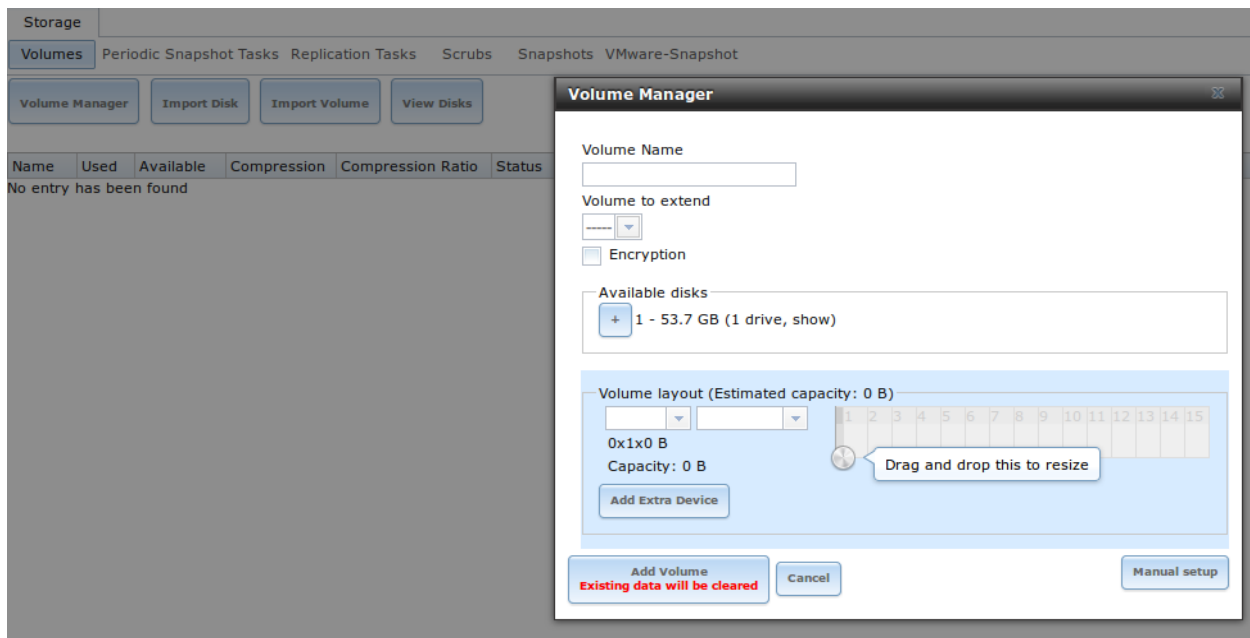


Fig. 7.1: Creating a ZFS Pool Using Volume Manager

Table 7.1 summarizes the configuration options of this screen.

Table 7.1: Options When Creating a ZFS Volume

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); it is recommended to choose a name that will stick out in the logs (e.g. not <code>data</code> or <code>freenas</code>)
Volume to extend	drop-down menu	used to extend an existing ZFS pool; see Extending a ZFS Volume (page 91) for instructions
Encryption	checkbox	read the section on Encryption (page 89) before choosing to use encryption
Available disks	display	displays the number and size of available disks; hover over <i>show</i> 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; once at least one disk is selected, the layouts supported by the selected number of disks will be added to the drop-down menu
Add Extra Device	button	used to configure multiple vdevs or to add log or cache devices during pool creation
Manual setup	button	used to make a non-optimal pool (not recommended); see Manual Setup (page 90) for details

Drag the slider to select the desired number of disks. *Volume Manager* displays the resulting storage capacity, including taking swap space into account. To change the layout or the number of disks, use the mouse to 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 in this situation by using the *Manual setup* button and following the instructions in *Manual Setup* (page 90).

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 242).

The *Add Volume* button warns that **existing data will be cleared**. In other words, creating a new volume reformats the selected disks. If the existing data is meant to be preserved, click the *Cancel* button and refer to *Import Disk* (page 98) and *Import Volume* (page 99) 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 will refresh and the new volume is listed in the tree under *Storage* → *Volumes*. Click the + next to the volume name to access its *Change Permissions* (page 92), *Create Dataset* (page 94), and *Create zvol* (page 97) options.

Encryption

TrueNAS® supports **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:

- This is **not** the encryption method used by Oracle's version of ZFS. That version is not open source and is the property of Oracle.
- This is full disk encryption and **not** per-filesystem encryption. The underlying drives are first encrypted, then the pool is created on top of the encrypted devices.
- This type of encryption is primarily targeted at users who store sensitive data and want to retain the ability to remove disks from the pool without having to first wipe the disk's contents.
- This 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.
- On the other hand, if the key is lost, the data on the disks is inaccessible. Always back up the key!
- The encryption key is per ZFS volume (pool). Multiple pools each have their own encryption key.

-
- Data in the ARC cache and the contents of RAM are unencrypted.
 - Swap is always encrypted, even on unencrypted volumes.
 - There is no way to convert 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. In other words, newly created vdevs must match the existing encryption scheme. When extending a volume, Volume Manager automatically encrypts the new vdev being added to the existing encrypted pool.

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.

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

Manual Setup

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

Note: When using disks of differing sizes, the volume is limited by the size of the smallest disk. For this reason, it is recommended to instead use *Volume Manager* with same-size disks.

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

Manual Setup

Volume name:

Encryption: ☐

Member disks (0):

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

Deduplication:

ZFS Extra:

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

Existing data will be cleared

Fig. 7.2: Creating a Non-Optimal ZFS Volume

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) ; it is recommended to choose a name that will stick out in the logs (e.g. not <code>data</code> or <code>freenas</code>)
Encryption	checkbox	read the section on Encryption (page 89) before choosing to use 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 your iXsystems support engineer
ZFS Extra	bullet selection	used to specify if disk is used for storage (<i>None</i>), 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](#), can be used to add additional disks to an existing ZFS volume. This drop-down menu will be empty if no ZFS volume exists.

Note: If the existing volume is encrypted, a warning message will remind you that the operation of extending a volume will reset the passphrase and recovery key. After extending the volume, you should 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 ZFS pool.

When adding disks to increase the capacity of a volume, ZFS supports the addition of virtual devices, known as 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, you can stripe a new vdev (and its disks) with another of the **same type of existing vdev** to increase the overall size of ZFS the pool. In other words, when you extend a ZFS volume, you are really striping similar vdevs. Here are some examples:

- to extend a ZFS stripe, add one or more disks. Since there is no redundancy, you do not have to add the same amount of disks 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 you have 10 drives, you could start by creating a mirror of two drives, extending this mirror by creating another mirror of two drives, and repeating three more times until all 10 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 you try to add an incorrect number of disks to the existing vdev, an error message will appear, indicating the number of disks that are needed. You will need to select the correct number of disks in order to continue.

Adding L2ARC or ZIL Devices

Storage → Volumes → Volume Manager (see [Figure 7.1](#)) is also used to add L2ARC or ZIL SSDs to improve specific types of volume performance. This is described in more detail in the [ZFS Primer](#) (page 242).

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 144) 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 18) or imported from a directory service using the instructions in [Directory Services](#) (page 130). 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* → *View 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 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 ☐

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 <i>Owner (user)</i> , 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 <i>Owner (group)</i> , 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 <i>Mode</i> , see Note below
Mode	checkboxes	only applies to the <i>Unix</i> or <i>Mac</i> "Permission Type" so will be grayed out if <i>Windows</i> is selected
Permission Type	bullet selection	choices are <i>Unix</i> , <i>Mac</i> or <i>Windows</i> ; select the type which matches the type of client accessing the volume/dataset

Continued on next page

Table 7.3 – continued from previous page

Setting	Value	Description
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. A dataset is similar to a folder in that you can set permissions; it is also similar to a filesystem in that you can set properties such as quotas and compression as well as create snapshots.

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](#).

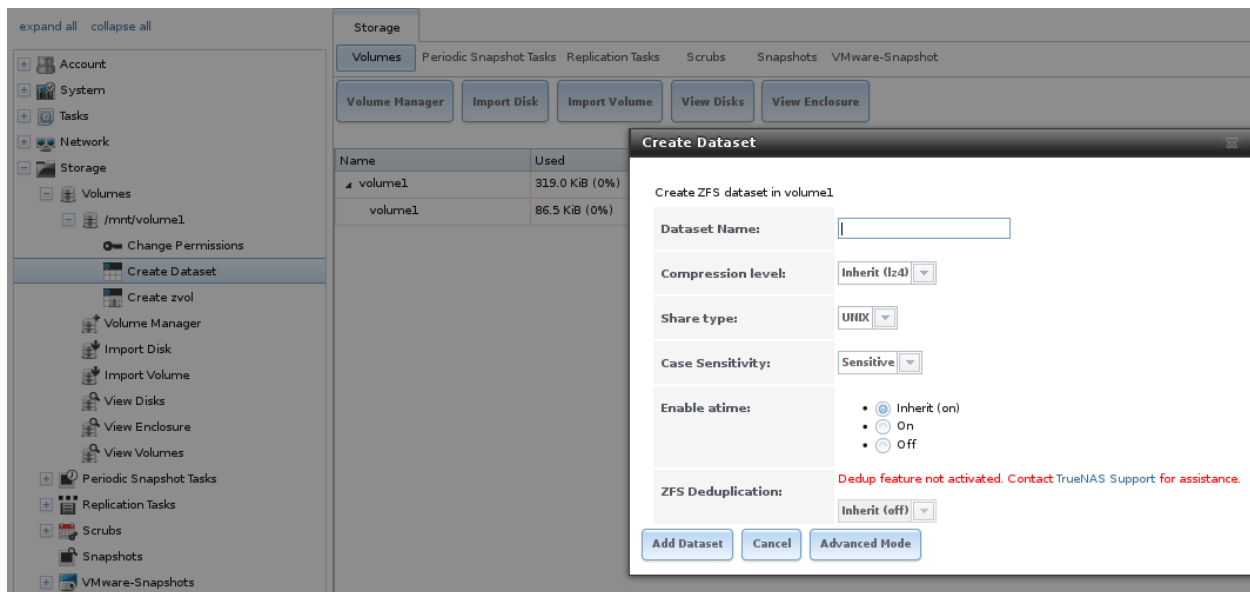


Fig. 7.4: Creating a ZFS Dataset

Table 7.4 summarizes the options available when creating a ZFS 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 these 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 its *Edit Options* button in *Storage* → *Volumes* → *View 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
Compression Level	drop-down menu	see the section on Compression (page 96) for a description of the available algorithms
Share type	drop-down menu	select the type of share that will be used on the dataset; choices are <i>UNIX</i> for an NFS share, <i>Windows</i> for a SMB share, or <i>Mac</i> for an AFP share
Case Sensitivity	drop-down menu	choices are <i>sensitive</i> (default, assumes filenames are case sensitive), <i>insensitive</i> (assumes filenames are not case sensitive), or <i>mixed</i> (understands both types of filenames)
Enable atime	Inherit, On, or Off	controls whether the access time for files is updated when they are read; setting this property to <i>Off</i> avoids producing log traffic when reading files and can result in significant performance gains
Quota for this dataset	integer	only available in <i>Advanced Mode</i> ; default of <i>0</i> 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 <i>Advanced Mode</i> ; a specified value applies to both this dataset and any child datasets

Continued on next page

Table 7.4 – continued from previous page

Setting	Value	Description
Reserved space for this dataset	integer	only available in <i>Advanced Mode</i> ; 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 <i>Advanced Mode</i> ; 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
Record Size	drop-down menu	only available in <i>Advanced Mode</i> ; 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

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 204) 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](#).

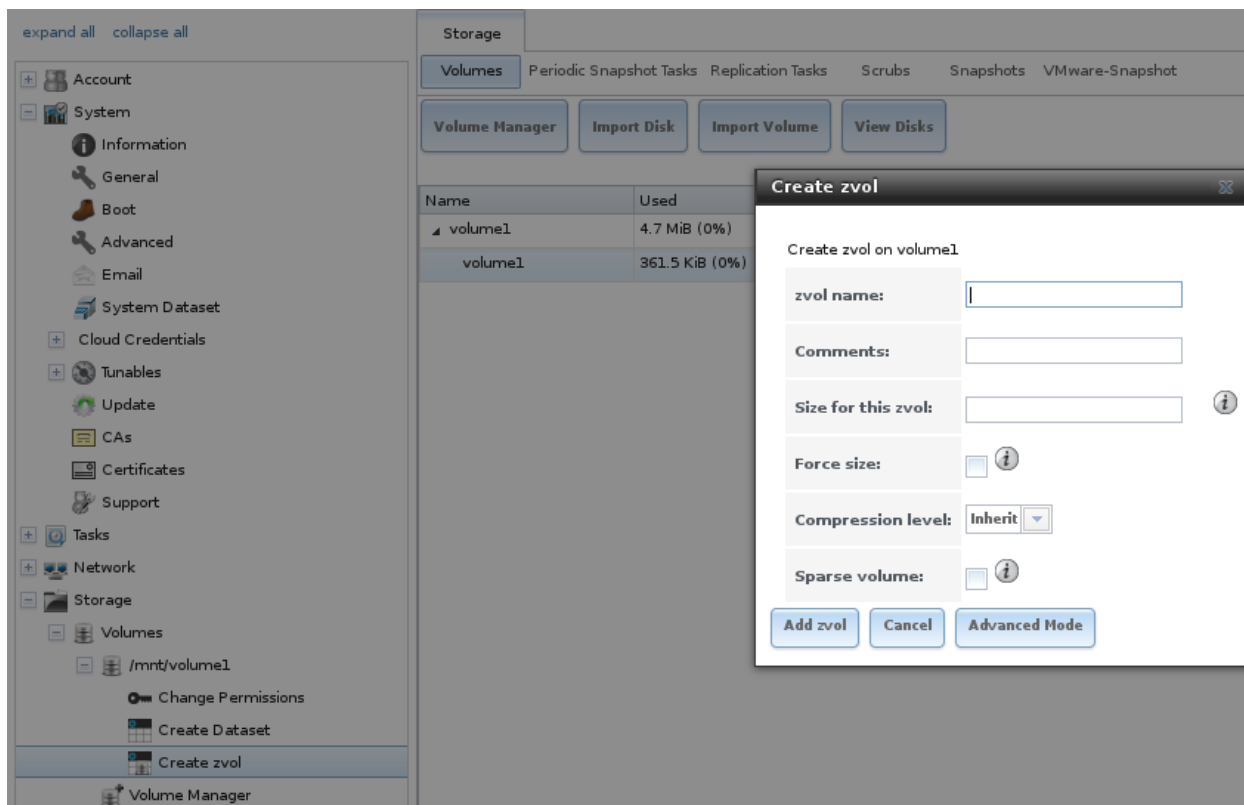


Fig. 7.5: Creating a zvol

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; input a name for the zvol
Comments	string	short comments or user notes about this zvol
Size for this zvol	integer	specify size and value such as <i>10Gib</i> ; if the size is more than 80% of the available capacity, the creation will fail with an “out of space” error unless <i>Force size</i> is checked
Force size	checkbox	by default, the system will not let you create a zvol if that operation will bring the pool to over 80% capacity; while NOT recommended , checking this box will force the creation of the zvol in this situation

Continued on next page

Table 7.5 – continued from previous page

Setting	Value	Description
Compression level	drop-down menu	see the section on Compression (page 96) 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 <i>Advanced Mode</i> 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.

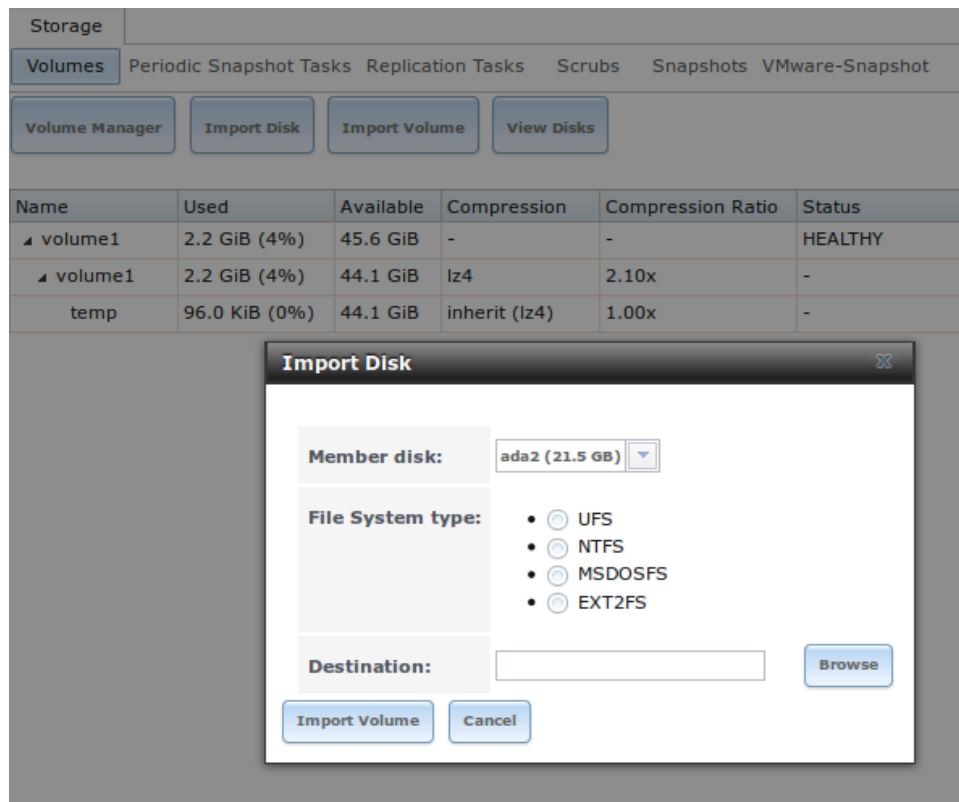


Fig. 7.6: Importing a Disk

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.

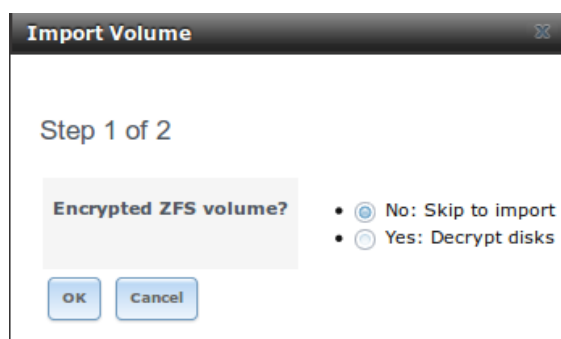


Fig. 7.7: Initial Import Volume Screen

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

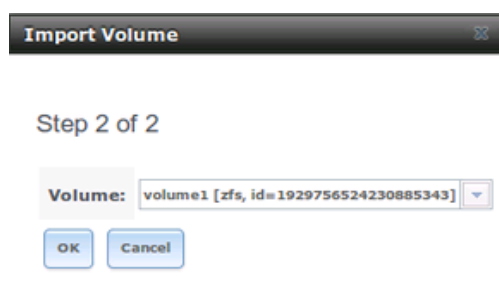


Fig. 7.8: Importing a Non-Encrypted Volume

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 237) 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

If you are importing an existing GELI-encrypted ZFS pool, you must decrypt the disks before importing the pool. In [Figure 7.7](#), select *Yes: Decrypt disks* to access the screen shown in [Figure 7.9](#).

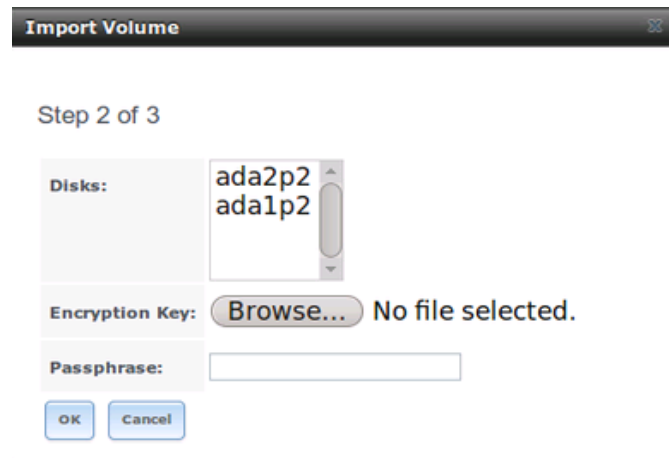


Fig. 7.9: Decrypting Disks Before Importing a ZFS Pool

Select the disks in the encrypted pool, browse to the location of the saved encryption key, input 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.

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

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 <i>Disabled</i> , can select a power management profile from the menu
Acoustic Level	drop-down menu	default is <i>Disabled</i> ; 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 71) for the disk
S.M.A.R.T. extra options	string	additional smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) options

Note: If a disk's serial number is not displayed in this screen, use the **smartctl** command from [Shell](#) (page 237). 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 View Enclosure

Click [Storage](#) → [Volumes](#) → [View Enclosure](#) to receive a status summary of the appliance's disks and hardware. An example is shown in [Figure 7.11](#).

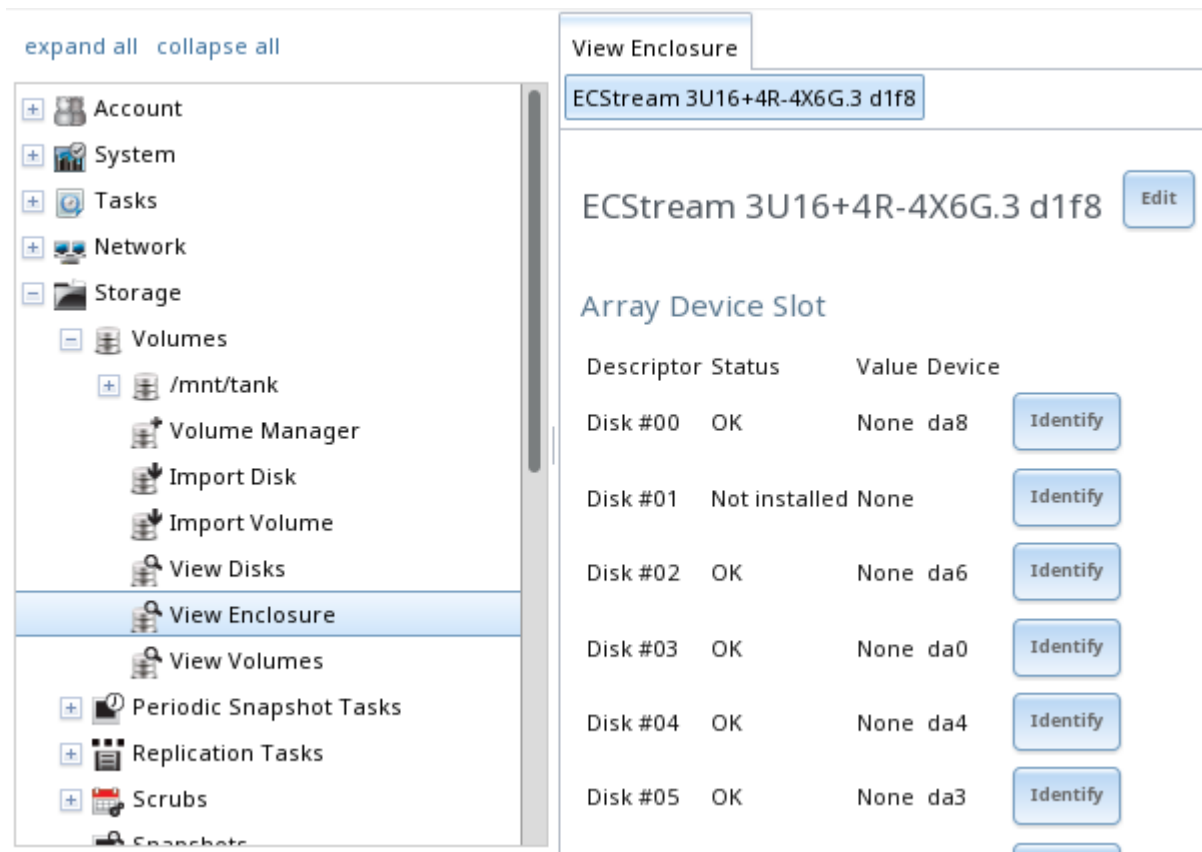


Fig. 7.11: View Enclosure

This screen is divided into the following sections:

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

Cooling: has an entry for each fan, its status, and its 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.

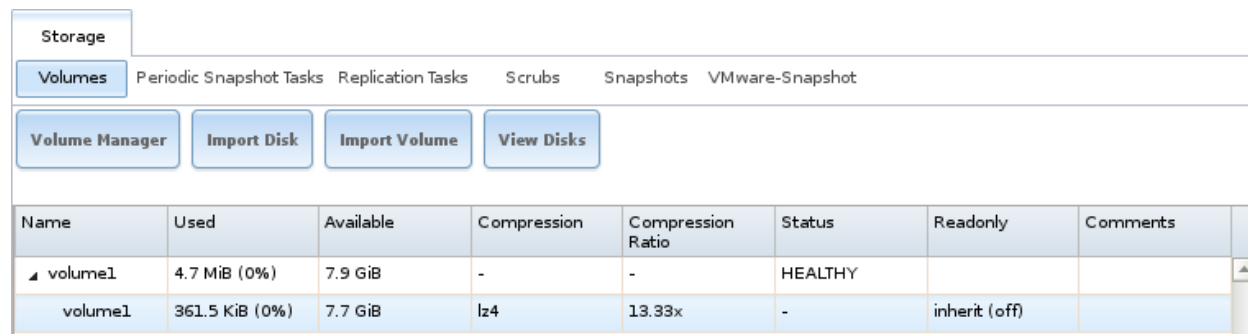
Voltage Sensor: shows the current voltage for each sensor, VCCP, and VCC.

7.1.9 View Volumes

Storage → Volumes → View Volumes is used to view and further configure existing ZFS pools, datasets, and zvols. The example shown in Figure 7.12 shows one ZFS pool (*volume1*) with two datasets (the one automatically created with the pool, *volume1*, and *dataset1*) and one zvol (*zvol1*).

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

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



Storage							
Volumes							
Periodic Snapshot Tasks Replication Tasks Scrubs Snapshots VMware-Snapshot							
Volume Manager Import Disk Import Volume View Disks							
Name	Used	Available	Compression	Compression Ratio	Status	Readonly	Comments
▲ volume1	4.7 MiB (0%)	7.9 GiB	-	-	HEALTHY		
volume1	361.5 KiB (0%)	7.7 GiB	lz4	13.33x	-	inherit (off)	

Fig. 7.12: Viewing Volumes

Clicking the entry for a pool causes several buttons to appear at the bottom of the screen. The buttons perform these actions:

Detach Volume: allows you to either export the pool or to delete the contents of the pool, depending upon the choice you make in the screen shown in Figure 7.13. The *Detach Volume* screen displays the current used space and indicates if there are any shares, provides checkboxes to *Mark the disks as new (destroy data)* and to *Also delete the share's configuration*, asks if you are sure that you want to do this, and the browser will turn red to alert you that you are about to do something that will make the data inaccessible. **If you do not check the box to mark the disks as new, the volume will be exported.** This means that the data is not destroyed and the volume can be re-imported at a later time. If you will be 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. **If you do check the box to mark the disks as new, the pool and all the data in its datasets, zvols, and shares will be destroyed and the underlying disks will be returned to their raw state.**

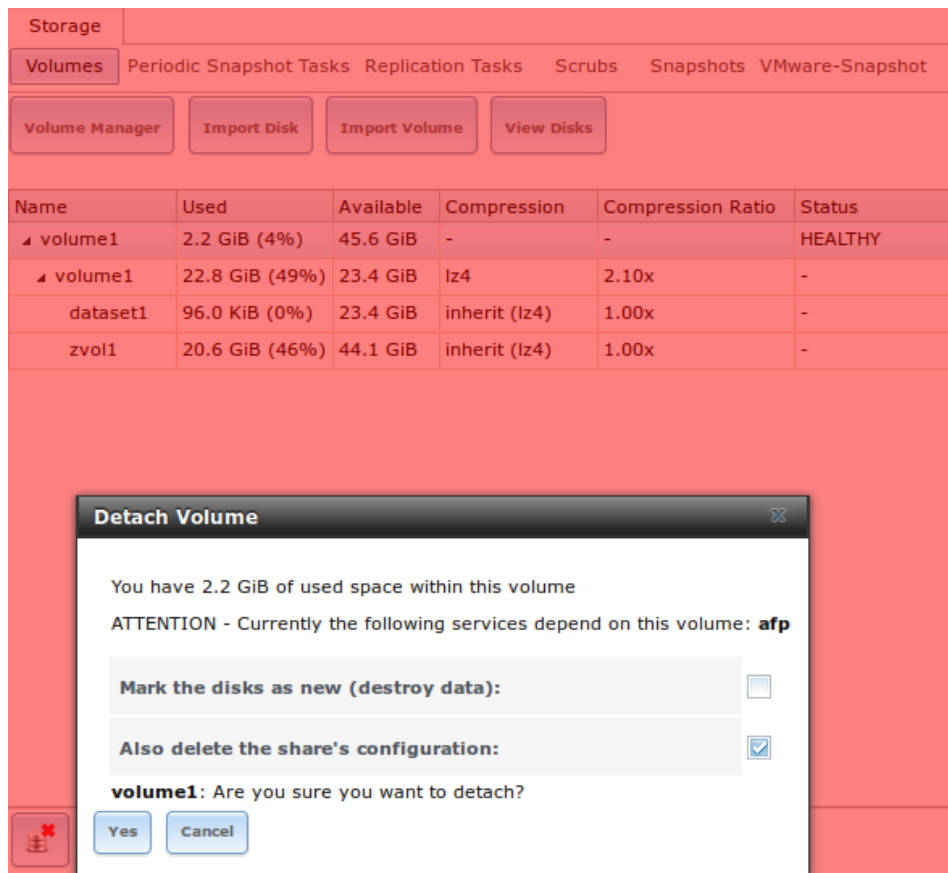


Fig. 7.13: Detach or Delete a Volume

Scrub Volume: scrubs and scheduling them are described in more detail in [Scrubs](#) (page 125). 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 7.14](#), 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 7.15](#)), offline or online the device, or replace the device (as described in [Replacing a Failed Drive](#) (page 109)).

Upgrade: used to upgrade the pool to the latest ZFS features, as described in [Upgrading a ZFS Pool](#) (page 43). 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. 7.14: Volume Status

Selecting a disk in *Volume Status* and clicking its *Edit Disk* button shows the screen in Figure 7.15. 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. 7.15: Editing a Disk

Clicking a dataset in *Storage* → *Volumes* → *View 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 92).

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use [Periodic Snapshot Tasks](#) (page 112).

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` → `View 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

If the *Encryption* box is checked during the creation of a pool, additional buttons appear in the entry for the volume in `Storage` → `Volumes` → `View Volumes`. An example is shown in Figure 7.16.

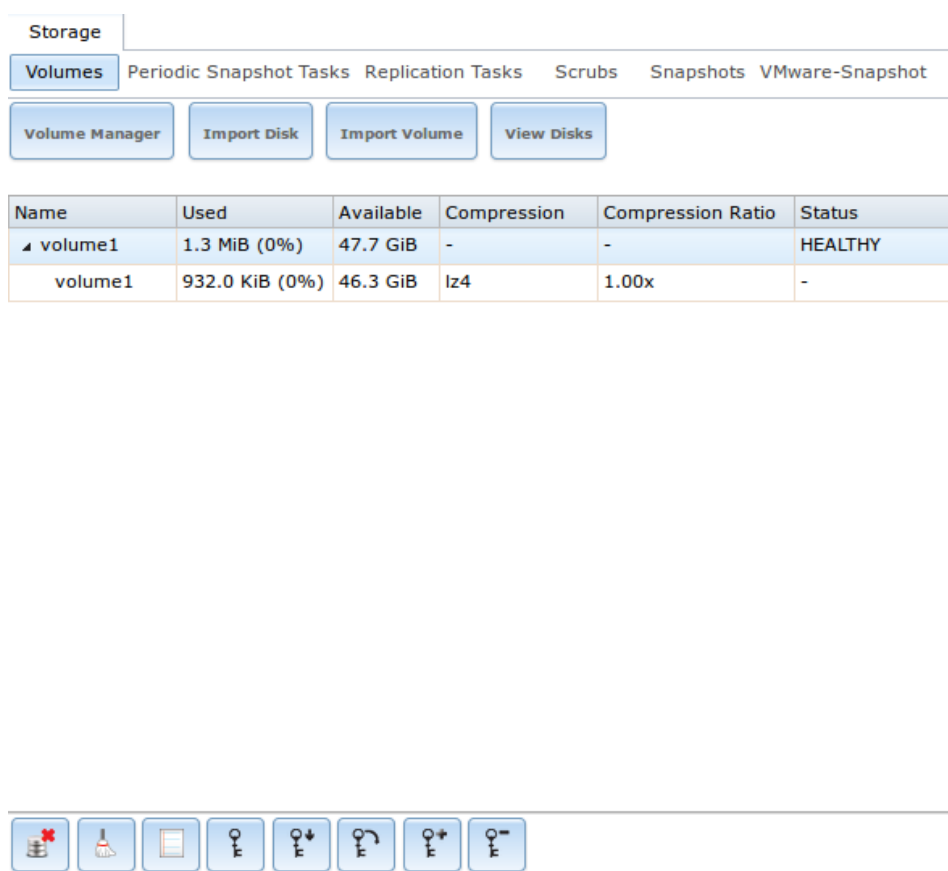


Fig. 7.16: Encryption Icons Associated with an Encrypted Volume

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.

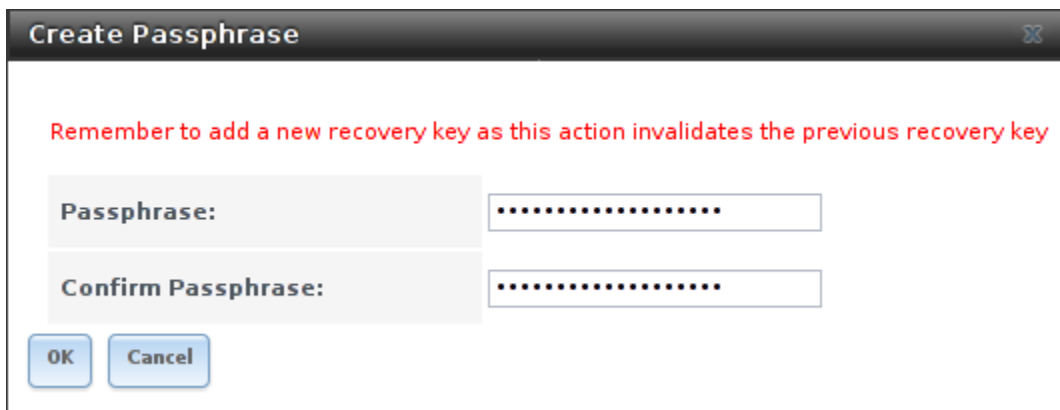


Fig. 7.17: Add or Change a Passphrase to an Encrypted Volume

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. 7.18: 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. 7.19: Unlock Button

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

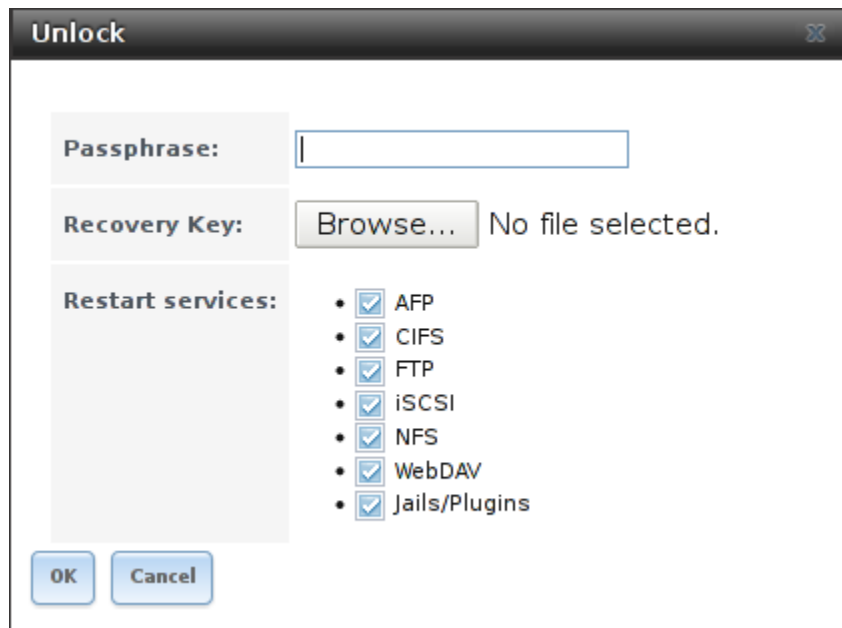


Fig. 7.20: Unlock Locked Volume

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 53) (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.

7.1.10 View Multipaths

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

7.1.11 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 your pool is encrypted with GELI, refer to [Replacing an Encrypted Drive](#) (page 111) before proceeding.

Before physically removing the failed device, go to *Storage → Volumes → View 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 offline and you can safely skip this step.

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 → View 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 starts to resilver and the status of the resilver is displayed.
3. After the drive replacement process is complete, re-add the replaced disk in the [S.M.A.R.T. Tests](#) (page 71) screen.

In the example shown in Figure 7.21, a failed disk is being replaced by disk *ada5* in the volume named *volume1*.

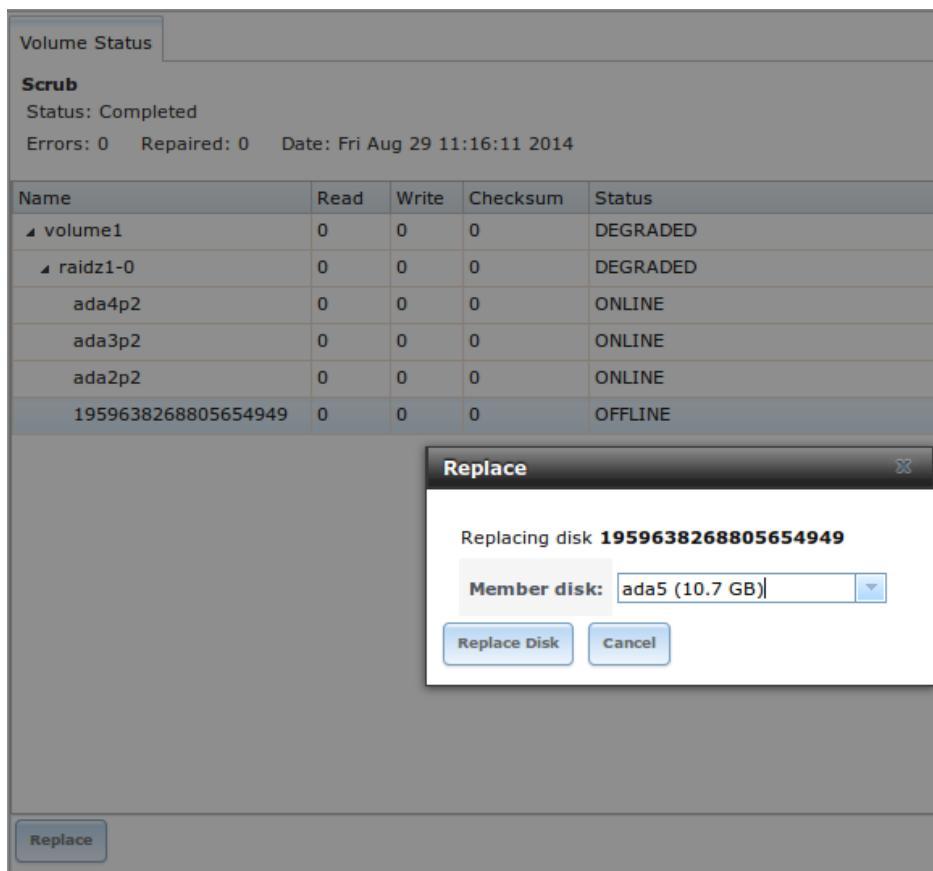


Fig. 7.21: Replacing a Failed Disk

After the resilver is complete, *Volume Status* shows a *Completed* resilver status and indicates any errors. Figure 7.22 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. 7.22: 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 89) **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 53) (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* → *View 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.

7.1.12 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 87) 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 109). 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.

7.2 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 114)). 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 87).

To create a periodic snapshot task, click *Storage* → *Periodic Snapshot Tasks* → *Add Periodic Snapshot* which opens the screen shown in *Figure 7.23*. *Table 7.7* summarizes the fields in this screen.

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

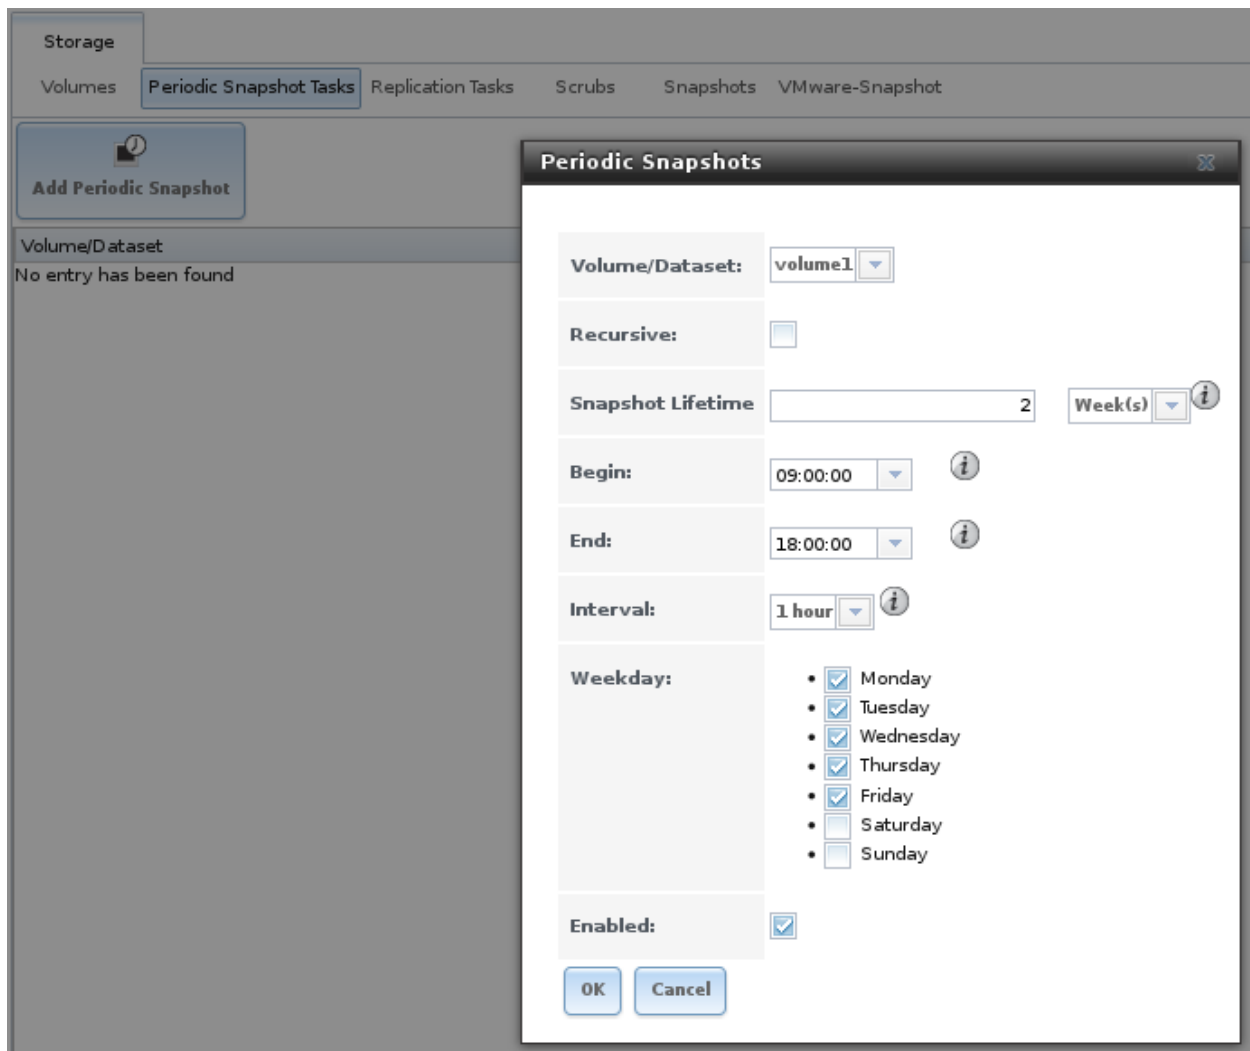


Fig. 7.23: Creating a Periodic Snapshot

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 Lifetime	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 <i>Begin</i> and <i>End</i> times
Weekday	checkboxes	which days of the week to take snapshots

Continued on next page

Table 7.7 – continued from previous page

Setting	Value	Description
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.

When the *OK* button is clicked, a snapshot is taken and the task will be repeated according to your settings.

After creating a periodic snapshot task, an entry for the snapshot task will be added to *View Periodic Snapshot Tasks*. Click an entry to access its *Edit* and *Delete* buttons.

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 112) 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.

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 87) named *alphavol* has already been created, and a *dataset* (page 94) 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* → *Volumes*. Click the *alphavol/alphadata* dataset to highlight it. Create a *periodic snapshot* (page 112) of it by clicking *Periodic Snapshot Tasks*, then *Add Periodic Snapshot* as shown in *Figure 7.24*.

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.

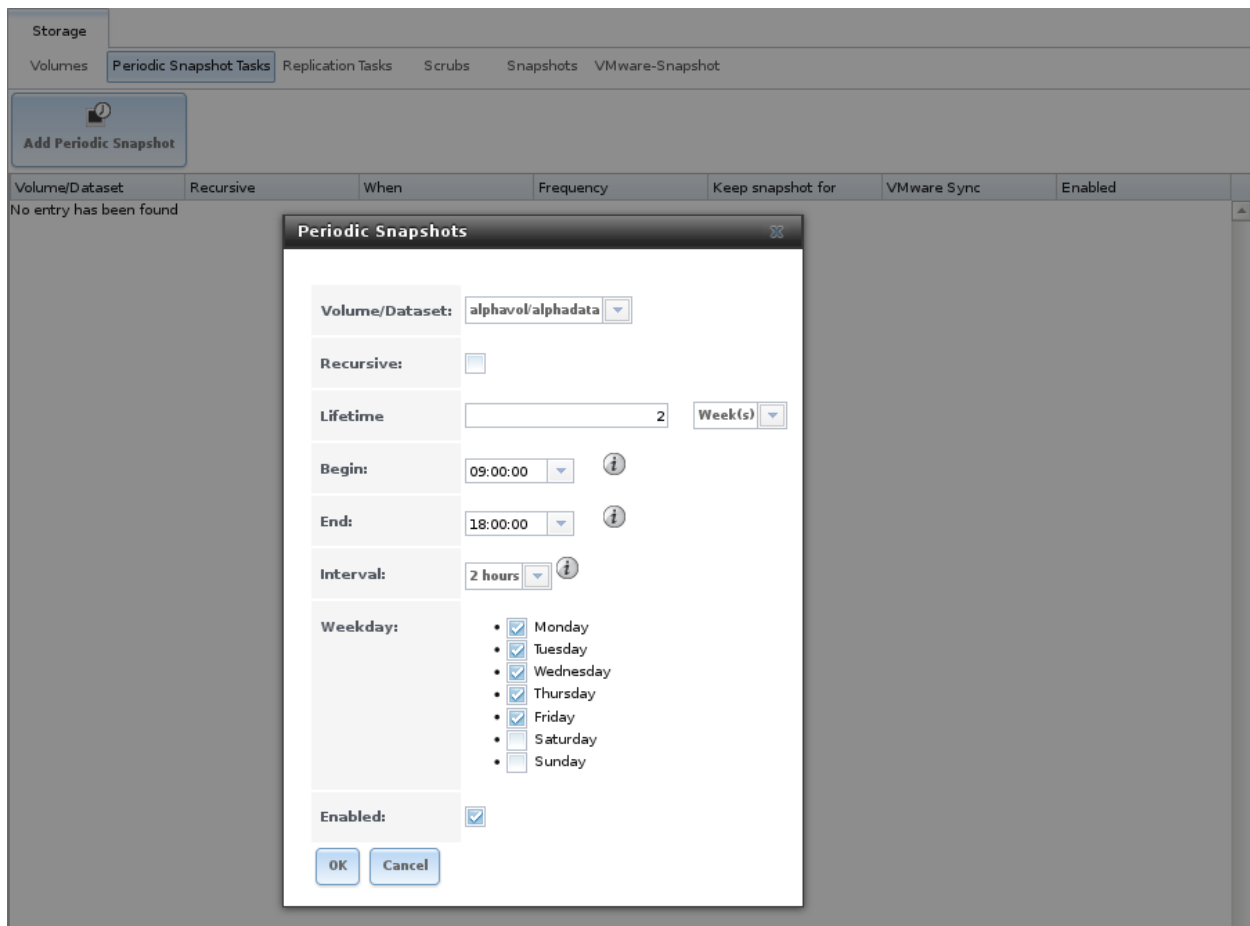


Fig. 7.24: Create a Periodic Snapshot for Replication

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 87) named *betavol* has already been created.

Snapshots are transferred with *SSH* (page 217). 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 7.25. The IP address of *Beta* is entered in the *Remote hostname* field. A hostname can be entered here if local DNS resolves for that hostname.

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

Add Replication

Volume/Dataset:

alphavol/alphadata

Remote ZFS Volume/Dataset:

betavol

Recursively replicate child dataset's snapshots:

☐

Delete stale snapshots on remote system:

☐

Replication Stream Compression:

lz4 (fastest)

Limit (kB/s):

0

Begin:

00:00:00

End:

23:59:00

Enabled:

☒

Setup mode:

Semi-automatic

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

Remote hostname:

10.0.0.118

Remote HTTP/HTTPS Port:

80

Remote HTTPS:

☐

Remote Auth Token:

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

Dedicated User Enabled:

☐

Dedicated User:

Encryption Cipher:

Standard

OK

Cancel

Fig. 7.25: 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 7.26](#).

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



Fig. 7.26: 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 7.27](#).

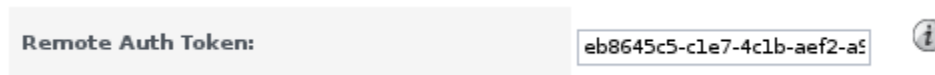


Fig. 7.27: Temporary Authentication Token Pasted to Source

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 122) 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® 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 7.28](#). Use the mouse to highlight the key data shown in the window, then copy it.

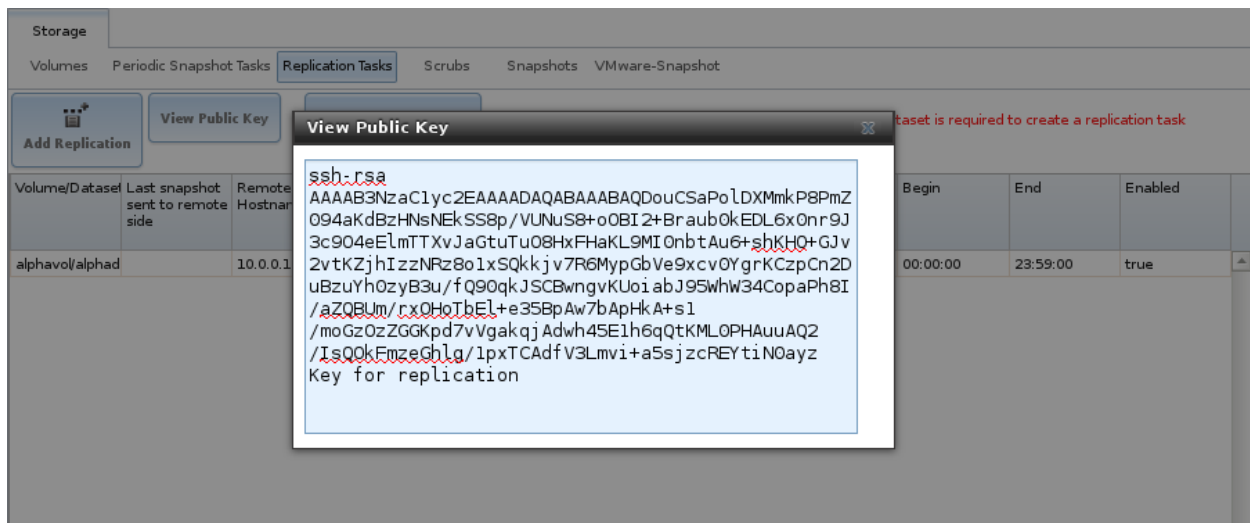


Fig. 7.28: Copy the Replication Key

On *Beta*, select **Account** → **Users** → **View Users**. Click the *root* account to select it, then click *Modify User*. Paste the copied key into the *SSH Public Key* field and click *OK* as shown in Figure 7.29.

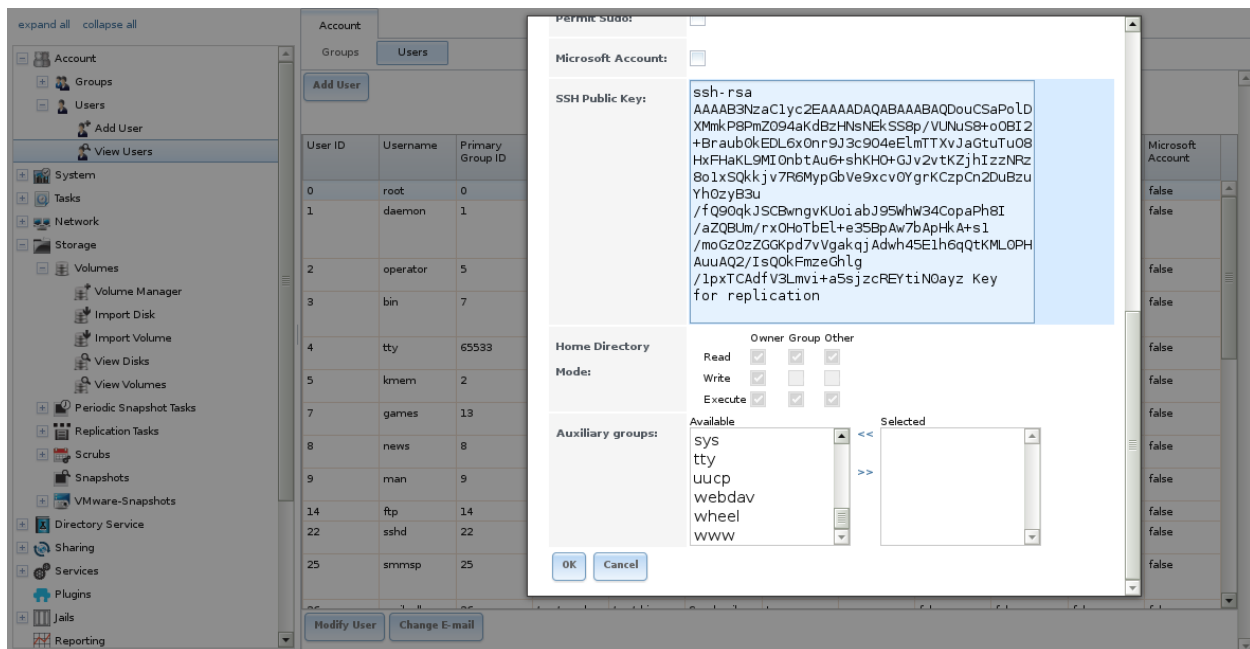


Fig. 7.29: Paste the Replication Key

Back on *Alpha*, create the replication task by clicking *Replication Tasks* and *Add Replication*. *alphavol/alphadata* is selected as the dataset to replicate. The destination volume is *betavol*. The *alphadata* dataset and snapshots are replicated there. The IP address of *Beta* is entered in the *Remote hostname* field as shown in Figure 7.30. A hostname can be entered here if local DNS resolves for that hostname.

Click the *SSH Key Scan* button to retrieve the SSH host keys from *Beta* and fill the *Remote hostkey* field. Finally, click *OK* to create the replication task. After each periodic snapshot is created, a replication task will copy it

to the destination system. See [Limiting Replication Times](#) (page 122) for information about restricting when replication is allowed to run.

Add Replication

Volume/Dataset:	<div> <div>alphavol/alphadata</div> <div></div> </div>
Remote ZFS Volume/Dataset:	<div> <div>betavol</div> <div></div> </div>
Recursively replicate child dataset's snapshots:	<div> <input type="checkbox"/> </div>
Delete stale snapshots on remote system:	<div> <input type="checkbox"/> </div>
Replication Stream Compression:	<div> <div>lz4 (fastest)</div> <div></div> </div>
Limit (kB/s):	<div> <div>0</div> <div></div> </div>
Begin:	<div> <div>00:00:00</div> <div></div> </div>
End:	<div> <div>23:59:00</div> <div></div> </div>
Enabled:	<div> <input checked="" type="checkbox"/> </div>
Setup mode:	<div> <div>Manual</div> <div></div> </div>
Remote hostname:	<div> <div>10.0.0.118</div> <div></div> </div>
Remote port:	<div> <div>22</div> <div></div> </div>
Dedicated User Enabled:	<div> <input type="checkbox"/> </div>
Dedicated User:	<div> <div></div> <div></div> </div>
Encryption Cipher:	<div> <div>Standard</div> <div></div> </div>
Remote hostkey:	<div> <div> 10.0.0.118 ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCA4WnS+kfJa CDL1SnPWEqHwuVjE0k8pl+kU8JlS8yyf0ALP1/aB c82DdZoNGwtJjn14xTyxA1XJKXio1YYkTnTiLj7M R+S905HLt+vwSUhkfs3EdD8/oOCFmeiw /00dzjT9oiCrqqnHiL+dySqBjAE0yfoQyTGfzbsy FYG9BZ6aLSzA+oEd7i+aJlE++n6oRCENUCopeFGF m9gADtWwETiHxJkY292JRqhy02k7JrhzyYPSLZvL Yy3mwObSG1Xjf8D2xGxs7qdiai3r6aKl+TRA4Bi /d8GxVAKwzJPgv /K/aWiibmaUcVBavUbM60yaRFg9uuhn43HYMHbJa 4fE/r1 10.0.0.118 ecdsa-sha2-nistp256 AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlz dHAyNTYAAABBBANGLomMyTZl/FplaScYX /8S/b3nvXibX /levDCDwJecuD1ASwY5Xx+Wp8YkraJzLv9bonf1w yc2fCL4gzFs0Ag= 10.0.0.118 ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIOZtUTtc59hv90WH 7nDoD4li3GdRkaZR/V70gzT8t7GE </div> </div>

OK

Cancel

SSH Key Scan

7.3.4 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 <i>lz4 (fastest)</i> , <i>pigz (all rounder)</i> , <i>plzip (best compression)</i> , or <i>Off</i> (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 <i>Begin</i> and <i>End</i> 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	<i>Manual</i> or <i>Semi-automatic</i>
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 <i>Dedicated User Enabled</i> is checked; select the user account to be used for replication
Encryption Cipher	drop-down menu	<i>Standard</i> , <i>Fast</i> , or <i>Disabled</i>
Remote hostkey	string	use the <i>SSH Key Scan</i> 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.

Selecting **Storage** → **Replication Tasks** displays [Figure 7.31](#), 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.



Storage										
Volumes Periodic Snapshot Tasks Replication Tasks Scrubs Snapshots VMware-Snapshot										
Add Replication View Public Key ATTENTION: A periodic snapshot of a given ZFS Volume/Dataset is required to create a replication task										
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. 7.31: 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.

7.3.5 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.

7.3.6 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 125). 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.

7.3.7 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.

7.3.8 Troubleshooting Replication

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

SSH

[SSH](#) (page 217) must be able to connect from the source system to the destination system with an encryption key. This can be tested from [Shell](#) (page 237) by making an [SSH](#) (page 217) 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 237) 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 preceeding information on the screen and type *yes*. At this point, an [SSH](#) (page 217) shell connection is open to the destination system, *Beta*.

If a password is requested, SSH authentication is not working. See [Figure 7.28](#) 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 237) 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 237) 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 237) with `zfs list -t snapshot` or by going to *Storage* → *Snapshots*.

Error messages here can indicate any remaining problems.

7.4 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 71) 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 71) 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 7.9](#) summarizes the options in this screen.

The screenshot shows the TrueNAS web interface. On the left, the 'Storage' menu is open, and the 'Scrubs' tab is selected. Below the menu, there is a table of scrub entries. The entry for 'volume1' is highlighted. On the right, the 'Edit' screen for 'volume1 (ZFS)' is displayed. The 'Threshold days' is set to 35. The 'Description' field is empty. The 'Minute' field is set to 'Every N minute' with a grid of options from 00 to 59. The 'Hour' field is set to 'Every N hour' with a grid of options from 00 to 23.

Volume	Threshold days	Description	Minute	Hour	Day of month	Month	Day of week	Enabled
volume1	35		00	00	Everyday	Every month	Sunday	true

Edit

Volume: volume1 (ZFS)

Threshold days: 35

Description:

Minute: Every N minute Each selected minute

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

Hour: Every N hour Each selected hour

00	01	02	03	04	05	06	07	08	09
10	11	12	13	14	15	16	17	18	19
20	21	22	23						

Fig. 7.32: Viewing a Volume's Default Scrub Settings

Table 7.9: ZFS Scrub Options

Setting	Value	Description
Volume	drop-down menu	select ZFS volume to scrub
Threshold days	integer	number of days since the last scrub completed before the next scrub can occur, regardless of the calendar schedule; the default is a multiple of 7 which should ensure that the scrub always occurs on the same day of the week
Description	string	optional
Minute	slider or minute selections	if use the slider, scrub occurs every N minutes; if use minute selections, scrub starts at the highlighted minutes
Hour	slider or hour selections	if use the slider, scrub occurs every N hours; if use hour selections, scrub occurs at the highlighted hours
Day of Month	slider or month selections	if use the slider, scrub occurs every N days; if use month selections, scrub occurs on the highlighted days of the selected months
Month	checkboxes	scrub occurs on the selected months
Day of week	checkboxes	scrub occurs on the selected days; default is <i>Sunday</i> to least impact users
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.

Scrubs can be deleted with the *Delete* button, but **deleting a scrub is not recommended as a scrub provides an early indication of disk issues that could lead to a disk failure**. If a scrub is too intensive for the hardware, consider unchecking the *Enabled* button for the scrub as a temporary measure until the hardware can be upgraded.






7.5 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 112) 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 237).

Storage

VolumesPeriodic Snapshot TasksReplication TasksScrubsSnapshotsVMware-Snapshot

Volume/Dataset	Snapshot Name	Used	Refer	Replication	Available Actions
No filter applied					
<input type="checkbox"/> volume1	auto-20150204.0959-2h	100.0 KiB	684.0 KiB		 
<input type="checkbox"/> volume1	auto-20150204.1004-2h	0	664.0 KiB		  

1 - 2 of 2 items

10 | 25 | 50 | 100 | All


 Destroy

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 237).

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.

Replication shows whether the snapshot has been replicated to a remote system.

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 really a dataset which can be mounted, the clone appears in the *Active Volumes* tab, instead of the *Periodic Snapshots* tab, and has the word *clone* in its name.

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 171). 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 your filter and exit the define filter screen. Alternately, click the + button to add another filter.

If you create multiple filters, select the filter to use before leaving the define filter screen. Once a filter is selected, the *No filter applied* text changes to *Clear filter*. If you click *Clear filter*, a pop-up message indicates that this removes the filter and all available snapshots are listed.

7.6 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 126).

Figure 7.34 shows the menu for adding a VMware snapshot and Table 7.10 summarizes the available options.

Add VMware-Snapshot

Hostname:

Username:
i

Password:

ZFS Filesystem:
volume1
▼

Datastore:
▼
i

OK
Cancel
Fetch Datastores

Fig. 7.34: Adding a VMware Snapshot

Table 7.10: 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 <i>Username</i>
ZFS Filesystem	drop-down menu	the filesystem to snapshot
Datastore	drop-down menu	after entering the <i>Hostname</i> , <i>Username</i> , and <i>Password</i> , click <i>Fetch Datastores</i> to populate the menu and select the datastore with which to synchronize

DIRECTORY SERVICES

TrueNAS® supports integration with these directory services:

- [Active Directory](#) (page 130) (for Windows 2000 and higher networks)
- [LDAP](#) (page 136)
- [NIS](#) (page 139)
- [NT4](#) (page 140) (for Windows networks older than Windows 2000)

It also supports [Kerberos Realms](#) (page 141), [Kerberos Keytabs](#) (page 142), 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.

Note: If the network has an NT4 domain controller, or any domain controller with a version earlier than Windows 2000, configure [NT4](#) (page 140) instead.

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® 9.10.2 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 237) 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*.

Fig. 8.1: Configuring Active Directory

Table 8.1: Active Directory Configuration Options

Setting	Value	Advanced Mode	Description
Domain Name	string		name of Active Directory domain (<i>example.com</i>) or child domain (<i>sales.example.com</i>); 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
Encryption Mode	drop-down menu	✓	choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i>
Continued on next page			

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Certificate	drop-down menu	✓	select the certificate of the LDAP server if SSL connections are used; if a certificate does not exist yet, create a CA (in CAs (page 44)), then create a certificate on the Active Directory server and import it to the TrueNAS® system with Certificates (page 47)
Verbose logging	checkbox	✓	when checked, logs attempts to join the domain to <code>/var/log/messages</code>
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 GUIDs, 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 domain/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	✓	only available in <i>Advanced Mode</i> ; when unchecked, the domain name is prepended to the username; if <i>Allow Trusted Domains</i> 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
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
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 142)
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 <i>Edit</i> link to configure that backend's editable options

Continued on next page

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Windbind NSS Info	drop-down menu	✓	defines the schema to use when querying AD for user/group info; <i>rfc2307</i> uses the RFC2307 schema support included in Windows 2003 R2, <i>sfu20</i> is for Services For Unix 3.0 or 3.5, and <i>sfu</i> is for Services For Unix 2.0
SASL wrapping	drop-down menu	✓	defines how LDAP traffic is transmitted; choices are <i>plain</i> (plain text), <i>sign</i> (signed only), or <i>seal</i> (signed and encrypted); Windows 2000 SP3 and higher can be configured to enforce signed LDAP connections
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 <i>Workgroup</i> name
NetBIOS Name (Node B)	string	✓	limited to 15 characters; when using <i>Failover</i> (page 53), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	✓	limited to 15 characters; when using <i>Failover</i> (page 53), this is the NetBIOS name that resolves to either node

Table 8.2 summarizes the backends which are available in the *ldmap backend* drop-down menu. Each backend has its own [man page](https://www.samba.org/samba/docs/man/manpages/) (<https://www.samba.org/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 <i>uidNumber</i> attributes for users and <i>gidNumber</i> attributes for groups in the AD
adex	AD server uses RFC2307 schema extensions and supports domain trusts as well as two-way cross-forest trusts; mappings must be provided in advance by adding the POSIX attribute information to the users and groups objects in AD using a tool such as “Identity Services for Unix” on Windows 2003 R2 and later
autorid	similar to <i>rid</i> , but automatically configures the range to be used for each domain, so there is no need to specify a specific range for each domain in the forest; the only needed configuration is the range of UID/GIDs to use for user/group mappings and an optional size for the ranges
hash	uses a hashing algorithm for mapping and can be used to support local name mapping files
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

Continued on next page

Table 8.2 – continued from previous page

Value	Description
rid	default for AD and NT4 directory services; 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
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®; otherwise 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 237). 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.

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 bugs.freenas.org (https://bugs.freenas.org/) 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 142):

```
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 must be imported with System → Certificates → Import Certificate.

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 Mac OS X environments](https://forums.freenas.org/index.php?threads/howto-freenas-with-open-directory-in-mac-os-x-environments.46493/) (<https://forums.freenas.org/index.php?threads/howto-freenas-with-open-directory-in-mac-os-x-environments.46493/>).

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

Directory Service

Active Directory **LDAP** NIS NT4 Kerberos Realms Kerberos Keytabs

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. <i>dc=test,dc=org</i>)
Bind DN	string		name of administrative account on LDAP server (e.g. <i>cn=Manager,dc=test,dc=org</i>)
Bind password	string		password for <i>Root bind DN</i>
Allow Anonymous Binding	checkbox	✓	instructs LDAP server to not provide authentication and to allow read and write access to any client
User Suffix	string	✓	optional; can be added to name when user account added to LDAP directory (e.g. dept. or company name)
Group Suffix	string	✓	optional; can be added to name when group added to LDAP directory (e.g. dept. or company name)
Password Suffix	string	✓	optional; can be added to password when password added to LDAP directory
Machine Suffix	string	✓	optional; can be added to name when system added to LDAP directory (e.g. server, accounting)
SUDO Suffix	string	✓	use if LDAP-based users need superuser access
Kerberos Realm	drop-down menu	✓	select the realm created using the instructions in Kerberos Realms (page 141)
Kerberos Keytab	drop-down menu	✓	browse to the location of the keytab created using the instructions in Kerberos Keytabs (page 142)

Continued on next page

Table 8.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Encryption Mode	drop-down menu	✓	choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i> ; note that either <i>SSL</i> or <i>TLS</i> and a <i>Certificate</i> must be selected in order for authentication to work
Certificate	drop-down menu	✓	select the certificate of the LDAP server or the CA that signed that certificate (required if authentication is used); if your LDAP server does not already have a certificate, create a CA using <i>CAs</i> (page 44), then the certificate using <i>Certificates</i> (page 47) and install the certificate on the LDAP server
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 <i>Edit</i> 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 sssds.conf(5) (https://jhrozek.fedorapeople.org/sssds/1.11.6/man/sssds.conf.5.html)
Schema	drop-down menu		if <i>Samba Schema</i> is checked, select the schema to use; choices are <i>rfc2307</i> and <i>rfc2307bis</i>
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 <i>Workgroup</i> name
NetBIOS Name (Node B)	string	✓	limited to 15 characters; when using <i>Failover</i> (page 53), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	✓	limited to 15 characters; when using <i>Failover</i> (page 53), 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 237) 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) (<http://www.openldap.org/doc/admin24/appendix-common-errors.html>) for common errors and how to fix them. When troubleshooting LDAP, open *Shell* (page 237) 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.

The screenshot shows the 'NIS' configuration tab under 'Directory Service'. It includes input fields for 'NIS domain' and 'NIS servers', and checkboxes for 'Secure mode', 'Manycast', and 'Enable'. Information icons are present next to the input fields and checkboxes. At the bottom, there are 'Save' and 'Rebuild Directory Service Cache' buttons.

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, ypbind(8) (http://www.freebsd.org/cgi/man.cgi?query=ybind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, ybind will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet
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 NT4

This service should only be configured if the Windows network's domain controller is running NT4. If the network's domain controller is running a more recent version of Windows, you should configure [Active Directory](#) (page 130) instead.

Figure 8.4 shows the configuration screen that appears when `Directory Service` → `NT4` is clicked. These options are summarized in Table 8.5. Some settings are only available in Advanced Mode. To see these settings, either click the *Advanced Mode* button or configure the system to always display these settings by checking the box *Show advanced fields by default* in `System` → `Advanced`.

Directory Service

Active Directory LDAP NIS **NT4** Kerberos Realms Kerberos Keytabs

Domain Controller: ⓘ

NetBIOS Name: ⓘ

Workgroup Name: ⓘ

Administrator Name: ⓘ

Administrator Password: ⓘ

Confirm Administrator Password:

Enable: ☐

Fig. 8.4: NT4 Configuration Options

Table 8.5: NT4 Configuration Options

Setting	Value	Advanced Mode	Description
Domain Controller	string		hostname of domain controller
NetBIOS Name	string		hostname of TrueNAS [®] system ; cannot be longer than 15 characters; cannot be the same as the <i>Workgroup Name</i>
Workgroup Name	string		name of Windows server's workgroup
Administrator Name	string		name of the domain administrator account
Administrator Password	string		input and confirm the password for the domain administrator account
Use default domain	checkbox		only available in <i>Advanced Mode</i> ; when unchecked, the domain name is prepended to the username

Continued on next page

Table 8.5 – continued from previous page

Setting	Value	Advanced Mode	Description
Idmap backend	drop-down and Edit menu	✓	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 <i>Edit</i> link to configure the backend's editable options
Enable	checkbox		uncheck to disable the configuration without deleting it

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

8.5 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 8.5](#).

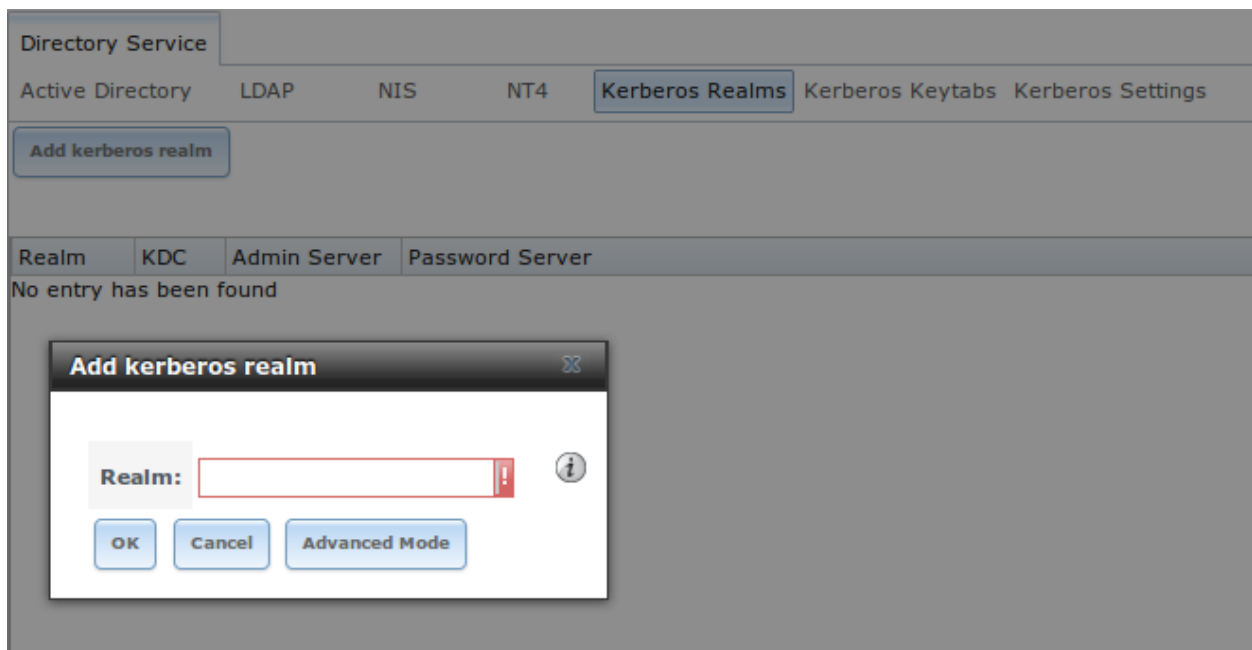


Fig. 8.5: Adding a Kerberos Realm

[Table 8.6](#) 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 8.6: 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

8.6 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 the keytab, select the installed keytab using the drop-down “Kerberos keytab” menu in *Directory Service* → *LDAP*.

8.7 Kerberos Settings

To configure additional Kerberos parameters, use *Directory Service* → *Kerberos Settings*. [Figure 8.6](#) 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).

The screenshot shows a web interface for configuring a Directory Service. At the top, there is a horizontal menu with tabs: 'Directory Service', 'Active Directory', 'LDAP', 'NIS', 'NT4', 'Kerberos Realms', 'Kerberos Keytabs', and 'Kerberos Settings'. The 'Kerberos Settings' tab is currently selected and highlighted. Below the menu, the main content area is divided into two sections. The first section is labeled 'Appdefaults auxiliary parameters:' and contains a large, empty text input field. The second section is labeled 'Libdefaults auxiliary parameters:' and also contains a large, empty text input field. At the bottom left of the configuration area, there is a blue 'Save' button.

Fig. 8.6: 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 229) for creating shares. The *Wizard* (page 229) 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 145): Apple File Protocol shares are often used when the client computers all run Mac OS X. Apple has slowly shifted to preferring *SMB* (page 162) for modern networks, although Time Machine still requires AFP.
- *Unix (NFS)* (page 153): 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 161): 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 162): 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 173): 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 198) service instead of a share and using a cross-platform FTP file manager application such as [Filezilla](https://filezilla-project.org/) (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) (http://winscp.net/eng/index.php), consider using the [SSH](#) (page 217) 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 a 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 191).

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 229). It then provides configuration examples for using the [Wizard](#) (page 229) 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.

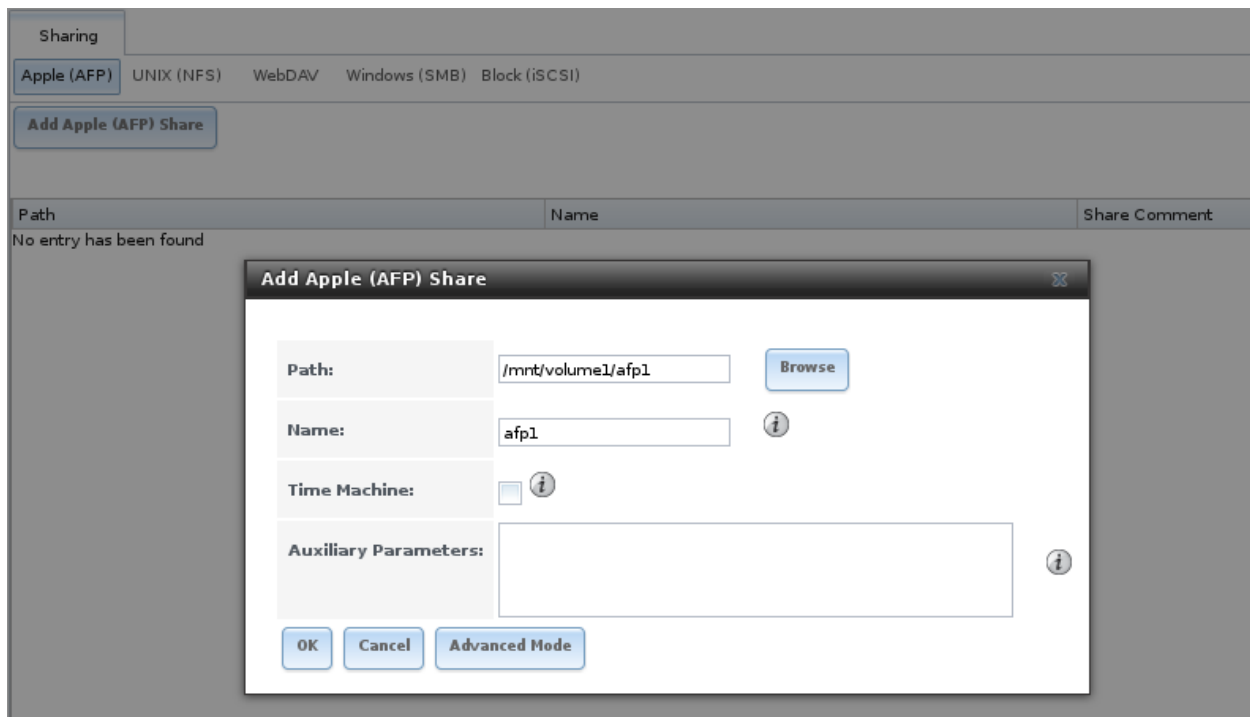


Fig. 9.1: Creating an AFP Share

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 <i>connect to server</i> 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

Continued on next page

Table 9.1 – continued from previous page

Setting	Value	Advanced Mode	Description
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 <i>Time Machine</i> on multiple shares could result in intermittent failed backups
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 a window titled "Wizard" with a close button in the top right corner. Inside the window, there is a "Share name:" label followed by a text input field containing "afp_guest". Below this is a "Purpose" section with four radio button options: "Windows (SMB)", "Mac OS X (AFP)" (which is selected), "Generic Unix (NFS)", and "Block Storage (iSCSI)". To the right of these options are two checkboxes: "Allow Guest" and "Time Machine". A button labeled "Ownership" is positioned to the right of the "Purpose" section. Below the "Purpose" section are three buttons: "Add", "Delete", and "Update". Underneath these buttons is a list box with the heading "Name" and a single 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.

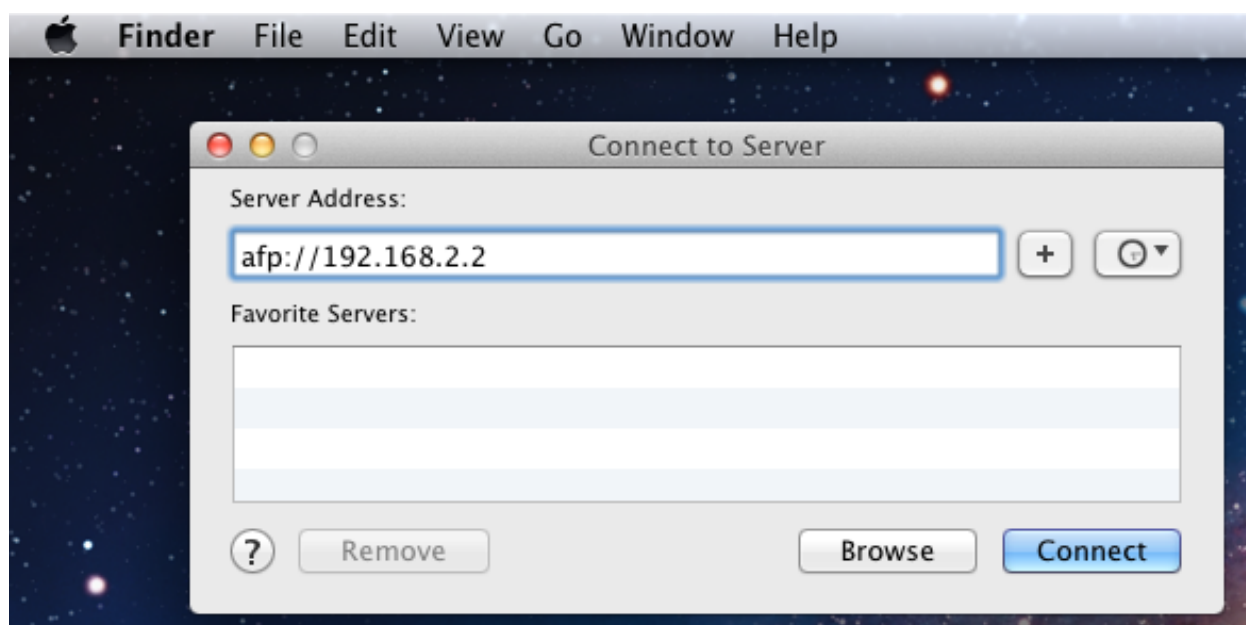


Fig. 9.3: Connect to Server Dialogue

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*. If the user will not be using Time Machine, leave the box unchecked.
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'. A list box labeled 'Name' contains the 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"/>

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.

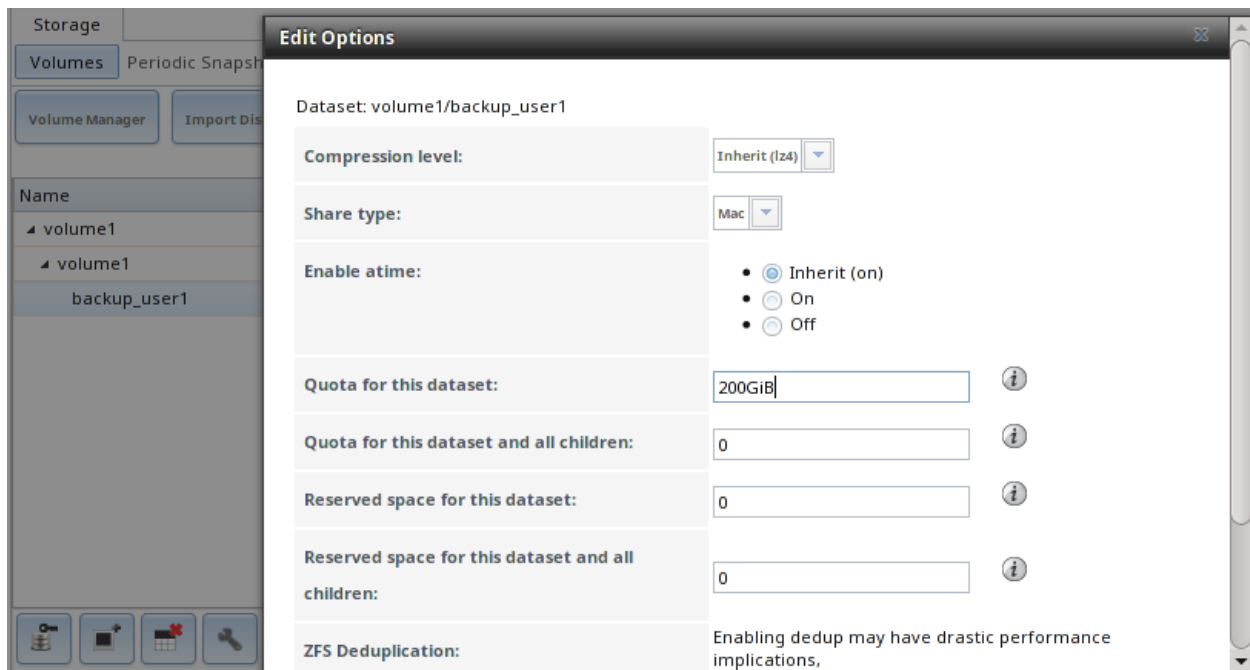


Fig. 9.6: Setting a Quota

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

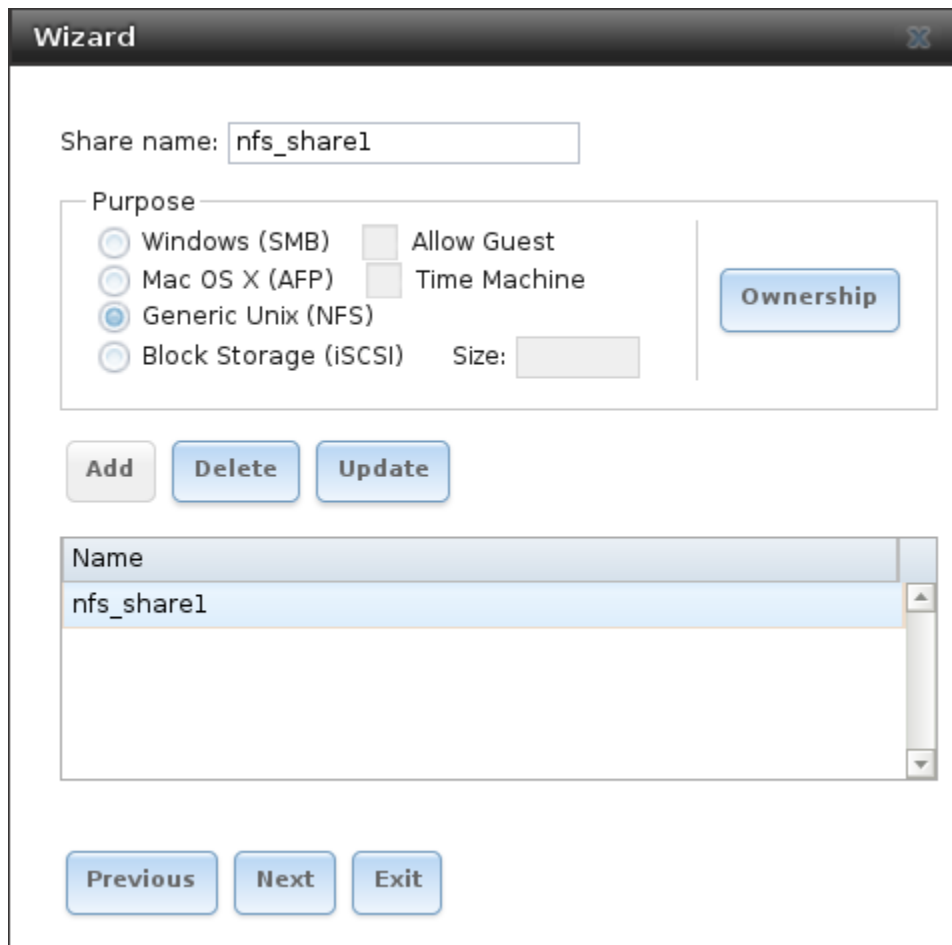
If *Time Machine* could not complete the backup. The backup disk image could not be created (error 45) is shown when backing up to the TrueNAS[®] system, a sparsebundle image must be created using [these instructions](http://forum1.netgear.com/showthread.php?t=49482) (<http://forum1.netgear.com/showthread.php?t=49482>).

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 229), 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' window for configuring an NFS share. 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 contains a single entry 'nfs_share1' under the header 'Name'. At the bottom of the window 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.

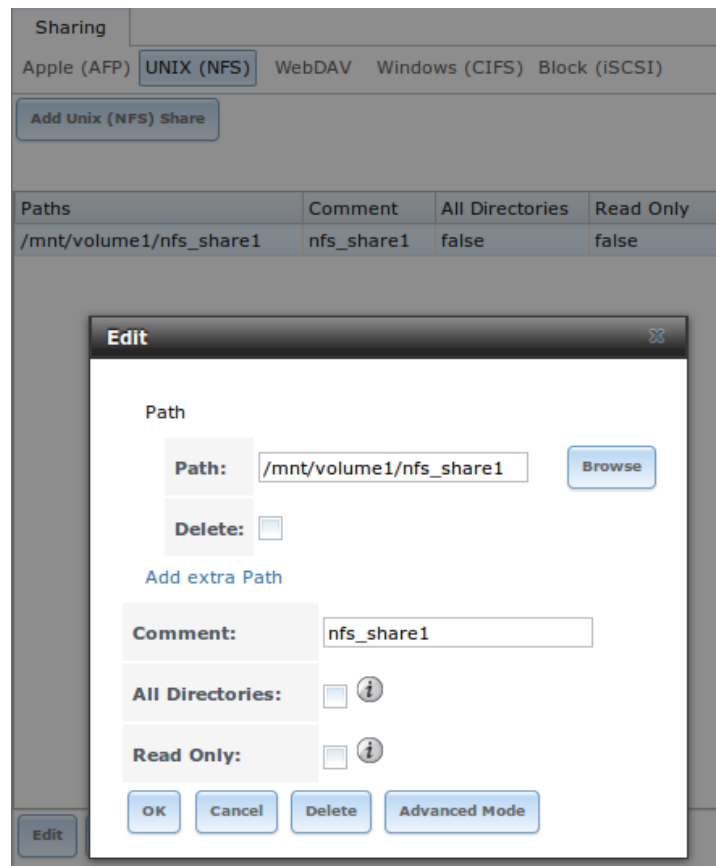


Fig. 9.9: NFS Share Settings

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 <i>Add extra path</i> to select multiple paths
Comment	string		set the share name; if left empty, share name is the list of selected <i>Path</i> entries
Authorized networks	string	✓	list of allowed networks in network/mask CIDR notation, like <i>1.2.3.0/24</i> , 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 <i>Path</i>
Read only	checkbox		prohibit writing to the share

Continued on next page

Table 9.2 – continued from previous page

Setting	Value	Advanced Mode	Description
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 <i>root</i> user is limited to that user's permissions
Maproot Group	drop-down menu	✓	when a group is selected, the <i>root</i> 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 <i>Enable NFSv4</i> 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 the following 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 the following scenario where there are:

- 2 networks named *10.0.0.0/8* and *20.0.0.0/8*
- a ZFS volume named *volume1* with 2 datasets named *dataset1* and *dataset2*
- *dataset1* has 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* options show as *N/A*. This means that when a user connects to the NFS share, they connect with 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 scenario is to do the following:

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

The NFS share is mounted on BSD or Linux systems needing access with this command executed as the superuser 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 162).

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.

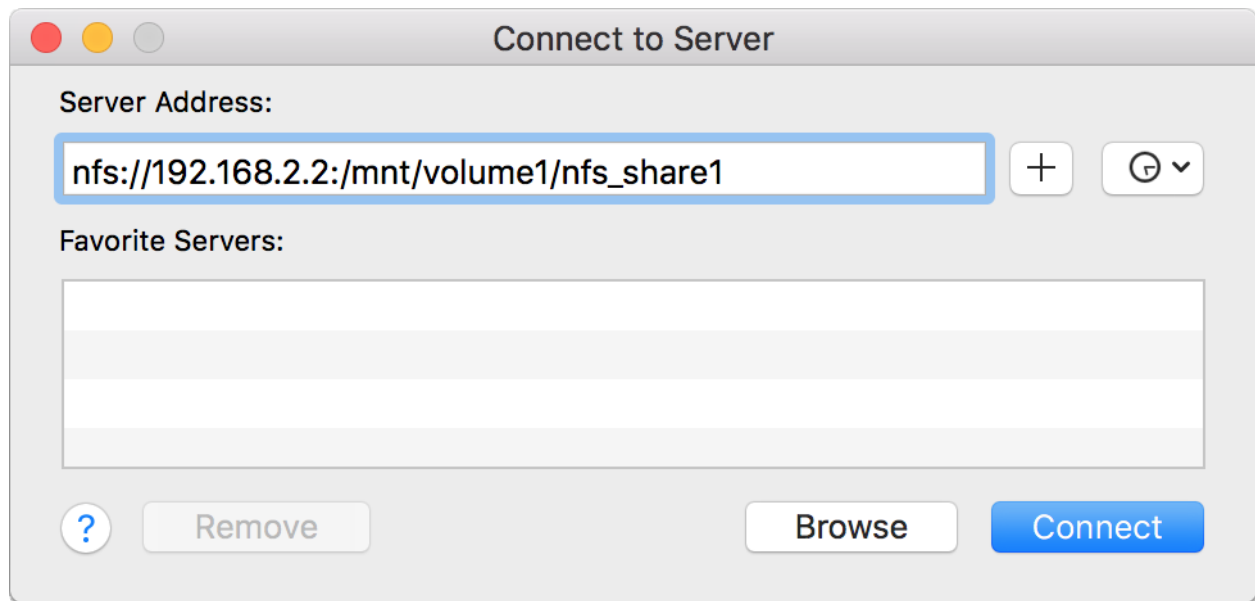


Fig. 9.10: Mounting the NFS Share from Mac OS X

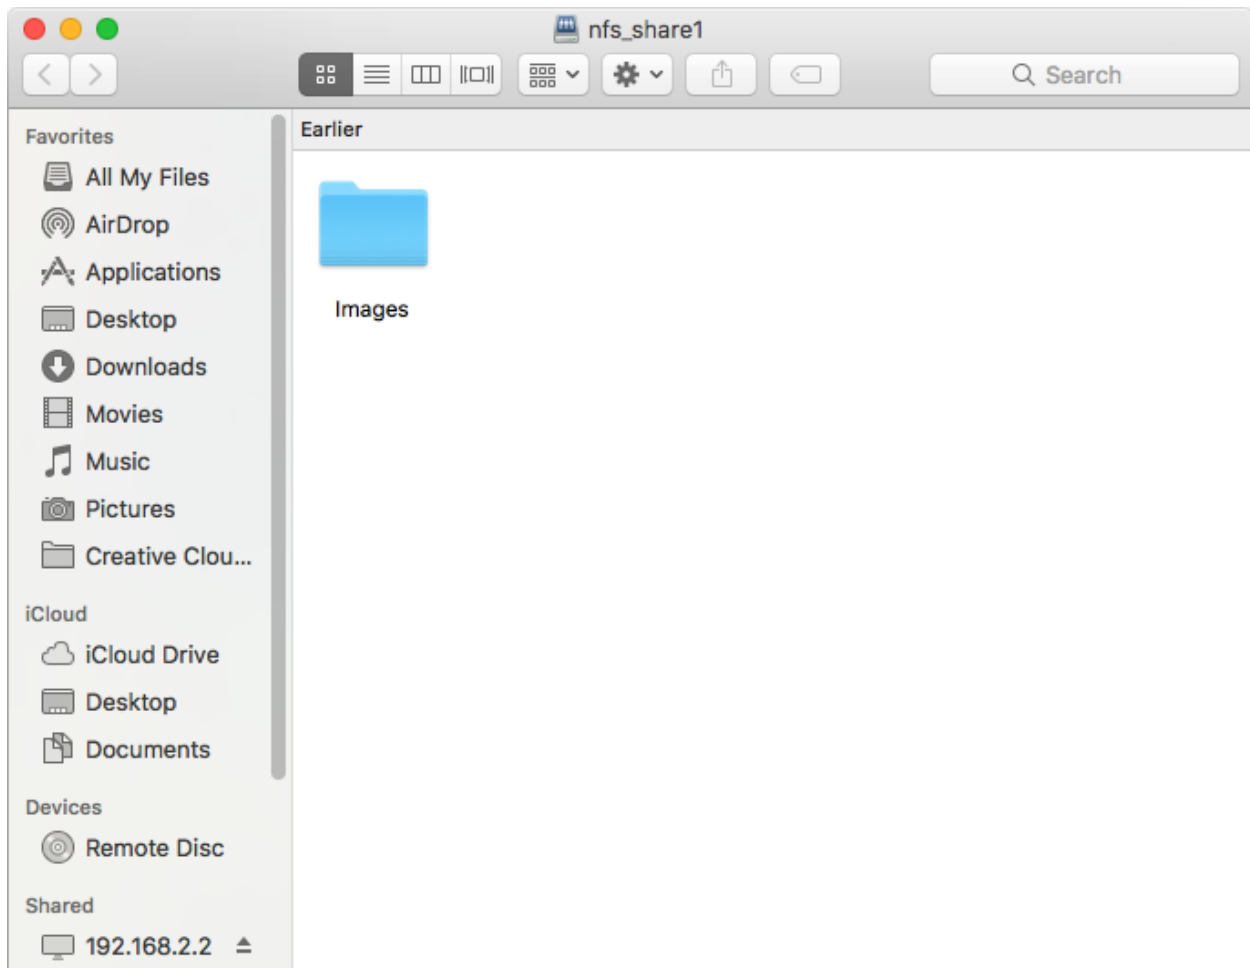


Fig. 9.11: Viewing the NFS Share in Finder

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 237). 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](#).

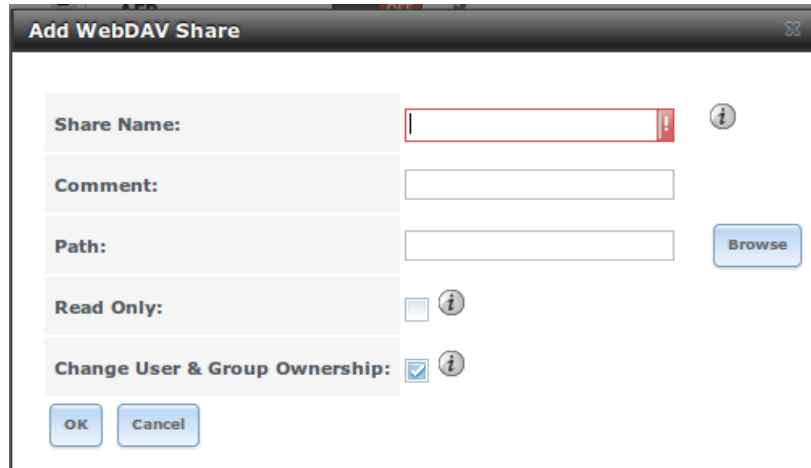


Fig. 9.12: Adding a WebDAV Share

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 <i>webdav</i> 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 223).

9.4 Windows (SMB) Shares

TrueNAS® uses [Samba](https://www.samba.org/) (<https://www.samba.org/>) to share volumes using Microsoft's SMB protocol. SMB is built into the Windows and Mac OS X operating systems and most Linux and BSD systems pre-install the Samba client in order to provide support for SMB. If your distro did not, install the Samba client using the distro's software repository.

The SMB protocol supports many different types of configuration scenarios, ranging from the very simple to quite complex. The complexity of the scenario depends upon the types and versions of the client operating systems that will connect to the share, whether the network has a Windows server, and whether Active Directory is being used. Depending on the authentication requirements, it might be necessary to create or import users and groups.

Samba supports server-side copy of files on the same share with clients from Windows 8 and higher. Copying between two different shares is not server-side. Windows 7 clients support server-side copying with [Robocopy](https://technet.microsoft.com/en-us/library/cc733145) (<https://technet.microsoft.com/en-us/library/cc733145>).

This chapter starts by summarizing the available configuration options. It demonstrates some common configuration scenarios as well as offering some troubleshooting tips. It is recommended to first read through this entire chapter before creating any SMB shares to get a better idea of the configuration scenario that best meets your network's needs.

Tip: [SMB Tips and Tricks](https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/) (<https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/>) shows helpful hints for configuring and managing SMB networking. The [FreeNAS and Samba \(CIFS\) permissions](https://www.youtube.com/watch?v=RxggaE935PM) (<https://www.youtube.com/watch?v=RxggaE935PM>) and [Advanced Samba \(CIFS\) permissions on FreeNAS](https://www.youtube.com/watch?v=QhwOyLtArw0) (<https://www.youtube.com/watch?v=QhwOyLtArw0>) videos clarify setting up permissions on SMB shares. Another helpful reference is [Methods For Fine-Tuning Samba Permissions](https://forums.freenas.org/index.php?threads/methods-for-fine-tuning-samba-permissions.50739/) (<https://forums.freenas.org/index.php?threads/methods-for-fine-tuning-samba-permissions.50739/>).

Tip: Run `smbstatus` from the *Shell* (page 237) 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*.

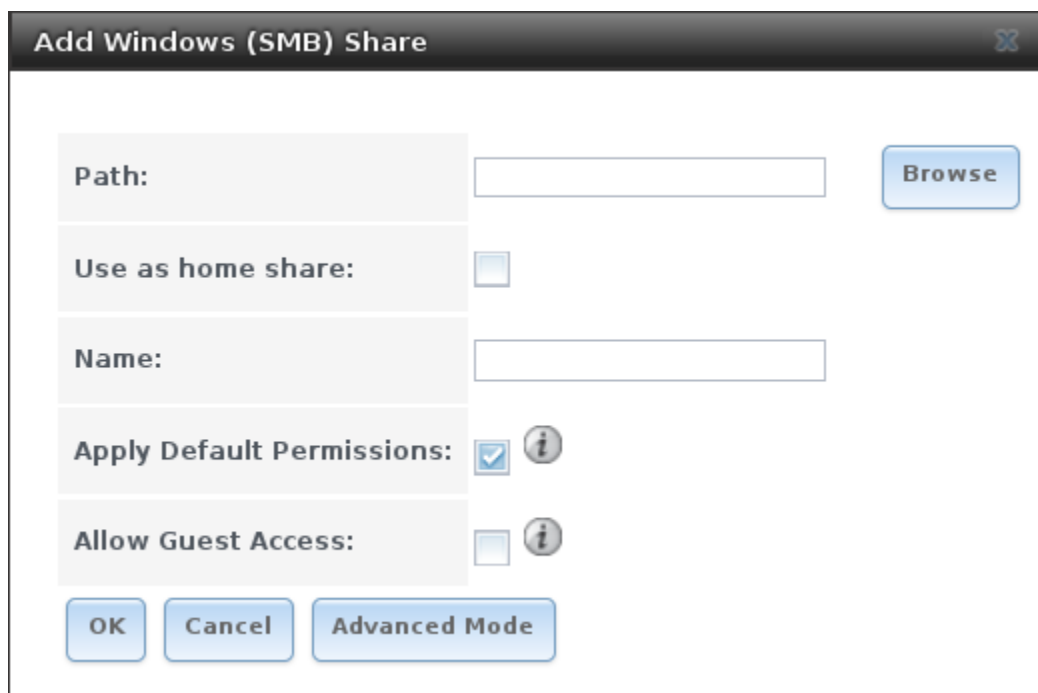


Fig. 9.13: Adding an SMB Share

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+10.3-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 <i>/homes</i> (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 <i>.recycle</i> in the root folder of the share; the <i>.recycle</i> directory can be deleted to reclaim space and is automatically recreated when a file is deleted
Show Hidden Files	checkbox	✓	if enabled, the Windows hidden attribute is not set when filenames that begin with a dot (a Unix hidden file) are created; existing files are not affected
Allow Guest Access	checkbox		if checked, no password is required to connect to the share and all users share the permissions of the guest user defined in the <i>SMB</i> (page 210) service
Only Allow Guest Access	checkbox	✓	requires <i>Allow guest access</i> to also be checked; forces guest access for all connections
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 <i>ALL</i> in this field and specify allowed hosts in <i>Hosts Allow</i>
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 home directory shadow copies on a per-share basis; select the pre-configured periodic snapshot task to use for the share's shadow copies
Auxiliary Parameters	string	✓	additional <i>smb4.conf</i> 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*.
- Be careful about unchecking the *Browsable to Network Clients* box. When this box is checked (the default), other users will see the names of every share that exists using Windows Explorer, but they will receive a permissions denied error message if they try to access someone else's share. If this box is unchecked, even the owner of the share will not see it or be able to create a drive mapping

for the share in Windows Explorer. However, they can still access the share from the command line. Unchecking this option provides limited security and is not a substitute for proper permissions and password control.

- 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+10.3-RELEASE+and+Ports) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+10.3-RELEASE+and+Ports>).

To configure support for OS/2 clients, add this line to *Auxiliary Parameters*:

```
lanman auth = yes
```

To configure lanman authentication for pre-NT authentication, add these lines instead:

```
client lanman auth = yes
client plaintext auth = yes
```

Samba disables NTLMv1 authentication by default for security. Standard configurations of Windows XP and some configurations of later clients like Windows 7 will not be able to connect with NTLMv1 disabled. [Security guidance for NTLMv1 and LM network authentication](https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-network-authentication) (<https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-network-authentication>) has information about the security implications and ways to enable NTLMv2. If changing the client configuration is not possible, NTLMv1 authentication can be enabled by adding this entry to *Auxiliary Parameters*:

```
ntlm auth = yes
```

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) (<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_posix	enables asynchronous I/O on systems running POSIX kernels
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
Continued on next page	

Table 9.5 – continued from previous page

Value	Description
catia	creates filenames that use characters that are illegal in SMB filenames
commit	tracks the amount of data written to a file and synchronizes it to disk when a specified amount accumulates
crossrename	allows server side rename operations even if source and target are on different physical devices
default_quota	stores the default quotas that are reported to a windows client in the quota record of a user
dfs_samba4	distributed file system for providing an alternative name space, load balancing, and automatic failover
dirsort	sorts directory entries alphabetically before sending them to the client
expand_msdfs	enables support for Microsoft Distributed File System (DFS)
extd_audit	sends <i>audit</i> logs to both syslog and the Samba log files
fake_acls	stores file ownership and ACLs as extended attributes
fake_perms	allows roaming profile files and directories to be set as read-only
fruit	enhances OS X support by providing the SMB2 AAPL extension and Netatalk interoperability (see 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
posix_eadb	provides Extended Attributes (EAs) support so they can be used on filesystems which do not provide native support for EAs
preopen	useful for video streaming applications that want to read one file per frame
readahead	useful for Windows Vista clients reading data using Windows Explorer
readonly	marks a share as read-only for all clients connecting within the configured time period
scannedonly	ensures that only files that have been scanned for viruses are visible and accessible
shadow_copy	allows Microsoft shadow copy clients to browse shadow copies on Windows shares
shadow_copy_test	shadow copy testing
Continued on next page	

Table 9.5 – continued from previous page

Value	Description
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)
smb_traffic_analyzer	logs Samba read and write operations through a socket to a helper application
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
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

Note: When using *fruit*, also add the *streams_xattr* and *catia* VFS objects and be sure to configure **all** SMB shares this way. Reboot the Mac client after making this change.

These VFS objects do not appear in the drop-down menu as they are always enabled:

- **recycle:** moves deleted files to the recycle directory instead of deleting them
- **shadow_copy2:** a more recent implementation of *shadow_copy* with some additional features
- **zfs_space:** correctly calculates ZFS space used by share, including any reservations or quotas
- **zfsacl:** provide ACL extensions for proper integration with ZFS.

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*.

The screenshot shows the 'Wizard' window in TrueNAS. At the top, the title bar says 'Wizard'. Below it, the 'Share name' field contains 'smb_insecure'. Under the 'Purpose' section, 'Windows (SMB)' is selected with a radio button, and the 'Allow Guest' checkbox is checked. Other options like 'Mac OS X (AFP)', 'Generic Unix (NFS)', and 'Block Storage (iSCSI)' are unselected. There is a 'Time Machine' checkbox which is unchecked. To the right of these options is an 'Ownership' button. Below the 'Purpose' section are three buttons: 'Add', 'Delete', and 'Update'. At the bottom of the window is a table with a header 'Name' and one row containing 'smb_insecure'. At the very bottom are three buttons: 'Previous', 'Next', and '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).

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_user1

Previous **Next** **Exit**

Fig. 9.15: Creating an Authenticated SMB Share

Wizard

User: ☒ Create User ⓘ

User Password:

Confirm User Password:

Group: ☒ Create Group ⓘ

Mode:

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

Return **Cancel**

Fig. 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) (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) (<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 169) 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.

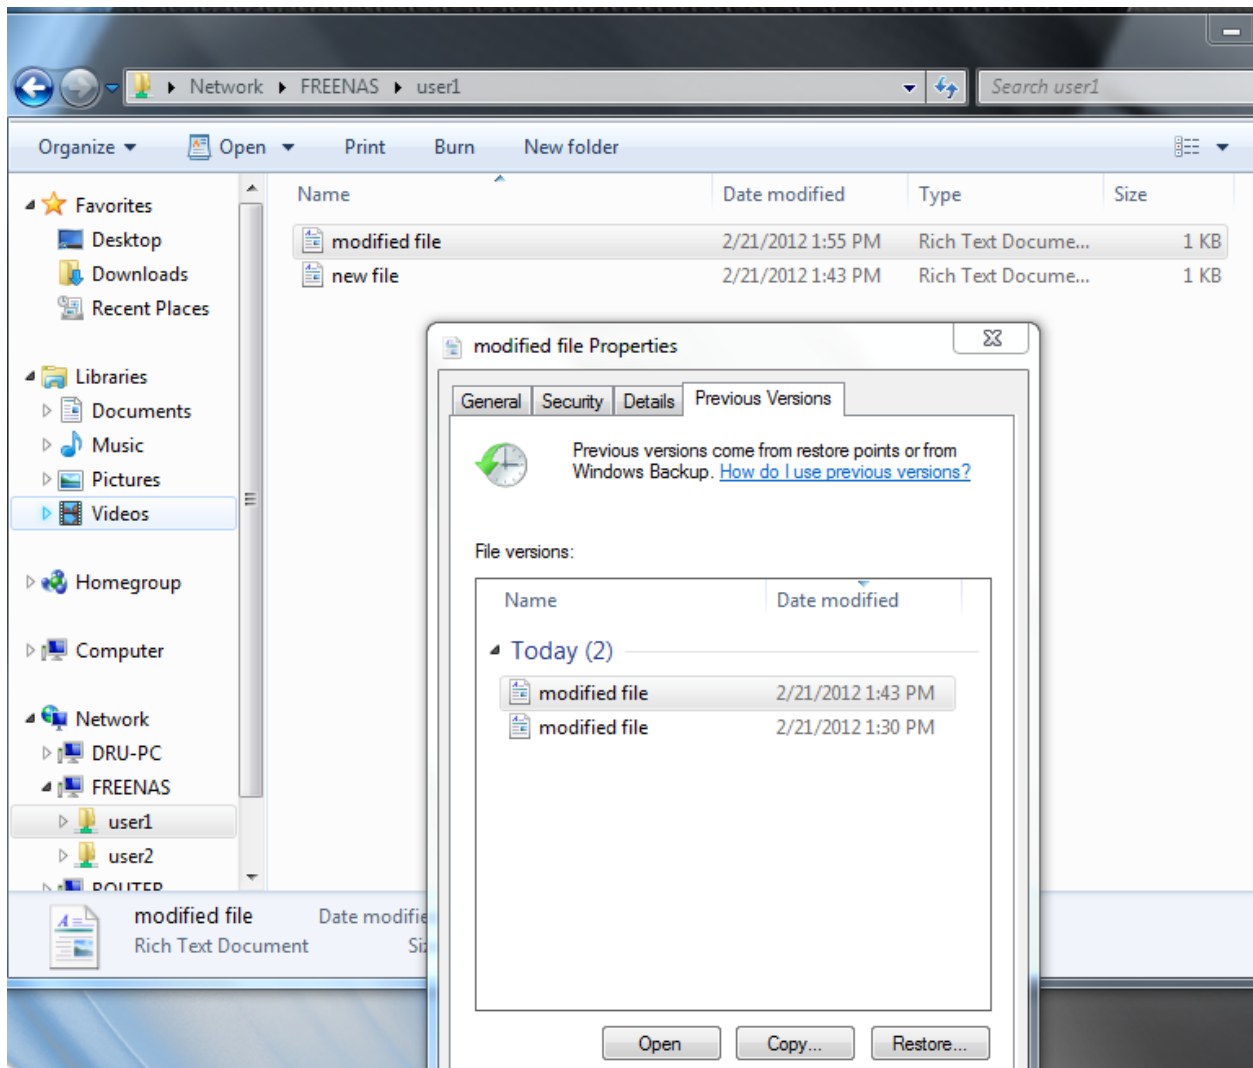


Fig. 9.17: Viewing Previous Versions within Explorer

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.

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 273) (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 97) and use it to create a device extent, as described in [Extents](#) (page 181).

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 184) 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.

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) 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 273) Threshold Warning

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*.

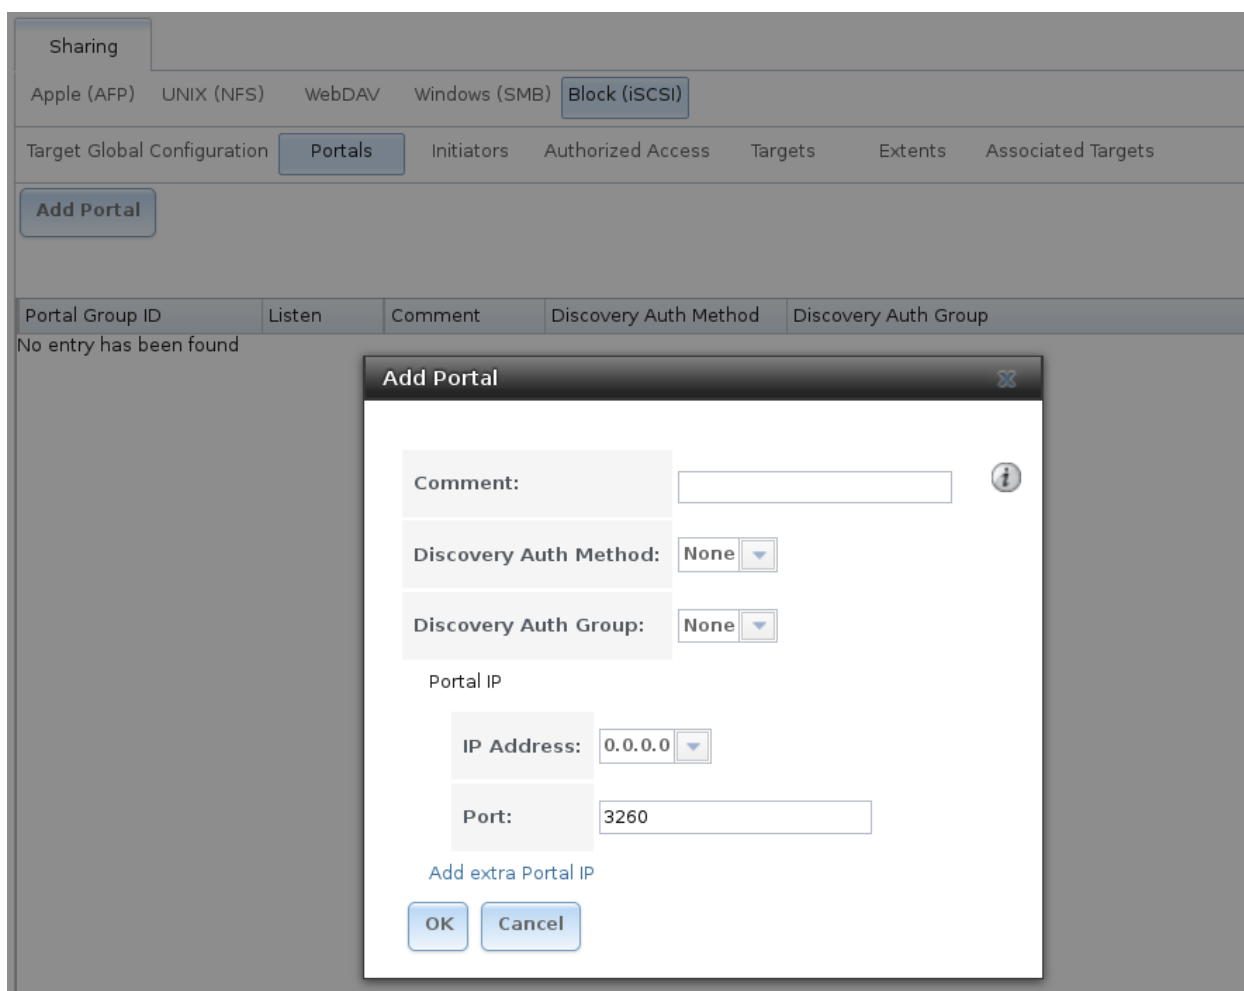


Fig. 9.19: Adding an iSCSI Portal

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 <i>None</i> will allow anonymous discovery while <i>CHAP</i> and <i>Mutual CHAP</i> require authentication
Discovery Auth Group	drop-down menu	select a user created in <i>Authorized Access</i> if the <i>Discovery Auth Method</i> is set to <i>CHAP</i> or <i>Mutual CHAP</i>
IP address	drop-down menu	select the IP address associated with an interface or the wildcard address of <i>0.0.0.0</i> (any interface)
Port	integer	TCP port used to access the iSCSI target; default is <i>3260</i>

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](#).



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 <i>ALL</i> keyword or a list of initiator hostnames separated by spaces
Authorized network	string	use <i>ALL</i> keyword or a network address with CIDR mask such as <i>192.168.2.0/24</i>
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 as well.

The screenshot shows a web interface for iSCSI configuration. At the top, there's a 'Sharing' tab with sub-tabs for 'Apple (AFP)', 'UNIX (NFS)', 'WebDAV', 'Windows (SMB)', and 'Block (iSCSI)'. Below this is a navigation bar with 'Target Global Configuration', 'Portals', 'Initiators', 'Authorized Access', 'Targets', 'Extents', and 'Associated Targets'. The 'Initiators' tab is selected. Below the navigation bar is an 'Add Initiator' button. The main content area contains a table with the following data:

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 174).

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 <i>User</i> ; 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 <i>User</i>
Peer Secret	string	the mutual secret password which must be different than the Secret ; required if <i>Peer User</i> is set

Note: CHAP does not work with GlobalSAN initiators on 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.

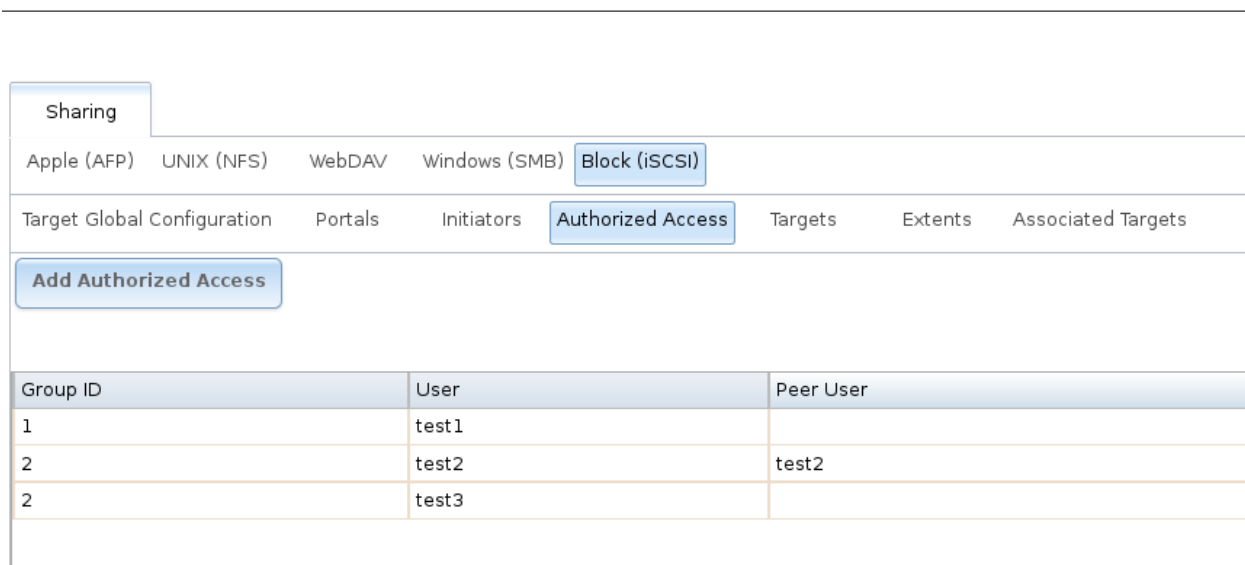


Fig. 9.23: Viewing Authorized Accesses

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 Groups

Portal Group ID:

Initiator Group ID:

Auth Method:

None

Authentication Group number:

None

Add extra iSCSI Groups

OK

Cancel

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 <i>iqn</i>
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 <i>None</i> , <i>Auto</i> , <i>CHAP</i> , or <i>Mutual CHAP</i>
Authentication Group number	drop-down menu	<i>None</i> or integer representing number of existing authorized access

9.5.6 Extents

In iSCSI, the target virtualizes something and presents it as a device to the iSCSI client. That something can be a device extent or a file extent:

Device extent: virtualizes an unformatted physical disk, RAID controller, zvol, zvol snapshot, or an existing [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).

Virtualizing a single disk is slow as there is no caching, but virtualizing a hardware RAID controller has higher performance due to its cache. This type of virtualization does a pass-through to the disk or hardware RAID

181

controller. None of the benefits of ZFS are provided and performance is limited to the capabilities of the disk or controller.

Virtualizing a zvol adds the benefits of ZFS, such as its read cache and write cache. Even if the client formats the device extent with a different filesystem, as far as TrueNAS® is concerned, the data benefits from ZFS features such as block checksums and snapshots.

When determining whether to use a file or a device extent, be aware that a zvol is required to take advantage of all VAAI primitives and is recommended when using 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.

File extent: allows you to export a portion of a ZFS volume. The advantage of a file extent is that you can create multiple exports per volume.

Warning: For performance reasons and to avoid excessive fragmentation, it is recommended to keep the used space of the pool below 50% when using iSCSI. As required, you can increase the capacity of an existing extent using the instructions in [Growing LUNs](#) (page 189).

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.**

Extent Name:	<input type="text"/>
Extent Type:	Device
Serial:	080027e8c63801
Device:	<input type="text"/>
Logical Block Size:	512
Disable Physical Block Size Reporting:	<input type="checkbox"/>
Comment:	<input type="text"/>
Enable TPC:	<input checked="" type="checkbox"/>
Xen initiator compat mode:	<input type="checkbox"/>

Fig. 9.25: Adding an iSCSI Extent

Table 9.11: Extent Configuration Settings

Setting	Value	Description
Extent Name	string	name of extent; if the <i>Extent size</i> is not 0, it cannot be an existing file within the volume/dataset
Extent Type	drop-down menu	select from <i>File</i> or <i>Device</i>
Serial	string	unique LUN ID; the default is generated from the system's MAC address
Path to the extent	browse button	only appears if <i>File</i> is selected; browse to an existing file and use 0 as the <i>Extent size</i> , or browse to the volume or dataset, click <i>Close</i> , append the <i>Extent Name</i> to the path, and specify a value in <i>Extent size</i> ; extents cannot be created inside the jail root directory
Device	drop-down menu	only appears if <i>Device</i> is selected; select the unformatted disk, controller, zvol, zvol snapshot, or HAST device
Extent size	integer	only appears if <i>File</i> is selected; if the size is specified as 0, the file must already exist and the actual file size will be used; otherwise, specify the size of the file to create
Logical Block Size	drop-down menu	only override the default if the initiator requires a different block size
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 <i>File</i> or a zvol is selected; when the specified percentage of free space is reached, the system issues an alert; see VAAI (page 273) Threshold Warning
Comment	string	optional
Enable TPC	checkbox	if checked, an initiator can bypass normal access control and access any scannable target; this allows xcopy 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.

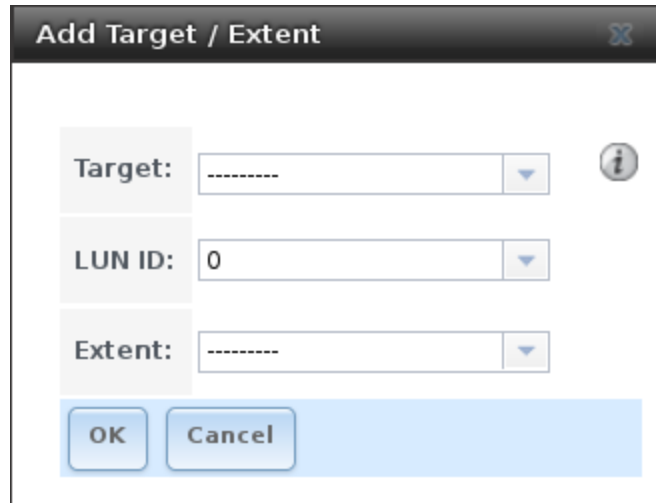


Fig. 9.26: Associating a Target With an Extent

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 of the ID or type in the desired value; TrueNAS® supports up to 1024 LUNs
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) WebDAV Windows (CIFS) **Block (iSCSI/FC)**

Target Global Configuration Portals (iSCSI) Initiators (iSCSI) Authorized Access (iSCSI) Targets Extents Associated Targets Fibre Channel Ports

Base Name: ⓘ

ISNS Servers: ⓘ

Pool Available Space Threshold (%): ⓘ

Fig. 9.27: Block (iSCSI) Screen

Otherwise, the *Target Global Configuration* screen is the same as described in [Target Global Configuration](#) (page 174).

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: ▼

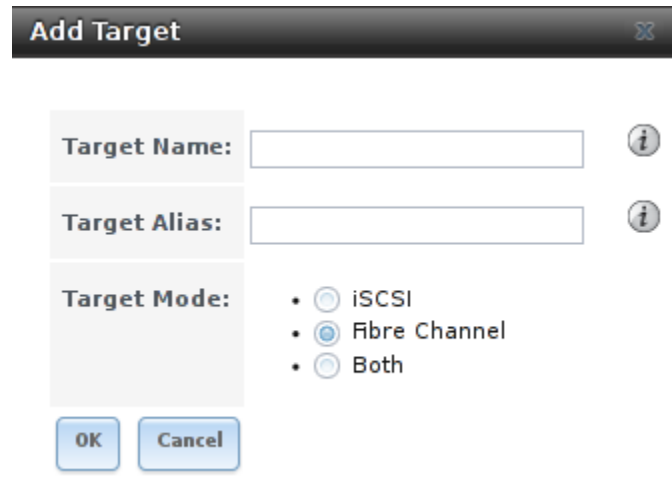
Initiator Group ID: ▼

Auth Method: ⓘ

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](#).



The image shows a dialog box titled "Add Target" with a close button (X) in the top right corner. The dialog contains three input fields, each with an information icon (i) to its right:

- Target Name:** A text input field.
- Target Alias:** A text input field.
- Target Mode:** A radio button selection area with three options:
 - ☐ iSCSI
 - ☒ Fibre Channel
 - ☐ Both

At the bottom of the dialog are two buttons: "OK" and "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 181) and [Target/Extents](#) (page 183).

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	
isp0 - Ready (8 Gbps) WWPN: naa.21000024ff4ce7ea <input type="radio"/> Initiator <input checked="" type="radio"/> Target fc-target <input type="radio"/> Disabled	Connected Initiators - naa.21000024ff5105c1 - naa.21000024ff5105c0 (Node B)
isp0/1 - Ready (8 Gbps) WWPN: naa.22000024ff4ce7ea <input type="radio"/> Initiator <input checked="" type="radio"/> Target fc-target2 <input type="radio"/> Disabled	Connected Initiators - naa.21000024ff5105c1 - naa.21000024ff5105c0 (Node B)
isp0/2 - No Link WWPN: naa.23000024ff4ce7ea <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled	
isp0/3 - No Link WWPN: naa.24000024ff4ce7ea <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled	
isp0/4 - No Link WWPN: naa.25000024ff4ce7ea <input type="radio"/> Initiator <input type="radio"/> Target <input checked="" type="radio"/> Disabled	

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 227) provides Fibre Channel port bandwidth graphs.

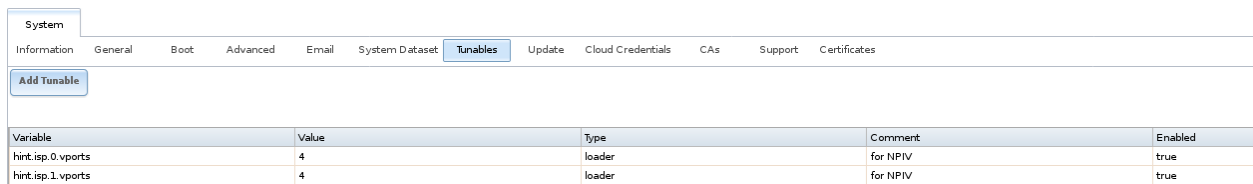
This example has also been configured for NPIV (N_Port ID Virtualization). Note that the physical interface *isp0* has two virtual ports (*isp0/1* and *isp0/2*) displayed in Figure 9.30:. NPIV allows the administrator to use switch zoning to configure each virtual port as if it was a physical port in 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 227) where its bandwidth usage can be viewed.



Variable	Value	Type	Comment	Enabled
hint.isp.0.vports	4	loader	for NPIV	true
hint.isp.1.vports	4	loader	for NPIV	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.

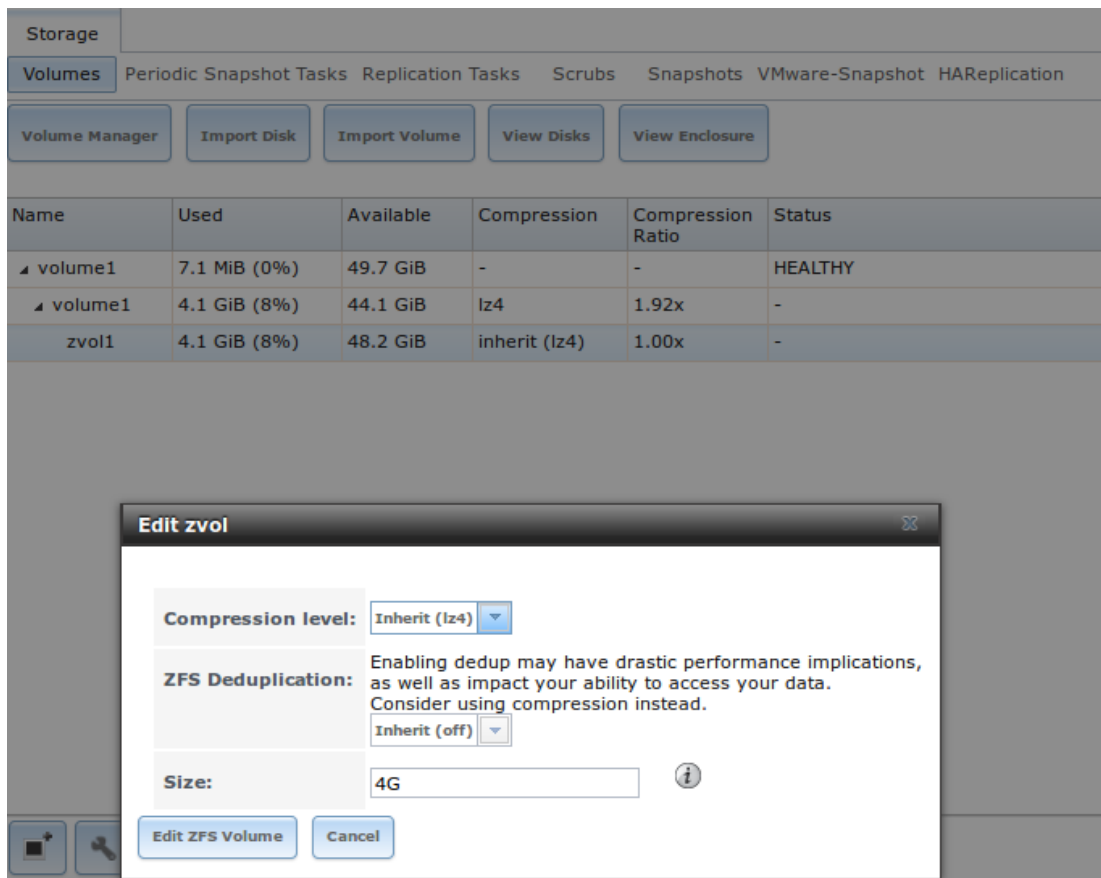


Fig. 9.32: Editing an Existing Zvol

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 193)
- *Domain Controller* (page 195)
- *Dynamic DNS* (page 197)
- *FTP* (page 198)
- *iSCSI* (page 204)
- *LLDP* (page 204)
- *NFS* (page 205)
- *Rsync* (page 206)
- *S.M.A.R.T.* (page 208)
- *SMB* (page 210)
- *SNMP* (page 215)
- *SSH* (page 217)
- *TFTP* (page 219)
- *UPS* (page 220)
- *WebDAV* (page 223)

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.

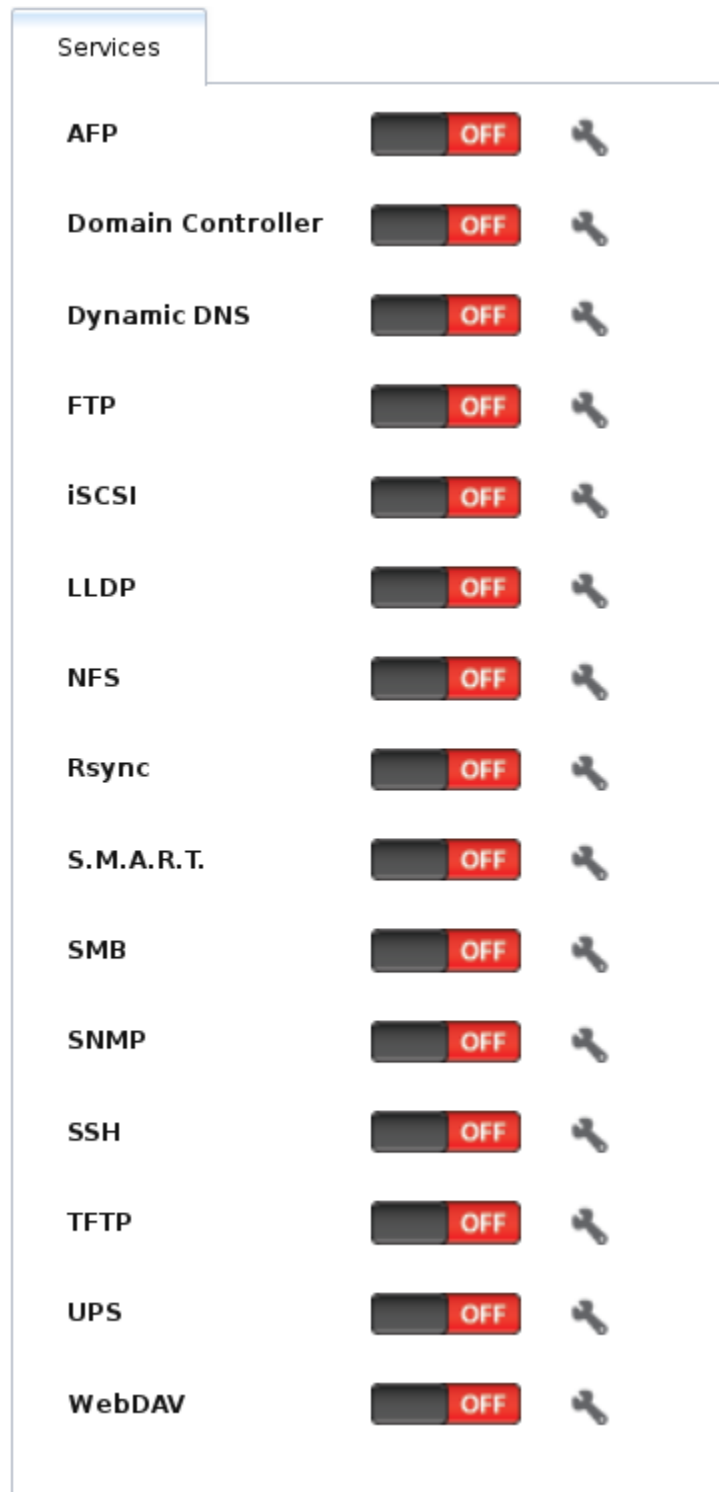


Fig. 10.1: Control Services

A service is stopped if its icon is a red *OFF*. A service is running if the icon is a blue *ON*. To start or stop a service, click the *ON/OFF* icon.

To configure a service, click the wrench icon associated with the service or click 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 will now show at the bottom of the browser. Clicking the console messages area will make it into a pop-up window, allowing scrolling through the output and copying 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 237) 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.

AFP Settings

Guest Access: ☐ ⓘ

Guest account: ▼

Max. Connections: ⓘ

Enable home directories: ☐ ⓘ

Home directories: **Browse**

Home share name: ⓘ

Database Path: **Browse** ⓘ

Global auxiliary parameters: ⓘ

Map ACLs: ⓘ

Bind IP Addresses: ⓘ

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
Enable home directories	checkbox	if checked, any user home directories located under <i>Home directories</i> will be available over the share

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Description
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
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/htmldocs/afp.conf.5.html) parameters not covered elsewhere in this screen
Map ACLs	drop-down menu	choose mapping of effective permissions for authenticated users; <i>Rights</i> (default, Unix-style permissions), <i>Mode</i> (ACLs), or <i>None</i>
Bind IP Addresses	selection	specify the IP addresses to listen for FTP connections; highlight the desired IP addresses in the <i>Available</i> list and use the >> button to add to the <i>Selected</i> 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 237), 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 130) 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 130) network to access its authentication and authorization services, configure [Active Directory](#) (page 130) instead.

Be aware that configuring a domain controller is a complex process that requires a good understanding of how [Active Directory](#) (page 130) works. While [Services → Domain Controller](#) makes it easy to input 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). 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.

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 <i>SAMBA_INTERNAL</i> is selected
Domain Forest Level	drop-down menu	choices are 2000, 2003, 2008, or 2008_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 <i>Active Directory</i> (page 130) administrator account
Kerberos Realm	drop-down menu	auto-populates with information from the <i>Realm</i> 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 237), 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 to enter will be provided by the DDNS provider. After configuring DDNS, remember to start the DDNS service in `Services → Control Services`.

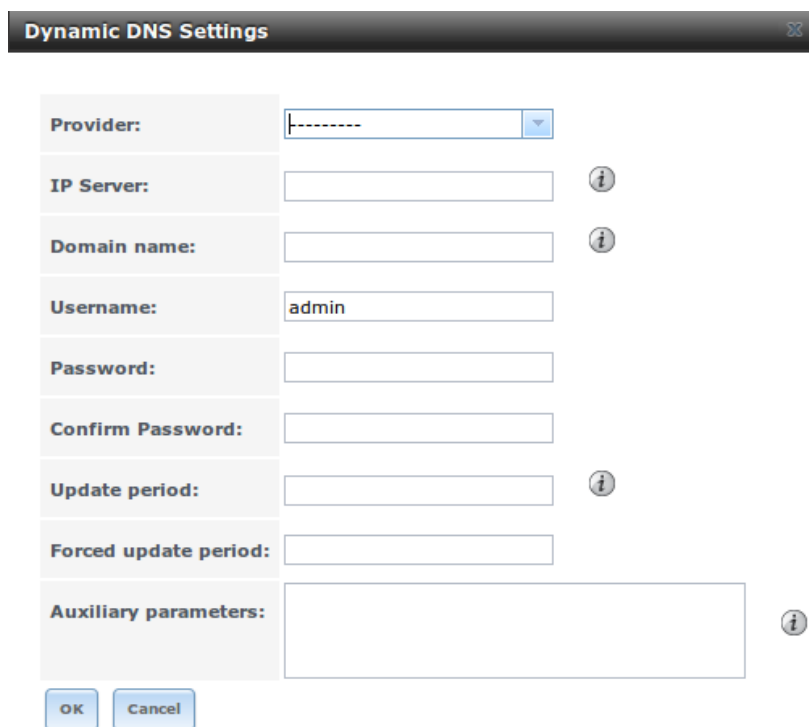


Fig. 10.4: Configuring DDNS

Table 10.3: DDNS Configuration Options

Setting	Value	Description
Provider	drop-down menu	several providers are supported; if your provider is not listed, leave this field blank and specify the custom provider in the <i>Auxiliary parameters</i> field
IP Server	string	can be used to specify the hostname and port of the IP check server
Domain name	string	fully qualified domain name (e.g. <i>yourname.dyndns.org</i>)
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
Forced update period	integer	how often the IP should be updated, even it has not changed, in seconds

Continued on next page

Table 10.3 – continued from previous page

Setting	Value	Description
Auxiliary parameters	string	additional parameters passed to the provider during record update; an example of specifying a custom provider is <i>dyn-dns_system default@provider.com</i>

When using “freedns.afraid.org”, see [this forum post](https://forums.freenas.org/index.php?threads/dynamic-dns-and-freedns-afraid-org.24455/#post-151746) (https://forums.freenas.org/index.php?threads/dynamic-dns-and-freedns-afraid-org.24455/#post-151746) for an example configuration.

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/) (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 the 'FTP Settings' window with the following configuration options:

- Port:** 21
- Clients:** 5
- Connections:** 2
- Login Attempts:** 1
- Timeout:** 600
- Allow Root Login:** ☐
- Allow Anonymous Login:** ☐
- Path:** (empty text box) **Browse** button
- Allow Local User Login:** ☐
- Display Login:** (empty text box)
- Allow Transfer Resumption:** ☐

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 <i>Path</i>
Path	browse button		root directory for anonymous FTP connections
Allow Local User Login	checkbox		required if <i>Anonymous Login</i> is disabled
Display Login	string		message displayed to local login users after authentication; not displayed to anonymous login users
Continued on next page			

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
File Permission	checkboxes	✓	sets default permissions for newly created files
Directory Permission	checkboxes	✓	sets default permissions for newly created directories
Enable FXP (https://en.wikipedia.org/wiki/File_eXchange_Protocol)	checkbox	✓	enables File eXchange Protocol which is discouraged (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 <i>wheel</i>
Require IDENT Authentication	checkbox	✓	will result in timeouts if identd 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 47)
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

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
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
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 <code>System → Certificates</code>
Auxiliary parameters	string	✓	used to add <code>proftpd(8)</code> (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 red *OFF* button next to *FTP*. After a second or so, it will change to a blue *ON*, indicating that the service has been enabled.
4. Test the connection from a client using a utility such as *Filezilla* (<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

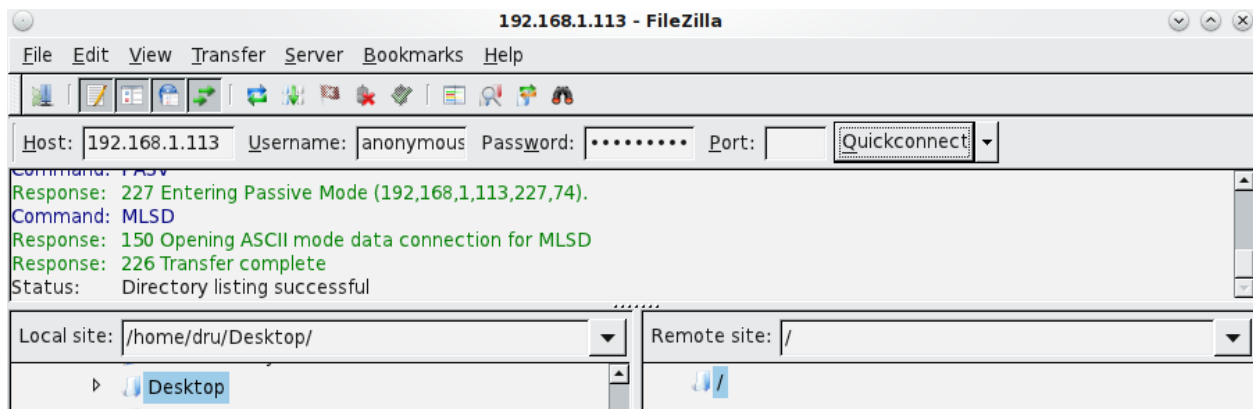


Fig. 10.6: Connecting Using Filezilla

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 130) 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 red *OFF* button next to FTP. After a second or so, it will change to a blue *ON*, indicating that the service has been enabled.
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 44). Then, import or create the certificate to use for encrypted connections using the instructions in [Certificates](#) (page 47).
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 input *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 237) 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

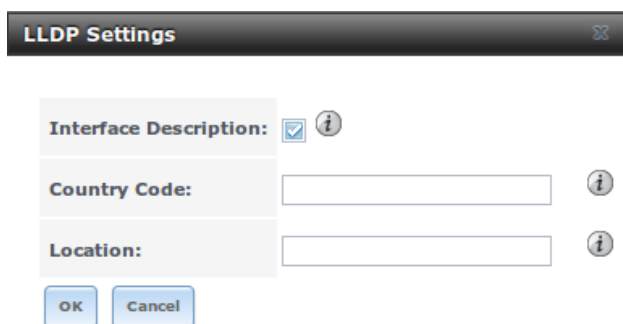
Refer to *Block (iSCSI)* (page 173) for instructions on configuring iSCSI. To start the iSCSI service, click its entry in *Services*.

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>) 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.



The screenshot shows the 'LLDP Settings' dialog box. It has a title bar with the text 'LLDP Settings' and a close button. Below the title bar, there are three configuration sections. The first section is 'Interface Description' with a checked checkbox and an information icon. The second section is 'Country Code' with an empty text input field and an information icon. The third section is 'Location' with an empty text input field and an information icon. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

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; input 2 letter ISO 3166 country code
Location	string	optional; specify the physical location of the host

10.8 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 273) for more details.

Figure 10.8 shows the configuration screen and Table 10.6 summarizes the configuration options for the NFS service.

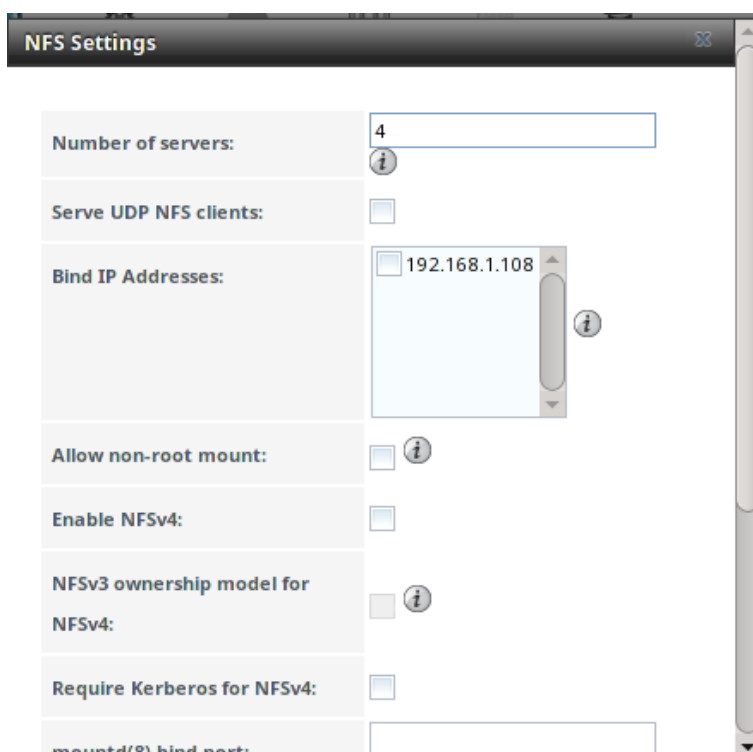


Fig. 10.8: Configuring NFS

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
Allow non-root mount	checkbox	check this box only if the NFS client requires it
Enable NFSv4	checkbox	NFSv3 is the default, check this box to switch to NFSv4
NFSv3 ownership model for NFSv4	checkbox	grayed out unless <i>Enable NFSv4</i> is checked and, in turn, will gray out <i>Support>16 groups</i> 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

Note: NFSv4 sets all ownership to *nobody:nobody* if user and group do not match on client and server.

10.9 Rsync

Services → **Rsync** is used to configure an rsync server when using rsync module mode. See the section on Rsync Module Mode for a configuration example.

This section describes the configurable options for the **rsyncd** service and rsync modules.

10.9.1 Configure Rsyncd

Figure 10.9 shows the rsyncd configuration screen which is accessed from **Services** → **Rsync** → **Configure Rsyncd**.

Configure Rsyncd

TCP Port

873

i

Auxiliary parameters

i

OK

Cancel

Fig. 10.9: 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 rsyncd to listen on, default is 873
Auxiliary pa- rameters	string	additional parameters from rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html)

10.9.2 Rsync Modules

Figure 10.10 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.

The screenshot shows a window titled "Add Rsync Module". It contains the following fields and controls:

- Module name:** A text input field.
- Comment:** A text input field.
- Path:** A text input field with a "Browse" button and an information icon to its right.
- Access Mode:** A dropdown menu currently showing "Read and Write" with an information icon to its right.
- Maximum connections:** A text input field containing the value "0" with an information icon to its right.
- User:** A dropdown menu currently showing "nobody" with an information icon to its right.
- Group:** A dropdown menu currently showing "nobody" with an information icon to its right.
- Hosts allow:** A large text input area with an information icon below it.
- Hosts deny:** A text input field.

Fig. 10.10: 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 <i>Read and Write</i> , <i>Read-only</i> , or <i>Write-only</i>
Maximum connections	integer	0 is unlimited
User	drop-down menu	select 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)

10.10 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 125).

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>) 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 Volumes`. 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 10.11 shows the configuration screen that appears after clicking `Services` → `S.M.A.R.T.`

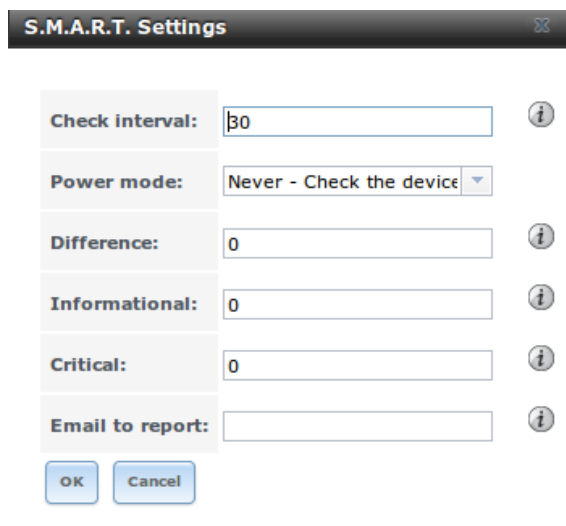


Fig. 10.11: S.M.A.R.T Configuration Options

Note: `smartd` wakes up at the configured *Check Interval*. It checks the times configured in `Tasks` → `S.M.A.R.T. Tests` to see 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 10.9 summarizes the options in the S.M.A.R.T configuration screen.

Table 10.9: S.M.A.R.T Configuration Options

Setting	Value	Description
Check interval	integer	in minutes, how often smartd 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: <i>Never</i> , <i>Sleep</i> , <i>Standby</i> , or <i>Idle</i>
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

10.11 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) (<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 10.12 shows the global SMB configuration options which are described in Table 10.10. This configuration screen is really a front-end to [smb4.conf](https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+10.3-RELEASE+and+Ports) (<https://www.freebsd.org/cgi/man.cgi?query=smb.conf&manpath=FreeBSD+10.3-RELEASE+and+Ports>).

SMB

NetBIOS name:

truenas

NetBIOS alias:

Workgroup:

WORKGROUP

Description:

TrueNAS Server

DOS charset:

CP437

UNIX charset:

UTF-8

Log level:

Minimum

Use syslog only:

☐

Local Master:

☒

Domain logons:

☐

Time Server for Domain:

☒

Guest account:

nobody

Fig. 10.12: Global SMB Configuration

Table 10.10: 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 <i>Workgroup</i> name
NetBIOS Name (Node B)	string	limited to 15 characters; when using Failover (page 53), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	limited to 15 characters; when using Failover (page 53), 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 130) or LDAP (page 136) service is running
Continued on next page		

Table 10.10 – continued from previous page

Setting	Value	Description
Description	string	optional
DOS charset	drop-down menu	the character set Samba uses when communicating with DOS and Windows 9x/ME clients; default is <i>CP437</i>
UNIX charset	drop-down menu	default is <i>UTF-8</i> which supports all characters in all languages
Log level	drop-down menu	choices are <i>Minimum</i> , <i>Normal</i> , or <i>Debug</i>
Use syslog only	checkbox	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 <i>nobody</i> ; account must have permission to access the shared volume/dataset; if Guest Account user is deleted, resets to <i>nobody</i>
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 <i>Hosts Allow</i> or <i>Hosts Deny</i> fields of a SMB share; uncheck if IP addresses are used to avoid the delay of a host lookup
Server minimum protocol	drop-down menu	the minimum protocol version the server will support where the default sets automatic negotiation; refer to Table 10.11 for descriptions
Server maximum protocol	drop-down menu	the maximum protocol version the server will support; refer to Table 10.11 for descriptions
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 130) users and groups to domain admins on another forest

Continued on next page

Table 10.10 – continued from previous page

Setting	Value	Description
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

Table 10.11: SMB Protocol Versions

Value	Description
CORE	used by DOS
COREPLUS	used by DOS
LANMAN1	used by Windows for Workgroups, OS/2, and Windows 9x
LANMAN2	used by Windows for Workgroups, OS/2, and Windows 9x
NT1	used by Windows NT
SMB2	used by Windows 7; same as SMB2_10
SMB2_02	used by Windows Vista
SMB2_10	used by Windows 7
SMB3	used by Windows 8
SMB3_00	used by Windows 8
SMB3_02	used by Windows 8.1 and Windows Server 2012
SMB3_11	used by Windows 10

Changes to SMB settings and SMB shares take effect immediately.

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 210) cannot be disabled while [Active Directory](#) (page 130) is enabled.

10.11.1 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) (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 237) 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 237) for usage instructions.

The [Common Errors](http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/domain-member.html#id2573692) (<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) (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.**

- Use at least the *SMB2* version of the protocol when possible. Enable this on the client if possible. The default settings for *Server minimum protocol* (—) and *Server maximum protocol* (*SMB3*) in the [global SMB service options](#) (page 211) allow clients to connect and negotiate higher and faster levels of the protocol. If these have been changed from the default, they might reduce performance. Note that

Windows XP does not support SMB2, so it is particularly important to leave *Server minimum protocol* at the default on networks with XP clients.

- *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 211).
- 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, 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 164).
- Disable as many *VFS Objects* as possible in the *share settings* (page 164). Many have performance overhead.

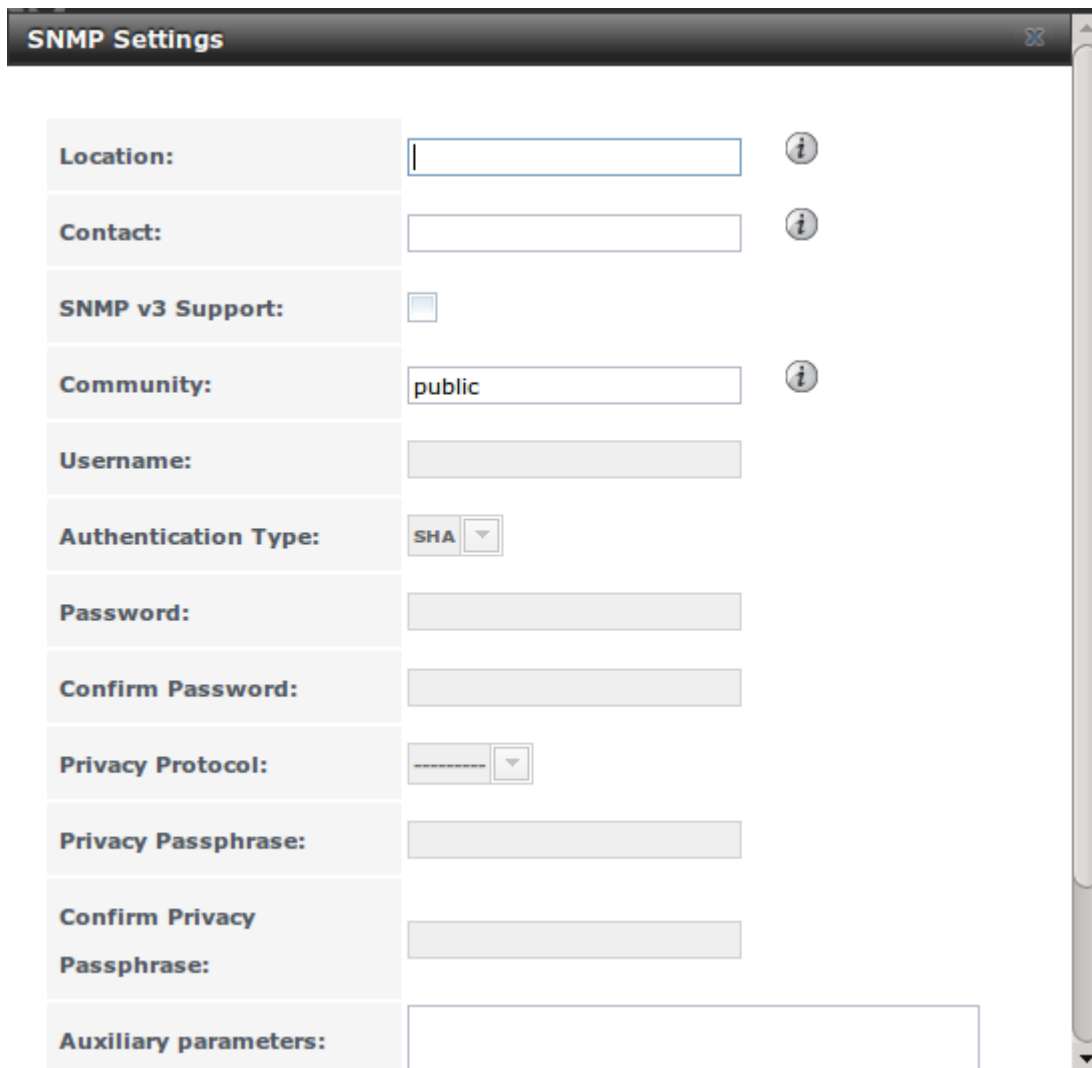
10.12 SNMP

SNMP (Simple Network Management Protocol) is used to monitor network-attached devices for conditions that warrant administrative attention. TrueNAS® uses *Net-SNMP* (<http://net-snmp.sourceforge.net/>) to provide SNMP. When 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 10.13 shows the SNMP configuration screen. Table 10.12 summarizes the configuration options.



The image shows a 'SNMP Settings' window with a dark title bar. It contains a list of configuration fields on the left and their corresponding input controls on the right. The fields are: Location (text input), Contact (text input), SNMP v3 Support (checkbox), Community (text input with 'public' entered), Username (text input), Authentication Type (dropdown menu showing 'SHA'), Password (text input), Confirm Password (text input), Privacy Protocol (dropdown menu), Privacy Passphrase (text input), Confirm Privacy Passphrase (text input), and Auxiliary parameters (text input). Information icons (i) are present next to Location, Contact, and Community. A vertical scrollbar is on the right side of the window.

Fig. 10.13: 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 <i>public</i> and should be changed for security reasons ; can only contain alphanumeric characters, underscores, dashes, periods, and spaces; this value can be empty for SNMPv3 networks
Continued on next page		

Table 10.12 – continued from previous page

Setting	Value	Description
Username	string	only applies if <i>SNMP v3 Support</i> is checked; specify the user-name 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 <i>Authentication Type</i> , <i>Password</i> , <i>Privacy Protocol</i> , and “Privacy Passphrase” fields
Authentication Type	drop-down menu	only applies if <i>SNMP v3 Support</i> is checked; choices are <i>MD5</i> or <i>SHA</i>
Password	string	only applies if <i>SNMP v3 Support</i> is checked; specify and confirm a password of at least eight characters
Privacy Protocol	drop-down menu	only applies if <i>SNMP v3 Support</i> is checked; choices are <i>AES</i> or <i>DES</i>
Privacy Passphrase	string	if not specified, <i>Password</i> is used
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

10.13 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 10.14 shows the `Services → SSH` configuration screen. After configuring SSH, remember to start it in `Services → Control Services`.

Fig. 10.14: 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 <i>Available</i> field and added to the <i>Selected</i> 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 <i>root</i> user in <i>View Users</i>
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 142) 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
Extra Options	string	✓	additional sshd_config(5) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options not covered in this screen, one per line; these options are case-sensitive and misspellings may prevent the SSH service from starting

A few [sshd_config\(5\)](#) (http://www.freebsd.org/cgi/man.cgi?query=sshd_config) options that are useful to enter in the *Extra Options* field include:

- increase the *ClientAliveInterval* if SSH connections tend to drop
- *ClientMaxStartup* defaults to 10; increase this value if you need more concurrent SSH connections

10.13.1 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**.

10.13.2 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 237) 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`.

10.14 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 10.15 shows the TFTP configuration screen and Table 10.14 summarizes the available options:

Fig. 10.15: 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 <i>Directory</i>
Umask	integer	umask for newly created files, default is 022 (everyone can read, nobody can write); some devices require a less strict umask
Extra options	string	additional tftpd(8) (http://www.freebsd.org/cgi/man.cgi?query=tftpd) options not shown in this screen, one per line

10.15 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 10.16 shows the UPS configuration screen:

UPS 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

Monitor User:

upsmon

Monitor Password:

fixmepass

Extra users (upsd.users):

Remote Monitor:

☐

Send Email Status Updates:

☐

To email:

Email Subject:

UPS report generated by %F

Power Off UPS:

☐

OK

Cancel

Fig. 10.16: 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 <i>Master</i> or <i>Slave</i>
Continued on next page		

Table 10.15 – continued from previous page

Setting	Value	Description
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 <i>UPS goes on battery</i> and <i>UPS reaches low battery</i>
Shutdown timer	integer	in seconds; will initiate shutdown after this many seconds after UPS enters <i>UPS goes on battery</i> , 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
Monitor User	string	default is <i>upsmon</i>
Monitor Password	string	default is known value <i>fixmepass</i> 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 <i>upsmon</i> and password <i>fixmepass</i>
Send Email Status Updates	checkbox	if checked, activates the <i>To email</i> field
To email	email address	if <i>Send Email</i> 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 console messages show the name of the */dev/ugenX.X* device, where the X's are the numbers that show on the console.

[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 adminis-

trative rights can use this command. These users are created in the *Extra users* field.

10.15.1 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).

10.16 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 161) 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 10.17](#) shows the WebDAV configuration screen. [Table 10.16](#) summarizes the available options.

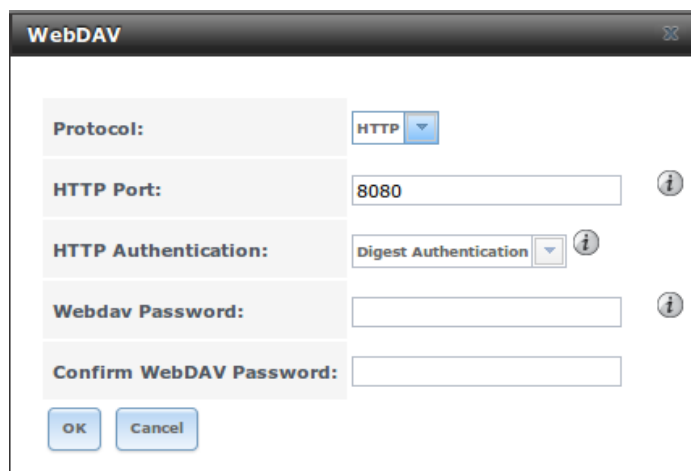
The image shows a 'WebDAV' configuration window. It has a title bar with the text 'WebDAV' and a close button. Inside the window, there are five labeled input fields: 'Protocol:' with a dropdown menu showing 'HTTP'; 'HTTP Port:' with a text box containing '8080'; 'HTTP Authentication:' with a dropdown menu showing 'Digest Authentication'; 'Webdav Password:' with an empty text box; and 'Confirm WebDAV Password:' with an empty text box. To the right of the 'HTTP Port:', 'HTTP Authentication:', and 'Webdav Password:' fields are small circular information icons. At the bottom left of the window are 'OK' and 'Cancel' buttons.

Fig. 10.17: WebDAV Configuration Screen

Table 10.16: WebDAV Configuration Options

Setting	Value	Description
Protocol	drop-down menu	choices are <i>HTTP</i> (connection always unencrypted), <i>HTTPS</i> (connection always encrypted), or <i>HTTP+HTTPS</i> (both types of connections allowed)
HTTP Port	string	only appears if the selected <i>Protocol</i> is <i>HTTP</i> or <i>HTTP+HTTPS</i> and is used to specify the port to be used for unencrypted connections; the default of <i>8080</i> should work, if you change it, do not use a port number already being used by another service

Continued on next page

Table 10.16 – continued from previous page

Setting	Value	Description
HTTPS Port	string	only appears if the selected <i>Protocol</i> is <i>HTTPS</i> or <i>HTTP+HTTPS</i> and is used to specify the port to be used for encrypted connections; the default of <i>8081</i> 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 <i>Protocol</i> is <i>HTTPS</i> or <i>HTTP+HTTPS</i> ; select the SSL certificate to be used for encrypted connections; to create a certificate, use <i>System</i> → <i>Certificates</i>
HTTP Authentication	drop-down menu	choices are <i>Basic Authentication</i> (unencrypted) or <i>Digest Authentication</i> (encrypted)
Webdav Password	string	default is <i>davtest</i> ; this should be changed as it is a known value

VCENTER

Administrators who use [VMware vCenter Server](https://www.vmware.com/products/vcenter-server) (<https://www.vmware.com/products/vcenter-server>) to administer their vSphere environments can use the TrueNAS® vCenter plugin to manage their TrueNAS® array from vCenter Server.

Note: At this time, the vCenter plugin only supports the creation of iSCSI and NFS datastores from vCenter Server.

Note: The current vCenter plugin works with versions of vCenter up to vCenter 6.

To configure the vCenter plugin, click *vCenter*. This opens the screen shown in [Figure 11.1](#).

The screenshot displays the 'vCenter Plugin Configuration' window. It features a table-like structure with the following fields and values:

Plugin Name:	TrueNAS vCenter Plugin
Available Plugin Version:	2.0.0
Installed Plugin Version:	2.0.0
TrueNAS Management IP Address:	<input type="text"/> <input style="border: none; background-color: #ccc; padding: 2px 5px;" type="button" value="v"/>
vCenter Hostname/IP Address:	<input style="background-color: #eee;" type="text" value="vcenter-01"/>
vCenter Port:	<input type="text" value="443"/>
vCenter Username:	<input style="background-color: #eee;" type="text" value="root"/>
vCenter Password:	<input style="background-color: #eee;" type="password" value="password"/>

At the bottom of the configuration area, there are four buttons: 'Install' (disabled), 'Uninstall' (active), 'Upgrade' (disabled), and 'Repair' (active).

Fig. 11.1: Configuring the vCenter Plugin

Table 11.1 summarizes the options in 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	input the IP address or resolveable hostname of the vCenter Server
vCenter Port	integer	input the port number the vCenter Server is listening on
vCenter Username	string	input the username for the vCenter Server
vCenter Password	string	input the password associated with <i>vCenter Username</i>

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 will be incremented. Click this button to upgrade to the newer version and access its features.

Repair: click this button if your iXsystems support engineer requests it. This reinstalls the TrueNAS® object to repair a corrupted object.

Note: In a HA-configured scenario, the *Upgrade* button can only be used from the system that originally installed the plugin. The *Upgrade* button is grayed out on the other system in the HA pair.

To configure the vCenter plugin for a secure connection, click *vCenter* → *vCenter Auxiliary Settings* in the left tree. In the screen shown in Figure 11.2, check the *Enable vCenter Plugin over https* box.

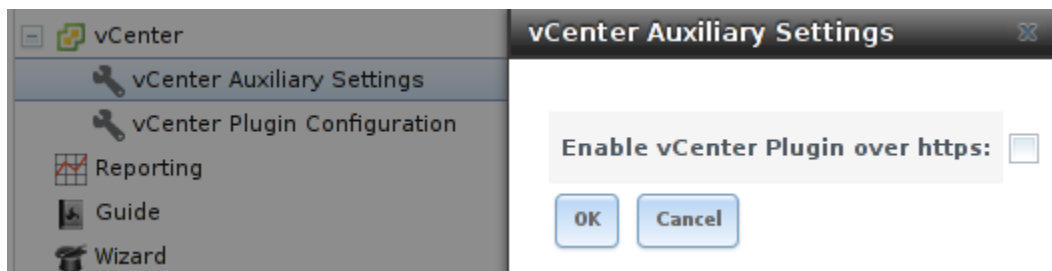


Fig. 11.2: Securing the vCenter Plugin Connection

REPORTING

Reporting displays several graphs, as seen in the example in [Figure 12.1](#). Click the tab for a device type to see its graphs.

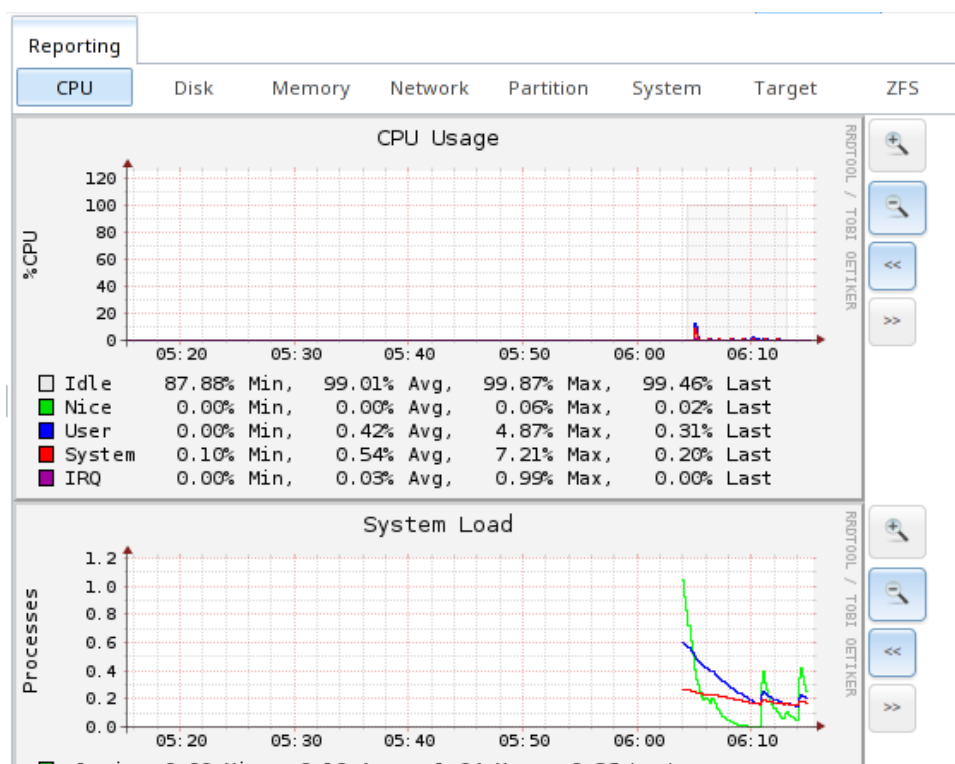


Fig. 12.1: Reporting Graphs

TrueNAS® uses [collectd](https://collectd.org/) (<https://collectd.org/>) to provide reporting statistics. `collectd` plugins are enabled in `/conf/base/etc/local/collectd.conf` to provide reporting graphs. These graphs are grouped into several tabs on the Reporting page:

- *CPU*
 - [CPU](https://collectd.org/wiki/index.php/Plugin:CPU) (<https://collectd.org/wiki/index.php/Plugin:CPU>) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle.
- *Disk*
 - [Disk](https://collectd.org/wiki/index.php/Plugin: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>) shows received and transmitted traffic in bits per second for each configured interface.
 - *Partition*
 - **Disk space** (<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>) displays the number of processes, grouped by state.
 - **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) shows ARC size, hit ratio, and requests.

Reporting data is saved, making it possible to view and monitor usage trends over time. By default, reporting data is saved to `/data/rrd_dir.tar.bz2` and is preserved across system upgrades and at shutdown. To instead save this data to the system dataset, check the *Reporting database* box in *System* → *System Dataset*.

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>) contains instructions for sending the collected information to a **Graphite** (<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 13.1 shows the first wizard configuration screen.

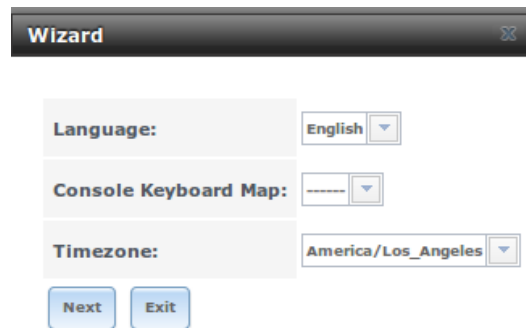


Fig. 13.1: Configuration Wizard

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 13.2 and Figure 13.3 will only appear if unformatted disks are available or the system has been reinstalled.

Figure 13.2 shows the configuration screen that appears if the storage disks have not yet been formatted.

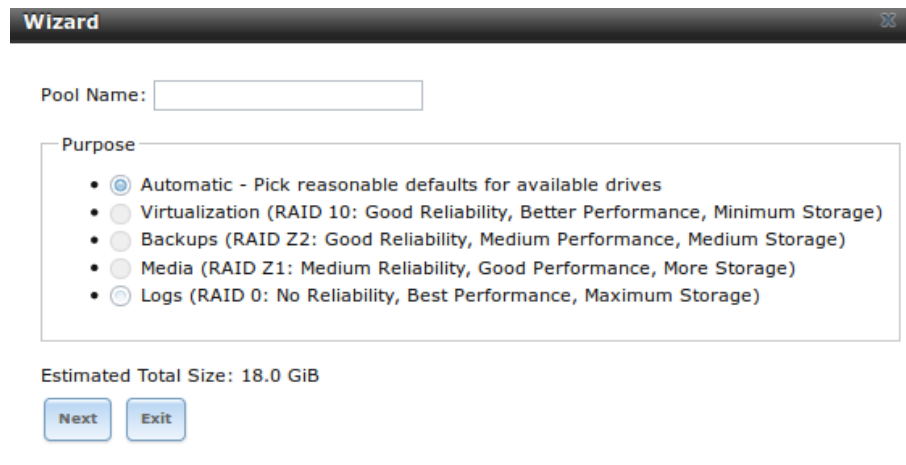


Fig. 13.2: Volume Creation Wizard

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 100) 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 242) 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 87).

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 13.3](#).



Fig. 13.3: 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 13.4](#).

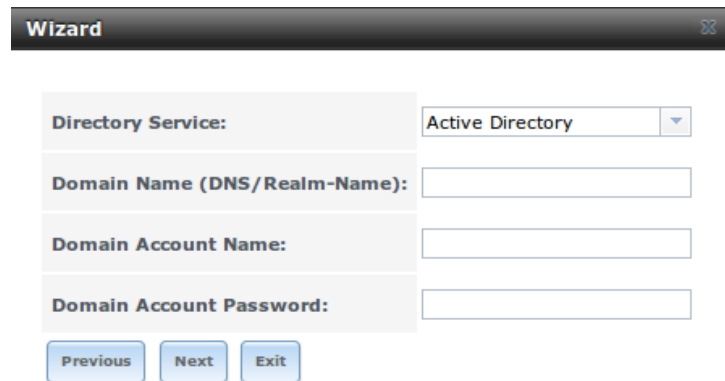


Fig. 13.4: Directory Service Selection

If the TrueNAS® system is on a network that does not contain an Active Directory, LDAP, NIS, or NT4 server, click *Next* to skip to the next screen.

However, if the TrueNAS® system is on a network containing an Active Directory, LDAP, NIS, or NT4 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 13.1 through 13.4](#).

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 130) to determine if additional configuration is required.

Table 13.1: Active Directory Options

Setting	Value	Description
Domain Name	string	name of Active Directory domain (e.g. <i>example.com</i>) or child domain (e.g. <i>sales.example.com</i>)
Domain Account Name	string	name of the Active Directory administrator account
Domain Account Password	string	password for the Active Directory administrator account

Table 13.2: LDAP Options

Setting	Value	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. <i>dc=test,dc=org</i>)
Bind DN	string	name of administrative account on LDAP server (e.g. <i>cn=Manager,dc=test,dc=org</i>)
Base password	string	password for

Table 13.3: NIS Options

Setting	Value	Description
NIS domain	string	name of NIS domain
NIS servers	string	comma delimited list of hostnames or IP addresses
Secure mode	checkbox	if checked, ypbind(8) (http://www.freebsd.org/cgi/man.cgi?query=ybind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024
Manycast	checkbox	if checked, ybind will bind to the server that responds the fastest; this is useful when no local NIS server is available on the same subnet

Table 13.4: NT4 Options

Setting	Value	Description
Domain Controller	string	hostname of domain controller
NetBIOS Name	string	hostname of TrueNAS [®] system; cannot be greater than 15 characters or the same as the <i>Workgroup Name</i>
Workgroup Name	string	name of Windows server's workgroup
Administrator Name	string	name of the domain administrator account
Administrator Password	string	input and confirm the password for the domain administrator account

The next configuration screen, shown in [Figure 13.5](#), is used to create network shares.

Fig. 13.5: 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 144).

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 162).
- **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 145).
- **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 153).
- **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 204).

After selecting the *Purpose*, click the *Ownership* button to see the screen shown in [Figure 13.6](#).

Wizard

User: root ☐ Create User ⓘ

Group: wheel ☐ Create Group ⓘ

Mode:

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

Fig. 13.6: 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 13.7](#).

Wizard

Console messages: ☐ ⓘ

Root E-mail: ⓘ

From email: root@freenas.local ⓘ

Outgoing mail server: ⓘ

Port to connect to: 25 ⓘ

TLS/SSL: Plain ⓘ

Use SMTP Authentication: ☐

Username: ⓘ

Password:

Password confirmation: ⓘ

Fig. 13.7: 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.

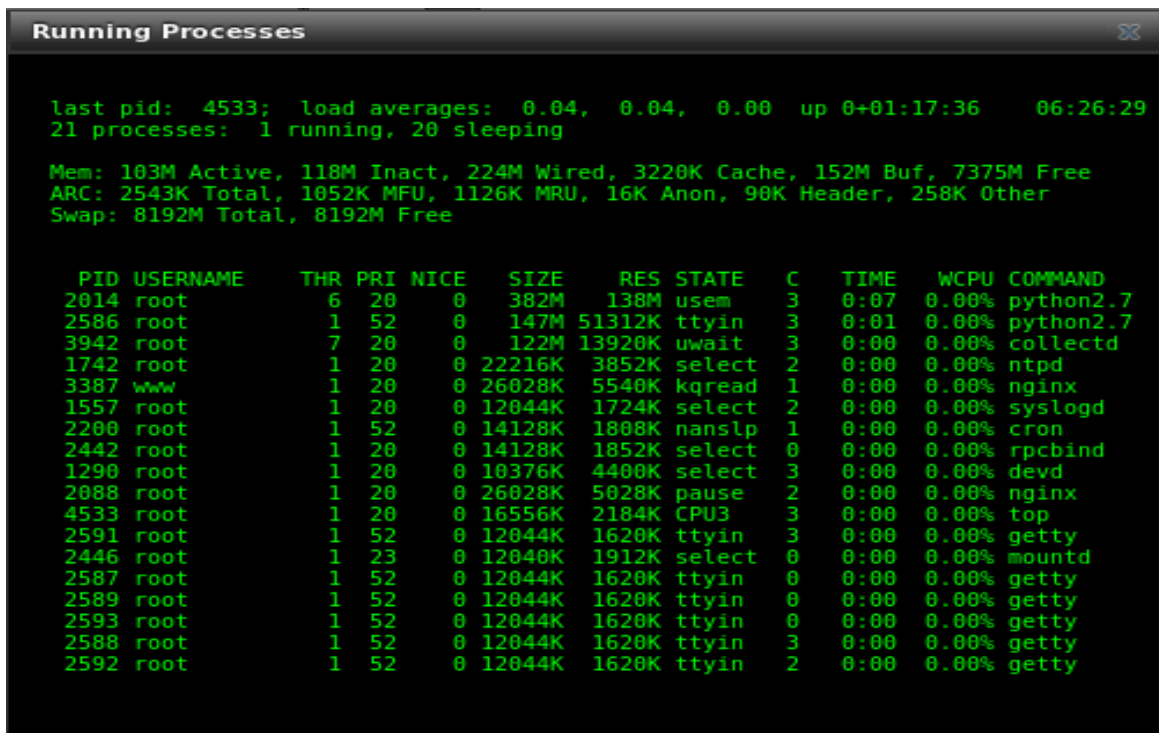
In addition to the settings that you specify, the wizard will automatically enable [S.M.A.R.T. Tests](#) (page 71), create a boot environment, and add the new boot environment to the boot menu. If you also wish to save a backup of the configuration database to the system being used to access the administrative 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.

14.1 Display System Processes

If you click Display System Processes, a screen will open showing the output of `top(1)` (<http://www.freebsd.org/cgi/man.cgi?query=top>). An example is shown in Figure 14.1.



```
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. 14.1: System Processes Running on TrueNAS®

The display will automatically refresh itself. Simply click the X in the upper right corner to close the display when you are finished. Note that the display is read-only, meaning that you won't be able to issue a `kill` command within it.

14.2 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 14.2](#), the link has been clicked and Shell is open.

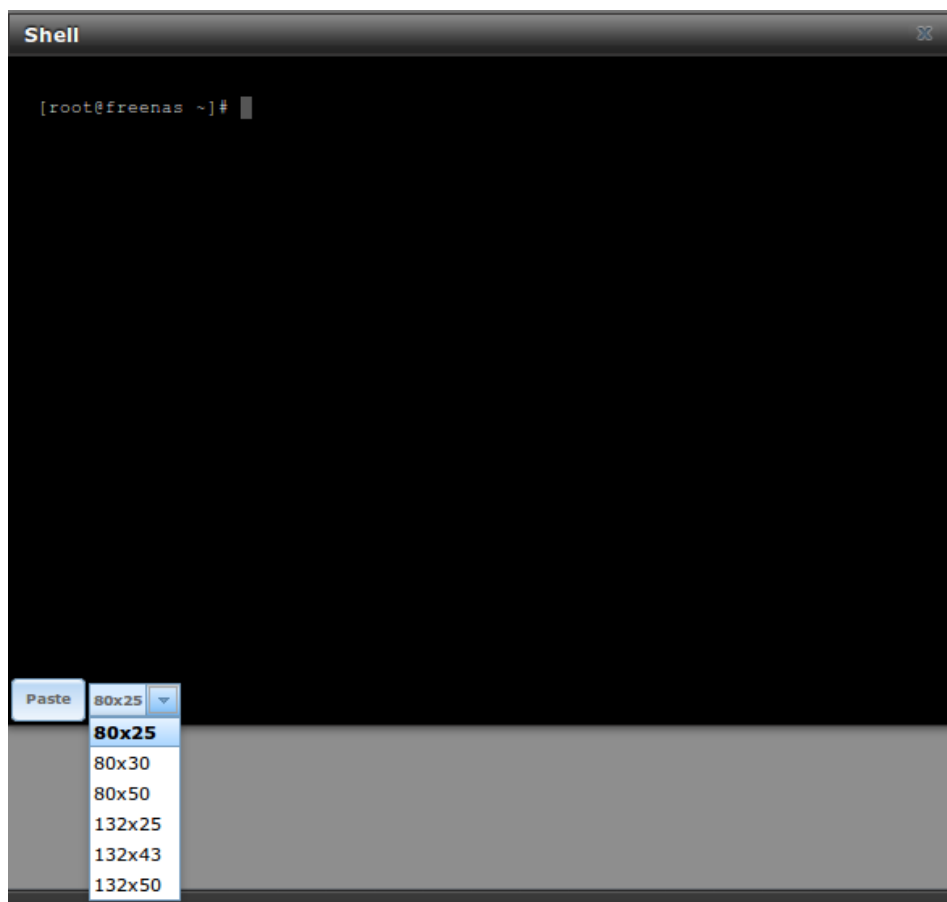


Fig. 14.2: Web Shell

The prompt indicates that the current user is *root*, the hostname is *trueenas*, 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.

14.3 Log Out

To log out of the TrueNAS® GUI, click the *Log Out* entry in the tree. You will immediately be logged out. An informational message will indicate that you are logged out and will provide a hyperlink which you can click on to log back in. When logging back in, you will be prompted for the *root* password.

14.4 Reboot

If you click *Reboot*, you will receive the warning message shown in Figure 14.3 and your browser window color will change to red to indicate that you have selected an option that will negatively impact users of the TrueNAS® system.

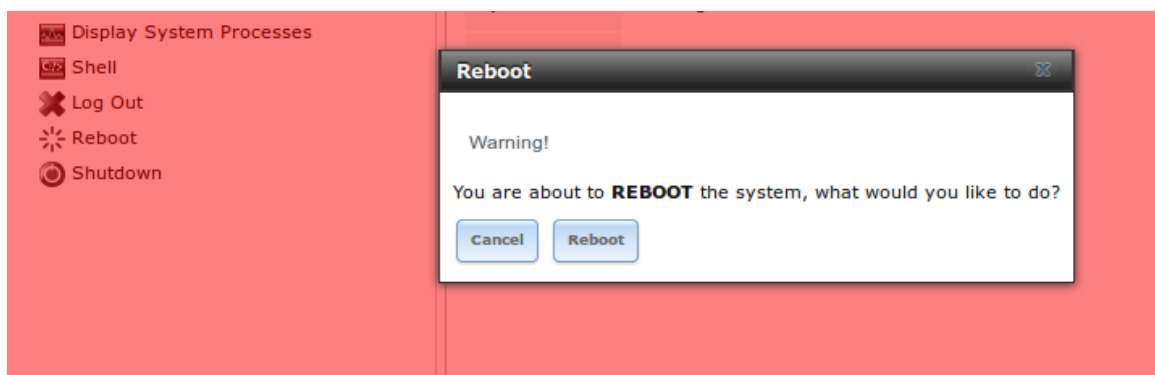


Fig. 14.3: Reboot Warning Message

If a scrub or resilver is in progress when a reboot is requested, an additional warning will ask you to make sure that you wish to proceed. In this case, it is recommended to *Cancel* the reboot request and to periodically run `zpool status` from *Shell* (page 237) 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 if you wish to cancel the reboot request. Otherwise, click the *Reboot* button to reboot the system. Rebooting the system will disconnect all clients, including the web administration GUI. The URL in your web browser will change to add `/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. If all goes well, the GUI login screen is displayed. If the login screen does not appear, access the system using IPMI in order to determine what problem is preventing the system from resuming normal operation.

14.5 Shutdown

If you click *Shutdown*, you will receive the warning message shown in Figure 14.4 and your browser color will change to red to indicate that you have selected an option that will negatively impact users of the TrueNAS®

system.



Fig. 14.4: Shutdown Warning Message

If a scrub or resilver is in progress when a shutdown is requested, an additional warning will ask you to make sure that you wish to proceed. In this case, it is recommended to *Cancel* the shutdown request and to periodically run `zpool status` from [Shell](#) (page 237) until it is verified that the scrub or resilver process is complete. Once complete, the shutdown request can be re-issued.

Click the *Cancel* button to cancel the shutdown request. Otherwise, click the *Shutdown* button to halt the system. Shutting down the system will disconnect all clients, including the web administration GUI, and will power off the TrueNAS® system.

14.6 Support Icon

The *Support* icon, located as the third icon from the left in the top 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.

14.7 Guide

The *Guide* icon, located as the second icon from the left in the top menubar, provides a built-in browser to the TrueNAS® Administrator Guide (this documentation).

14.8 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 14.5](#), the system is warning that the S.M.A.R.T. service is not running.

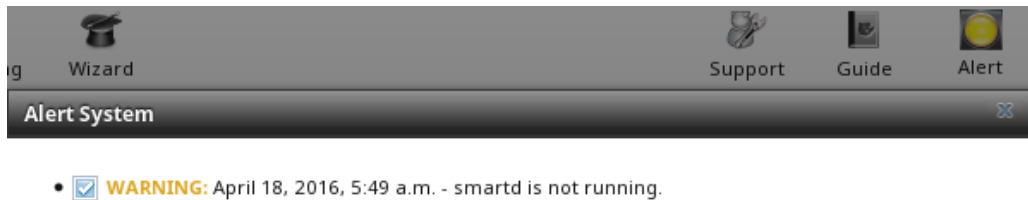


Fig. 14.5: Example Alert Message

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 you deal with it, 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 will change the solid green alert icon to flashing red if 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 14.2) by running `alertcli.py`. This can be useful when the alert originates from the standby node of a *High Availability (HA)* (page 53) system.

Some of the conditions that trigger an alert include:

- used space on a volume goes over 80%
- new OpenZFS feature flags are available for the pool; this alert can be unchecked if you choose not to upgrade the pool at this time
- a new update is available
- non-optimal multipath states 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 128) 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
- if LDAP failed to bind to the domain
- 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 52) is enabled but any of the configuration fields are empty
 - if VMware failed to log in (usually preceding a VMware snapshot)

Note: Alerts which could be related to a hardware issue automatically create a support ticket if the system is connected to the internet. These include a ZFS pool status change, a multipath failure, a failed S.M.A.R.T. test, and a failed re-key operation.

An alert is also generated when the Avago HBA firmware version does not match the driver version. To resolve this alert, download the IT (integrated target) firmware, not the IR (integrated RAID) firmware, from the Avago website. Specify the name of the firmware image and BIOS as well as the controller to flash:

```
sas2flash -o -f firmwareimagename -b biosname -c controllernumber
```

When finished, reboot the system. The new firmware version will appear in the system messages and the alert will be cleared.

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 in [View Volumes](#) (page 103). 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 87) 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 237) 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 237)

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 a adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of *random* reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 64 GB of RAM, and the size of an L2ARC should not exceed five 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. See [ZFS RAIDZ stripe width, or: How I Learned to Stop Worrying and Love RAIDZ](http://blog.delphix.com/matt/2014/06/06/zfs-stripe-width/) (<http://blog.delphix.com/matt/2014/06/06/zfs-stripe-width/>) for details.

These resources can also help determine the RAID configuration best suited to your storage needs:

- [Getting the Most out of ZFS Pools](https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/) (<https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/>)

-
- [A Closer Look at ZFS, Vdevs and Performance](http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance) (<http://constantin.glez.de/blog/2010/06/closer-look-zfs-vdevs-and-performance>)

Warning: RAID AND DISK REDUNDANCY ARE NOT A SUBSTITUTE FOR A RELIABLE BACKUP STRATEGY. BAD THINGS HAPPEN AND A GOOD BACKUP STRATEGY IS STILL REQUIRED TO PROTECT VALUABLE DATA. See [Periodic Snapshot Tasks](#) (page 112) and [Replication Tasks](#) (page 114) 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>)

HARDWARE SETUP

TrueNAS® hardware consists of one or two main *Unified Storage Array* units. Optional *Expansion Shelves* can be added to expand storage capacity. Racking and connection of these units is described individually below.

If the TrueNAS® Storage Array does not arrive in perfect condition, or if any parts are missing, please take photos and contact iXsystems support immediately.

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

16.1 TrueNAS® Unified Storage Array

The TrueNAS® Storage Array is shipped with several accessories. Please verify that the shipment includes these items:

- TrueNAS® Storage Array



- Up to 16 Populated 3.5" drive trays



- One pair of outer rails, left and right



- Eight thumbscrews



- Two short screws



- Two long screws



- Two power cables



- One serial to 3.5mm cable



- One faceplate



- One printed guide



Network cables are highly configuration-dependent. Please contact your iXsystems Sales Representative for any questions about the included cables.

Unused drive bays are populated with drive tray blanks to maintain proper airflow.

Layout of the storage controller varies with configuration. [Figure 16.1](#) provides an example of the front view of the TrueNAS® Storage Array.

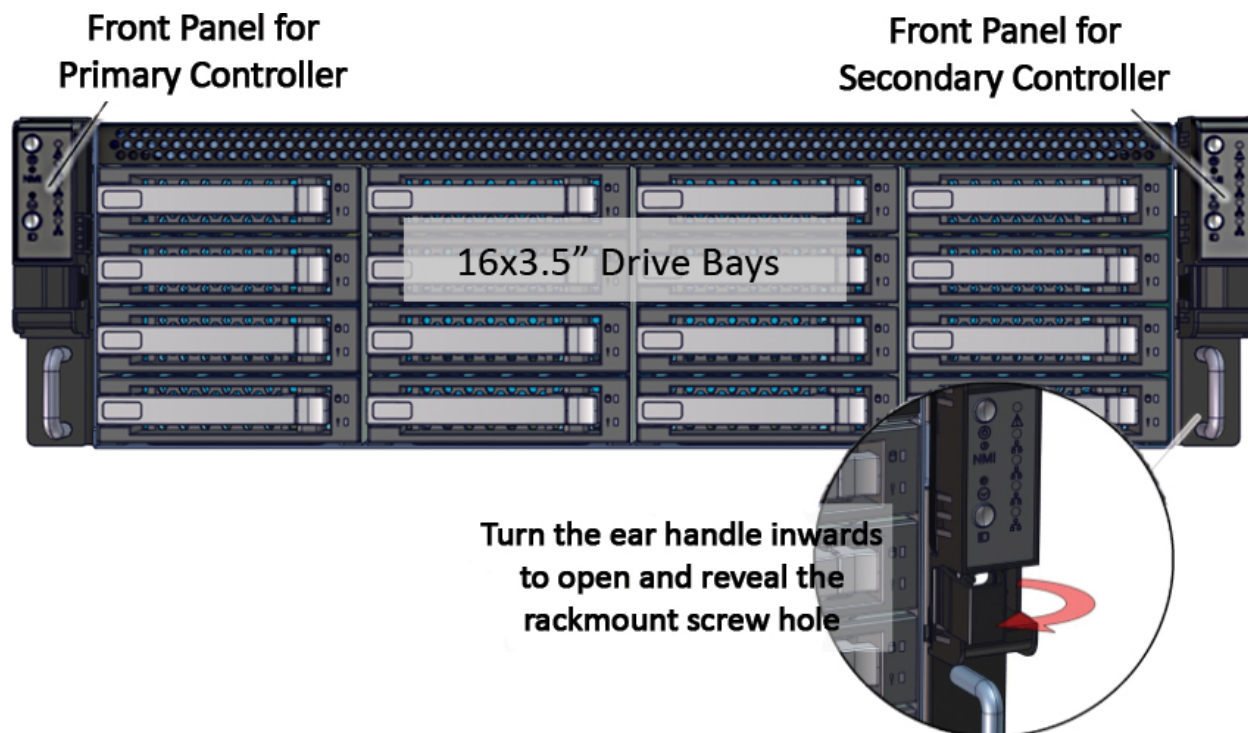


Fig. 16.1: Front View

There are two control panels, one on each side of the front of the array. The one on the left controls the primary storage controller, and the one on the right controls the secondary storage controller in High Availability models.

[Figure 16.2](#) shows the layout of the front panel buttons and indicators.

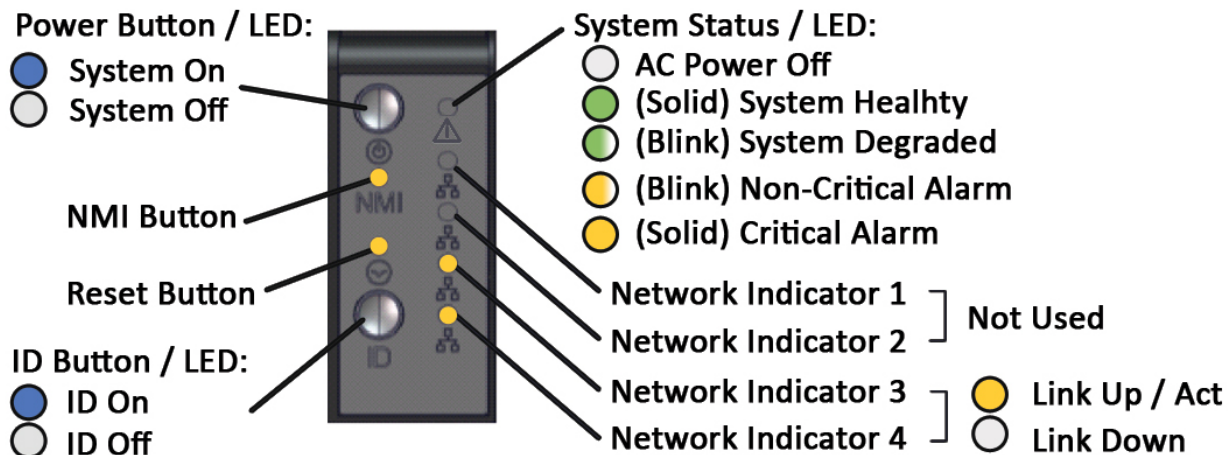


Fig. 16.2: Front Panel Buttons and Indicators

Figure 16.3 shows the rear view of the array. If the TrueNAS® Storage Array is configured for High Availability, both storage controller slots are populated. In a single-controller model, the bottom slot is covered with a blank panel.

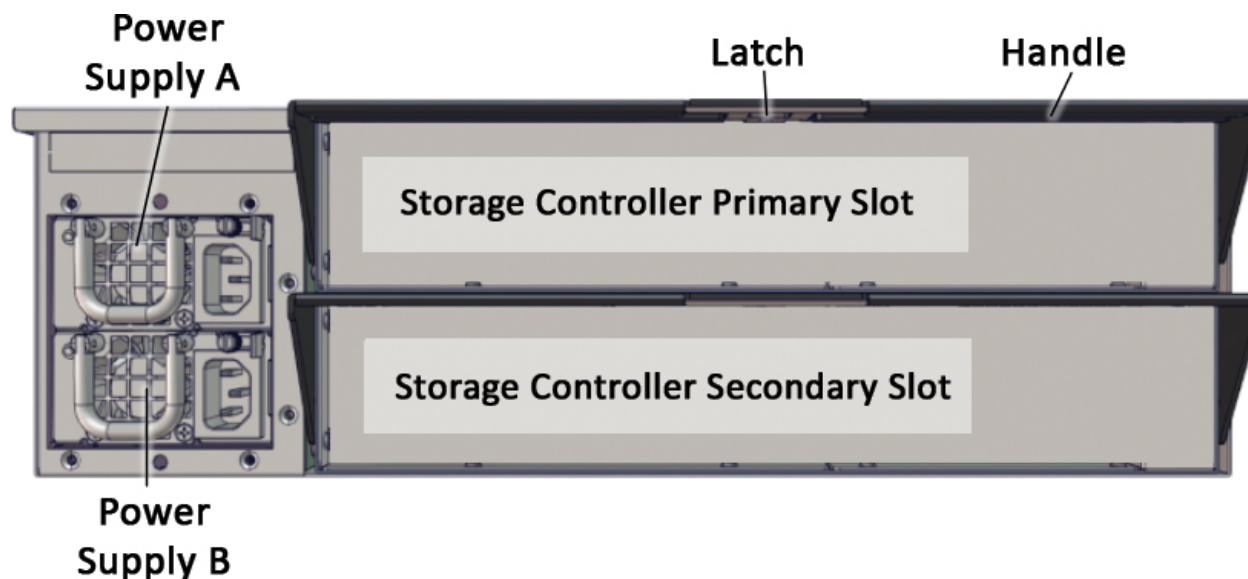


Fig. 16.3: Rear View

Figure 16.4 shows a drive tray and the LED color indications.

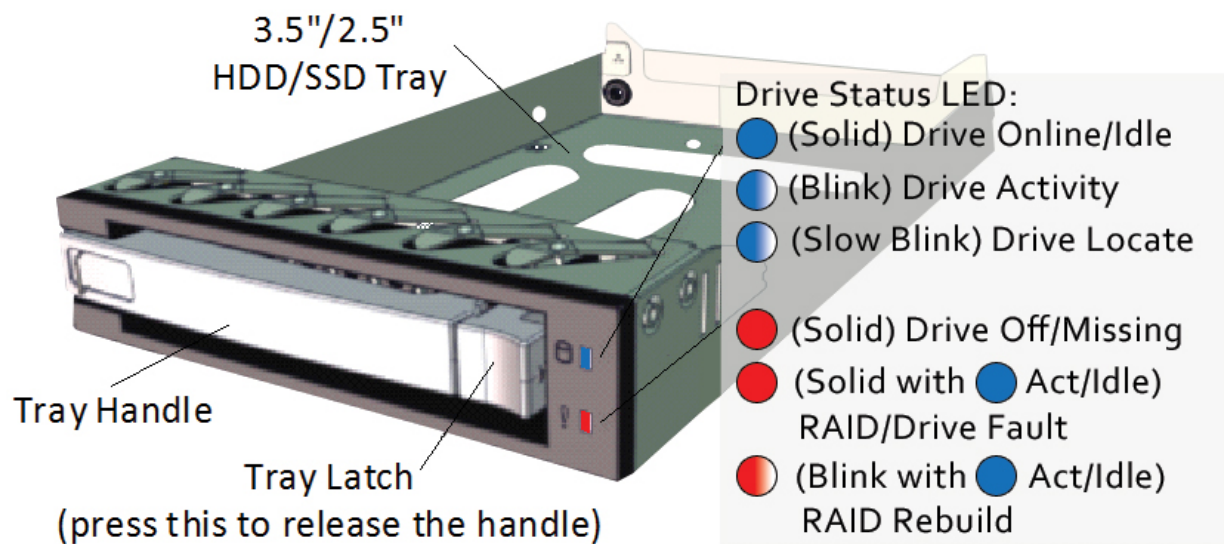


Fig. 16.4: Drive Tray

16.2 Hardware Installation

The TrueNAS® Storage Array slide rails work on racks with either square or circular hole types. Set the mounting brackets into the correct position for the rack type by pressing the button on the mounting bracket and rotating them, as shown in [Figure 16.5](#). The square rack style brackets are the default. The circular hole style has a flat surface and screw holes.

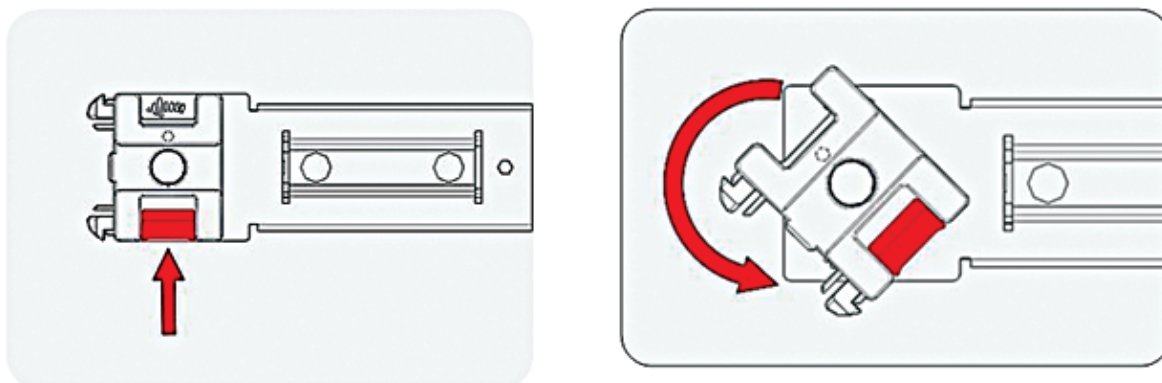


Fig. 16.5: Rotate Rackmount Bracket

Before installing, confirm that the rails included are long enough for the rack. Examine each rail to find the sides labeled *Front* and *Rear*.

For racks with square holes, snap the mounting brackets into the holes at either end of the rail into the mounting holes. Make sure to install the rails with the end labeled *Front* toward the front of the rack. Refer to [Figure 16.6](#) for a detailed view.

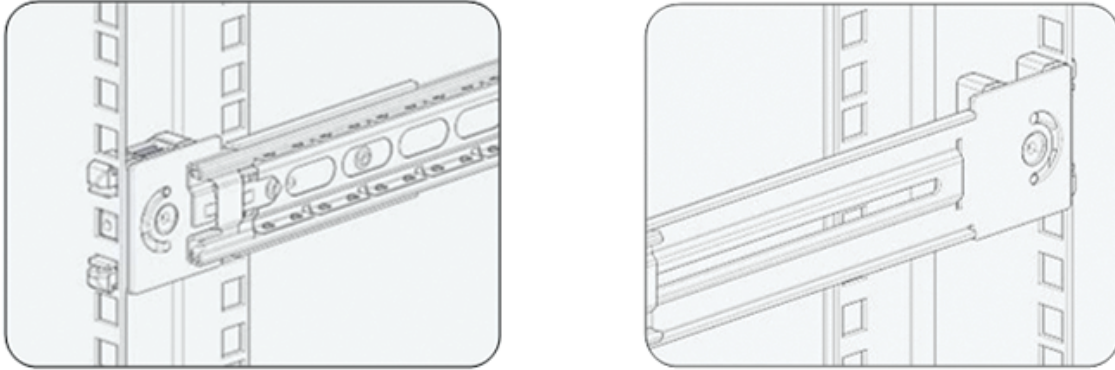


Fig. 16.6: Installing Rails in Racks with Square Holes

For racks with round holes, secure the rails into the rack at the desired position using the eight thumbscrews included. Make sure to install the rails with the end labeled *Front* toward the front of the rack. Refer to [Figure 16.7](#) for a detailed view.

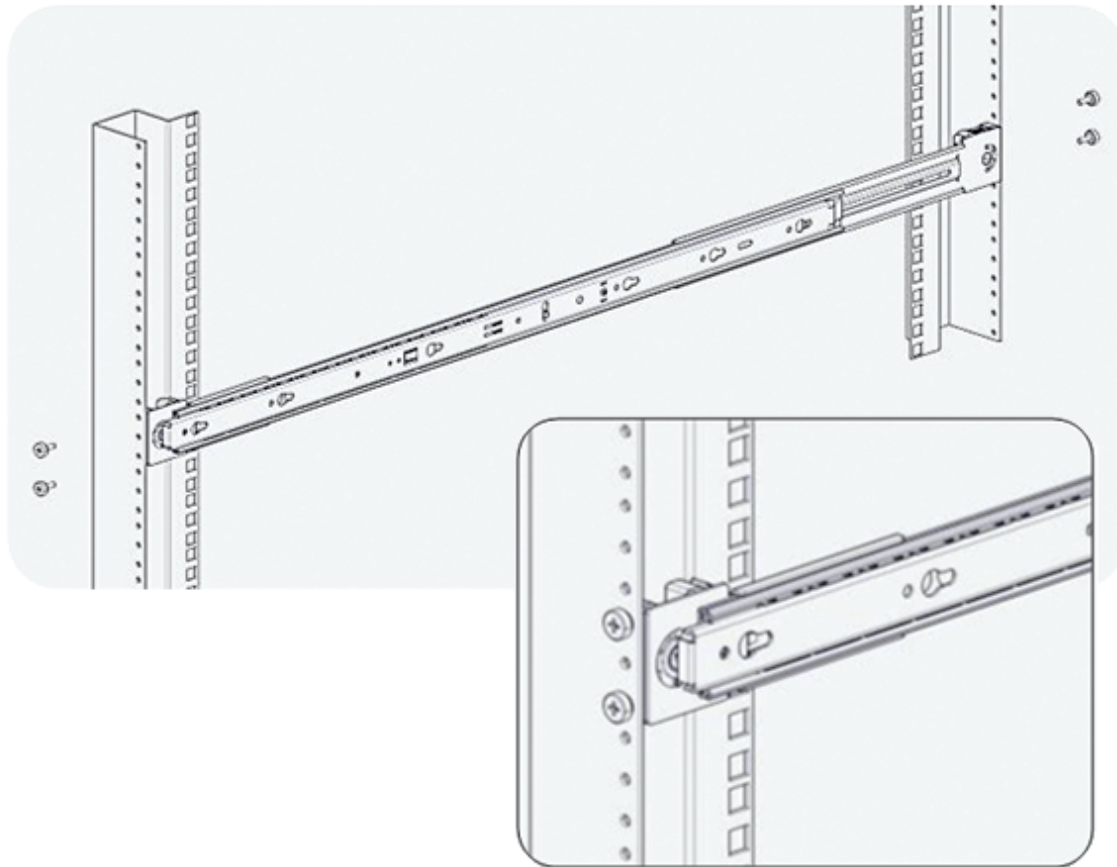


Fig. 16.7: Installing Rails in Racks with Round Holes

The TrueNAS® Storage Array can now be installed into the rack.

Warning: Two people are required to lift a TrueNAS® Storage Array.

Carefully align the TrueNAS® Storage Array inner rail with the notches in the outer rail. When the rails are aligned, slide the array toward the rack. When the array stops moving, move the pin-lock latches to allow the array to slide the rest of the way into the rack. Refer to [Figure 16.8](#) for a detailed view.

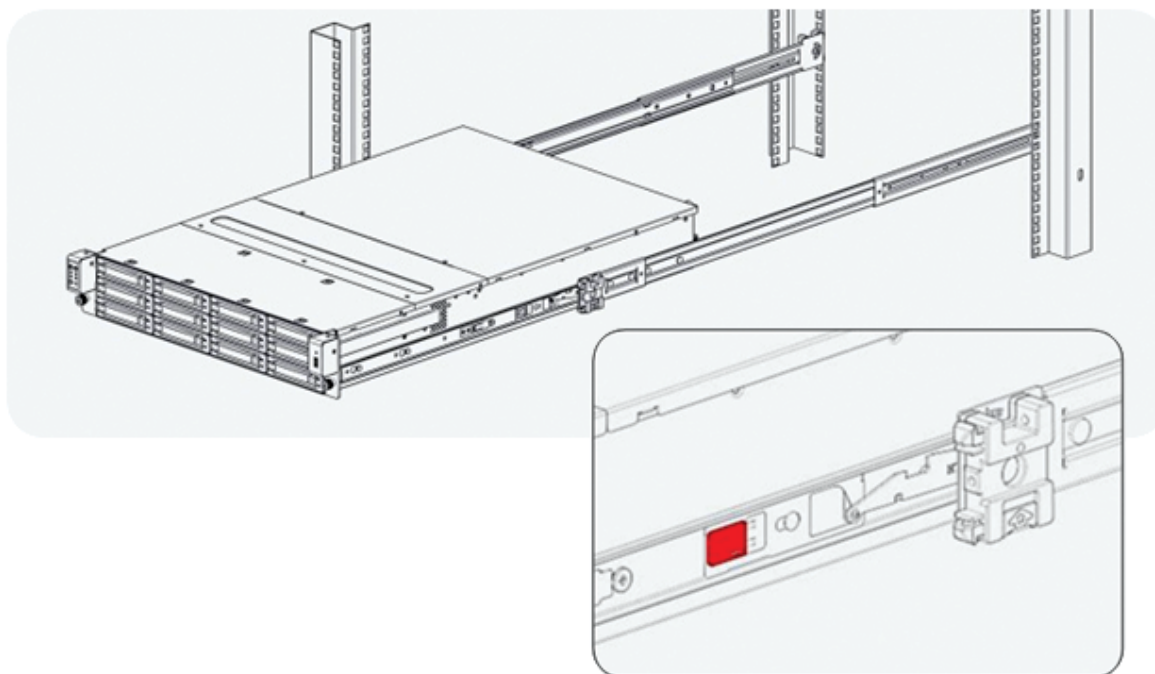


Fig. 16.8: Push Array into Rack and Release Pin-lock Latches

Install all of the populated drive trays into the front of the array. Refer to [Figure 16.9](#) for a detailed view.

Note: To avoid personal injury, do not install drives into the TrueNAS® Storage Array before racking.

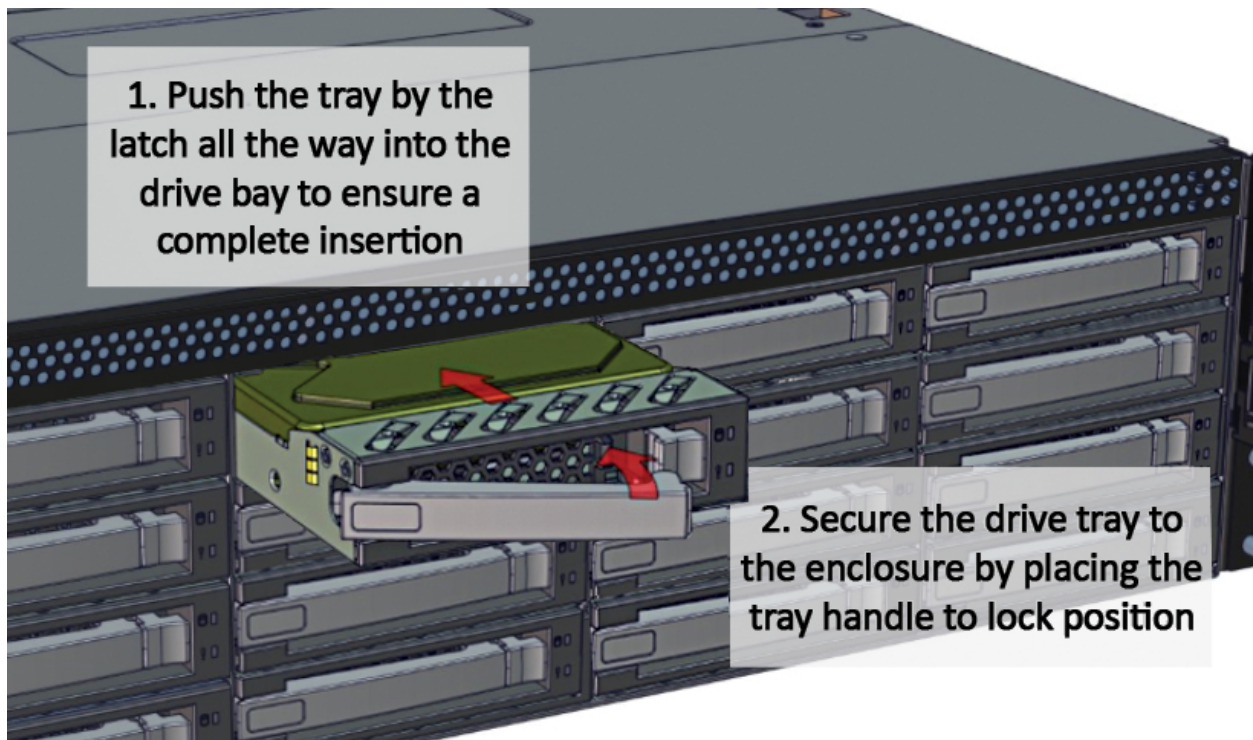


Fig. 16.9: Drive Installation Instructions

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.

Using the *Out-of-Band Management* (page 4) feature requires connecting and configuring the out-of-band management port before turning on the TrueNAS® Storage Array. Refer to [Figure 16.10](#) or the sticker on the storage controller handle for the location of the out-of-band management port.



Fig. 16.10: Back Panel Layout

Storage cabling instructions are shown in the *E16/E16F Expansion Shelf* (page 255) and *E24 Expansion Shelf* (page 264) sections.

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 will determine 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. Refer to the [Alert](#) (page 239) section of the TrueNAS® administrative graphical interface for more details about the error condition.

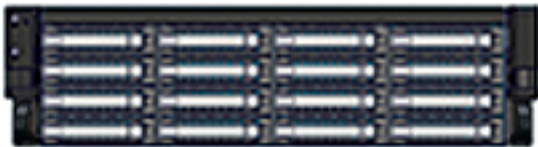
16.3 E16/E16F Expansion Shelf

Note: This product is End Of Life (EOL). Support ends January 31, 2020.

The TrueNAS® E16 expansion shelf is a 3U, 16-bay storage expansion unit designed specifically to work with the TrueNAS® Storage Array. This section will cover setting up an E16 expansion shelf and connecting it to a TrueNAS® Storage Array.

The E16 expansion shelf comes with a number of necessary accessories. If anything is missing or your E16 expansion shelf arrived in less than pristine condition, immediately take pictures and contact iXsystems support.

- TrueNAS® E16 expansion shelf



- Up to 16 populated 3.5" drive trays



- Two power cables



- Two host expansion cables (SAS 8088)



- Inner and outer rails, left and right



- Two sets of screws



- One printed guide



Unused drive bays are populated with drive tray blanks to maintain proper airflow.

Figure 16.11 shows the front view of the TrueNAS® E16 expansion shelf.

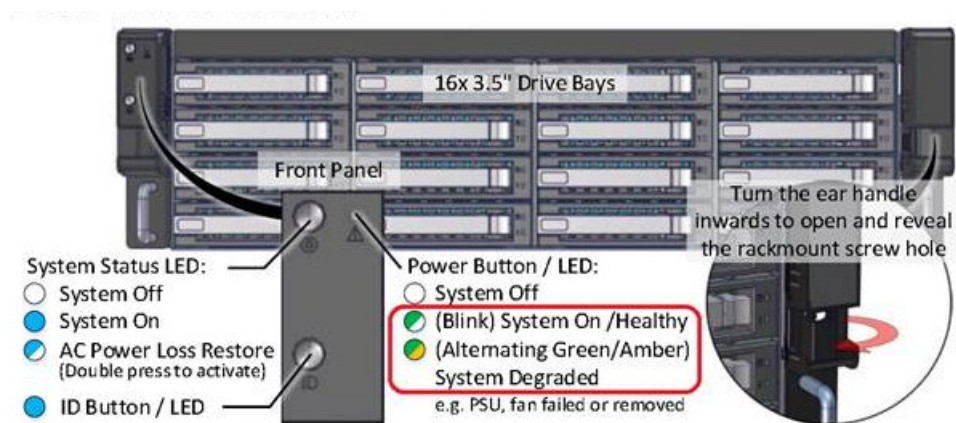


Fig. 16.11: Front View

Figure 16.12 shows the rear view of the TrueNAS® E16 expansion shelf.

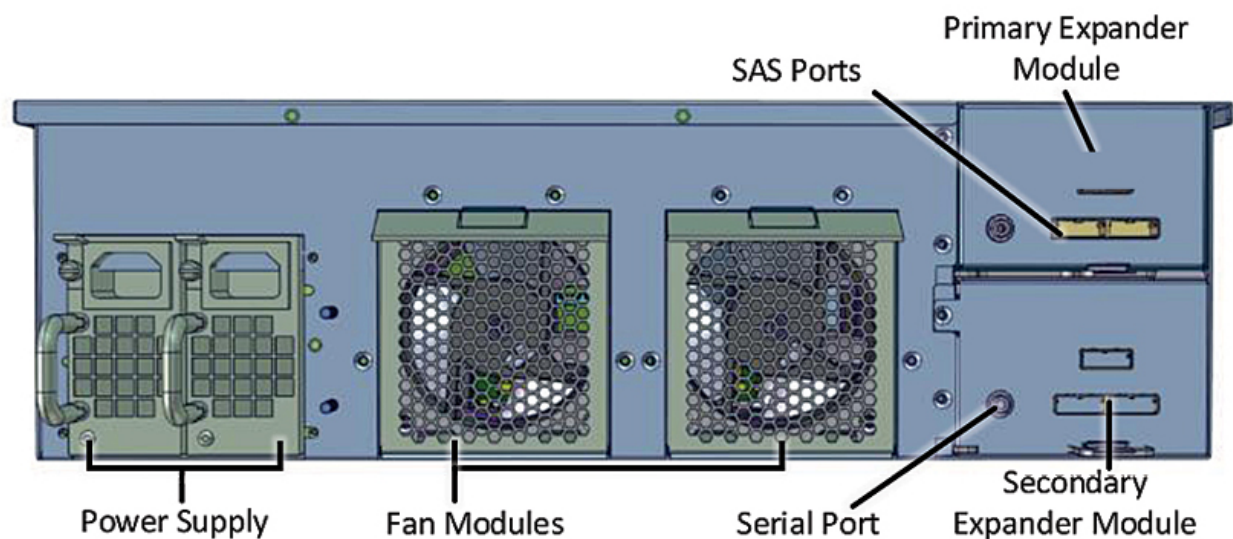


Fig. 16.12: Rear View

Figure 16.13 provides a detailed view of a drive tray and the possible statuses for the LED.

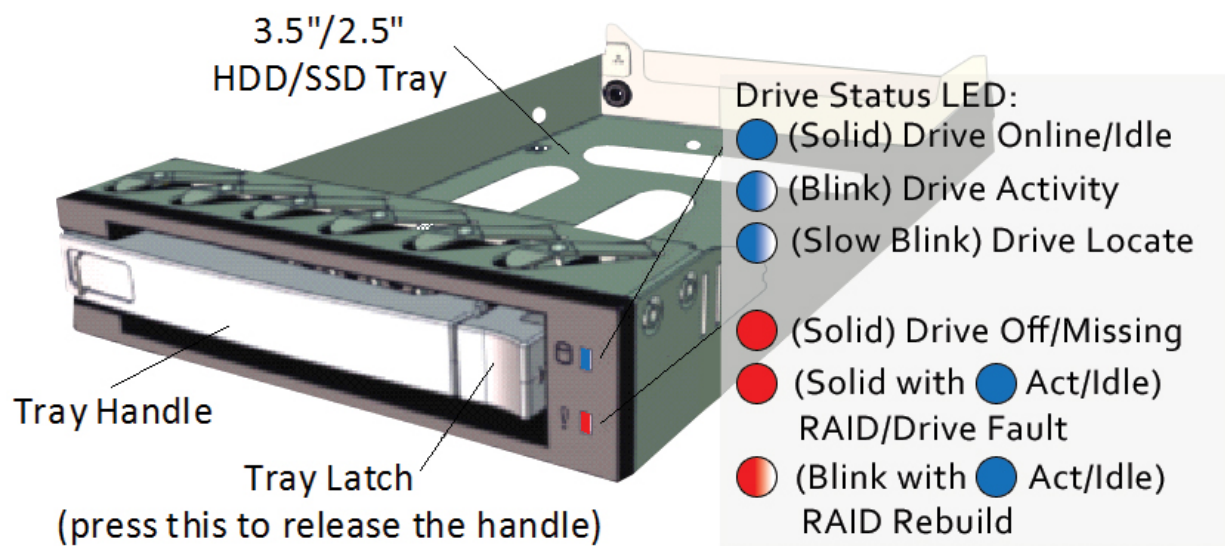


Fig. 16.13: Drive Tray

To attach the E16 expansion shelf inner rail to the chassis, remove the inner rail from both rails. Slide the inner and outer rails apart, and then push the pin-lock latch outward to allow the rails to separate completely, as shown in Figure 16.14.

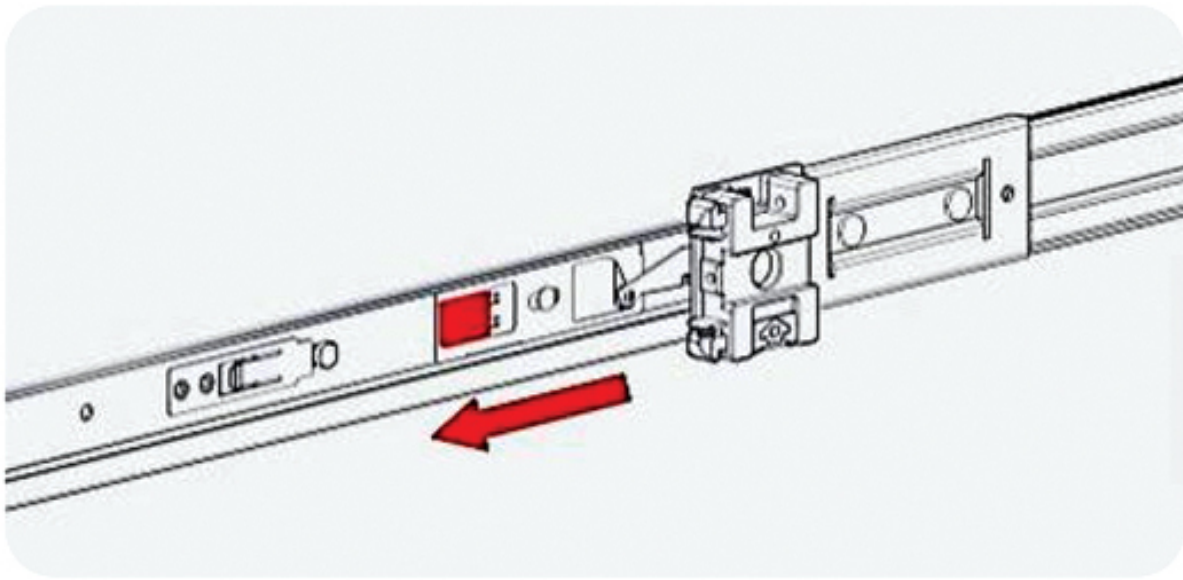


Fig. 16.14: Separate Inner and Outer Rails

Align the inner rail keyholes to the two hooks near the front of the chassis, then slide the rails forward into place as shown in [Figure 16.15](#).

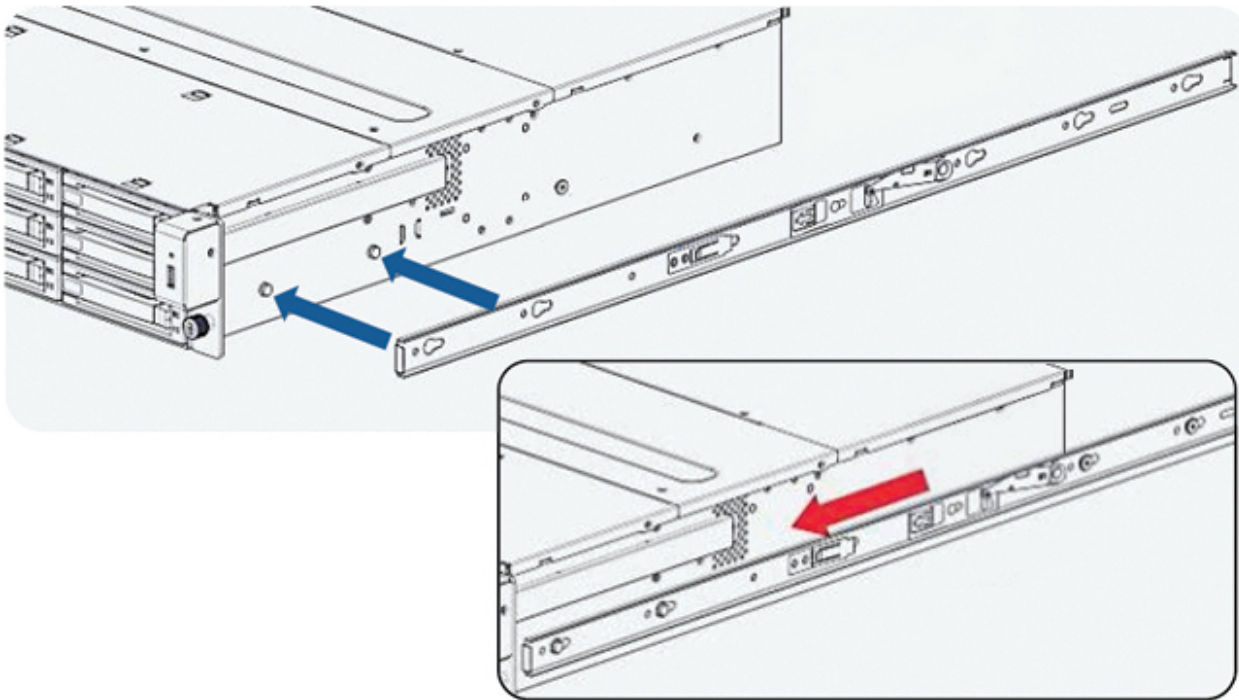


Fig. 16.15: Attach Inner Rail to Chassis

Secure the inner rail in place with a small screw from the rail kit. Refer to [Figure 16.16](#) for a detailed view.

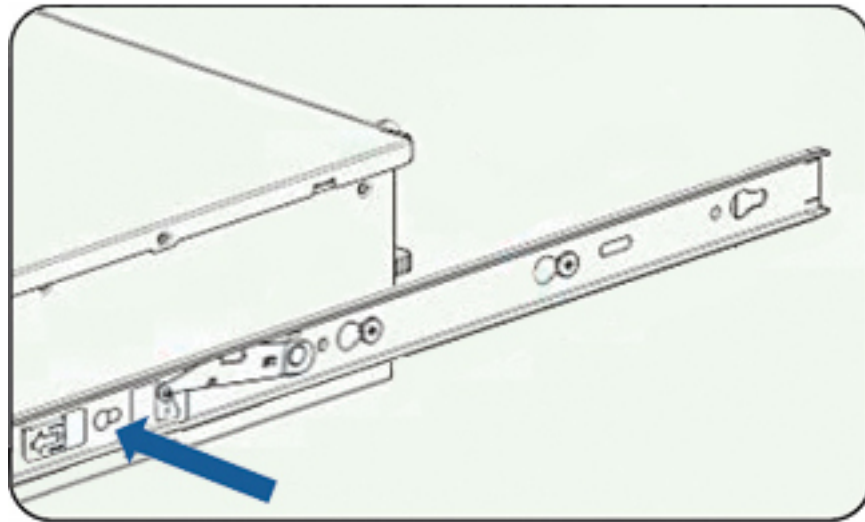


Fig. 16.16: Secure Inner Rail in Place

The TrueNAS® E16 expansion shelf slide rails support racks with both square and circular hole types. Set the mounting brackets into the correct position for your rack type by pressing the button on the mounting bracket and rotating them in place, as shown in [Figure 16.17](#). The square rack style brackets are the default. The circular hole style is the one with a flat surface and screw holes.

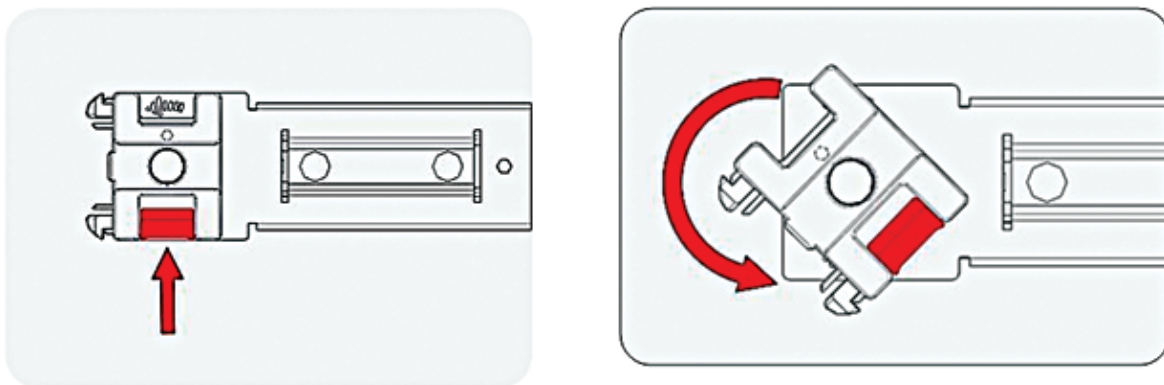


Fig. 16.17: Rotate Rackmount Bracket

Before installing, confirm that the rails included with the TrueNAS® E16 expansion shelf are long enough for your rack. Examine each rail to find the sides labeled *Front* and *Rear*.

For racks with square holes, snap the mounting brackets into the holes at either end of the rail into the mounting holes. Make sure to install the rails with the end labeled *Front* toward the front of the rack. Refer to [Figure 16.18](#) for a detailed view.

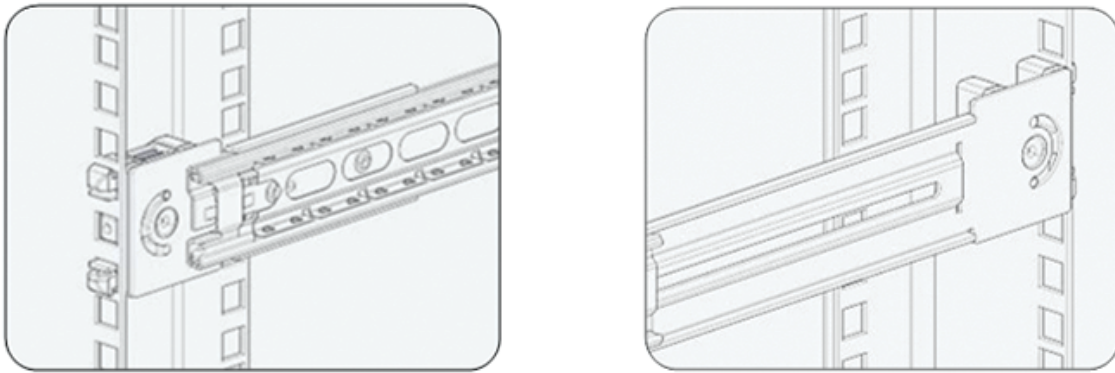


Fig. 16.18: Installing Rails in Racks with Square Holes

For racks with round holes, secure the rails into the rack at the desired position using the eight thumbscrews included with the rails. Make sure to install the rails with the end labeled *Front* toward the front of the rack. Refer to [Figure 16.19](#) for a detailed view.

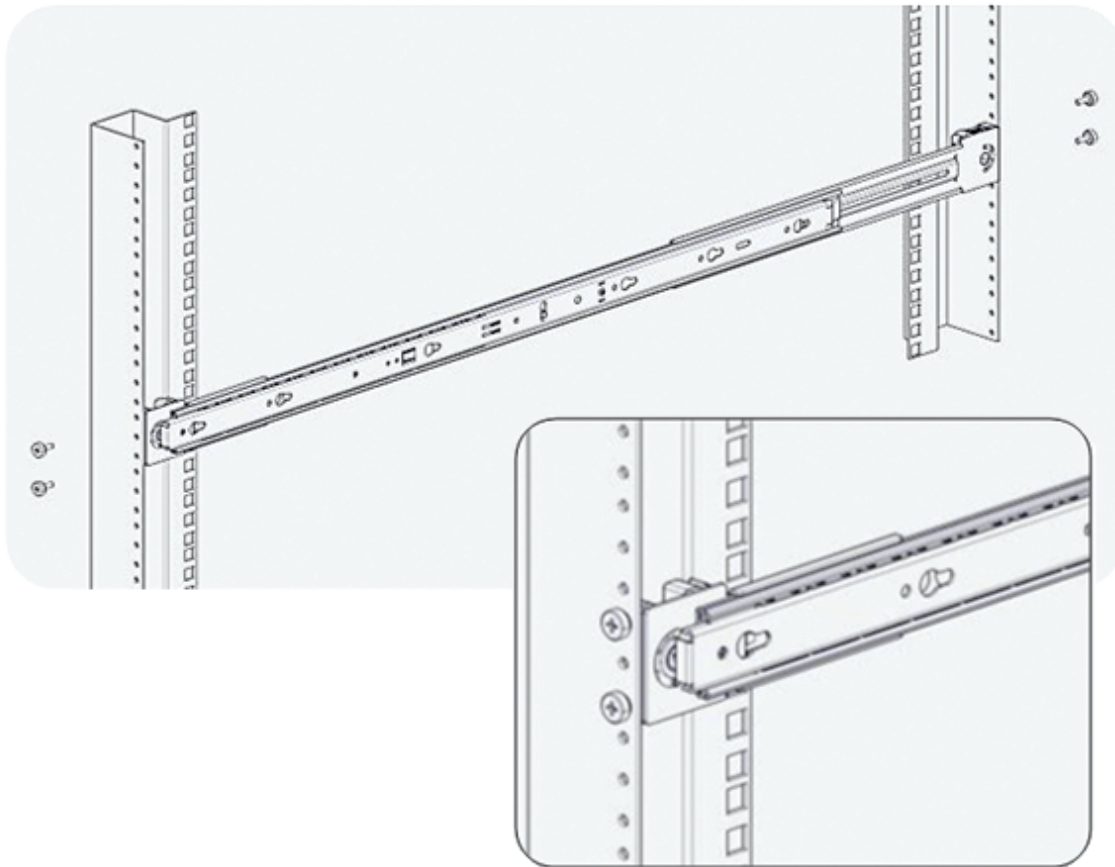


Fig. 16.19: Installing Rails in Racks with Round Holes

You are now ready to install the E16 expansion shelf into the rack.

Warning: Two people are required to lift a TrueNAS® E16 expansion shelf.

Carefully align the TrueNAS® E16 expansion shelf inner rail with the notches in the outer rail. Once the rails are aligned, slide the array toward the rack. When the array stops moving, move the pin-lock latches to allow the array to slide the rest of the way into the rack. Refer to [Figure 16.20](#) for a detailed view.

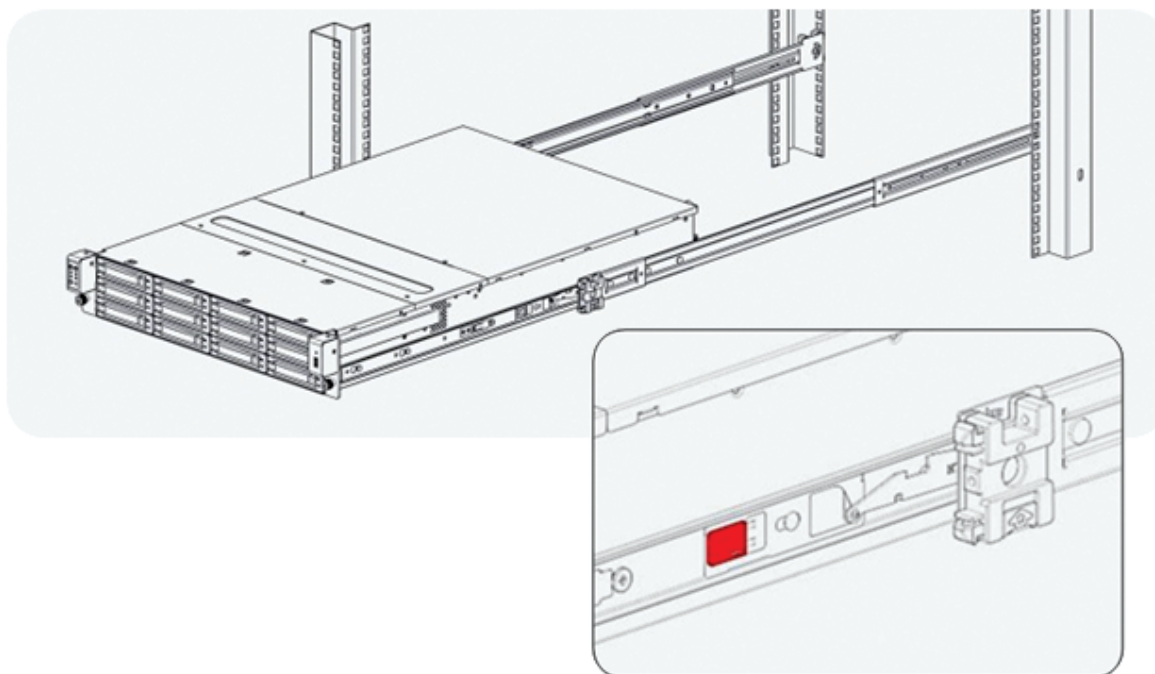


Fig. 16.20: Push Expansion Shelf into Rack and Release pin-lock Latches

Next, install all populated drive trays into the front of the expansion shelf as shown in [Figure 16.21](#).

Note: to avoid personal injury, do not install drives into the E16 expansion shelf before racking.

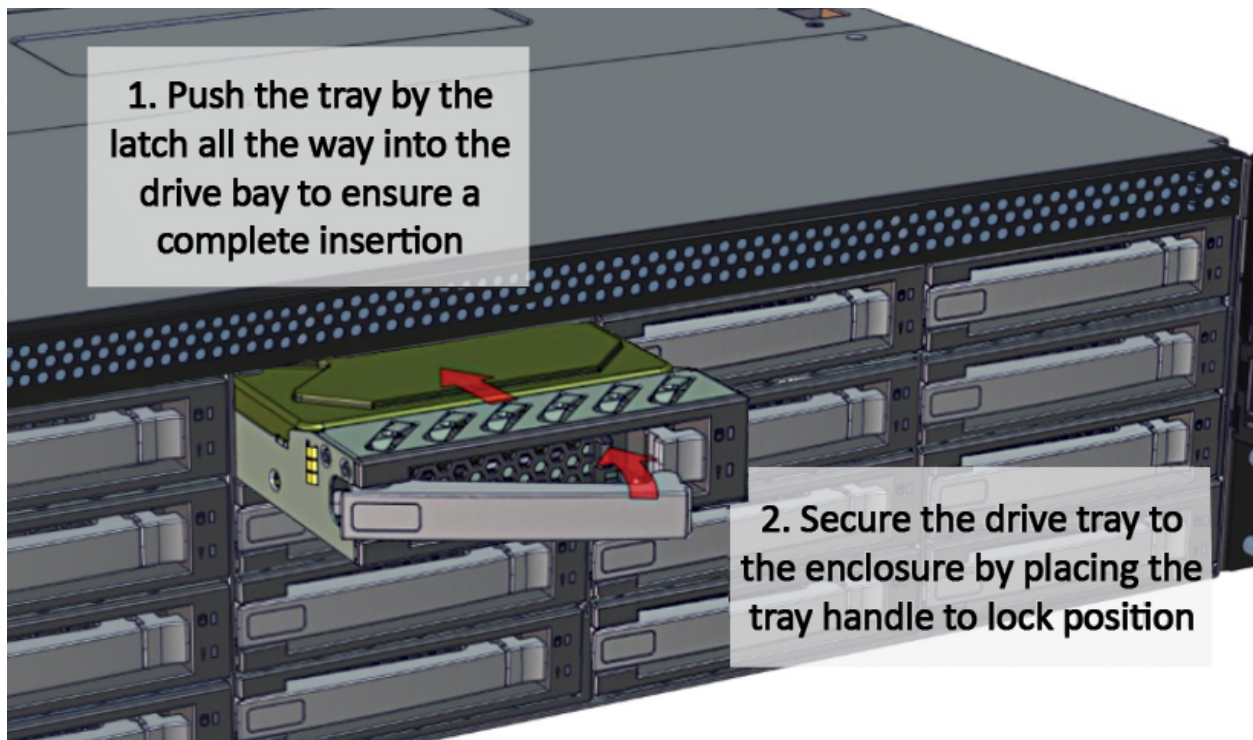


Fig. 16.21: Drive Installation Instructions

Note the labels on the SAS ports on the back of the TrueNAS® Storage Array and the letter label on the back of the expansion shelf. Using the included SAS cables, connect the *In* SAS port of the top expander on the E16 expansion shelf to the SAS port with the same letter on the TrueNAS® Storage Array's primary storage controller (the one in the top slot). If you have a secondary storage controller, connect the *In* SAS port of the bottom expander to the port with the same letter on the secondary storage controller. Refer to [Figure 16.22](#) for a detailed view.

Expansion Shelf D

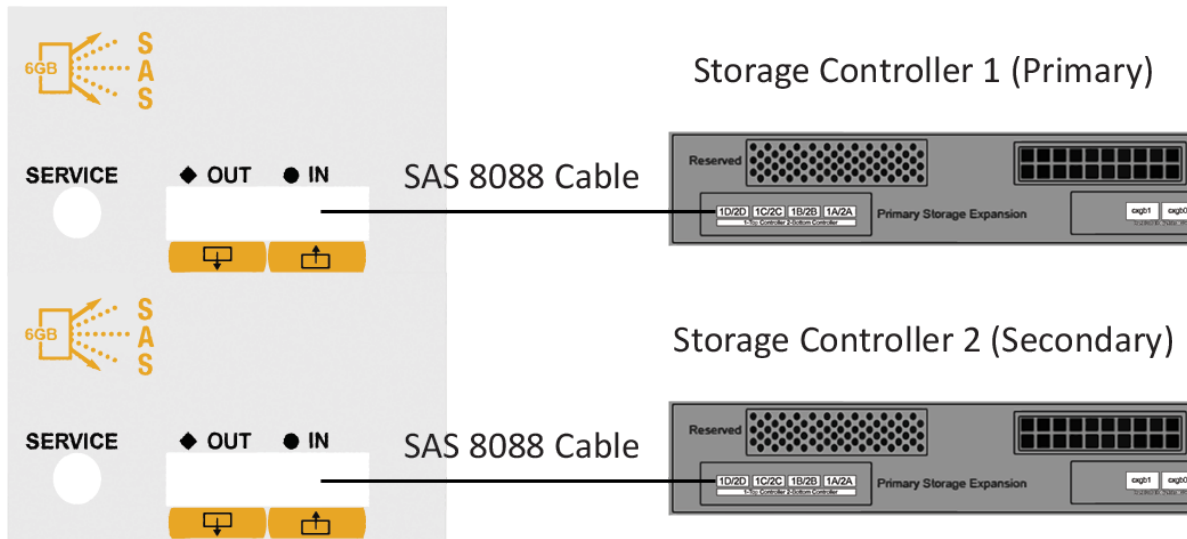


Fig. 16.22: Connecting an E16 Expansion Shelf to a TrueNAS® Storage Array

Once all the other hardware setup steps are complete, plug the power cords into the AC receptacles on the back of the E16 expansion shelf and secure them in place with the wire locks. Power on the E16 expansion shelf by pressing the top left button on the control panel.

If you are setting up a TrueNAS® Storage Array for the first time, wait two minutes after powering on all expansion shelves before turning on the TrueNAS® Storage Array.

16.4 E24 Expansion Shelf

The TrueNAS® E24 expansion shelf is a 4U, 24-bay storage expansion unit designed specifically for use with the TrueNAS® Storage Array. This section will cover setting up an E24 expansion shelf and connecting it to a TrueNAS® Storage Array.

The E24 expansion shelf comes with a number of necessary accessories. If anything is missing or your E24 expansion shelf arrived in less than pristine condition, immediately take pictures and contact iXsystems support.

- TrueNAS® E24 expansion shelf



- Up to 24 populated drive trays



- Two power cables



- Two host expansion cables (SAS 8088)



- One rail kit



- One printed guide



Unused drive bays are populated with drive tray blanks to maintain proper airflow.

Figure 16.23 shows the front of the TrueNAS® E24 expansion shelf.



Fig. 16.23: Front View

Figure 16.24 shows the rear view of the TrueNAS® E24 expansion shelf.



Fig. 16.24: Rear View

Figure 16.25 provides a detailed view of a 3.5" drive tray.



Fig. 16.25: Drive Tray

Two rails and three sets of screws are included in the rail kit. Use only the screws labeled for use in the type of rack you have. Take note of the engraved rails at either end of each rail specifying whether they are for the Left (L) or Right (R) and which end is the front and which is the back. With two people, attach each rail to the rack using the topmost and bottommost screw holes. The folded ends of the rails should be inside the corners of the rack. [Figure 16.26](#) shows the front left attachments for an L-type rack.

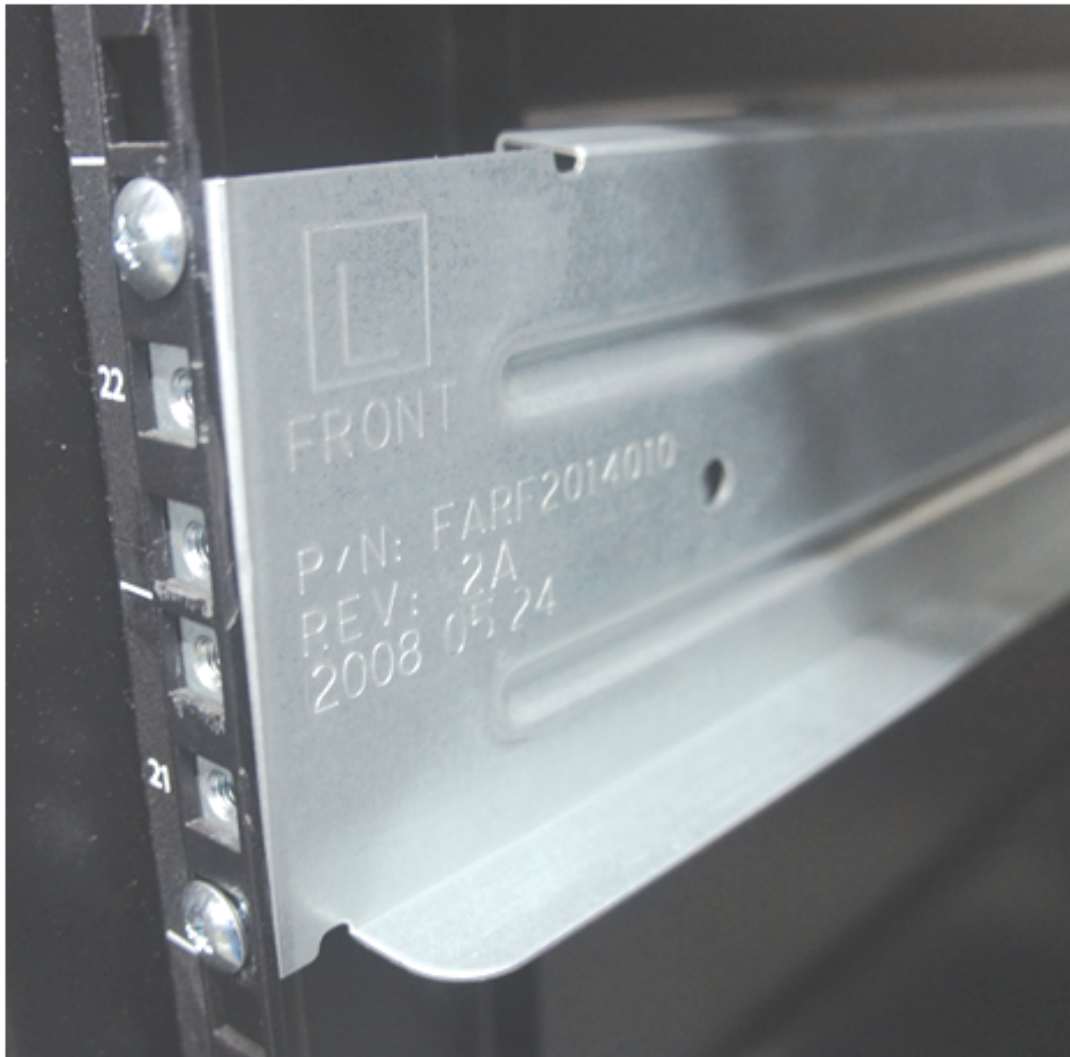


Fig. 16.26: Front Left Rail

Figure 16.27 shows the rear right attachments for an L-type rack.



Fig. 16.27: Rear Right Rail

Next, install the E24 expansion shelf into the rack.

Note: To avoid personal injury, do not install drives into the E24 expansion shelf before racking.

With two people, place the back of the expansion shelf on the rack. Gently push it backwards until the front panels of the expansion shelf are pressed against the front of the rack.

Secure the expansion shelf to the rack by pushing down and tightening the two built-in thumbscrews as indicated in [Figure 16.28](#).



Fig. 16.28: Secure E24 Expansion Shelf to the Rack

Once the E24 expansion shelf is secured into the rack, insert the included hard drives. To insert a drive, release the handle with the tab on the right side, push it into the drive bay until the handle starts to be pulled back, and then push the handle the rest of the way forward to secure the drive in place.

To connect the E24 expansion shelf to the TrueNAS® Storage Array, note the labels on the SAS ports on the back of the TrueNAS® Storage Array and the letter label on the back of the expansion shelf. Using the included SAS cables, connect the left *In* SAS port of the left side expander on the E24 expansion shelf to the SAS port with the same letter on the TrueNAS® Storage Array's primary storage controller (the one in the top slot). If you have a secondary storage controller, connect the left *In* SAS port of the right side expander to the port with the same letter on the secondary storage controller. Refer to [Figure 16.29](#) for a detailed view.

Expansion Shelf C

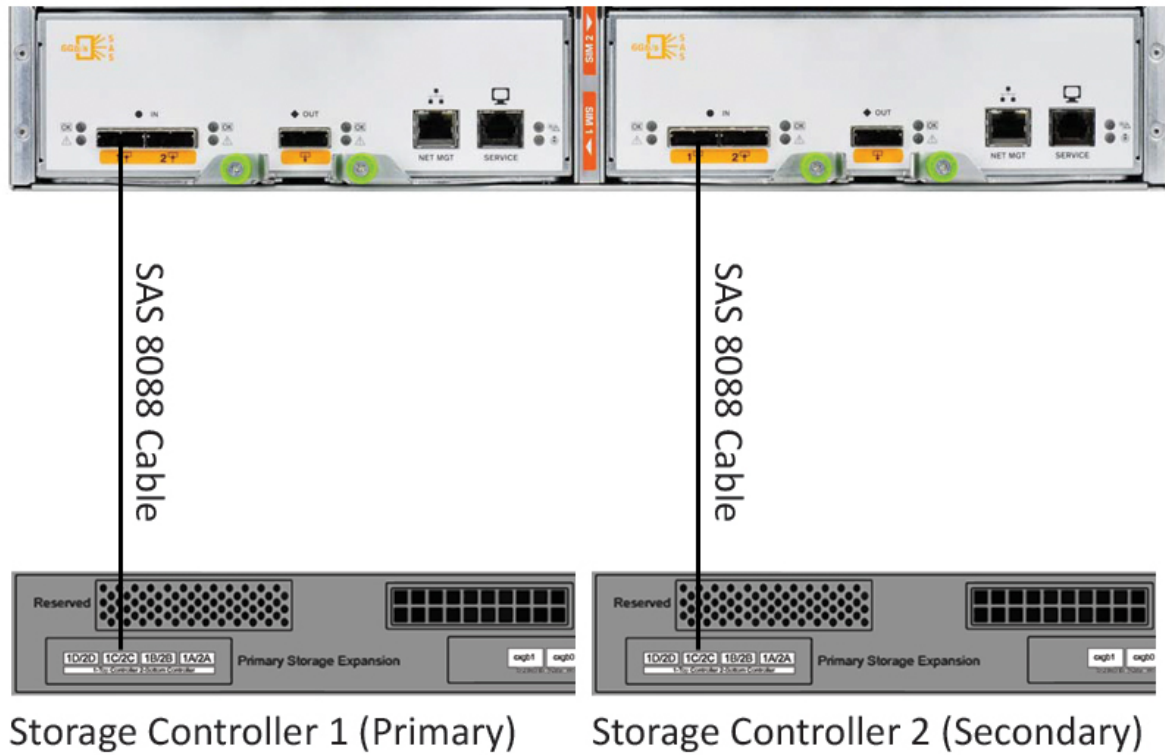


Fig. 16.29: Example connection between E24 Expansion Shelf and TrueNAS® Storage Array

Note: If you only have one storage controller, retain your second SAS cable. If you later upgrade TrueNAS® with a second storage controller, you will need it to connect to the E24 expansion shelf.

Before you plug in and power on the E24 expansion shelf, make sure the power switches on both power supplies are set to the Off (Circle) position shown in [Figure 16.30](#). Using the power cables provided, connect both power supplies to appropriate power sources. Secure the power cables in place with the plastic locks.



Fig. 16.30: E24 Power Supply

Once all the power and storage connections are set up, turn on the expansion shelf by moving the power switches on both power supplies to the On (line) position.

When setting up a TrueNAS® Storage Array for the first time, wait two minutes after powering on all expansion shelves before turning on the TrueNAS® Storage Array.

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.

17.1 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.

17.2 VAAI for NAS

VAAI for NAS (<https://code.vmware.com/programs/vaa-nas>) is automatically enabled on TrueNAS® when the *NFS* (page 205) 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.

-
- *Reserve Space* reserves space on the NAS.

USING THE API

A **REST** (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>), such as GET, PUT, POST, or DELETE.

As shown in [Figure 18.1](#), an online version of the API is available at api.freenas.org (<http://api.freenas.org>).



Fig. 18.1: 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.

18.1 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 18.1](#) summarizes acceptable values. This resource uses JSON, so the boolean values are *True* or *False*.

Table 18.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
bsdusr_username	string	maximum 32 characters, though a maximum of 8 is recommended for interoperability; can include numerals but cannot include a space
bsdusr_full_name	string	may contain spaces and uppercase characters
bsdusr_password	string	can include a mix of upper and lowercase letters, characters, and numbers
bsdusr_uid	integer	by convention, user accounts have an ID greater than 1000 with a maximum allowable value of 65,535
bsdusr_group	integer	if <i>bsdusr_creategroup</i> is set to <i>False</i> , specify the numeric ID of the group to create
bsdusr_creategroup	boolean	if set to <i>True</i> , a primary group with the same numeric ID as <i>bsdusr_uid</i> 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 <i>True</i> , user is not allowed to log in
bsdusr_locked	boolean	if set to <i>True</i> , user is not allowed to log in
bsdusr_sudo	boolean	if set to <i>True</i> , sudo 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*.

18.2 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

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.freebsd.org>

[//www.FreeNAS.org](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, [84](#)

A

Add Group, [19](#)

Add User, [22](#)

AFP, [145](#), [193](#)

Alert, [239](#)

API, [274](#)

Apple Filing Protocol, [145](#), [193](#)

Attach E16 Expansion Shelf Inner Rail to Chassis, [257](#)

Attach the TrueNAS\ :sup:'@' Faceplate, [254](#)

Autotune, [34](#)

B

Boot Environments, [29](#)

C

CA, [44](#)

Certificate Authority, [44](#)

Certificates, [47](#)

CIFS, [162](#), [210](#)

Cloud Credentials, [44](#)

Cloud Sync, [57](#)

Compression, [96](#)

Connect E16 Expansion Shelf to TrueNAS Array, [263](#)

Connect E24 Expansion Shelf to TrueNAS Array, [270](#)

Console Setup Menu, [14](#)

Create Dataset, [94](#)

Create Group, [19](#)

Create User, [22](#)

Cron Jobs, [62](#)

D

DC, [195](#)

DDNS, [196](#)

Delete Group, [20](#)

Delete User, [24](#)

Domain Controller, [195](#)

Dynamic DNS, [196](#)

E

E16 Expansion Shelf, [255](#)

E16 Expansion Shelf Contents, [255](#)

E16 Expansion Shelf Layout, [256](#)

E24 Expansion Shelf, [264](#)

Email, [34](#)

Encryption, [89](#)

EtherChannel, [78](#)

EULA, [278](#)

F

Failover, [53](#), [55](#)

File Transfer Protocol, [198](#)

FTP, [198](#)

G

Groups, [18](#)

GUI Access, [15](#)

Guide, [239](#)

H

Hardware Installation, [250](#)

I

Install Array into Rack, [252](#)

Install Drive Trays into a TrueNAS Array, [253](#)

Install Drives into the E24 Expansion Shelf, [270](#)

Install E24 Expansion Shelf into Rack, [269](#)

Install E24 Expansion Shelf Rails, [267](#)

Install TrueNAS\ :sup:'@' Outer Rail in Rack, [250](#)

Internet Small Computer System Interface, [173](#)

iSCSI, [173](#)

L

LACP, [78](#)

LAGG, [78](#)

Link Aggregation, [78](#)

Link Layer Discovery Protocol, [204](#)

LLDP, [204](#)

Log Out, [238](#)

M

Multiple Boot Environments, [29](#)

N

Network File System, [153](#), [205](#)
Network Settings, [73](#)
New Group, [19](#)
New User, [22](#)
NFS, [153](#), [205](#)

O

Out-of-Band Management, [4](#)

P

Periodic Snapshot, [112](#)
Plug in and Power on E16 Expansion Shelf, [264](#)
Plug in and Power on E24 Expansion Shelf, [271](#)
Plug in and Power on your TrueNAS\ :sup:® array, [254](#)
Proactive Support, [52](#)
Processes, [236](#)

R

Reboot, [238](#)
Remove Group, [20](#)
Remove User, [24](#)
Replace Failed Drive, [109](#)
Replication, [114](#)
Reporting, [226](#)
RFC
 RFC 2616, [275](#)
 RFC 3721, [175](#)
Route, [84](#)
Rsync, [206](#)
Rsync Tasks, [64](#)

S

S.M.A.R.T., [208](#)
S.M.A.R.T. Tests, [71](#)
Samba, [162](#), [210](#)
SCP, [218](#)
Scrub, [125](#)
Secure Copy, [218](#)
Secure Shell, [217](#)
Services, [190](#)
Shadow Copies, [171](#)
Shell, [236](#)
Shutdown, [238](#)
Simple Network Management Protocol, [215](#)
SMB, [162](#), [210](#)
Snapshot, [112](#)
Snapshots, [126](#)
SNMP, [215](#)
SSH, [217](#)
Start Service, [191](#)
Static Route, [84](#)

Stop Service, [191](#)
Support, [50](#), [239](#)
System Dataset, [36](#)

T

Tasks, [56](#)
TFTP, [219](#)
Time Machine, [149](#)
Trivial File Transfer Protocol, [219](#)
TrueNAS E24 Expansion Shelf Contents, [264](#)
TrueNAS E24 Expansion Shelf Layout, [265](#)
TrueNAS\ :sup:® Unified Storage Array, [246](#)
Trunking, [84](#)
Tunables, [37](#)

U

Uninterruptible Power Supply, [220](#)
Upgrade ZFS Pool, [43](#)
UPS, [220](#)
Users, [20](#)

V

VAAI, [272](#)
VAAI for iSCSI, [273](#)
VAAI for NAS, [273](#)
vCenter, [224](#)
VLAN, [84](#)
VMware Snapshot, [128](#)
Volumes, [87](#)

W

WebDAV, [161](#), [223](#)
Windows File Share, [210](#)
Windows Shares, [162](#)
Wizard, [228](#)

Z

ZVOL, [96](#)