TrueNAS® 11.2-U7 User Guide

CONTENTS

Welcome

Welcome to the TrueNAS® User Guide.

Typographic Conventions

Typographic Conventions

The TrueNAS® Administrator Guide uses these typographic conventions:

Table 1: Text Format Examples

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

CHAPTER

ONE

INTRODUCTION

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

1.1 Contacting iXsystems

For assistance, please contact iX Support:

Contact Method	Contact Options
Web	https://support.ixsystems.com
Email	support@iXsystems.com
Telephone	Monday - Friday, 6:00AM to 6:00PM Pacific Standard Time: • US-only toll-free: 855-473-7449 option 2 • Local and international: 408-943-4100 option 2
Telephone	After Hours (24x7 Gold Level Support only): • US-only toll-free: 855-499-5131 • International: 408-878-3140 (international calling rates will apply)

1.2 Path and Name Lengths

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

Table 1.2: Path and Name Lengths

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

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

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

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

- **Automatic mount:** ZFS temporarily mounts a snapshot whenever a user attempts to view or search the files within the snapshot. The mountpoint used will be in the hidden directory `.zfs/snapshot/name` within the same ZFS dataset. For example, the snapshot `mypool/dataset/snap1@snap2` is mounted at `/mnt/mypool/dataset/.zfs/snapshot/snap2/`. If the length of this path exceeds 88 bytes the snapshot will not be automatically mounted by ZFS and the snapshot contents will not be visible or searchable. This can be resolved by renaming the ZFS pool or dataset containing the snapshot to shorter names (`mypool` or `dataset`), or by shortening the second part of the snapshot name (`snap2`), so that the total mounted path length does not exceed 88 bytes. ZFS will automatically perform any necessary unmount or remount of the file system as part of the rename operation. After renaming, the snapshot data will be visible and searchable again.
- **Manual mount:** If the same example snapshot is mounted manually from the `CLI`, using `mount -t zfs mypool/dataset/snap1@snap2 /mnt/mymountpoint` the path `/mnt/mountpoint/` must not exceed 88 bytes, but the length of the snapshot name will be *irrelevant*. When renaming a manual mountpoint, any object mounted on the mountpoint must be manually unmounted (using the `umount` command in the `CLI`) before renaming the mountpoint and can be remounted afterwards.

Note: A snapshot that cannot be mounted automatically by ZFS, can still be mounted manually from the `CLI` using a shorter mountpoint path. This makes it possible to mount and access snapshots that cannot be accessed automatically in other ways, such as from the GUI or from features such as “File History” or “Versions”.

CHAPTER

TWO

INITIAL SETUP

Please set up the TrueNAS® hardware before beginning software configuration. Basic Setup Guides for TrueNAS® systems and expansion shelves are available in the iX Information Library (https://www.ixsystems.com/blog/knowledgebase_category/truenas/). These guides provide detailed instructions on rack installation, drive tray loading and LED behavior, cable connections, and other important setup information.

Depending on the degree of pre-configuration requested from iXsystems, most of the initial TrueNAS® software 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 Console Setup Menu

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

Note: When connecting to the TrueNAS® system with SSH or the web *Shell* (page 244), the Console Setup menu is not shown by default. It can be started by the *root* user or another user with root permissions by typing `/etc/netcli`.

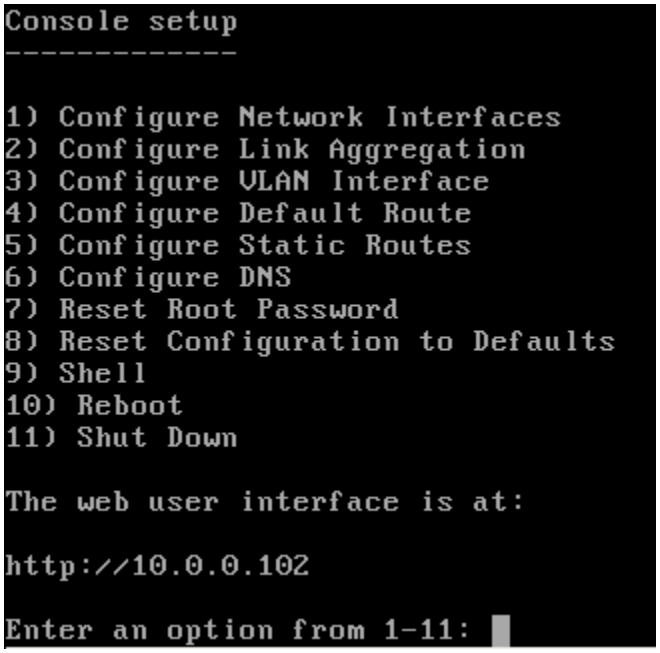

Fig. 2.1: Console Setup Menu

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

The menu provides these options:

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

The wizard also prompts marking an interface as critical for failover. This allows logging in to the web interface available at the virtual IP address after a failover.

2) Configure Link Aggregation is for creating or deleting link aggregations.

3) Configure VLAN Interface is used to create or delete VLAN interfaces.

4) Configure Default Route is used to set the IPv4 or IPv6 default gateway. When prompted, enter the IP address of the default gateway.

5) Configure Static Routes prompts for the destination network and gateway IP address. Re-enter this option for each static route needed.

6) Configure DNS prompts for the name of the DNS domain and the IP address of the first DNS server. When adding multiple DNS servers, press `Enter` to enter the next one. Press `Enter` twice to leave this option.

7) Reset Root Password is used to reset a lost or forgotten `root` password. Select this option and follow the prompts to set the password.

8) Reset Configuration to Defaults **Caution**! This option deletes *all* of the configuration settings made in the administrative GUI and is used to reset a TrueNAS® system back to defaults. **Before selecting this option, make a full backup of all data and make sure all encryption keys and passphrases are known!** After this option is selected, the configuration is reset to defaults and the system reboots. *Storage* → *Volumes* → *Import Volume* can be used to re-import volumes.

9) Shell starts a shell for running FreeBSD commands. To leave the shell, type `exit`.

10) Reboot reboots the system.

11) Shut Down shuts down the system.

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

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

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

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

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

...

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

2.2 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 2.2.

Fig. 2.2: Enter the Root Password

Enter the default password of *abcd1234*.

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 *Modify User*, 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 End User License Agreement (EULA) found in *Appendix A: TrueNAS Software End User License Agreement* (page 260) is displayed. To accept the EULA, click *I agree*.

Next, a box for the license key is displayed. Paste in the license key to access the web interface.

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.

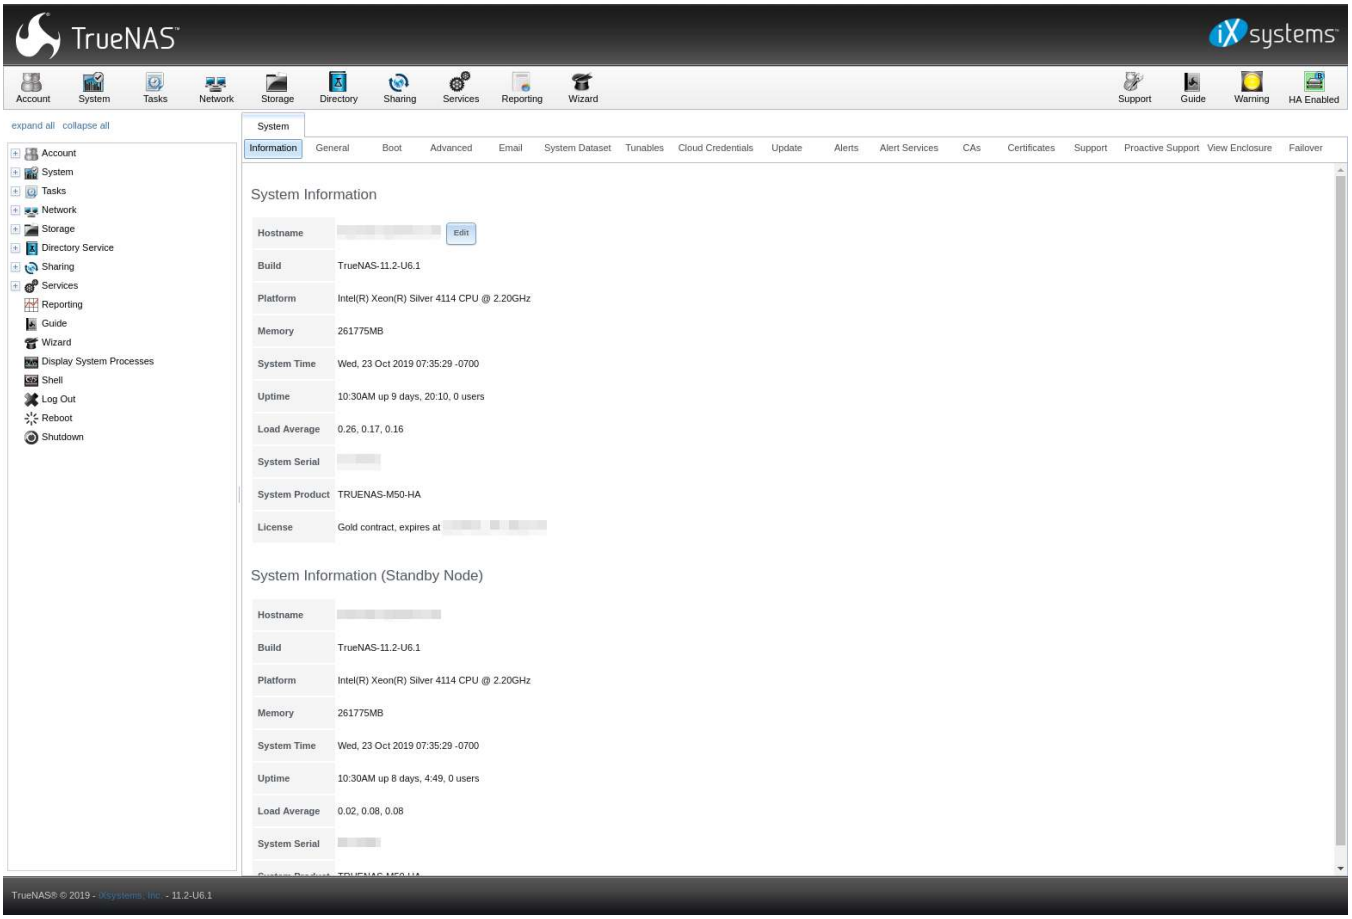

Fig. 2.3: TrueNAS® Graphical Configuration Menu

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/) 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) 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.

Tip: iXsystems recommends *contacting an iXsystems Support Representative* (page 9) 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.

Warning: 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.

CHAPTER

THREE

ACCOUNT

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

- *Groups* (page 16): used to manage UNIX-style groups on the TrueNAS® system.
- *Users* (page 19): 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: It is unnecessary to recreate the network users or groups when a directory service is running on the same network. Instead, import the existing account information into TrueNAS®. Refer to *Directory Services* (page 141) for details.

This section describes how to create a group and assign user accounts to it. The next section, *Users* (page 19), describes creating user accounts.

Click *Groups* → *View Groups* to see a screen like Figure 3.1.

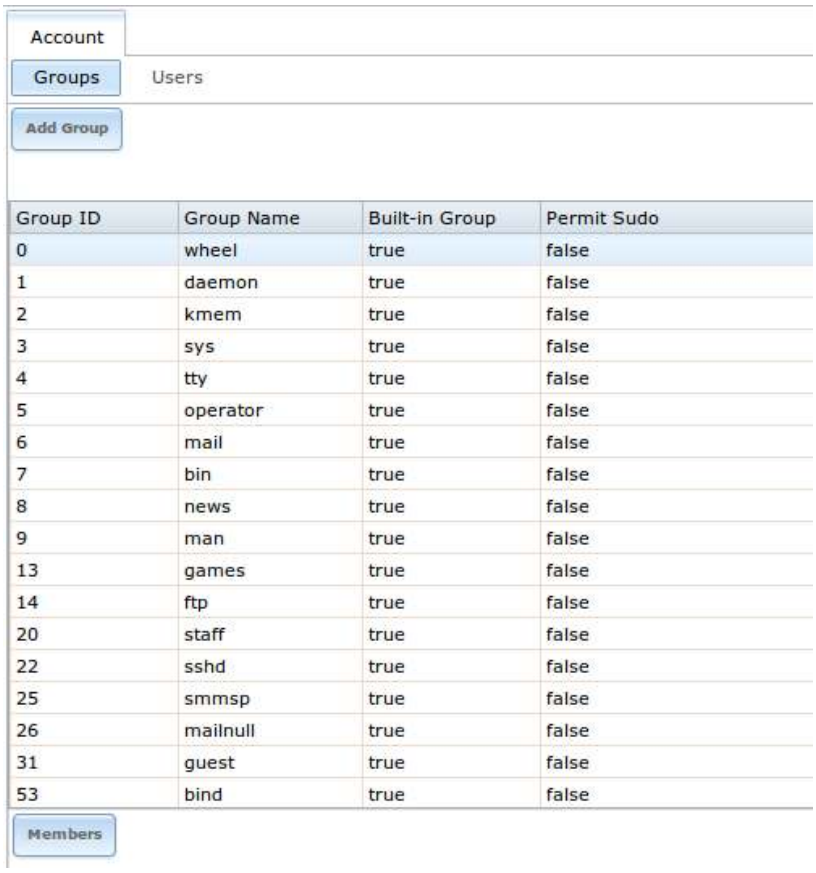

Fig. 3.1: Group Management

The *Groups* page lists all groups, including those built-in and used by the operating system. The table displays group names, group IDs (GID), built-in groups, and if `sudo` is permitted. 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.

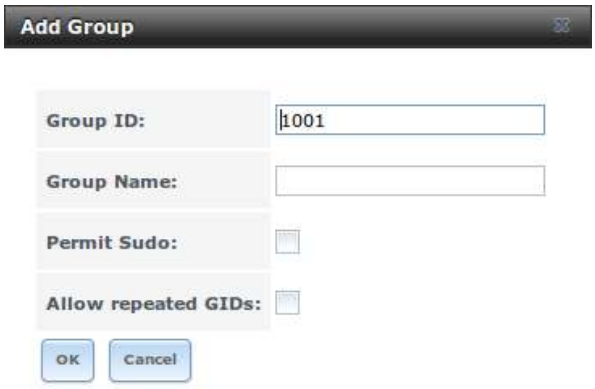

Fig. 3.2: Creating a New Group

Table 3.1: Group Creation Options

Setting	Value	Description
Group ID	string	The next available group ID is suggested. UNIX groups containing user accounts typically have an ID greater than 1000 and groups required by a service have an ID equal to the default port number used by the service. Example: the sshd group has an ID of 22.
Group Name	string	Enter an alphanumeric name for the new group. The period (.), hyphen (-), and underscore (_) characters are allowed as long as the group name does not begin with a period (.) or hyphen (-).
Permit Sudo	checkbox	Set to allow group members to use sudo (https://www.sudo.ws/). When using `sudo`, a user is prompted for their own password.
Allow repeated GIDs	checkbox	Set to allow multiple groups to share the same group id (GID). This is useful when a GID is already associated with the UNIX permissions for existing data, but is generally not recommended.

After a group and users are created, users can be added to a group. Highlight the group where users will be assigned, then click the *Members* button. Highlight the user in the *Member users* list. This shows all user accounts on the system. 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.

Figure 3.3, shows user1 added as a member of group data1.

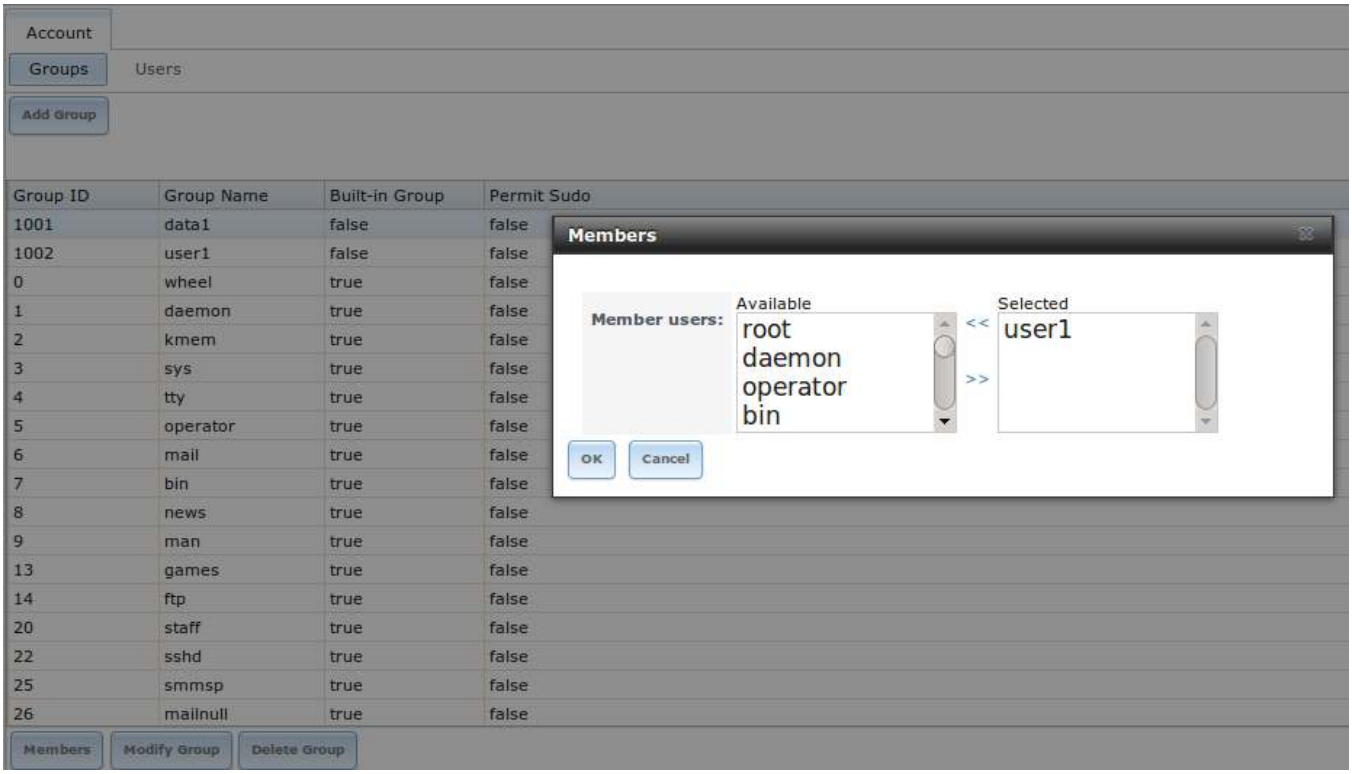

Fig. 3.3: Assigning a User to a Group

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 flexibility in configuring which users have access to the data stored on TrueNAS®. To assign permissions to shares, **one of these options** must be done:

1. Create a guest account for all users, or create a user account for every user in the network where the name of each account is the same as a login name used on a computer. For example, if a Windows system has a login name of *bobsmith*, create a user account with the name *bobsmith* on TrueNAS®. A common strategy is to create groups with different sets of permissions on shares, then assign users to those groups.
2. If the network uses a directory service, import the existing account information using the instructions in *Directory Services* (page 141).

Account → *Users* lists all system accounts installed with the TrueNAS® operating system, as shown in Figure 3.4.

Account

Groups | Users

Add User

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

Modify User

Fig. 3.4: Managing User Accounts

Each account entry indicates the user ID, username, primary group ID, home directory, default shell, full name, whether it is a built-in user that came with the TrueNAS® installation, the email address, if logins are disabled, if the user account is locked, whether the user is allowed to use `sudo`, and if the user connects from a Windows 8 or newer

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

Except for the *root* user, the accounts that come with TrueNAS® are system accounts. Each system account is used by a service and should not be used as a login account. For this reason, the default shell on system accounts is nologin(8) (https://www.freebsd.org/cgi/man.cgi?query=nologin). For security reasons and to prevent breakage of system services, 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 *Advanced Mode* or configure the system to always display these settings by setting *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 141), Windows user passwords must be set from within Windows.

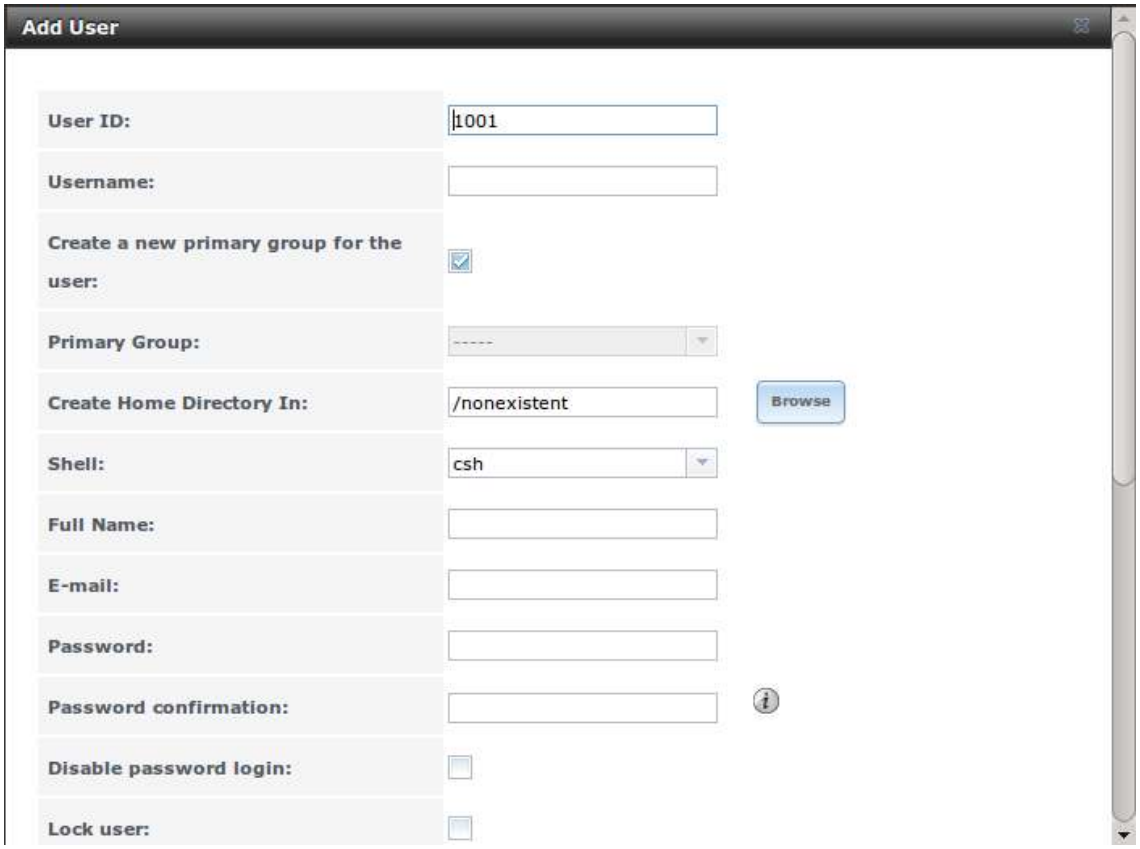

Fig. 3.5: Adding or Editing a User Account

Table 3.2: User Account Configuration

Setting	Value	Advanced Mode	Description
User ID	integer		Grayed out if the user already exists. When creating an account, the next numeric ID is suggested. User accounts typically have an ID greater than 1000 and system accounts have an ID equal to the default port number used by the service.
Username	string		Usernames can be up to 16 characters long. When using NIS or other legacy software with limited username lengths, keep usernames to eight characters or less for compatibility. Usernames cannot begin with a hyphen (-) or contain a space, tab, or these characters: *, : + & # % ^ ( ) ! @ ~ * ? < > = . $* can only be used as the last character of the username.
Create a new primary group	checkbox		A primary group with the same name as the user is created automatically. Unset to select a different primary group name.
Primary Group	drop-down menu		Unset *Create a new primary group* to access this menu. For security reasons, FreeBSD does not give a user `su` permissions if *wheel* is their primary group. To give a user `su` access, add them to the *wheel* group in *Auxiliary groups*.
Create Home Directory In	browse button		Choose a path to the user's home directory. If the directory exists and matches the username, it is set as the user's home directory. When the path does not end with a subdirectory matching the username, a new subdirectory is created. The full path to the user's home directory is shown here when editing a user.
Home Directory Mode	checkboxes	✓	Sets default Unix permissions of the user's home directory. This is read-only for built-in users.
Shell	drop-down menu		Select the shell to use for local and SSH logins. The *root* user shell is used for web interface *Shell* (page 244) sessions. See Table 3.3 for an overview of available shells.
Full Name	string		Required. This field may contain spaces.
E-mail	string		The email address associated with the account.
Password	string		Required unless *Disable password login* is set. Cannot contain a `?`.
Password confirmation	string		This must match the value of *Password*.
Disable password login	checkbox		Set to disable password logins and authentication to SMB shares. To undo this setting, create a password for the user by clicking *Modify User* for the user in the *View Users* screen. Setting this grays out *Lock user* and *Permit Sudo*.
Lock user	checkbox		Set to prevent the user from logging in until this box is unset. Setting this grays out *Disable password login*.
Permit Sudo	checkbox		Set to give group members permission to use sudo (https://www.sudo.ws/). When using `sudo`, a user is prompted for their own password.
Microsoft Account	checkbox		Set this when the user is connecting from a Windows 8 or newer system or when using a Microsoft cloud service.
SSH Public Key	string		Enter or 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 groups to add the user. Click the >> to add the user to the highlighted groups.

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

Table 3.3: Available Shells

Shell	Description
netcli.sh	User is shown the Console Setup menu (Figure 2.1) on connection, even if it is disabled in *System* → *Advanced* → *Enable Console Menu*. The user must be *root* or have root permissions (effective user ID 0, like *toor*).
csh	C shell (https://en.wikipedia.org/wiki/C_shell)
sh	Bourne shell (https://en.wikipedia.org/wiki/Bourne_shell)
tcsh	Enhanced C shell (https://en.wikipedia.org/wiki/Tcsh)
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 (https://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 were added by the system administrator. If the user to be removed is the last user in a custom group, an option is offered to keep the user primary group after deleting the user.

CHAPTER

FOUR

SYSTEM

The System section of the web interface contains these entries:

- *Information* (page 23) provides general TrueNAS® system information such as hostname, operating system version, platform, and uptime
- *General* (page 24) configures general settings such as HTTPS access, the language, and the timezone
- *Boot* (page 27) creates, renames, and deletes boot environments. It also shows the condition of the Boot Volume
- *Advanced* (page 29) 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 35) configures the location where logs and reporting graphs are stored
- *Tunables* (page 36) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- *Cloud Credentials* (page 39) is used to enter connection credentials for remote cloud service providers
- *Update* (page 42) performs upgrades and checks for system updates
- *Alerts* (page 46) lists the available *Alert* (page 247) conditions and provides configuration of the notification frequency for each alert
- *Alert Services* (page 47) configures services used to notify the administrator about system events
- *CAs* (page 49): import or create internal or intermediate CAs (Certificate Authorities)
- *Certificates* (page 51): import existing certificates or create self-signed certificates
- *Support* (page 54): view licensing information or create a support ticket
- *Proactive Support* (page 56): enable and configure automatic proactive support (Silver or Gold support coverage only)
- *View Enclosure* (page 57): view status of disk enclosures
- *Failover* (page 58): 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 hostname, build version, type of CPU (platform), amount of memory, current system time, system uptime, number of users connected at the console or by serial, telnet, or SSH connections, and current load average. On systems supplied or certified by iXsystems, an additional *Serial Number* field showing the hardware serial number is displayed.

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

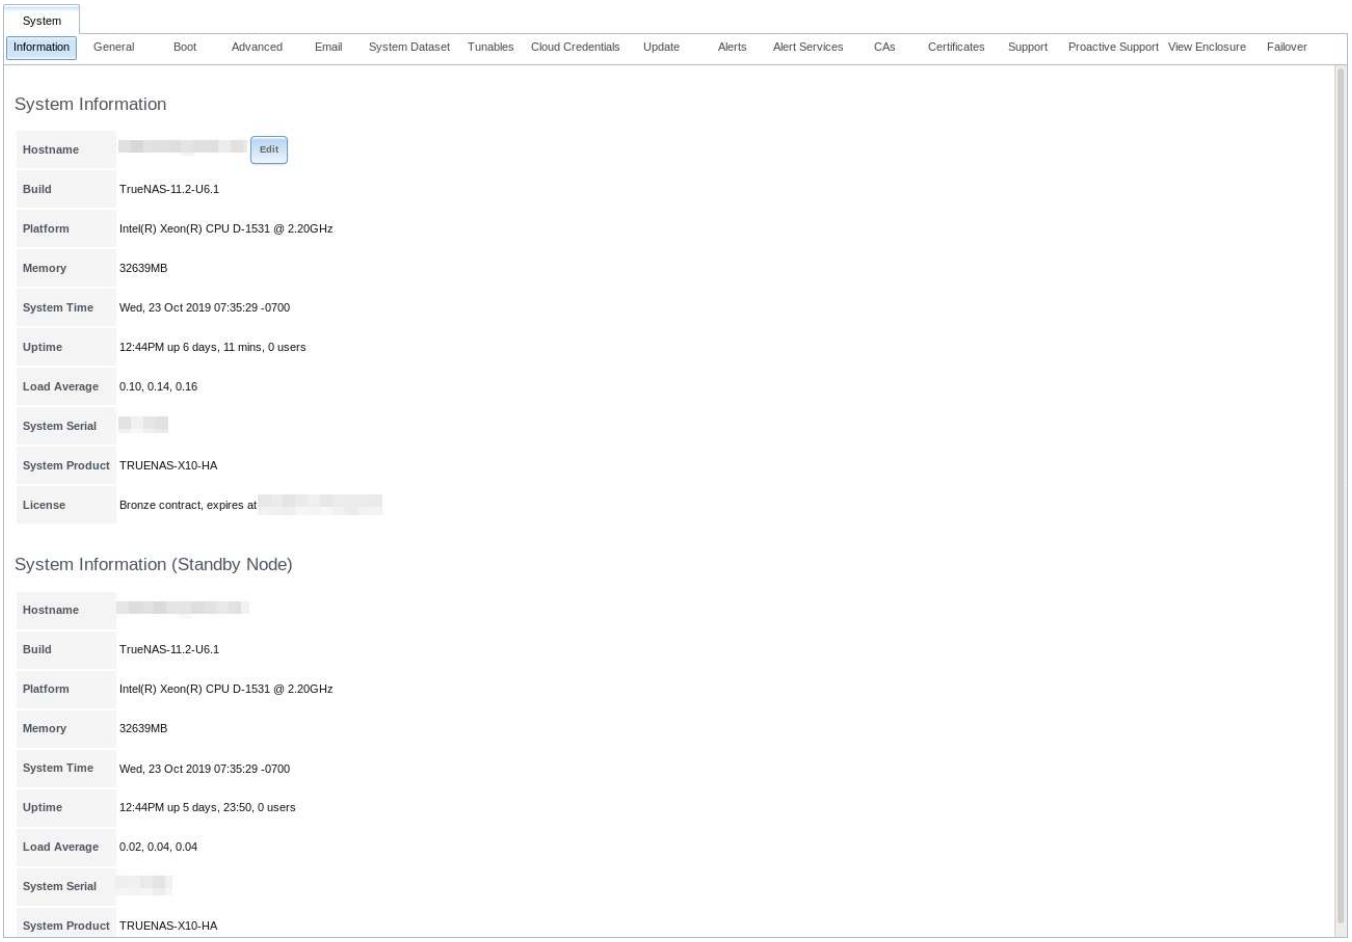

Fig. 4.1: System Information Tab

4.2 General

System → *General* is shown in Figure 4.2.

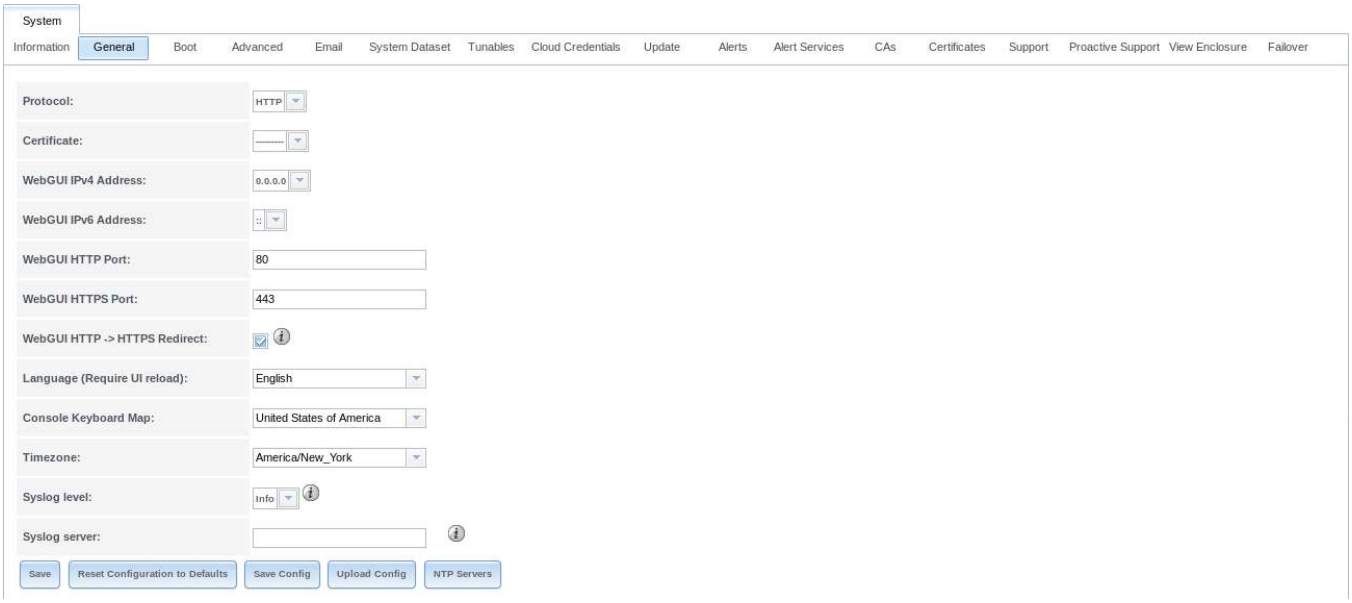

Fig. 4.2: General Screen

Table 4.1 summarizes the configurable settings in the General tab:

Table 4.1: General Configuration Settings

Setting	Value	Description
Protocol	drop-down menu	Set the web protocol to use when connecting to the web interface from a browser. To change the default *HTTP* to *HTTPS* or to *HTTP+HTTPS*, select a certificate to use in *Certificate*. If there are no certificates, first create a *CA* (page 49) then a *certificate* (page 51).
Certificate	drop-down menu	Required for *HTTPS*. Select a certificate to use for encrypted connections.
WebGUI IPv4 Address	drop-down menu	Choose a recent IP address to limit the usage when accessing the web interface. The built-in HTTP server binds to the wildcard address of *0.0.0.0* (any address) and issues an alert if the specified address becomes unavailable.
WebGUI IPv6 Address	drop-down menu	Choose a recent IPv6 address to limit the usage when accessing the web interface. The built-in HTTP server binds to any address issues an alert if the specified address becomes unavailable.
WebGUI HTTP Port	integer	Allow configuring a non-standard port for accessing the web interface over HTTP. Changing this setting can also require changing a Firefox configuration setting (https://www.redbrick.dcu.ie/~d_fens/articles/Firefox:_This_Address_is_Restricted).
WebGUI HTTPS Port	integer	Allow configuring a non-standard port for accessing the web interface over HTTPS.
WebGUI HTTP –> HTTPS Redirect	checkbox	Set to redirect *HTTP* connections to *HTTPS*. *HTTPS* must be selected in *Protocol*.
Language	drop-down menu	Select a localization.
Console Keyboard Map	drop-down menu	Select a keyboard layout.

Continued on next page

Table 4.1 – continued from previous page

Setting	Value	Description
Timezone	drop-down menu	Select a timezone.
Syslog level	drop-down menu	When *Syslog server* is defined, only logs matching this level are sent.
Syslog server	string	Enter an *IP address_or_hostname:optional_port_number* to send logs to. Configure to write log entries to both the console and the remote server.

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

This screen also contains these buttons:

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

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

Note: *SSH* (page 225) keys are not stored in the configuration database and must be backed up separately. System host keys are files with names beginning with `ssh_host_` in `/usr/local/etc/ssh/`. The root user keys are stored in `/root/.ssh`.

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

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

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 the network is using a directory service, ensure that the TrueNAS® system and the server running the directory service have been configured to use the same NTP servers.

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

Click *NTP Servers* → *Add NTP Server* to add an NTP server. Figure 4.3 shows the screen that appears. Table 4.2 summarizes the options available when adding an NTP server. ntp.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=ntp.conf) explains these options in more detail.

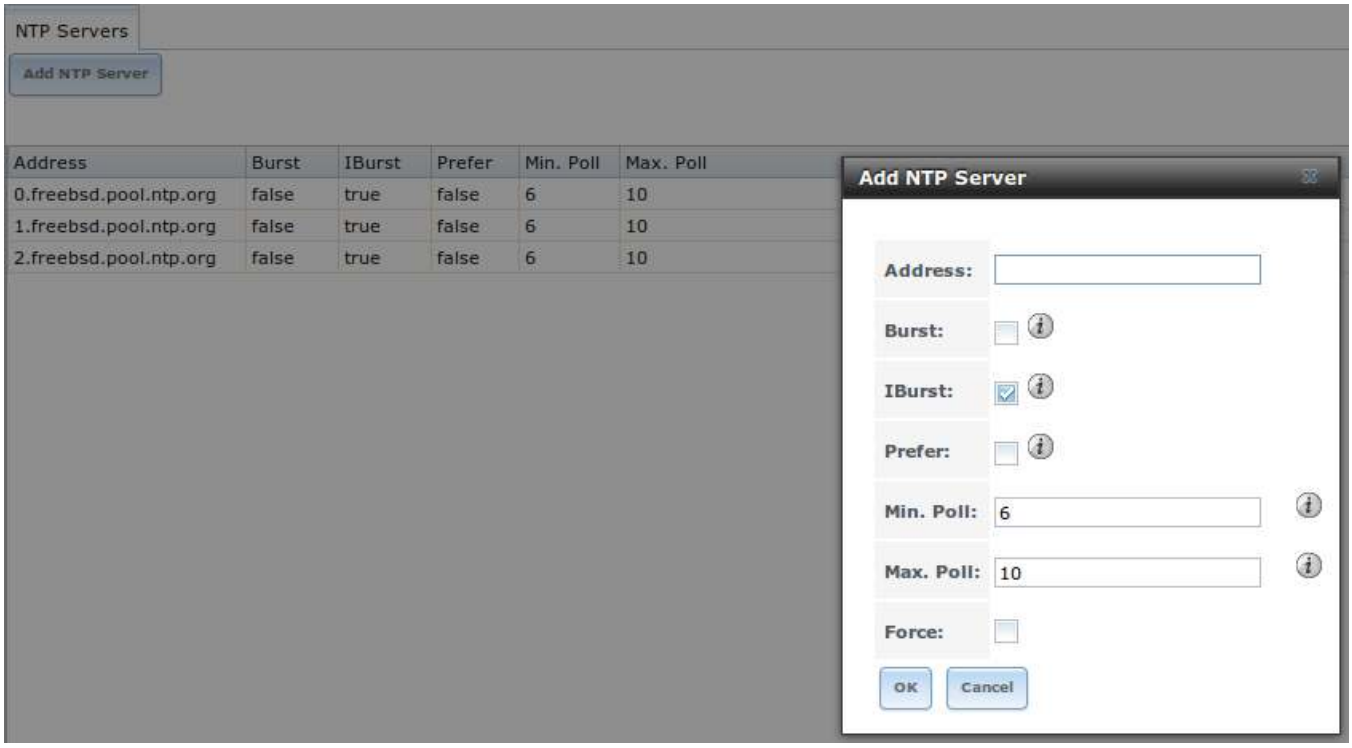

Fig. 4.3: Add an NTP Server

Table 4.2: NTP Servers Configuration Options

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

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

As seen in Figure 4.4, TrueNAS® displays the condition and statistics of the *Boot Volume*. It also shows the two boot environments that 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 *Initial-Install* boot environment can be booted into if the system needs to be returned to a non-configured version of the installation.

If the *Wizard* (page 237) was used, a third boot environment called `Wizard-date` is also created, indicating the date and time the *Wizard* (page 237) was run.

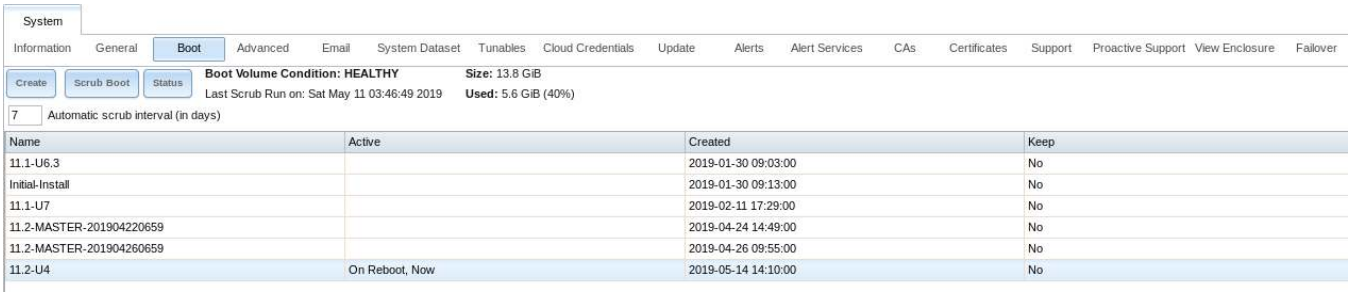

Fig. 4.4: Viewing Boot Environments

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 *Keep* for an entry if that boot environment should not be automatically pruned.

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

- **Clone:** makes a new boot environment from the selected boot environment.
- **Delete:** used to delete the highlighted entry, which also removes that entry from the boot menu. Since an activated entry cannot be deleted, this button does not appear for the active boot environment. To delete an entry that is currently activated, first activate another entry, which will clear the *On reboot* field of the currently activated entry. Note that this button does not appear for the *default* boot environment as this entry is needed to return the system to the original installation state.
- **Activate:** only appears on entries which are not currently set to *Active*. Changes the selected entry to the default boot entry on next boot. The status changes to *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.
- **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.

The buttons above the boot entries can be used to:

- **Create:** makes a new boot environment from the active environment. The active boot environment contains the text `On Reboot, Now` in the *Active* column. Only alphanumeric characters, underscores, and dashes are allowed in the name.

- **Scrub Boot:** can be used to perform a manual scrub of the boot devices. By default, the operating system device is scrubbed every 7 days. To change the default interval, change the number in the *Automatic scrub interval (in days)* field. The date and results of the last scrub are also listed in this screen. The condition of the operating system device should be listed as *HEALTHY*.
- **Status:** click this button to see the status of the operating system device. Figure 4.5, shows only one operating system device, which is *ONLINE*.

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

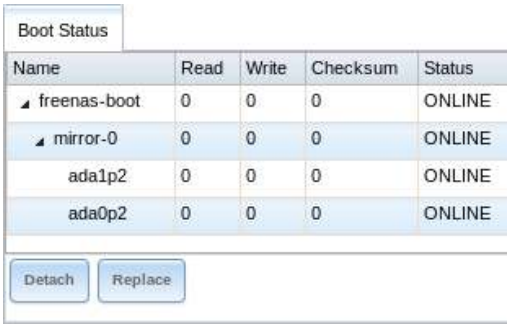

Fig. 4.5: Viewing the Status of the Operating System Device

If one of the operating system 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.

4.4 Advanced

System → *Advanced* is shown in Figure 4.6. The configurable settings are summarized in Table 4.3.

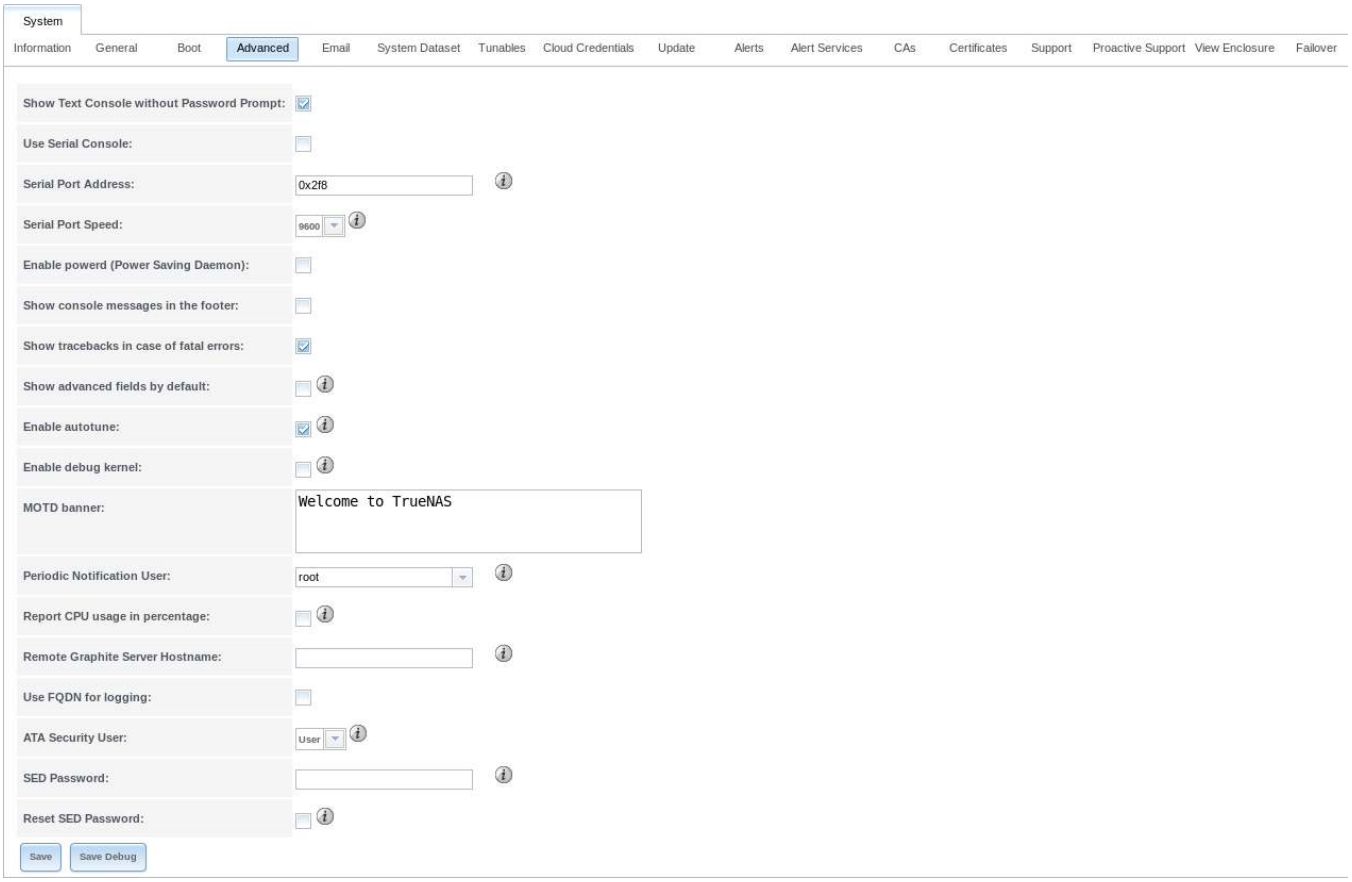

Fig. 4.6: Advanced Screen

Table 4.3: Advanced Configuration Settings

Setting	Value	Description
Show Text Console without Password Prompt	checkbox	Set for the system to immediately display the text console after booting. Unset to require logging into the system before the console menu is shown.
Use Serial Console	checkbox	**Do not** enable this option if the serial port is disabled.
Serial Port Address	string	Enter a serial port address in hex.
Serial Port Speed	drop-down menu	Select the speed used by the serial port.
Enable powerd (Power Saving Daemon)	checkbox	powerd(8) (https://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	Set to display console messages in real time at the bottom of the browser. Click the console to bring up a scrollable screen. Set *Stop refresh* in the scrollable screen to pause updating, and deselect the option to continue to watch the messages as they occur.
Show tracebacks in case of fatal errors	checkbox	Open a pop-up of diagnostic information when a fatal error occurs.
Show advanced fields by default	checkbox	Show *Advanced Mode* fields by default.
Enable autotune	checkbox	Enable an *Autotune* (page 31) script which attempts to optimize the system based on the installed hardware. *Warning*: Autotuning is only used as a temporary measure and is not a permanent fix for system hardware issues.

Continued on next page

Table 4.3 – continued from previous page

Setting	Value	Description
Enable debug kernel	checkbox	Use a debug version of the kernel on the next boot.
MOTD banner	string	This message is shown when a user logs in with SSH.
Periodic Notification User	drop-down menu	Choose a user to receive security output emails. This output runs nightly but only sends email when the system reboots or encounters an error.
Report CPU usage in percentage	checkbox	Display CPU usage as percentages in *Reporting* (page 235).
Remote Graphite Server hostname	string	IP address or hostname of a remote server running Graphite (http://graphiteapp.org/).
Use FQDN for logging	checkbox	Include the Fully-Qualified Domain Name in logs to precisely identify systems with similar hostnames.
ATA Security User	drop-down menu	User passed to `camcontrol security -u` for unlocking *Self-Encrypting Drives* (page 31). Values are *User* or *Master*.
SED Password	string	Global password used to unlock *Self-Encrypting Drives* (page 31).
Reset SED Password	checkbox	Select to clear the *Password for SED* column of *Storage* → *View Disks*.

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 compressed .tgz text file.

4.4.1 Autotune

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

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

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

4.4.2 Self-Encrypting Drives

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

These SED specifications are supported:

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

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

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

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

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

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

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

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

4.4.2.1 Deploying SEDs

Run `sedutil-cli --scan` in the *Shell* (page 244) to detect and list devices. The second column of the results identifies the drive type:

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

Example:

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

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

Setting a global password for SEDs

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

Now the SEDs must be configured with this password. Go to the *Shell* (page 244) and enter `sedhelper setup password`, where *password* is the global password entered in *System* → *Advanced* → *SED Password*.

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

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

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

Creating separate passwords for each SED

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

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

The SED must be configured to use the new password. Go to the *Shell* (page 244) and enter `sedhelper setup --disk da1 password`, where *da1* is the SED to configure and *password* is the created password from *Storage* → *Volumes* → *View Disks* → *Edit* → *Password for SED*.

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

Danger: Remember SED passwords! If the SED password is lost, SEDs cannot be unlocked and their data is unavailable. While it is possible to specify the PSID number on the label of the device with `sedutil-cli`, doing so **erases the contents** of the device rather than unlock it. Always record SED passwords whenever they are configured or modified and store them in a secure place!

4.4.2.2 Check SED Functionality

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

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

To verify SED locking is working correctly, go to the *Shell* (page 244). Enter `sedutil-cli --listLockingRange 0 password dev/da1`, where *da1* is the SED and *password* is the global or individual password for that SED. The command returns `ReadLockEnabled: 1`, `WriteLockEnabled: 1`, and `LockOnReset: 1` for drives with locking enabled:

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

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

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 247) events are also emailed to the *root* user account. Problems with *Scrubs* (page 134) are reported separately in an email sent at 03:00AM.

Note: *S.M.A.R.T.* (page 218) reports are mailed separately to the address configured in that service.

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

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

Additional configuration is performed with *System* → *Email*, shown in Figure 4.7.

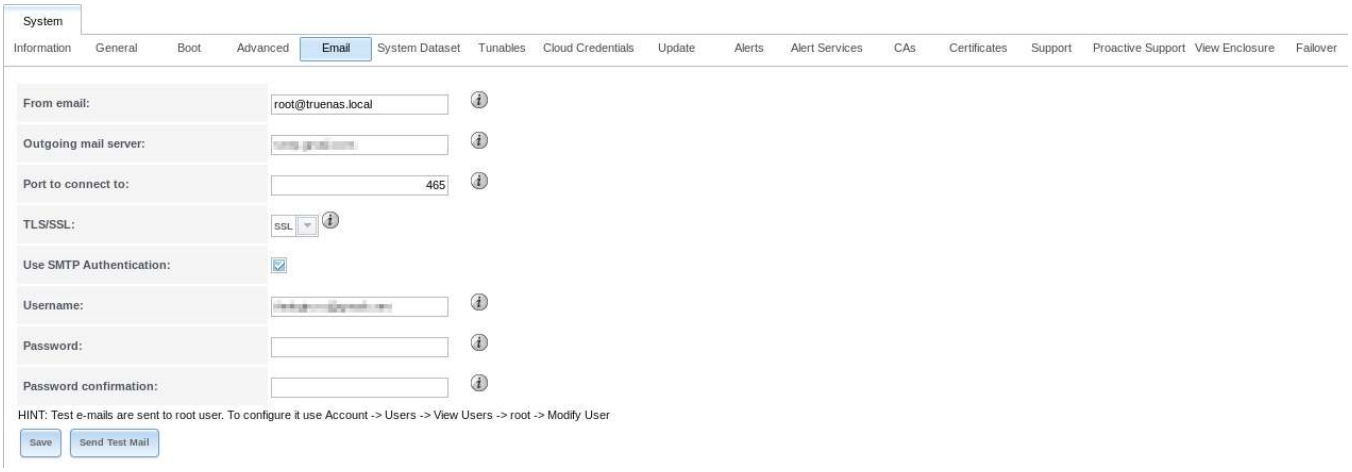

Fig. 4.7: Email Screen

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 make filtering mail on the receiving system easier. The friendly name is set like this: `Friendly Name <address@example.com>`
Outgoing mail server	string or IP address	Hostname or IP address of SMTP server used 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	Choose an encryption type. Choices are *Plain*, *SSL*, or *TLS*

Continued on next page

Table 4.4 – continued from previous page

Setting	Value	Description
Use SMTP Authentication	checkbox	Enable or disable SMTP AUTH (https://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN SASL. If enabled, enter the required *Username* and *Password*.
Username	string	Enter the SMTP username if the SMTP server requires authentication.
Password	string	Enter the SMTP password if the SMTP server requires authentication. Only plain text characters (7-bit ASCII) are allowed in passwords. UTF or composed characters are not allowed.
Password Confirmation	string	Confirm the SMTP password.

Click the *Send Test Mail* button to verify that the configured email settings are working. If the test email fails, double-check that the *E-mail* field of the *root* user is correctly configured by clicking the *Modify User* button for the *root* account in *Account* → *Users* → *View Users*.

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

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

4.6 System Dataset

System → *System Dataset*, shown in Figure 4.8, is used to select the pool which contains the persistent system dataset. The system dataset stores debugging core files and Samba4 metadata such as the user or 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.

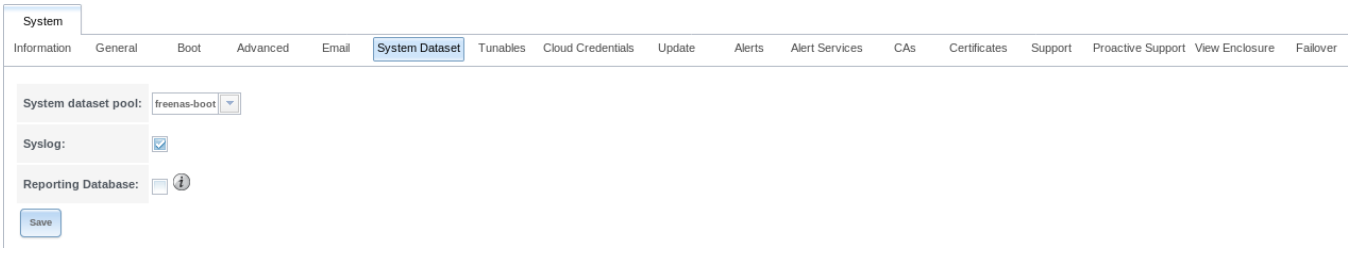

Fig. 4.8: System Dataset Screen

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

Moving the system dataset also requires rebooting the passive storage controller for *High Availability* (page 58) TrueNAS® systems and restarting the *SMB* (page 219) service. A dialog warns that the SMB service must be restarted, causing a temporary outage of any active SMB connections.

System logs can also be stored on the system dataset. Storing this information on the system dataset is recommended when large amounts of data is being generated and the system has limited memory or a limited capacity operating system device. Set *Syslog* to store system logs on the system dataset. Leave unset to store system logs in `/var` on the operating system device.

Set *Reporting Database* to store *Reporting* (page 235) data on the system dataset. Leave unset to create a `/temp` disk in RAM to store the reporting database.

Click *Save* to save changes.

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

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

4.7 Tunables

System → *Tunables* can be used to manage:

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

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

Since sysctl, loader, and rc.conf values are specific to the kernel parameter to be tuned, the driver to be loaded, or the service to configure, descriptions and suggested values can be found in the man page for the specific driver and in many sections of the FreeBSD Handbook (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/).

To add a loader, sysctl, or `rc.conf` option, go to *System* → *Tunables* → *Add Tunable*, to access the screen shown in Figure 4.9.

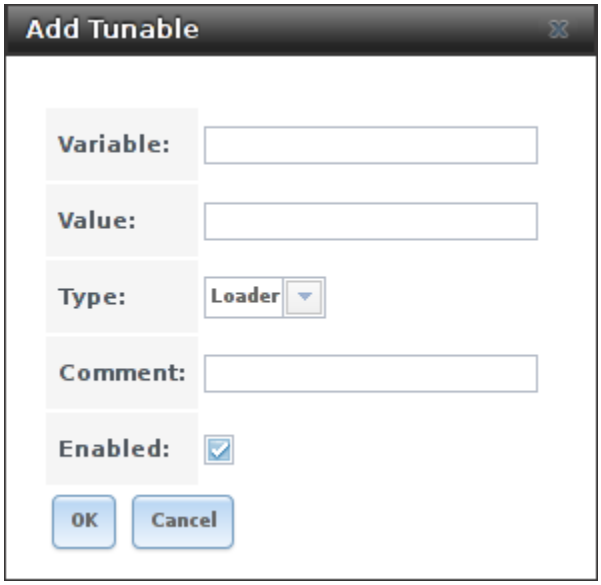

Fig. 4.9: Adding a Tunable

Table 4.5 summarizes the options when adding a tunable.

Table 4.5: Adding a Tunable

Setting	Value	Description
Variable	string	The name of the sysctl or driver to load.
Value	integer or string	Set a value for the *Variable*. Refer to the man page for the specific driver or the FreeBSD Handbook (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/) for suggested values.
Type	drop-down menu	Choices are *Loader*, *rc.conf*, or *Sysctl*.
Comment	string	Enter a userful description of this tunable.
Enabled	checkbox	Unset this option to disable the tunable without deleting it.

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

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

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

The web interface does not display the sysctls that are pre-set when TrueNAS® is installed. TrueNAS® 11.2 ships with the sysctls set:

```
kern.metadelay=3
kern.dirdelay=4
kern.filedelay=5
kern.coredump=1
net.inet.carp.preempt=1
debug.ddb.textdump.pending=1
vfs.nfsd.tcpcachetimeo=300
```

```
vfs.nfsd.tcphighwater=150000
vfs.zfs.vdev.larger_ashift_minimal=0
net.inet.carp.senderr_demotion_factor=0
net.inet.carp.ifdown_demotion_factor=0
```

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

The web interface does not display the loaders that are pre-set when TrueNAS® is installed. TrueNAS® 11.2 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"
freenas_sysctl_load="YES"
hint.isp.0.topology="nport-only"
hint.isp.1.topology="nport-only"
hint.isp.2.topology="nport-only"
hint.isp.3.topology="nport-only"
module_path="/boot/kernel;/boot/modules;/usr/local/modules"
net.inet6.ip6.auto_linklocal="0"
net.inet.tcp.reass.maxqueuelen=1436
vfs.zfs.vol.mode=2
kern.geom.label.disk_ident.enable=0
kern.geom.label.ufs.enable=0
kern.geom.label.ufsid.enable=0
kern.geom.label.reiserfs.enable=0
kern.geom.label.ntfs.enable=0
kern.geom.label.msdosfs.enable=0
kern.geom.label.ext2fs.enable=0
hint.ahciem.0.disabled="1"
hint.ahciem.1.disabled="1"
kern.msgbufsize="524288"
hw.mfi.mrsas_enable="1"
hw.usb.no_shutdown_wait=1
vfs.nfsd.fha.write=0
vfs.nfsd.fha.max_nfsds_per_fh=32
kern.ipc.nmbclusters="262144"
kern.hwpmc.nbuffers="4096"
kern.hwpmc.nsamples="4096"
hw.memtest.tests="0"
vfs.zfs.trim.enabled="0"
kern.cam.ctl.ha_mode=2
hint.ntb_hw.0.config="ntb_pmem:1:4:0,ntb_transport"
hint.ntb_transport.0.config=":3"
hw.ntb.msix_mw_idx="-1"
```

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

The ZFS version used in 11.2 deprecates these tunables:

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

```
vfs.zfs.no_write_throttle
```

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

4.8 Cloud Credentials

TrueNAS® can use cloud services for features like *Cloud Sync* (page 61). The credentials to provide secure connections with cloud services are entered here. Amazon S3, Backblaze B2, Box, Dropbox, FTP, Google Cloud Storage, Google Drive, HTTP, hubiC, Mega, Microsoft Azure Blob Storage, Microsoft OneDrive, pCloud, SFTP, WebDAV, and Yandex are supported.

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

Warning: Cloud Credentials are stored in encrypted form. To be able to restore Cloud Credentials from a *saved configuration* (page 24), *Export Password Secret Seed* must be set when saving that configuration.

Select *System* → *Cloud Credentials* to see the screen shown in Figure 4.10.

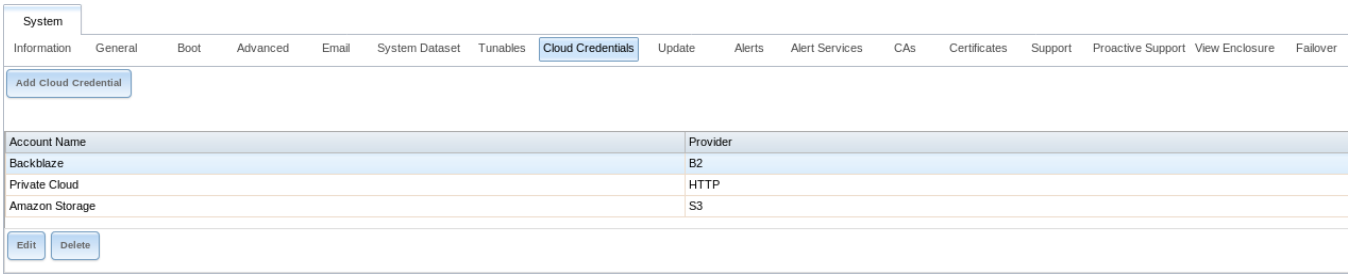

Fig. 4.10: Cloud Credentials List

The list shows the *Account Name* and *Provider* for each credential. There are options to *Edit* and *Delete* a credential after selecting it. Click *Add Cloud Credential* to display the dialog shown in Figure 4.11.

Fig. 4.11: Adding Cloud Credentials

Amazon S3 options are shown by default. Enter a descriptive and unique name for the cloud credential in the *Account Name* field, then select a *Provider*. The remaining options vary by provider, and are shown in Table 4.6.

Table 4.6: Cloud Credential Options

Provider	Setting	Description
Amazon S3	Access Key ID	Enter the Amazon Web Services Key ID. This is found on Amazon AWS (https://aws.amazon.com) by going through My account –> Security Credentials –> Access Keys.
Amazon S3	Secret Access Key	Enter the Amazon Web Services password. If the Secret Access Key cannot be found or remembered, go to My Account –> Security Credentials –> Access Keys and create a new key pair.
Amazon S3	Endpoint URL	Leave blank when using AWS as the available buckets are fetched dynamically. Only enter an Endpoint URL (https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteEndpoints.html) if using *custom* S3 API. URL general format: *bucket-name.s3-website-region.amazonaws.com*. Refer to the AWS Documentation for a list of Simple Storage Service Websites Endpoints (https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_website_region_enc
Amazon S3	Endpoint does not support regions	Skip automatic detection of the *Endpoint URL* region. Set this when configuring a custom *Endpoint URL*.
Amazon S3	Use v2 signatures	Force using Signature Version 2 (https://docs.aws.amazon.com/general/latest/gr/signature-version-2.html) to sign API requests. Set this when configuring a custom *Endpoint URL*.

Continued on next page

Table 4.6 – continued from previous page

Provider	Setting	Description
Backblaze B2	Account ID or Application Key ID, Application Key	Enter the Account ID and Master Application Key (https://help.backblaze.com/hc/en-us/articles/224991568-Where-can-I-find-my-Account-ID-and-Application-Key-) for the Backblaze B2 account. These are visible after logging into the account, clicking *Buckets*, and clicking *Show Account ID and Application Key*. An *Application Key* with limited permissions can be used in place of the *Account ID*. Create a new Application Key, enter the key string in the *Application Key* field, and replace the *Account ID* with the *keyID*.
Box	Automatic config, OAuth Client ID, OAuth Client Secret, Access Token	Configured with *Open Authentication* (page 42).
Dropbox	Automatic config OAuth Client ID, OAuth Client Secret, Access Token	Configured with *Open Authentication* (page 42). The access token can be manually created by going to the Dropbox App Console (https://www.dropbox.com/developers/apps). After creating an app, go to *Settings* and click *Generate* under the Generated access token field.
FTP	Host, Port	Enter the FTP host and port.
FTP	Username, Password	Enter the FTP username and password.
Google Cloud Storage	Service Account	*Browse* to the location of the saved Google Cloud Storage key and select it.
Google Drive	Automatic config, OAuth Client ID, OAuth Client Secret, Access Token, Team Drive ID	*OAuth Client ID*, *OAuth Client Secret*, and *Access Token* are configured with *Open Authentication* (page 42). The *Team Drive ID* is only used when connecting to a Team Drive (https://developers.google.com/drive/api/v3/reference/teamdrives). The ID is also the ID of the top level folder of the Team Drive.
HTTP	URL	Enter the URL.
hubiC	Access Token	Enter the access token. See the Hubic guide (https://api.hubic.com/sandbox/) for instructions to obtain an access token.
Mega	Username, Password	Enter the Mega (https://mega.nz) username and password.
Microsoft Azure Blob Storage	Account Name, Account Key	Enter the Azure Blob Storage account name and key.
Microsoft OneDrive	Automatic config, OAuth Client ID, OAuth Client Secret, Access Token, Drive Account Type, Drive ID	*OAuth Client ID*, *OAuth Client Secret*, and *Access Token* are configured with *Open Authentication* (page 42). Choose the account type: *PERSONAL*, *BUSINESS*, or SharePoint (https://products.office.com/en-us/sharepoint/collaboration) *DOCUMENT_LIBRARY*. To find the *Drive ID*, log in to the OneDrive account (https://onedrive.live.com) and copy the string that appears in the browser address bar after `cid=`. Example: `https://onedrive.live.com/?id=root&cid=12A34567B89C10D1`, where *12A34567B89C10D1* is the drive ID.
pCloud	Automatic config, OAuth Client ID, OAuth Client Secret, Access Token	Configured with *Open Authentication* (page 42).
SFTP	Host, Port	Enter the SFTP host and port.

Continued on next page

Table 4.6 – continued from previous page

Provider	Setting	Description
SFTP	Username, Password, PEM-encoded private key file path	Enter the SFTP username, password, and PEM-encoded private key file path.
WebDAV	URL, WebDAV Service	Enter URL and use the dropdown to select the WebDAV service.
WebDAV	Username, Password	Enter the username and password.
Yandex	Automatic config, OAuth Client ID, OAuth Client Secret, Access Token	Configured with *Open Authentication* (page 42).

For Amazon S3, *Access Key* and *Secret Key* are shown. These values are found on the Amazon AWS website by clicking on the account name, then *My Security Credentials* and *Access Keys (Access Key ID and Secret Access Key)*. Copy the Access Key value to the TrueNAS® Cloud Credential *Access Key* field, then enter the *Secret Key* value saved when the key pair was created. If the Secret Key value is unknown, a new key pair can be created on the same Amazon screen.

The Google Cloud Storage *JSON Service Account Key* is found on the Google Cloud Platform Console (https://console.cloud.google.com/apis/credentials). Open Authentication (OAuth) (https://openauthentication.org/) is used with some cloud providers. These providers have an *Automatic config* link that opens a dialog to log in to that provider and fill the TrueNAS® *OAuth Client ID*, *OAuth Client Secret*, and *Access Token* fields with valid credentials.

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

4.9 Update

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

4.9.1 Preparing for Updates

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

For individual support during an upgrade, open a ticket with or call *iXsystems Support* (page 9) to schedule an upgrade. 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 iXsystems 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 27) to remove unneeded boot environments.

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

Available ZFS version upgrades are indicated by an *Alert* (page 247) 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®.

4.9.2 Updates and Trains

Cryptographically signed update files are used to update TrueNAS®. Update files provide flexibility in deciding when to upgrade the system. *Boot environments* (page 45) make it possible to test an update.

Figure 4.12 shows an example of the *System* → *Update* screen.

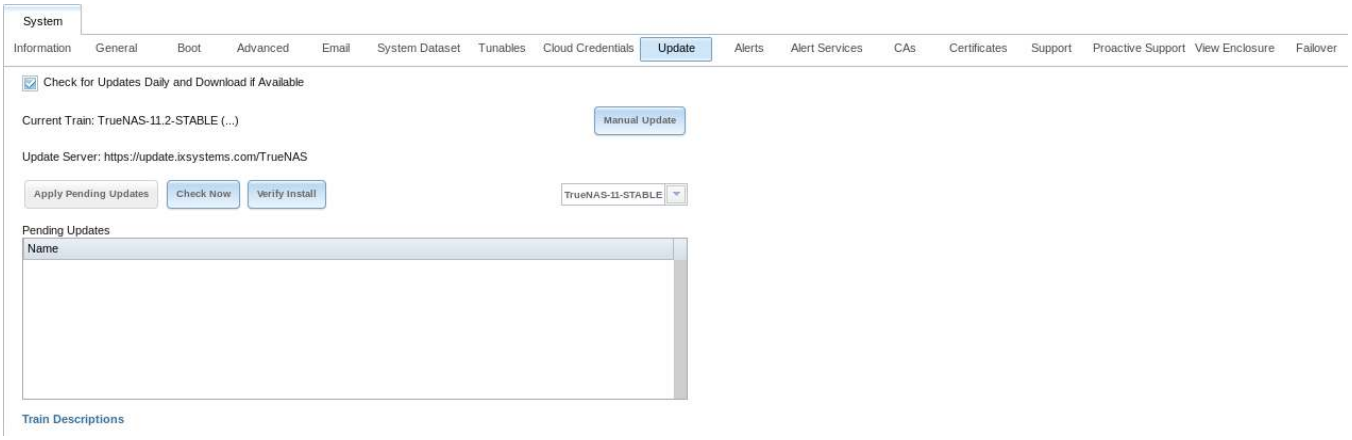

Fig. 4.12: Update Options

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

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

Several trains are available for updates. Update trains are labeled with a numeric version and a short description.

These update trains are available:

For Production Use

- **TrueNAS-11-STABLE** (Recommended)

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

Legacy Versions

- **TrueNAS-9.10-STABLE**

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

- **TrueNAS-9.3-STABLE**

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

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

4.9.3 Checking for Updates

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

4.9.4 Applying Updates

Make sure the system is in a low-usage state as described above in *Preparing for Updates* (page 42).

Click the *OK* button to immediately download and install an update. 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 27) 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.

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

Updates can also be downloaded and applied later. To do so, unset the *Apply updates after downloading* option 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 4.12. When ready to apply the previously downloaded updates, click the *Apply Pending Updates* button. Remember that the system reboots 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.

4.9.5 Manual Updates

Updates can be manually downloaded as a file with a name ending in `-manual-update-unsigned.tar`. Find a `.tar` file with the desired version at https://download.freenas.org/. 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*.

There is also an option to back up the system configuration before updating. Click *Click here* and select any options to export in the configuration file. Click *OK* to open a popup window to save the system configuration. A progress dialog is displayed during the update. **Do not** interrupt the update.

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

4.9.6 Updating from the Shell

Updates can also be performed from the *Shell* (page 244) 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`.

4.9.7 Updating an HA System

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

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

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

4.9.8 If Something Goes Wrong

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

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

4.9.9 Upgrading a ZFS Pool

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

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

- the pool upgrade is a one-way street, meaning that **if you change your mind you cannot go back to an earlier ZFS version or downgrade to an earlier version of the software that does not support those ZFS features.**
- before performing any operation that may affect the data on a storage disk, **always back up all data first and verify the integrity of the backup.** While it is unlikely that the pool upgrade will affect the data, it is always better to be safe than sorry.
- upgrading a ZFS pool is **optional**. Do not upgrade the pool if the the possibility of reverting to an earlier version of TrueNAS® or repurposing the disks in another operating system that supports ZFS is desired. It is not necessary to upgrade the pool unless the end user has a specific need for the newer *ZFS Feature Flags* (page 253). If a pool is upgraded to the latest feature flags, it will not be possible to import that pool into another operating system that does not yet support those feature flags.

To perform the ZFS pool upgrade, go to *Storage* → *Volumes* → *View Volumes* and highlight the volume (ZFS pool) to upgrade. Click the *"Up Arrow" (Upgrade)* button as shown in Figure 4.13.

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

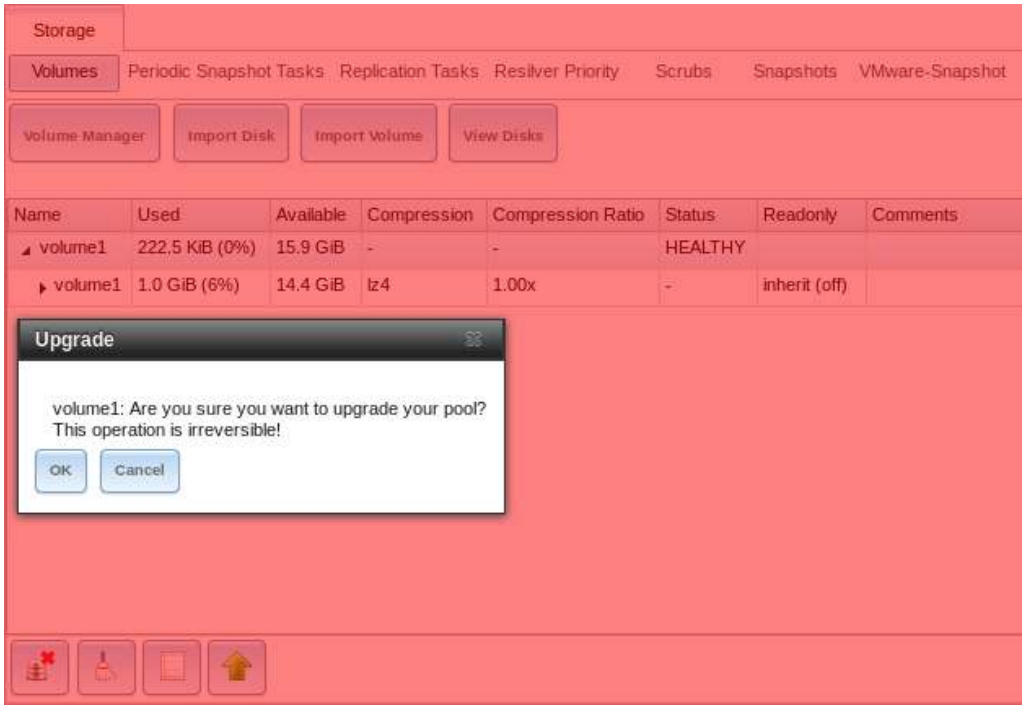

Fig. 4.13: Upgrading a ZFS Pool

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

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

4.10 Alerts

System → *Alerts* displays the default notification frequency for each type of *Alert* (page 247). An example is seen in Figure 4.14.

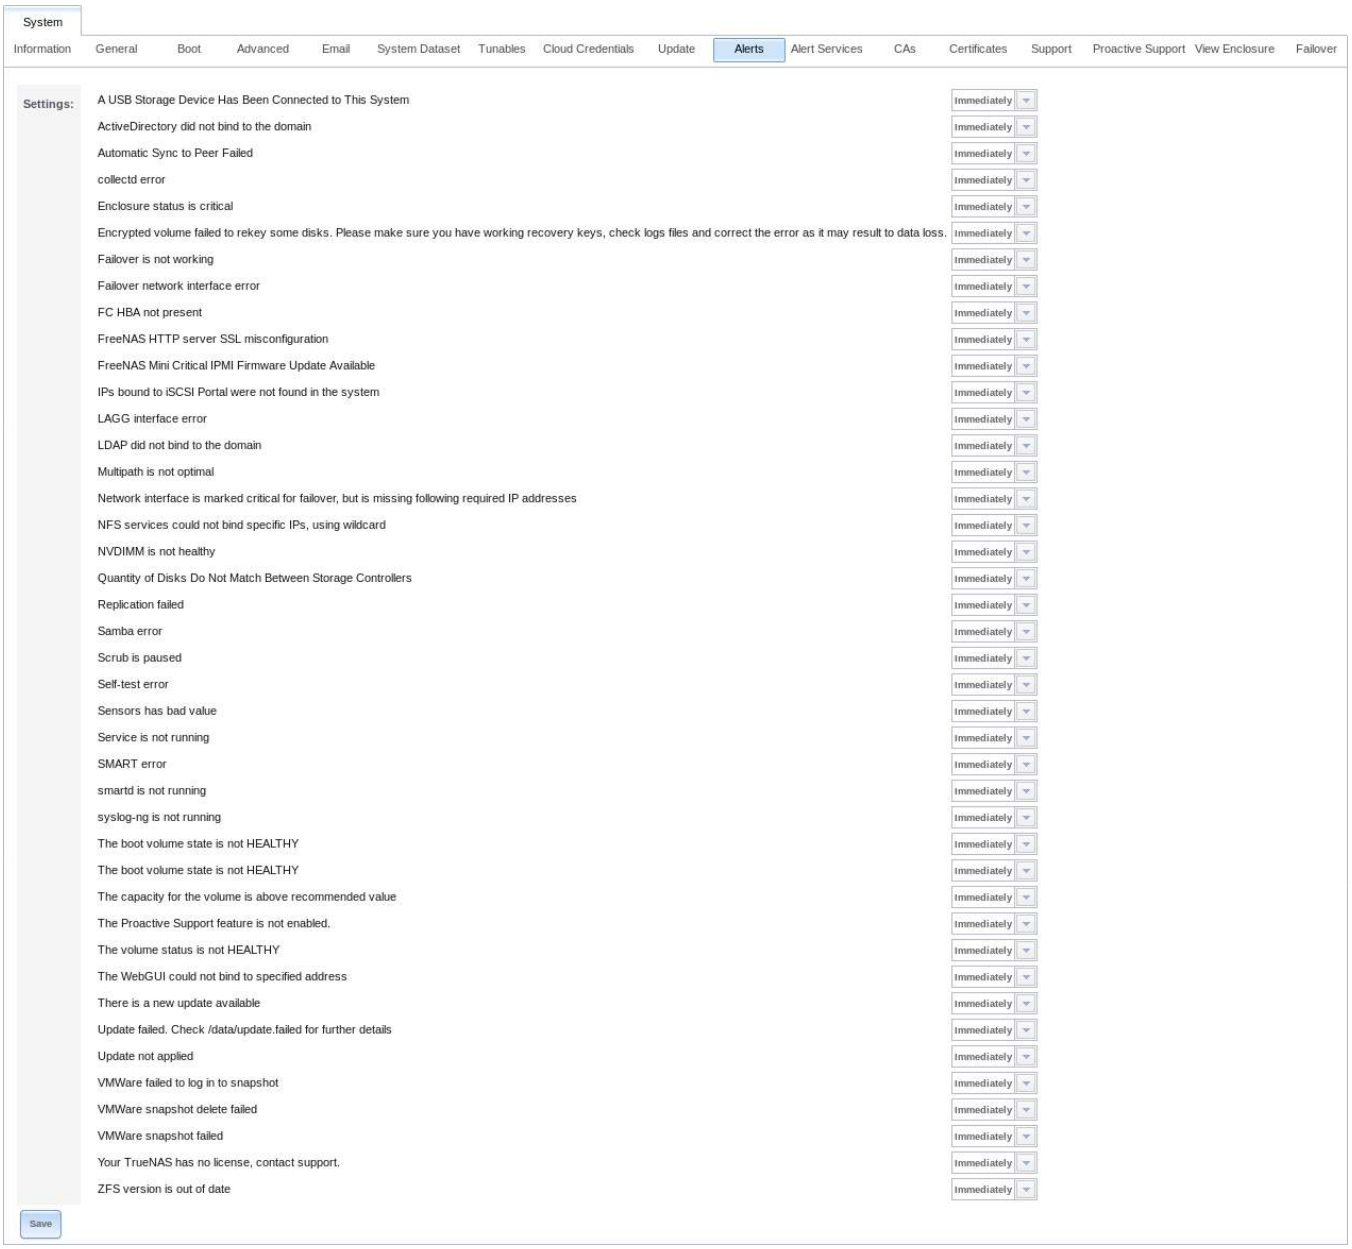

Fig. 4.14: Configure Alert Notification Frequency

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

To configure where to send alerts, use *Alert Services* (page 47).

4.11 Alert Services

TrueNAS® can use a number of methods to notify the administrator of system events that require attention. These events are system *Alerts* (page 247) marked *WARN* or *CRITICAL*.

Currently available alert services:

- AWS-SNS (https://aws.amazon.com/sns/)
- E-Mail

- Hipchat (https://www.atlassian.com/software/hipchat)
- InfluxDB (https://www.influxdata.com/)
- Mattermost (https://about.mattermost.com/)
- OpsGenie (https://www.opsgenie.com/)
- PagerDuty (https://www.pagerduty.com/)
- SNMP Trap (https://www.freebsd.org/cgi/man.cgi?query=snmptrap)
- Slack (https://slack.com/)
- VictorOps (https://victorops.com/)

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

Select *System* → *Alert Services* to show the Alert Services screen. Click *Add Service* to display the dialog shown in Figure 4.15.

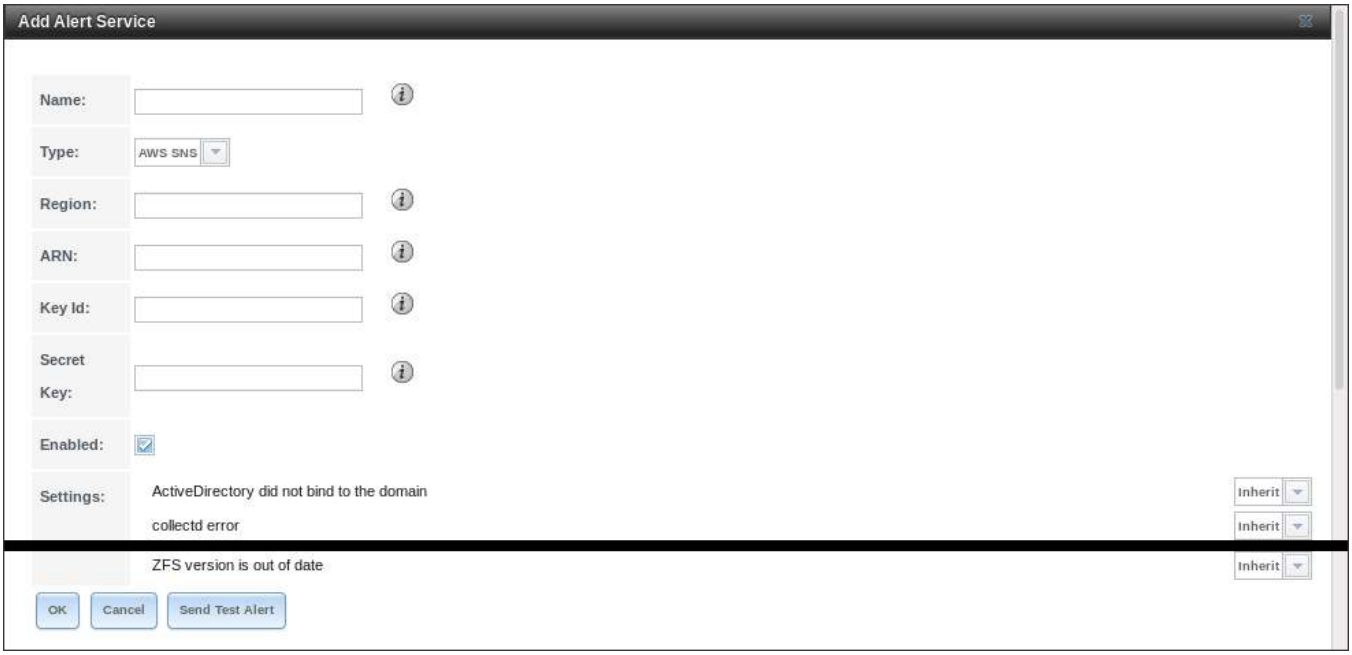

Fig. 4.15: Add Alert Service

Enter a specific *Name* for the new alert service. The *Type* drop-down menu is used to pick a specific alert service. The *Settings* area allows configuring when specific alerts will trigger. Options are to *Inherit* the setting from *Alerts* (page 46) or generate the alert *Immediately*, *Hourly*, *Daily*, or *Never*. The fields shown in the rest of the dialog change to those required by that service.

Click *Send Test Alert* to test the current selections. Click *OK* to save the new alert service. To send a test alert using an existing service, highlight an alert entry, click *Edit*, and click *Send Test Alert*.

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

Alert services are deleted from this list by clicking them and then clicking *Delete* at the bottom of the window. To disable an alert service, click *Edit* and unset *Enabled*.

4.12 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 is imported with *Import CA*, or a new CA created on the TrueNAS® system and used on the LDAP server also.

Figure 4.16 shows the screen after clicking *System* → *CAs*.

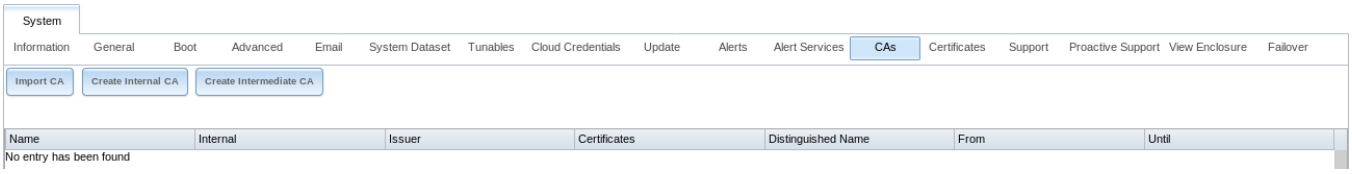

Fig. 4.16: Initial CA Screen

If the organization already has a CA, the CA certificate and key can be imported. Click the *Import CA* button to open the configuration screen shown in Figure 4.17. The configurable options are summarized in Table 4.7.

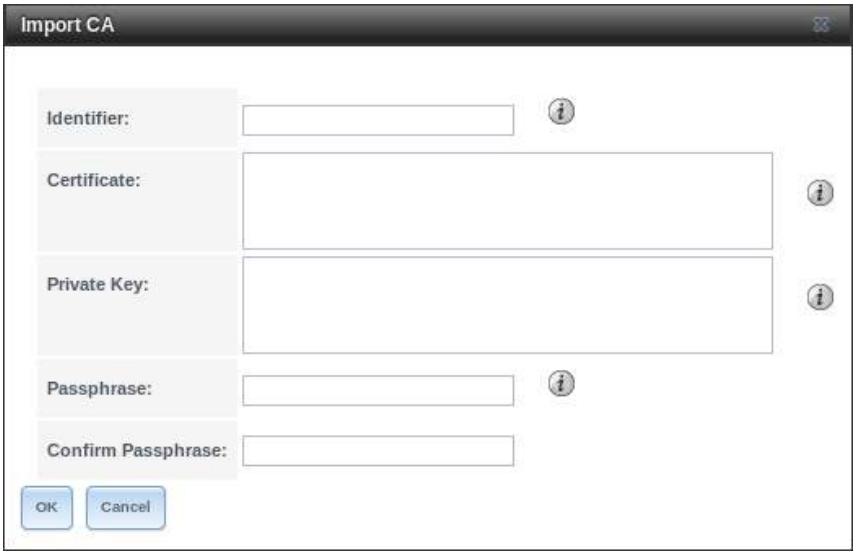

Fig. 4.17: Importing a CA

Table 4.7: Importing a CA Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore (_), and dash (-) characters.
Certificate	string	Paste in the certificate for the CA.
Private Key	string	If there is a private key associated with the *Certificate*, paste it here. Private keys must be at least 1024 bits long.
Passphrase	string	If the *Private Key* is protected by a passphrase, enter it here and repeat it in the *Confirm Passphrase* field.

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

To create a CA for internal use only, click the *Create Internal CA* button which will open the screen shown in Figure 4.18.

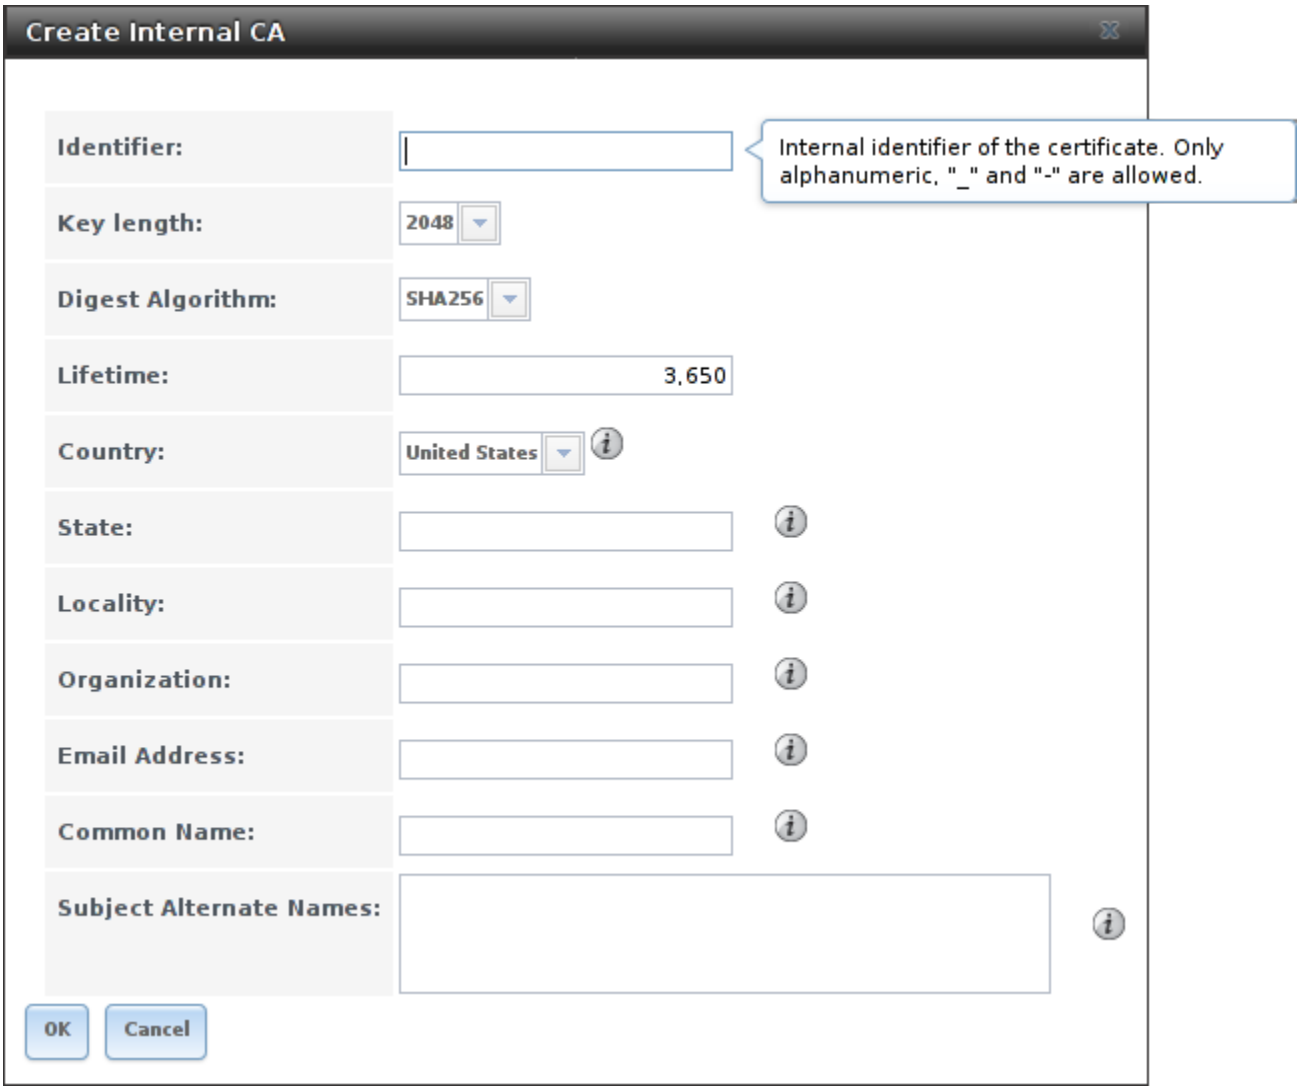

Fig. 4.18: Creating an Internal CA

The configurable options are described in Table 4.8. When completing the fields for the certificate authority, supply the information for the organization.

Table 4.8: Internal CA Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore (_), and dash (–) characters.
Key Length	drop-down menu	For security reasons, a minimum of *2048* is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algorithm.
Lifetime	integer	The lifetime of the CA is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	Enter the state or province of the organization.
Locality	string	Enter the location of the organization.
Organization	string	Enter the name of the company or organization.

Continued on next page

Table 4.8 – continued from previous page

Setting	Value	Description
Email Address	string	Enter the email address for the person responsible for the CA.
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The *Common Name* **must** be unique within a certificate chain.
Subject Alternate Names	string	Multi-domain support. Enter additional domain names and separate them with a space.

To create an intermediate CA which is part of a certificate chain, click *Create Intermediate CA*. This screen adds one more option to the screen shown in Figure 4.18:

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

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

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

- **Sign CSR:** used to sign internal Certificate Signing Requests created using *System → Certificates → Create Certificate Signing Request*.
- **Export Certificate:** prompts to browse to the location to save a copy of the CA 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 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.

4.13 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 49).

Figure 4.19 shows the initial screen after clicking *System → Certificates*.

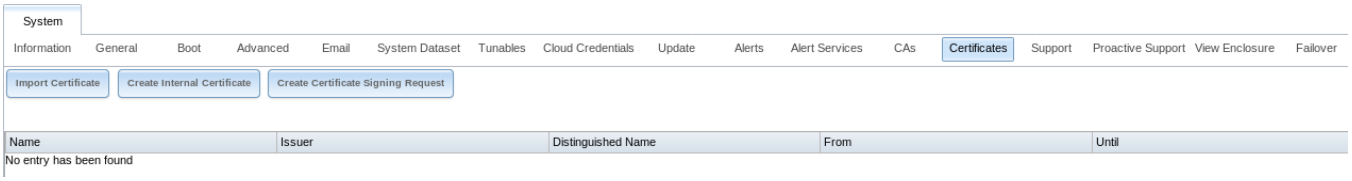

Fig. 4.19: Initial Certificates Screen

To import an existing certificate, click *Import Certificate* to open the configuration screen shown in Figure 4.20. When importing a certificate chain, paste the primary certificate, followed by any intermediate certificates, followed by the root CA certificate.

On TrueNAS® *High Availability (HA)* (page 58) systems, the imported certificate must include the IP addresses or DNS hostnames of both nodes and the CARP virtual IP address. These IP addresses or DNS hostnames can be placed in the *Subject Alternative Name* (SAN) x509 extension field of the certificate being imported.

The configurable options are summarized in Table 4.9.

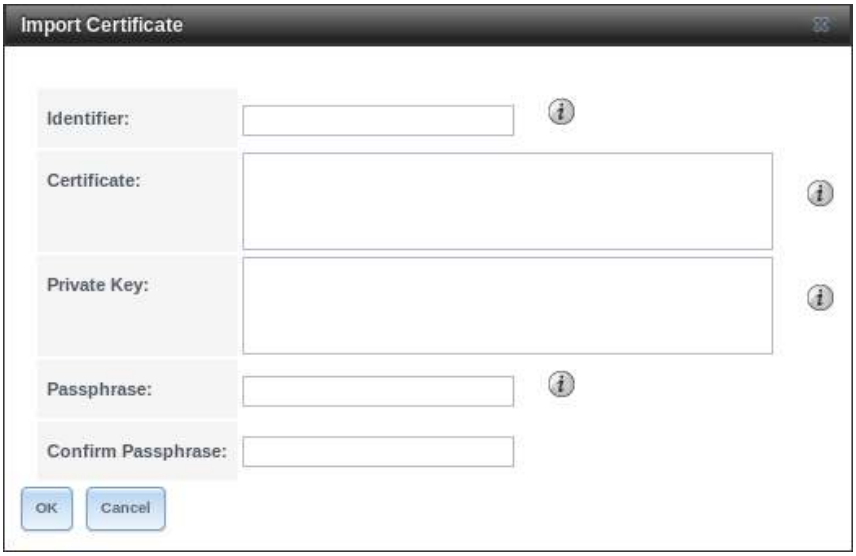

Fig. 4.20: Importing a Certificate

Table 4.9: Certificate Import Options

Setting	Value	Description
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric, underscore (_), and dash (-) characters.
Certificate	string	Paste the contents of the certificate.
Private Key	string	Paste the private key associated with the certificate. Private keys must be at least 1024 bits long.
Passphrase	string	If the private key is protected by a passphrase, enter it here and repeat it in the *Confirm Passphrase* field.

To create a new self-signed certificate, click the *Create Internal Certificate* button to see the screen shown in Figure 4.21. The configurable options are summarized in Table 4.10. When completing the fields for the certificate authority, use the information for the organization. Since this is a self-signed certificate, use the CA that was imported or created with *CAs* (page 49) as the signing authority.

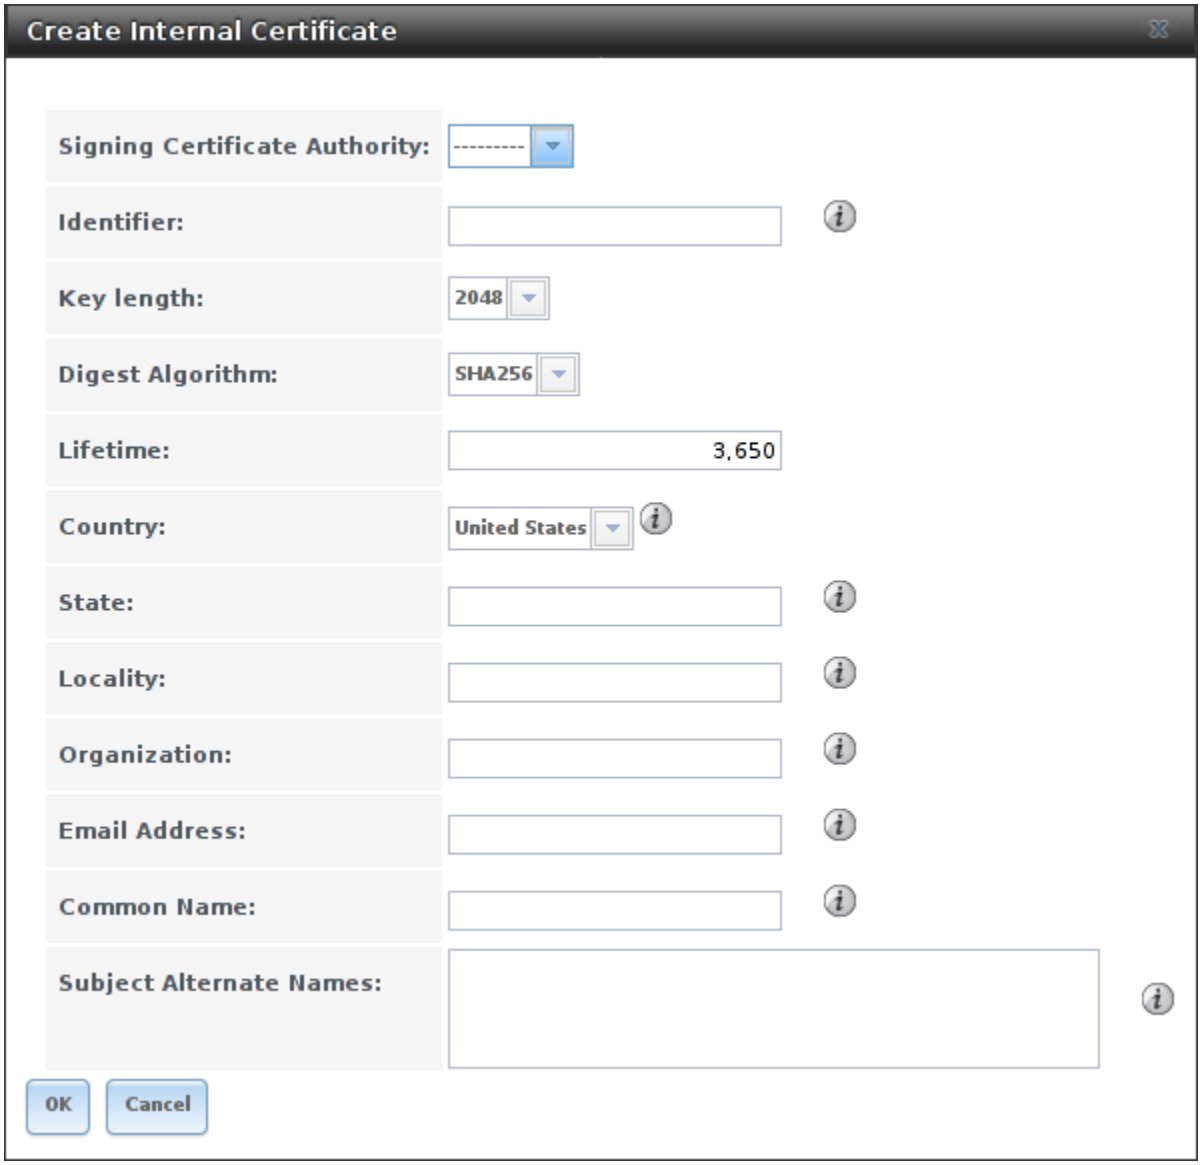

Fig. 4.21: Creating a New Certificate

Table 4.10: Certificate Creation Options

Setting	Value	Description
Signing Certificate Authority	drop-down menu	Select the CA which was previously imported or created using *CAs* (page 49).
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric, underscore (_), and dash (-) characters.
Key Length	drop-down menu	For security reasons, a minimum of *2048* is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algorithm.
Lifetime	integer	The lifetime of the certificate is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	State or province for the organization.
Locality	string	Location of the organization.
Organization	string	Name of the company or organization.
Email Address	string	Email address for the person responsible for the CA.

Continued on next page

Table 4.10 – continued from previous page

Setting	Value	Description
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The *Common Name* **must** be unique within a certificate chain.
Subject Alternate Names	string	Multi-domain support. Enter additional domain names and separate them with a space.

If the certificate is signed by an external CA, such as Verisign, instead create a certificate signing request. To do so, click *Create Certificate Signing Request*. A screen like the one in Figure 4.21 opens, but without the *Signing Certificate Authority* field.

Certificates that are imported, self-signed, or for which a certificate signing request is created are added as entries to *System → Certificates*. In the example shown in Figure 4.22, 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 *Import Certificate* so it is available as a configurable option for encrypting connections.

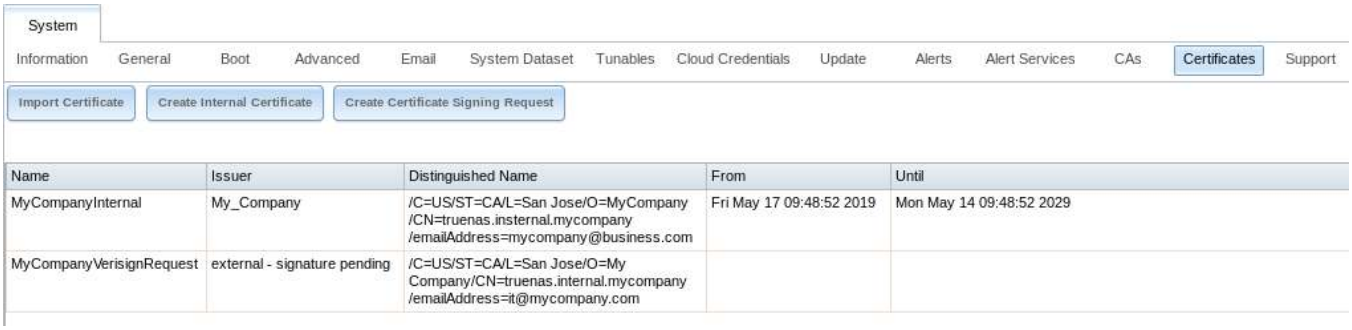

Fig. 4.22: Managing Certificates

Clicking an entry activates these configuration buttons:

- **View:** use this option to view the contents of an existing certificate or to edit the *Identifier*.
- **Export Certificate** saves a copy of the certificate or certificate signing request to the system being used to access the TrueNAS® system. For a certificate signing request, send the exported certificate to the external signing authority so that it can be signed.
- **Export Private Key** saves a copy of the private key associated with the certificate or certificate signing request to the system being used to access the TrueNAS® system.
- **Edit** shows the details for an existing certificate signing request and includes an area to paste a *Certificate*.
- **Delete** is used to delete a certificate or certificate signing request.

4.14 Support

The TrueNAS® *Support* tab, shown in Figure 4.23, is used to view or update the system license information. It also provides a built-in ticketing system for generating support requests.

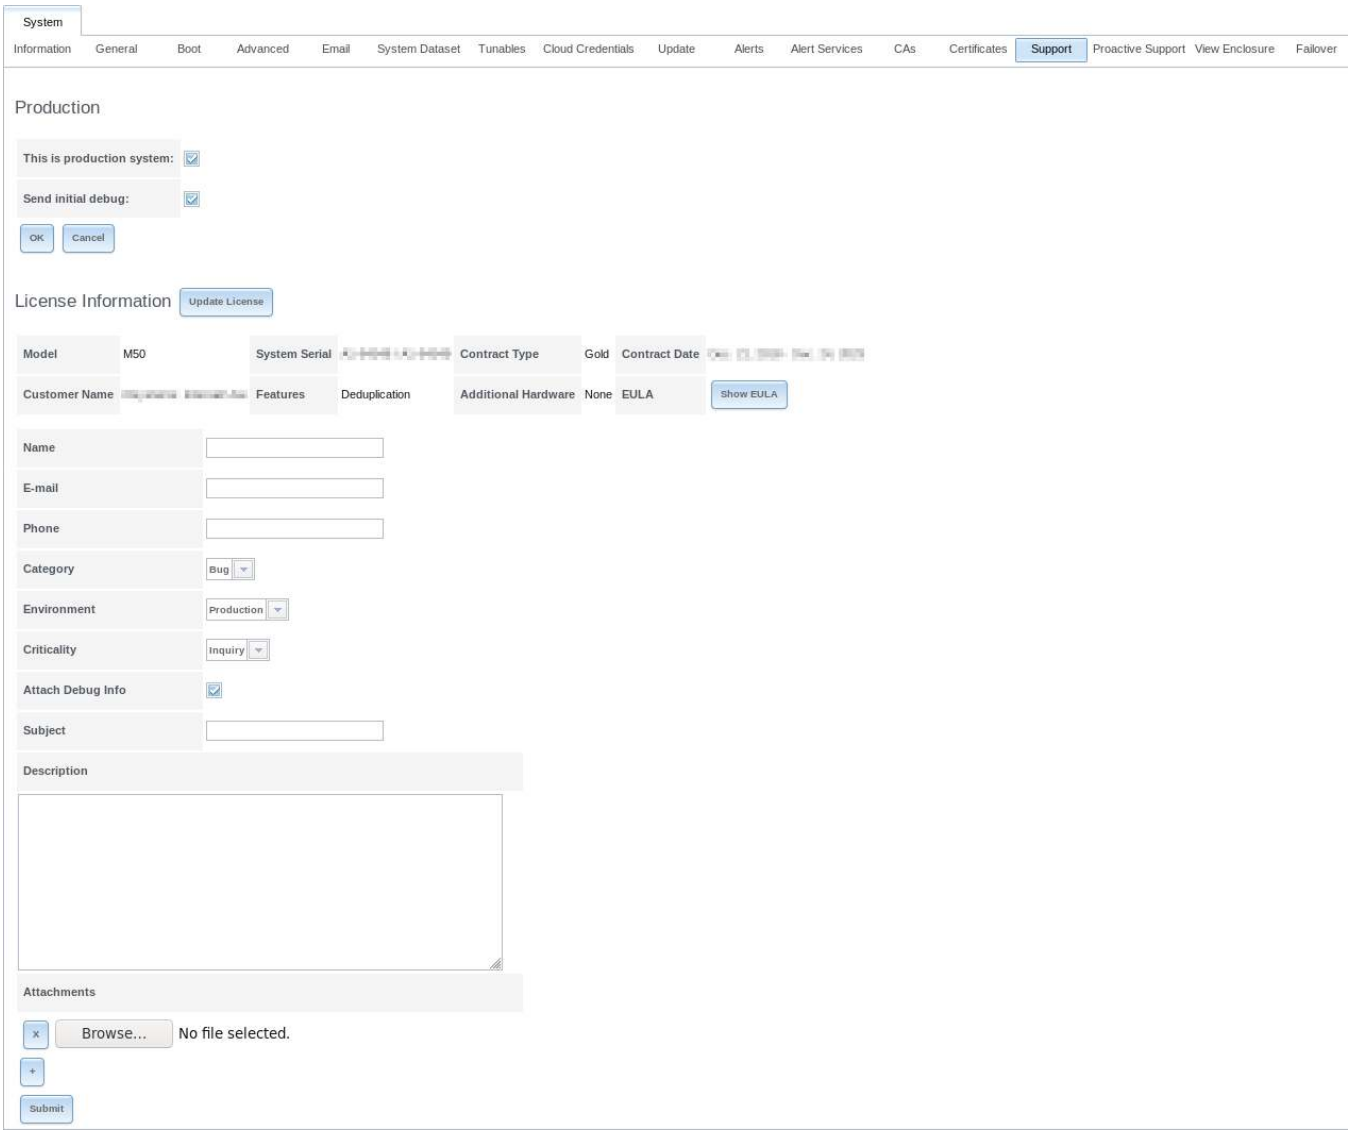

Fig. 4.23: Support Tab

This example shows a system that is used in production with an initial debug sent to iXsystems Support.

The system has a valid license which indicates the hardware model, system serial number, support contract type, licensed period, customer name, licensed features, additional supported hardware, and a *Show EULA* button.

If the license expires or additional hardware, features, or contract type are required, contact an iXsystems support engineer. After a new license string has been provided, 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.

Table 4.11: Environment Options :class: longtable

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

- **Criticality** is a drop-down menu to indicate the criticality level. Choices are *Inquiry*, *Loss of Functionality*, or *Total Down*.
- **Attach Debug Info** leaving this option selected is recommended so an overview of the system hardware and configuration to be automatically generated and included with the ticket.
- **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. Click the + button to add more attachments.

Click *Submit* to generate and send the support ticket to iXsystems. This process can take several minutes while information is collected and sent.

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

4.15 Proactive Support

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

Note: The fields on this tab are only enabled for Silver and Gold support coverage level customers. Please *contact iXsystems* (page 9) for information on upgrading from other support levels.

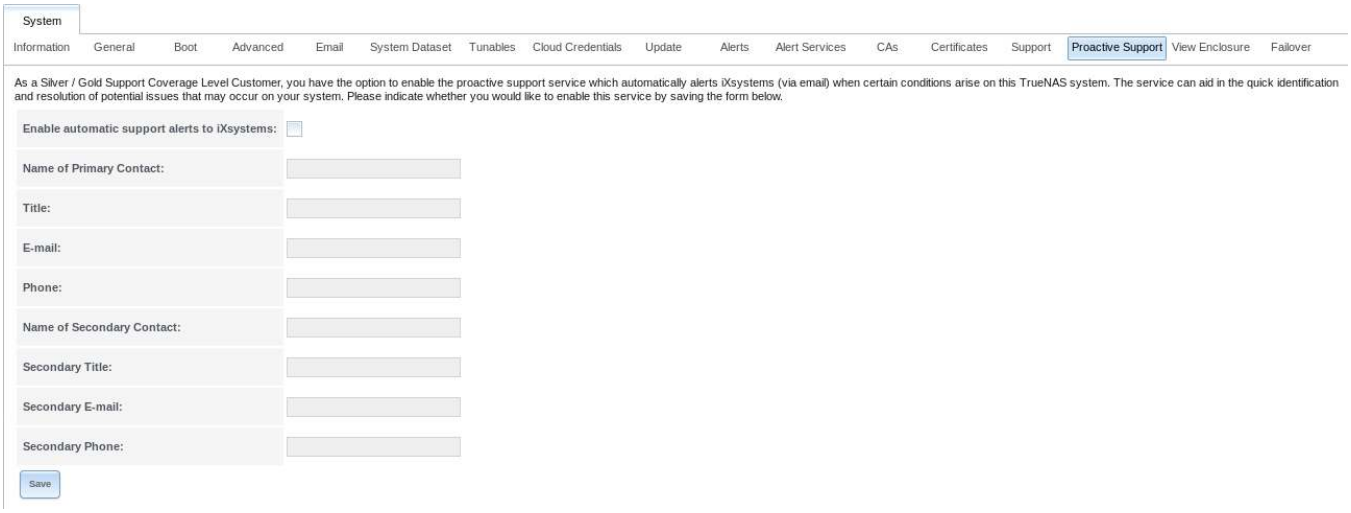

Fig. 4.24: Proactive Support Tab

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* option is enabled, then click *Save*.

TrueNAS® sends an email alert if ticket creation fails while Proactive Support is active.

4.16 View Enclosure

Click *Storage* → *Volumes* → *View Enclosure* to display a status summary of the connected disks and hardware. An example is shown in Figure 4.25.

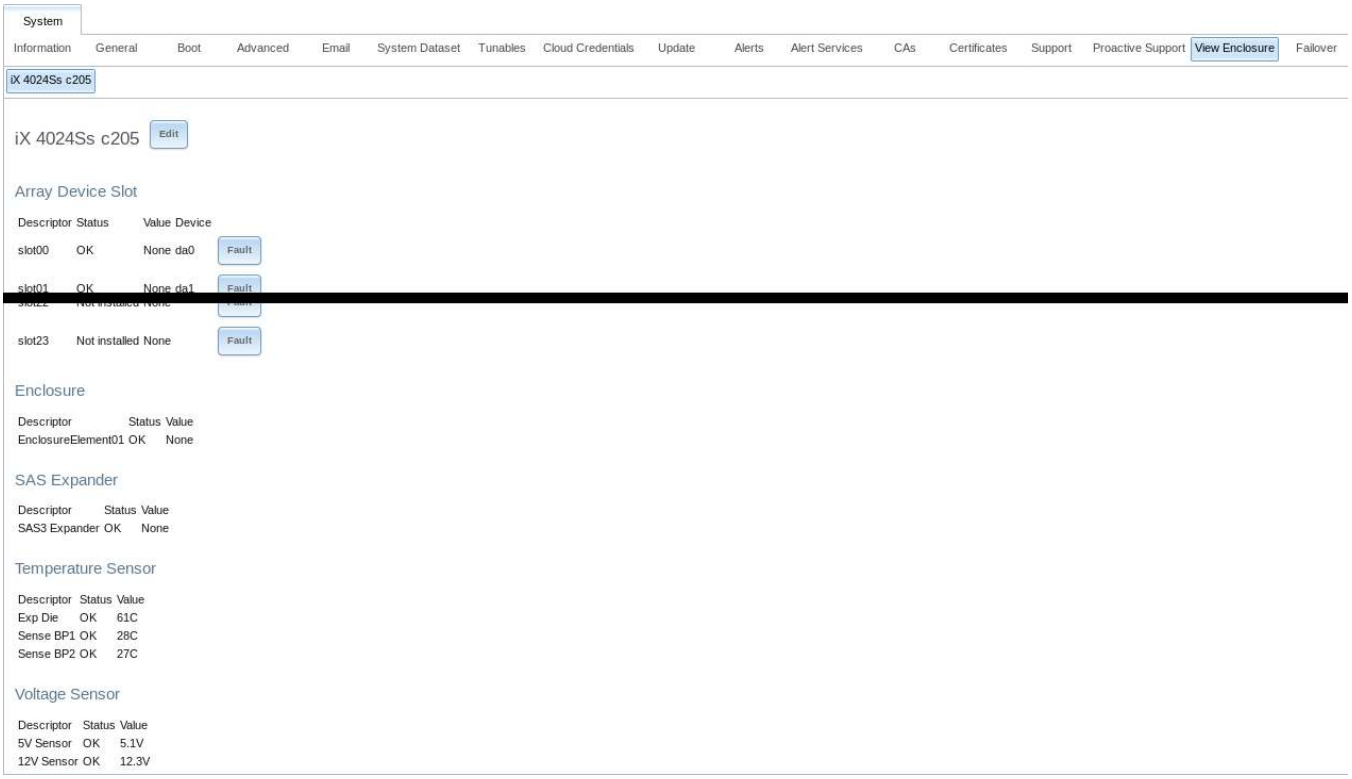

Fig. 4.25: View Enclosure

The screen is divided into these sections:

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

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

Enclosure: shows the status of the enclosure.

Power Supply: shows the status of each power supply.

SAS Expander: shows the status of the expander.

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

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

4.17 Failover

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

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

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

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

To configure HA, turn on both units in the array. Use the instructions in the *Console Setup Menu* (page 11) 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 86) 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 three 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.
- **Hostname (Virtual):** Enter the fully qualified hostname plus the domain name. When using a virtualhost, this is also used as the Kerberos principal name.

Next, go to *Network* → *Interfaces* → *Add Interface*. The HA license adds several fields to the usual *Interfaces* (page 82) screen:

- **IPv4 Address (Node A/B):** if the other node will use a static IP address, rather than DHCP, set it here.
- **IPv4 Address (This Node):** if this node will use a static IP address, rather than DHCP, set it here.
- **Virtual IP:** enter the IP address to use for administrative access to the array.
- **Virtual Host ID:** use a unique Virtual Host ID (VHID) on the broadcast segment of the network. Configuring multiple Virtual IP addresses requires a separate VHID for each address.
- **Critical for Failover:** enable this option 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 4.12. This option is interface-specific, allowing different settings for a management network and a data network. Note that enabling this option 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* option is enabled. This option allows grouping multiple, critical-for-failover interfaces. Groups apply to single systems. A failover occurs when every interface in the group fails. Groups with a single interface trigger a failover when that interface fails. Configuring the system to failover when any interface fails requires marking each interface as critical and placing them in separate groups.

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

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

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

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

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

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

The options available in *System* → *Failover* are shown in Figure 4.26: and described in Table 4.12.

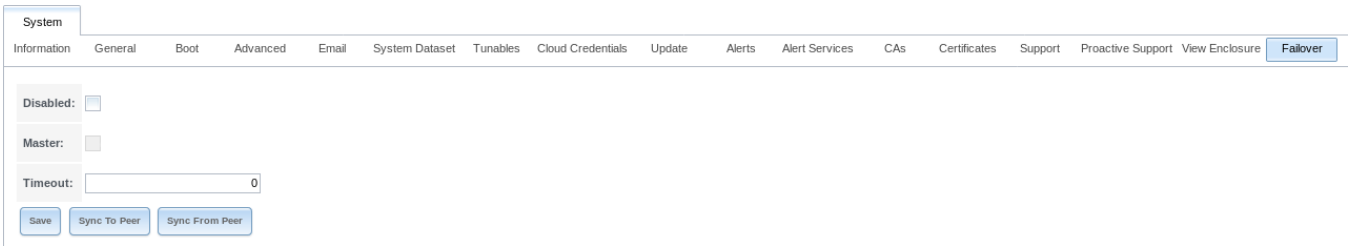

Fig. 4.26: Example Failover Screen

Table 4.12: Failover Options

Setting	Value	Description
Disabled	checkbox	Set to disable failover. The *HA Enabled* icon changes to *HA Disabled* and activates the *Master* field. An error message is generated if the standby node is not responding or failover is not configured.
Master	checkbox	Grayed out unless *Disabled* is selected. In that case, this option is automatically enabled on the master system, allowing the master to automatically take over when the *Disabled* option is deselected.
Timeout	integer	Specify, in seconds, how quickly failover occurs after a network failure. The default of *0* indicates that failover either occurs immediately or, if the system is using a link aggregation, after 2 seconds.
Sync to Peer	button	Open a dialog window to force the TrueNAS® configuration to sync from the active node to the standby node. After the sync, the standby node must be rebooted (enabled by default) to load the new configuration. *Do not use this unless requested by an iXsystems support engineer, the HA daemon normally handles configuration sync automatically*.
Sync From Peer	button	Open a dialog window to force the TrueNAS® configuration to sync from the standby node to the active node. *Do not use this unless requested by an iXsystems support engineer, the HA daemon normally handles configuration sync automatically*.

Notes about High Availability and failovers:

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

Failover is not allowed if both storage controllers have the same CARP state. A critical *Alert* (page 247) is generated and the HA icon shows *HA Unavailable*.

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

If both nodes reboot simultaneously, the GELI passphrase for an *encrypted* (page 97) volume must be entered at the web interface login screen.

If there are a different number of disks connected to each node, an *Alert* (page 247) is generated and the HA icon switches to *HA Unavailable*.

CHAPTER

FIVE

TASKS

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

- *Cloud Sync* (page 61) schedules data synchronization to cloud providers
- *Cron Jobs* (page 67) schedules a command or script to automatically execute at a specified time
- *Init/Shutdown Scripts* (page 69) configures a command or script to automatically execute during system startup or shutdown
- *Rsync Tasks* (page 70) schedules data synchronization to another system
- *S.M.A.R.T. Tests* (page 77) schedules disk tests

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

Note: By default, *Scrubs* (page 134) are run once a month by an automatically-created task. *S.M.A.R.T. Tests* (page 77) and *Periodic Snapshot Tasks* (page 120) 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.

Cloud Credentials (page 39) must be pre-defined before a cloud sync is created. One set of credentials can be used for more than one cloud sync. For example, a single set of credentials for Amazon S3 can be used for separate cloud syncs that push different sets of files or directories.

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

After the credentials and receiving bucket have been configured, *Tasks* → *Cloud Syncs* → *Add Cloud Sync* is used to define the schedule for running a cloud sync task. An example is shown in Figure 5.1.

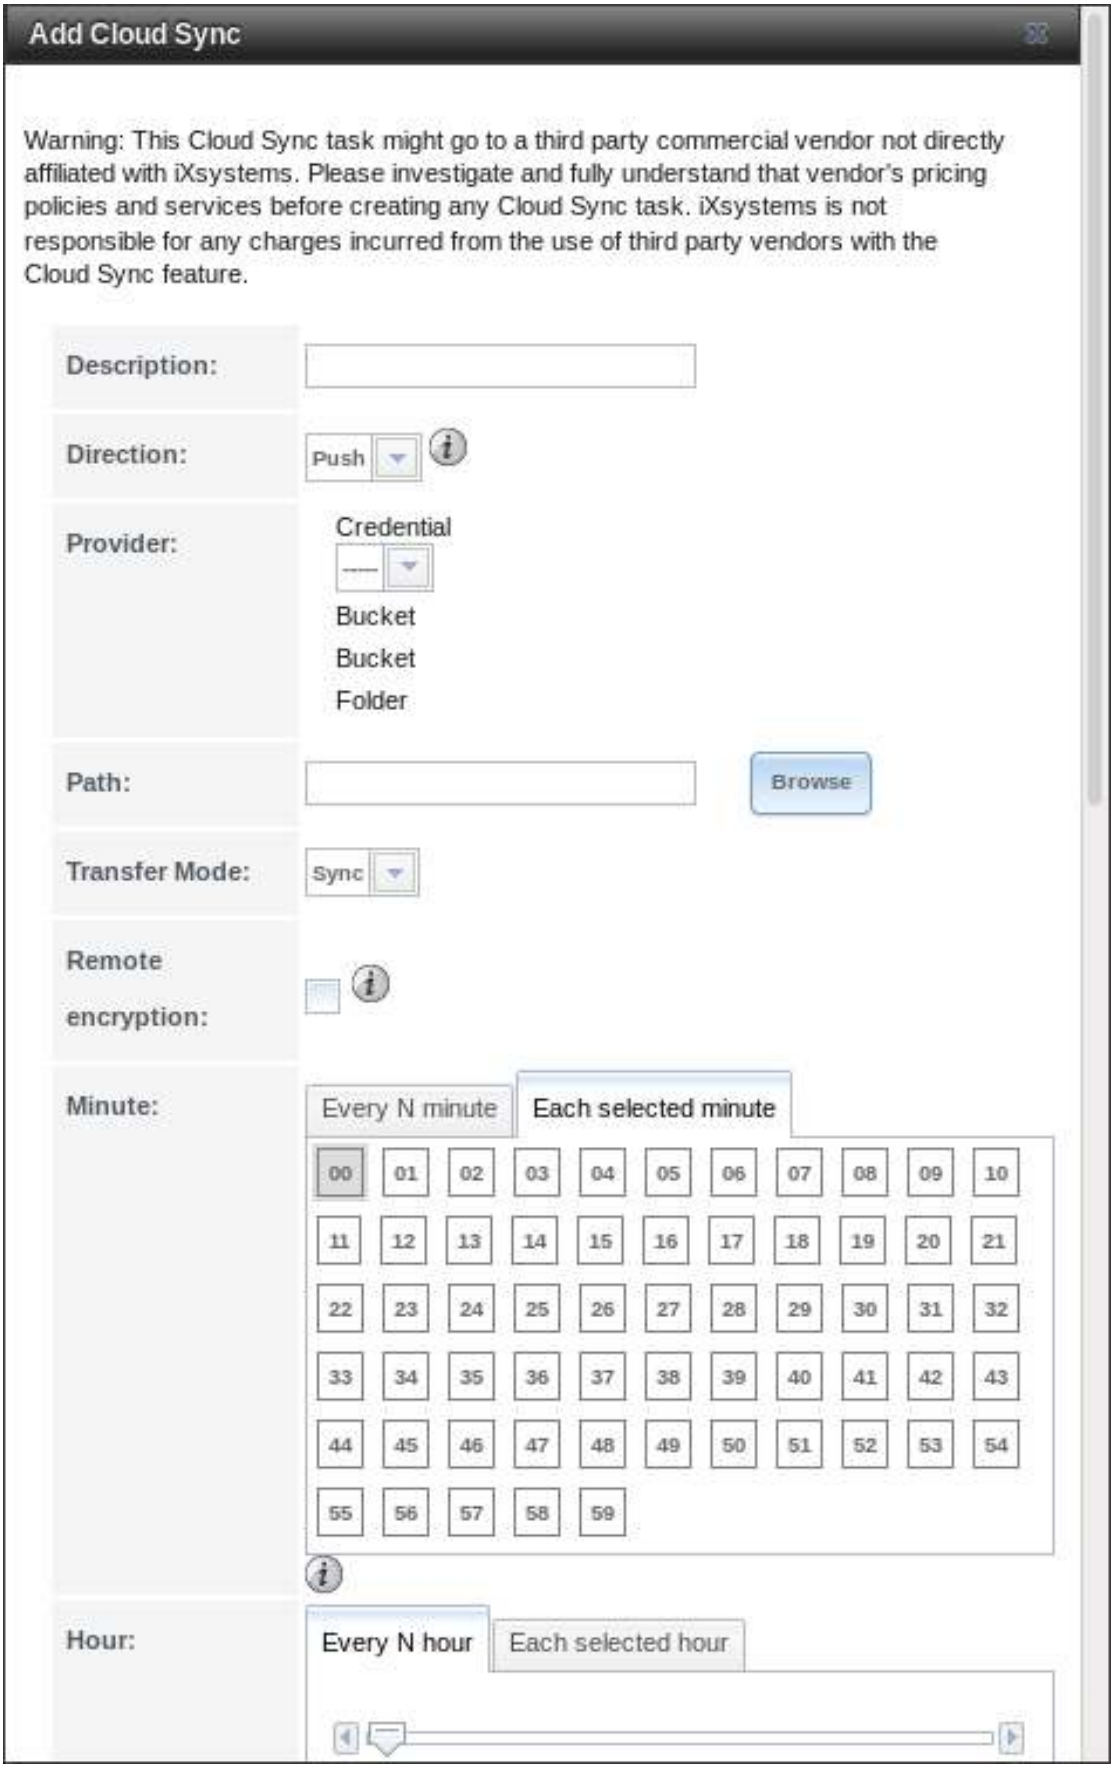

Fig. 5.1: Adding a Cloud Sync

Table 5.1 shows the configuration options for Cloud Syncs.

Table 5.1: Cloud Sync Options

Setting	Value Type	Description
Description	string	Enter a descriptive name for this Cloud Sync.
Direction	string	*Push* sends data to cloud storage. *Pull* receives data from cloud storage.
Provider	drop-down menu	Choose the cloud storage provider credentials from the list of entered *Cloud Credentials* (page 39). The UI tests the credential and displays an error if a connection cannot be made.
Amazon S3 Buckets	drop-down menu	Only appears when a valid S3 credential is the *Provider*. Select the pre-defined S3 bucket to use.
Folder	string	Only appears when an S3 credential is the *Provider*. Optionally enter the name of the folder within the selected bucket.
Server Side Encryption	drop-down menu	Only appears when an S3 credential is the *Provider*. Choices are *None* (no encryption) or *AES-256* (encrypted).
Path	browse button	Select the directories or files to be sent to the cloud for *Push* syncs, or the destination to be written as the destinations for *Pull* syncs. Be cautious about the destination of *Pull* jobs to avoid overwriting existing files.
Transfer Mode	drop-down menu	*Sync* (default) makes files on destination system identical to those on the source. Files removed from the source are also removed from the destination, similar to `rsync --delete`. *Copy* copies files from the source to the destination and skips files that are identical, similar to `rsync`. *Move* copies files from the source to the destination and deletes the source files after the copy, similar to `mv`.
Remote encryption	checkbox	Use rclone crypt (https://rclone.org/crypt/) to manage data encryption during *PUSH* or *PULL* transfers: *PUSH:* Encrypt files before transfer and store the encrypted files on the remote system. Files are encrypted using the *Encryption password* and *Encryption salt* values. *PULL:* Decrypt files that are being stored on the remote system before the transfer. Transferring the encrypted files requires entering the same *Encryption Password* and *Encryption salt* that was used to encrypt the files. Adds the *Filename encryption*, *Encryption password*, and *Encryption salt* options. Additional details about the encryption algorithm and key derivation are available in the rclone crypt File formats documentation (https://rclone.org/crypt/#file-formats).
Filename encryption	checkbox	Encrypt (*PUSH*) or decrypt (*PULL*) file names with the rclone "Standard" file name encryption mode (https://rclone.org/crypt/#file-name-encryption-modes). The original directory structure is preserved. A filename with the same name always has the same encrypted filename. *PULL* tasks that have *Filename encryption* enabled and an incorrect *Encryption password* or *Encryption salt* will not transfer any files but still report that the task was successful. To verify that files were transferred successfully, click the finished *task status* (page 65) to see a list of transferred files.
Encryption password	string	Password to encrypt and decrypt remote data. **Warning**: Always securely back up this password! Losing the encryption password will result in data loss.

Continued on next page

Table 5.1 – continued from previous page

Setting	Value Type	Description
Encryption salt	string	Enter a long string of random characters for use as salt (https://searchsecurity.techtarget.com/definition/salt) for the encryption password. **Warning**: Always securely back up the encryption salt value! Losing the salt value will result in data loss.
Minute	slider or minute selections	Select *Every N minutes* and use the slider to choose a value, or select *Each selected minute* and choose specific minutes to run the task.
Hour	slider or hour selections	Select *Every N hours* and use the slider to choose a value, or select *Each selected hour* and choose specific hours to run the task.
Day of month	slider or day of month selections	Select *Every N days of month* and use the slider to choose a value, or select *Each selected day of month* and choose specific days to run the task.
Month	checkboxes	Months when the task runs.
Day of week	checkboxes	Days of the week to run the task.
Enabled	checkbox	Unset to temporarily disable this Cloud Sync.

The time selected is when the Cloud Sync task is allowed to begin. The cloud sync runs until finished, even after the time selected.

Note: Files that have completed the sync process are not deleted from the destination if the rclone sync (https://rclone.org/commands/rclone_sync/) is interrupted or encounters an error. This includes a common error when the Dropbox copyright detector (https://techcrunch.com/2014/03/30/how-dropbox-knows-when-youre-sharing-copyrighted-stuff-without-actually-looking-at-your-stuff/) identifies a copyrighted file.

Figure 5.2 shows a cloud sync called *backup-acctg* that "pushes" a file to cloud storage. The last run finished with a status of *SUCCESS*.

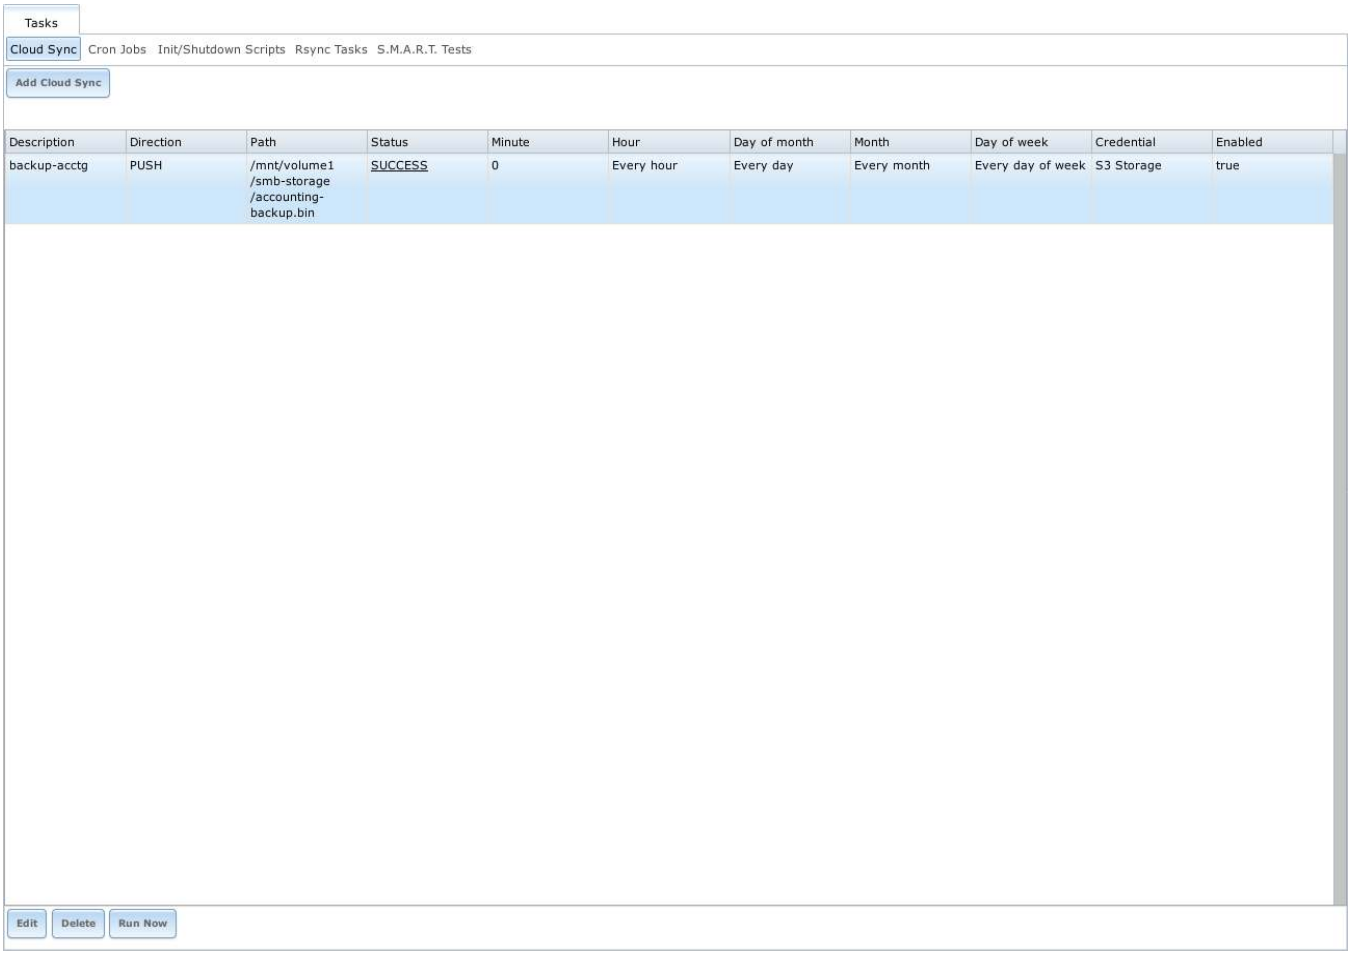

Fig. 5.2: Cloud Sync Status

To modify an existing cloud sync, click the entry to access the *Edit*, and *Delete*, and *Run Now* buttons.

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

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:

Fig. 5.3: Example: Adding Cloud Credentials

The local data to be sent to the cloud is in a dataset called `acctg-backups`. The cloud sync task is created by going to *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* is selected.

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* option is set by default, so this cloud sync will run at the next scheduled time.

The completed dialog is shown in Figure 5.4:

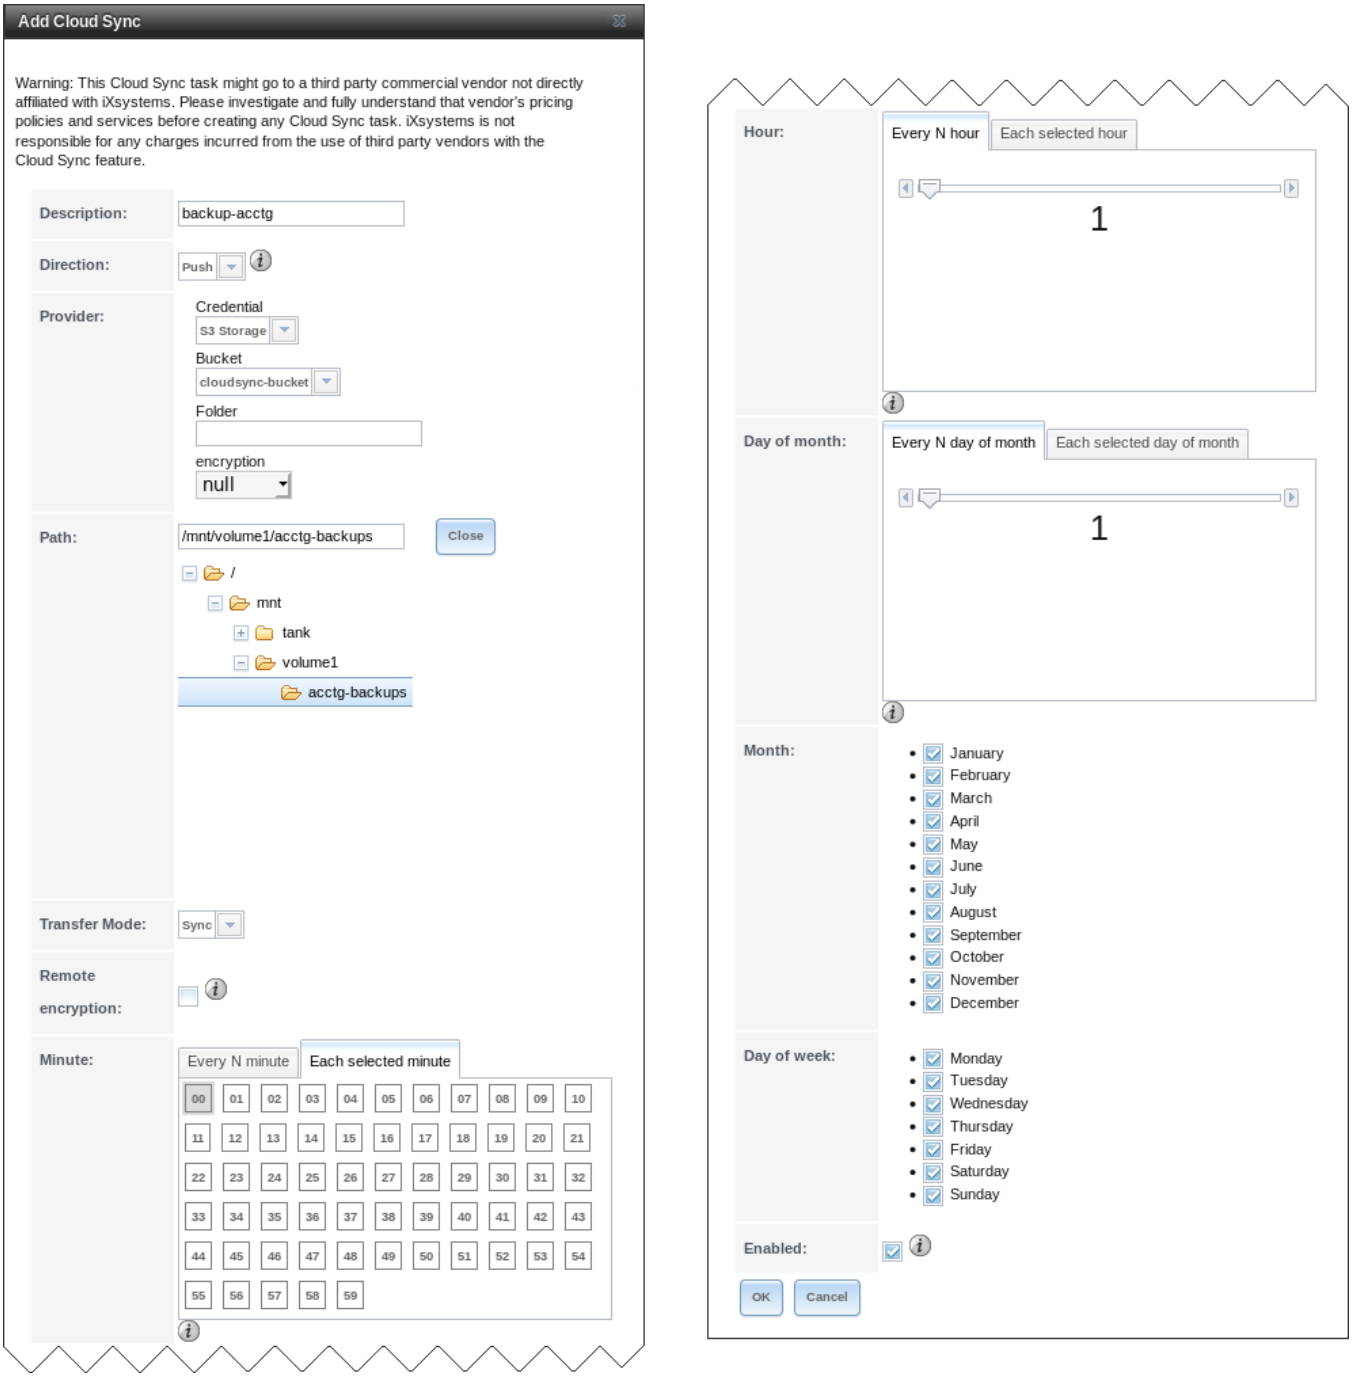

Fig. 5.4: Example: Adding a Cloud Sync

5.2 Cron Jobs

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

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

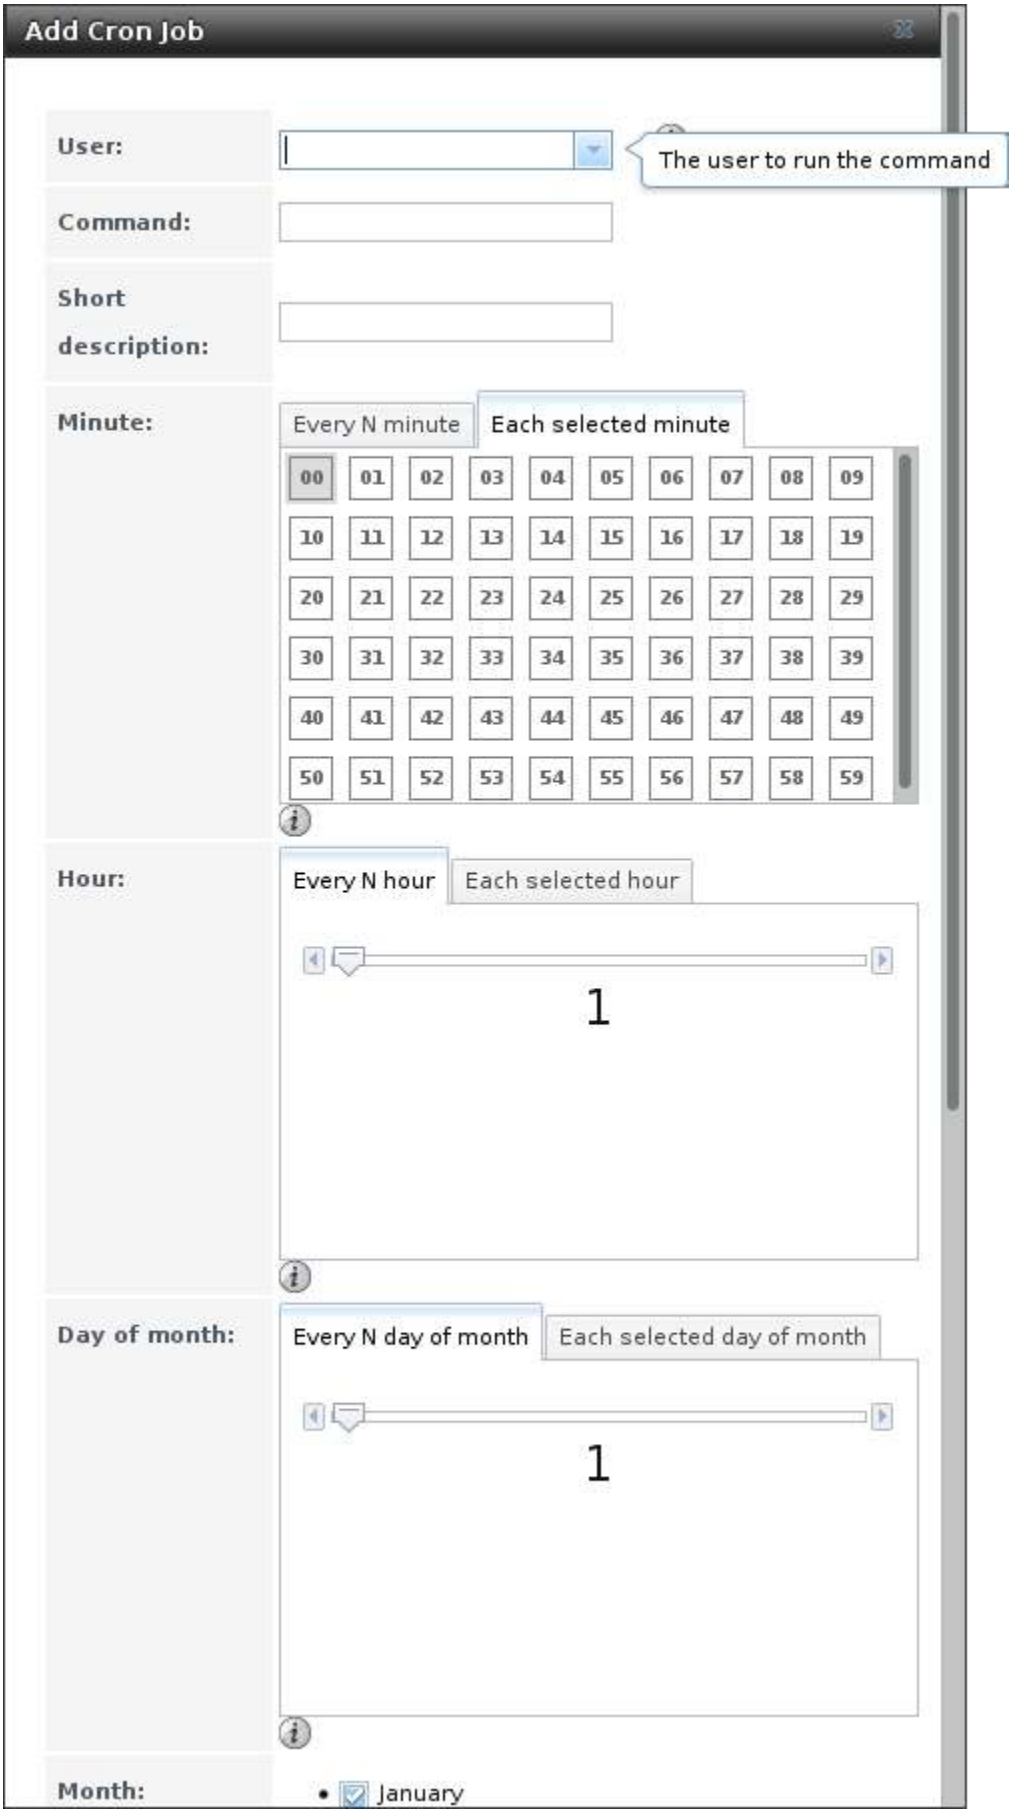

Fig. 5.5: Creating a Cron Job

Table 5.2 lists the configurable options for a cron job.

Table 5.2: Cron Job Options

Setting	Value	Description
User	drop-down menu	Choose a user account to run the command or script. The user must have permissions to run the command.
Command	string	Enter the **full path** to the command or script to be run. Test a script at the command line first to make sure it works as expected.
Short description	string	Optional. Describe the new cron job.
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, the cron job occurs every N days. With day selections, the cron job occurs on the highlighted days each month.
Month	checkboxes	Cron job occurs on the selected months.
Day of week	checkboxes	Cron job occurs on the selected days.
Redirect Stdout	checkbox	Disables emailing standard output to the *root* user account.
Redirect Stderr	checkbox	Disables emailing errors to the *root* user account.
Enabled	checkbox	Deselect disable the cron job without deleting it.

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

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

5.3 Init/Shutdown Scripts

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

Go to *Tasks* → *Init/Shutdown Scripts* and click *Add Init/Shutdown Script*.

Fig. 5.6: Add an Init/Shutdown Command or Script

Table 5.3: Init/Shutdown Command or Script Options

Setting	Value	Description
Type	drop-down menu	Select *Command* for an executable or *Script* for an executable script.
Command or Script	string	If *Command* is selected, enter the command with any options. When *Script* is selected, click *Browse* to select the script from an existing pool.
When	drop-down menu	Select when the *Command* or *Script* runs: • *Pre Init*: early in the boot process, after mounting filesystems and starting networking • *Post Init*: at the end of the boot process, before TrueNAS® services start • *Shutdown*: during the system power off process.
Enabled	checkbox	Enable this task. Unset to disable the task without deleting it.

Scheduled commands must be in the default path. The full path to the command can also be included in the entry. The path can be tested with `which {commandname}` in the *Shell* (page 244). When available, the path to the command is shown:

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

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

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

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

5.4 Rsync Tasks

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

Rsync is most effective when only a relatively small amount of the data has changed. There are also some limitations when using Rsync with Windows files (https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/). For large amounts of data, data that has many changes from the previous copy, or Windows files, *Replication Tasks* (page 122) are often the faster and better solution.

Rsync is single-threaded and gains little from multiple processor cores. To see whether rsync is currently running, use `pgrep rsync` from the *Shell* (page 244).

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

This section summarizes the options when creating an rsync task. It then provides a configuration example between two 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.

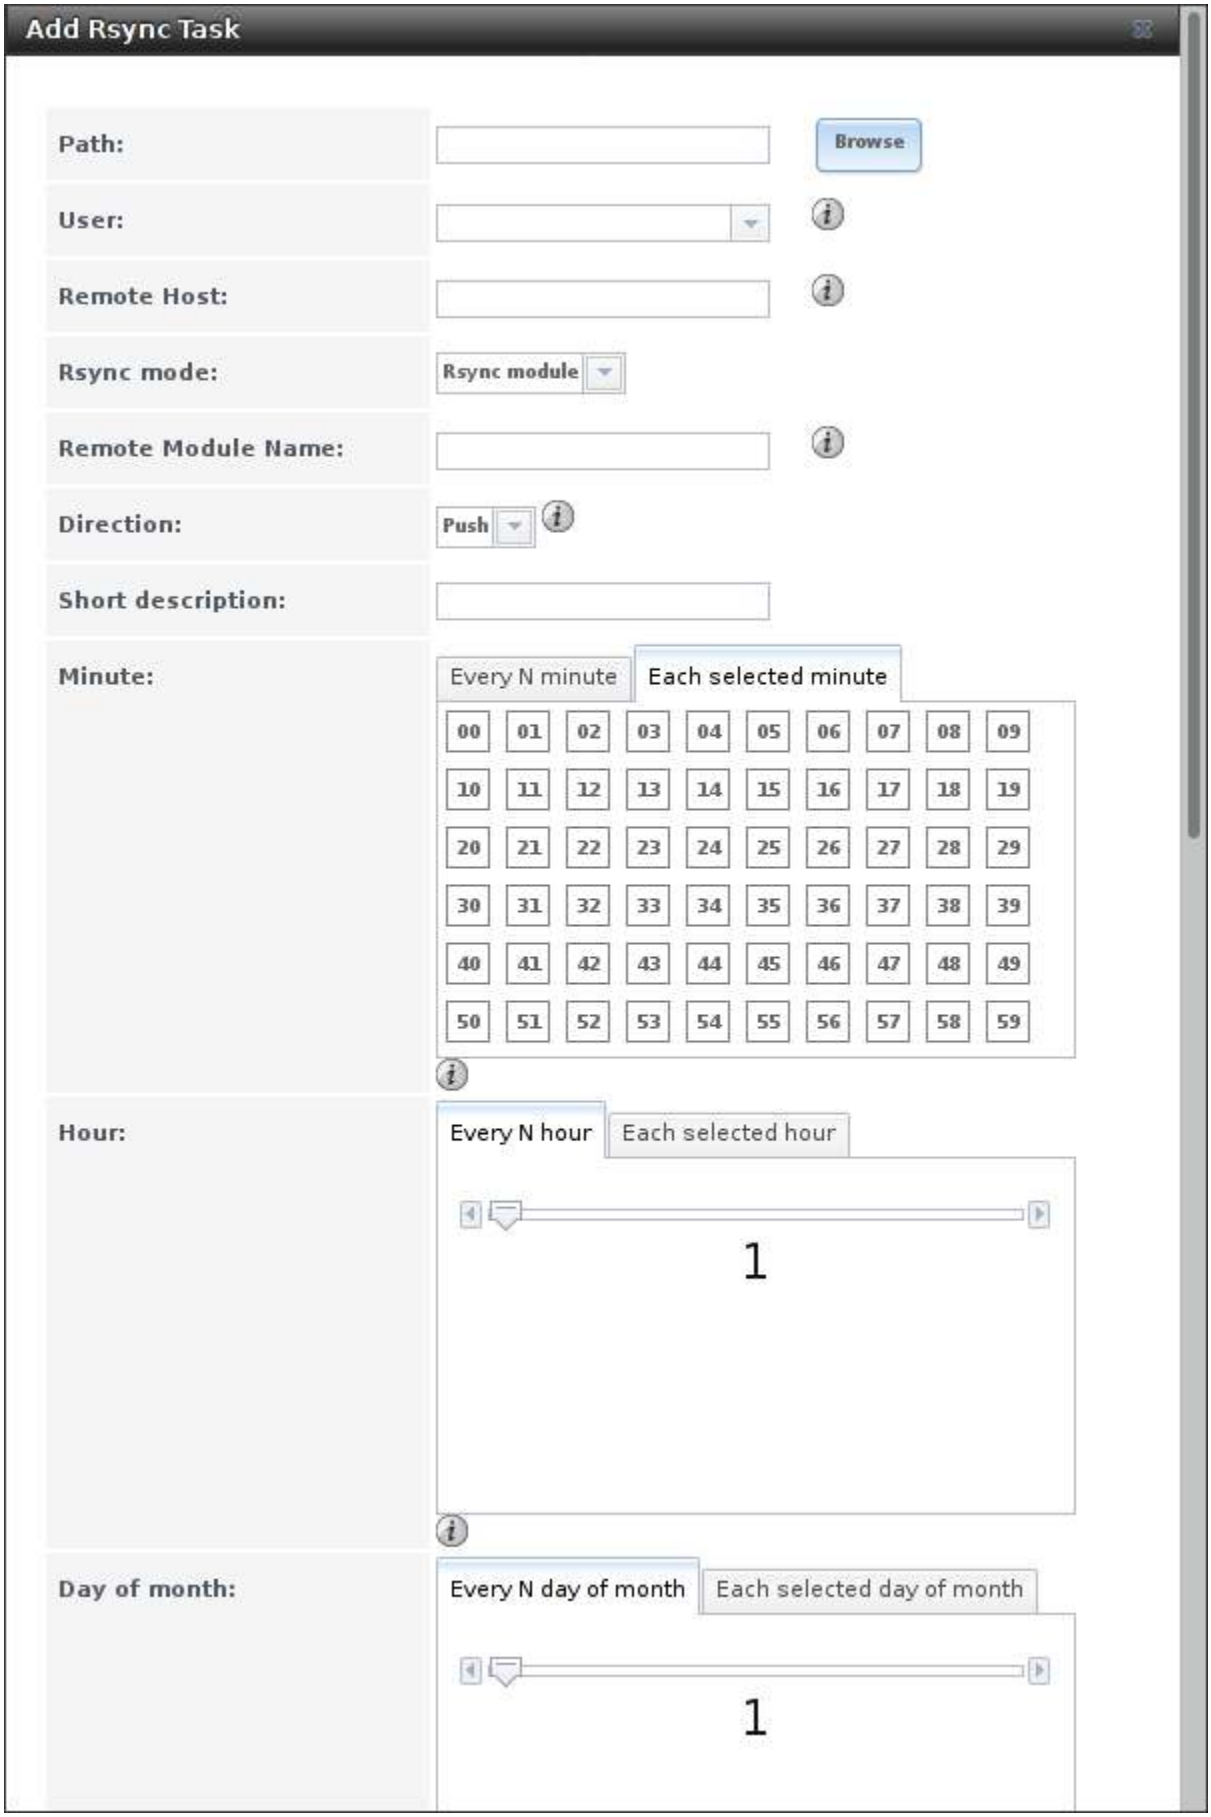

Fig. 5.7: Adding an Rsync Task

Table 5.4: Rsync Configuration Options

Setting	Value	Description
Path	browse button	*Browse* to the path to be copied. Path lengths cannot be greater than 255 characters.
User	drop-down menu	The chosen user must have write permissions for the specified remote directory. The user name cannot contain spaces or exceed 17 characters.
Remote Host	string	Enter the IP address or hostname of the remote system that will store the copy. Use the format *username@remote_host* if the username differs on the remote host.
Remote SSH Port	integer	Only available in *Rsync over SSH* mode. Allows specifying an SSH port other than the default of *22*.
Rsync mode	drop-down menu	Choices are *Rsync module* or *Rsync over SSH*.
Remote Module Name	string	At least one module must be defined in rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html) of the rsync server or in the *Rsync Modules* of another system.
Remote Path	string	Only appears when using *Rsync over SSH* mode. Enter the **existing** path on the remote host to sync with. Example: */mnt/volume*. Note that maximum path length is 255 characters.
Validate Remote Path	checkbox	Verifies the existence of the *Remote Path*.
Direction	drop-down menu	Direct the flow of the data to the remote host. Choices are *Push* or *Pull*. Default is to *Push* to a remote host.
Short Description	string	Enter an optional description of the new rsync task.
Minute	slider or minute selections	When the slider is used the sync occurs every N minutes. Use *Each selected minute* for the sync to occur at the highlighted minutes.
Hour	slider or hour selections	When the slider is used the sync occurs every N hours. Use *Each selected hour* for the sync to occur at the highlighted hours.
Day of month	slider or day selections	When the slider is used the sync occurs every N days. Use *Each selected day of the month* for the sync to occur on the highlighted days.
Month	checkboxes	Define which months to run the task.
Day of week	checkboxes	Define which days of the week to run the task.
Recursive	checkbox	Set to include all subdirectories of the specified volume during the rsync task.
Times	checkbox	Set to preserve the modification times of the files.
Compress	checkbox	Set to reduce the size of data to transmit. Recommended for slower connections.
Archive	checkbox	Equivalent to `-rlptgoD`. This will run the task as recursive, copy symlinks as symlinks, preserve permissions, preserve modification times, preserve group, preserve owner (root only), and preserve device and special files.
Delete	checkbox	Set to delete files in the destination directory that do not exist in the sending directory.
Quiet	checkbox	Set to suppresses informational messages from the remote server.
Preserve permissions	checkbox	Set to preserve original file permissions. Useful if User is set to *root*.
Preserve extended attributes	checkbox	Both systems must support extended attributes. (https://en.wikipedia.org/wiki/Xattr).
Delay Updates	checkbox	Set to save the temporary file from each updated file to a holding directory. At the end of the transfer, all transferred files are renamed into place and temporary files deleted.

Continued on next page

Table 5.4 – continued from previous page

Setting	Value	Description
Extra options	string	Add any other rsync(1) (http://rsync.samba.org/ftp/rsync/rsync.html) options. The `*` character must be escaped with a backslash (`\*.txt`) or used inside single quotes (`'*.txt'`).
Enabled	checkbox	Unset to disable the rsync task without deleting it.

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

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

5.4.1 Rsync Module Mode

This configuration example configures rsync module mode between these two 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 *User* is set to *root* so it has permission to write anywhere
- 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 *Preserve permissions* option is enabled so that the original permissions are not overwritten by the *root* user

On *PULL*, an rsync module is defined in *Services* → *Rsync 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* (page 215).

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 244) on *PUSH* and run `ssh-keygen`. This example generates an RSA type public/private key pair for the *root* user. When creating the key pair, do not enter the passphrase as the key is meant to be used for an automated task.

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

TrueNAS® supports RSA keys for SSH. When creating the key, use `-t rsa` to specify this type of key. Refer to Key-based Authentication (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/openssh.html#security-ssh-keygen) for more information.

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

```
su user1
```

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

```
more .ssh/id_rsa.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQC1lBEXRgw1W8y8k+lXPlVR3xsmVSjtsoyIzV/PlQPo
SrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4h
dcD7Y5mvU3MAEeDClt02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmasti00qmDDcp/kO
xT+S6DFNDBy6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+l9RaEXMRuTyQgqJB/rsRcmJX5fApd
DmNfwrRSxLjDvUzfywnjFHlKk/+TQIT1gg1QQaj21PJD9pnDVF0AiJrWyWnR root@freenas.local
```

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

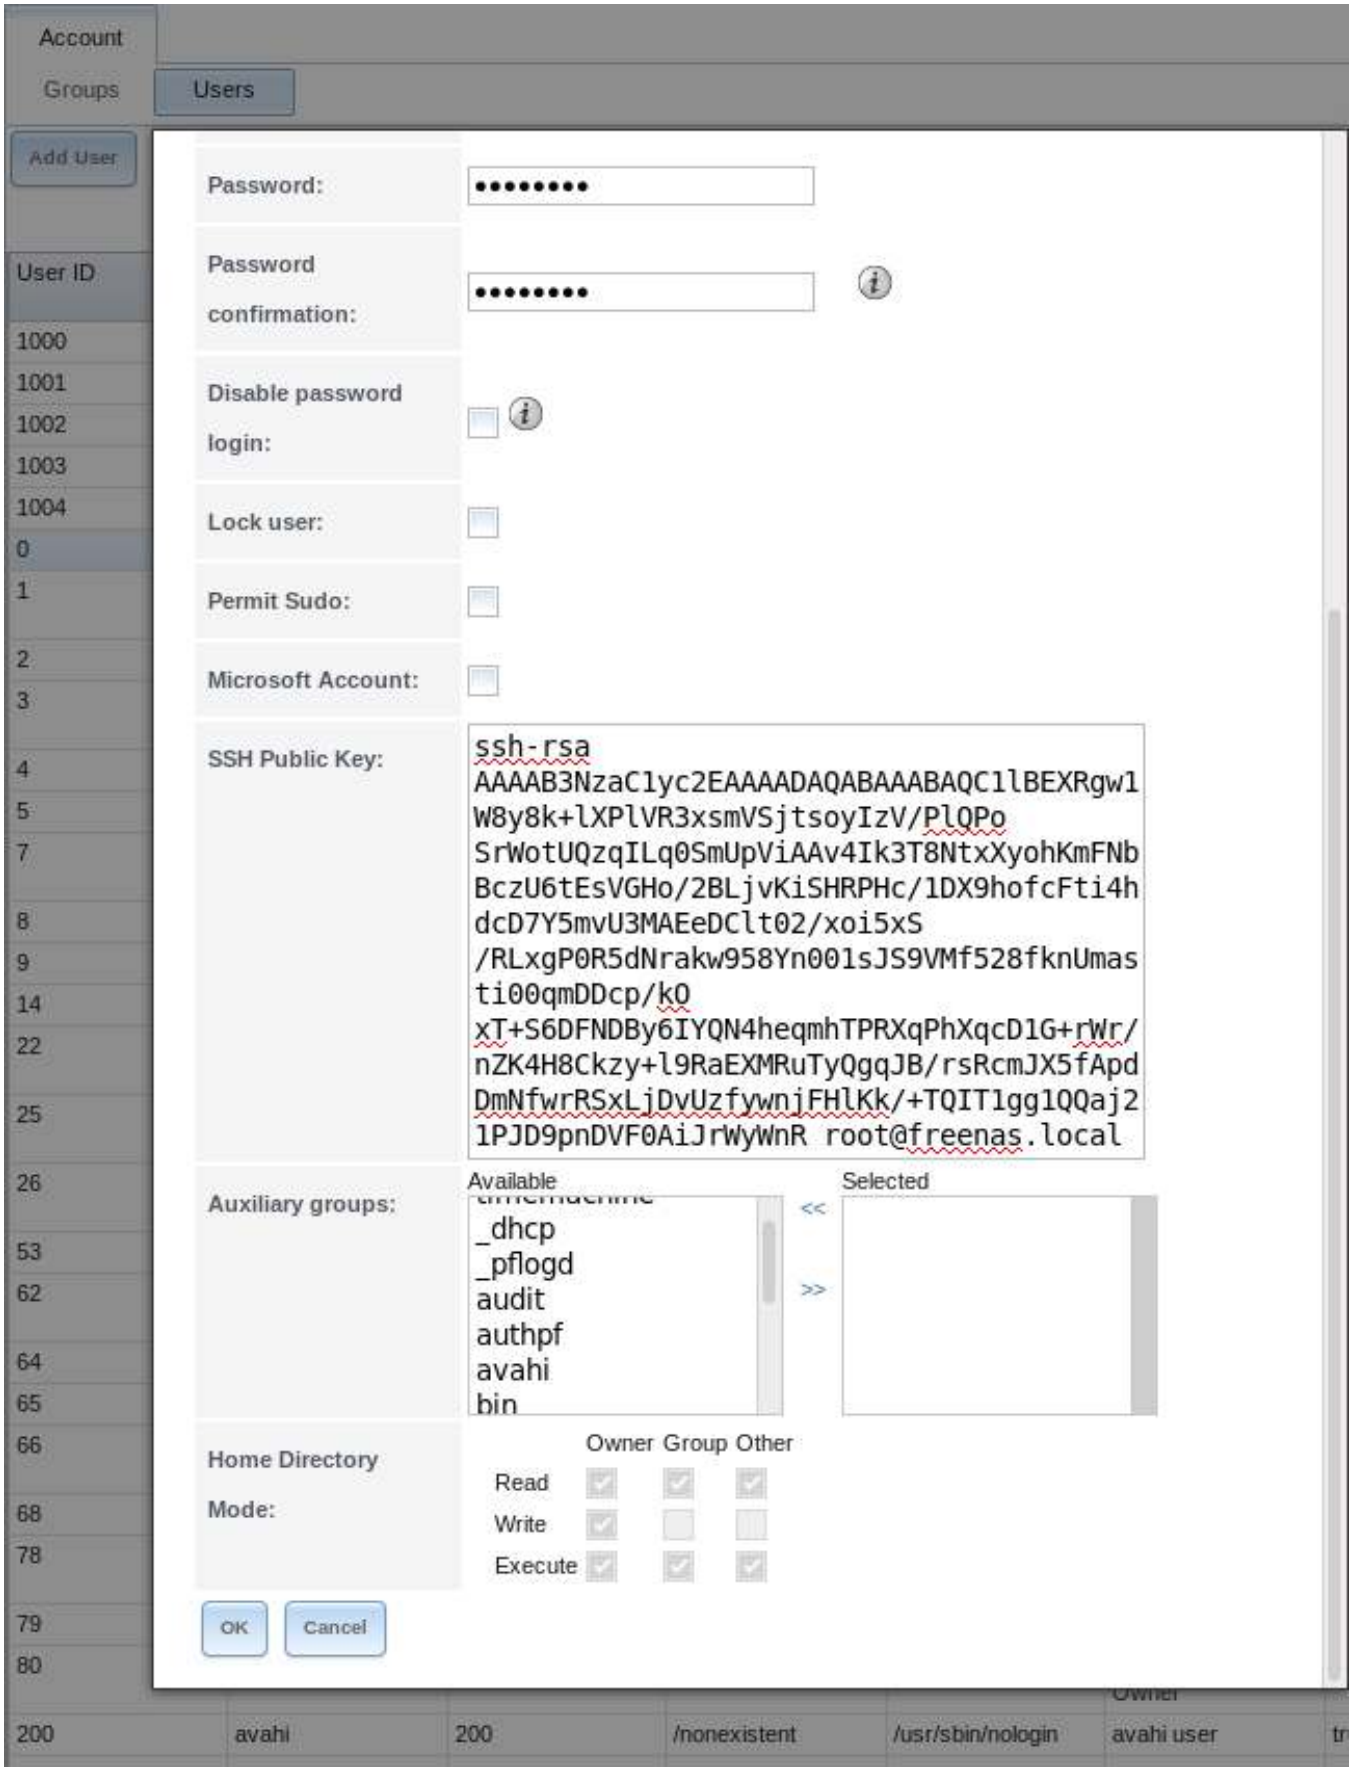

Fig. 5.8: Pasting the User SSH Public Key

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 command below 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 the previous example, use this configuration:

- the *Path* points to `/mnt/local/images`, the directory to be copied
- 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 *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 *Preserve Permissions* option is enabled so that the original permissions are not overwritten by the *root* user

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

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

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

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.

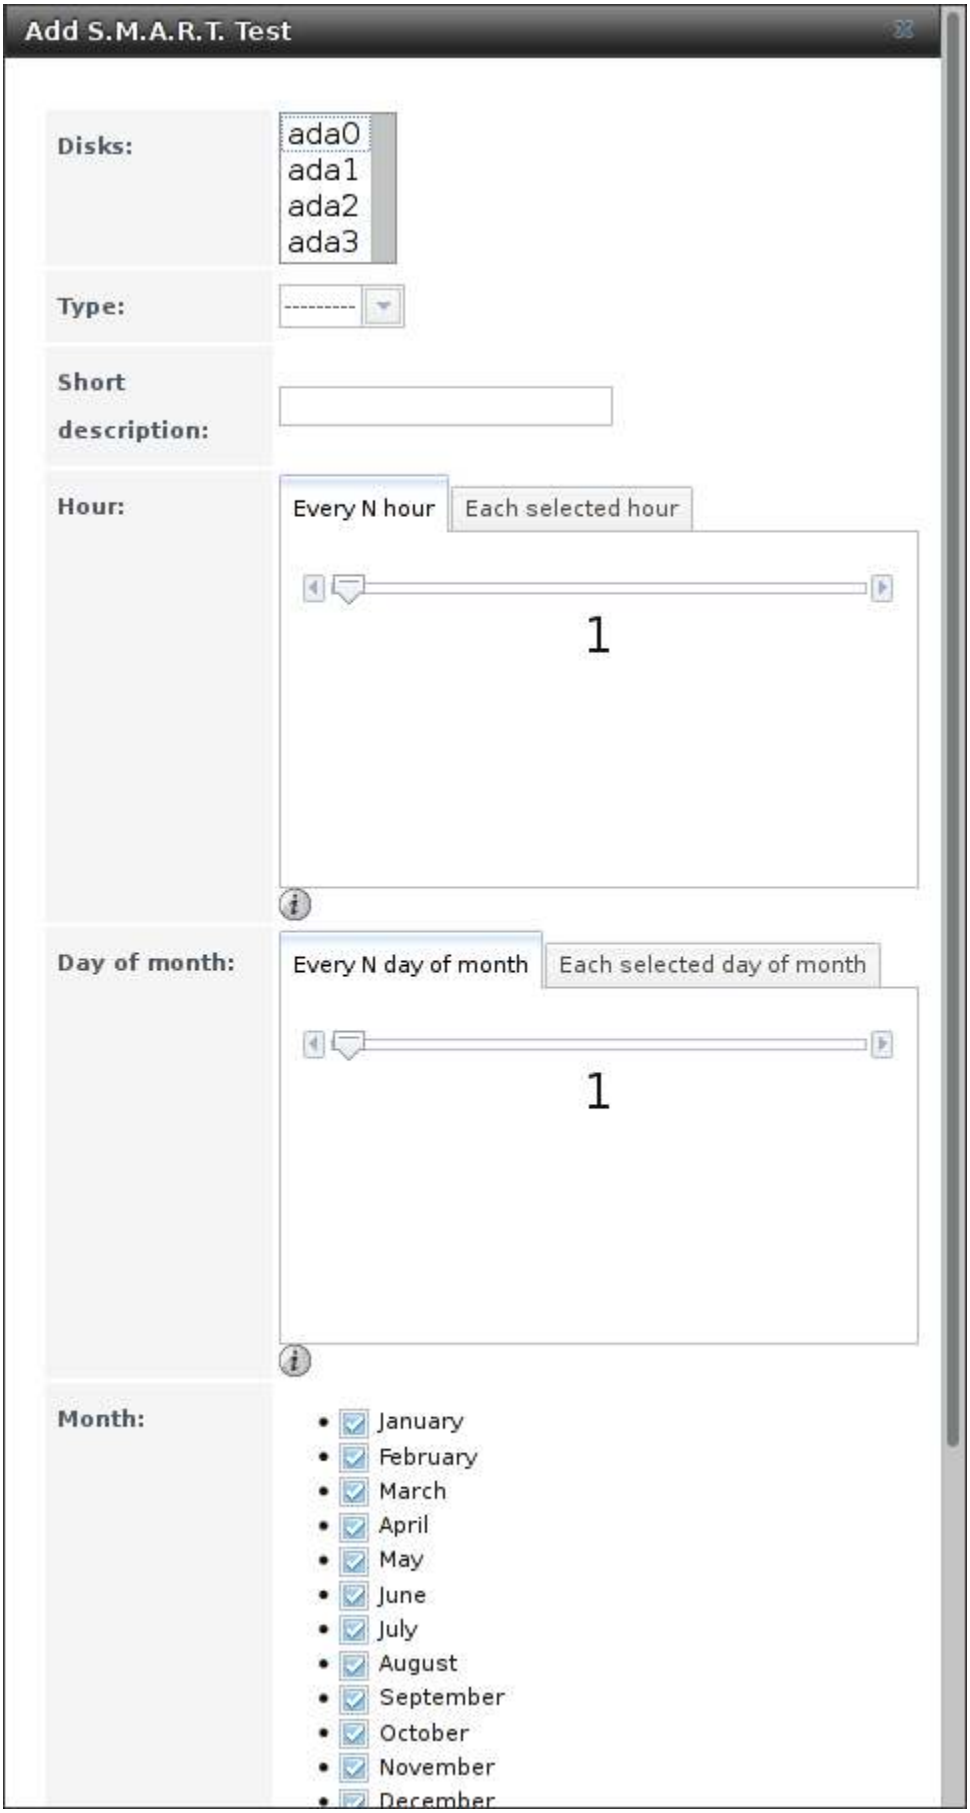

Fig. 5.9: Adding a S.M.A.R.T. Test

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	Select the disks to monitor.
Type	drop-down menu	Choose the test type. See smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) for descriptions of each type of test. Some test types will degrade performance or take disks offline. Avoid scheduling S.M.A.R.T. tests simultaneously with scrub or resilver operations.
Short description	string	Optional. Enter a short description of this test.
Hour	slider or hour selections	When the slider is used the sync occurs every N hours. Use *Each selected hour* for the test to occur at the highlighted hours.
Day of month	slider or day selections	When the slider is used the sync occurs every N days. Use *Each selected day of the month* for the sync to occur on the highlighted days.
Month	checkboxes	Select which months to run the test.
Day of week	checkboxes	Select which days of the week to run the test.

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

An example configuration is to schedule a *Short Self-Test* once a week and a *Long Self-Test* once a month. These tests do not have a performance impact, as the disks prioritize normal I/O over the tests. If a disk fails a test, even if the overall status is *Passed*, 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 244).

The results of a test can be checked from *Shell* (page 244) 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 an email to that address when a test fails. Logging information for S.M.A.R.T. tests can be found in `/var/log/daemon.log`.

CHAPTER

SIX

NETWORK

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

- *Global Configuration* (page 80): general network settings.
- *Interfaces* (page 82): settings for each network interface.
- *IPMI* (page 86): settings controlling connection to the appliance through the hardware side-band management interface if the graphical user interface becomes unavailable.
- *Link Aggregations* (page 87): settings for network link aggregation and link failover.
- *Network Summary* (page 92): display an overview of the current network settings.
- *Static Routes* (page 92): add static routes.
- *VLANs* (page 92): configure IEEE 802.1q tagging for virtual LANs.

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

Warning: Making changes to the network interface the web interface uses can result in losing connection to the TrueNAS® system! Misconfiguring network settings might require command line knowledge or physical access to the TrueNAS® system to fix. Be very careful when configuring *Interfaces* (page 82) and *Link Aggregations* (page 87).

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.

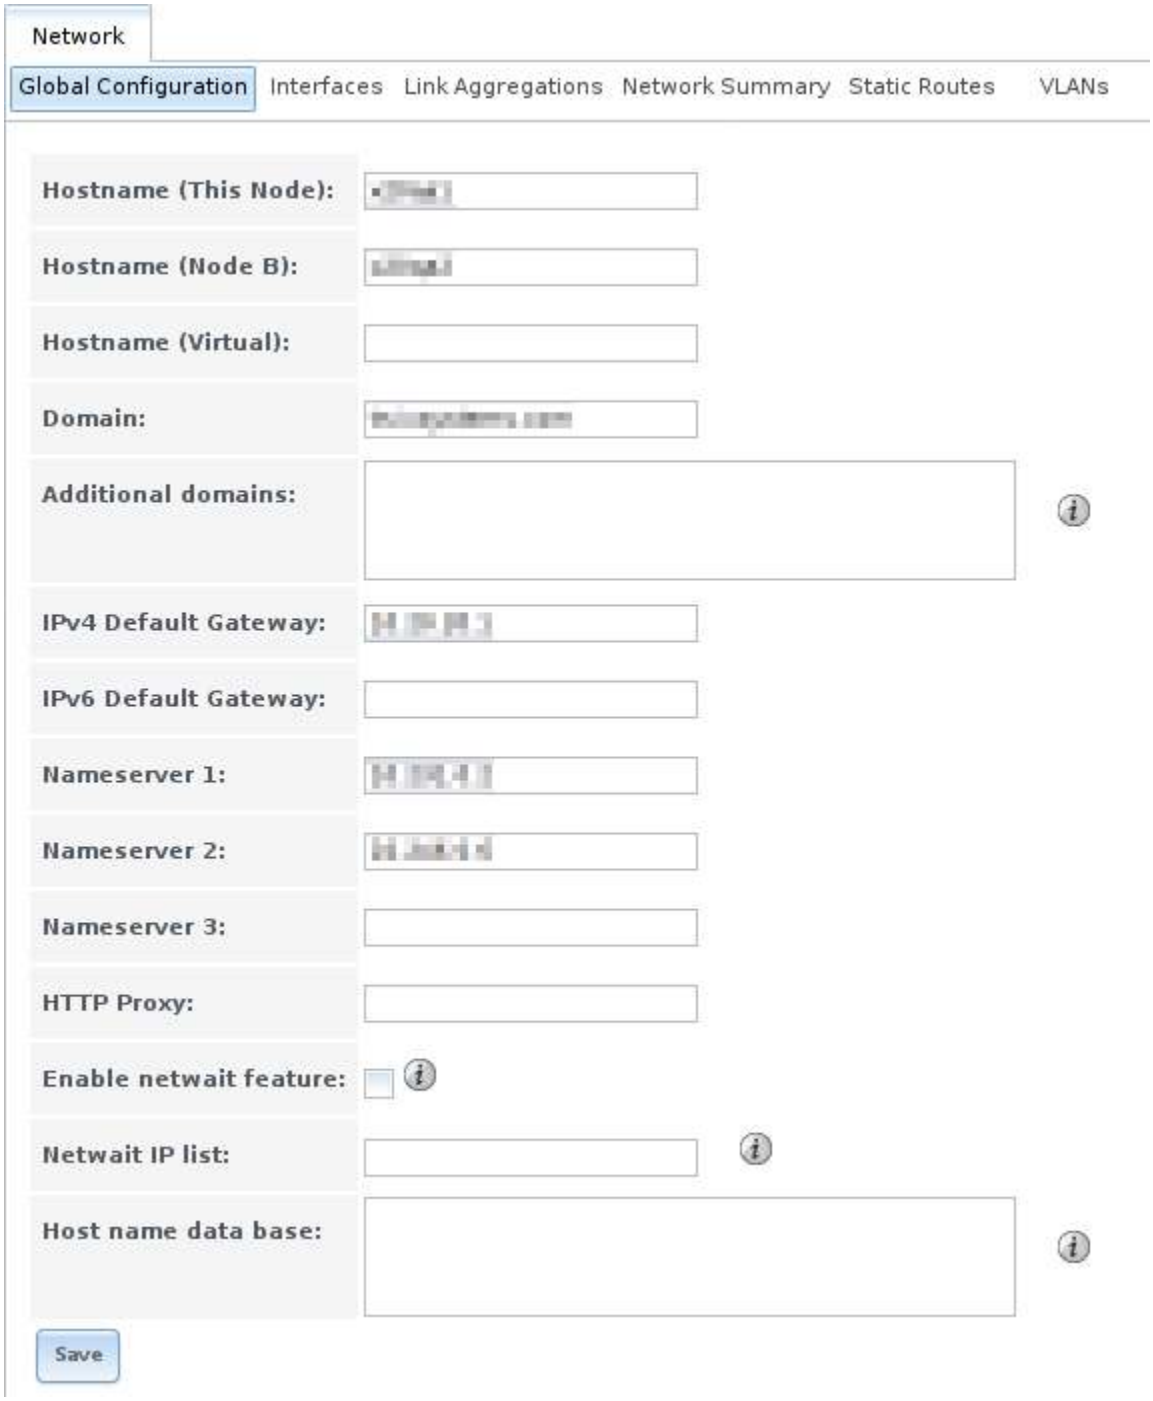

Fig. 6.1: Global Network Configuration

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

Table 6.1: Global Configuration Settings

Setting	Value	Description
Hostname (This Node)	string	Host name of first storage controller. Upper and lower case alphanumeric, . , and – characters are allowed.

Continued on next page

Table 6.1 – continued from previous page

Setting	Value	Description
Hostname (Node B)	string	Host name of second storage controller. Upper and lower case alphanumeric, ., and – characters are allowed.
Hostname (Virtual)	string	Virtual host name. When using a virtualhost, this is also used as the Kerberos principal name. Enter the fully qualified hostname plus the domain name. Upper and lower case alphanumeric, ., and – characters are allowed.
Domain	string	System domain name.
Additional domains	string	Can enter up to 6 space delimited search domains. Adding multiple domains may result in slower DNS lookups.
IPv4 Default Gateway	IP address	Typically not set. See *this note about Gateways* (page 82). If set, used instead of default gateway provided by DHCP.
IPv6 Default Gateway	IP address	Typically not set. See *this note about Gateways* (page 82).
Nameserver 1	IP address	Primary DNS server.
Nameserver 2	IP address	Secondary DNS server.
Nameserver 3	IP address	Tertiary DNS server.
HTTP Proxy	string	Enter the proxy information for the network in the format *http://my.proxy.server:3128* or *http://user:password@my.proxy.server:3128*.
Enable netwait feature	checkbox	If enabled, network services do not start at boot until the interface is able to ping the addresses listed in the *Netwait IP list*.
Netwait IP list	string	If *Enable netwait feature* is unset, 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 `/etc/hosts`. Use the format *IP_address space hostname* where multiple hostnames can be used if separated by a space.

When using Active Directory, 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 92) to be able to reach external DNS, NTP, and mail servers which are configured with static IP addresses. When a gateway to the Internet is added, make sure the TrueNAS® system is protected by a properly configured firewall.

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.

Creating a Link Aggregation (page 88) that does **not** include the NIC used to access the TrueNAS® administrative GUI may require adding an *Interfaces* entry for this interface with DHCP enabled. See this *warning* (page 80) about changing the interface that the web interface uses.

Figure 6.2 shows the screen that opens on clicking *Interfaces* → *Add Interface*. Table 6.2 summarizes the configuration options shown when adding an interface or editing an already configured interface. Note that if any changes to this screen require a network restart, the screen will turn red when the *OK* button is clicked and a pop-up message will point out that network connectivity to the TrueNAS® system will be interrupted while the changes are applied.

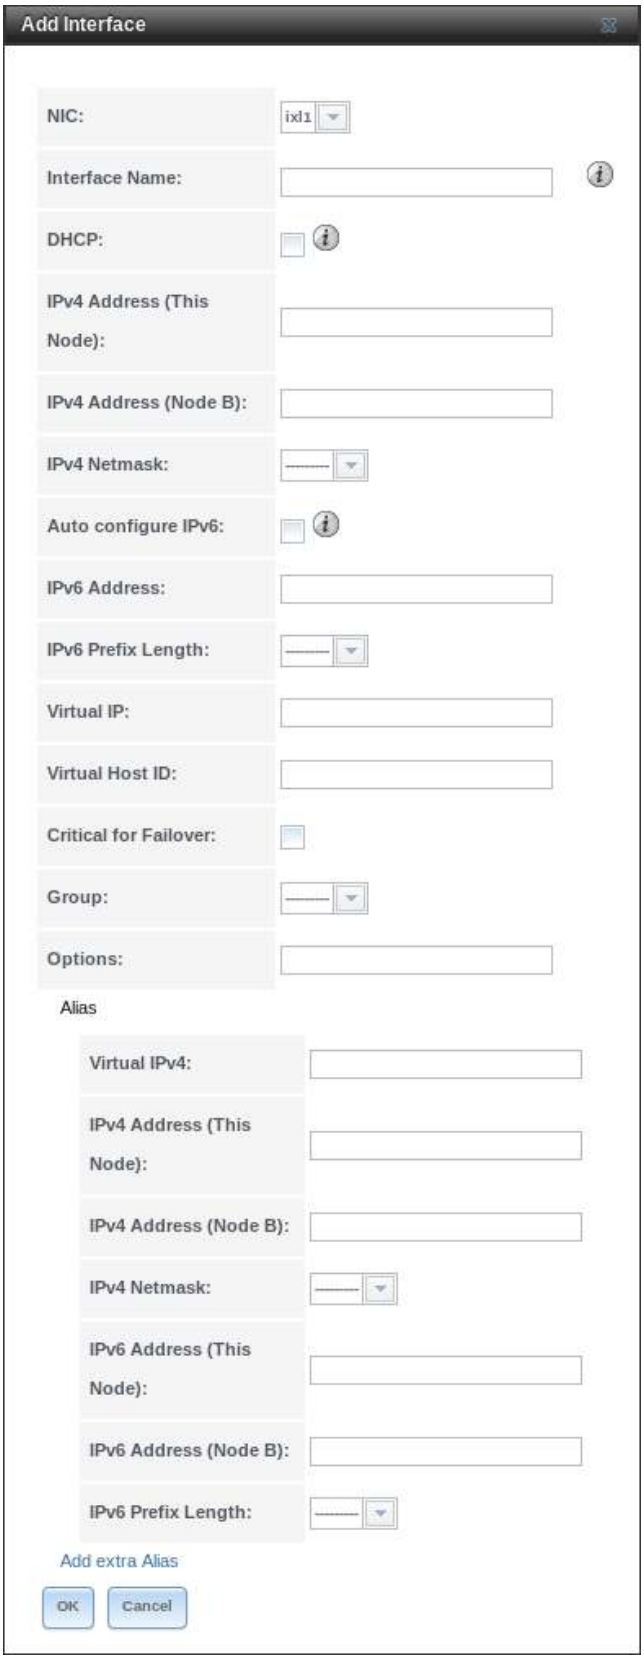

Fig. 6.2: Adding or Editing an Interface

Table 6.2: Interface Configuration Settings

Setting	Value	Description
NIC	drop-down menu	The FreeBSD device name of the interface. This is a read-only field when editing an interface.
Interface Name	string	Description of interface.
DHCP	checkbox	Requires static IPv4 or IPv6 configuration if unselected. Only one interface can be configured for DHCP.
IPv4 Address (This Node)	IP address	Enter a static IP address for the active storage controller if *DHCP* is unset.
IPv4 Address (Node B)	IP address	Enter a static IP address for the inactive storage controller if *DHCP* is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if *DHCP* is unset.
Auto configure IPv6	checkbox	Only one interface can be configured for this option. If unset, manual configuration is required to use IPv6.
IPv6 Address	IPv6 address	Must be unique on the network.
IPv6 Prefix Length	drop-down menu	Match the prefix used on the network.
Virtual IP	IP address	IP address for the virtual host. This is used to log in to the web interface.
Virtual Host ID	string	Unique identifier for the virtual host.
Critical for Failover	checkbox	Sets this interface as critical. This allows logging in to the web interface available at the *Virtual IP* address after a failover. **Warning:** At least one interface must have this option set or the web interface will become unavailable. This can also be set when configuring network interfaces in the *Console Setup Menu* (page 11).
Options	string	Additional parameters from ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig). Separate multiple parameters with a space. For example: *mtu 9000* increases the MTU for interfaces which support jumbo frames (but see *this note* (page 91) about MTU and lagg interfaces).

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

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

Note: Interfaces cannot be edited or deleted when *High Availability (HA)* (page 58) has been 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/) 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 to assist with a configuration or troubleshooting issue.

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

IPMI is configured from *Network* → *IPMI*. The IPMI configuration screen, shown in Figure 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.

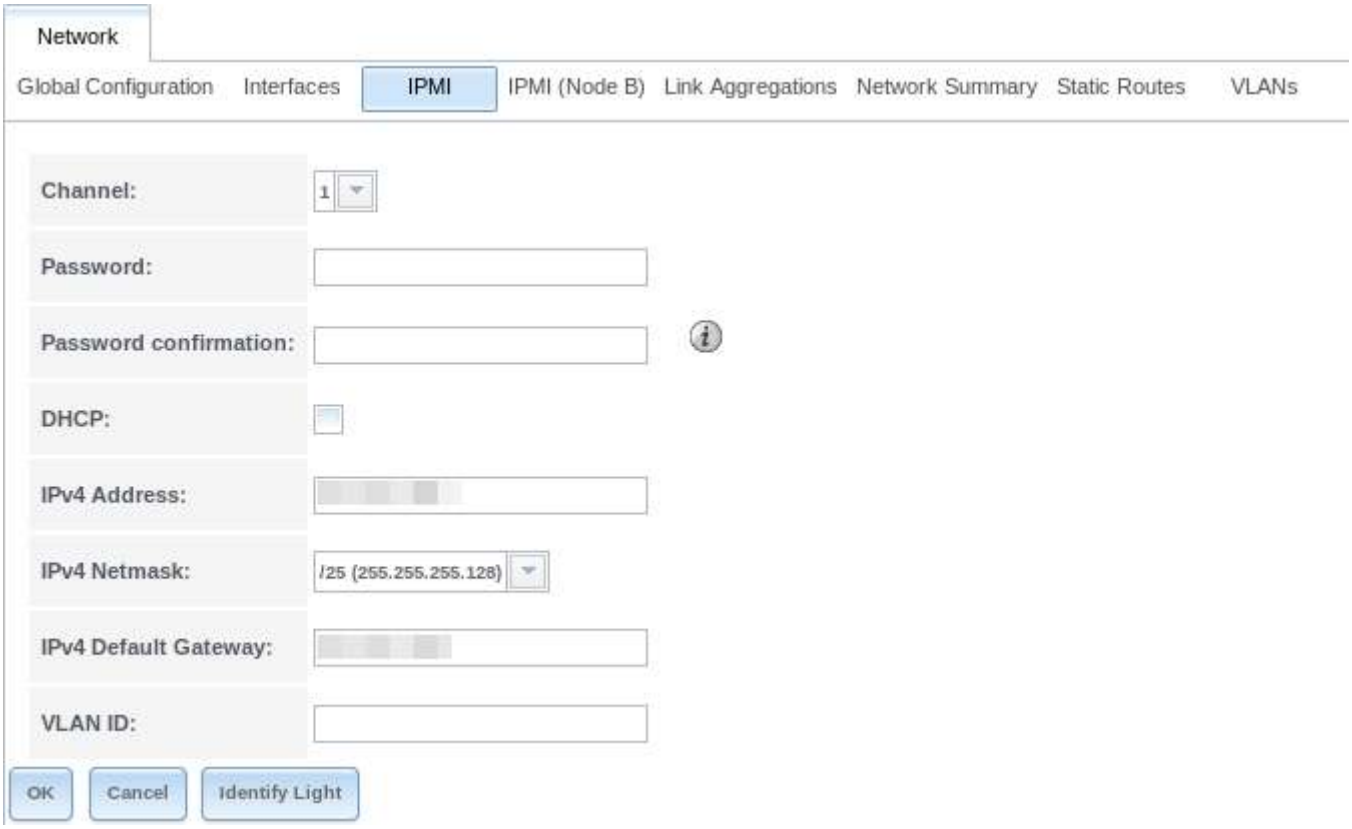

Fig. 6.3: IPMI Configuration

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. The maximum length is 20 characters.
DHCP	checkbox	If left unset, the next 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.

Continued on next page

Table 6.3 – continued from previous page

Setting	Value	Description
IPv4 Default Gateway	string	Default gateway associated with the IP address.
VLAN ID	string	Enter the VLAN identifier if the IPMI out-of-band management interface is not on the same VLAN as management networking.

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

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

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

6.4 Link Aggregations

TrueNAS® uses the FreeBSD lagg(4) (https://www.freebsd.org/cgi/man.cgi?query=lagg) interface to provide link aggregation and link failover support. A lagg interface allows combining multiple network interfaces into a single virtual interface. This provides fault-tolerance and high-speed multi-link throughput. The aggregation protocols supported by lagg both determines the ports to use for outgoing traffic and if a specific port accepts incoming traffic. The link state of the lagg interface is used to validate whether the port is active.

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

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

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

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

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

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

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

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

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

6.4.1 LACP, MPIO, NFS, and ESXi

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

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

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

LACP and other forms of link aggregation generally do not work well with virtualization solutions. In a virtualized environment, consider the use of iSCSI MPIO through the creation of an iSCSI Portal with at least two network cards on different networks. This allows an iSCSI initiator to recognize multiple links to a target, using them for increased bandwidth or redundancy. This how-to (https://fojta.wordpress.com/2010/04/13/iscsi-and-esxi-multipathing-and-jumbo-frames/) 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, make sure that all interfaces to use in the lagg are not manually configured in *Network* → *Interfaces*.

Lagg creation fails if any of the included interfaces are manually configured. See this *warning* (page 80) about changing the interface that the web interface uses.

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

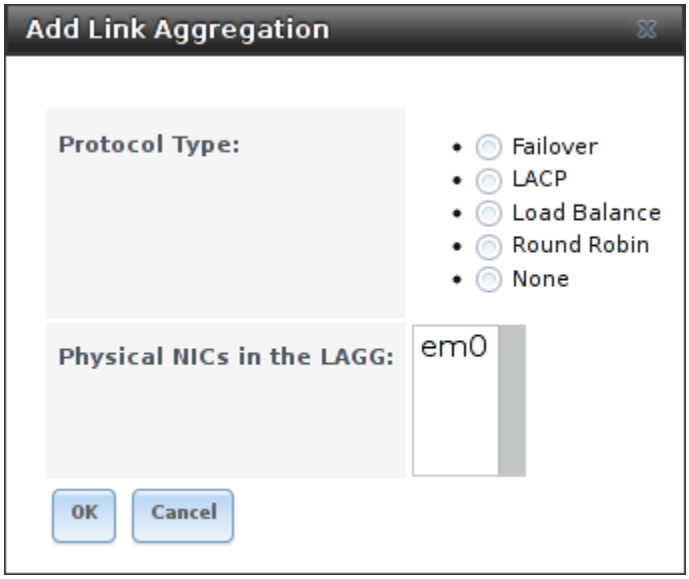

Fig. 6.4: Creating a lagg Interface

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.

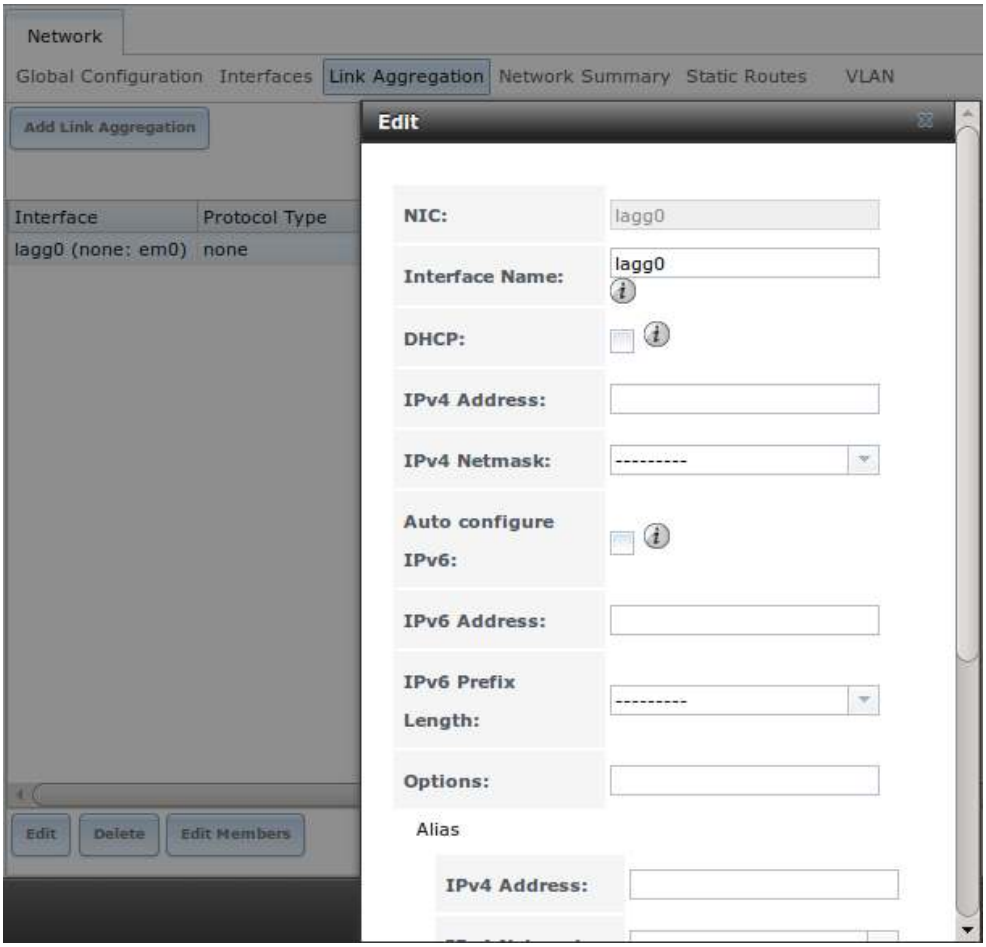

Fig. 6.5: Editing a lagg

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, this is the same as device (NIC) name. This can be changed to a more descriptive value.
DHCP	checkbox	Enable if the lagg device will get IP address info from DHCP server. The IP address of the new lagg can be set to DHCP only if no other interface uses DHCP.
IPv4 Address	string	Enter a static IP address if *DHCP* is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if *DHCP* is unset.
Auto configure IPv6	checkbox	Set only if DHCP server available to provide IPv6 address info
IPv6 Address	string	This is optional.
IPv6 Prefix Length	drop-down menu	Required if an *IPv6 address* is entered.
Options	string	Additional ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig) options.

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

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

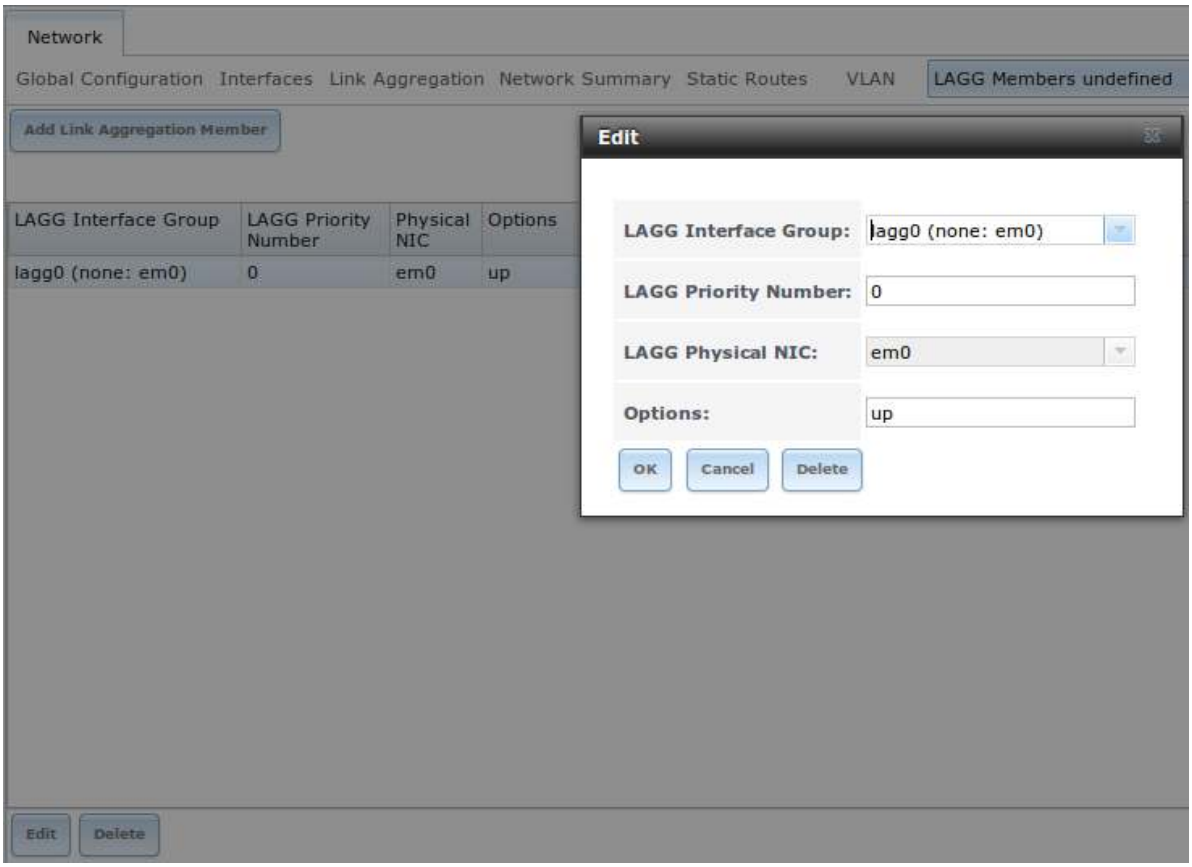

Fig. 6.6: Editing a Member Interface

Table 6.5: Configuring a Member Interface

Setting	Value	Description
LAGG Interface group	drop-down menu	Select the member interface to configure.
LAGG Priority Number	integer	Order of selected interface within the lagg. Configure a failover to set the master interface to *0* and the other interfaces to *1*, *2*, etc.
LAGG Physical NIC	drop-down menu	Physical interface of the selected member. The drop-down is empty when no NICs are available.
Options	string	Additional parameters from ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig).

Click *Add Link Aggregation Member* to see the same options. Click *OK* to add the new member to the list.

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, repeat the configuration for each interface within the lagg. Some options can only be set on the parent interfaces and are inherited by the lagg interface. For example, to set the MTU on a lagg, use *Edit Members* to set the MTU for each parent interface.

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

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

Link aggregation load balancing can be tested with:

```
systat -ifstat
```

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

6.5 Network Summary

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

6.6 Static Routes

No static routes are defined on a default TrueNAS® system. If a static route is required to reach portions of the network, add the route with *Network* → *Static Routes* → *Add Static Route*, shown in Figure 6.7.

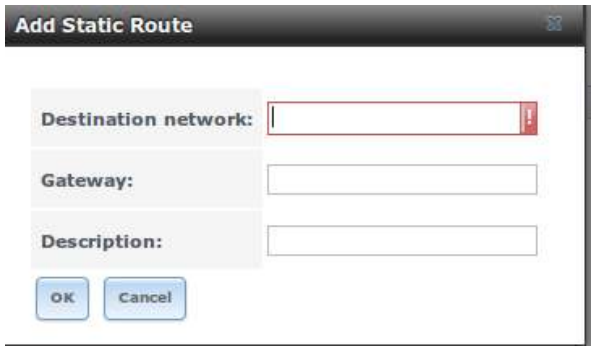

Fig. 6.7: Adding a Static Route

The available options are summarized in Table 6.6.

Table 6.6: Static Route Options

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

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) (https://www.freebsd.org/cgi/man.cgi?query=vlan) interface to demultiplex frames with IEEE 802.1q tags. This allows nodes on different VLANs to communicate through a layer 3 switch or router. A vlan interface must be assigned a parent interface and a numeric VLAN tag. A single parent can be assigned to multiple vlan interfaces provided they have different tags.

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

Click *Network* → *VLANs* → *Add VLAN*, to see the screen shown in Figure 6.8.

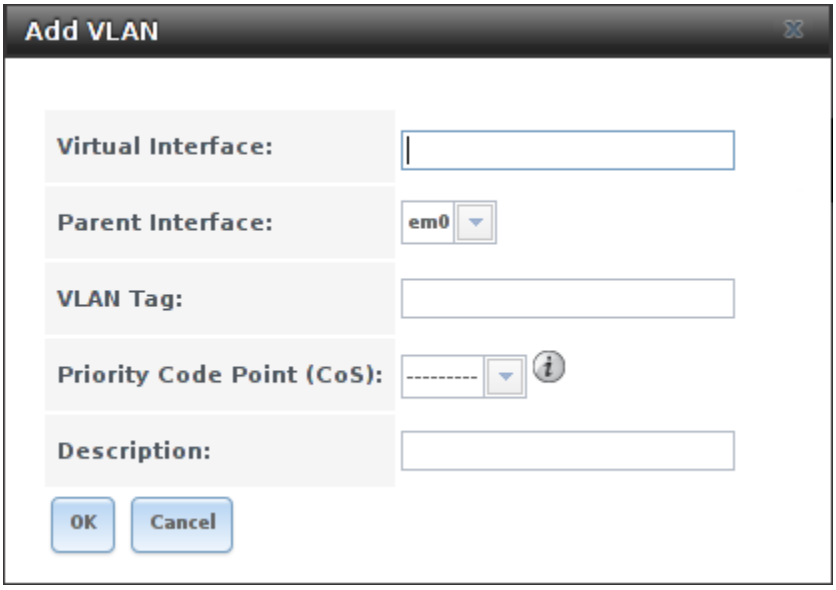

Fig. 6.8: Adding a VLAN

Table 6.7 summarizes the configurable fields.

Table 6.7: Adding a VLAN

Setting	Value	Description
Virtual Interface	string	Use the format *vlanX* where *X* is a number representing a vlan interface not currently being used as a parent.
Parent Interface	drop-down menu	Usually an Ethernet card connected to a properly configured switch port. Newly created *Link Aggregations* (page 87) do not appear in the drop-down until the system is rebooted.
VLAN Tag	integer	Enter a number between *1* and *4095* which matches a numeric tag set up in the switched network.
Priority Code Point	drop-down menu	Available 802.1p Class of Service ranges from *Best Effort (default)* to *Network Control (highest)*.
Description	string	Optional. Enter any notes about this VLAN.

The parent interface of a VLAN must be up, but it can either have an IP address or be unconfigured, depending upon the requirements of the VLAN configuration. This makes it difficult for the GUI to do the right thing without trampling the configuration. To remedy this, add the VLAN, then select *Network* → *Interfaces* → *Add Interface*. Choose 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 causes an interruption to network connectivity and, if *Failover* (page 58) is configured, a failover event. The GUI provides a warning and an opportunity to cancel the VLAN creation.

CHAPTER

SEVEN

STORAGE

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

- *Volumes* (page 94) create and manage storage volumes.
- *Periodic Snapshot Tasks* (page 120) schedule automatic creation of filesystem snapshots.
- *Replication Tasks* (page 122) automate the replication of snapshots to a remote system.
- *Resilver Priority* (page 133) control the priority of resilvers.
- *Scrubs* (page 134) schedule scrubs as part of ongoing disk maintenance.
- *Snapshots* (page 137) manage local snapshots.
- *VMware-Snapshot* (page 139) coordinate OpenZFS snapshots with a VMware datastore.

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

7.1 Swap Space

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

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

Swap space is allocated when drives are partitioned before being added to a *vdev* (page 250). A 2 GiB partition for swap space is created on each data drive by default. The size of space to allocate can be changed in *System* → *Advanced* in the *Swap size on each drive in Gib, affects new disks only. Setting this to 0 disables swap creation completely (STRONGLY DISCOURAGED)* field. Changing the value does not affect the amount of swap on existing disks, only disks added after the change. This does not affect log or cache devices, which are created without swap. Swap can be disabled by entering *0*, but that is **strongly discouraged**.

7.2 Volumes

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

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

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

7.2.1 Volume Manager

Before creating a volume, determine the level of required redundancy, how many disks will be added, and if any data exists on those disks. Creating a volume overwrites disk data, so save any required data to different media before adding disks to a pool. Refer to the *ZFS Primer* (page 250) for information on ZFS redundancy with multiple disks before using *Volume Manager*. It is important to realize that different layouts of virtual devices (*vdevs*) affect which operations can be performed on that volume later. For example, drives can be added to a mirror to increase redundancy, but that is not possible with RAIDZ arrays.

To create a volume, click *Storage* → *Volumes* → *Volume Manager*. This opens a screen like the example shown in Figure 7.1.

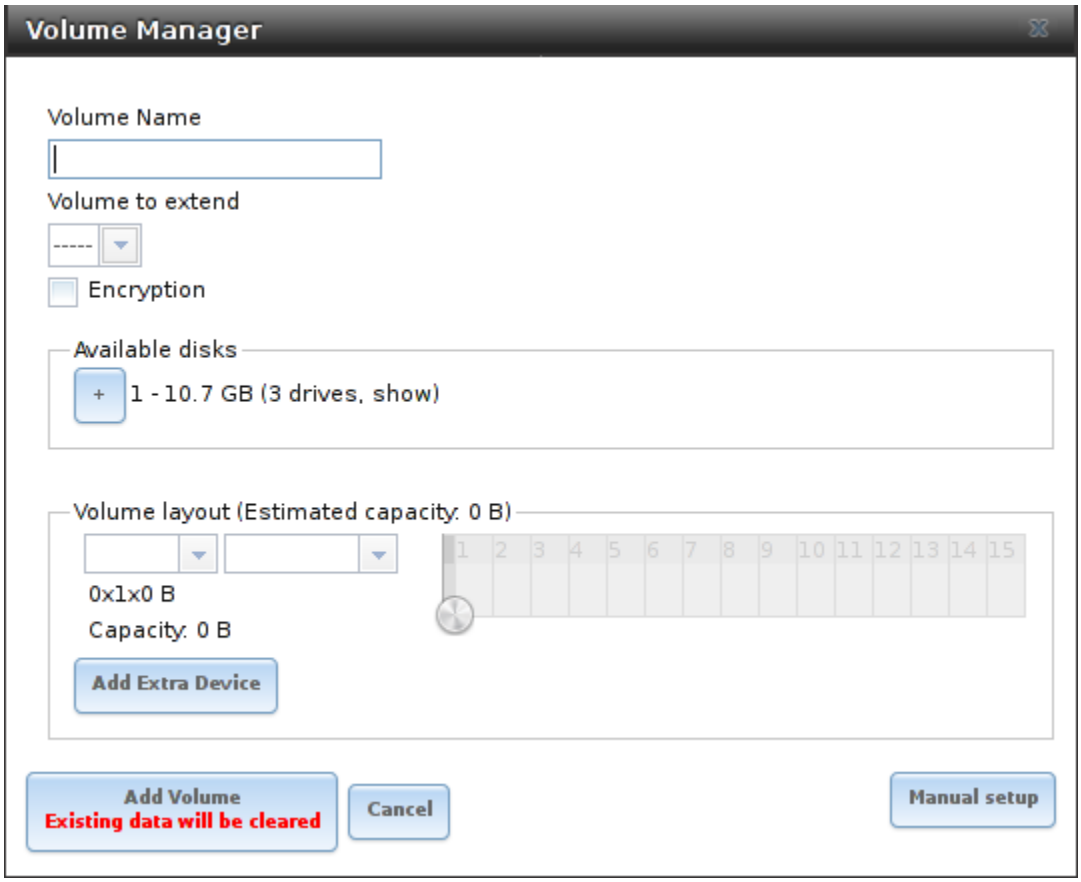

Fig. 7.1: Creating a ZFS Pool Using Volume Manager

Table 7.1 summarizes the configuration options of this screen.

Table 7.1: ZFS Volume Creation Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (https://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html) Choose a memorable name that sticks out in the logs and avoid generic names.

Continued on next page

Table 7.1 – continued from previous page

Setting	Value	Description
Volume to extend	drop-down menu	Extend an existing ZFS pool. See *Extending a ZFS Volume* (page 99) for more details.
Encryption	checkbox	See the warnings in *Encryption* (page 97) before enabling encryption.
Available disks	display	Display the number and size of available disks. Hover over *show* to list the available device names, and click the + to add all of the disks to the pool.
Volume layout	drag and drop	Click and drag the icon to select the desired number of disks for a vdev. When at least one disk is selected, the layouts supported by the selected number of disks are added to the drop-down menu.
Add Extra Device	button	Configure multiple vdevs or add log or cache devices during pool creation.
Manual setup	button	Create a pool manually, which is not recommended. See *Manual Setup* (page 98) for more details.

Click the *Volume name* field and enter a name for the pool. Ensure that the chosen name conforms to these naming conventions (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html).

If the underlying disks need to be encrypted as a protection against physical theft, enable the *Encryption* option.

Warning: Refer to the warnings in *Encryption* (page 97) before enabling encryption! Be aware that this form of encryption will be replaced by OpenZFS native encryption in a future version. Volumes created with the current encryption mechanism will have to be backed up and destroyed to be recreated with native encryption when it becomes available.

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

Note: For performance and capacity reasons, this screen does not allow creating a volume from disks of differing sizes. While it is not recommended, it is possible to create a volume of differently-sized disks with the *Manual setup* button. Follow the instructions in *Manual Setup* (page 98).

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 250).

The *Add Volume* button warns that **existing data will be cleared**. In other words, creating a new volume **reformats the selected disks**. To preserve existing data, click the *Cancel* button and refer to *Import Disk* (page 106) and *Import Volume* (page 107) to see if the existing format is supported. If so, perform that action instead. If the current storage format is not supported, it is necessary to back up the data to external media, format the disks, then restore the data to the new volume.

Depending on the size and number of disks, the type of controller, and whether encryption is selected, creating the volume may take some time. After the volume is created, the screen refreshes and the new volume is listed in the tree under *Storage* → *Volumes*. Click the + next to the volume name to access *Change Permissions* (page 100), *Create Dataset* (page 102), and *Create zvol* (page 105) options for that volume.

7.2.1.1 Encryption

Note: TrueNAS® uses GELI (https://www.freebsd.org/cgi/man.cgi?query=geli) full disk encryption for ZFS volumes. This type of encryption is primarily intended to protect data against the risks of data being read or copied when the system is powered down, when the volume is locked, or when disks are physically stolen.

Because data cannot be read without the key, encrypted disks containing sensitive data can be safely removed, reused, or discarded without secure wiping or physical destruction of the media.

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

TrueNAS® encrypts disks and volumes, not individual filesystems. The partition table on each disk is not encrypted, but only identifies the location of partitions on the disk. On an encrypted volume, the data in each partition is encrypted.

Encrypted volumes which do not have a passphrase are unlocked at startup. Volumes with a passphrase remain locked until the user enters the passphrase to unlock them.

Encrypted volumes can be locked on demand by the user. They are automatically locked when the system is shut down.

Understanding the details of TrueNAS® encryption is required to be able to use it effectively:

- TrueNAS® encryption differs from the encryption used in Oracle's proprietary version of ZFS. To convert between these formats, both volumes must be unlocked, and the data copied between them.
- When discarding disks that still contain encrypted sensitive data, the encryption key must also be destroyed or securely deleted. If the encryption key is not destroyed, it must be stored securely and kept physically separate from the discarded disks. If the encryption key is present on or with the discarded disks, or can be obtained by the same person who gains access to the disks, the data will be vulnerable to decryption.
- Protect the key with a strong passphrase and store all key backups securely. If the encryption key is lost, the data on the disks is inaccessible. Always back up the key!
- Encryption keys are per ZFS volume. Each volume has a separate encryption key. Technical details about how encryption key use, storage, and management are described in this forum post (https://forums.freenas.org/index.php?threads/recover-encryption-key.16593/#post-85497).
- All drives in an encrypted volume are encrypted, including L2ARC (read cache) and SLOG (write intent log). Drives added to an existing encrypted volume are encrypted with the same method specified when the volume was created. Swap data on disk is always encrypted. Data in memory (RAM), including ARC, is not encrypted.
- At present, there is no one-step way to encrypt an existing volume. The data must be copied to an existing or new encrypted volume. After that, the original volume and any unencrypted backup should be destroyed to prevent unauthorized access and any disks that contained unencrypted data should be wiped.
- Hybrid volumes are not supported. Added vdevs must match the existing encryption scheme. *Volume Manager* (page 95) automatically encrypts new vdevs added to an existing encrypted volume.

To create an encrypted volume, enable the *Encryption* option shown in Figure 7.1. A pop-up message shows a reminder that **it is extremely important to back up the key**. Without the key, the data on the disks is inaccessible. See *Managing Encrypted Volumes* (page 114) for instructions.

7.2.1.2 Encryption Performance

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

7.2.1.3 Manual Setup

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

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

Figure 7.2 shows the *Manual Setup* screen. Table 7.2 shows the available options.

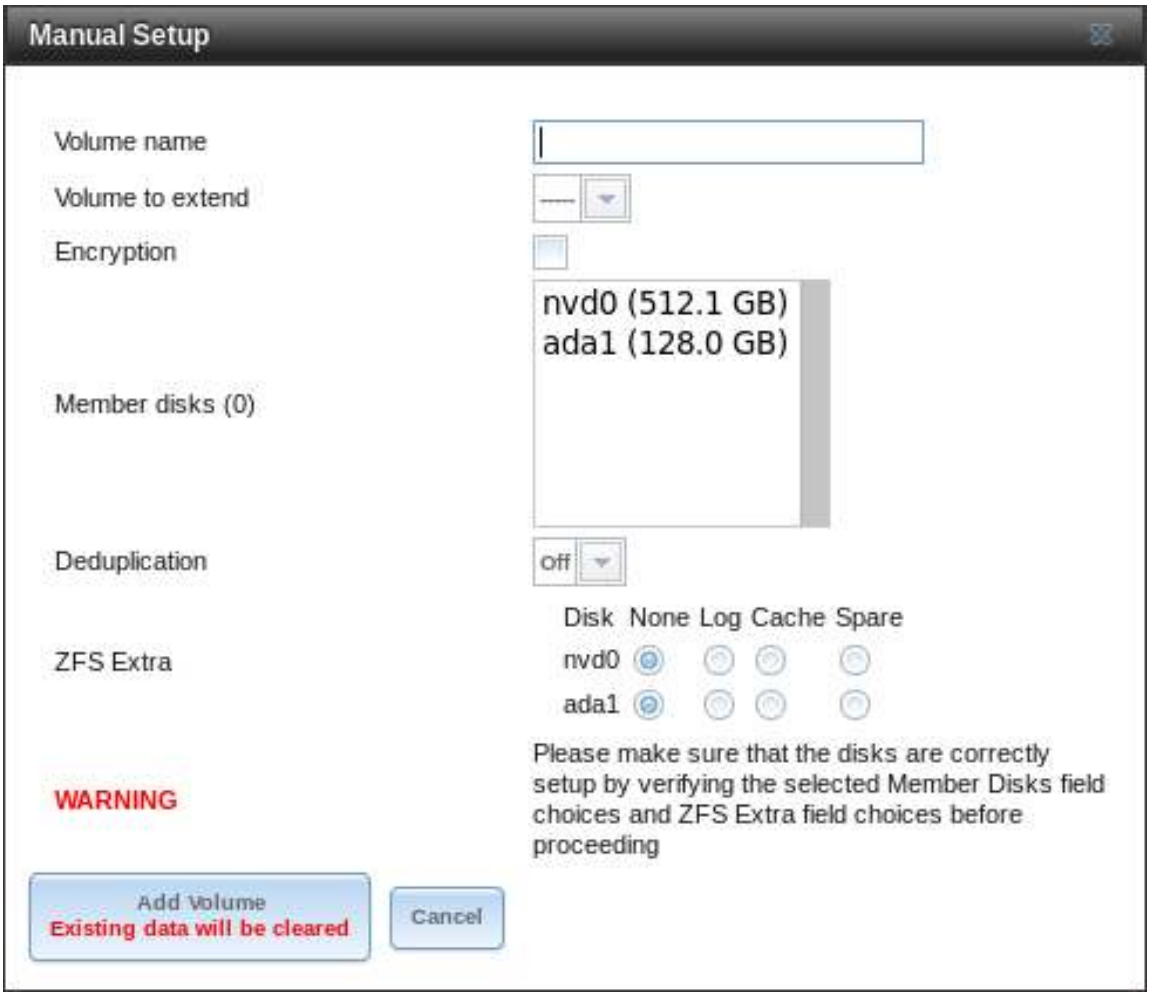

Fig. 7.2: Manually Creating a ZFS Volume

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

Table 7.2: Manual Setup Options

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions (https://docs.oracle.com/cd/E53394_01/index.html). Choosing a unique, memorable name is recommended.
Volume to extend	drop-down menu	Extend an existing ZFS pool. See *Extending a ZFS Volume* (page 99) for more details.
Encryption	checkbox	See the warnings in *Encryption* (page 97) before using encryption.
Member disks	list	Highlight desired number of disks from list of available disks. Hold `Ctrl` and click a highlighted item to de-select it. Selecting a member disk removes it from the *ZFS Extra* list.
Deduplication	drop-down menu	Do not change this setting unless instructed to do so by an iXsystems support engineer.
ZFS Extra	bullet selection	Specify disk usage: storage (*None*), a log device, a cache device, or a spare. Choosing a value other than *None* removes the disk from the *Member disks* list'.

7.2.1.4 Extending a ZFS Volume

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

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

Note: If the existing volume is encrypted, a warning message shows a reminder that **extending a volume resets the passphrase and recovery key**. After extending the volume, immediately recreate both using the instructions in *Managing Encrypted Volumes* (page 114).

After an existing volume has been selected from the drop-down menu, drag and drop the desired disks and select the desired volume layout. For example, disks can be added to increase the capacity of the volume.

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

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

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

Adding L2ARC or SLOG Devices

Storage → *Volumes* → *Volume Manager* (see Figure 7.1) is also used to add L2ARC or SLOG SSDs to improve volume performance for specific use cases. Refer to the *ZFS Primer* (page 250) to determine if the system will benefit or suffer from the addition of the device.

Once the SSD has 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.

Removing L2ARC or SLOG Devices

Cache or log devices can be removed by going to *Storage* → *Volumes*. Choose the desired pool and click *Volume Status*. Choose the log or cache device to remove, then click *Remove*.

7.2.2 Change Permissions

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

Sharing (page 153) contains configuration examples for several types of permission scenarios. This section provides an overview of the options available for configuring the initial set of permissions.

Note: For users and groups to be available, they must either be first created using the instructions in *Account* (page 16) or imported from a directory service using the instructions in *Directory Services* (page 141). If more than 50 users or groups are available, the drop-down menus described in this section will automatically truncate their display to 50 for performance reasons. In this case, start to type in the desired user or group name so that the display narrows its search to matching results.

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

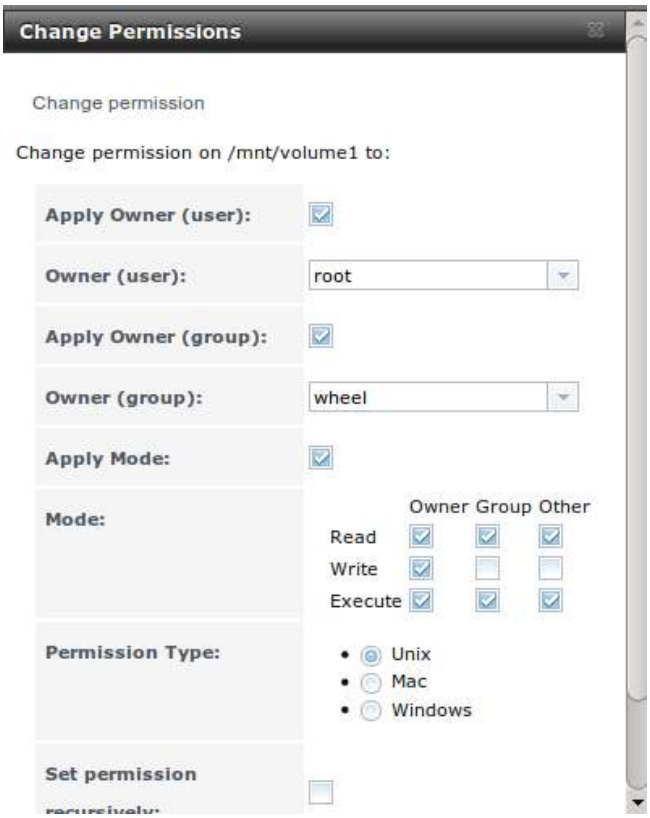

Fig. 7.3: Changing Permissions on a Volume or Dataset

Table 7.3: Options When Changing Permissions

Setting	Value	Description
Apply Owner (user)	checkbox	Deselect to prevent new permission change from being applied to *Owner (user)*, see Note below.
Owner (user)	drop-down menu	Select the user to control the volume or dataset. Users manually created or imported from a directory service will appear in the drop-down menu.
Apply Owner (group)	checkbox	Deselect to prevent new permission change from being applied to *Owner (group)*, see Note below for more information.
Owner (group)	drop-down menu	Select the group to control the volume or dataset. Groups manually created or imported from a directory service will appear in the drop-down menu.
Apply Mode	checkbox	Deselect to prevent new permission change from being applied to *Mode*, see Note below.
Mode	checkboxes	Only applies to the *Unix* or *Mac* "Permission Type". Will be grayed out if *Windows* is selected.
Permission Type	bullet selection	Select the type which matches the type of client accessing the volume or dataset. Choices are *Unix*, *Mac*, or *Windows*.
Set permission recursively	checkbox	If enabled, permissions will also apply to subdirectories of the volume or dataset. If data already exists on the volume or dataset, change the permissions on the **client side** to prevent a performance lag.

Note: The *Apply Owner (user)*, *Apply Owner (group)*, and *Apply Mode* options allow fine-tuning of the change permissions behavior. By default, all options are enabled and TrueNAS® resets the owner, group, and mode when the

Change button is clicked. These optionss allow choosing which settings to change. For example, to change just the *Owner (group)* setting, deselect the *Apply Owner (user)* and *Apply Mode* options.

The *Windows Permission Type* is used for *Windows (SMB) Shares* (page 166) or when the TrueNAS® system is a member of an Active Directory domain. This type adds ACLs to traditional *Unix* permissions. When the *Windows Permission Type* is set, ACLs are set to the Windows defaults for new files and directories. A Windows client can be used to further fine-tune permissions as needed.

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

The *Unix Permission Type* is usually used with *Unix (NFS) Shares* (page 158). Unix permissions are compatible with most network clients and generally work well with a mix of operating systems or clients. However, *Unix* permissions do not support Windows ACLs. Do not use them with *Windows (SMB) Shares* (page 166).

The *Mac Permission Type* can be used with *Apple (AFP) Shares* (page 154).

7.2.3 Create Dataset

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

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

Selecting an existing ZFS volume in the tree and clicking *Create Dataset* shows the screen in Figure 7.2.3.

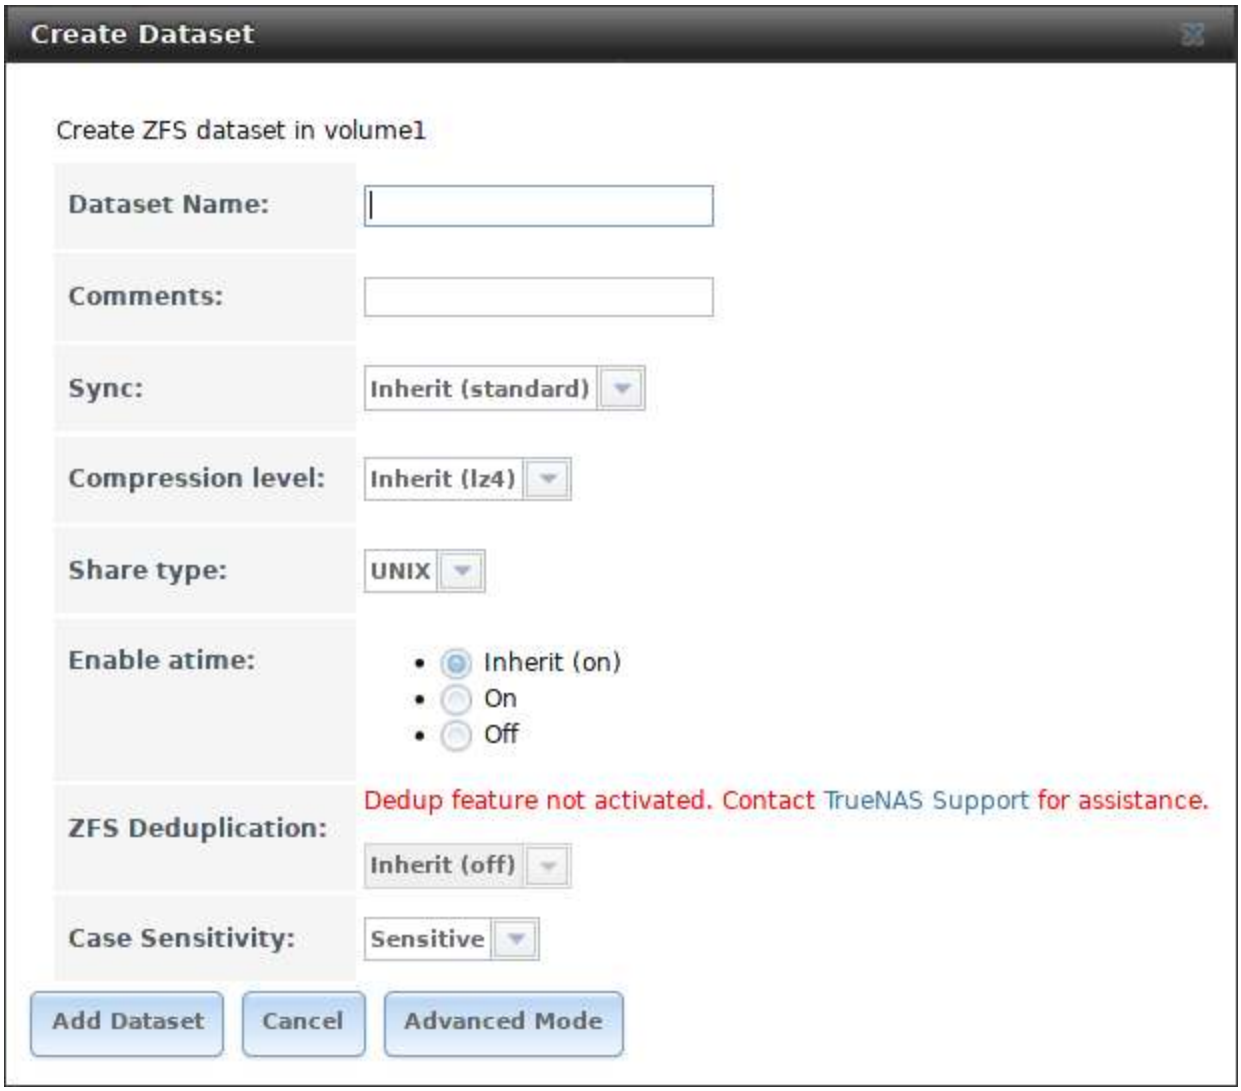

Fig. 7.4: Creating a ZFS Dataset

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

Table 7.4: ZFS Dataset Options

Setting	Value	Description
Dataset Name	string	Enter a mandatory unique name for the dataset.
Comments	string	Enter optional comments or notes about this dataset.
Sync	drop-down menu	Sets the data write synchronization. *Inherit* inherits the sync settings from the parent dataset. *Always* always waits. *Standard* uses the sync settings that are requested by the client software for data writes to complete. *Disabled* never waits for writes to complete.
Compression Level	drop-down menu	Refer to the section on *Compression* (page 104) for a description of the available algorithms.

Continued on next page

Table 7.4 – continued from previous page

Setting	Value	Description
Share type	drop-down menu	Select the type of share that will be used on the dataset. Choices are *UNIX* for an NFS share, *Windows* for a SMB share, or *Mac* for an AFP share.
Enable atime	Inherit, On, or Off	Choose *On* to update the access time for files when they are read. Choose *Off* to prevent producing log traffic when reading files. This can result in significant performance gains.
Quota for this dataset	integer	Only available in *Advanced Mode*. Default of *0* disables quotas. Specifying a value uses no more than the specified size and is suitable for user datasets to prevent users from taking all available space.
Quota for this dataset and all children	integer	Only available in *Advanced Mode*. A specified value applies to both this dataset and any child datasets.
Reserved space for this dataset	integer	Only available in *Advanced Mode*. Default of *0* is unlimited. Specifying a value keeps at least this much space free and is suitable for datasets with logs that could take all free space.
Reserved space for this dataset and all children	integer	Only available in *Advanced Mode*. A specified value applies to both this dataset and any child datasets.
ZFS Deduplication	drop-down menu	Do not change this setting unless instructed to do so by an iXsystems support engineer.
Read-Only	drop-down menu	Only available in *Advanced Mode*. Choices are *Inherit (off)*, *On*, or *Off*.
Exec	drop-down menu	Only available in *Advanced Mode*. Choices are *Inherit (on)*, *On*, or *Off*.
Record Size	drop-down menu	Only available in *Advanced Mode*. While ZFS automatically adapts the record size dynamically to adapt to data, if the data has a fixed size, matching that size can result in better performance.
Case Sensitivity	drop-down menu	*Sensitive* is the default and assumes filenames are case sensitive. *Insensitive* assumes filenames are not case sensitive. *Mixed* understands both types of filenames.

Create a nested dataset by clicking on an existing dataset and selecting *Create Dataset*. A zvol can also be created within a dataset.

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

7.2.3.1 Compression

When selecting a compression type, try 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:** default and 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.
- **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 which eliminates runs of zeroes.

- **lzjb:** provides decent data compression, but is considered deprecated as *lz4* provides much better performance.

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

7.2.4 Create zvol

A zvol is a feature of ZFS that creates a raw block device over ZFS. The zvol can be used as an *iSCSI* (page 210) 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.

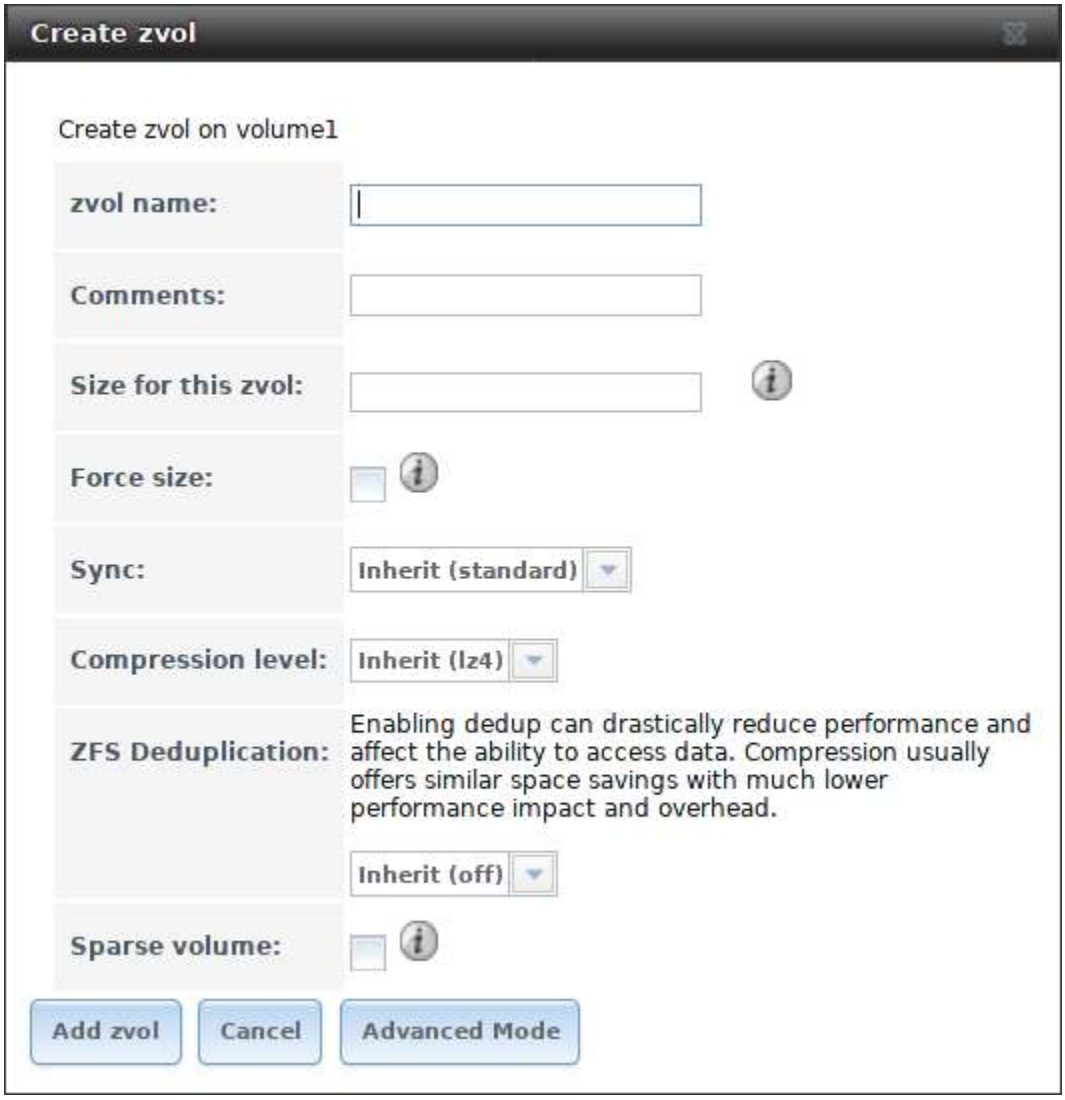

Fig. 7.5: Creating a Zvol

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 enabling *Show advanced fields by default* in *System* → *Advanced*.

Table 7.5: zvol Configuration Options

Setting	Value	Description
zvol Name	string	Enter a short name for the zvol. Using a zvol name longer than 63-characters can prevent accessing zvols as devices. For example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent. This setting is mandatory.
Comments	string	Enter any notes about this zvol.
Size for this zvol	integer	Specify size and value such as *10Gib*. If the size is more than 80% of the available capacity, the creation will fail with an "out of space" error unless *Force size* is also enabled.
Force size	checkbox	By default, the system does not create a zvol when it brings the pool above 80% capacity. **While NOT recommended**, enabling this option will force the creation of the zvol.
Compression level	drop-down menu	Refer to the section on *Compression* (page 104) for a description of the available algorithms.
ZFS Deduplication	drop-down menu	Do not change this setting unless instructed to do so by an iXsystems support engineer.
Sparse volume	checkbox	Used to provide thin provisioning. **Caution:** when this option is set, writes will fail when the pool is low on space.
Block size	drop-down menu	Only available in *Advanced Mode*. The default is based on the number of disks in the pool. Can be set to match the block size of the filesystem to be formatted onto the iSCSI target.

7.2.5 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 (BSD Unix), FAT or NTFS (Windows), or EXT2 (Linux) filesystems. 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/), 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.

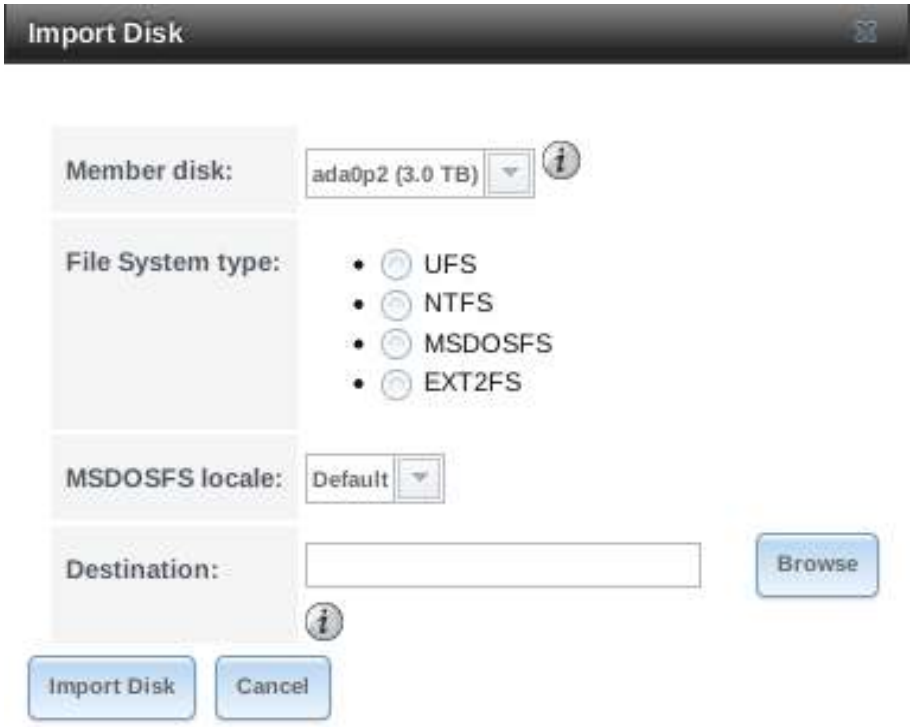

Fig. 7.6: Importing a Disk

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. If the *MSDOSFS* filesystem is selected, the *MSDOSFS locale* drop-down menu can be used to select the locale when non-ascii characters are present on the disk.

Once *Import Disk* is clicked, the disk is mounted, its contents are copied to the specified ZFS dataset, and the disk is unmounted after the copy operation completes.

7.2.6 Import Volume

Click *Storage* → *Volumes* → *Import Volume*, to 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 a volume is imported.

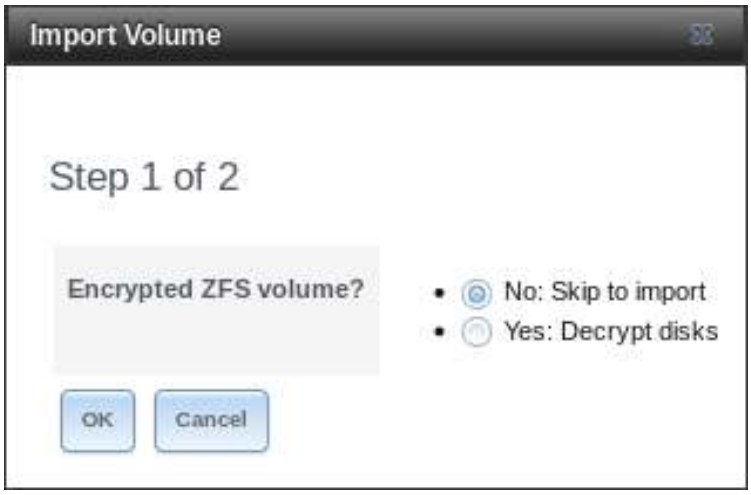

Fig. 7.7: Initial Import Volume Screen

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

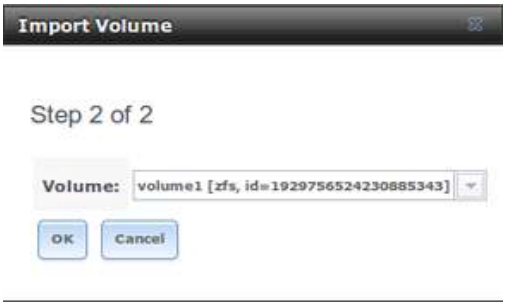

Fig. 7.8: Importing a Non-Encrypted Volume

Existing volumes are 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 244) to import the pool.

If physically installing ZFS formatted disks from another system, ensure to export the drives on that system to prevent an "in use by another machine" error during the import.

7.2.6.1 Importing an Encrypted Volume

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

Fig. 7.9: Decrypting Disks Before Importing a Volume

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

Note: The encryption key is required to decrypt the volume. If the volume 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 114) for instructions on how to manage the keys for encrypted volumes.

After the volume is decrypted, it appears in the drop-down menu of Figure 7.8. Click the *OK* button to finish the volume import.

Note: For security reasons, GELI keys for encrypted volumes are not saved in a configuration backup file. When TrueNAS® has been installed to a new device and a saved configuration file restored to it, the GELI keys for encrypted disks will not be present, and the system will not request them. To correct this, export the encrypted volume with Detach Volume, making sure that the options *Mark the disks as new (destroy data)* or *Also delete the share's configuration* are **not** selected. Then import the volume again. During the import, the GELI keys can be entered as described above.

7.2.7 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	Password for SED
ada0	WD-WMC4N2694653	3.0 TB		Auto	Always On	Disabled	Disabled	true		
ada1	WD-WCC4N4DDKSEH	3.0 TB		Auto	Always On	Disabled	Disabled	true		
ada2	WD-WCC4N5SDAAF3	3.0 TB		Auto	Always On	Disabled	Disabled	true		
ada3	WD-WCC4N2PNTD34	3.0 TB		Auto	Always On	Disabled	Disabled	true		
ada4	E011734600000000159	32.0 GB		Auto	Always On	Disabled	Disabled	true		

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.

To bulk edit disks, hold `Shift` and click each disk to edit. *Edit* changes to *Edit In Bulk*. Click it to open the *Edit In Bulk* window. This window displays which disks are being edited and a short list of configurable options. The *Disk Options table* (page 110) indicates the options available when editing multiple disks.

Table 7.6: Disk Options

Setting	Value	Bulk Edit	Description
Name	string		This is the FreeBSD device name for the disk.
Serial	string		This is the serial number of the disk.
Description	string		Enter any notes about this disk.
HDD Standby	drop-down menu	✓	Indicates the time of inactivity in minutes before the drive enters standby mode to conserve energy. This forum post (https://forums.freenas.org/index.php?threads/how-to-find-out-if-a-drive-is-spinning-down-properly.2068/) demonstrates how to determine if a drive has spun down.
Advanced Power Management	drop-down menu	✓	Select a power management profile from the menu. The default value is *Disabled*.
Acoustic Level	drop-down menu	✓	Default is *Disabled*. Other values can be selected for disks that understand AAM (https://en.wikipedia.org/wiki/Automatic_acoustic_management).
Enable S.M.A.R.T.	checkbox	✓	Enabled by default when the disk supports S.M.A.R.T. Disabling S.M.A.R.T. tests prevents collecting new temperature data for this disk. Historical temperature data is still displayed in *Reporting* (page 235).
S.M.A.R.T. extra options	string	✓	Enter additional smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in) options.
Password for SED	string		Enter and confirm the password which will be used for this device instead of the global SED password. Refer to *Self-Encrypting Drives* (page 31) for more information.
Reset Password	checkbox		Set to clear the SED password.

Note: If the serial number of a disk is not displayed in this screen, use the `smartctl` command from *Shell* (page 244). 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.

7.2.8 Volumes

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

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

Buttons are provided for quick access to *Volume Manager*, *Import Disk*, *Import Volume*, and *View Disks*. If the system has multipath-capable hardware, a *View Multipaths* button is also shown. 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.

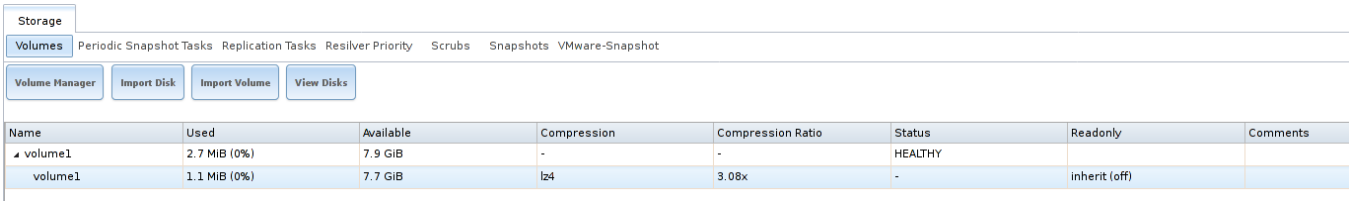

Fig. 7.11: Viewing Volumes

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

Note: When the system has *High Availability (HA)* (page 58) active, volumes cannot be exported or destroyed.

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

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

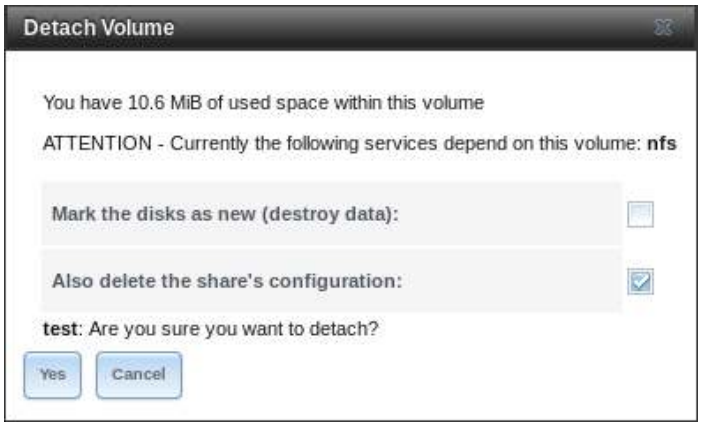

Fig. 7.12: Detach or Delete a Volume

Scrub Volume: scrubs and scheduling them are described in more detail in *Scrubs* (page 134). This button allows manually initiating a scrub. Scrubs are I/O intensive and can negatively impact performance. Avoid initiating a scrub when the system is busy.

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

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

Volume Status: as shown in the example in Figure 7.13, this screen shows the device name and status of each disk in the ZFS pool as well as any read, write, or checksum errors. It also indicates the status of the latest ZFS scrub. Clicking the entry for a device causes buttons to appear to edit the device options (shown in Figure 7.14), offline or online the device, or replace the device (as described in *Replacing a Failed Drive* (page 117)).

Upgrade: used to upgrade the pool to the latest *ZFS Feature Flags* (page 253). See the warnings in *Upgrading a ZFS Pool* (page 45) before selecting this option. This button does not appear when the pool is running the latest version of the feature flags.

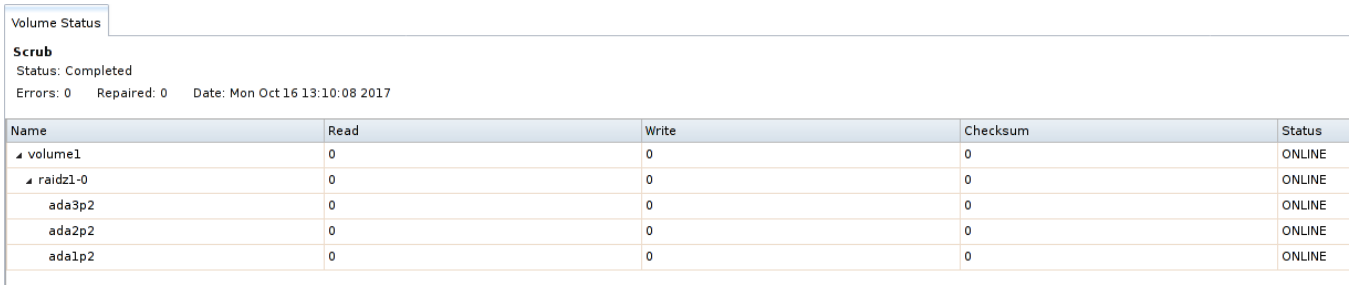

Fig. 7.13: Volume Status

Selecting a disk in *Volume Status* and clicking its *Edit Disk* button shows the screen in Figure 7.14. Table 7.6 summarizes the configurable options.

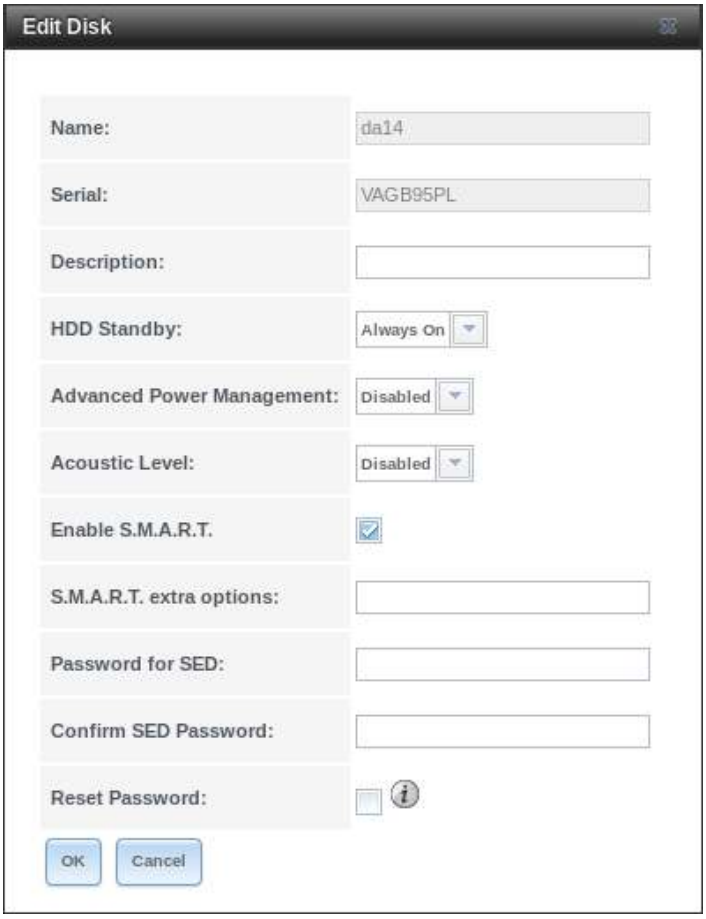

Fig. 7.14: Editing a Disk

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

Change Permissions: edit the dataset permissions as described in *Change Permissions* (page 100).

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use *Periodic Snapshot Tasks* (page 120).

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

Destroy Dataset: clicking the *Destroy Dataset* button causes the browser window to turn red to indicate that this is a destructive action. Clicking *Yes* proceeds with the deletion.

Edit Options: edit the volume properties described in Table 7.2.3. Note that it will not allow changing the dataset name.

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

Create zvol: create a child zvol within this dataset.

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

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

7.2.8.1 Managing Encrypted Volumes

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

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

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

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

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

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

7.2.8.2 Additional Controls for Encrypted Volumes

If the *Encryption* option is enabled during the creation of a pool, additional buttons appear in the entry for the volume in *Storage* → *Volumes*. An example is shown in Figure 7.15.

Fig. 7.15: Encryption Icons Associated with an Encrypted Volume

These additional encryption buttons are used to:

Create/Change Passphrase: set and confirm a passphrase associated with the GELI encryption key. The desired passphrase is entered and repeated for verification. A red warning is a reminder to *Remember to add a new recovery key as this action invalidates the previous recovery key*. Unlike a password, a passphrase can contain spaces and is typically a series of words. A good passphrase is easy to remember (like the line to a song or piece of literature) but hard to guess. **Remember this passphrase. An encrypted volume cannot be reimported without it.** In other words, if the passphrase is forgotten, the data on the volume can become inaccessible if it becomes necessary to

reimport the pool. Protect this passphrase, as anyone who knows it could reimport the encrypted volume, thwarting the reason for encrypting the disks in the first place.

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

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

Encrypted volumes with a passphrase display an additional lock button:

Fig. 7.17: Lock Button

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

Fig. 7.18: Unlock Button

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

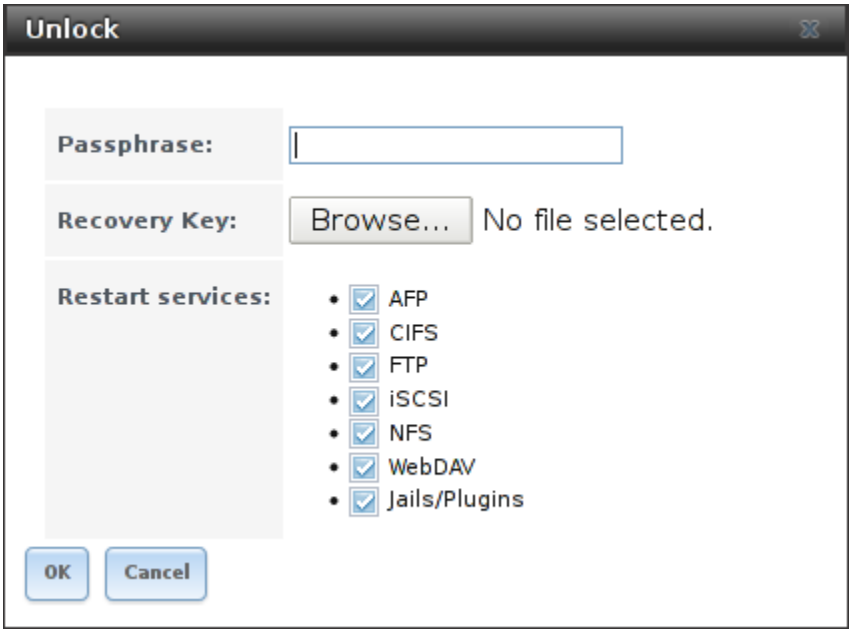

Fig. 7.19: Unlock Locked Volume

Unlock the volume by entering a passphrase *or* using the *Browse* button to load the recovery key. Only the passphrase is used when both a passphrase and a recovery key are entered. The services listed in *Restart Services* will restart when the pool is unlocked. This allows them to see the new volume and share or access data on it. Individual services can be prevented from restarting by deselecting 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 58) (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.

7.2.9 View Multipaths

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

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

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

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

7.2.10 Replacing a Failed Drive

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

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

Note: If the volume is encrypted with GELI, refer to *Replacing an Encrypted Drive* (page 119) before proceeding.

Before physically removing the failed device, go to *Storage* → *Volumes*. Select the volume 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 entry, then its *Offline* button to change that disk status to OFFLINE. This step is needed to properly remove the device from the ZFS pool and to prevent swap issues. Click the disk *Offline* button and pull the disk. If there is no *Offline* button but only a *Replace* button, the disk is already offlined and this step can be skipped.

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

2. After the disk has been replaced and is showing as OFFLINE, click the disk again and then click its *Replace* button. Select the replacement disk from the drop-down menu and click the *Replace Disk* button. After clicking the *Replace Disk* button, the ZFS pool begins resilvering.
3. After the drive replacement process is complete, re-add the replaced disk in the *S.M.A.R.T. Tests* (page 77) screen.

In the example shown in Figure 7.20, a failed disk is being replaced by disk *ada5* in the volume named `volume1`.

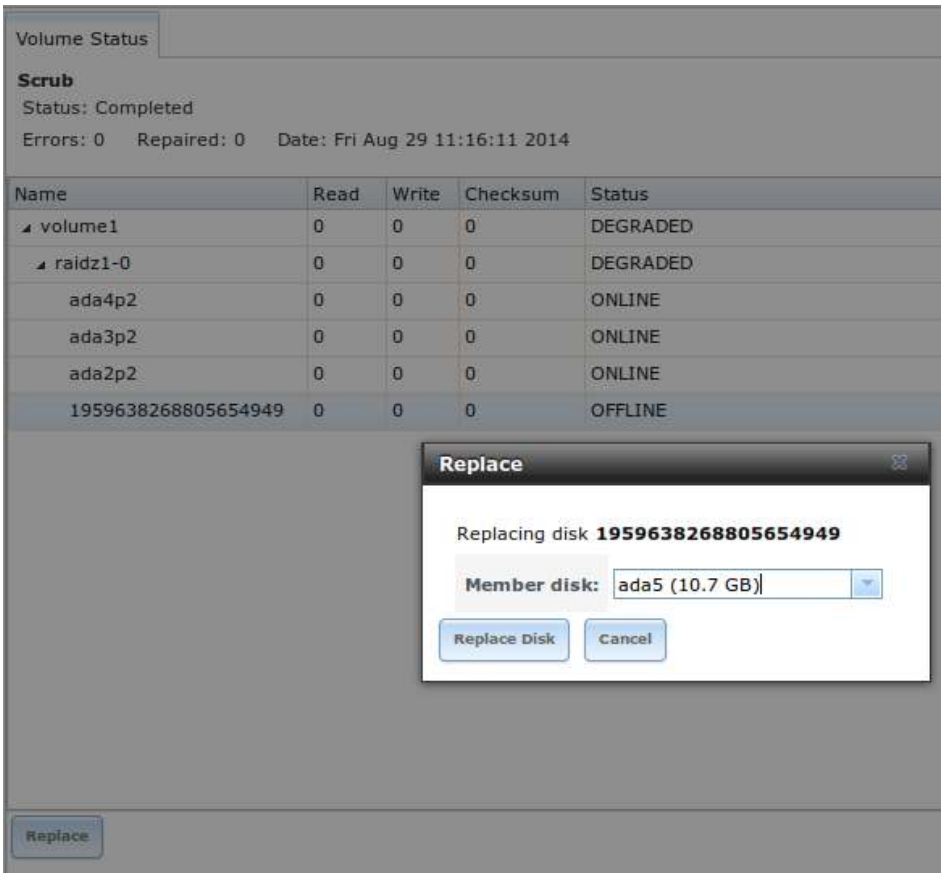

Fig. 7.20: Replacing a Failed Disk

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

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

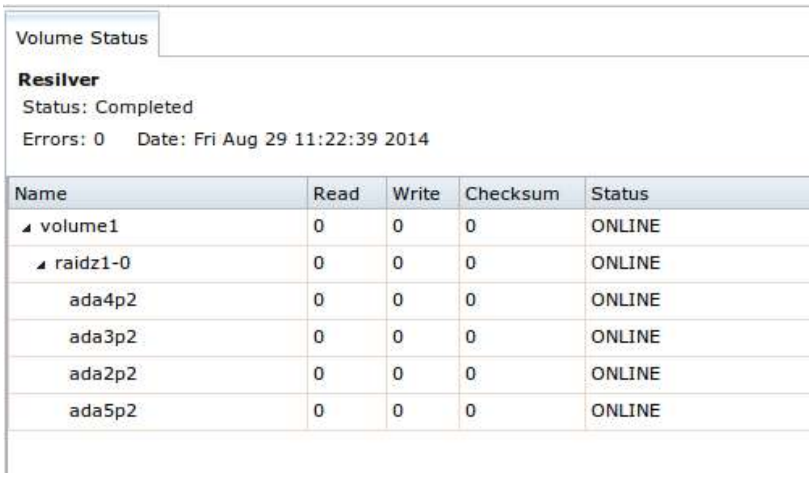

Fig. 7.21: Disk Replacement is Complete

7.2.10.1 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 97) **before** attempting to replace the failed drive. Then, follow the steps 1 and 2 as described above. During step 3, a prompt will appear to input and confirm the passphrase for the pool. Enter this information then click *Replace Disk*.

Wait until resilvering is complete before *restoring the encryption keys to the pool* (page 114). **Restore the encryption keys before the next reboot or access to the pool will be permanently lost**.

Warning: Access to the pool will be permanently lost unless the encryption keys are restored to the pool before the next system reboot!

1. 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.

7.2.10.2 Removing a Log or Cache Device

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

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

7.2.11 Replacing Drives to Grow a ZFS Pool

The recommended method for expanding the size of a ZFS pool is to pre-plan the number of disks in a vdev and to stripe additional vdevs from *Volumes* (page 94) as additional capacity is needed.

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

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

Note: A volume that is configured as a stripe (https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_0) can only be increased by following the steps in *Extending a ZFS Volume* (page 99).

1. Connect the new, larger disk to the unused disk port or drive bay.
2. Go to *Storage* → *Volumes*.
3. Select the volume and click the *Volume Status* button.
4. Select one of the old, smaller disks in the volume. Click the *Replace* button. Choose the new disk as the replacement.

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

If a unused disk port or bay is not available, a drive can be replaced with a larger one as shown in *Replacing a Failed Drive* (page 117). This process is slow and places the system in a degraded state. Since a failure at this point could be disastrous, **do not attempt this method unless the system has a reliable backup.** Replace one drive at a time and wait for the resilver process to complete on the replaced drive before replacing the next drive. After all the drives are replaced and the final resilver completes, the added space appears in the volume.

7.2.12 Adding Spares

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

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

Add a spare by going to *Storage* → *Volume Manager*. Select the volume to extend from the *Volume to extend* drop-down. Choose a disk from the list of *Available disks* and click + to add that disk to the volume. Select *spare* in the *Volume layout* drop down. Click *Extend Volume* to add the hot spare.

Danger: When adding a spare disk to an encrypted volume, **the passphrase and recovery key are reset**. Click *Download Key* to download the new recovery key. To create a new passphrase, click *Create Passphrase*.

7.3 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 MiB 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 122)). 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 95).

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

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

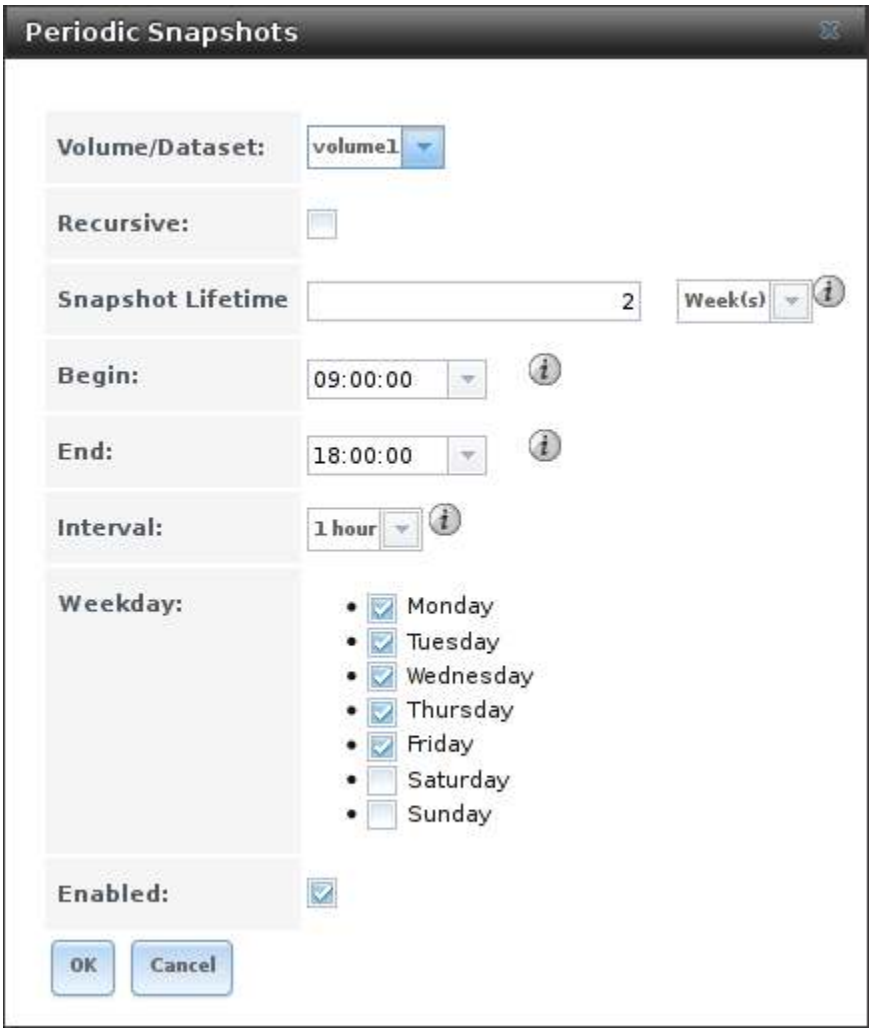

Fig. 7.22: Creating a Periodic Snapshot

Table 7.7: Options When Creating a Periodic Snapshot

Setting	Value	Description
Volume/Dataset	drop-down menu	Select an existing ZFS volume, dataset, or zvol.
Recursive	checkbox	Set to take separate snapshots of the volume or dataset and each of its child datasets. Unset to take a single snapshot of only the specified volume or dataset.
Snapshot Lifetime	integer and drop-down menu	Define a length of time to retain the snapshot on this system. After the time expires, the snapshot is removed. Snapshots replicated to other systems are not affected.
Begin	drop-down menu	Choose the hour and minute when the system can begin taking snapshots.
End	drop-down menu	Choose the hour and minute when the system will stop taking snapshots.

Continued on next page

Table 7.7 – continued from previous page

Setting	Value	Description
Interval	drop-down menu	Define how often the system takes snapshots between the *Begin* and *End* times.
Weekday	checkboxes	Choose the days of the week to take snapshots.
Enabled	checkbox	Unset to disable this task without deleting it.

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

Click the *OK* button to save the task. Entries for each task are shown in *View Periodic Snapshot Tasks*. Click an entry to display *Edit* and *Delete* buttons for it.

7.4 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. When a *periodic snapshot* (page 120) of the selected dataset occurs, the replication task copies the data to the destination system.

When snapshots are automatically created on the source computer, they are replicated to the destination computer. First-time replication tasks can take a long time to complete as the entire snapshot must be copied to the destination system. Replicated data is not visible on the receiving system until the replication task completes. Later replications only send the snapshot changes to the destination system. Interrupting a running replication requires the replication task to restart from the beginning.

The target dataset on the receiving system is automatically created in read-only mode to protect the data. To mount or browse the data on the receiving system, create a clone of the snapshot and use the clone. Clones are created in read/write mode, making it possible to browse or mount them. See *Snapshots* (page 137) for more information on creating clones.

7.4.1 Examples: Common Configuration

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

7.4.1.1 *Alpha* (Source)

Alpha is the source computer with the data to be replicated. It is at IP address *10.0.0.102*. A *volume* (page 94) named *alphavol* has already been created, and a *dataset* (page 102) named *alphadata* has been created on that volume. This dataset contains the files which will be snapshotted and replicated onto *Beta*.

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

Create a periodic snapshot of the source dataset by selecting *Storage* → *Periodic Snapshot Tasks*. Click the *alphavol/alphadata* dataset to highlight it. Create a *periodic snapshot* (page 120) of it by clicking *Periodic Snapshot Tasks*, then *Add Periodic Snapshot* as shown in Figure 7.23.

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

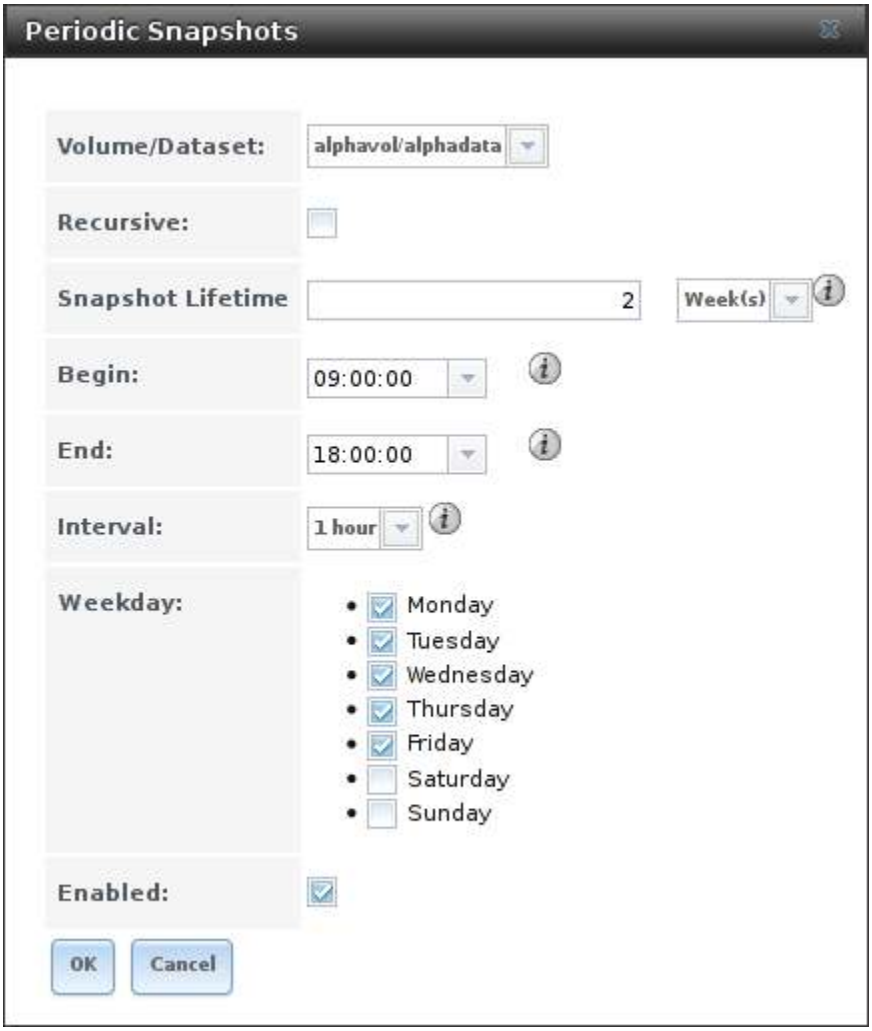

Fig. 7.23: Create a Periodic Snapshot for Replication

7.4.1.2 *Beta* (Destination)

Beta is the destination computer where the replicated data will be copied. It is at IP address *10.0.0.118*. A *volume* (page 94) named *betavol* has already been created.

Snapshots are transferred with *SSH* (page 225). 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*.

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

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

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

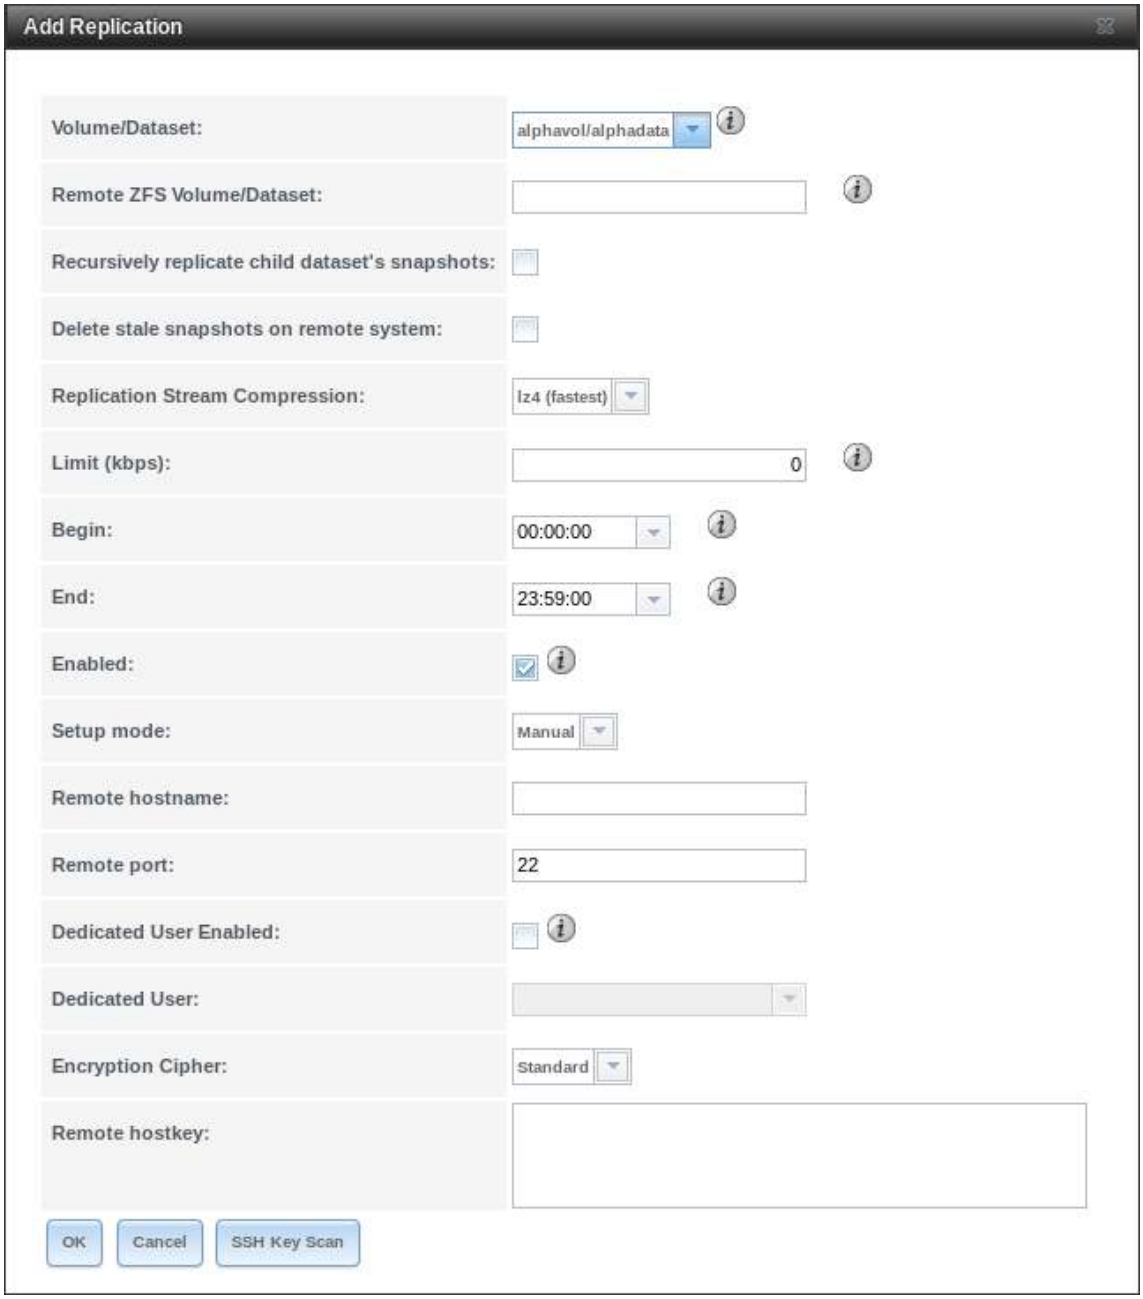

Fig. 7.24: Add Replication Dialog, Semi-Automatic

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

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

Fig. 7.25: Temporary Authentication Token on Destination

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

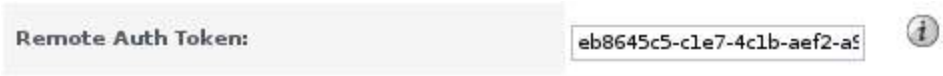

Fig. 7.26: Temporary Authentication Token Pasted to Source

Finally, click the *OK* button to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See *Limiting Replication Times* (page 131) 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.

7.4.3 Example: TrueNAS® to TrueNAS® Dedicated User Replication

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

In this example, the periodic snapshot task has not been created yet. If the periodic snapshot shown in the *example configuration* (page 122) has already been created, go to *Storage* → *Periodic Snapshot Tasks*, click on the task to select it, and click *Delete* to remove it before continuing.

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

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

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

The replication user must be given permissions to the destination dataset. Still on *Beta*, open a *Shell* (page 244) and enter this command:

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

The destination dataset must also be set to read-only. Enter this command in the *Shell* (page 244):

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

Close the *Shell* (page 244) by typing `exit` and pressing `Enter`.

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

Back on *Alpha*, create a periodic snapshot of the source dataset by selecting *Storage* → *Periodic Snapshot Tasks*. Click the *alphavol/alphadata* dataset to highlight it. Create a *periodic snapshot* (page 120) of it by clicking *Periodic Snapshot Tasks*, then *Add Periodic Snapshot* as shown in Figure 7.23.

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

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

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

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

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

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

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

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

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

Replication will begin when the periodic snapshot task runs.

Additional replications can use the same dedicated user that has already been set up. The permissions and read only settings made through the *Shell* (page 244) must be set on each new destination dataset.

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

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

7.4.4.1 Encryption Keys

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

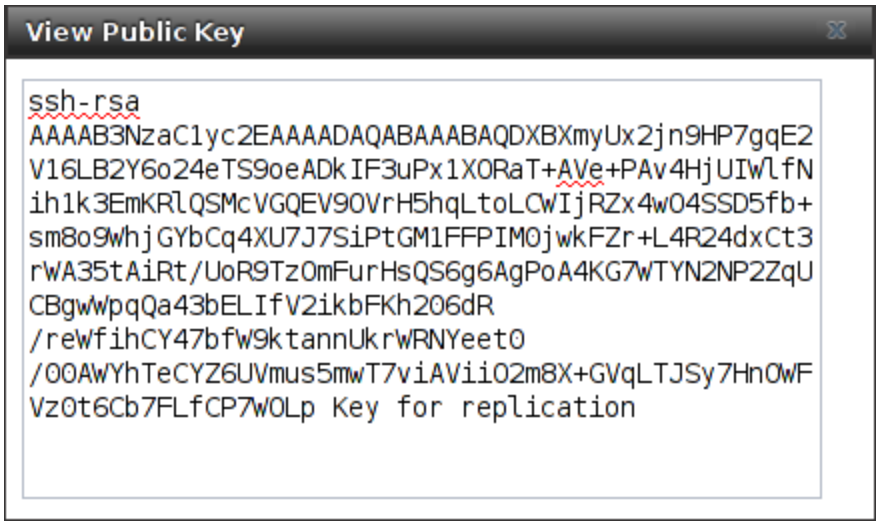

Fig. 7.27: Copy the Replication Key

On *Beta*, select *Account* → *Users* → *View Users*. Click the *root* account to select it, then click *Modify User*. Paste the copied key into the *SSH Public Key* field and click *OK* as shown in Figure 7.28.

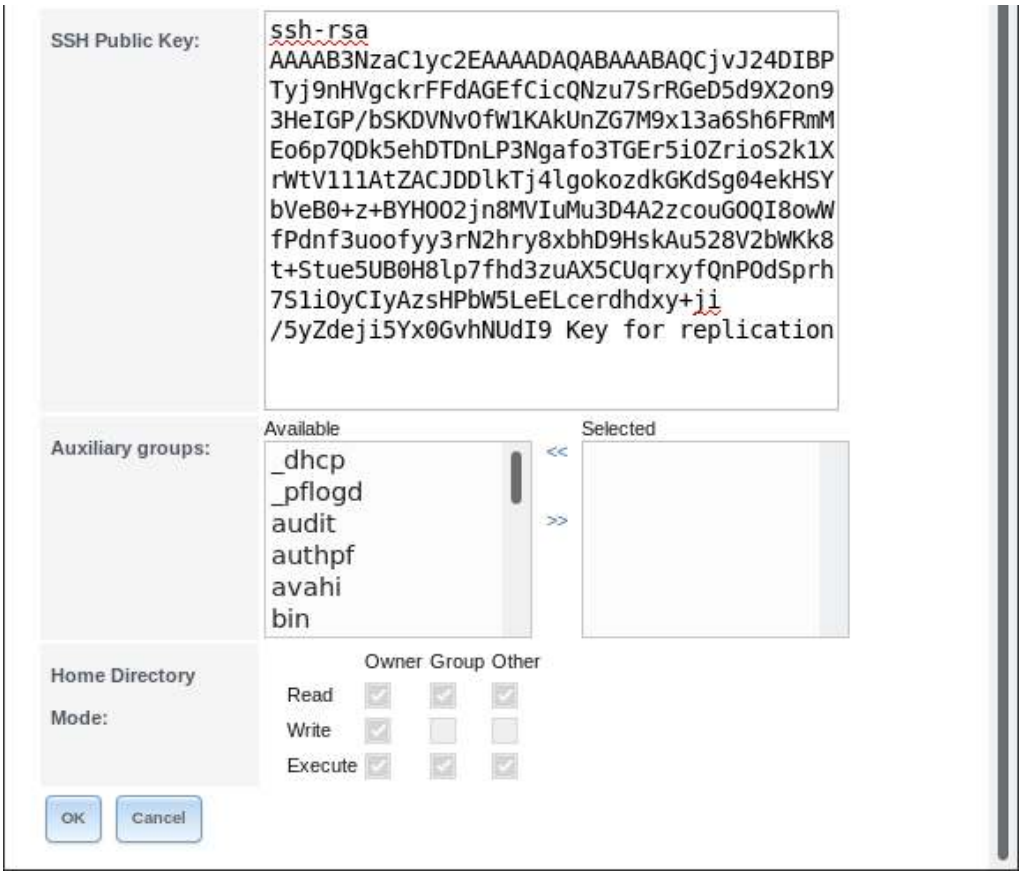

Fig. 7.28: Paste the Replication Key

Back on *Alpha*, create the replication task by clicking *Replication Tasks* and *Add Replication*. *alphavol/alphadata* is selected as the dataset to replicate. The destination volume is *betavol*. The *alphadata* dataset and snapshots are replicated there. The IP address of *Beta* is entered in the *Remote hostname* field as shown in Figure 7.29. A hostname can be entered here if local DNS resolves for that hostname.

Click the *SSH Key Scan* button to retrieve the SSH host keys from *Beta* and fill the *Remote hostkey* field. Finally, click *OK* to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See *Limiting Replication Times* (page 131) for information about restricting when replication is allowed to run.

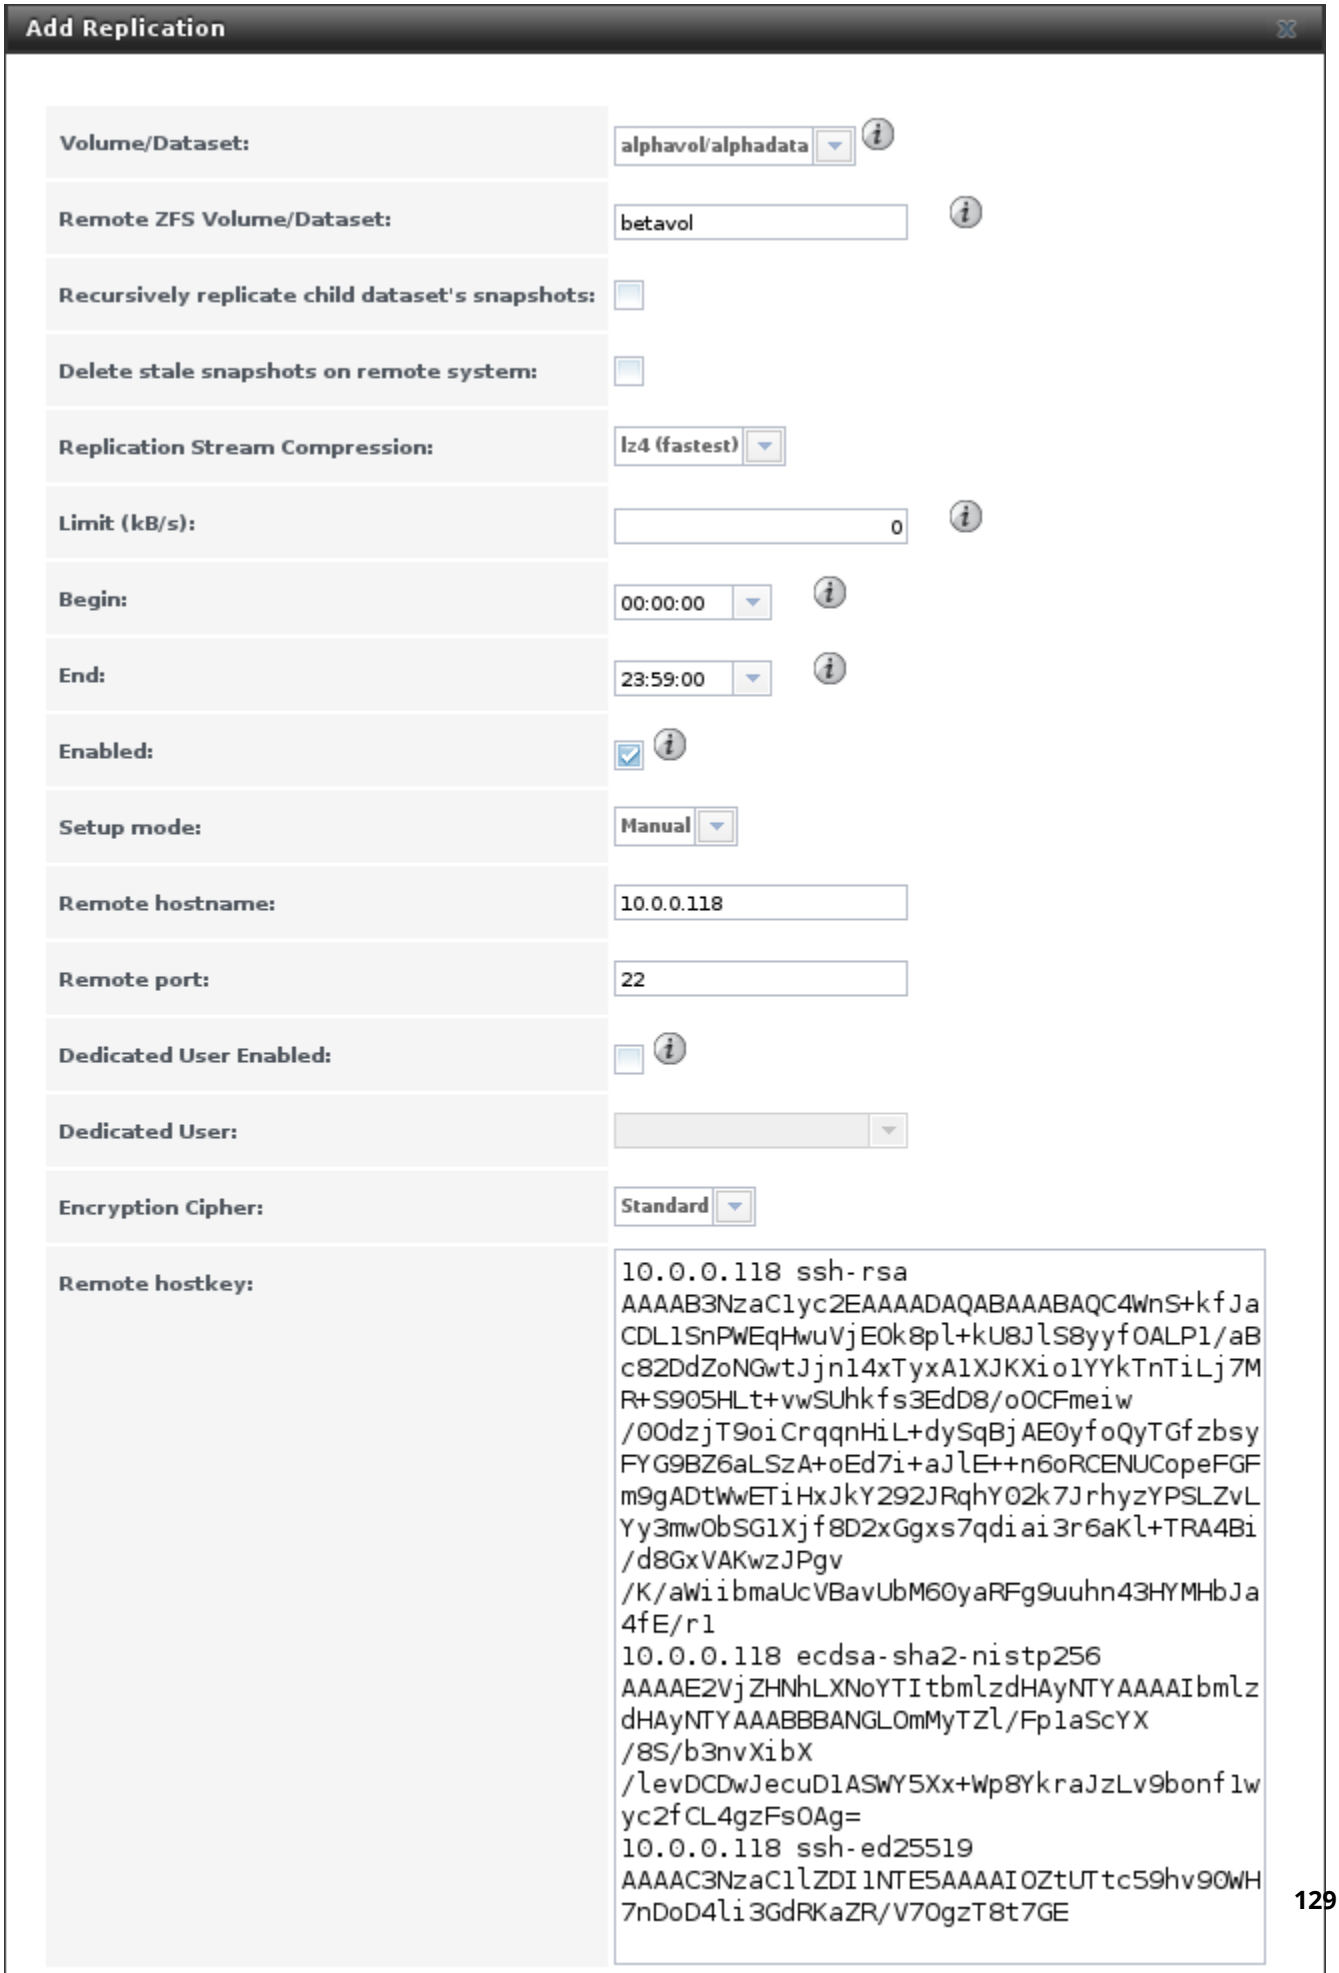
Add Replication
Volume/Dataset:
alphavol/alphadata
Remote ZFS Volume/Dataset:
betavol
Recursively replicate child dataset's snapshots:
Delete stale snapshots on remote system:
Replication Stream Compression:
lz4 (fastest)
Limit (kB/s):
0
Begin:
00:00:00
End:
23:59:00
Enabled:
Setup mode:
Manual
Remote hostname:
10.0.0.118
Remote port:
22
Dedicated User Enabled:
Dedicated User:
Encryption Cipher:
Standard
Remote hostkey:
10.0.0.118 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQC4WnS+kfJa
CDL1SnPWEqHwuVjE0k8pl+kU8JlS8yyf0ALP1/aB
c82DdZoNGwtJjn14xTyxA1XJKXio1YYkTnTiLj7M
R+S905HLt+vwSUhkfs3EdD8/o0CFmeiw
/O0dzjT9oiCrqqnHiL+dySqBjAE0yfoQyTGfzbsy
FYG9BZ6aLSzA+oEd7i+aJlE++n6oRCENUCopeFGF
m9gADtWwETiHxJkY292JRqhY02k7JrhyzYPSLZvL
Yy3mwObSG1Xjf8D2xGgxs7qdiai3r6aKl+TRA4Bi
/d8GxVAKwzJPgv
/K/aWiibmaUcVBavUbM60yaRFg9uuhn43HYMHbJa
4fE/r1
10.0.0.118 ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlz
dHAyNTYAAABBBANGLOmMyTZl/Fp1aScYX
/8S/b3nvXibX
/levDCDwJecuD1ASWY5Xx+Wp8YkraJzLv9bonf1w
yc2fCL4gzFs0Ag=
10.0.0.118 ssh-ed25519
AAAAC3NzaC1lZDI1NTE5AAAAI0ZtUTtc59hv90WH
7nDoD4li3GdRKaZR/V70gzT8t7GE

7.4.5 Replication Options

Table 7.8 describes the options in the replication task dialog.

Table 7.8: Replication Task Options

Setting	Value	Description
Volume/Dataset	drop-down menu	On the source computer with snapshots to replicate, choose an existing ZFS pool or dataset with an active periodic snapshot task.
Remote ZFS Volume/Dataset	string	Enter the ZFS volume or dataset on the remote or destination computer which will store the snapshots. Example: poolname/datasetname, not the mount point or filesystem path.
Recursively replicate child dataset snapshots	checkbox	When enabled, include snapshots of child datasets from the primary dataset.
Delete stale snapshots	checkbox	Set to delete previous snapshots from the remote or destination system which are no longer present on the source computer.
Replication Stream Compression	drop-down menu	Choices are *lz4 (fastest)*, *pigz (all rounder)*, *plzip (best compression)*, or *Off* (no compression). Selecting a compression algorithm can reduce the size of the data being replicated.
Limit (kbps)	integer	Limit replication speed to the specified value in kilobits/second. Default of *0* is unlimited.
Begin	drop-down menu	Define a time to start the replication task.
End	drop-down menu	Define the point in time by which replication must start. A started replication task conitinues until it is finished.
Enabled	checkbox	Deselect to disable the scheduled replication task without deleting it.
Setup mode	drop-down menu	Choose the configuration mode for the remote. Choices are *Manual* or *Semi-automatic*. Note semi-automatic only works with remote version 9.10.2 or later.
Remote hostname	string	Enter the IP address or DNS name of remote system to receive the replication data.
Remote port	string	Enter the port number used by the SSH server on the remote or destination computer.
Dedicated User Enabled	checkbox	Select the user account other than root to be used for replication.
Dedicated User	drop-down menu	Only available if *Dedicated User Enabled* is enabled. Select the user account to be used for replication.
Encryption Cipher	drop-down menu	*Standard*, *Fast*, or *Disabled*.
Remote hostkey	string	Click *SSH Key Scan* to retrieve the public host key of the remote or destination computer and populate this field with that key.

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

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

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

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

Selecting *Storage* → *Replication Tasks* displays Figure 7.30, the list of replication tasks. The *Last snapshot sent to remote side* column shows the name of the last snapshot that was successfully replicated, and *Status* shows the current status of each replication task. The display is updated every five seconds, always showing the latest status.

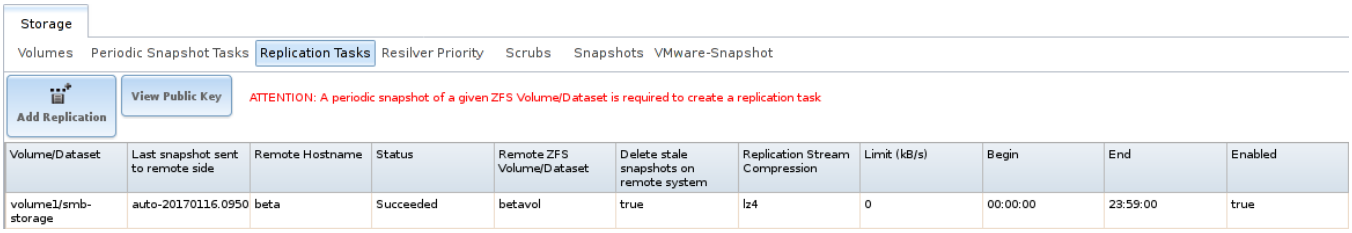

Fig. 7.30: Replication Task List

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

7.4.6 Replication Encryption

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

7.4.7 Limiting Replication Times

The *Begin* and *End* times in a replication task make it possible to restrict when replication is allowed. These times can be set to only allow replication after business hours, or at other times when disk or network activity will not slow down other operations like snapshots or *Scrubs* (page 134). 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.

7.4.8 Replication Topologies and Scenarios

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

7.4.8.1 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.

7.4.8.2 Tiered Replication

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

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

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

7.4.8.3 N-way Replication

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

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

7.4.8.4 Disaster Recovery

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

Recovering data onto a replacement computer 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.

7.4.9 Troubleshooting Replication

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

7.4.9.1 SSH

SSH (page 225) must be able to connect from the source system to the destination system with an encryption key. This can be tested from *Shell* (page 244) by making an *SSH* (page 225) 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 244) on the source machine (*Alpha*), then enter this command:

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

On the first connection, the system might say

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

Verify that this is the correct destination computer from the preceding information on the screen and type `yes`. At this point, an *SSH* (page 225) shell connection is open to the destination system, *Beta*.

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

7.4.9.2 Compression

Matching compression and decompression programs must be available on both the source and destination computers. This is not a problem when both computers are running TrueNAS®, but other operating systems might not have *lz4*, *pigz*, or *plzip* compression programs installed by default. An easy way to diagnose the problem is to set *Replication Stream Compression* to *Off*. If the replication runs, select the preferred compression method and check `/var/log/debug.log` on the TrueNAS® system for errors.

7.4.9.3 Manual Testing

On *Alpha*, the source computer, the `/var/log/messages` file can also show helpful messages to locate the problem.

On the source computer, *Alpha*, open a *Shell* (page 244) 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 244) 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 244) with `zfs list -t snapshot` or by going to *Storage* → *Snapshots*.

Error messages here can indicate any remaining problems.

7.5 Resilver Priority

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

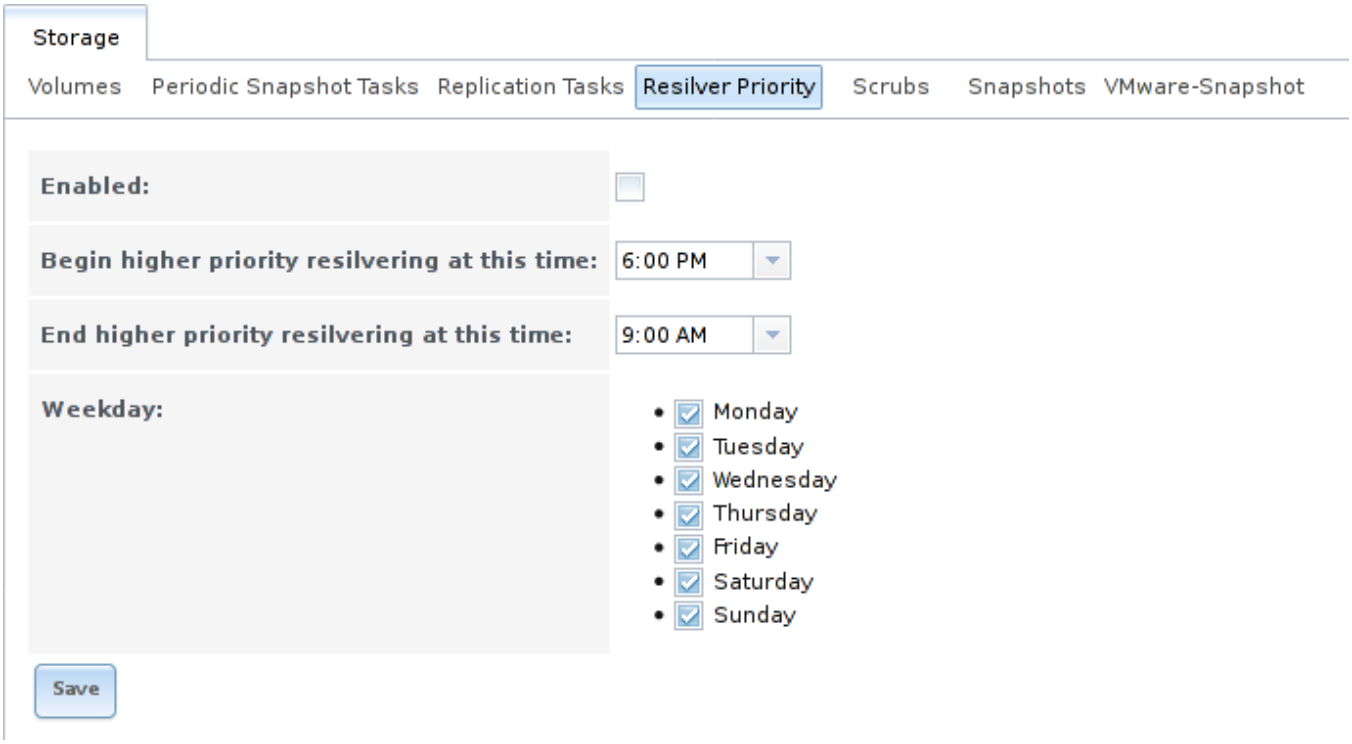

Fig. 7.31: Resilver Priority

Table 7.9: Resilver Priority Options

Setting	Value	Description
Enabled	checkbox	Set to enable higher-priority resilvering.
Begin higher priority resilvering at this time	drop-down	Start time to begin higher-priority resilvering.
End higher priority resilvering at this time	drop-down	End time to begin higher-priority resilvering.
Weekday	checkboxes	Use higher-priority resilvering on these days of the week.

7.6 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 77) 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 77) of *Type Long Self-Test* to run once or twice a month.

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

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

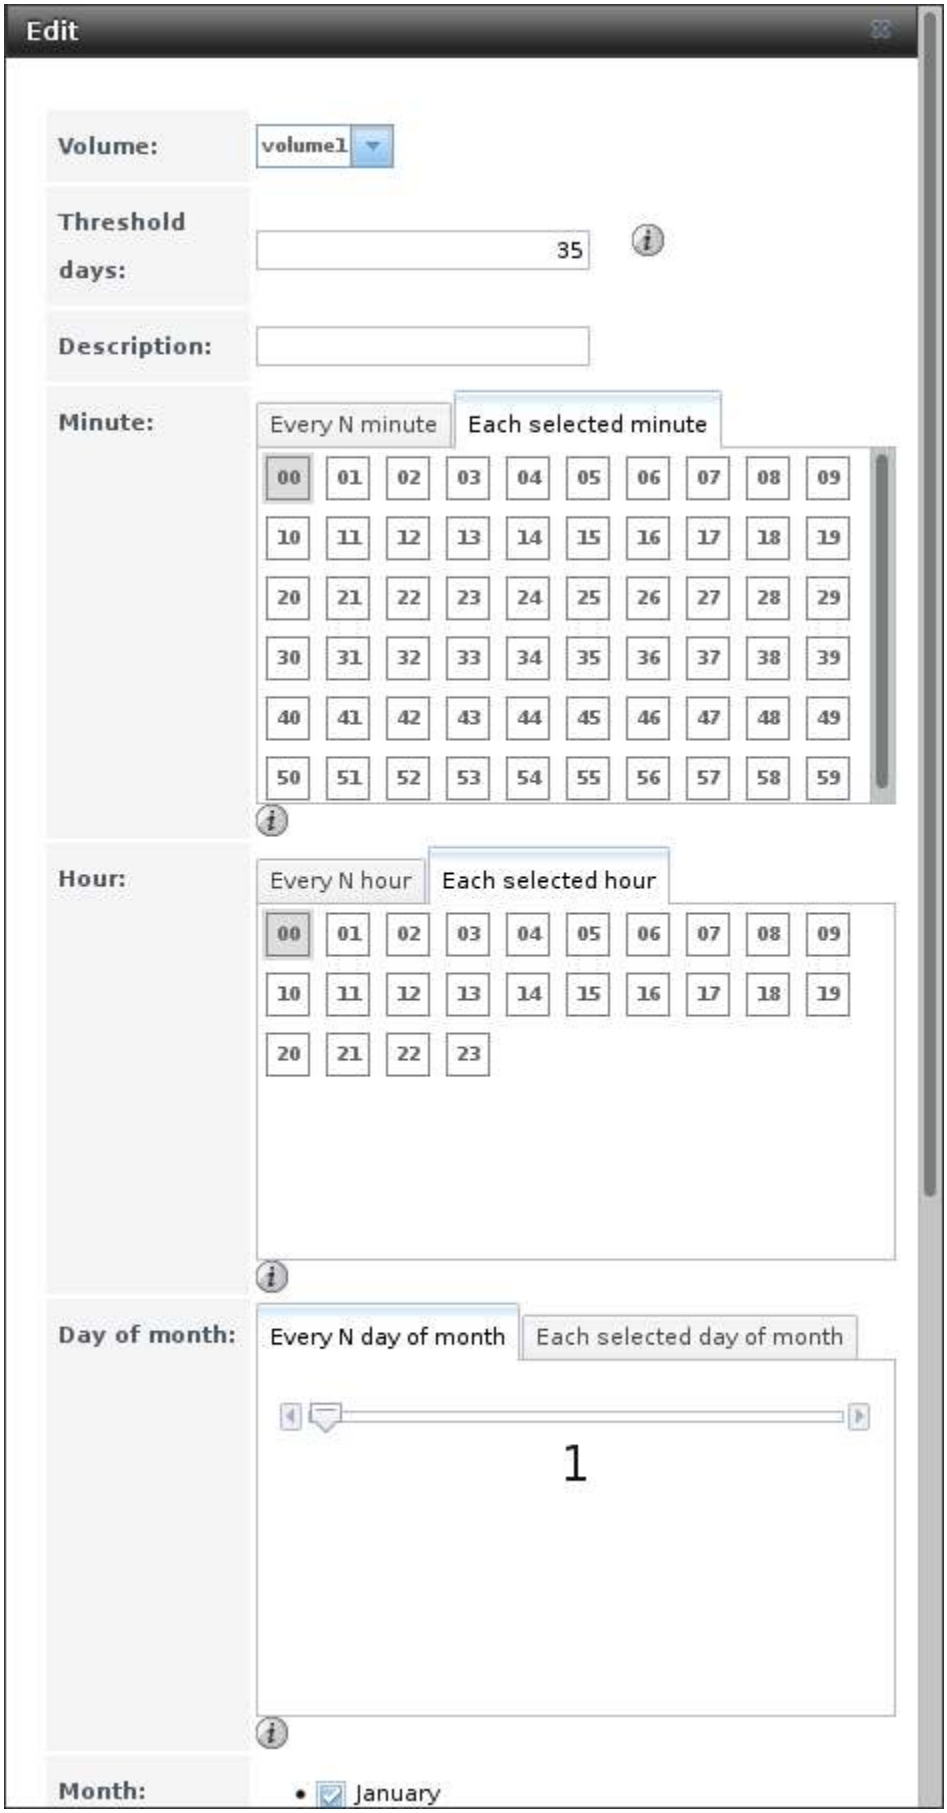

Fig. 7.32: Viewing Volume Default Scrub Settings

Table 7.10: ZFS Scrub Options

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

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

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

7.7 Snapshots

Snapshots are scheduled using *Storage* → *Periodic Snapshot Tasks*. To view and manage the listing of created snapshots, use *Storage* → *Snapshots*. An example listing is shown in Figure 7.33.

Note: If snapshots do not appear, check that the current time configured in *Periodic Snapshot Tasks* (page 120) 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 244).

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 quota and reservation. The space used does not include the dataset reservation, but does take into account the reservations of any descendent datasets. The amount of space that a dataset consumes from its parent, as well as the amount of space freed if this dataset is recursively deleted, is the greater of its space used and its reservation. When a snapshot is created, the space is initially shared between the snapshot and the filesystem, and possibly with previous snapshots. As the filesystem changes, space that was previously shared becomes unique to the snapshot, and is counted in the used space of the snapshot. 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 244).

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

Snapshots have icons on the right side for several actions.

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

Destroy Snapshot a pop-up message asks for confirmation. Child clones must be deleted before their parent snapshot can be deleted. While creating a snapshot is instantaneous, deleting a snapshot can be I/O intensive and can take a long time, especially when deduplication is enabled. 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 used, it can be freed.

The most recent snapshot also has a **Rollback Snapshot** icon. Clicking the icon asks for confirmation before rolling back to the chosen 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 option in the left column of the first snapshot, then press and hold `Shift` and click on the option 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 175). 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 the view criteria.
- Click the *Filter* button to save the filter and exit the define filter screen. Alternately, click the + button to add another filter.

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

Warning: A snapshot and any files it contains will not be accessible or searchable if the mount path of the snapshot is longer than 88 ascii characters. The data within the snapshot will be safe, and the snapshot will become accessible again when the mount path is shortened. For details of this limitation, and how to shorten a long mount path, see *Path and Name Lengths* (page 9).

7.7.1 Browsing a snapshot collection

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

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

The effect will be that any user who can access the dataset contents, will also be able to view the list of snapshots by navigating to the `.zfs` directory of the dataset, and to browse and search any files they have permission to access throughout the entire snapshot collection of the dataset. A user's ability to view files within a snapshot will be limited by any permissions or ACLs set on the files when the snapshot was taken. Snapshots are fixed as "read-only", so this access does not permit the user to change any files in the snapshots, or to modify or delete any snapshot, even if they had write permission at the time when the snapshot was taken.

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

7.8 VMware-Snapshot

Storage → *VMware-Snapshots* is used to coordinate ZFS snapshots when using TrueNAS® as a VMware datastore. When a ZFS snapshot is created, TrueNAS® automatically snapshots any running VMware virtual machines before

taking a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore. Virtual machines **must be powered on** for TrueNAS® snapshots to be copied to VMware. The temporary VMware snapshots are then deleted on the VMware side but still exist in the ZFS snapshot and can be used as stable resurrection points in that snapshot. These coordinated snapshots are listed in *Snapshots* (page 137).

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

Fig. 7.34: Adding a VMware Snapshot

Table 7.11: VMware Snapshot Options

Setting	Value	Description
Hostname	string	Enter the IP address or hostname of VMware host. When clustering, this is the vCenter server for the cluster.
Username	string	Enter the username on the VMware host with permission to snapshot virtual machines.
Password	string	Enter the password associated with *Username*.
ZFS Filesystem	drop-down menu	Select the filesystem to snapshot.
Datastore	drop-down menu	Enter the *Hostname*, *Username*, and *Password*. Click *Fetch Datastores* to populate the menu and select the datastore with which to synchronize.

CHAPTER

EIGHT

DIRECTORY SERVICES

TrueNAS® supports integration with these directory services:

- *Active Directory* (page 141) (for Windows 2000 and higher networks)
- *LDAP* (page 147)
- *NIS* (page 150)

It also supports *Kerberos Realms* (page 151), *Kerberos Keytabs* (page 151), and the ability to add more parameters to *Kerberos Settings* (page 152).

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

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/Setting_up_Samba_as_an_Active_Directory_Domain_Controller#Provisioning_a_Samba_Active
Since AD provides authentication and authorization services for the users in a network, it is not necessary to recreate these user accounts on the TrueNAS® system. Instead, configure the Active Directory service so that it can import the account information and imported users can be authorized to access the SMB shares on the TrueNAS® system.

Many changes and improvements have been made to Active Directory support within TrueNAS®. It is strongly recommended to update the system to the latest TrueNAS® 11.2 before attempting Active Directory integration.

Ensure name resolution is properly configured before configuring the Active Directory service. `ping` the domain name of the Active Directory domain controller from *Shell* (page 244) on the TrueNAS® system. If the `ping` fails, check the DNS server and default gateway settings in *Network* → *Global Configuration* on the TrueNAS® system.

By default, *Allow DNS updates* in the *Active Directory options* (page 142) is enabled. This adds TrueNAS® *SMB 'Bind IP Addresses'* (page 221) DNS records to the Active Directory DNS when the domain is joined. Disabling *Allow DNS updates* means that the Active Directory DNS records must be updated manually.

Active Directory relies on Kerberos, a time-sensitive protocol. The time on the TrueNAS® system and the Active Directory Domain Controller cannot be out of sync by more than five minutes in a default Active Directory environment.

To ensure both systems are set to the same time:

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

Using a TrueNAS® system as an AD server and connecting to it with a TrueNAS® client requires additional configuration. On the AD server, go to *System* → *CAs* and create a new internal or intermediate *Certificate Authority (CA)* (page 49). Highlight the created CA and click *Export Certificate* and *Export Private Key* to save these values.

On the client web interface, select *Directory Service* → *Active Directory* → *Advanced Mode*. Set *Encryption Mode* to *TLS* and *SASL wrapping* to *sign*. Go to *System* → *CAs* and click *Import CA*. Create a unique *Identifier* and paste the AD server CA certificate and private keys in those fields. Click *OK* and continue configuring AD.

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 *Advanced Mode* or configure the system to always display these settings by checking *Show advanced fields by default* in *System* → *Advanced*.

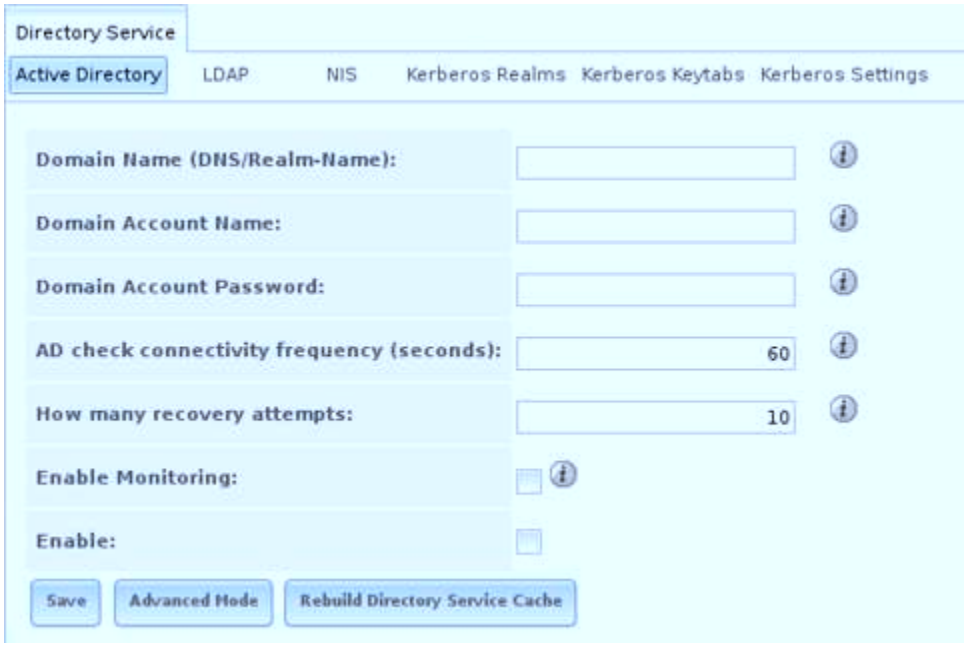

Fig. 8.1: Configuring Active Directory

Table 8.1: Active Directory Configuration Options

Setting	Value	Advanced Mode	Description
Domain Name (DNS/Realm-Name)	string		Name of Active Directory domain (*example.com*) or child domain (*sales.example.com*). This setting is mandatory and the GUI will refuse to save the settings if the domain controller for the specified domain cannot be found.
Domain Account Name	string		Name of the Active Directory administrator account. This setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this account name.
Domain Account Password	string		Password for the Active Directory administrator account. This setting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this password.
AD check connectivity frequency (seconds)	integer		How often to verify that Active Directory services are active.
How many recovery attempts	integer		Number of times to attempt reconnecting to the Active Directory server. Tries forever when set to *0*.
Enable Monitoring	checkbox		Restart Active Directory automatically if the service is disconnected. Setting this prevents configuring the *Domain Controller* (page 202) service.

Continued on next page

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Encryption Mode	drop-down	✓	Choices are *Off*, *SSL (LDAPS protocol port 636)*, or *TLS (LDAP protocol port 389)*. See http://info.ssl.com/article.aspx?id=10241 and https://hpbn.co/transport-layer-security-tls/ for more information about SSL and TLS.
Certificate	drop-down menu	✓	Select the Active Directory server certificate if SSL connections are used. If a certificate does not exist, create a *Certificate Authority* (page 49), then create a certificate on the Active Directory server. Import the certificate to the TrueNAS® system using the *Certificates* (page 51) menu. To clear a saved certificate, choose the blank entry and click *Save*.
Verbose logging	checkbox	✓	Set to log attempts to join the domain to `/var/log/messages`.
UNIX extensions	checkbox	✓	Deprecated. Use the System Security Services Daemon (SSSD) for retrieving RFC2307 (https://tools.ietf.org/html/rfc2307) extensions from an Active Directory domain. Use the *ad idmap backend* (page 144) to enable this feature.
Allow Trusted Domains	checkbox	✓	Do not set this unless the network has active domain/forest trusts (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc757352(v=ws.10)) and managing files on multiple domains is required. Setting this option generates more winbindd traffic and slows down filtering with user and group information. If enabled, also configuring the idmap ranges and a backend for each trusted domain in the environment is recommended.
Use Default Domain	checkbox	✓	Unset to prepend the domain name to the username. If *Allow Trusted Domains* is set and multiple domains use the same usernames, unset to prevent name collisions.
Allow DNS updates	checkbox	✓	Unset to disable Samba from doing DNS updates when joining a domain.
Disable Active Directory user/group cache	checkbox	✓	Disable caching AD users and groups. Setting this hides all AD users and groups from web interface drop-down menus and auto-completion suggestions, but manually entering names is still allowed. This can help when unable to bind to a domain with a large number of users or groups.
Site Name	string	✓	Auto-detected site name. Do not change this unless the detected site name is incorrect for the particular AD environment.
Domain Controller	string	✓	The server that manages user authentication and security as part of a Windows domain. Leave empty for TrueNAS® to use the DNS SRV records to automatically detect and connect to the domain controller. If the domain controller must be set manually, enter the server hostname or IP address.

Continued on next page

Table 8.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Global Catalog Server	string	✓	The global catalog server holds a full set of attributes for the domain in which it resides and a subset of attributes for all objects in the Microsoft Active Directory Forest. See the IBM Knowledge Center (https://www.ibm.com/support/knowledgecenter/en/SSEQTP_9.0.0/com. Leave empty for TrueNAS® to use the DNS SRV records to automatically detect and connect to the server. If the global catalog server must be entered manually, enter the server hostname or IP address.
Kerberos Realm	drop-down menu	✓	Select the realm created using the instructions in *Kerberos Realms* (page 151).
Kerberos Principal	drop-down menu	✓	Browse to the location of the keytab created using the instructions in *Kerberos Keytabs* (page 151).
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 *Edit* to configure the backend.
Windbind NSS Info	drop-down menu	✓	Defines the schema to use when querying AD for user/group info. *rfc2307* uses the RFC2307 schema included in Windows 2003 R2, *sfu20* is for Services For Unix 3.0 or 3.5, and *sfu* is for Services For Unix 2.0.
SASL wrapping	drop-down menu	✓	Defines how LDAP traffic is transmitted. Choices are *plain* (plain text), *sign* (signed only), or *seal* (signed and encrypted). Windows 2000 SP3 and newer can be configured to enforce signed LDAP connections.
Enable	checkbox		Activate the Active Directory service.
NetBIOS Name (This Node)	string	✓	Name for the computer object generated in AD. Limited to 15 characters. Automatically populated with the original hostname of the system. This **must** be different from the *Workgroup* name.
NetBIOS Name (Node A/B)	string	✓	Name for the computer object generated in AD. Limited to 15 characters. When using *Failover* (page 58), set a unique NetBIOS name for the standby node.
NetBIOS Alias	string	✓	Limited to 15 characters. When using *Failover* (page 58), this is the NetBIOS name that resolves to either node.

Table 8.2 summarizes the backends which are available in the *Idmap backend* drop-down menu. Each backend has its own man page (http://samba.org.ru/samba/docs/man/manpages/) which gives implementation details.

Changing idmap backends requires refreshing the `windbind` resolver cache by sending SIGHUP (signal hang up) to the parent `windbindd` process. To find this parent process, start an *SSH* (page 225) session with the TrueNAS® system and enter `service samba_server status`. To send the SIGHUP, enter `kill -HUP pid`, where *pid* is the parent process ID.

Table 8.2: ID Mapping Backends

Value	Description
ad	AD server uses RFC2307 or Services For Unix schema extensions. Mappings must be provided in advance by adding the *uidNumber* attributes for users and *gidNumber* attributes for groups in the AD.

Continued on next page

Table 8.2 – continued from previous page

Value	Description
autorid	Similar to *rid*, but automatically configures the range to be used for each domain, so there is no need to specify a specific range for each domain in the forest. The only needed configuration is the range of UID/GIDs to use for user/group mappings and an optional size for the ranges.
fruit	Generate IDs the way Apple Mac OS X does, so UID and GID can be identical on all TrueNAS® servers on the network. For use in *LDAP* (page 147) environments where Apple Open Directory is the authoritative LDAP server.
ldap	Stores and retrieves mapping tables in an LDAP directory service. Default for LDAP directory service.
nss	Provides a simple means of ensuring that the SID for a Unix user is reported as the one assigned to the corresponding domain user.
rfc2307	IDs for AD users stored as RFC2307 (https://tools.ietf.org/html/rfc2307) ldap schema extensions. This module can either look up the IDs in the AD LDAP servers or an external (non-AD) LDAP server.
rid	Default for AD. Requires an explicit idmap configuration for each domain, using disjoint ranges where a writeable default idmap range is to be defined, using a backend like *tdb* or *ldap*.
script	Stores mapping tables for clustered environments in `winbind_cache.tdb`.
tdb	Default backend used by `winbindd` for storing mapping tables.
tdb2	Substitute for *tdb* used by `winbindd` in clustered environments.

Rebuild Directory Service Cache immediately refreshes the web interface directory service cache. This occurs automatically once a day as a cron job.

If there are problems connecting to the realm, verify (https://support.microsoft.com/en-us/help/909264/naming-conventions-in-active-directory-for-computers-domains-sites-and) the settings do not include any disallowed characters. Active Directory does not allow `$` characters in Domain or NetBIOS names. The length of those names is also limited to 15 characters. The Administrator account password cannot contain the *$* character. If a `$` exists in the domain administrator password, kinit (https://www.freebsd.org/cgi/man.cgi?query=kinit) reports a `Password Incorrect` error and ldap_bind (https://www.freebsd.org/cgi/man.cgi?query=ldap_bind) reports 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 are imported to the TrueNAS® system are shown by typing commands in the TrueNAS® *Shell* (page 244):

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

In addition, `wbinfo -m` shows the domains and `wbinfo -t` tests the connection. When successful, `wbinfo -t` shows a message similar to:

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

To manually check that a specified user can authenticate, open the *Shell* (page 244) and enter `smbclient//127.0.0.1/SHARE -U DOMAIN\username`, where *SHARE* is the SMB share name, *DOMAIN* is the name of the trusted domain, and *username* is the user account for authentication testing.

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

Tip: Sometimes network users do not appear in the drop-down menu of a *Permissions* screen but the `wbinfo`

commands display these users. This is typically due to the TrueNAS® system taking longer than the default ten seconds to join Active Directory. Increase the value of *AD timeout* to 60 seconds.

To change a certificate, set the *Encryption Mode* to *Off* and unset *Enable* to disable AD. Click *Save*. Select the new *Certificate*, set the *Encryption Mode* as desired, set *Enable* to re-enable AD, and click *Save* to restart AD.

8.1.1 Troubleshooting Tips

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

Active Directory uses DNS to determine the location of the domain controllers and global catalog servers in the network. Use `host -t srv _ldap._tcp.domainname.com` to determine the SRV records of the network and change the weight and/or priority of the SRV record to reflect the fastest server. More information about SRV records can be found in the Microsoft article How DNS Support for Active Directory Works (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759550(v=ws.10)).

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

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. Ensure the password is still valid. Also, double-check the password on the AD account being used does not include any spaces, special symbols, and is not unusually long.

If the Windows server version is lower than 2008 R2, try creating a *Computer* entry on the Windows server's OU. When creating this entry, enter the TrueNAS® hostname in the *name* field. Make sure it is under 15 characters, the same name as the one set in the *Hostname* field in *Network* → *Global Configuration*, and the same *NetBIOS 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 Does not Join the Domain

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

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

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

```
sqlite3 /data/freenas-v1.db "UPDATE directoryservice_activedirectory SET ad_enable=1"
service ix-hostname start
service ix-kerberos start
service ix-kinit start
klist
service ix-pre-samba start
net -k -d 5 ads join [this generates verbose output of the domain join]
service samba_server restart
service ix-nsswitch start
service ix-pam start
service ix-cache start
```

Next, only run these two commands **if** *UNIX extensions* is set in *Advanced Mode* and a keytab has been uploaded using *Kerberos Keytabs* (page 151):

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

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

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

8.2 LDAP

TrueNAS® includes an OpenLDAP (http://www.openldap.org/) client for accessing information from an LDAP server. An LDAP server provides directory services for finding network resources such as users and their associated permissions. Examples of LDAP servers include Microsoft Server (2000 and newer), Mac OS X Server, Novell eDirectory, and OpenLDAP running on a BSD or Linux system. If an LDAP server is running on the network, configure the TrueNAS® LDAP service so network users can authenticate to the LDAP server and have authorized access to the data stored on the TrueNAS® system.

Tip: TrueNAS® can also integrate with the Apple Open Directory (https://manuals.info.apple.com/MANUALS/0/MA954/en_US/Op LDAP-compatible directory service. 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/).

LDAP authentication for SMB shares is disabled unless the LDAP directory has been configured for and populated with Samba attributes. The most popular script for performing this task is smbldap-tools (https://wiki.samba.org/index.php/4.1_smbldap-tools). In addition, the LDAP server must support SSL/TLS and the certificate for the LDAP server CA must be imported with *System* → *CAs* → *Import CA*. Note that non-CA certificates are not supported at this time.

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

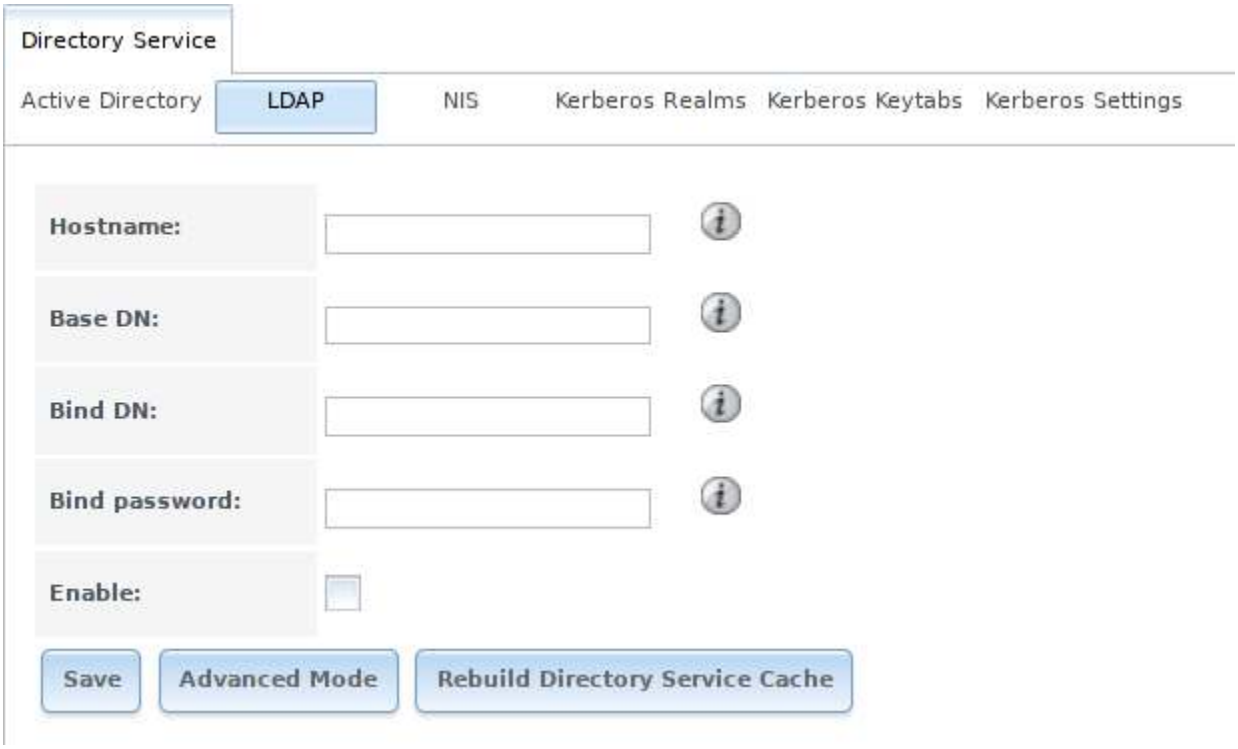

Fig. 8.2: Configuring LDAP

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

Table 8.3: LDAP Configuration Options

Setting	Value	Advanced Mode	Description
Hostname	string		Hostname or IP address of the LDAP server.
Base DN	string		Top level of the LDAP directory tree to be used when searching for resources. Example: *dc=test,dc=org*.
Bind DN	string		Name of administrative account on the LDAP server. Example: *cn=Manager,dc=test,dc=org*.
Bind password	string		Password for *Root bind DN*.
Allow Anonymous Binding	checkbox	✓	Instructs the LDAP server to not provide authentication and to allow read and write access to any client.
User Suffix	string	✓	Optional. Can be added to the name when the user account is added to the LDAP directory. Example: dept. or company name.
Group Suffix	string	✓	Optional. Can be added to the name when the group is added to the LDAP directory. Example: dept. or company name.
Password Suffix	string	✓	Optional. Can be added to the password when the password is added to LDAP directory.
Machine Suffix	string	✓	Optional. Can be added to the name when the system added to the LDAP directory. Example: server, accounting.
SUDO Suffix	string	✓	Use if LDAP-based users need superuser access.

Continued on next page

Table 8.3 – continued from previous page

Setting	Value	Advanced Mode	Description
Kerberos Realm	drop-down menu	✓	Select the realm created using the instructions in *Kerberos Realms* (page 151).
Kerberos Principal	drop-down menu	✓	Browse to the location of the principal in the keytab created as described in *Kerberos Keytabs* (page 151).
Encryption Mode	drop-down menu	✓	Choices are *Off*, *SSL (LDAPS, port 636)*, or *TLS (LDAP, port 389)*. Note that either *SSL* or *TLS* and a *Certificate* must be selected for authentication to work.
Certificate	drop-down menu	✓	Select the certificate of the LDAP CA (required if authentication is used). The certificate for the LDAP server CA must first be imported with *System* → *Certificates* → *Import Certificate*.
LDAP timeout	integer	✓	Increase this value (in seconds) if obtaining a Kerberos ticket times out.
DNS timeout	integer	✓	Increase this value (in seconds) if DNS queries timeout.
Idmap Backend	drop-down menu and Edit button	✓	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 *Edit* to configure the selected backend.
Samba Schema	checkbox	✓	Set if LDAP authentication for SMB shares is needed **and** the LDAP server is **already** configured with Samba attributes.
Auxiliary Parameters	string	✓	Additional options for sssd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=sssd.conf).
Schema	drop-down menu	✓	If *Samba Schema* is set, select the schema to use. Choices are *rfc2307* and *rfc2307bis*.
Enable	checkbox		Unset to disable the configuration without deleting it.
NetBIOS name (This Node)	string	✓	Limited to 15 characters. Automatically populated with the original hostname of the system. This **must** be different from the *Workgroup* name.
NetBIOS name (Node A/B)	string	✓	Limited to 15 characters. When using *Failover* (page 58), set a unique NetBIOS name for the standby node.
NetBIOS alias	string	✓	Limited to 15 characters. When using *Failover* (page 58), 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 the scope and root DN are not to be included when configuring the user, group, password, and machine suffixes.

LDAP users and groups appear in the drop-down menus of the *Permissions* screen of a dataset after configuring the LDAP service. Type `getent passwd` from *Shell* (page 244) to verify the users have been imported. Type `getent group` to verify the groups have been imported.

If the users and groups are not listed, refer to Common errors encountered when using OpenLDAP Software (http://www.openldap.org/doc/admin24/appendix-common-errors.html) for common errors and how to fix them. When troubleshooting LDAP, open *Shell* (page 244) and look for error messages in `/var/log/auth.log`.

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

8.3 NIS

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

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/).

Figure 8.3 shows the configuration screen which opens after going to *Directory Service* → *NIS*. Table 8.4 summarizes the configuration options.

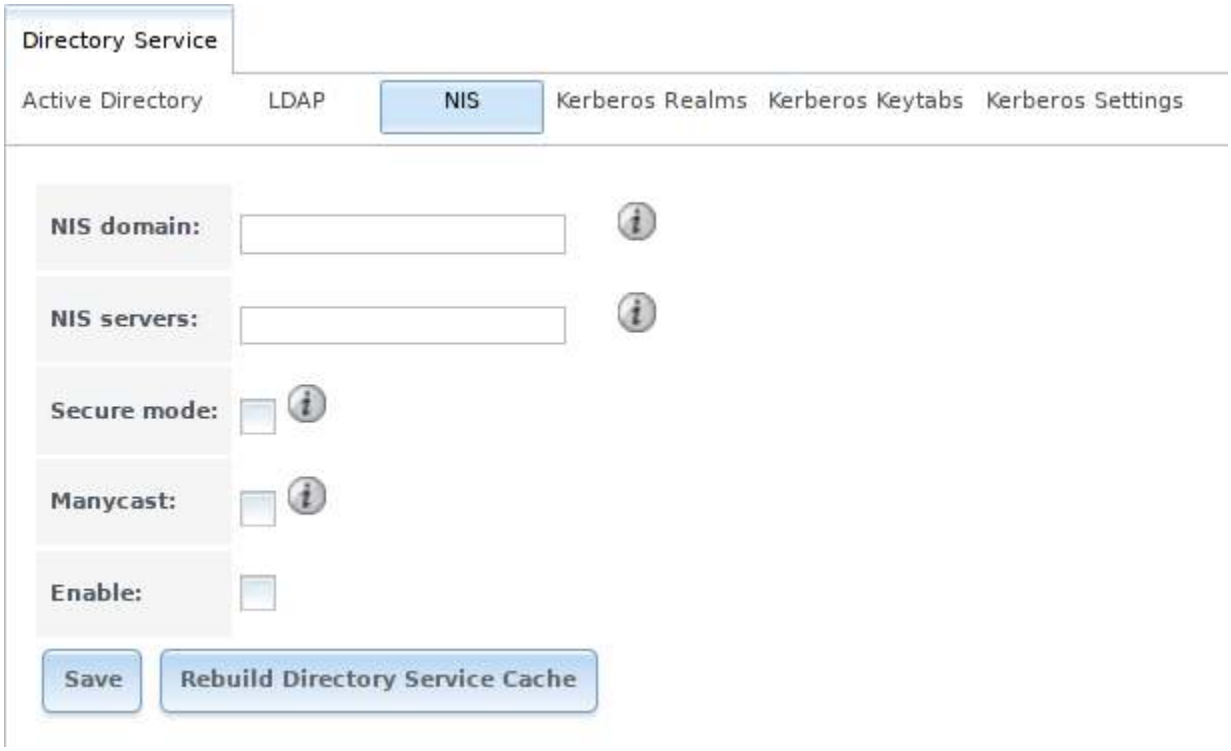

Fig. 8.3: NIS Configuration

Table 8.4: NIS Configuration Options

Setting	Value	Description
NIS domain	string	Name of NIS domain.
NIS servers	string	Comma-delimited list of hostnames or IP addresses.
Secure mode	checkbox	If set, ypbind(8) (https://www.freebsd.org/cgi/man.cgi?query=ypbind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024.
Manycast	checkbox	If set, `ypbind` 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	Unset to disable the configuration without deleting it.

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

8.4 Kerberos Realms

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

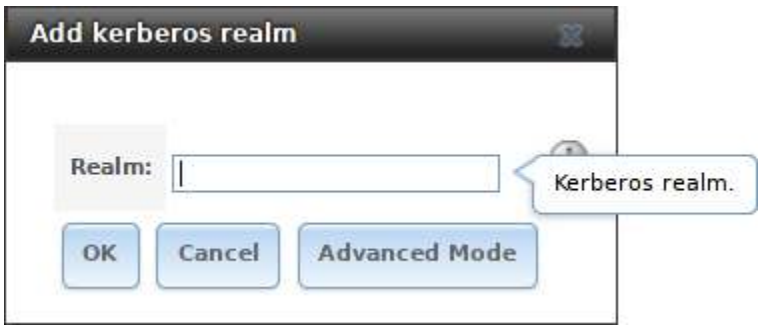

Fig. 8.4: Adding a Kerberos Realm

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

Table 8.5: Kerberos Realm Options

Setting	Value	Advanced Mode	Description
Realm	string		Mandatory. Name of the Kerberos realm.
KDC	string	✓	Name of the Key Distribution Center.
Admin Server	string	✓	Server where all changes to the database are performed.
Password Server	string	✓	Server where all password changes are performed.

8.5 Kerberos Keytabs

Kerberos keytabs are used to do Active Directory or LDAP joins without a password. This means the password for the Active Directory or LDAP administrator account does not need to be saved into the TrueNAS® configuration database, which is a security risk in some environments.

When using a keytab, it is recommended to create and use a less privileged account for performing the required queries as the password for that account will be stored in the TrueNAS® configuration database. To create the keytab on a Windows system, use the ktpass (https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass) command:

```
ktpass.exe /out freenas.keytab /princ http/useraccount@EXAMPLE.COM /mapuser useraccount
/ptype KRB5_NT_PRINCIPAL /crypto ALL /pass userpass
```

where:

- `freenas.keytab` is the file to upload to the TrueNAS® server.
- `http/useraccount@EXAMPLE.COM` is the principal name written in the format *host/user.account@KERBEROS.REALM*. By convention, the kerberos realm is written in all caps, but make sure the case used for the *Kerberos Realm* (page 151) matches the realm name. See this note (https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass#BKMK_remarks) about using `/princ` for more details.
- `useraccount` is the name of the user account for the TrueNAS® server generated in Active Directory Users and Computers (https://technet.microsoft.com/en-us/library/aa998508(v=exchg.65).aspx).

- `userpass` is the password associated with `useraccount`.

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

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

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

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

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

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

8.6 Kerberos Settings

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

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

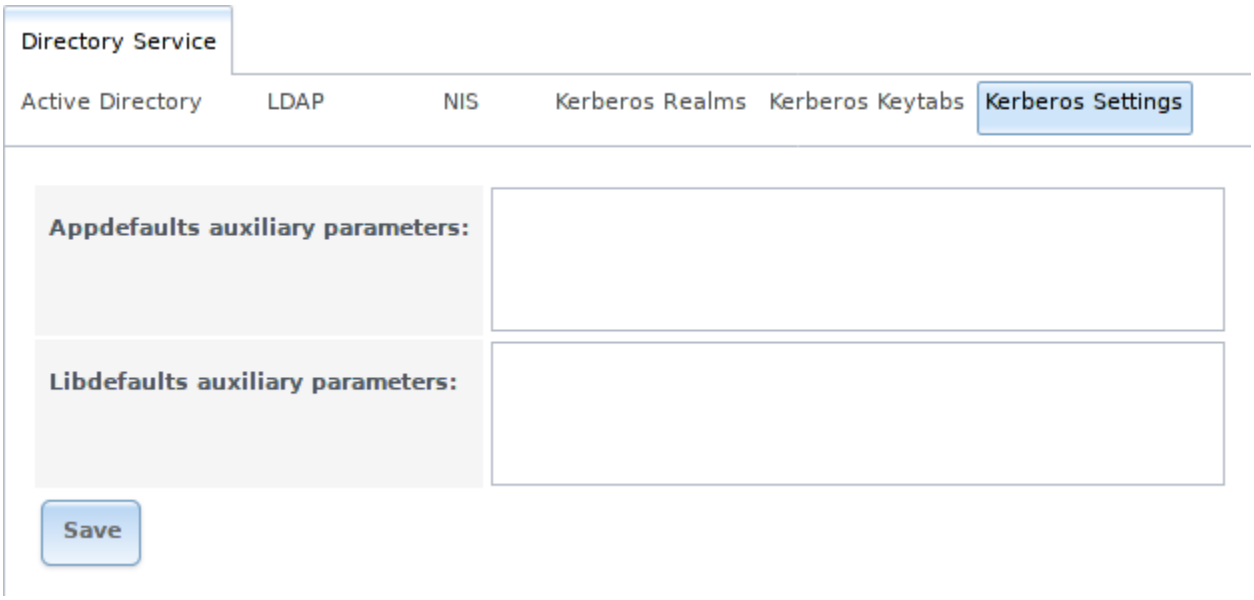

Fig. 8.5: Additional Kerberos Settings

CHAPTER

NINE

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 237) for creating shares. The *Wizard* (page 237) 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, making a list of the users that need access to storage data, which operating systems these users are using, whether all users should have the same permissions to the stored data, and whether these users should authenticate before accessing the data is recommended. This information can help determine which type of shares are needed, whether multiple datasets are needed to divide the storage into areas with different access and permissions, and how complex it will be to set up those permission requirements. Note that shares are used to provide access to data. When a share is deleted, it removes access to data but does not delete the data itself.

These types of shares and services are available:

- *AFP* (page 154): Apple Filing Protocol shares are used when the client computers all run macOS. Apple has deprecated AFP in favor of *SMB* (page 166). Using AFP in modern networks is no longer recommended.
- *Unix (NFS)* (page 158): Network File System shares are accessible from macOS, Linux, BSD, and the professional and enterprise versions (but not the home editions) of Windows. This can be a good choice when the client computers do not all run the same operating system but NFS client software is available for all of them.
- *WebDAV* (page 165): WebDAV shares are accessible using an authenticated web browser (read-only) or WebDAV client (https://en.wikipedia.org/wiki/WebDAV#Client_support) running on any operating system.
- *SMB* (page 166): Server Message Block shares, also known as Common Internet File System (CIFS) shares, are accessible by Windows, macOS, Linux, and BSD computers. Access is slower than an NFS share due to the single-threaded design of Samba. SMB provides more configuration options than NFS and is a good choice on a network for Windows or Mac systems. However, it is a poor choice if the CPU on the TrueNAS® system is limited. If it is maxed out, upgrade the CPU or consider a different type of share.
- *Block (iSCSI)* (page 177): 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 205) service instead of a share and using a cross-platform FTP file manager application such as Filezilla (https://filezilla-project.org/). Secure FTP can be configured if the data needs to be encrypted.

When data security is a concern and the network users are familiar with SSH command line utilities or WinSCP (https://winscp.net/eng/index.php), consider using the *SSH* (page 225) service instead of a share. It is slower than unencrypted FTP due to the encryption overhead, but the data passing through the network is encrypted.

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

This section demonstrates configuration and fine-tuning of AFP, NFS, SMB, WebDAV, and iSCSI shares. FTP and SSH configurations are described in *Services* (page 198).

9.1 Apple (AFP) Shares

TrueNAS® uses the Netatalk (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 237). It then provides configuration examples for using the *Wizard* (page 237) 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 macOS 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.

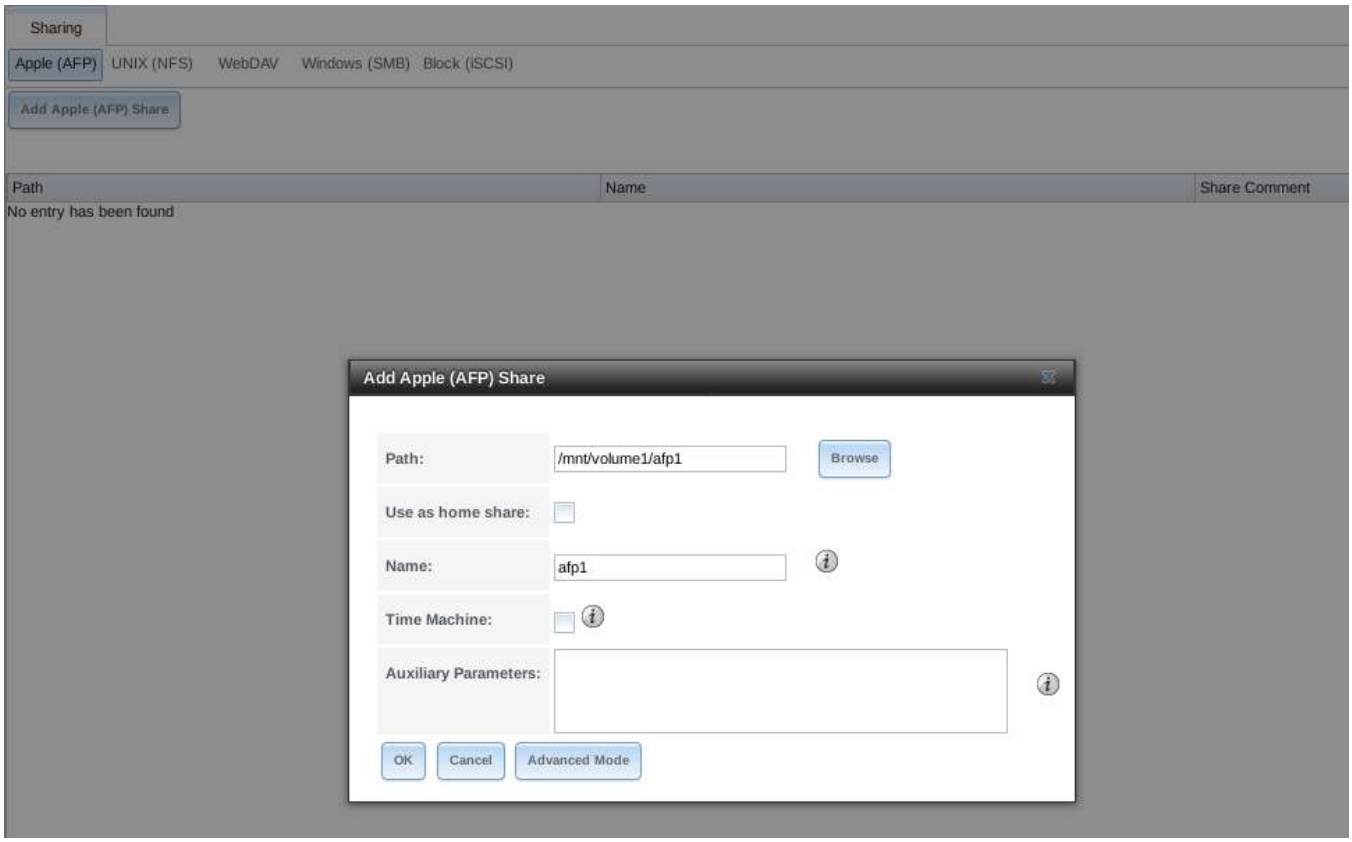

Fig. 9.1: Creating an AFP Share

Note: Table 9.1 summarizes the options available to fine-tune an AFP share. Leaving these options at the de-

fault settings is recommended as changing them can cause unexpected behavior. Most settings are only available with *Advanced Mode*. Do **not** change an advanced option without fully understanding the function of that option. Refer to Setting up Netatalk (http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html) 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. Netatalk does not fully support nesting functionality.
Use as home share	checkbox		Set to allow the share to host user home directories. Only one share can be used as the home share.
Name	string		Enter the volume name that appears in in macOS after selecting *Go* → *Connect to server* in the Finder menu. Limited to 27 characters and cannot contain a period.
Share Comment	string	✓	Enter an optional comment.
Allow List	string	✓	Comma-delimited list of allowed users and/or groups where groupname begins with a `@`. Note that adding an entry will deny any user/group that is not specified.
Deny List	string	✓	Comma-delimited list of denied users and/or groups where groupname begins with a `@`. Note that adding an entry will allow all users/groups that are not specified.
Read-only Access	string	✓	Comma-delimited list of users and/or groups who only have read access where groupname begins with a `@`.
Read-write Access	string	✓	Comma-delimited list of users and/or groups who have read and write access where groupname begins with a `@`.
Time Machine	checkbox		Set to advertise TrueNAS® as a Time Machine disk so it can be found by Macs. Setting multiple shares for Time Machine use is not recommended. When multiple Macs share the same pool, low diskspace issues and intermittently failed backups can occur.
Time Machine Quota, GiB	integer		Appears when *Time Machine* is set. Enter a storage quota for each Time Machine backup on this share. The share must be remounted for any changes to this value to take effect.
Zero Device Numbers	checkbox	✓	Enable when the device number is not constant across a reboot.
No Stat	checkbox	✓	If enabled, 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	✓	Set to enable Unix privileges supported by Mac OS X 10.5 and higher. Do not enable if the network has Mac OS X 10.4 or lower clients. Those systems do not support this feature.
Default file 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 is used for newly created files. Default is *000* (anyone can read, write, and execute).
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma, space, or tab.

Continued on next page

Table 9.1 – continued from previous page

Setting	Value	Advanced Mode	Description
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Separate entries with a comma, space, or tab.
Auxiliary Parameters	string		Additional afp.conf (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered by other option fields.

9.1.1 Creating AFP Guest Shares

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

Note: When a guest share is created along with a share that requires authentication, AFP only maps users who log in as *guest* to the guest share. If a user logs in to the share that requires authentication, permissions on the guest share can prevent that user from writing to the guest share. The only way to allow both guest and authenticated users to write to a guest share is to set the permissions on the guest share to *777* or to add the authenticated users to a guest group and set the permissions to *77x*.

Before creating a guest share, go to *Services* → *AFP* and make sure that the *Guest Access* option is enabled.

To create the AFP guest share, click *Wizard*, then click the *Next* button three times 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*.

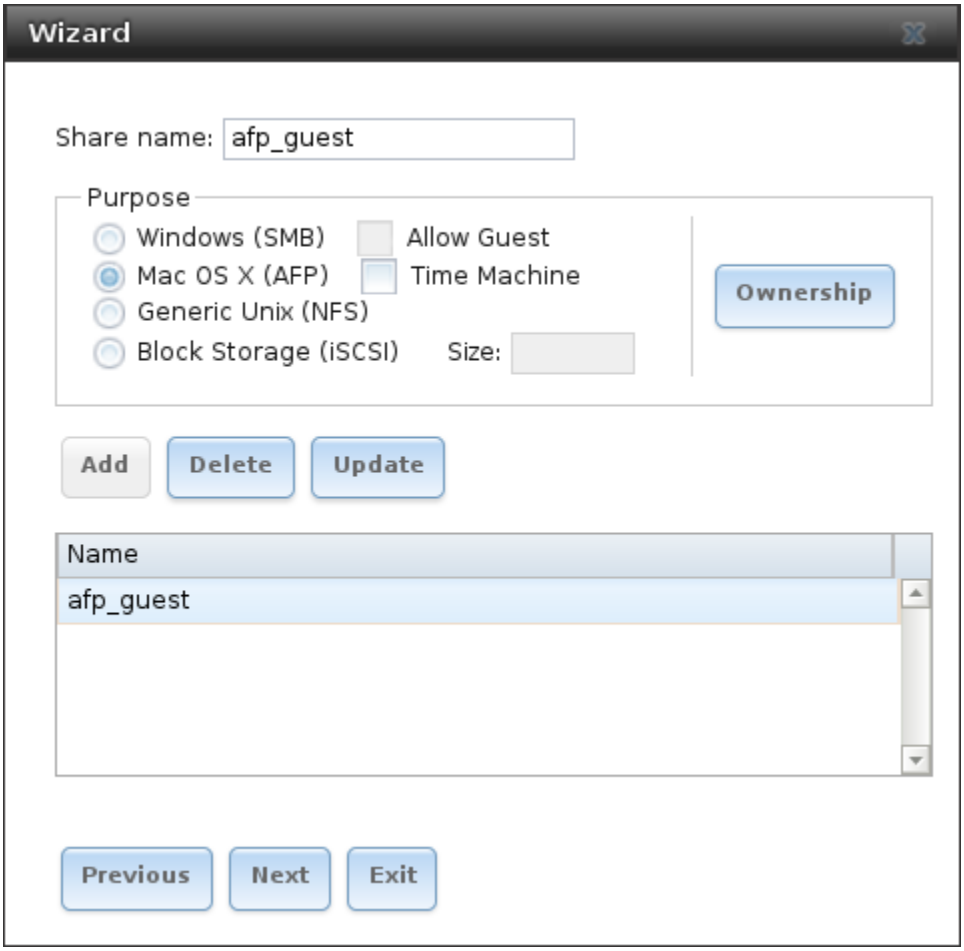

Fig. 9.2: Creating a Guest AFP Share

Click the *Next* button three times, 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)*.

macOS users can use Finder to connect to the guest AFP share by clicking *Go* → *Connect to Server*. In the example shown in Figure 9.3, the user entered `afp://` followed by the IP address of the TrueNAS® system.

Click the *Connect* button. Once connected, Finder opens automatically. The name of the AFP share is displayed in the SHARED section in the left frame and the contents of any data saved in the share is displayed in the right frame.

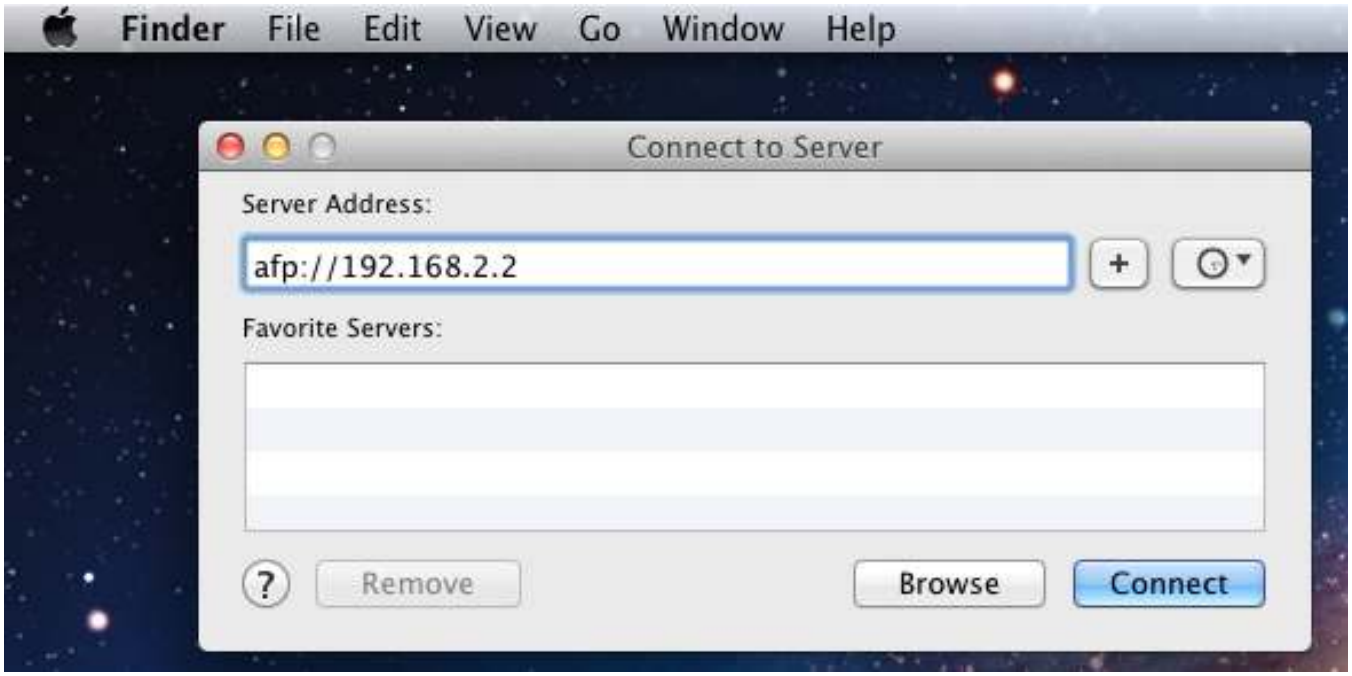

Fig. 9.3: Connect to Server Dialogue

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

9.2 Unix (NFS) Shares

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

To create an NFS share using the *Wizard* (page 237), click the *Next* button three times to display the screen shown in Figure 9.4. 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 the requirements, the IP addresses that are allowed to access the NFS share can be restricted, or the permissions adjusted.

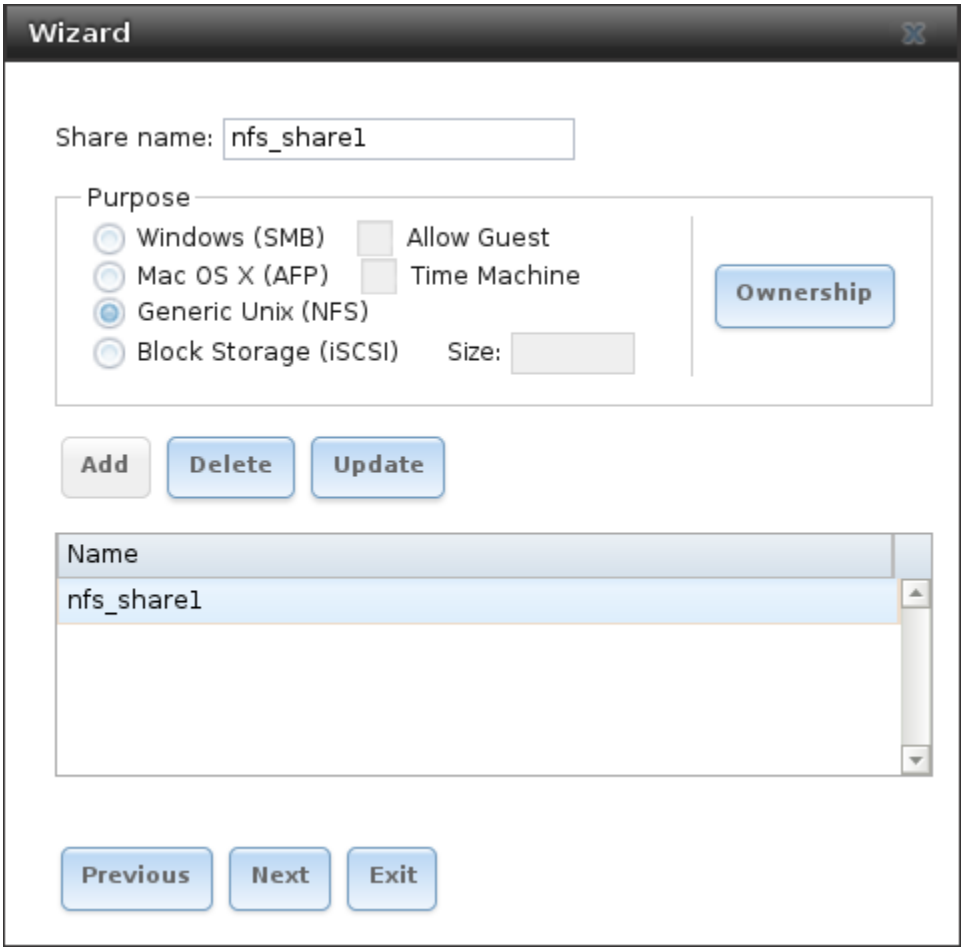

Fig. 9.4: NFS Share Wizard

NFS shares are edited by clicking *Sharing* → *Unix (NFS)*, highlighting the entry for the share, and clicking the *Edit* button. In the example shown in Figure 9.5, the configuration screen is open for the *nfs_share1* share.

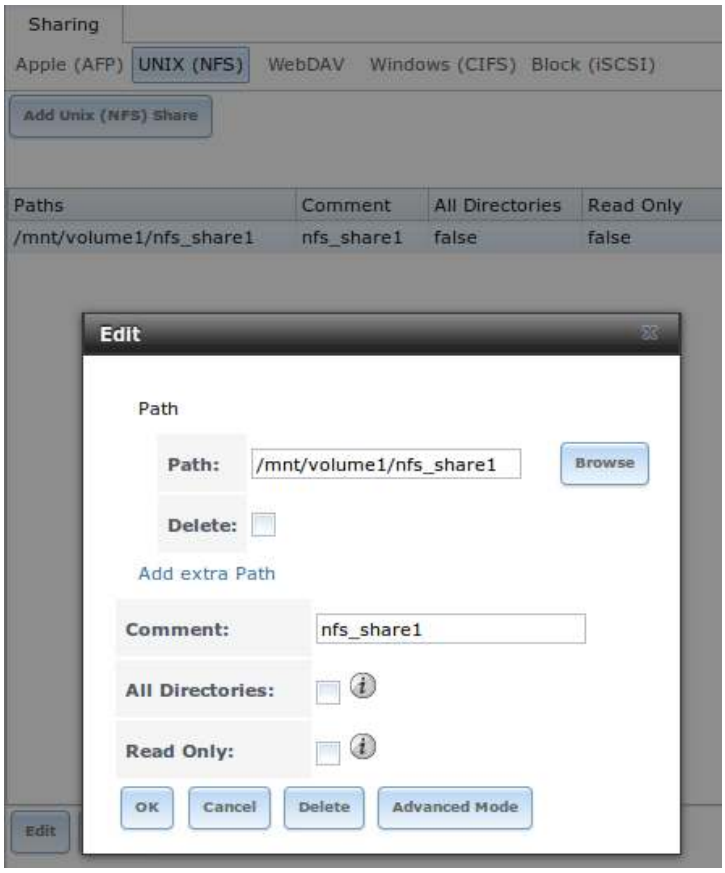

Fig. 9.5: NFS Share Settings

Remember these points when creating NFS shares:

1. Clients specify the *Path* when mounting the share.
2. The *Maproot* and *Mapall* options cannot both be enabled. The *Mapall* options supersede the *Maproot* options. To restrict only the *root* user permissions, set the *Maproot* option. To restrict permissions of all users, set the *Mapall* options.
3. Each volume or dataset is considered to be a unique filesystem. Individual NFS shares cannot cross filesystem boundaries. Adding paths to share more directories only works if those directories are within the same filesystem.
4. The network and host must be unique to both each created share and the filesystem or directory included in that share. Because `/etc/exports` is not an access control list (ACL), the rules contained in `/etc/exports` become undefined with overlapping networks or when using the same share with multiple hosts.
5. The *All dirs* option can only be used once per share per filesystem.

To better understand these restrictions, consider a scenario where there are:

- two networks, *10.0.0.0/8* and *20.0.0.0/8*
- a ZFS volume named `volume1` with 2 datasets named `dataset1` and `dataset2`
- `dataset1` contains directories named `directory1`, `directory2`, and `directory3`

Because of restriction #3, an error is shown when trying to create one NFS share like this:

- *Authorized networks* set to *10.0.0.0/8 20.0.0.0/8*
- *Path* set to `/mnt/volume1/dataset1` and `/mnt/volume1/dataset1/directory1`

The correct method to configure this share is to set the *Path* to `/mnt/volume1/dataset1` and set *All Directories*. This allows the client to also mount `/mnt/volume1/dataset1/directory1` when `/mnt/volume1/dataset1` is mounted.

Additional paths are used to define specific directories to be shared. For example, `dataset1` has three directories. To share only `/mnt/volume1/dataset1/directory1` and `/mnt/volume1/dataset1/directory2`, create paths for `directory1` and `directory2` within the share. This excludes `directory3` from the share.

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

First NFS share:

- *Authorized networks* set to *10.0.0.0/8*
- *Path* set to `/mnt/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 creating two shares. It cannot be done with only one share.

Table 9.2 summarizes the available configuration options in *NFS Share Settings* (page 160). Click *Advanced Mode* to see all settings.

Table 9.2: NFS Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		*Browse* to the volume, dataset, or directory to be shared. Click *Add extra Path* to add multiple directories to this share.
Comment	string		Text describing the share. Typically used to name the share. If left empty, this shows the *Path* entries of the share.
Authorized networks	string	✓	Space-delimited list of allowed networks in network/mask CIDR notation. Example: *1.2.3.0/24*. Leave empty to allow all.
Authorized IP addresses or hosts	string	✓	Space-delimited list of allowed IP addresses or hostnames. Leave empty to allow all.
All directories	checkbox		Allow the client to also mount any subdirectories of the selected pool or dataset.
Read only	checkbox		Prohibit writing to the share.
Quiet	checkbox	✓	Restrict some syslog diagnostics to avoid some error messages. See exports(5) (https://www.freebsd.org/cgi/man.cgi?query=exports) for examples.
Maproot User	drop-down menu	✓	When a user is selected, the *root* user is limited to permissions of that user.
Maproot Group	drop-down menu	✓	When a group is selected, the *root* user is also limited to permissions of that group.
Mapall User	drop-down menu	✓	All clients use the permissions of the specified user.
Mapall Group	drop-down menu	✓	All clients use the permissions of the specified group.

Continued on next page

Table 9.2 – continued from previous page

Setting	Value	Advanced Mode	Description
Security	selection	✓	Only appears if *Enable NFSv4* is enabled in *Services* → *NFS*. Choices are *sys* or these Kerberos options: *krb5* (authentication only), *krb5i* (authentication and integrity), or *krb5p* (authentication and privacy). If multiple security mechanisms are added to the *Selected* column using the arrows, use the *Up* or *Down* buttons to list in order of preference.

9.2.1 Example Configuration

By default, the *Mapall* fields are not set. This means that when a user connects to the NFS share, the user has the permissions associated with their user account. This is a security risk if a user is able to connect as *root* as they will have complete access to the share.

A better option is to do this:

1. Specify the built-in *nobody* account to be used for NFS access.
2. In the *Change Permissions* screen of the volume/dataset that is being shared, change the owner and group to *nobody* and set the permissions according to the desired requirements.
3. Select *nobody* in the *Mapall User* and *Mapall Group* drop-down menus for the share in *Sharing* → *Unix (NFS) Shares*.

With this configuration, it does not matter which user account connects to the NFS share, as it will be mapped to the *nobody* user account and will only have the permissions that were specified on the 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* option is enabled
   - *MapAll User* is set to *nobody*
   - *MapAll Group* is set to *nobody*

9.2.2.1 From BSD or Linux

NFS shares are mounted on BSD or Linux clients with this command executed as the superuser (*root*) or with `sudo`:

```
mount -t nfs 192.168.2.2:/mnt/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.

Successfully mounting the share returns to the command prompt without any status or error messages.

Note: If this command fails on a Linux system, make sure that the nfs-utils (https://sourceforge.net/projects/nfs/files/nfs-utils/) package is installed.

This configuration allows users on the client system to copy files to and from `/mnt` (the mount point). All files are owned by *nobody:nobody*. Changes to any files or directories in `/mnt` write to the TrueNAS® system `/mnt/volume1/nfs_share1` dataset.

NFS share settings cannot be changed when the share is mounted on a client computer. The `umount` command is used to unmount the share on BSD and Linux clients. Run it as the superuser or with `sudo` on each client computer:

```
umount /mnt
```

9.2.2.2 From Microsoft

Windows NFS client support varies with versions and releases. For best results, use *Windows (SMB) Shares* (page 166).

9.2.2.3 From macOS

A macOS client uses Finder to mount the NFS volume. Go to *Go* → *Connect to Server*. In the *Server Address* field, enter *nfs://* followed by the IP address of the TrueNAS® system and the name of the volume/dataset being shared by NFS. The example shown in Figure 9.6 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 displays in the SHARED section in the left frame and the contents of the share display in the right frame. Figure 9.7 shows an example where `/mnt/data` has one folder named `images`. The user can now copy files to and from the share.

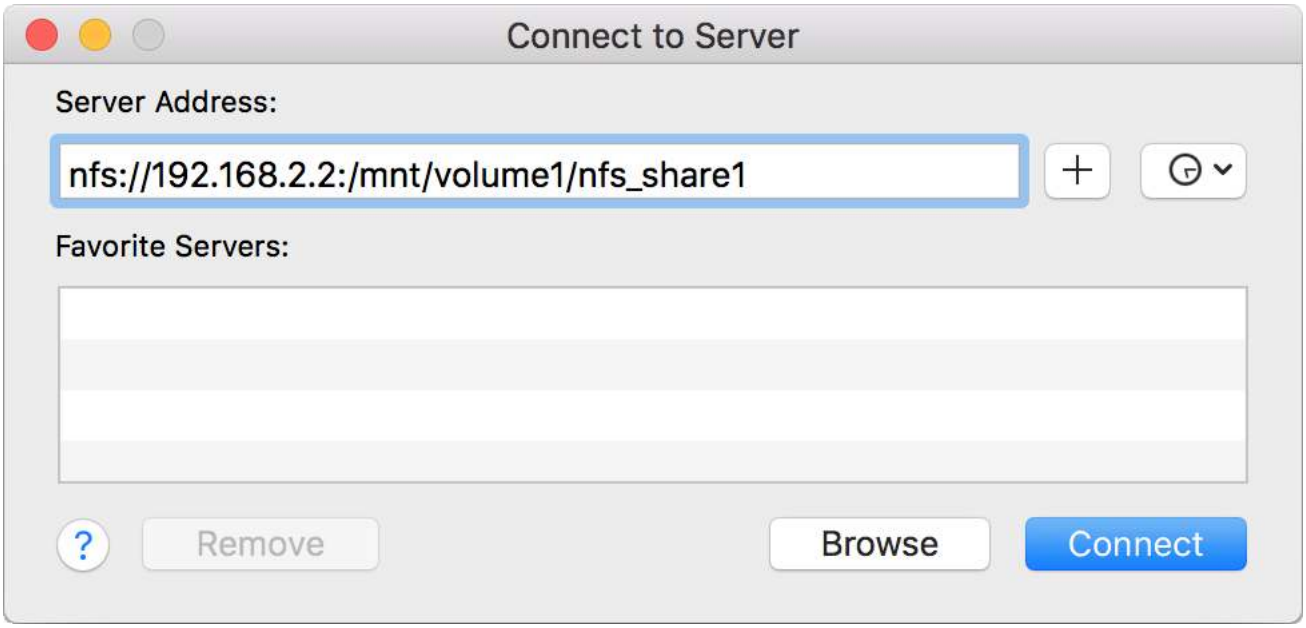

Fig. 9.6: Mounting the NFS Share from macOS

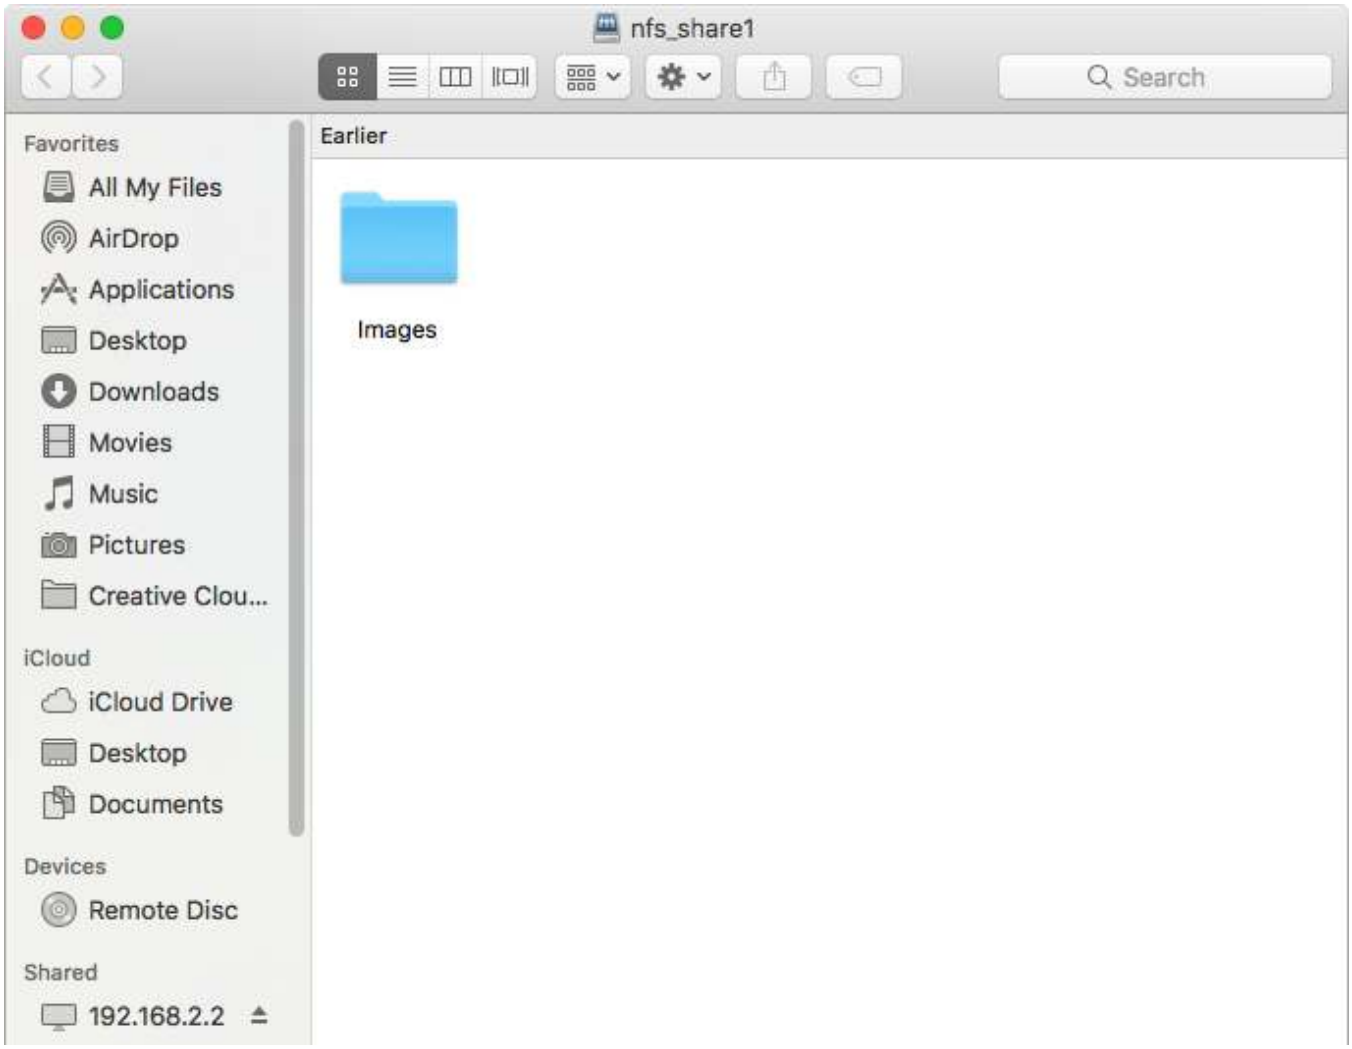

Fig. 9.7: Viewing the NFS Share in Finder

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 data base* 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 enable the *Serve UDP NFS clients* option.

The `nfsstat -c` or `nfsstat -s` commands can be helpful to detect problems from the *Shell* (page 244). 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 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 9.8.

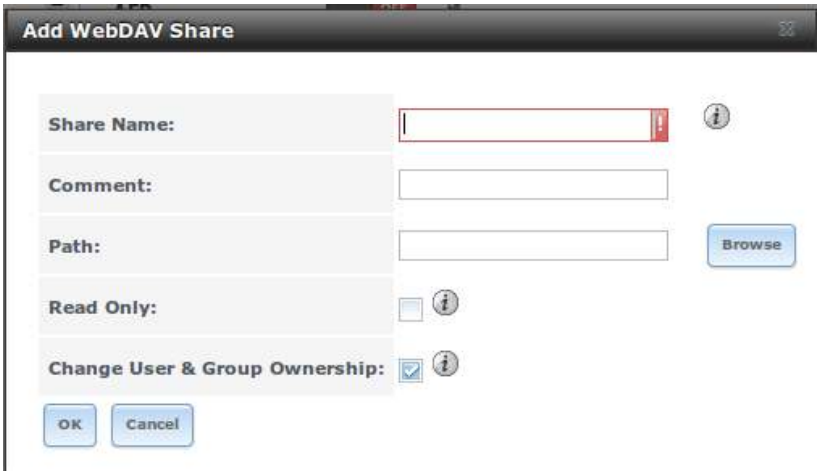

Fig. 9.8: Adding a WebDAV Share

Table 9.3 summarizes the available options.

Table 9.3: WebDAV Share Options

Setting	Value	Description
Share Path Name	string	Enter a name for the share.
Comment	string	Optional.

Continued on next page

Table 9.3 – continued from previous page

Setting	Value	Description
Path	browse button	*Browse* to the volume/dataset to share.
Read Only	checkbox	Set to prohibit users from writing to the share.
Change User & Group Ownership	checkbox	Enable to automatically set the share contents to the *webdav* user and group.

After clicking *OK*, a pop-up asks about enabling the service. Once the service starts, review the settings in *Services* → *WebDAV* as they are used to determine which URL is used to access the WebDAV share and whether or not authentication is required to access the share. These settings are described in *WebDAV* (page 232).

9.4 Windows (SMB) Shares

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

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

Samba supports server-side copy of files on the same share with clients from Windows 8 and higher. Copying between two different shares is not server-side. Windows 7 clients support server-side copying with Robocopy (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc733145(v=ws.11)).

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

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

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

Figure 9.9 shows the configuration screen that appears after clicking *Sharing* → *Windows (SMB Shares)* → *Add Windows (SMB) Share*.

Fig. 9.9: 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) provides more details for each configurable option.

Table 9.4: SMB Share Options

Setting	Value	Advanced Mode	Description
Path	browse button		Select the volume, dataset, or directory to share. The same path can be used by more than one share.
Name	string		Enter a name for this share. An existing SMB share name can not be reused.
Use as home share	checkbox		Set to allow this share to hold user home directories. Only one share can be the home share. Note that lower case names for user home directories are strongly recommended, as Samba maps usernames to all lower case. For example, the username John will be mapped to a home directory named `john`. If the *Path* to the home share includes an upper case username, delete the existing user and *recreate* (page 19) it in *Accounts* → *Users* with an all lower case *Username*. Return to *Sharing* → *SMB* to create the home share, and select the *Path* that contains the new lower case username.
Time Machine	checkbox		Enable Time Machine (https://developer.apple.com/library/archive/releasenotes/NetworkingIn CH1-SW1) backups for this share. See *Configuring Time Machine Backups* (page 196).

Continued on next page

Table 9.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Name	string		Name the new share. Each share name must be unique.
Apply Default Permissions	checkbox		ACLs grant *read* and *write* for *owner* or *group* and *read-only* for others. Leave this unset when creating shares on a system with custom ACLs.
Comment	string	✓	Optional description.
Export Read Only	checkbox	✓	Prohibit write access to the share.
Browsable to Network Clients	checkbox	✓	Determine whether this share name is included when browsing shares. Home shares are only visible to the owner regardless of this setting.
Export Recycle Bin	checkbox	✓	Files that are deleted from the same dataset are moved to the Recycle Bin and do not take any additional space. When the files are in a different dataset or a child dataset, they are copied to the dataset where the Recycle Bin is located. To prevent excessive space usage, files larger than 20 MiB are deleted rather than moved. Adjust the *Auxiliary Parameter* `crossrename:sizelimit=` setting to allow larger files. For example, `crossrename:sizelimit=50` allows moves of files up to 50 MiB in size.
Show Hidden Files	checkbox	✓	Disable the Windows *hidden* attribute on a new Unix hidden file. Unix hidden filenames start with a dot: `.foo`. Existing files are not affected.
Allow Guest Access	checkbox		Privileges are the same as the guest account. Guest access is disabled by default in Windows 10 version 1709 and Windows Server version 1903. Additional client-side configuration is required to provide guest access to these clients.
Only Allow Guest Access	checkbox	✓	Requires *Allow guest access* to also be enabled. Forces guest access for all connections.
Access Based Share Enumeration	checkbox	✓	Restrict share visibility to users with a current Windows Share ACL access of read or write. Use Windows administration tools to adjust the share permissions. See smb.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=smb.conf).
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma (,), space, or tab.
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Separate entries with a comma (,), space, or tab. Specify `ALL` and list any hosts from *Hosts Allow* to have those hosts take precedence.
VFS Objects	selection	✓	Add virtual file system modules to enhance functionality. Table 9.5 summarizes the available modules.
Periodic Snapshot Task	drop-down menu	✓	Used to configure directory shadow copies on a per-share basis. Select the pre-configured periodic snapshot task to use for the shadow copies of the share. Periodic snapshots must be recursive.
Auxiliary Parameters	string	✓	Additional smb4.conf (https://www.freebsd.org/cgi/man.cgi?query=smb.conf) parameters not covered by other option fields.

Here are some notes about *ADVANCED MODE* settings:

- Hostname lookups add some time to accessing the SMB share. If only using IP addresses, unset the *Hostnames lookups* option in *Services* → *SMB*.
- When the *Browsable to Network Clients* option is enabled (the default), the share is visible through Windows File

Explorer or through `net view`. When the *Use as a home share* option is selected, deselecting the *Browsable to Network Clients* option hides the share named *homes* so that only the dynamically generated share containing the authenticated user home directory will be visible. By default, the *homes* share and the user home directory are both visible. Users are not automatically granted read or write permissions on browsable shares. This option provides no real security because shares that are not visible in Windows File Explorer can still be accessed with a *UNC* path.

- If some files on a shared 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).

Samba disables NTLMv1 authentication by default for security. Standard configurations of Windows XP and some configurations of later clients like Windows 7 will not be able to connect with NTLMv1 disabled. Security guidance for NTLMv1 and LM network authentication (https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-network-authentication) has information about the security implications and ways to enable NTLMv2 on those clients. If changing the client configuration is not possible, NTLMv1 authentication can be enabled by enabling the *NTLMv1 auth* option in *Services* → *SMB*.

Table 9.5 provides an overview of the available VFS modules. Be sure to research each module **before** adding or deleting it from the *Selected* column of the *VFS Objects* field of the share. Some modules need additional configuration after they are added. Refer to Stackable VFS modules (https://www.samba.org/samba/docs/old/Samba3-HOWTO/VFS.html) and the vfs_* man pages (https://www.samba.org/samba/docs/current/man-html/) for more details.

Table 9.5: Available VFS Modules

Value	Description
acl_tdb	Store NTFS ACLs in a tdb file to enable full mapping of Windows ACLs.
acl_xattr	Store NTFS ACLs in Extended Attributes (EAs) to enable the full mapping of Windows ACLs.
aio_fork	Enable async I/O.
audit	Log share access, connects/disconnects, directory opens/creates/removes, and file opens/closes/renames/unlinks/chmods to syslog.
cacheprime	Prime the kernel file data cache.
cap	Translate filenames to and from the CAP encoding format, commonly used in Japanese language environments.
catia	Improve Mac interoperability by translating characters that are unsupported by Windows.
commit	Track the amount of data written to a file and synchronize it to disk when a specified amount accumulates.
crossrename	Allow server side rename operations even if source and target are on different physical devices. Required for the recycle bin to work across dataset boundaries. Automatically added when *Export Recycle Bin* is enabled.
default_quota	**Deprecated: use the ixnas module instead.** Store the default quotas that are reported to a Windows client in the quota record of a user.
dirsort	Sort directory entries alphabetically before sending them to the client.
expand_msdfs	Enable support for Microsoft Distributed File System (DFS).
extd_audit	Send audit logs to both syslog and the Samba log files.

Continued on next page

Table 9.5 – continued from previous page

Value	Description
fake_perms	Allow roaming profile files and directories to be set to read-only.
fruit	Enhance macOS support by providing the SMB2 AAPL extension and Netatalk interoperability. Automatically loads *catia* and *streams_xattr*, but see the *warning* (page 171) below.
full_audit	Record selected client operations to the system log.
ixnas	Experimental module to improve ACL compatibility with Windows, store DOS attributes as file flags, and enable *User Quota Administration* (page 175) from Windows. Several *Auxiliary Parameters* are available with *ixnas*. Userspace Quota Settings: • *ixnas:base_user_quota* = sets a ZFS user quota on every user that connects to the share. Example: `ixnas:base_user_quota = 80G` sets the quota to *80 GiB*. • *ixnas:zfs_quota_enabled* = enables support for userspace quotas. Choices are *True* or *False*. Default is *True*. Example: `ixnas:zfs_quota_enabled = True`. Home Dataset Settings: • *ixnas:chown_homedir* = changes the owner of a created home dataset to the currently authenticated user. *ixnas:zfs_auto_homedir* must be set to *True*. Choices are *True* or *False*. Example: `ixnas:chown_homedir = True`. • *ixnas:homedir_quota* = sets a quota on new ZFS datasets. *ixnas:zfs_auto_homedir* must be set to *True*. Example: `ixnas:homedir_quota = 20G` sets the quota to *20 GiB*. • *ixnas:zfs_auto_homedir* = creates new ZFS datasets for users connecting to home shares instead of folders. Choices are *True* or *False*. Default is *False*. Example: `ixnas:zfs_auto_homedir = False`.
linux_xfs_sgid	Used to work around an old Linux XFS bug.
media_harmony	Allow Avid editing workstations to share a network drive.
netatalk	Ease the co-existence of SMB and AFP shares.
noacl	Disable NT ACL support. If an extended ACL is present in the share connection path, all access to this share will be denied. When the Read-only attribute (https://www.oreilly.com/openbook/samba/book/ch05_03.html) is set, all write bits are removed. Disabling the *Read-only* attribute adds the write bits back to the share, up to *create mask* (*umask*). Adding *noacl* requires adding the *zfsacl* object. *noacl* is incompatible with the *ixnas* VFS object.

Continued on next page

Table 9.5 – continued from previous page

Value	Description
offline	Mark all files in the share with the DOS *offline* attribute. This can prevent Windows Explorer from reading files just to make thumbnail images.
posix_eadb	Provide 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	Mark a share as read-only for all clients connecting within the configured time period.
shadow_copy	Allow Microsoft shadow copy clients to browse shadow copies on Windows shares.
shadow_copy_zfs	Allow Microsoft shadow copy clients to browse shadow copies on Windows shares. This object uses *ZFS snapshots* (page 250) of the shared pool or dataset to create the shadow copies.
shell_snap	Provide shell-script callouts for snapshot creation and deletion operations issued by remote clients using the File Server Remote VSS Protocol (FSRVP).
streams_depot	**Experimental** module to store alternate data streams in a central directory. The association with the primary file can be lost due to inode numbers changing when a directory is copied to a new location See https://marc.info/?l=samba&m=132542069802160&w=2.
streams_xattr	Enable storing NTFS alternate data streams in the file system. Enabled by default.
syncops	Ensure metadata operations are performed synchronously.
time_audit	Log system calls that take longer than the defined number of milliseconds.
unityed_media	Allow multiple Avid clients to share a network drive.
virusfilter	This extremely **experimental** module is still under development and does not work at this time.
winmsa	Emulate the Microsoft *MoveSecurityAttributes=0* registry option. Moving files or directories sets the ACL for file and directory hierarchies to inherit from the destination directory.
worm	Control the writability of files and folders depending on their change time and an adjustable grace period.
xattr_tdb	Store Extended Attributes (EAs) in a tdb file so they can be used on filesystems which do not provide support for EAs.
zfs_space	Correctly calculate ZFS space used by the share, including space used by ZFS snapshots, quotas, and resevations. Enabled by default.
zfsacl	Provide ACL extensions for proper integration with ZFS. Enabled by default.

Warning: Be careful when using multiple SMB shares, some with and some without *fruit*. macOS clients negotiate SMB2 AAPL protocol extensions on the first connection to the server, so mixing shares with and without fruit will globally disable AAPL if the first connection occurs without fruit. To resolve this, all macOS clients need to disconnect from all SMB shares and the first reconnection to the server has to be to a fruit-enabled share.

These VFS objects do not appear in the selection box:

- **recycle:** moves deleted files to the recycle directory instead of deleting them. Controlled by *Export Recycle Bin* in the *SMB share options* (page 167).
- **shadow_copy2:** a more recent implementation of *shadow_copy* with some additional features. *shadow_copy2* and the associated parameters are automatically added to the `smb4.conf` when a *Periodic Snapshot Task* is selected.

To view all active SMB connections and users, enter `smbstatus` in the *Shell* (page 244).

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.

Note: Windows 10, Windows Server 2016 version 1709, and Windows Server 2019 disable SMB2 guest access. Read the Microsoft security notice (https://support.microsoft.com/en-hk/help/4046019/guest-access-in-smb2-disabled-by-default-in-windows-10-and-windows-ser) for details about security vulnerabilities with SMB2 guest access and instructions to re-enable guest logins on these Microsoft systems.

To configure an unauthenticated SMB share, click *Wizard*, then click the *Next* button three times to display the screen shown in Figure 9.10. Complete the following fields in this screen:

1. **Share name:** enter a name for the share that is useful. In this example, the share is named *smb_insecure*.
2. Click the button for *Windows (SMB)* and enable the *Allow Guest* option.
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 this step is forgotten, 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*.

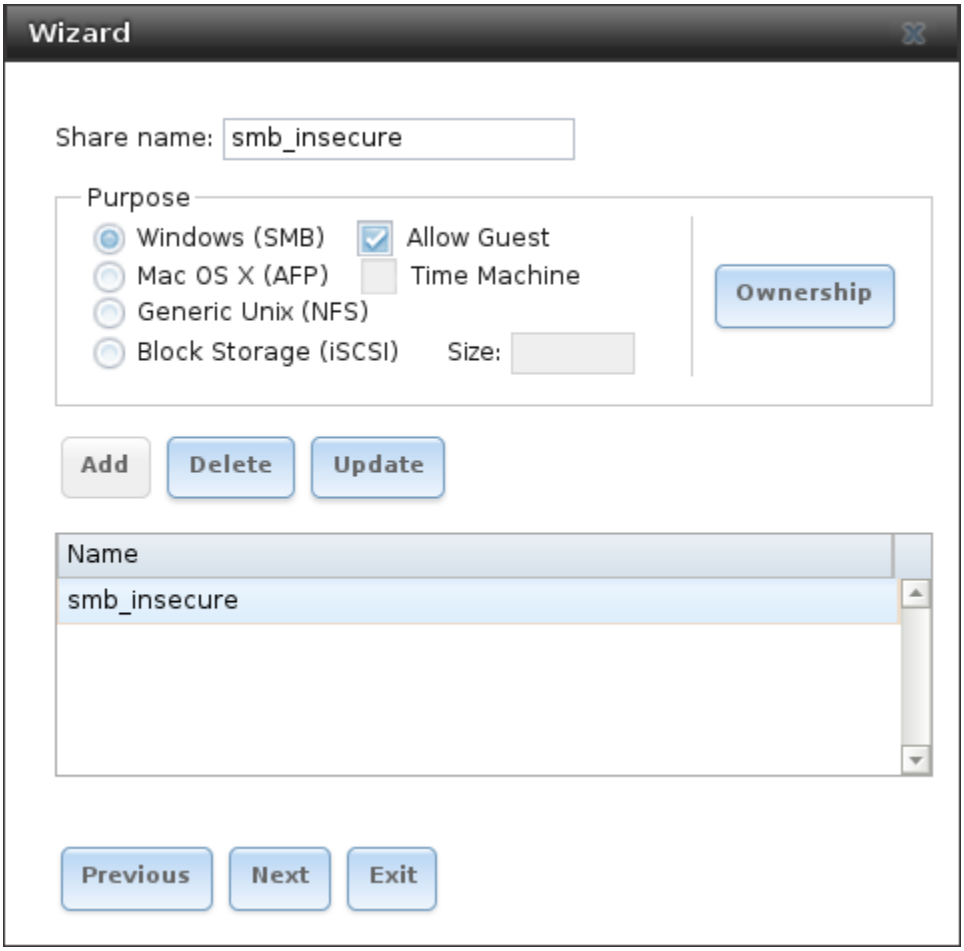

Fig. 9.10: Creating an Unauthenticated SMB Share

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 will appear in *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 With Local Users

Most configuration scenarios require each user to have their own user account and to authenticate before accessing the share. This allows the administrator to control access to data, provide appropriate permissions to that data, and to determine who accesses and modifies stored data. A Windows domain controller is not needed for authenticated SMB shares, which means that additional licensing costs are not required. However, because there is no domain controller to provide authentication for the network, each user account must be created on the TrueNAS® system. This type of configuration scenario is often used in home and small networks as it does not scale well if many user accounts are needed.

Before configuring this scenario, determine which users 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 9.11.

1. **Share name:** enter a name for the share that is useful. 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 enable the *Create User* option. 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 enable the *Create Group* option. In the example shown in Figure 9.12, *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 9.11.
4. Click the *Add* button. **If this step is forgotten, 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*.

When configuring 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)*.

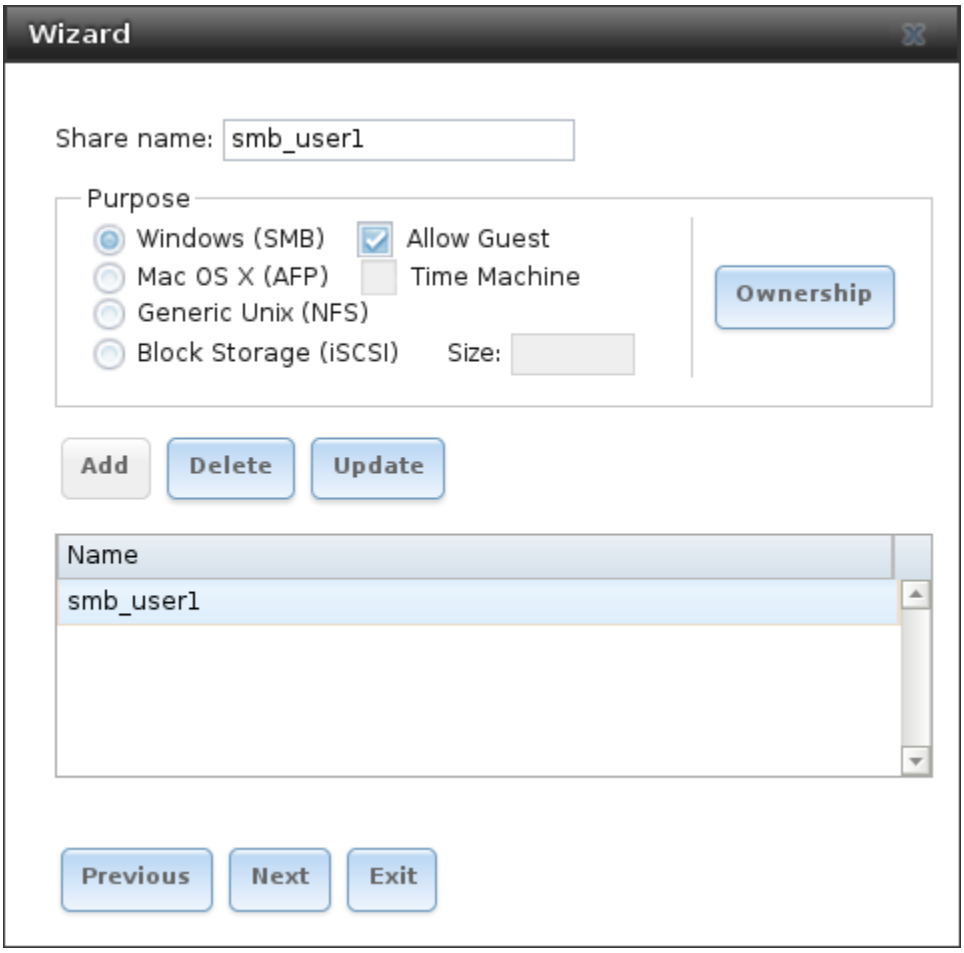

Fig. 9.11: Creating an Authenticated SMB Share

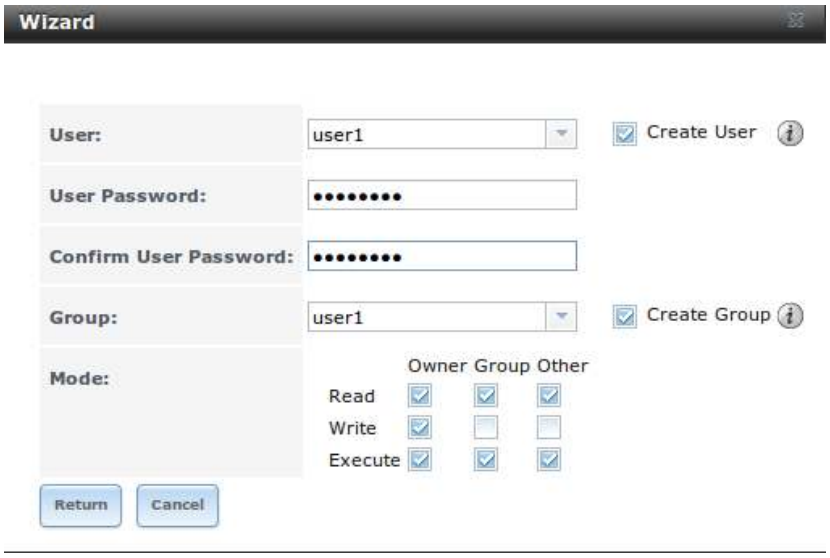

Fig. 9.12: Creating the User and Group

The authenticated share can now be tested from any SMB client. For example, to test an authenticated share from a Windows system with network discovery enabled, open Explorer and click on *Network*. If network discovery is disabled, open Explorer and enter `\HOST` in the address bar, where *HOST* is the IP address or hostname of the share system. This example shows a system named *FREENAS* with a share named *smb_user1*.

After clicking *smb_user1*, a Windows Security dialog prompts for the username and password of the user associated with *smb_user1*. After authenticating, the user can copy data to and from the SMB share.

Map the share as a network drive to prevent Windows Explorer from hanging when accessing the share. Right-click the share and select *Map network drive...*. Choose a drive letter from the drop-down menu and click *Finish*.

Windows caches user account credentials with the authenticated share. This sometimes prevents connection to a share, even when the correct username and password are provided. Logging out of Windows clears the cache. The authentication dialog reappears the next time the user connects to an authenticated share.

9.4.3 User Quota Administration

File Explorer can manage quotas on SMB shares connected to an *Active Directory* (page 141) server. Both the share and dataset being shared must be configured to allow this feature:

- Create an authenticated share with `domain admins` as both the user and group name in *Ownership*.
- Edit the SMB share and add *ixnas* to the list of selected *VFS Object* (page 169).
- In Windows Explorer, connect to and map the share with a user account which is a member of the `domain admins` group. The *Quotas* tab becomes active.

9.4.4 Configuring Shadow Copies

Shadow Copies (https://en.wikipedia.org/wiki/Shadow_copy), also known as the Volume Shadow Copy Service (VSS) or Previous Versions, is a Microsoft service for creating volume snapshots. Shadow copies can be used to restore previous versions of files from within Windows Explorer. Shadow Copy support is built into Vista and Windows 7. Windows XP or 2000 users need to install the Shadow Copy client (http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=16220).

When a periodic snapshot task is created on a ZFS 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 no previous versions of files to restore are visible, 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. To see the shadow copies in the child datasets, create separate shares for them.
- Shadow copies will not work with a manual snapshot. Creating a periodic snapshot task for the pool or dataset being shared by SMB or a recursive task for a parent dataset is recommended.
- The periodic snapshot task should be created and at least one snapshot should exist **before** creating the SMB share. If the SMB share was created first, restart the SMB service in *Services* → *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 With Local Users* (page 173) 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. There are two options for snapshot tasks. One is to create a snapshot task for each user's dataset. In this example the datasets are `/mnt/volume1/user1` and `/mnt/volume1/user2`. Another option is to create one periodic snapshot task for the entire volume,:file*:/mnt/volume1* in this case. **Before continuing to the next step,** confirm that at least one snapshot for each defined task is displayed in the *Storage* → *Snapshots* tab. When creating the schedule for the periodic snapshot tasks, keep in mind how often the users need to access modified files and during which days and time of day they are likely to make changes.
2. Go to *Sharing* → *Windows (SMB) Shares*. 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 9.13 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.

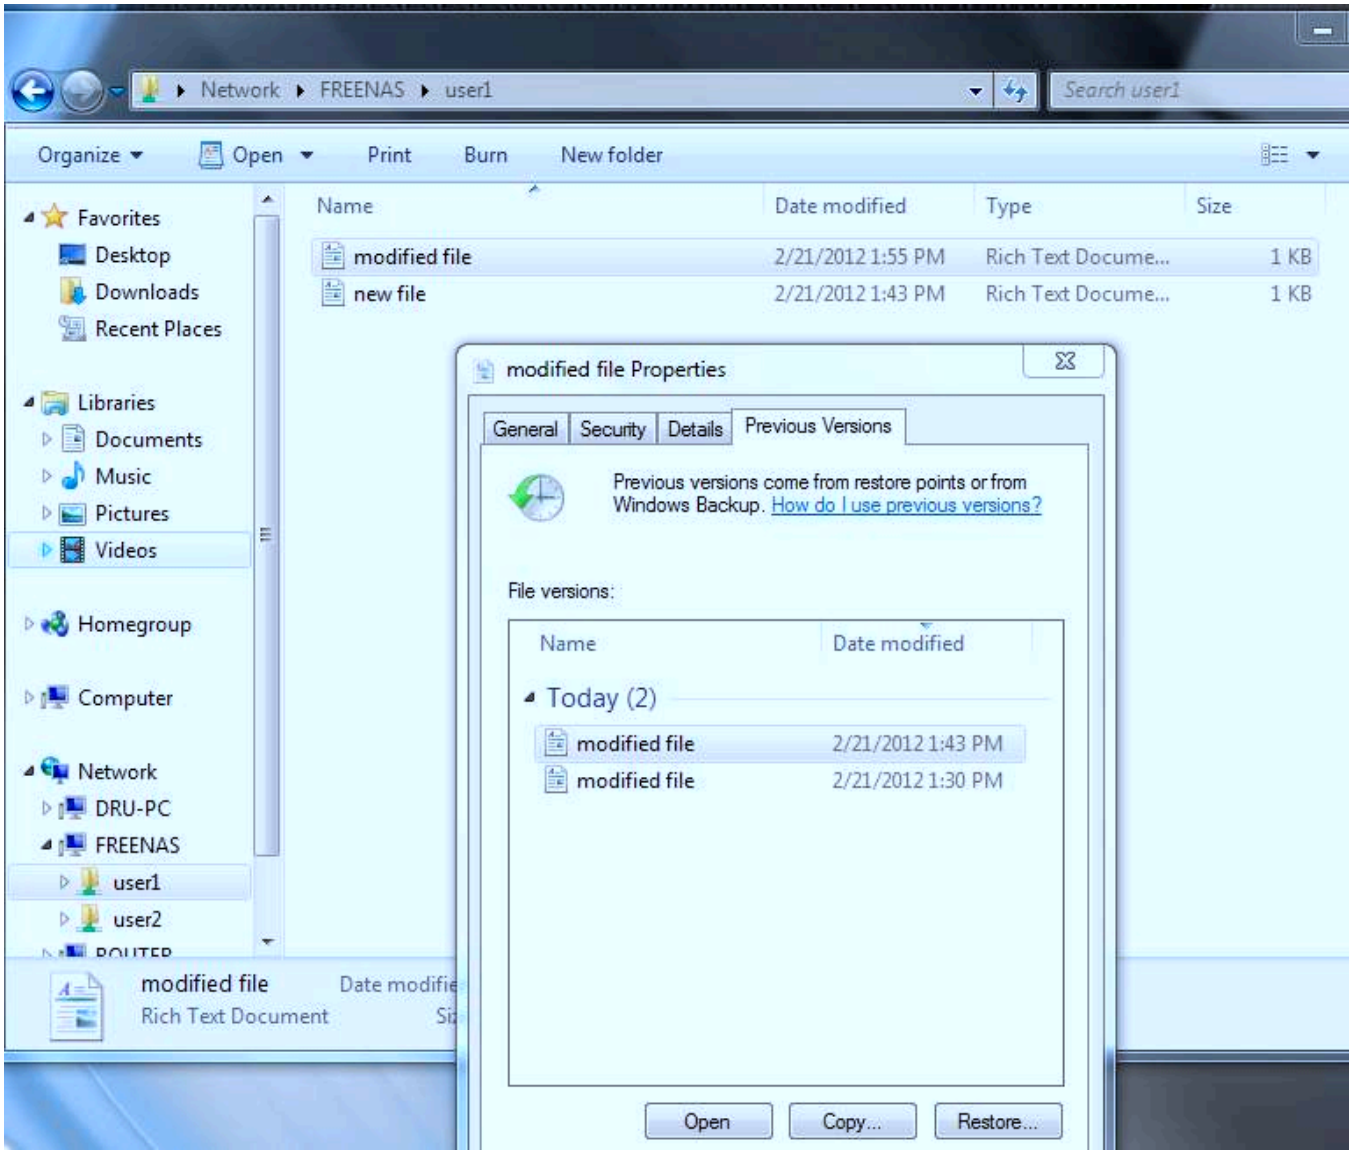

Fig. 9.13: Viewing Previous Versions within Explorer

9.5 Block (iSCSI)

iSCSI is a protocol standard for the consolidation of storage data. iSCSI allows TrueNAS® to act like a storage area network (SAN) over an existing Ethernet network. Specifically, it exports disk devices over an Ethernet network that iSCSI clients (called initiators) can attach to and mount. Traditional SANs operate over fibre channel networks which require a fibre channel infrastructure such as fibre channel HBAs, fibre channel switches, and discrete cabling. iSCSI can be used over an existing Ethernet network, although dedicated networks can be built for iSCSI traffic in an effort to boost performance. iSCSI also provides an advantage in an environment that uses Windows shell programs; these programs tend to filter "Network Location" but iSCSI mounts are not filtered.

Before configuring the iSCSI service, be familiar with this iSCSI terminology:

CHAP: an authentication method which uses a shared secret and three-way authentication to determine if a system is authorized to access the storage device and to periodically confirm that the session has not been hijacked by another system. In iSCSI, the initiator (client) performs the CHAP authentication.

Mutual CHAP: a superset of CHAP in that both ends of the communication authenticate to each other.

Initiator: a client which has authorized access to the storage data on the TrueNAS® system. The client requires initiator software to initiate the connection to the iSCSI share.

Target: a storage resource on the TrueNAS® system. Every target has a unique name known as an iSCSI Qualified Name (IQN).

Internet Storage Name Service (iSNS): protocol for the automated discovery of iSCSI devices on a TCP/IP network.

Extent: the storage unit to be shared. It can either be a file or a device.

Portal: indicates which IP addresses and ports to listen on for connection requests.

LUN: *Logical Unit Number* representing a logical SCSI device. An initiator negotiates with a target to establish connectivity to a LUN. The result is an iSCSI connection that emulates a connection to a SCSI hard disk. Initiators treat iSCSI LUNs as if they were a raw SCSI or SATA hard drive. Rather than mounting remote directories, initiators format and directly manage filesystems on iSCSI LUNs. When configuring multiple iSCSI LUNs, create a new target for each LUN. Since iSCSI multiplexes a target with multiple LUNs over the same TCP connection, there can be TCP contention when more than one target accesses the same LUN. TrueNAS® supports up to 1024 LUNs.

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

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

In TrueNAS®, iSCSI is built into the kernel. This version of iSCSI supports Microsoft Offloaded Data Transfer (ODX) (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)), meaning that file copies happen locally, rather than over the network. It also supports the *VAAI* (page 254) (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 105) and use it to create a device extent, as described in *Extents* (page 185).

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 189) for details.

9.5.1 Target Global Configuration

Sharing → *Block (iSCSI)* → *Target Global Configuration*, shown in Figure 9.14, contains settings that apply to all iSCSI shares. Table 9.6 summarizes the settings that are configured in the Target Global Configuration screen.

Some built-in values affect iSNS usage. Fetching of allowed initiators from iSNS is not implemented, so target ACLs must be configured manually. To make iSNS registration useful, iSCSI targets should have explicitly configured port IP addresses. This avoids initiators attempting to discover unconfigured target portal addresses like *0.0.0.0*.

The iSNS registration period is *900* seconds. Registered Network Entities not updated during this period are unregistered. The timeout for iSNS requests is *5* seconds.

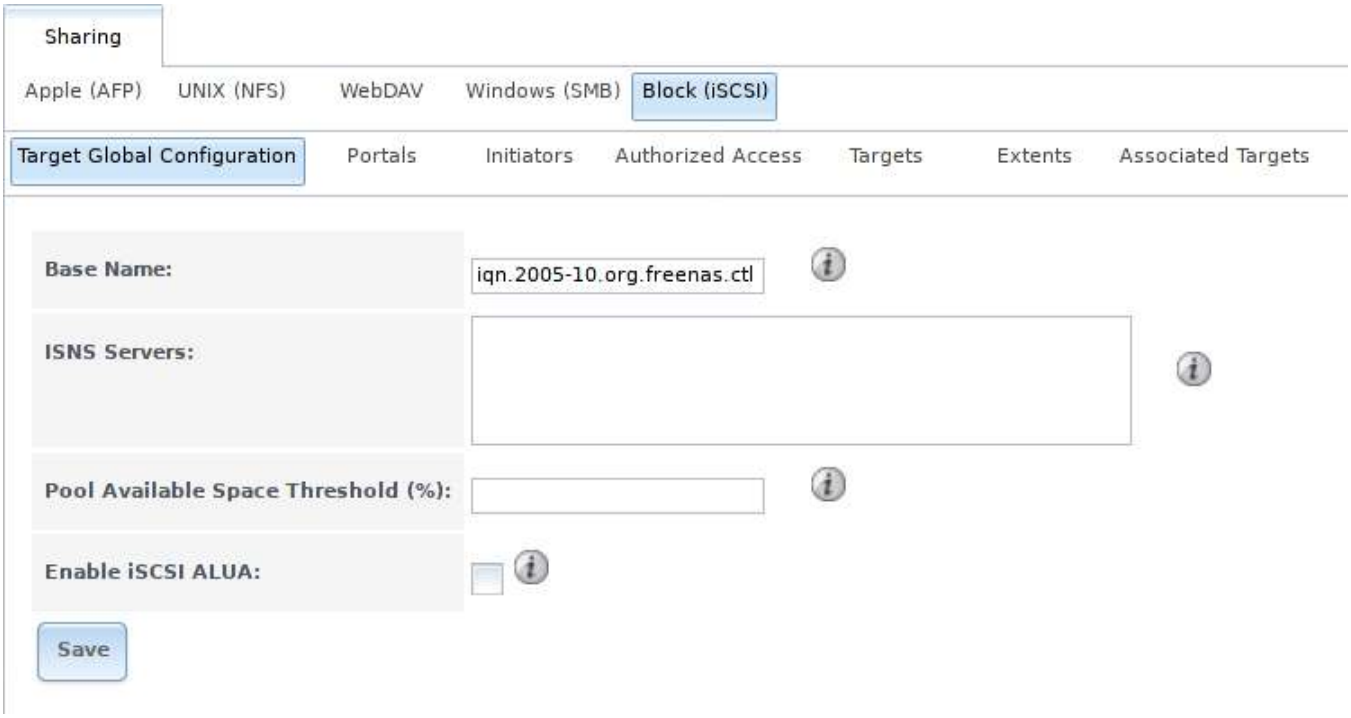

Fig. 9.14: iSCSI Target Global Configuration Variables

Table 9.6: Target Global Configuration Settings

Setting	Value	Description
Base Name	string	Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the "Constructing iSCSI names using the iqn. format" section of **RFC 3721** (https://tools.ietf.org/html/rfc3721.html).
ISNS Servers	string	Enter the hostnames or IP addresses of ISNS servers to be registered with iSCSI targets and portals of the system. Separate each entry with a space.
Pool Available Space Threshold	integer	Enter the percentage of free space to remain in the pool. When this percentage is reached, the system issues an alert, but only if zvols are used. See *VAAI* (page 254) Threshold Warning for more information.
Enable iSCSI ALUA	checkbox	Enable ALUA for automatic best path discovery when supported by clients. This option is only available on HA systems.

9.5.2 Portals

A portal specifies the IP address and port number to be used for iSCSI connections. *Sharing* → *Block (iSCSI)* → *Portals* → *Add Portal* brings up the screen shown in Figure 9.15.

Table 9.15 summarizes the settings that can be configured when adding a portal. To assign additional IP addresses to the portal, click the link *Add extra Portal IP*.

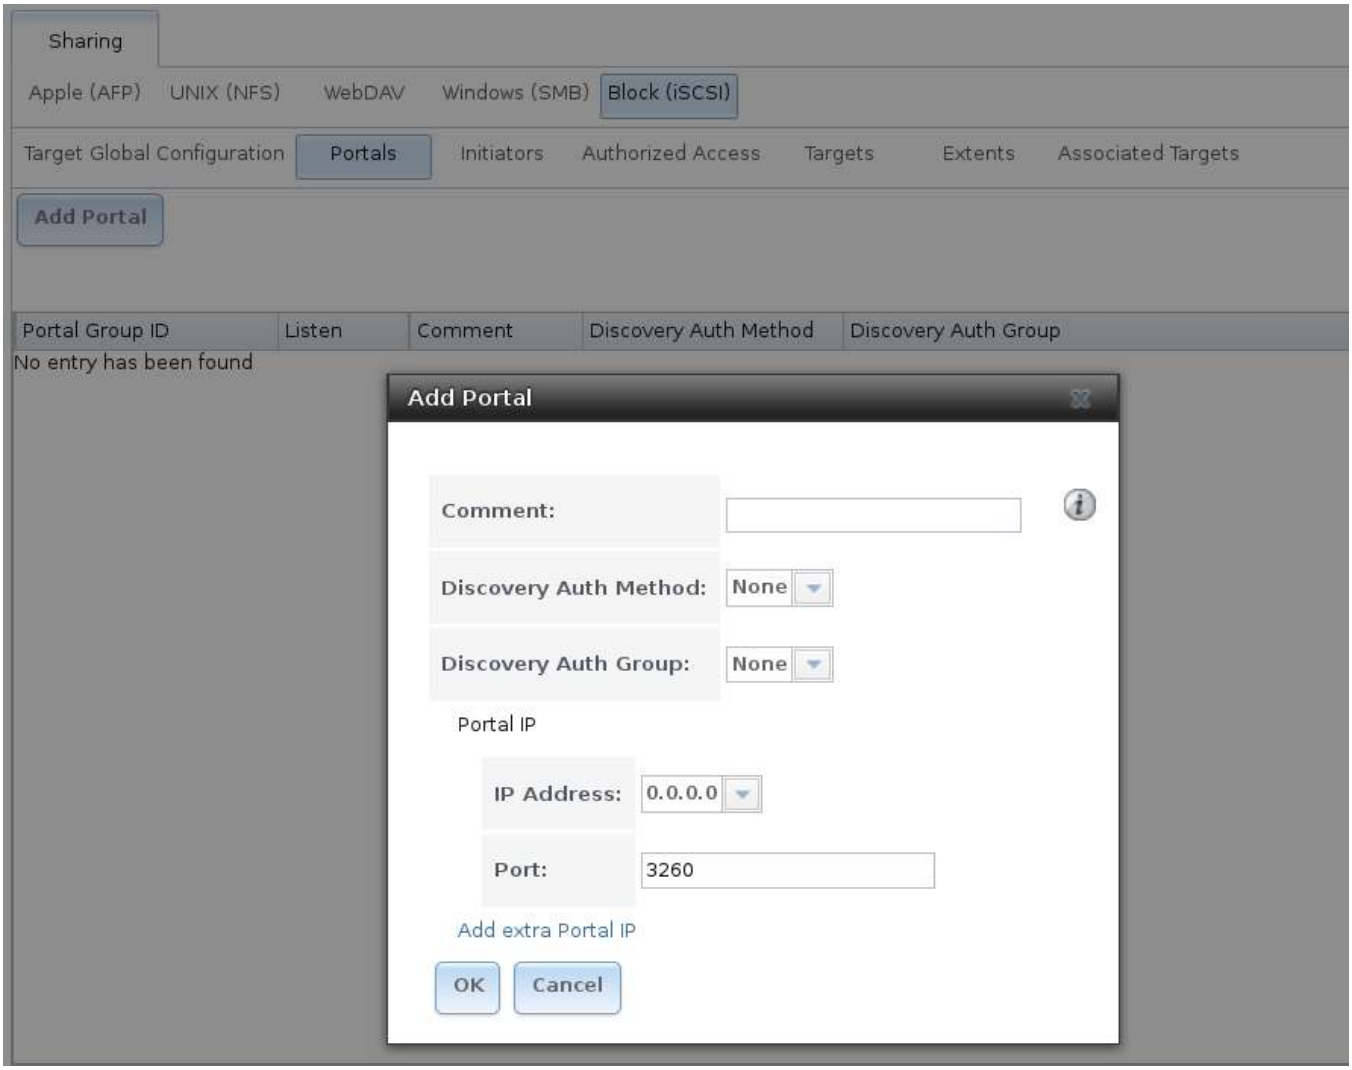

Fig. 9.15: Adding an iSCSI Portal

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	*iSCSI* (page 210) supports multiple authentication methods that are used by the target to discover valid devices. *None* allows anonymous discovery while *CHAP* and *Mutual CHAP* both require authentication.
Discovery Auth Group	drop-down menu	Select a user created in *Authorized Access* if the *Discovery Auth Method* is set to *CHAP* or *Mutual CHAP*.

Continued on next page

Table 9.7 – continued from previous page

Setting	Value	Description
IP address	drop-down menu	Select the IPv4 or IPv6 address associated with an interface or the wildcard address of *0.0.0.0* (any interface). Choose only physical interface IP addresses when configuring iSCSI ALUA. Do not use Virtual IP addresses with an ALUA configuration.
Port	integer	TCP port used to access the iSCSI target. Default is *3260*.

TrueNAS® systems with multiple IP addresses or interfaces can use a portal to provide services on different interfaces or subnets. This can be used to configure multi-path I/O (MPIO). MPIO is more efficient than a link aggregation.

If the TrueNAS® system has multiple configured interfaces, portals can also be used to provide network access control. For example, consider a system with four interfaces configured with these addresses:

192.168.1.1/24

192.168.2.1/24

192.168.3.1/24

192.168.4.1/24

A portal containing the first two IP addresses (group ID 1) and a portal containing the remaining two IP addresses (group ID 2) could be created. Then, a target named A with a Portal Group ID of 1 and a second target named B with a Portal Group ID of 2 could be created. In this scenario, the iSCSI service would listen on all four interfaces, but connections to target A would be limited to the first two networks and connections to target B would be limited to the last two networks.

Another scenario would be to create a portal which includes every IP address **except** for the one used by a management interface. This would prevent iSCSI connections to the management interface.

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 9.16.

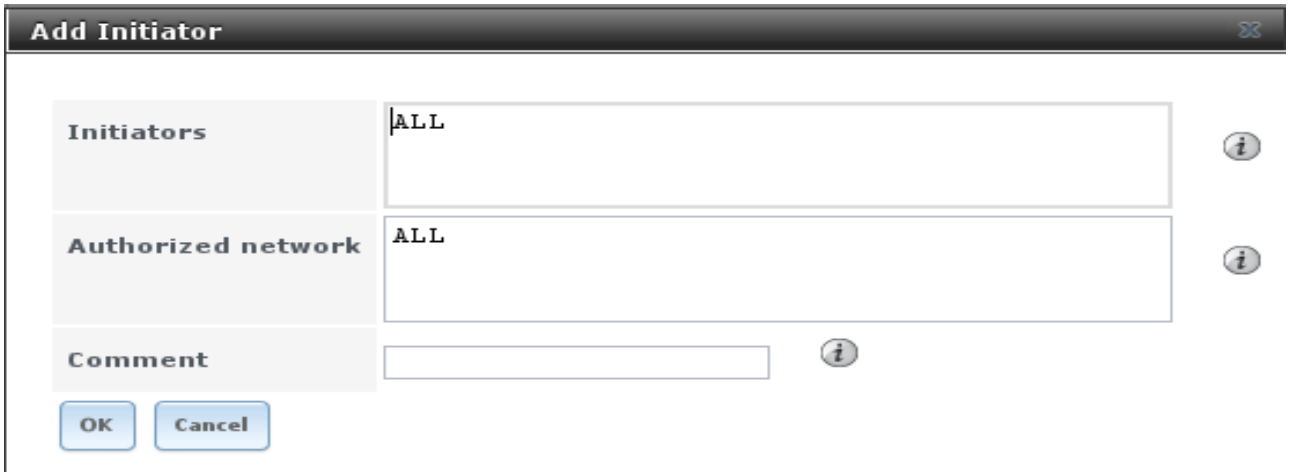

Fig. 9.16: Adding an iSCSI Initiator

Table 9.8 summarizes the settings that can be configured when adding an initiator.

Table 9.8: Initiator Configuration Settings

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

In the example shown in Figure 9.17, two groups are created. Group 1 allows connections from any initiator on any network. Group 2 allows connections from any initiator on the *10.10.1.0/24* network. Click an initiator's entry to display its *Edit* and *Delete* buttons.

Note: Attempting to delete an initiator causes a warning that indicates if any targets or target/extent mappings depend upon the initiator. Confirming the delete causes these to be deleted also.

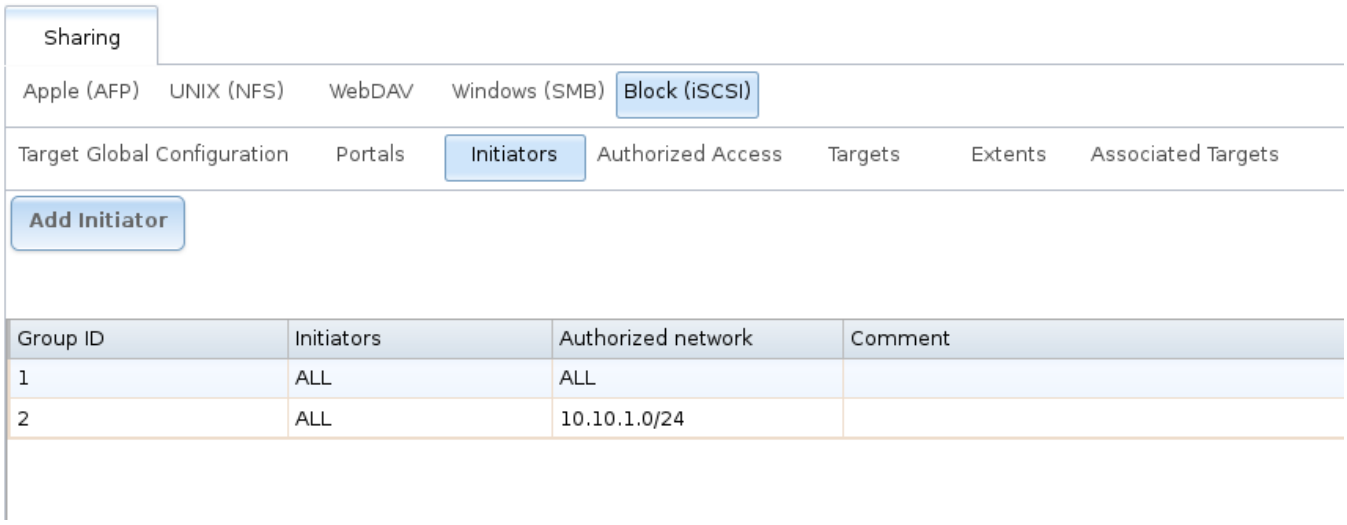

Fig. 9.17: Sample iSCSI Initiator Configuration

9.5.4 Authorized Accesses

When using CHAP or mutual CHAP to provide authentication, creating an authorized access in *Sharing* → *Block (iSCSI)* → *Authorized Accesses* → *Add Authorized Access* is recommended. This screen is shown in Figure 9.18.

Note: This screen sets login authentication. This is different from discovery authentication which is set in *Target Global Configuration* (page 179).

Fig. 9.18: 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	Allow different groups to be configured with different authentication profiles. Example: enter *1* for all users in Group *1* to inherit the Group *1* authentication profile. Group IDs that are already configured with authorized access cannot be reused.
User	string	Enter name of user account to create for CHAP authentication with the user on the remote system. Many initiators default to using the initiator name as the user.
Secret	string	Enter and confirm a password for *User*. Must be between 12 and 16 characters.
Peer User	string	Only input when configuring mutual CHAP. In most cases it will need to be the same value as *User*.
Peer Secret	string	Enter and confirm the mutual secret password which **must be different than the** *Secret*. Required if *Peer User* is set.

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

As authorized accesses are added, they will be listed under *View Authorized Accesses*. In the example shown in Figure 9.19, three users (*test1*, *test2*, and *test3*) and two groups (*1* and *2*) are 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.

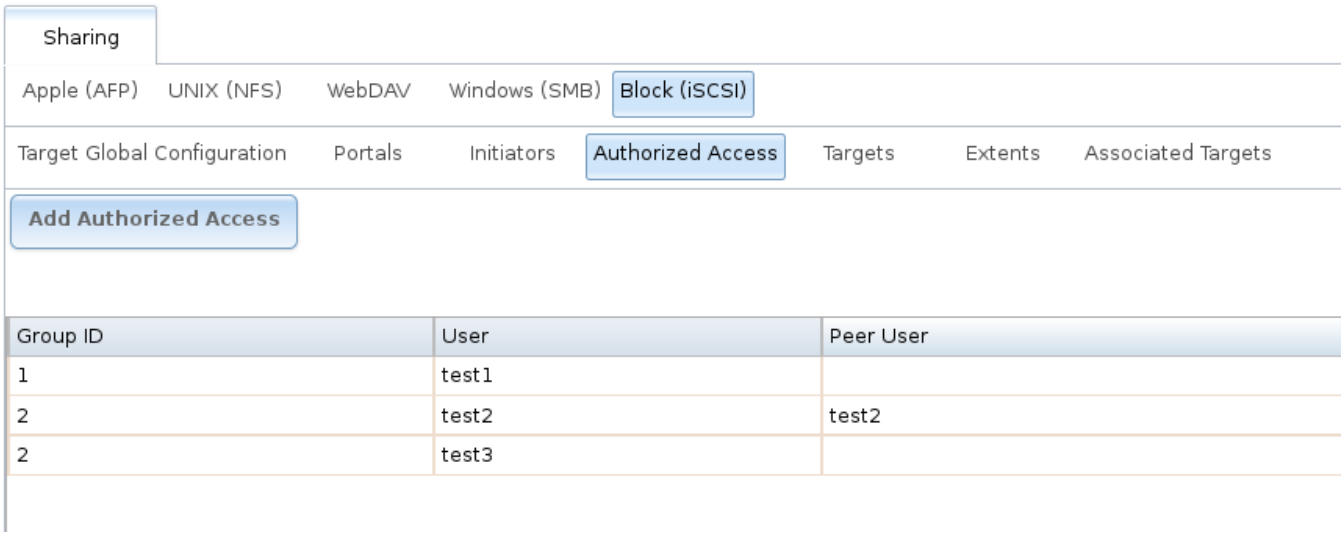

Fig. 9.19: Viewing Authorized Accesses

9.5.5 Targets

Next, create a Target using *Sharing* → *Block (iSCSI)* → *Targets* → *Add Target*, as shown in Figure 9.20. 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).

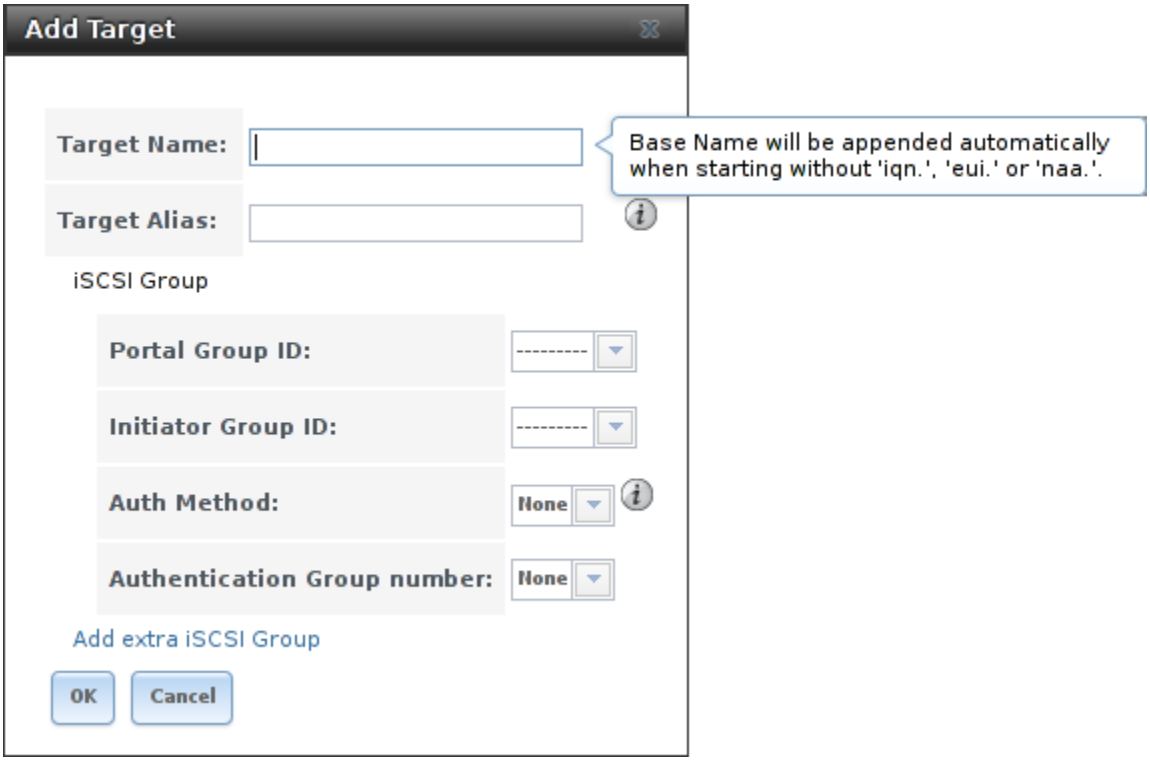

Fig. 9.20: Adding an iSCSI Target

Table 9.10: Target Settings

Setting	Value	Description
Target Name	string	Required. The base name is automatically prepended if the target name does not start with *iqn*. Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the "Constructing iSCSI names using the iqn. format" section of **RFC 3721** (https://tools.ietf.org/html/rfc3721.html).
Target Alias	string	Enter an optional user-friendly name.
Portal Group ID	drop-down menu	Leave empty or select number of existing portal to use.
Initiator Group ID	drop-down menu	Select which existing initiator group has access to the target.
Auth Method	drop-down menu	Choices are: *None*, *Auto*, *CHAP*, or *Mutual CHAP*.
Authentication Group number	drop-down menu	Select *None* or an integer. This number represents the number of existing authorized accesses.

9.5.6 Extents

iSCSI targets provide virtual access to resources on the TrueNAS® system. *Extents* are used to define resources to share with clients. There are two types of extents: *device* and *file*.

Device extents provide virtual storage access to zvols, zvol snapshots, or physical devices like a disk, an SSD, a hardware RAID volume, or a HAST device (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-

hast.html).

File extents provide virtual storage access to an individual file.

Tip: For typical use as storage for virtual machines where the virtualization software is the iSCSI initiator, device extents with zvols provide the best performance and most features. For other applications, device extents sharing a raw device can be appropriate. File extents do not have the performance or features of device extents, but do allow creating multiple extents on a single filesystem.

Virtualized zvols support all the TrueNAS® *VAAI* (page 254) primitives and are recommended for use with virtualization software as the iSCSI initiator.

The ATS, WRITE SAME, XCOPY and STUN, primitives are supported by both file and device extents. The UNMAP primitive is supported by zvols and raw SSDs. The threshold warnings primitive is fully supported by zvols and partially supported by file extents.

Virtualizing a raw device like a single disk or hardware RAID volume limits performance to the abilities of the device. Because this bypasses ZFS, such devices do not benefit from ZFS caching or provide features like block checksums or snapshots.

Virtualizing a zvol adds the benefits of ZFS, such as read and write cache. Even if the client formats a device extent with a different filesystem, the data still resides on a ZFS volume and benefits from ZFS features like block checksums and snapshots.

Warning: For performance reasons and to avoid excessive fragmentation, keep the used space of the pool below 80% when using iSCSI. The capacity of an existing extent can be increased as shown in *Growing LUNs* (page 192).

To add an extent, go to *Sharing* → *Block (iSCSI)* → *Extents* → *Add Extent*. In the example shown in Figure 9.21, 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 fails when the name of the file to be created to the volume/dataset name.** is not appended.

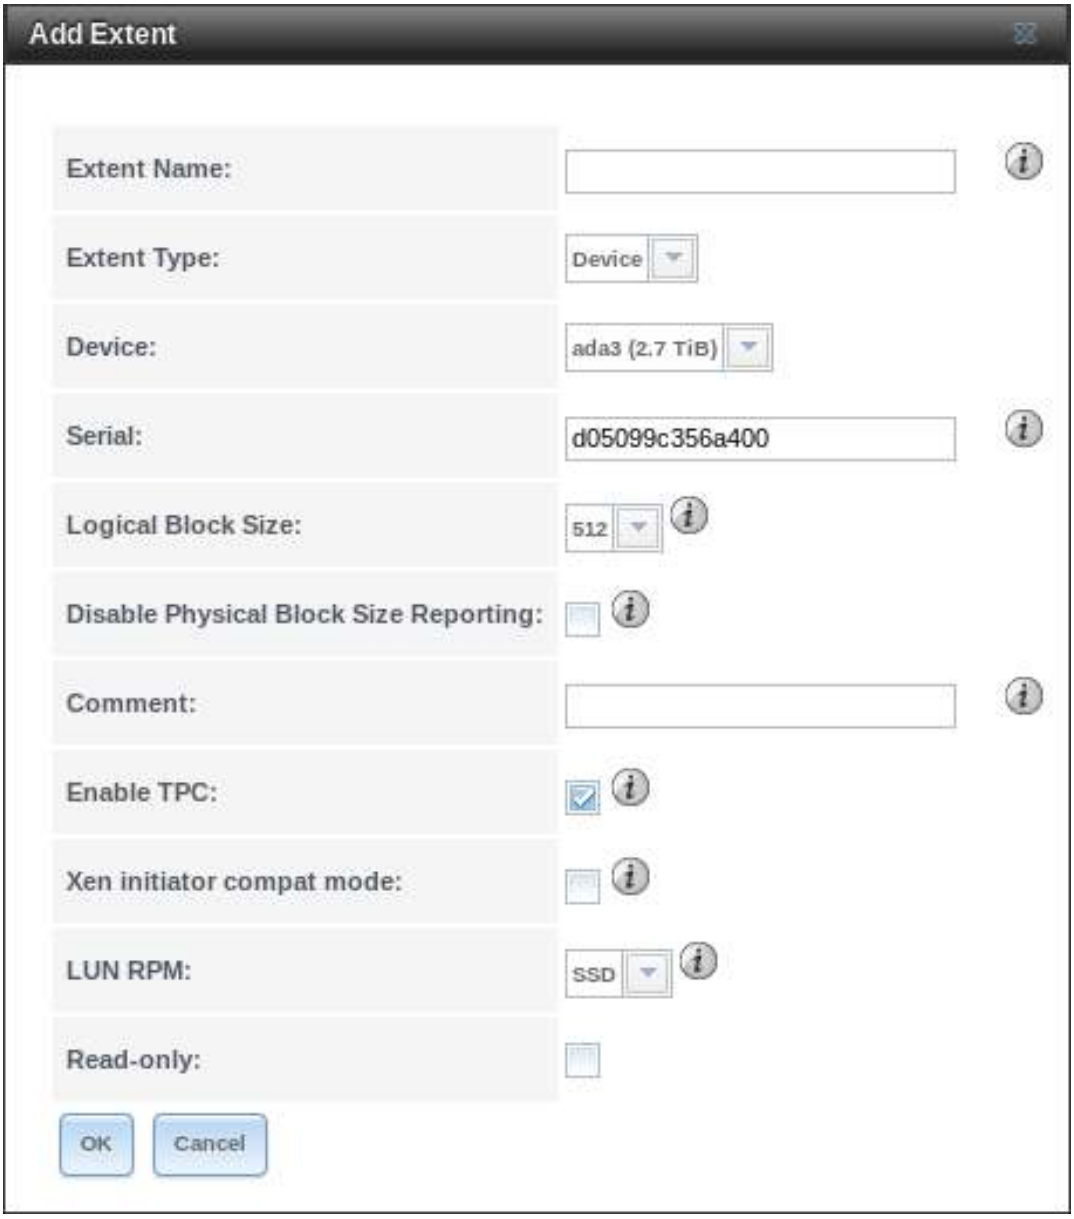

Fig. 9.21: Adding an iSCSI Extent

Table 9.11: Extent Configuration Settings

Setting	Value	Description
Extent Name	string	Enter the extent name. If the *Extent size* is not *0*, it cannot be an existing file within the volume/dataset.
Extent Type	drop-down menu	Select from *File* or *Device*.
Device	drop-down menu	Only appears if *Device* is selected. Select the unformatted disk, controller, zvol, zvol snapshot, or HAST device.
Serial	string	Unique LUN ID. The default is generated from the system MAC address.

Continued on next page

Table 9.11 – continued from previous page

Setting	Value	Description
Path to the extent	browse button	Only appears if *File* is selected. Browse to an existing file and use *0* as the *Extent size*, **or** browse to the volume or dataset, click *Close*, append the *Extent Name* to the path, and specify a value in *Extent size*. Extents cannot be created inside the jail root directory.
Extent size	integer	Only appears if *File* is selected. If the size is specified as *0*, the file must already exist and the actual file size will be used. Otherwise, specify the size of the file to create.
Logical Block Size	drop-down menu	Leave at the default of 512 unless the initiator requires a different block size.
Disable Physical Block Size Reporting	checkbox	Set if the initiator does not support physical block size values over 4K (MS SQL). Setting can also prevent constant block size warnings (https://www.virten.net/2016/12/the-physical-block-size-reported-by-the-device-is-not-supported/) when using this share with ESXi.
Available Space Threshold	string	Only appears if *File* or a zvol is selected. When the specified percentage of free space is reached, the system issues an alert. See *VAAI* (page 254) Threshold Warning for more information.
Comment	string	Enter an optional comment.
Enable TPC	checkbox	If enabled, an initiator can bypass normal access control and access any scannable target. This allows `xcopy` operations otherwise blocked by access control.
Xen initiator compat mode	checkbox	Set this option when using Xen as the iSCSI initiator.
LUN RPM	drop-down menu	Do **NOT** change this setting when using Windows as the initiator. Only needs to be changed in large environments where the number of systems using a specific RPM is needed for accurate reporting statistics.
Read-only	checkbox	Set 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 9.22. Use the drop-down menus to select the existing target and extent. Click *OK* to add an entry for the LUN.

Fig. 9.22: 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 an existing target.
LUN ID	integer	Select or enter a value between *0* and *1023*. Some initiators expect a value less than *256*. Use unique LUN IDs for each associated target.
Extent	drop-down menu	Select an existing extent.

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

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 9.23.

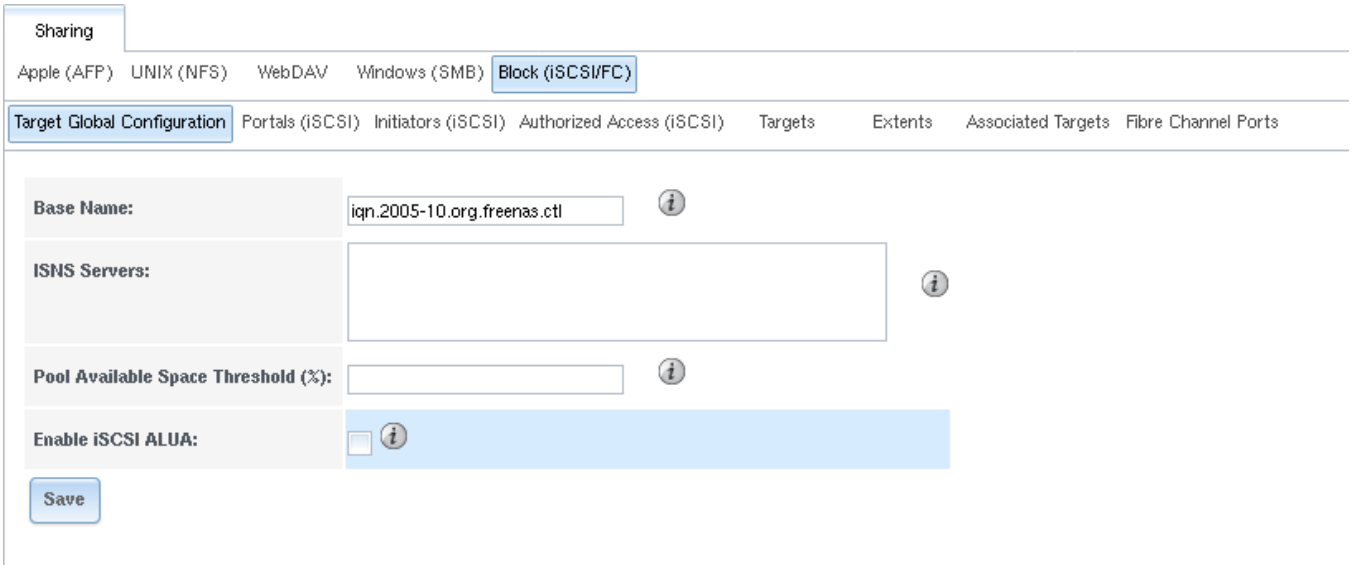

Fig. 9.23: Block (iSCSI) Screen

Otherwise, the *Target Global Configuration* screen is the same as described in *Target Global Configuration* (page 179).

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 9.24, the *Targets* → *Add Target* screen has an extra *Target Mode* option for indicating whether the target to create is iSCSI, Fibre Channel, or both.

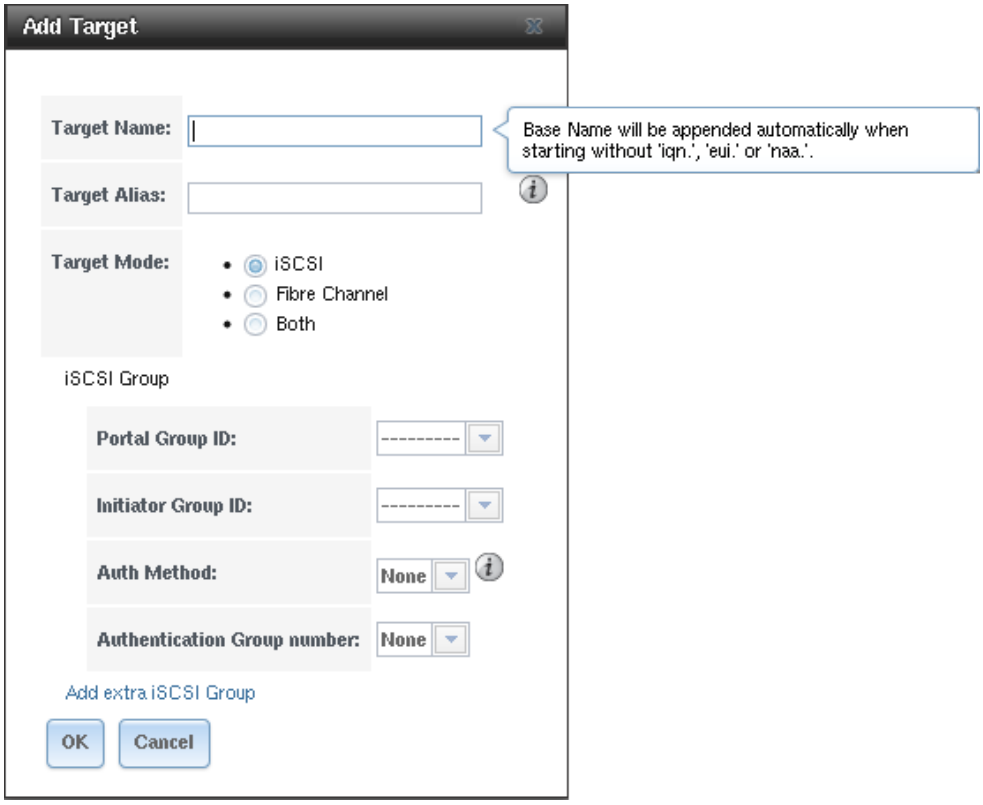

Fig. 9.24: Add Target Screen

After selecting *Fibre Channel*, this screen changes 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 9.25.

Fig. 9.25: Configuring a Fibre Channel Target

The screens for adding an extent and associating a target are the same as described in *Extents* (page 185) and *Target/Extents* (page 188).

An example of the *Fibre Channel Ports* screen is shown in Figure 9.26.

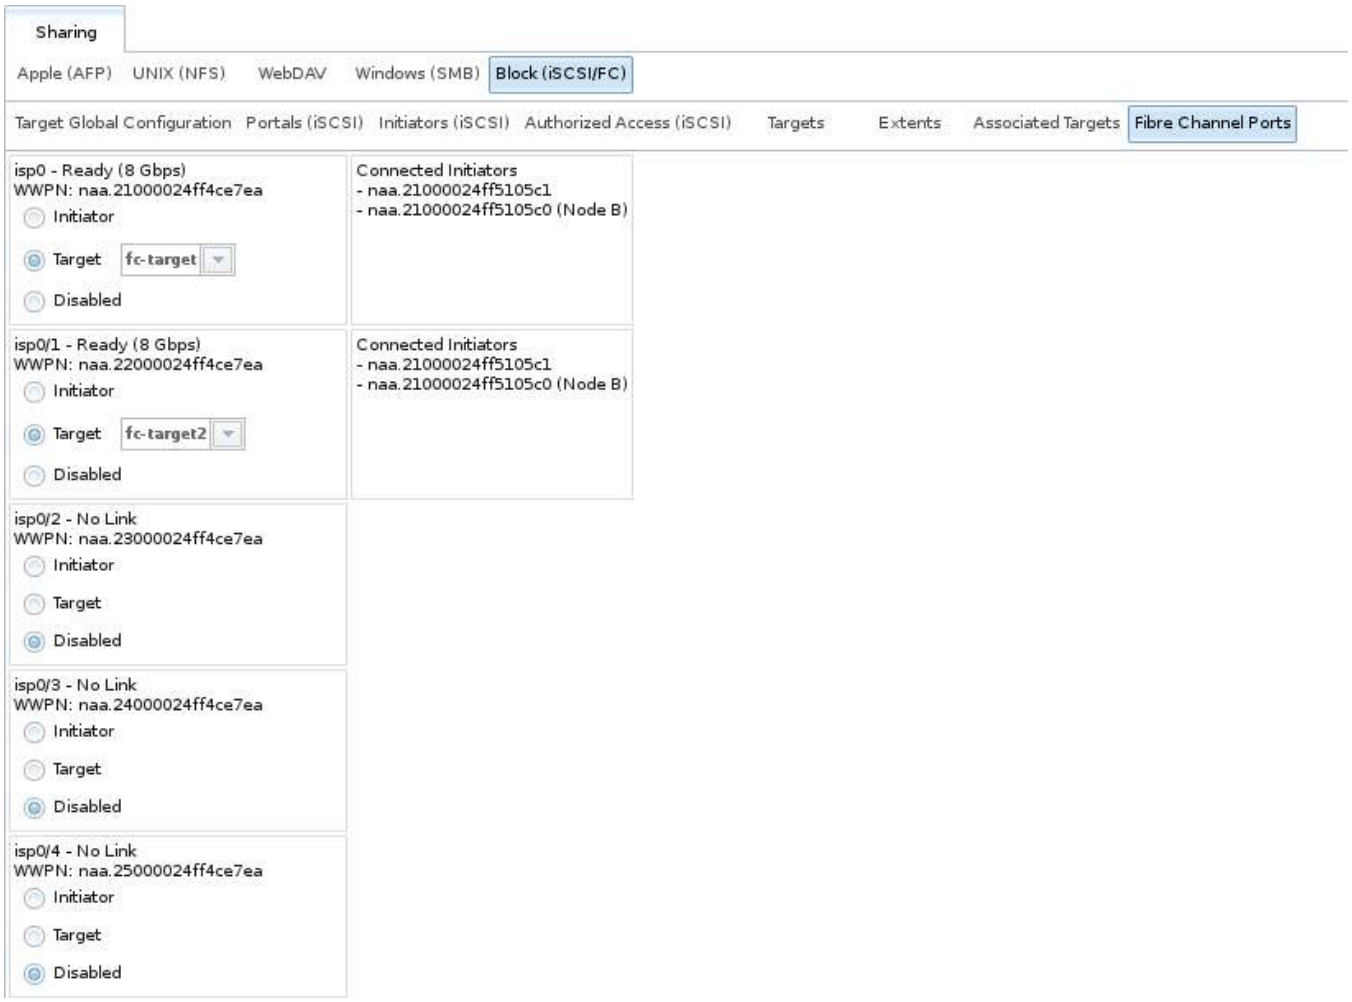

Fig. 9.26: Configuring a Fibre Channel Port

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 235) 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 9.26:. NPIV allows the administrator to use switch zoning to configure each virtual port as if it was a physical port in order to provide access control. This is important in an environment with a mix of Windows systems and virtual machines in order to prevent automatic or accidental reformatting of targets containing unrecognized filesystems. It can also be used to segregate data; for example, to prevent the engineering department from accessing data from the human resources department. Refer to the switch documentation for details on how to configure zoning of virtual ports.

To create the virtual ports on the TrueNAS® system, go to *System* → *Tunables* → *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 then 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 9.27, 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 235) where its bandwidth usage can be viewed.

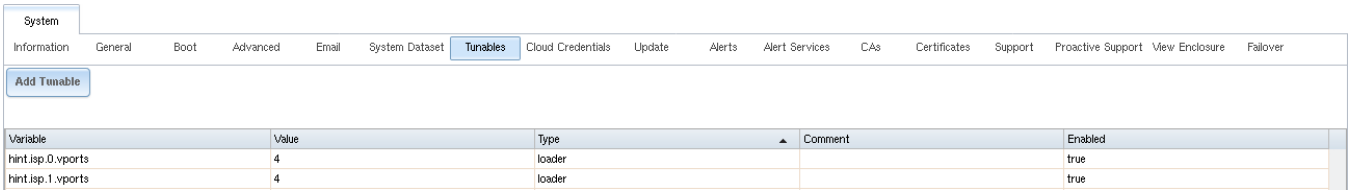

Fig. 9.27: Adding Virtual Ports

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://techgenix.com/Connecting-Windows-7-iSCSI-SAN/). A client for Windows 2000, XP, and 2003 can be found here (http://www.microsoft.com/en-us/download/details.aspx?id=18986). This how-to (https://www.pluralsight.com/blog/software-development/freenas-8-iscsi-target-windows-7) shows how to create an iSCSI target for a Windows 7 system.

macOS does not include an initiator. globalSAN (http://www.studionetworksolutions.com/globalsan-iscsi-initiator/) is a commercial, easy-to-use Mac initiator.

BSD systems provide command line initiators: iscontrol(8) (https://www.freebsd.org/cgi/man.cgi?query=iscontrol) comes with FreeBSD versions 9.x and lower, iscsictl(8) (https://www.freebsd.org/cgi/man.cgi?query=iscsictl) comes with FreeBSD versions 10.0 and higher, iscsi-initiator(8) (http://netbsd.gw.com/cgi-bin/man-cgi?iscsi-initiator++NetBSD-current) comes with NetBSD, and iscsid(8) (http://man.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man8/iscsid.8?query=iscsid) comes with OpenBSD.

Some Linux distros provide the command line utility `iscsiadm` from Open-iSCSI (http://www.open-iscsi.com/). Use a web search to see if a package exists for the distribution should the command not exist on the Linux system.

If a LUN is added while `iscsiadm` is already connected, it will not see the new LUN until rescanned with `iscsiadm -m node -R`. Alternately, use `iscsiadm -m discovery -t st -p portal_IP` to find the new LUN and `iscsiadm -m node -T LUN_Name -l` to log into the LUN.

Instructions for connecting from a VMware ESXi Server can be found at How to configure FreeNAS 8 for iSCSI and connect to ESX(i) (https://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/). Note that the requirements for booting vSphere 4.x off iSCSI differ between ESX and ESXi. ESX requires a hardware iSCSI adapter while ESXi requires specific iSCSI boot firmware support. The magic is on the booting host side, meaning that there is no difference to the TrueNAS® configuration. See the iSCSI SAN Configuration Guide (https://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) 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.

9.5.10.1 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 9.28, the current size of the zvol named *zvol1* is 10 GiB.

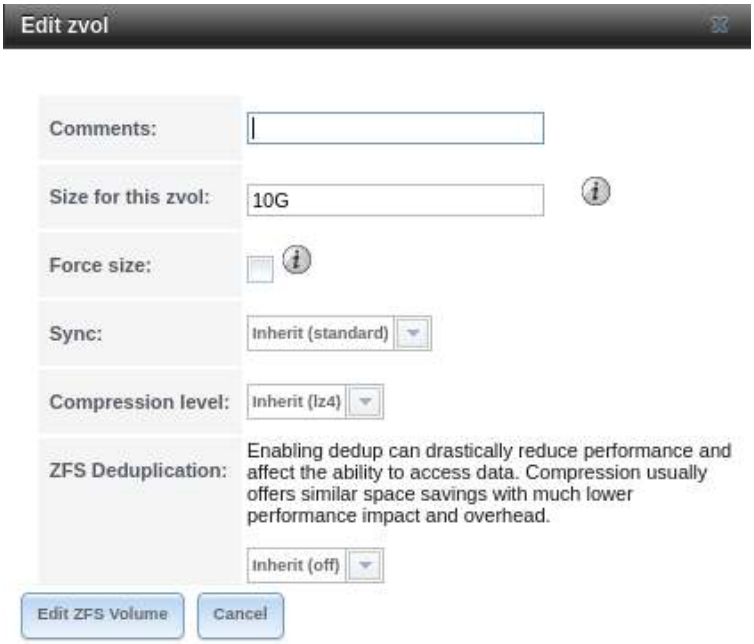

Fig. 9.28: Editing an Existing Zvol

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 web interface 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.

9.5.10.2 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.

9.6 Creating Authenticated and Time Machine Shares

macOS includes the Time Machine feature which performs automatic back ups. TrueNAS® supports Time Machine backups for both *SMB* (page 166) and *AFP* (page 154) shares. This section has instructions to create Time Machine SMB and AFP shares, using the *Wizard* to create an AFP Time Machine share. The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

9.6.1 Manual Creation of Authenticated or Time Machine Shares

Create Time Machine and authenticated shares on a *new dataset* (page 102).

Change permissions on the new dataset by going to *Storage* → *Volumes*. Select the dataset and click *Change Permissions*. Enter these settings:

1. **Permission Type:** Select *Mac*.
2. **Owner (user):** Use the drop-down to select the desired user account. If the user does not yet exist on the TrueNAS® system, create one with *Account* → *Users*. See *users* (page 19) for more information.
3. **Owner (group):** Select the desired group name. If the group does not yet exist on the TrueNAS® system, create one with *Account* → *Groups*. See *groups* (page 16) for more information.
4. Click *Change*.

Create the authenticated or Time Machine share:

1. Go to *Sharing* → *Windows (SMB)* or *Sharing* → *Apple (AFP)* and click *Add Share*. Apple deprecated the AFP protocol (https://support.apple.com/en-us/HT207828) and recommends using SMB.
2. *Browse* to the dataset created for the share.
3. When creating a Time Machine share, set the *Time Machine* option.
4. Fill out the other required fields.
5. Click *OK*.

9.6.2 Create AFP Time Machine Share with the Wizard

To use the Wizard to create an AFP authenticated or Time Machine share, enter the following information, as seen in the example in Figure 9.29.

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 enable the *Time Machine* option.
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 enable the *Create User* option. If the user is 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 set the *Create Group* option. Otherwise, enter the same name as the user. In the example shown in Figure 9.30, both a new *user1* user and a new *tm_backups* group are 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 9.29.
4. Click the *Add* button.

When creating multiple authenticated or Time Machine shares, repeat this process for each user. Give 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 creates a dataset for each share with the correct ownership and starts the AFP service so the shares are immediately available. The new shares appear in *Sharing* → *Apple (AFP)*.

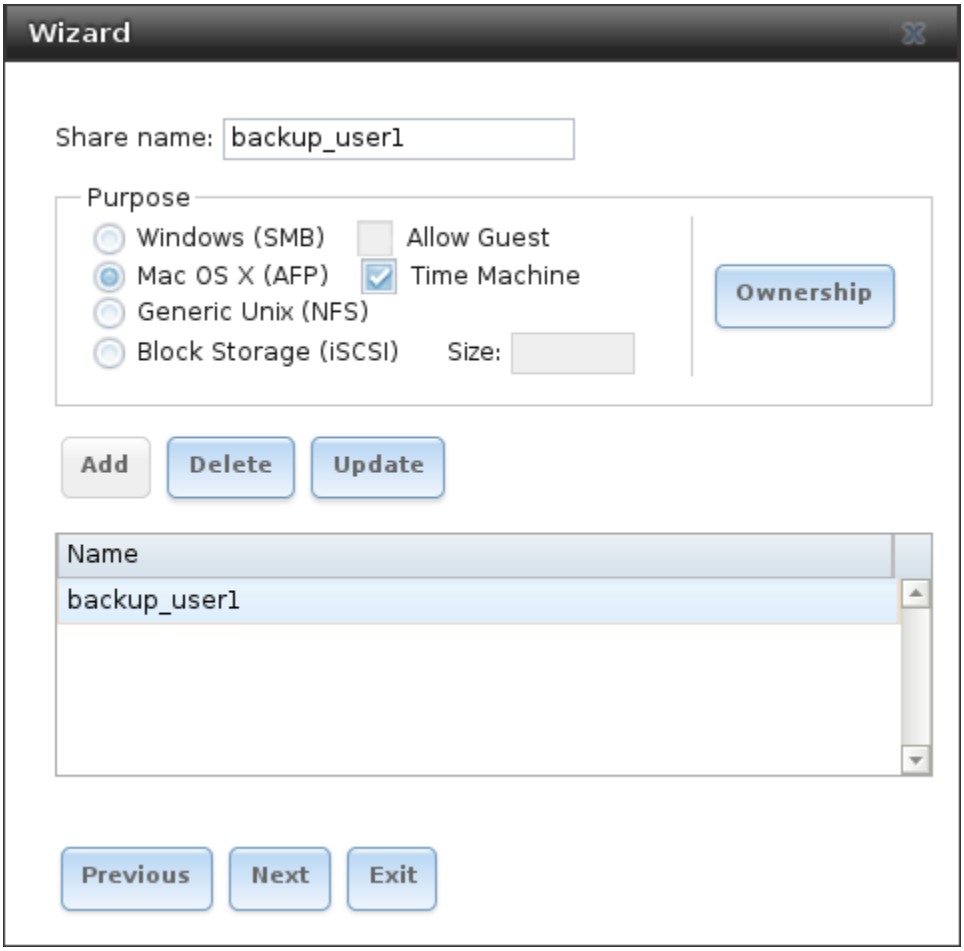

Fig. 9.29: Creating a Time Machine Share

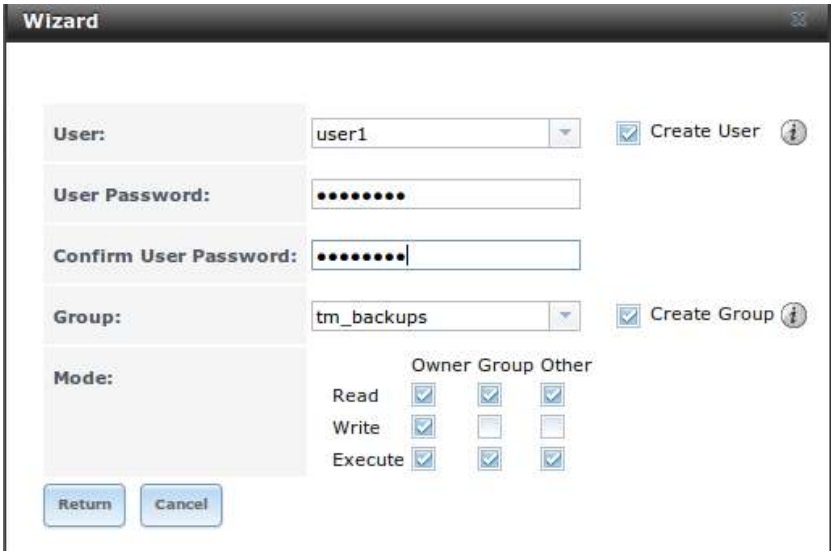

Fig. 9.30: Creating an Authenticated User

9.6.3 Configuring Time Machine Backups

Configuring a quota for each Time Machine share helps prevent backups from using all available space on the TrueNAS® system. Time Machine creates ongoing hourly, daily, weekly, and monthly backups. **The oldest backups are deleted when a Time Machine share fills up, so make sure that the quota size is large enough to hold the desired number of backups.** Note that a default installation of macOS is over 20 GiB.

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

To configure a quota, go to *Storage* → *Volumes* and select the share dataset. In the example shown in Figure 9.31, 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 GiB.

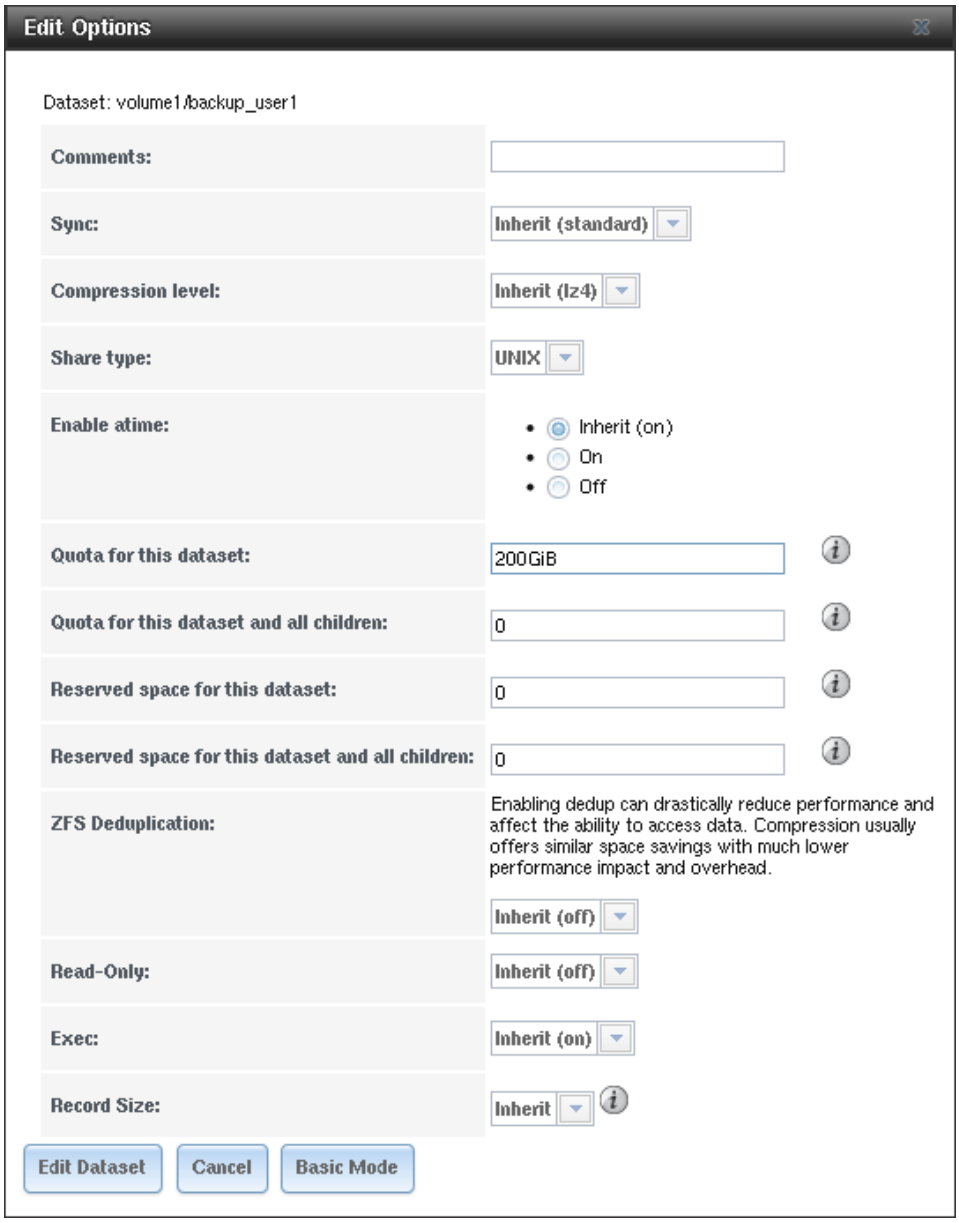

Fig. 9.31: Setting a Quota

Note: The example shown here is intended to show the general process of adding a TrueNAS® share in Time Machine. The example might not reflect the exact process to configure Time Machine on a specific version of macOS. See the Apple documentation (https://support.apple.com/en-us/HT201250) for detailed Time Machine configuration instructions.

To configure Time Machine on the macOS client, go to *System Preferences* → *Time Machine*, and click *ON* in the left panel.

Fig. 9.32: Configuring Time Machine on Mac OS X Lion

Click *Select Disk...* in the right panel to find the TrueNAS® system with the share. Highlight the share and click *Use Backup Disk*. A connection dialog prompts to log in to the TrueNAS® system.

If `Time Machine could not complete the backup. The backup disk image could not be created (error 45)` is shown when backing up to the TrueNAS® system, a sparsebundle image must be created using these instructions (https://community.netgear.com/t5/Stora-Legacy/Solution-to-quot-Time-Machine-could-not-complete-the-backup/td-p/294697).

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) to avoid making another backup or losing past backups.

CHAPTER

TEN

SERVICES

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

- *AFP* (page 200)
- *Asigra DS-System* (page 201)
- *Domain Controller* (page 202)
- *Dynamic DNS* (page 204)
- *FTP* (page 205)
- *iSCSI* (page 210)
- *LLDP* (page 211)
- *Netdata* (page 211)
- *NFS* (page 213)
- *Rsync* (page 214)
- *S3* (page 216)
- *S.M.A.R.T.* (page 218)
- *SMB* (page 219)
- *SNMP* (page 223)
- *SSH* (page 225)
- *TFTP* (page 227)
- *UPS* (page 229)
- *WebDAV* (page 232)

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, lists all services. It also shows where to start, stop, or configure the available services. The S.M.A.R.T. service is enabled by default, but only runs if the storage devices support S.M.A.R.T. data (https://en.wikipedia.org/wiki/S.M.A.R.T.) Other services default to off until started.

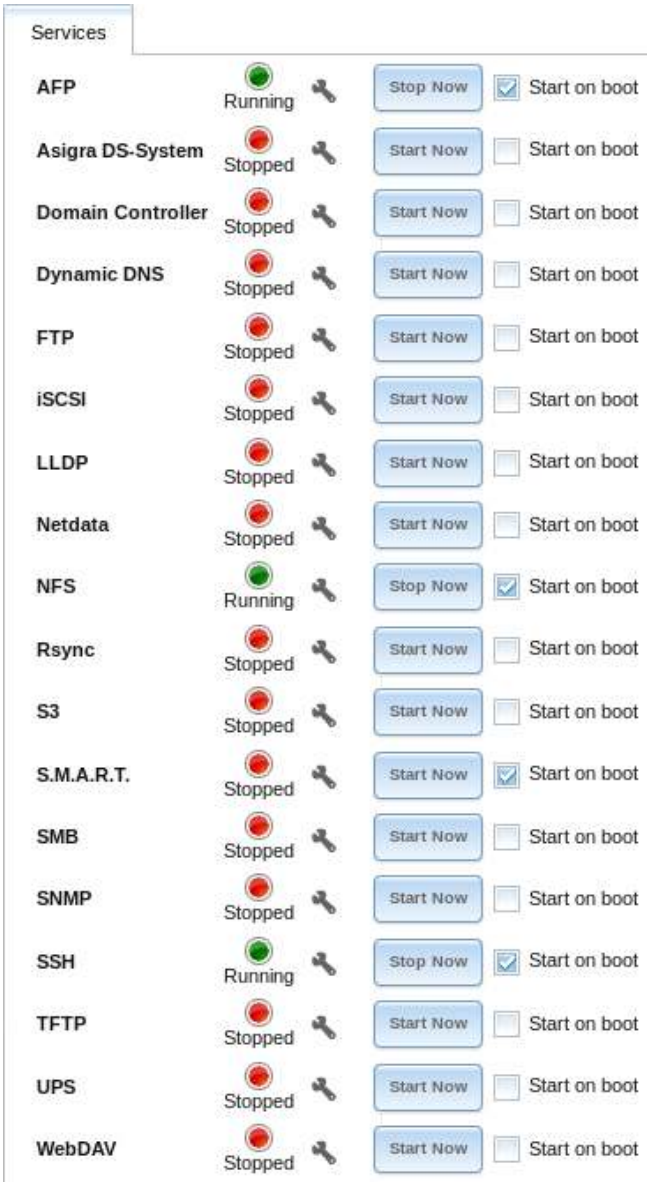

Fig. 10.1: Control Services

Stopped services show a red stop symbol and a *Start Now* button. Running services show a green light with a *Stop Now* button.

Tip: Using a proxy server can prevent the list of services from being displayed. If a proxy server is used, do not configure it to proxy local network connections or websocket connections. VPN software can also cause problems. If the list of services is displayed when connecting on the local network but not when connecting through the VPN, check the VPN software configuration.

Services are configured by clicking the wrench icon or the name of the service in the *Services* section of the tree menu.

If a service does not start, go to *System* → *Advanced* and enable *Show console messages in the footer*. Console messages appear at the bottom of the browser. Clicking the console message area makes it into a pop-up window, allowing scrolling through or copying the messages. Watch these messages for errors when stopping or starting the problematic service.

To read the system logs for more information about a service failure, open *Shell* (page 244) 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.

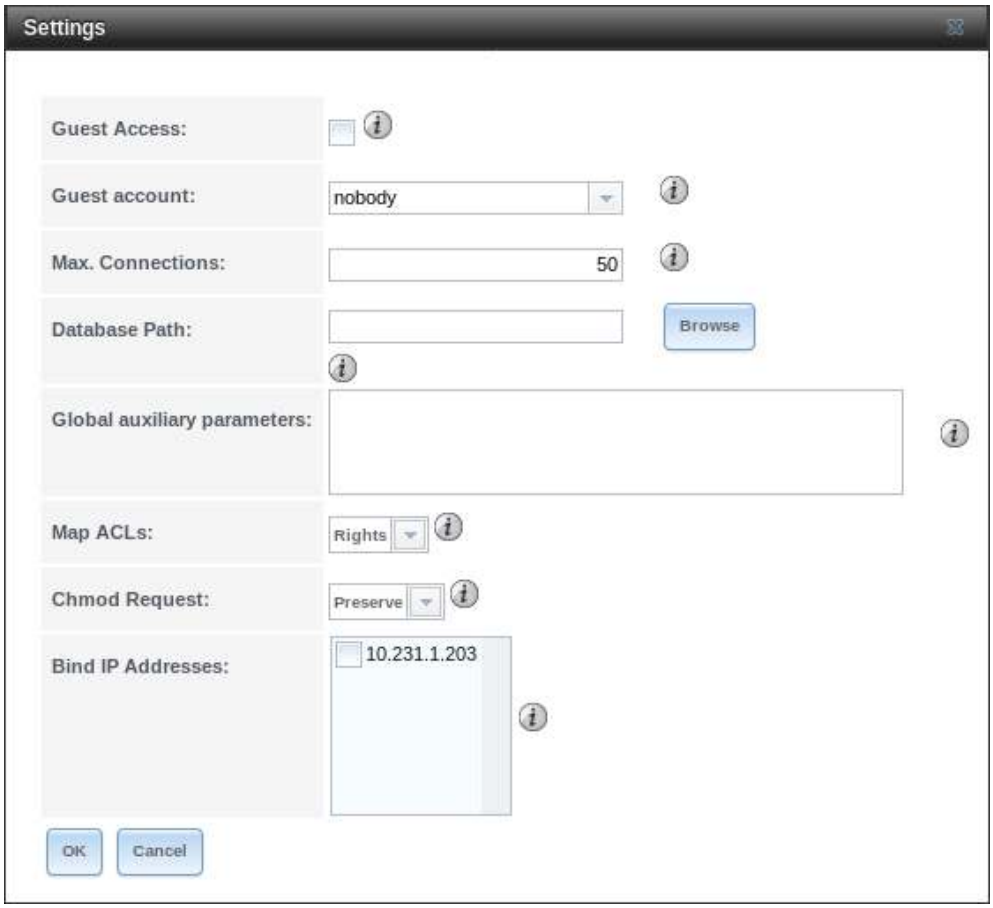

Fig. 10.2: Global AFP Configuration

Table 10.1: Global AFP Configuration Options

Setting	Value	Description
Guest Access	checkbox	Set to disable the password prompt that appears before clients access AFP shares.
Guest account	drop-down menu	Select an account to use for guest access. The account must have permissions to the volume or dataset being shared.
Max Connections	integer	Maximum number of simultaneous connections.
Database Path	browse button	Sets the database information to be stored in the path. Default is the root of the volume. The path must be writable even if the volume is read only.

Continued on next page

Table 10.1 – continued from previous page

Setting	Value	Description
Global auxiliary parameters	string	Add any additional afp.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered elsewhere in this screen.
Map ACLs	drop-down menu	Choose mapping of effective permissions for authenticated users. Choices are: *Rights* (default, Unix-style permissions), *Mode* (ACLs), or *None*
Chmod Request	drop-down menu	Sets how Access Control Lists are handled. *Ignore*: ignores requests and gives the parent directory ACL inheritance full control over new items. *Preserve*: preserves ZFS Access Control Entries for named users and groups or the POSIX ACL group mask. *Simple*: is set to chmod() as requested without any extra steps.
Bind IP Addresses	selection	Specify the IP addresses to listen for FTP connections. Highlight the desired IP addresses in the *Available* list and use the >> button to add to the *Selected* list.

10.2.1 Troubleshooting AFP

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

Determine which users are connected to an AFP share by typing `afpusers`.

If *Something wrong with the volume's CNID DB* is shown, run this command from *Shell* (page 244), replacing the path to the problematic AFP share:

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

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

10.3 Asigra DS-System

Asigra Backup allows administrators to back up data from network-connected computers and mobile devices. Asigra leverages standard API calls from a single on-site Asigra service (DS-Client) to reach into these devices and does not require any agent software on the endpoints to access the data.

Licensed Asigra Backup software can use TrueNAS® as the storage backend.

Note: To learn more about Asigra or to enquire about licensing, contact sales@ixsystems.com.

For the initial backend configuration, click *Services* → *Asigra DS-System*. When prompted to choose the *Base Filesystem*, select the dataset to store the Asigra backups, then click *OK*. Any required database entries are created and the service is started.

Note: Asigra DS-Operator requires a working installation of Java JRE (https://www.oracle.com/technetwork/java/javase/downloads-2133155.html) and a security exception for the TrueNAS® system. To add the exception, use *Configure Java* → *Security* → *Edit Site List* → *Add* and enter the URL to the TrueNAS® system. If the browser prompts to open `DSOP.jnlp` with an application, select Java Web Start Launcher (`javaws`).

While the service is running, the *Launch DS Operator* button appears in *Services* → *Asigra DS-System*. Click *Launch DS Operator* to download and launch the Asigra management application.

Fig. 10.3: Asigra settings

Contact Asigra (https://www.asigra.com/contact-us) for further documentation on using DS Operator.

10.4 Domain Controller

TrueNAS® can be configured to act either as the domain controller for a network or to join an existing *Active Directory* (page 141) network as a domain controller.

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 141) network to access its authentication and authorization services, configure *Active Directory* (page 141) instead.

Note: The Domain Controller service cannot be configured when *Enable Monitoring* is set in *Directory Services* → *Active Directory*

Configuring a domain controller is a complex process that requires a good understanding of how *Active Directory* (page 141) works. While *Services* → *Domain Controller* makes it easy to enter the needed settings into the web interface, it is important to understand what those settings should be. Before beginning configuration, read through the Samba AD DC HOWTO (https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO). After 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 10.4 shows the configuration screen for creating a domain controller and Table 10.2 summarizes the available options.

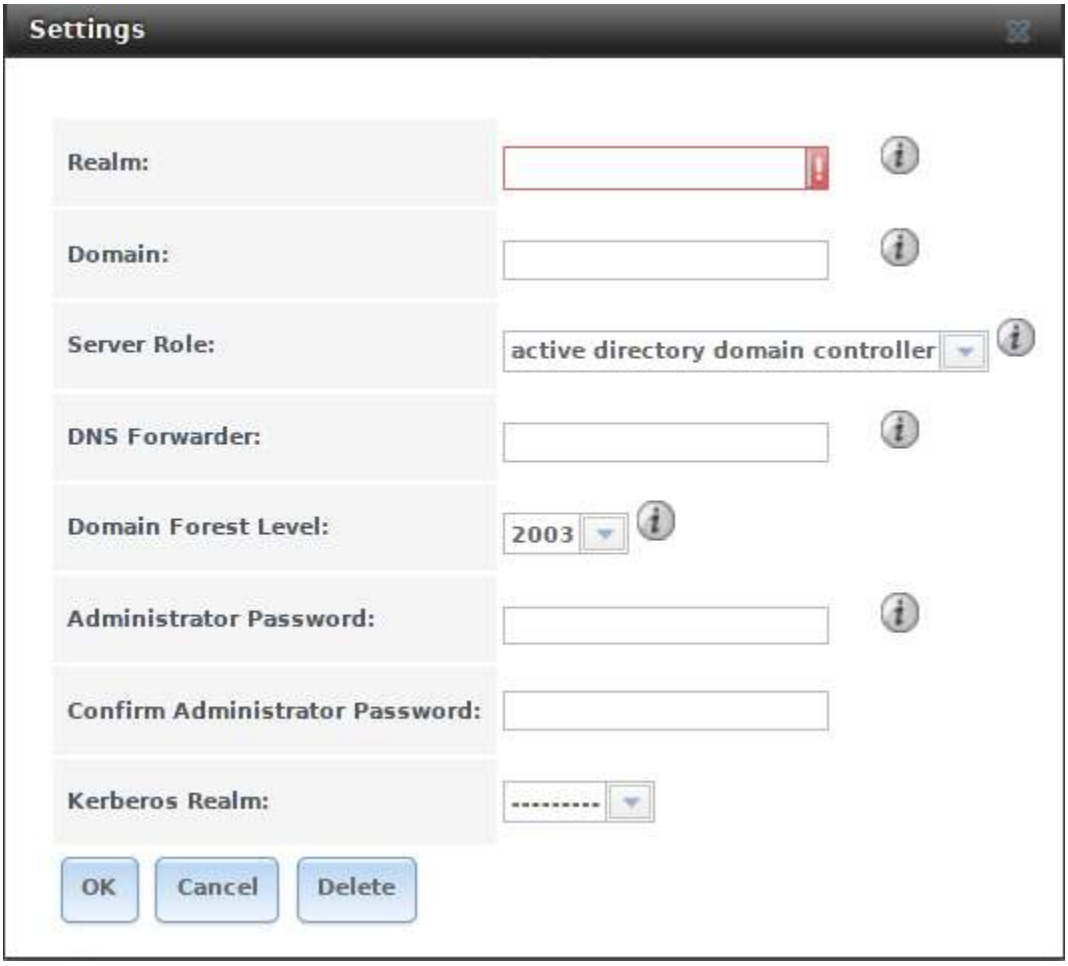

Fig. 10.4: Domain Controller Settings

Table 10.2: Domain Controller Configuration Options

Setting	Value	Description
Realm	string	Enter a capitalized DNS realm name.
Domain	string	Enter a capitalized domain name.
Server Role	drop-down menu	At this time, the only supported role is as the domain controller for a new domain.
DNS Forwarder	string	Enter the IP address of the DNS forwarder. Required for recursive queries when *SAMBA_INTERNAL* is selected.
Domain Forest Level	drop-down menu	Choices are *2000*, *2003*, *2008*, *2008_R2*, *2012*, or *2012_R2*. Refer to Understanding Active Directory Domain Services (AD DS) Functional Levels (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc754918(v=ws.10)).
Administrator password	string	Enter the password to be used for the *Active Directory* (page 141) administrator account.
Kerberos Realm	drop-down menu	Auto-populates with information from the *Realm* when the settings in this screen are saved.

10.4.1 Samba Domain Controller Backup

A `samba_backup` script is available to back up Samba4 domain controller settings is available. From the *Shell* (page 244), run `/usr/local/bin/samba_backup --usage` to show the input options.

10.5 Dynamic DNS

Dynamic DNS (DDNS) is useful if the TrueNAS® system is connected to an ISP that periodically changes the IP address of the system. With dynamic DNS, the system can automatically associate its current IP address with a domain name, allowing access to the TrueNAS® system even if the IP address changes. DDNS requires registration with a DDNS service such as DynDNS (https://dyn.com/dns/).

Figure 10.5 shows the DDNS configuration screen and Table 10.3 summarizes the configuration options. The values for these fields are provided by the DDNS provider. After configuring DDNS, remember to start the DDNS service in *Services* → *Control Services*.

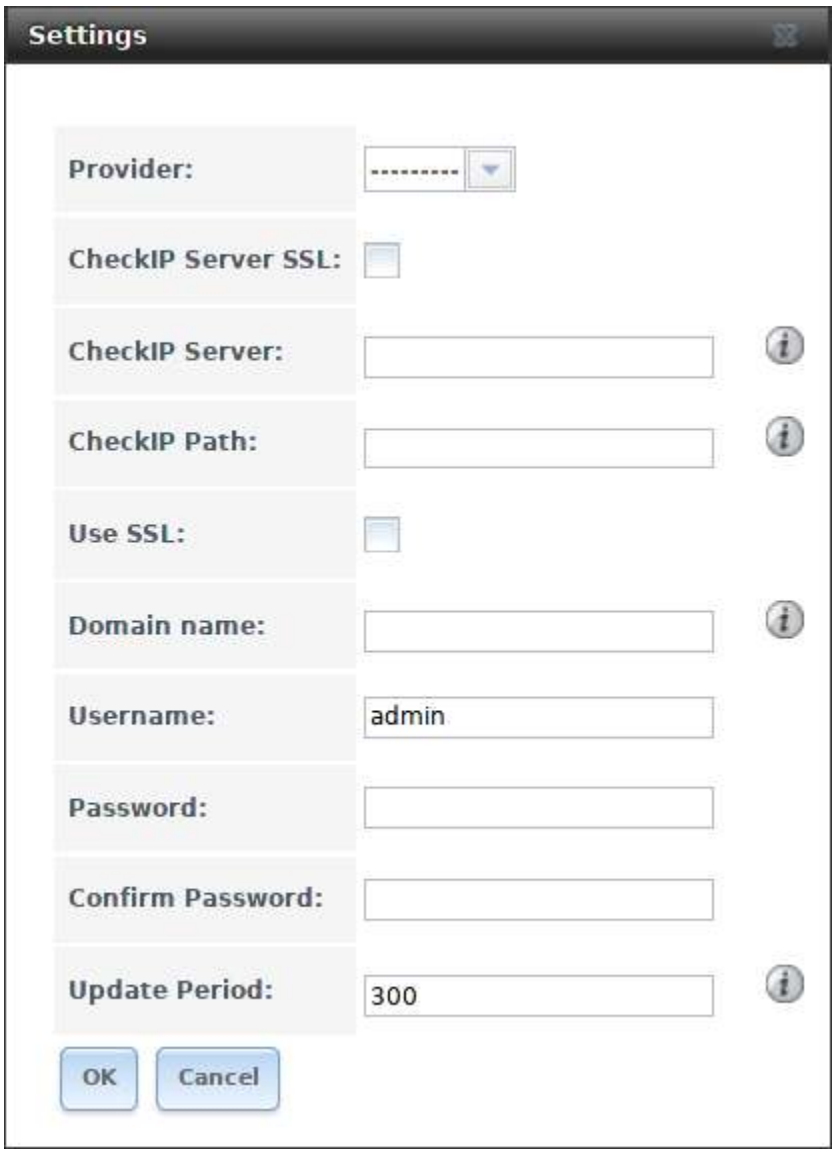

Fig. 10.5: Configuring DDNS

Table 10.3: DDNS Configuration Options

Setting	Value	Description
Provider	drop-down menu	Several providers are supported. If a specific provider is not listed, select *Custom Provider* and enter the information in the *Custom Server* and *Custom Path* fields.
CheckIP Server SSL	string	Set to use HTTPS for the connection to the *CheckIP Server*.
CheckIP Server	string	Enter the name and port of the server that reports the external IP address. Example: *server.name.org:port*.
CheckIP Path	string	Enter the path that is requested by the *CheckIP Server* to determine the user IP address.
Use SSL	checkbox	Set to use HTTPS for the connection to the server that updates the DNS record.
Domain name	string	Enter a fully qualified domain name. Separate multiple domains with a space, comma (`,`), or semicolon (`;`). Example: *yourname.dyndns.org;myname.dyndns.org*
Username	string	Enter the username used to log in to the provider and update the record.
Password	string	Enter the password used to log in to the provider and update the record.
Update period	integer	How often the IP is checked in seconds.

When using `he.net`, enter the domain name for *Username* and enter the DDNS key generated for that domain's A entry at the he.net (https://he.net) website for *Password*.

10.6 FTP

TrueNAS® uses the proftpd (http://www.proftpd.org/) FTP server to provide FTP services. Once the FTP service is configured and started, clients can browse and download data using a web browser or FTP client software. The advantage of FTP is that easy-to-use cross-platform utilities are available to manage uploads to and downloads from the TrueNAS® system. The disadvantage of FTP is that it is considered to be an insecure protocol, meaning that it should not be used to transfer sensitive files. If concerned about sensitive data, see *Encrypting FTP* (page 210).

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 10.6 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 enabling the *Show advanced fields by default* setting in *System* → *Advanced*.

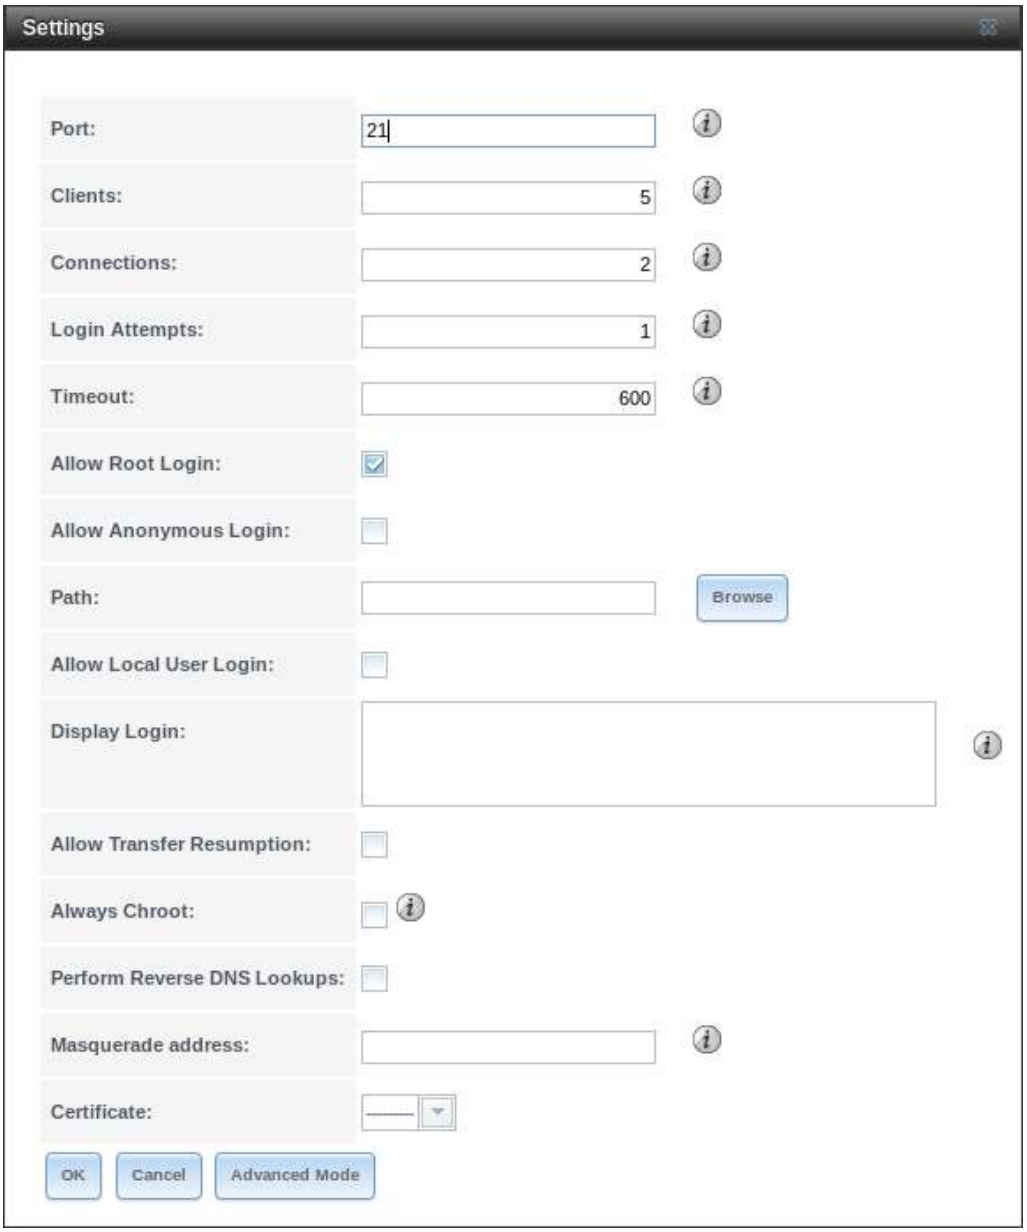

Fig. 10.6: Configuring FTP

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		Set the port the FTP service listens on.
Clients	integer		Set the maximum number of simultaneous clients.
Connections	integer		Set the maximum number of connections per IP address where *0* means unlimited.
Login Attempts	integer		Enter the maximum number of attempts before client is disconnected. Increase this if users are prone to typos.
Timeout	integer		Enter the maximum client idle time in seconds before client is disconnected.

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
Allow Root Login	checkbox		Enabling this option is discouraged as increases security risk.
Allow Anonymous Login	checkbox		Set to enable anonymous FTP logins with access to the directory specified in *Path*.
Path	browse button		Set the root directory for anonymous FTP connections.
Allow Local User Login	checkbox		Required if *Anonymous Login* is disabled.
Display Login	string		Specify the message displayed to local login users after authentication. Not displayed to anonymous login users.
File Permission	checkboxes	✓	Set the default permissions for newly created files.
Directory Permission	checkboxes	✓	Set the default permissions for newly created directories.
Enable FXP (https://en.wikipedia.org/wiki/File_eXchange_Protocol)	checkbox	✓	Set to enable the File eXchange Protocol. This setting makes the server vulnerable to FTP bounce attacks so it is not recommended
Allow Transfer Resumption	checkbox		Set to allow FTP clients to resume interrupted transfers.
Always Chroot	checkbox		When set, a local user is only allowed access to their home directory unless the user is a member of group *wheel*.
Require IDENT Authentication	checkbox	✓	Setting this option results in timeouts if `identd` is not running on the client.
Perform Reverse DNS Lookups	checkbox		Set to perform reverse DNS lookups on client IPs. Can cause long delays if reverse DNS is not configured.
Masquerade address	string		Public IP address or hostname. Set if FTP clients cannot connect through a NAT device.
Minimum passive port	integer	✓	Used by clients in PASV mode, default of *0* means any port above 1023.
Maximum passive port	integer	✓	Used by clients in PASV mode, default of *0* means any port above 1023.
Local user upload bandwidth	integer	✓	Defined in KiB/s, default of *0* means unlimited.
Local user download bandwidth	integer	✓	Defined in KiB/s, default of *0* means unlimited.
Anonymous user upload bandwidth	integer	✓	Defined in KiB/s, default of *0* means unlimited.
Anonymous user download bandwidth	integer	✓	Defined in KiB/s, default of *0* means unlimited.
Enable TLS	checkbox	✓	Set to enable encrypted connections. Requires a certificate to be created or imported using *Certificates* (page 51).
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	✓	Enabling this option is **not** recommended as it breaks several security measures. For this and the rest of the TLS fields, refer to mod_tls (http://www.proftpd.org/docs/contrib/mod_tls.html) for more details.

Continued on next page

Table 10.4 – continued from previous page

Setting	Value	Advanced Mode	Description
TLS allow dot login	checkbox	✓	If set, the user home directory is checked for a `.tlslogin` file which contains one or more PEM-encoded certificates. If not found, the user is prompted for password authentication.
TLS allow per user	checkbox	✓	If set, the user password can be sent unencrypted.
TLS common name required	checkbox	✓	Set to require the certificate common name to match the FQDN of the host.
TLS enable diagnostics	checkbox	✓	If set when troubleshooting a connection, logs more verbosely.
TLS export certificate data	checkbox	✓	If set, exports the certificate environment variables.
TLS no certificate request	checkbox	✓	Try enabling this option if the client cannot connect and it is suspected the client software is not properly handling server certificate requests.
TLS no empty fragments	checkbox	✓	Enabling this is **not** recommended as it bypasses a security mechanism.
TLS no session reuse required	checkbox	✓	Enabling this reduces the security of the connection. Only use this if the client does not understand reused SSL sessions.
TLS export standard vars	checkbox	✓	If enabled, sets several environment variables.
TLS DNS name required	checkbox	✓	If set, the client DNS name must resolve to its IP address and the cert must contain the same DNS name.
TLS IP address required	checkbox	✓	If set, the client certificate must contain the IP address that matches the IP address of the client.
Certificate	drop-down menu		The SSL certificate to be used for TLS FTP connections. To create a certificate, use *System* → *Certificates*.
Auxiliary parameters	string	✓	Add any additional proftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=proftpd) parameters not covered elsewhere in this screen.

This example demonstrates the auxiliary parameters that prevent all users from performing the FTP DELETE command:

```
<Limit DELE>
DenyAll
</Limit>
```

10.6.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the TrueNAS® system is not accessible from the Internet and everyone in the internal network needs easy access to the stored data. Anonymous FTP does not require a user account for every user. In addition, passwords are not required so it is not necessary to manage changed passwords on the TrueNAS® system.

To configure anonymous FTP:

1. Give the built-in ftp user account permissions to the 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 Unix ACLs are always used, even if Windows clients are accessing TrueNAS® via FTP.

2. Configure anonymous FTP in *Services* → *FTP* by setting these attributes:
   - *Allow Anonymous Login*: enable this option
   - *Path*: browse to the volume/dataset/directory to be shared
3. Start the FTP service in *Services* → *Control Services*. Click the *Start Now* button next to *FTP*. The FTP service takes a second or so to start. The indicator changes to green when the service is running, and the button changes to *Stop Now*.
4. Test the connection from a client using a utility such as Filezilla (https://filezilla-project.org/).

In the example shown in Figure 10.7, the user has entered this information into the Filezilla client:

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

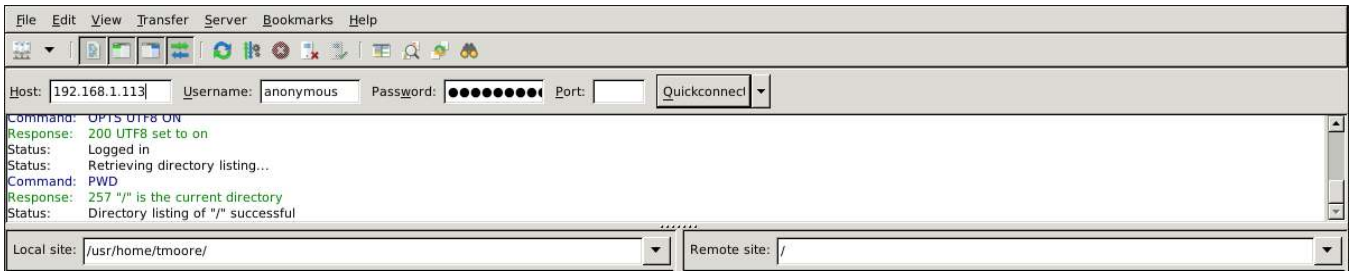

Fig. 10.7: Connecting Using Filezilla

The messages within the client indicate the FTP connection is successful. The user can now navigate the contents of the root folder on the remote site. This is the pool or dataset specified in the FTP service configuration. The user can also transfer files between the local site (their system) and the remote site (the TrueNAS® system).

10.6.2 FTP in chroot

If users are required to authenticate before accessing the data on the TrueNAS® system, either create a user account for each user or import existing user accounts using *Active Directory* (page 141) or *LDAP* (page 147). Then create a ZFS dataset for *each* user. Next, chroot each user so they are limited to the contents of their own home directory. Datasets provide the added benefit of configuring a quota so that the size of a user home directory is limited to the size of the quota.

To configure this scenario:

1. Create a ZFS dataset for each user in *Storage* → *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. When 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 Unix ACLs are always used, even if Windows clients will be accessing TrueNAS® with FTP.

4. Configure FTP in *Services* → *FTP* with these attributes:
   - *Path*: browse to the parent volume containing the datasets.
   - Make sure the options for *Allow Anonymous Login* and *Allow Root Login* are **unselected**.
   - Select the *Allow Local User Login* option to enable it.
   - Enable the *Always Chroot* option.
5. Start the FTP service in *Services* → *Control Services*. Click the *Start Now* button next to *FTP*. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to *Stop Now*.
6. Test the connection from a client using a utility such as Filezilla.

To test this configuration in Filezilla, use the *IP address* of the TrueNAS® system, the *Username* of a user that is associated with a dataset, and the *Password* for that user. The messages will indicate the authorization and the FTP connection are successful. The user can now navigate the contents of the root folder on the remote site. This time it is not the entire pool but the dataset created for that user. The user can transfer files between the local site (their system) and the remote site (their dataset on the TrueNAS® system).

10.6.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

1. Import or create a certificate authority using the instructions in *CAs* (page 49). Then, import or create the certificate to use for encrypted connections using the instructions in *Certificates* (page 51).
2. In *Services* → *FTP*, choose the certificate in the *Certificate*, and set the *Enable TLS* option.
3. Specify secure FTP when accessing the TrueNAS® system. For example, in Filezilla enter *ftps://IP_address* (for an implicit connection) or *ftpes://IP_address* (for an explicit connection) as the Host when connecting. The first time a user connects, they will be presented with the certificate of the TrueNAS® system. Click *OK* to accept the certificate and negotiate an encrypted connection.
4. To force encrypted connections, select *on* for the *TLS Policy*.

10.6.4 Troubleshooting FTP

The FTP service will not start if it cannot resolve the system hostname to an IP address with DNS. To see if the FTP service is running, open *Shell* (page 244) 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 *Show console messages in the footer*, and click *Save*. 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 hostname and IP address, or add an entry for the IP address of the TrueNAS® system in the *Network* → *Global Configuration Host name data base* field.

10.7 iSCSI

Refer to *Block (iSCSI)* (page 177) for instructions on configuring iSCSI. To start the iSCSI service, click its entry in *Services*.

Note: A warning message is shown if the iSCSI service is stopped when initiators are connected. Open the *Shell* (page 244) and type `ctladm islist` to determine the names of the connected initiators.

10.8 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/sspans/ladvd) LLDP implementation. If the network contains managed switches, configuring and starting the LLDP service will tell the TrueNAS® system to advertise itself on the network.

Figure 10.8 shows the LLDP configuration screen and Table 10.5 summarizes the configuration options for the LLDP service.

Fig. 10.8: Configuring LLDP

Table 10.5: LLDP Configuration Options

Setting	Value	Description
Interface Description	checkbox	Set to enable receive mode and to save received peer information in interface descriptions.
Country Code	string	Required for LLDP location support. Enter a two-letter ISO 3166 country code.
Location	string	Optional. Specify the physical location of the host.

10.9 Netdata

Netdata is a real-time performance and monitoring system. It displays data as web dashboards.

Start the Netdata service from the *Services* (page 198) screen. Click the wrench icon to display the Netdata settings dialog shown in Figure 10.9.

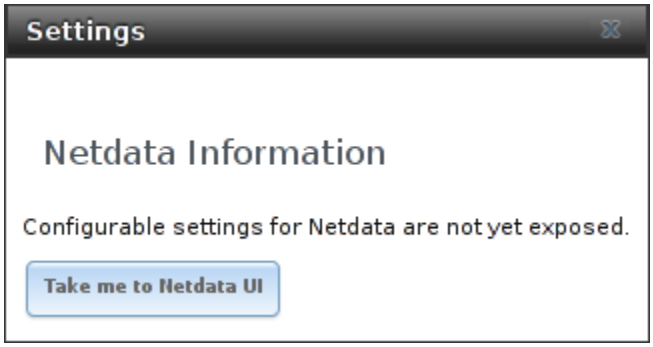

Fig. 10.9: Netdata Settings Dialog

Click the *Take me to the Netdata UI* button to view the web dashboard as shown in Figure 10.10.

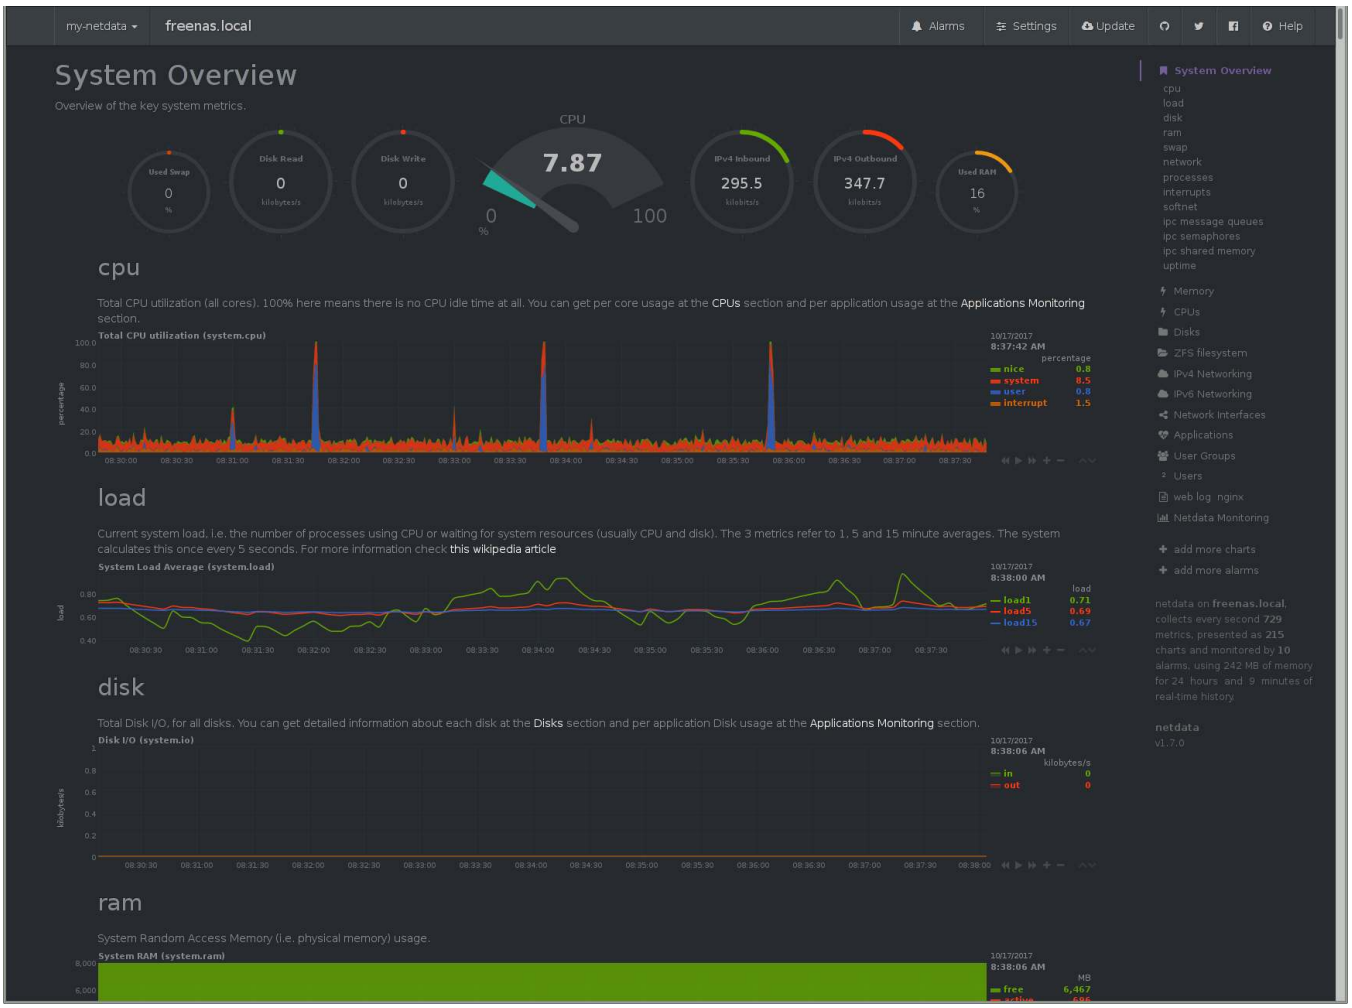

Fig. 10.10: Netdata Web Dashboard

More information on configuring and using Netdata is available at the Netdata website (https://my-netdata.io/).

10.10 NFS

The settings that are configured when creating NFS Shares in *Sharing* → *Unix (NFS) Shares* → *Add Unix (NFS) Share* are specific to each configured NFS Share. In contrast, global settings which apply to all NFS shares are configured in *Services* → *NFS*.

Figure 10.11 shows the configuration screen and Table 10.6 summarizes the configuration options for the NFS service.

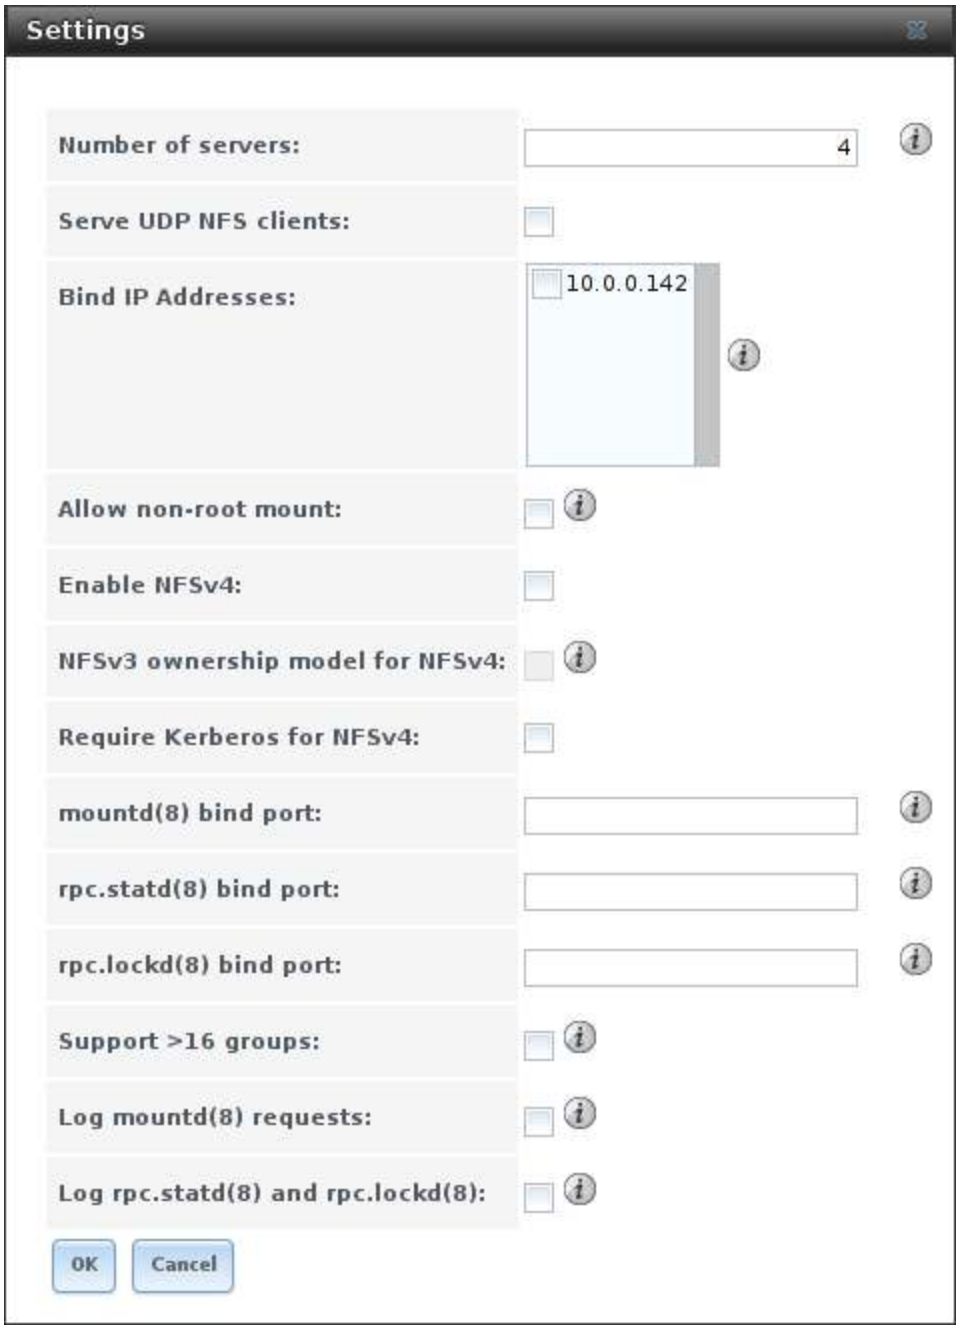

Fig. 10.11: Configuring NFS

Table 10.6: NFS Configuration Options

Setting	Value	Description
Number of servers	integer	Specify how many servers to create. Increase if NFS client responses are slow. To limit CPU context switching, keep this number less than or equal to the number of CPUs reported by `sysctl -n kern.smp.cpus`.
Serve UDP NFS clients	checkbox	Set if NFS clients need to use UDP.
Bind IP Addresses	checkboxes	Select the IP addresses to listen on for NFS requests. When unselected, NFS listens on all available addresses.
Allow non-root mount	checkbox	Set only if the NFS client requires it.
Enable NFSv4	checkbox	Set to switch from NFSv3 to NFSv4. The default is NFSv3.
NFSv3 ownership model for NFSv4	checkbox	Grayed out unless *Enable NFSv4* is checked and, in turn, grays out *Support>16 groups* which is incompatible. Set this option if NFSv4 ACL support is needed without requiring the client and the server to sync users and groups.
Require Kerberos for NFSv4	checkbox	Set to force NFS shares to fail if the Kerberos ticket is unavailable.
mountd(8) bind port	integer	Optional. Specify the port that mountd(8) (https://www.freebsd.org/cgi/man.cgi?query=mountd) binds to.
rpc.statd(8) bind port	integer	Optional. Specify the port that rpc.statd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) binds to.
rpc.lockd(8) bind port	integer	Optional. Specify the port that rpc.lockd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) binds to.
Support>16 groups	checkbox	Set this option if any users are members of more than 16 groups (useful in AD environments). Note this assumes group membership is configured correctly on the NFS server.
Log mountd(8) requests	checkbox	Enable logging of mountd(8) (https://www.freebsd.org/cgi/man.cgi?query=mountd) requests by syslog.
Log rpc.statd(8) and rpc.lockd(8)	checkbox	Enable logging of rpc.statd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) and rpc.lockd(8) (https://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) requests by syslog.

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

10.11 Rsync

Services → *Rsync* is used to configure an rsync server when using rsync module mode. Refer to *Rsync Module Mode* (page 74) for a configuration example.

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

10.11.1 Configure Rsyncd

Figure 10.12 shows the rsyncd configuration screen which is accessed from *Services* → *Rsync*.

Fig. 10.12: Rsyncd Configuration

Table 10.7 summarizes the configuration options 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 parameters	string	Enter any additional parameters from rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

10.11.2 Rsync Modules

Figure 10.13 shows the configuration screen that appears after clicking *Services* → *Rsync* → *Rsync Modules* → *Add Rsync Module*.

Table 10.8 summarizes the configuration options available when creating a rsync module.

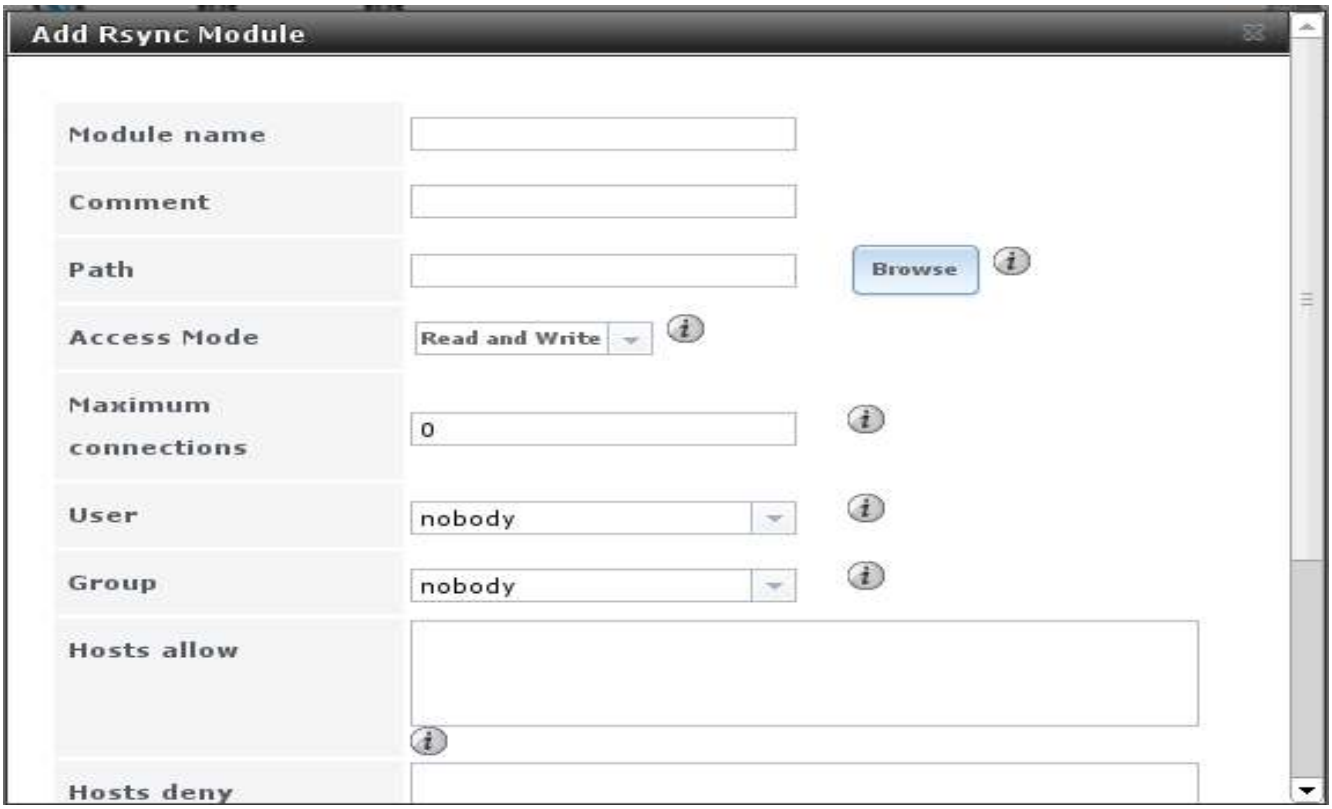

Fig. 10.13: Adding an Rsync Module

Table 10.8: Rsync Module Configuration Options

Setting	Value	Description
Module name	string	Mandatory. This is required to match the setting on the rsync client.
Comment	string	Optional description.
Path	browse button	Browse to the volume or dataset to hold received data.
Access Mode	drop-down menu	Choices are *Read and Write*, *Read-only*, or *Write-only*.
Maximum connections	integer	*0* is unlimited.
User	drop-down menu	Select the user to control file transfers to and from the module.
Group	drop-down menu	Select the group to control file transfers to and from the module.
Hosts allow	string	Optional patterns to match to allow hosts access. See rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf). Separate patterns with a space or newline. Defaults to empty, allowing all.
Hosts deny	string	Optional patterns to match to deny hosts access. See rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf). Separate patterns with a space or newline. Defaults to empty, denying none.
Auxiliary parameters	string	Enter any additional parameters from rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

10.12 S3

S3 is a distributed or clustered filesystem protocol compatible with Amazon S3 cloud storage. The TrueNAS® S3 service uses Minio (https://minio.io/) to provide S3 storage hosted on the TrueNAS® system itself. Minio also provides features beyond the limits of the basic Amazon S3 specifications.

Figure 10.14 shows the S3 service configuration screen and Table 10.9 summarizes the configuration options. After configuring the S3 service, start it in *Services* → *Control Services*.

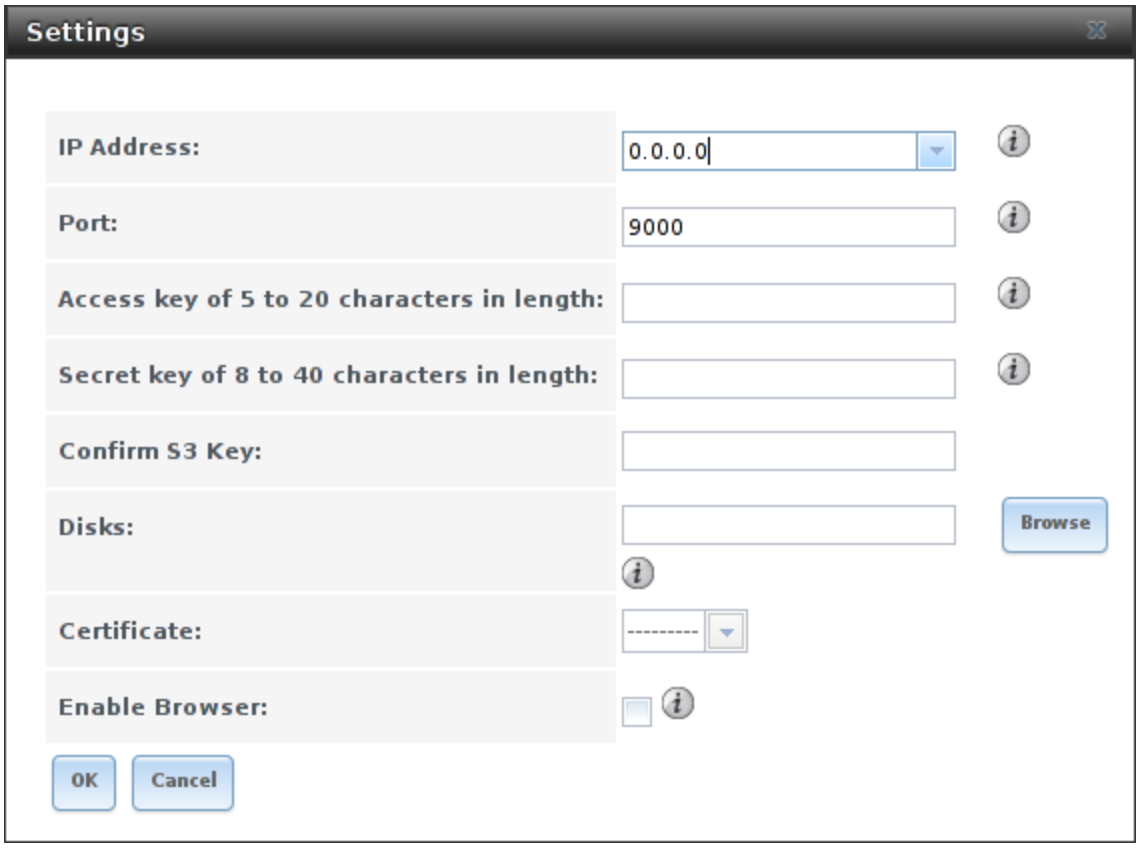

Fig. 10.14: Configuring S3

Table 10.9: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	Enter the IP address to run the S3 service. *0.0.0.0* sets the server to listen on all addresses.
Port	string	Enter the TCP port on which to provide the S3 service. Default is *9000*.
Access Key	string	Enter the S3 user name. This username must contain **only** alphanumeric characters and be between 5 and 20 characters long.
Secret Key	string	Enter the password to be used by connecting S3 systems. The key must contain **only** alphanumeric characters and be at least 8 but no more than 40 characters long.
Confirm S3 Key	string	Re-enter the S3 password to confirm.
Disks	string	Required. Directory where the S3 filesystem will be mounted. Ownership of this directory and all subdirectories is set to *minio:minio*. *Create a separate dataset* (page 102) for Minio to avoid issues with conflicting directory permissions or ownership.
Certificate	drop-down menu	The SSL certificate to be used for secure S3 connections. To create a certificate, use *System* → *Certificates*.
Enable Browser	checkbox	Set to enable the web user interface for the S3 service.

10.13 S.M.A.R.T.

S.M.A.R.T., or Self-Monitoring, Analysis, and Reporting Technology (https://en.wikipedia.org/wiki/S.M.A.R.T.), 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 134).

Of particular interest in a NAS environment are the *Short* and *Long* S.M.A.R.T. tests. Details vary between drive manufacturers, but a *Short* test generally does some basic tests of a drive that takes a few minutes. The *Long* test scans the entire disk surface, and can take several hours on larger drives.

TrueNAS® uses the smartd(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in) service to monitor S.M.A.R.T. information, including disk temperature. A complete configuration consists of:

1. Scheduling when S.M.A.R.T. tests are run in *Tasks* → *S.M.A.R.T. Tests* → *Add S.M.A.R.T. Test*.
2. Enabling or disabling S.M.A.R.T. for each disk member of a volume in *Volumes* → *View Disks*. This setting is enabled by default for disks that support S.M.A.R.T.
3. Checking the configuration of the S.M.A.R.T. service as described in this section.
4. Starting the S.M.A.R.T. service with *Services* → *Control Services*.

Figure 10.15 shows the configuration screen that appears after clicking *Services* → *S.M.A.R.T.*

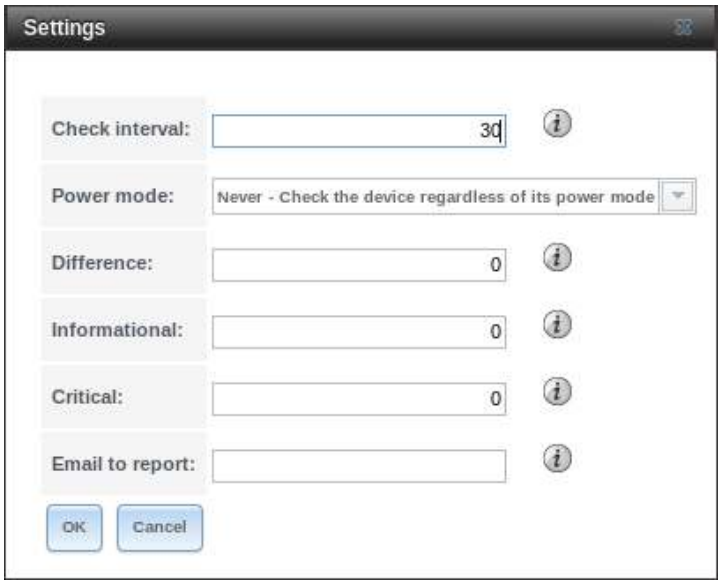

Fig. 10.15: S.M.A.R.T Configuration Options

Note: `smartd` wakes up at the configured *Check Interval*. It checks the times configured in *Tasks* → *S.M.A.R.T. Tests* to see if a test must begin. Since the smallest time increment for a test is an hour, it does not make sense to set a *Check Interval* value higher than 60 minutes. For example, if the *Check Interval* is set to *120* minutes and the smart test to every hour, the test will only be run every two hours because `smartd` only activates every two hours.

Table 10.10 summarizes the options in the S.M.A.R.T configuration screen.

Table 10.10: S.M.A.R.T Configuration Options

Setting	Value	Description
Check interval	integer	Define in minutes how often `smartd` activates to check if any tests are configured to run.
Power mode	drop-down menu	Tests are not performed if the system enters the specified power mode: *Never*, *Sleep*, *Standby*, or *Idle*.
Difference	integer in degrees Celsius	Enter number of degrees in Celsius. S.M.A.R.T reports if the temperature of a drive has changed by N degrees Celsius since the last report. Default of *0* disables this option.
Informational	integer in degrees Celsius	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a log level of LOG_INFO if the temperature is higher than specified degrees in Celsius. Default of *0* disables this option.
Critical	integer in degrees Celsius	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a log level of LOG_CRIT and send an email if the temperature is higher than specified degrees in Celsius. Default of *0* disables this option.
Email to report	string	Email address to receive S.M.A.R.T. alerts. Use a space to separate multiple email addresses.

10.14 SMB

The settings 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 (https://www.samba.org/samba/docs/old/Samba3-HOWTO/NetworkBrowsing.html#id2581357) to occur and for the TrueNAS® system to become available in Windows Explorer.

Figure 10.16 shows some of the global SMB configuration options described in Table 10.11. This configuration screen is really a front-end to smb4.conf (https://www.freebsd.org/cgi/man.cgi?query=smb4.conf).

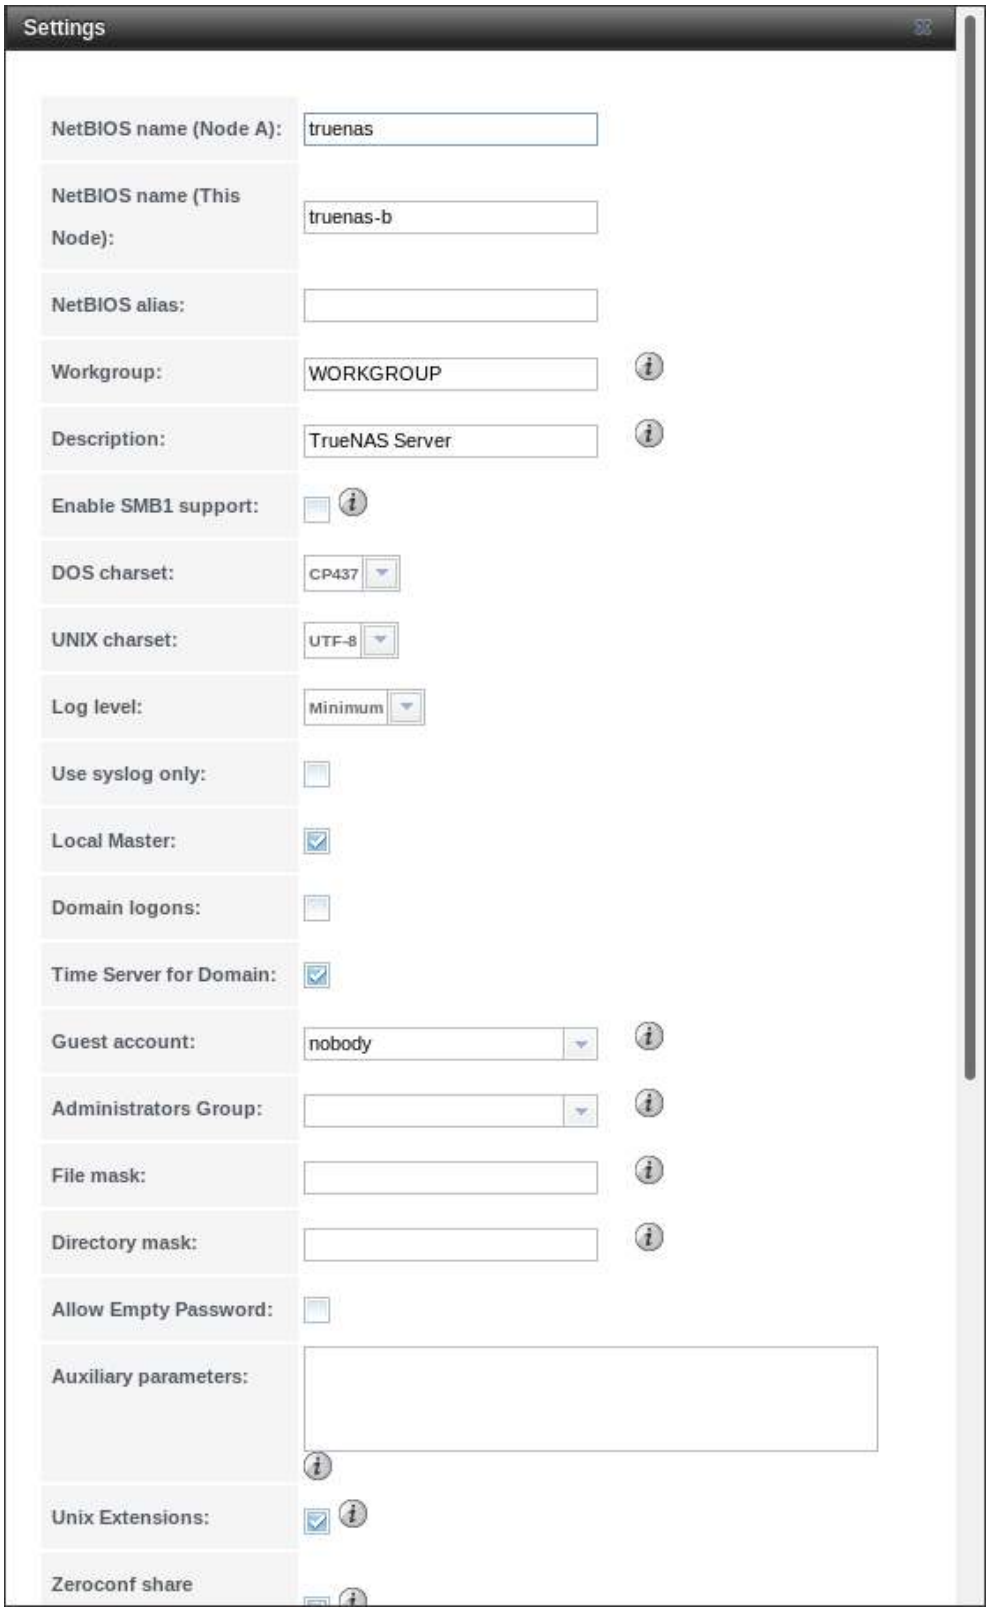

Fig. 10.16: Global SMB Configuration

Table 10.11: Global SMB Configuration Options

Setting	Value	Description
NetBIOS Name (This Node)	string	Automatically populated with the original hostname of the system. Limited to 15 characters. It **must** be different from the *Workgroup* name.
NetBIOS Name (Node B)	string	Limited to 15 characters. When using *Failover* (page 58), set a unique NetBIOS name for the standby node
NetBIOS Alias	string	Limited to 15 characters. When using *Failover* (page 58), 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 141) or *LDAP* (page 147) service is running.
Description	string	Enter an optional server description.
Enable SMB1 support	checkbox	Allow legacy SMB clients to connect to the server. **Warning:** SMB1 is not secure and has been deprecated by Microsoft. See Do Not Use SMB1 (https://www.ixsystems.com/blog/library/do-not-use-smb1/).
DOS charset	drop-down menu	The character set Samba uses when communicating with DOS and Windows 9x/ME clients. Default is *CP437*.
UNIX charset	drop-down menu	Default is *UTF-8* which supports all characters in all languages.
Log level	drop-down menu	Choices are *Minimum*, *Normal*, or *Debug*.
Use syslog only	checkbox	Set to log authentication failures to `/var/log/messages` instead of the default of `/var/log/samba4/log.smbd`.
Local Master	checkbox	Set to determine if the system will participate in a browser election. Disable when network contains an AD or LDAP server or Vista or Windows 7 machines are present.
Domain logons	checkbox	Set if it is necessary to provide the netlogin service for older Windows clients.
Time Server for Domain	checkbox	Determines if the system advertises itself as a time server to Windows clients. Disable when network contains an AD or LDAP server.
Guest Account	drop-down menu	Select the account to be used for guest access. Default is *nobody*. Account must have permission to access the shared volume/dataset. If Guest Account user is deleted, resets to *nobody*.
Administrators Group	drop-down menu	Members of this group are local admins and automatically have privileges to take ownership of any file in an SMB share, reset permissions, and administer the SMB server through the Computer Management MMC snap-in.
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	Set to allow users to press `Enter` when prompted for a password. Requires the username/password to be the same as the Windows user account.
Auxiliary parameters	string	Add any `smb.conf` 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	Set to allow non-Windows SMB clients to access symbolic links and hard links, has no effect on Windows clients.
Zeroconf share discovery	checkbox	Enable if Mac clients will be connecting to the SMB share.
Hostname lookups	checkbox	Set to allow using hostnames rather than IP addresses in the *Hosts Allow* or *Hosts Deny* fields of a SMB share. Unset if IP addresses are used to avoid the delay of a host lookup.
Allow execute always	checkbox	If set, Samba will allow the user to execute a file, even if that user's permissions are not set to execute.

Continued on next page

Table 10.11 – continued from previous page

Setting	Value	Description
Obey pam restrictions	checkbox	Unset this option to allow: Cross-domain authentication. Users and groups to be managed on another forest. Permissions to be delegated from *Active Directory* (page 141) users and groups to domain admins on another forest.
NTLMv1 auth	checkbox	Set to allow NTLMv1 authentication. Required by Windows XP clients and sometimes by clients in later versions of Windows.
Bind IP Addresses	checkboxes	Select the IPv4 and IPv6 addresses SMB will listen on. Always add the loopback interface *127.0.0.1* as Samba utilities connect to the loopback IP (https://wiki.samba.org/index.php/Configure_Sama_to_Bind_to_Specific_Interfaces) if no host name is provided.
Idmap Range Low	integer	The beginning UID/GID for which this system is authoritative. Any UID/GID lower than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs.
Idmap Range High	integer	The ending UID/GID for which this system is authoritative. Any UID/GID higher than this value is ignored, providing a way to avoid accidental UID/GID overlaps between local and remotely defined IDs.

Changes to SMB settings take effect immediately. Changes to share settings only take effect after the client and server negotiate a new session.

Note: Do not set the *directory name cache size* as an *Auxiliary parameter*. Due to differences in how Linux and BSD handle file descriptors, directory name caching is disabled on BSD systems to improve performance.

Note: *SMB* (page 219) cannot be disabled while *Active Directory* (page 141) is enabled.

10.14.1 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 in to the Windows system. Alternately, users can type `net use /delete` from the command line to clear their SMB sessions.

Windows also automatically caches login information. To require users to log in every time they access they system, reduce the cache settings on the client computers.

Where possible, avoid using a mix of case in filenames as this can cause confusion for Windows users. Representing and resolving filenames with Samba (http://www.oreilly.com/openbook/samba/book/ch05_04.html) explains in more detail.

If a particular user cannot connect to a SMB share, ensure their password does not contain the `?` character. If it does, have the user change the password and try again.

If permissions work for Windows users but not for macOS users, try disabling *Unix Extensions* and restarting the SMB service.

If the SMB service will not start, run this command from *Shell* (page 244) 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*.

Using a dataset for SMB sharing is recommended. When creating the dataset, make sure that the *Share type* is set to Windows.

Do not use `chmod` to attempt to fix the permissions on a SMB share as it destroys the Windows ACLs. The correct way to manage permissions on a SMB share is to manage the share security from a Windows system as either the owner of the share or a member of the group that owns the share. To do so, right-click on the share, click *Properties* and navigate to the *Security* tab. If the ACLs are already destroyed by using `chmod`, `winacl` can be used to fix them. Type `winacl` from *Shell* (page 244) for usage instructions.

The Common Errors (https://www.samba.org/samba/docs/old/Samba3-HOWTO/domain-member.html#id2573692) section of the Samba documentation contains additional troubleshooting tips.

The Samba Performance Tuning (https://wiki.samba.org/index.php/Performance_Tuning) page describes options to improve performance.

Directory listing speed in folders with a large number of files is sometimes a problem. A few specific changes can help improve the performance. However, changing these settings can affect other usage. In general, the defaults are adequate. **Do not change these settings unless there is a specific need.**

- *Hostname Lookups* and *Log Level* can also have a performance penalty. When not needed, they can be disabled or reduced in the *global SMB service options* (page 221).
- Make Samba datasets case insensitive by setting *Case Sensitivity* to *Insensitive* when creating them. This ZFS property is only available when creating a dataset. It cannot be changed on an existing dataset. To convert such datasets, back up the data, create a new case-insensitive dataset, create an SMB share on it, set the share level auxiliary parameter *case sensitive = true*, then copy the data from the old one onto it. After the data has been checked and verified on the new share, the old one can be deleted.
- If present, remove options for extended attributes and DOS attributes in *Auxiliary Parameters* (page 167) for the share.
- Disable as many *VFS Objects* as possible in the *share settings* (page 167). Many have performance overhead.

The SMB1 protocol is deprecated and vulnerable. Before enabling it, see Do Not Use SMB1 (https://www.ixsystems.com/blog/library/do-not-use-smb1/).

10.15 SNMP

SNMP (Simple Network Management Protocol) is used to monitor network-attached devices for conditions that warrant administrative attention. TrueNAS® uses Net-SNMP (http://net-snmp.sourceforge.net/) to provide SNMP. When starting the SNMP service, this port will be enabled on the TrueNAS® system:

- UDP 161 (listens here for SNMP requests)

Available MIBS are located in `/usr/local/share/snmp/mibs`.

Figure 10.17 shows the SNMP configuration screen. Table 10.12 summarizes the configuration options.

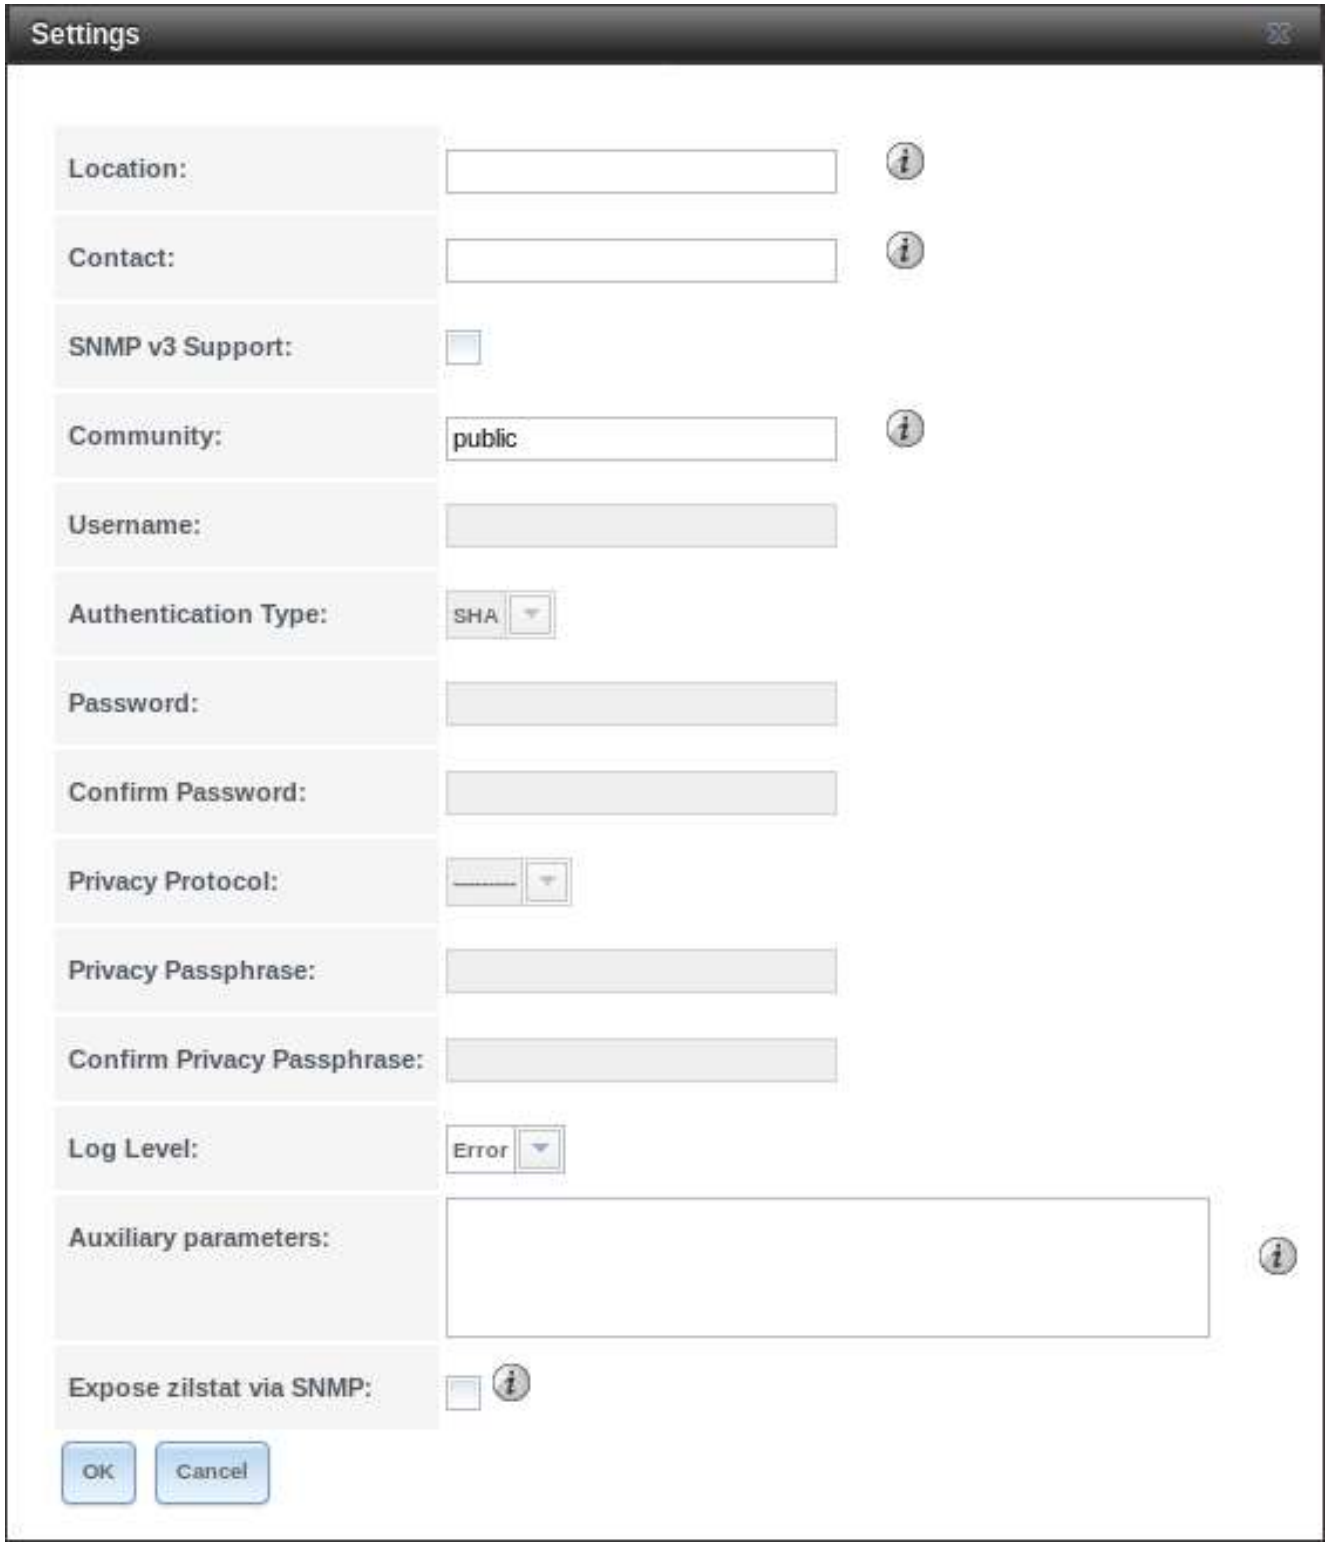

Fig. 10.17: Configuring SNMP

Table 10.12: SNMP Configuration Options

Setting	Value	Description
Location	string	Optional description of the system location.

Continued on next page

Table 10.12 – continued from previous page

Setting	Value	Description
Contact	string	Optional. Enter the administrator email address.
SNMP v3 Support	checkbox	Set to enable support for SNMP version 3.
Community	string	Default is *public*. **Change this for security reasons!** The value can only contain alphanumeric characters, underscores, dashes, periods, and spaces. This value can be empty for SNMPv3 networks.
Username	string	Only applies if *SNMP v3 Support* is set. Specify the username to register with this service. Refer to snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) for more information about configuring this and the *Authentication Type*, *Password*, *Privacy Protocol*, and *Privacy Passphrase* fields.
Authentication Type	drop-down menu	Only applies if *SNMP v3 Support* is enabled. Choices are: *MD5* or *SHA*.
Password	string	Only applies if *SNMP v3 Support* is enabled. Specify and confirm a password of at least eight characters.
Privacy Protocol	drop-down menu	Only applies if *SNMP v3 Support* is enabled. Choices are: *AES* or *DES*.
Privacy Passphrase	string	If not specified, *Password* is used.
Log Level	drop-down menu	Choices range from fewest log entries (*Emergency*) to the most (*Debug*).
Auxiliary Parameters	string	Enter additional snmpd.conf(5) (http://net-snmp.sourceforge.net/docs/man/snmpd.conf.html) options not covered in this screen. One option per line.
Expose zilstat via SNMP	checkbox	Gather ZFS Intent Log (ZIL) statistics. Enabling this option slows down pool performance.

Zenoss (https://www.zenoss.com/) provides a seamless monitoring service through SNMP for TrueNAS® called TrueNAS ZenPack (https://www.zenoss.com/product/zenpacks/truenas).

10.16 SSH

Secure Shell (SSH) is used to transfer files securely over an encrypted network. When a TrueNAS® system is used as an SSH server, the users in the network must use SSH client software (https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) 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 10.18 shows the *Services* → *SSH* configuration screen. After configuring SSH, remember to start it in *Services* → *Control Services*.

Fig. 10.18: 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 enabling the *Show advanced fields by default* option in *System* → *Advanced*.

Table 10.13: SSH Configuration Options

Setting	Value	Advanced Mode	Description
Bind Interfaces	selection	✓	By default, SSH listens on all interfaces unless specific interfaces are highlighted in the *Available* field and added to the *Selected* field.
TCP Port	integer		Port to open for SSH connection requests. *22* by default.
Login as Root with password	checkbox		**As a security precaution, root logins are discouraged and disabled by default.** If enabled, a password must be set for the *root* user in *View Users*.
Allow Password Authentication	checkbox		Unset to require key-based authentication for all users. Requires additional setup (http://the.earth.li/~sgtatham/putty/0.55/htmldoc/Chapter8.html) on both the SSH client and server.
Allow Kerberos Authentication	checkbox	✓	Before setting this option, ensure *Kerberos Realms* (page 151) and *Kerberos Keytabs* (page 151) are configured and TrueNAS® can communicate with the Kerberos Domain Controller (KDC).
Allow TCP Port Forwarding	checkbox		Set to allow users to bypass firewall restrictions using the SSH port forwarding feature (https://www.symantec.com/connect/articles/ssh-port-forwarding).
Compress Connections	checkbox		Set to attempt to reduce latency over slow networks.
SFTP Log Level	drop-down menu	✓	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) level of the SFTP server.
SFTP Log Facility	drop-down menu	✓	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) facility of the SFTP server.

Continued on next page

Table 10.13 – continued from previous page

Setting	Value	Advanced Mode	Description
Extra Options	string	✓	Add any additional sshd_config(5) (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) options not covered in this screen, one per line. These options are case-sensitive and misspellings can prevent the SSH service from starting.

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

- increase the *ClientAliveInterval* if SSH connections tend to drop
- *ClientMaxStartup* defaults to *10*. Increase this value if more concurrent SSH connections are required.

10.16.1 SCP Only

When SSH is configured, authenticated users with a user account created using *Account* → *Users* → *Add User* can use `ssh` to log into the TrueNAS® system over the network. The user home directory is the pool or dataset specified in the *Home Directory* field of the TrueNAS® account for that user. While the SSH login defaults to the user home directory, users are able to navigate outside their home directory, which can pose a security risk.

It is possible to allow users to use `scp` and `sftp` to transfer files between their local computer and their home directory on the TrueNAS® system, while restricting them from logging into the system using `ssh`. To configure this scenario, go to *Account* → *Users* → *View Users*, select the user, and click *Modify User*. Change the *Shell* to *scponly*. Repeat for each user that needs restricted SSH access.

Test the configuration from another system by running the `sftp`, `ssh`, and `scp` commands as the user. `sftp` and `scp` will work but `ssh` will fail.

Note: Some utilities like WinSCP and Filezilla can bypass the scponly shell. This section assumes that users are accessing the system using the command line versions of `scp` and `sftp`.

10.16.2 Troubleshooting SSH

Keywords listed in sshd_config(5) (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) are case sensitive. This is important to remember when adding any *Extra options*. The configuration will not function as intended if the upper and lowercase letters of the keyword are not an exact match.

If clients are receiving "reverse DNS" or timeout errors, add an entry for the IP address of the TrueNAS® system in the *Host name database* field of *Network* → *Global Configuration*.

When configuring SSH, always test the configuration as an SSH user account to ensure the user is limited by the configuration and they have permission to transfer files within the intended directories. If the user account is experiencing problems, the SSH error messages are specific in describing the problem. Type this command within *Shell* (page 244) to read these messages as they occur:

```
tail -f /var/log/messages
```

Additional messages regarding authentication errors are found in `/var/log/auth.log`.

10.17 TFTP

Trivial File Transfer Protocol (TFTP) is a light-weight version of FTP typically used to transfer configuration or boot files between machines, such as routers, in a local environment. TFTP provides an extremely limited set of commands and provides no authentication.

If the TrueNAS® system will be used to store images and configuration files for network devices, configure and start the TFTP service. Starting the TFTP service opens UDP port 69.

Figure 10.19 shows the TFTP configuration screen and Table 10.14 summarizes the available options.

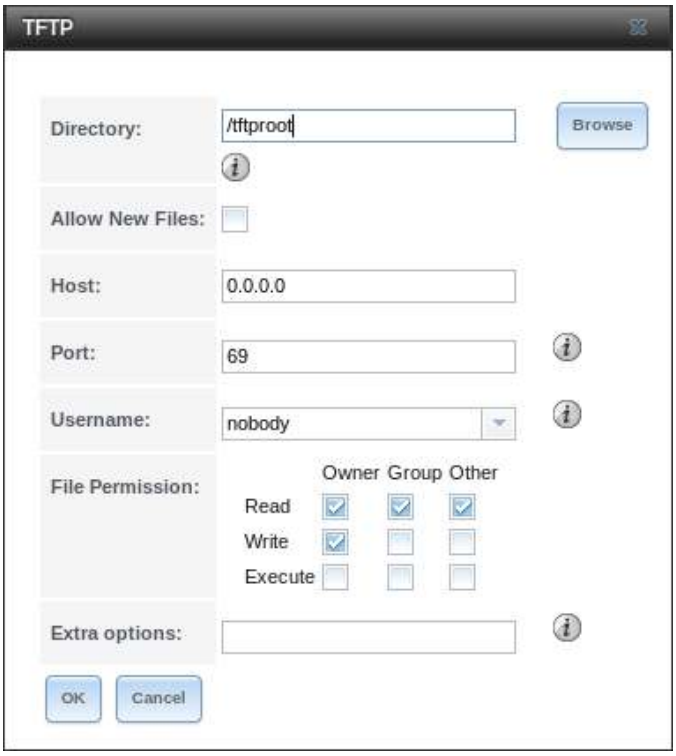

Fig. 10.19: TFTP Configuration

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 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).
Host	IP address	The default host to use for TFTP transfers. Enter an IP address. Example: *192.0.2.1*.
Port	integer	The UDP port number that listens for TFTP requests. Example: *8050*.
Username	drop-down menu	Select the account to be used for TFTP requests. The account must have permission to access the *Directory*.
File Permissions	checkboxes	Set permissions for newly created files. The default is everyone can read and only the owner can write. Some devices require less strict permissions.
Extra options	string	Add any additional tftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=tftpd) options not shown in this screen. Add one option on each line.

10.18 UPS

TrueNAS® uses NUT (http://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 10.20 shows the UPS configuration screen:

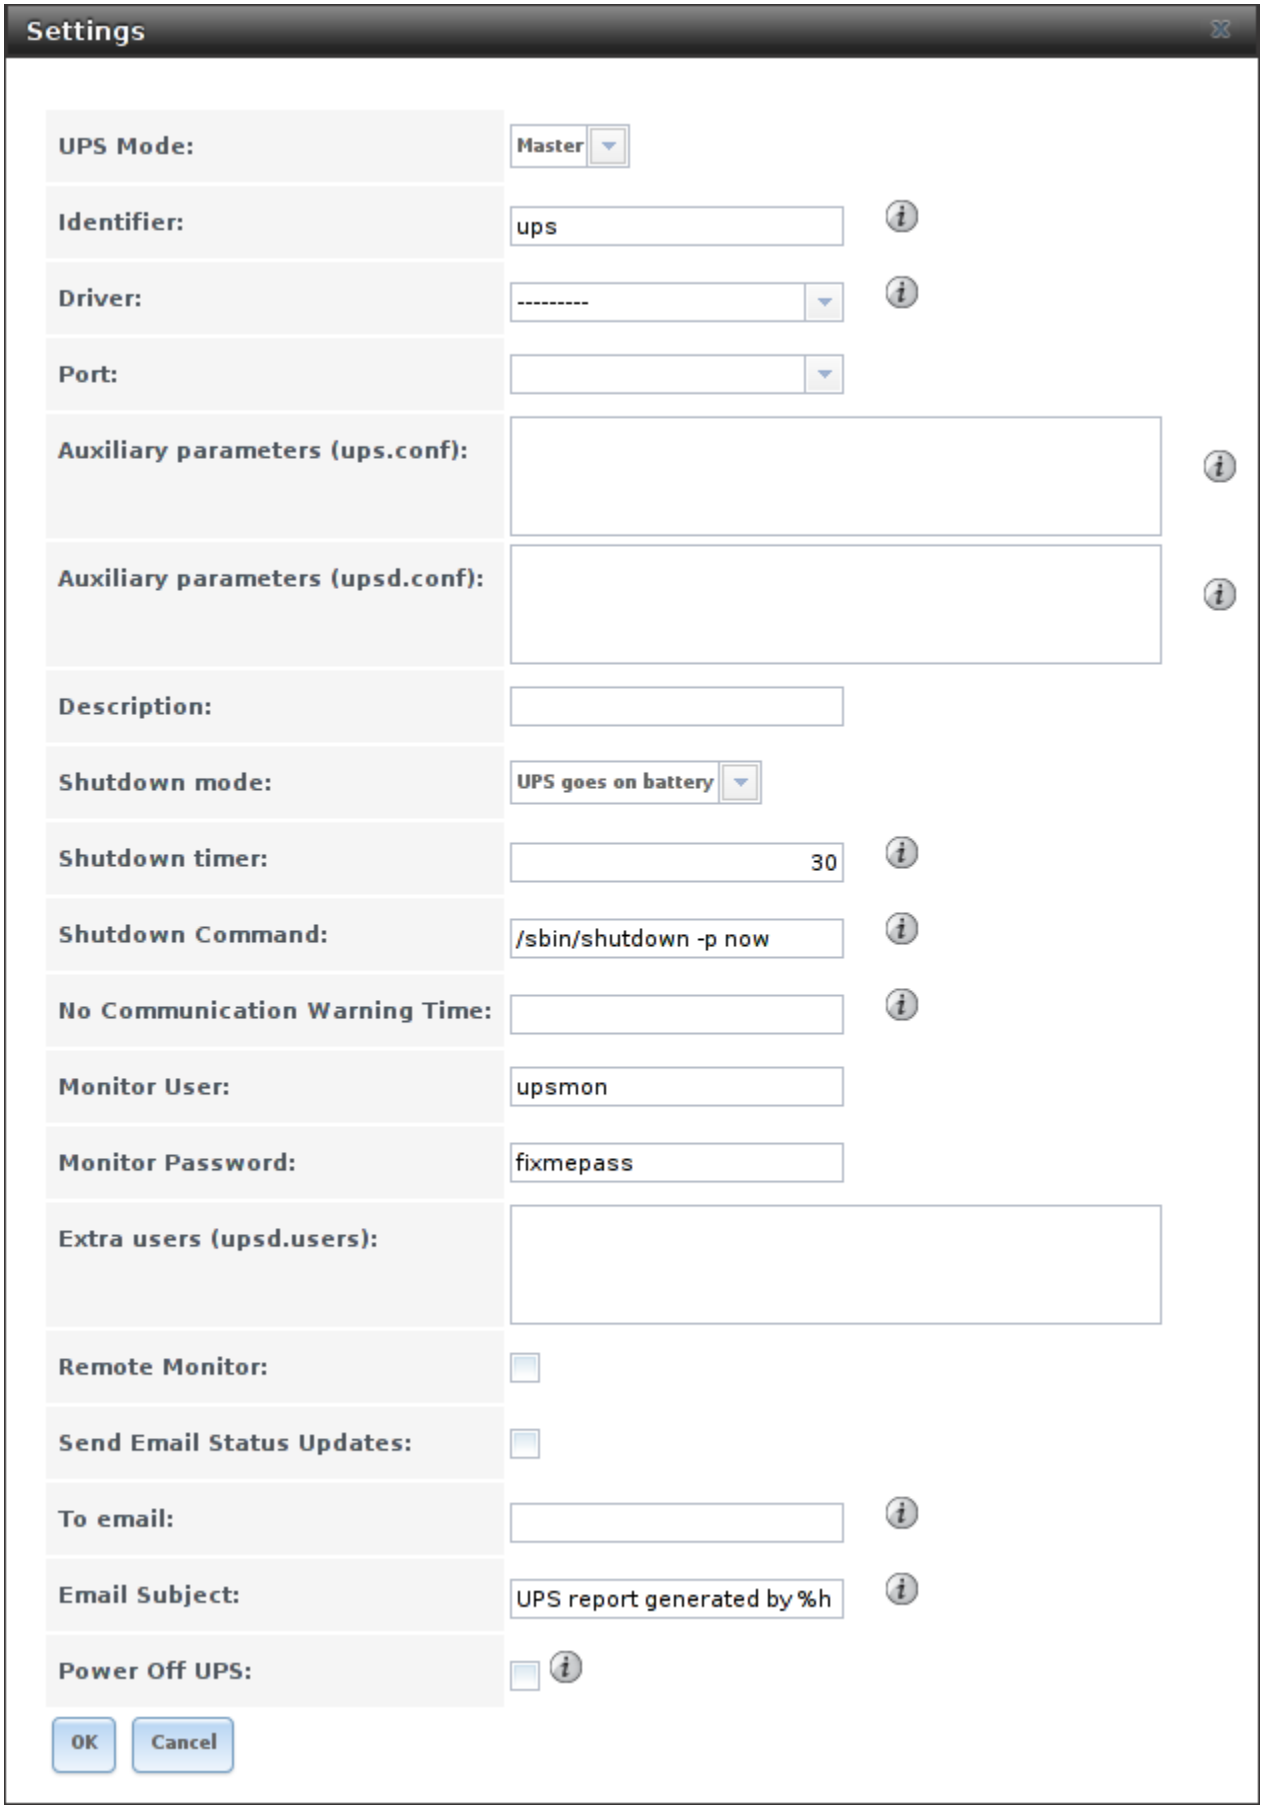

Fig. 10.20: UPS Configuration Screen

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 *Master* if the UPS is plugged directly into the system serial port. The UPS will remain the last item to shut down. Select *Slave* to have the system shut down before *Master*.
Identifier	string	Required. Describe the UPS device. Can contain alphanumeric, period, comma, hyphen, and underscore characters.
Driver / Remote Host	drop-down menu	Required. For a list of supported devices, see the Network UPS Tools compatibility list (https://networkupstools.org/stable-hcl.html). The *Driver* field changes to *Remote Host* when *UPS Mode* is set to *Slave*. Enter the IP address of the system configured as the UPS *Master* system. See this post (https://forums.freenas.org/index.php?resources/configuring-ups-support-for-single-or-multiple-freenas-servers.30/) for more details about configuring multiple systems with a single UPS.
Port / Remote Port	drop-down menu	Required. Enter the serial or USB port connected to the UPS (see *NOTE* (page 231)). Enter the IP address or hostname of the SNMP UPS device when an SNMP driver is selected. *Port* becomes *Remote Port* when the *UPS Mode* is set to *Slave*. Enter the open network port number of the UPS *Master* system. The default port is *3493*.
Auxiliary Parameters (ups.conf)	string	Enter any additional options from ups.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=ups.conf).
Auxiliary Parameters (upsd.conf)	string	Enter any additional options from upsd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=upsd.conf).
Description	string	Optional. Enter any notes about the UPS service.
Shutdown mode	drop-down menu	Choose when the UPS initiates shutdown. Choices are *UPS goes on battery* and *UPS reaches low battery*.
Shutdown timer	integer	Select a value in seconds for the UPS to wait before initiating shutdown. Shutdown will not occur if the power is restored while the timer is counting down. The value only applies when *Shutdown Mode* is set to *UPS goes on battery*.
Shutdown Command	string	Required. Enter the command to run to shut down the computer when battery power is low or shutdown timer runs out.
No Communication Warning Time	string	Enter a value in seconds to wait before alerting that the service cannot reach any UPS. Warnings continue until the situation is fixed.
Monitor User	string	Required. Enter a user to associate with this service. The recommended default user is *upsmon*.
Monitor Password	string	Required. Default is the known value *fixmepass*. Change this to enhance system security. Cannot contain a space or `#`.
Extra users (upsd.users)	string	Enter accounts that have administrative access. See upsd.users(5) (https://www.freebsd.org/cgi/man.cgi?query=upsd.users) for examples.
Remote monitor	checkbox	Set for the default configuration to listen on all interfaces using the known values of user *upsmon* and password *fixmepass*.
Send Email Status Updates	checkbox	Set to enable the TrueNAS® system to send email updates to the configured *To email* address.
To email	email address	Enter the email address to receive status updates. Separate multiple email addresses with a semicolon (`;`).
Email Subject	string	Enter a subject line to be used in email status updates.
Power Off UPS	checkbox	Set to power off the UPS after shutting down the FreeNAS system.

Note: For USB devices, the easiest way to determine the correct device name is to enable the *Show console messages* option in *System* → *Advanced*. Plug in the USB device and look for a */dev/ugen* or */dev/uhid* device name in the console messages.

Tip: Some UPS models might be unresponsive with the default polling frequency. This can show in TrueNAS® logs as a recurring error like: `libusb_get_interrupt: Unknown error`.

If this error occurs, decrease the polling frequency by adding an entry to *Auxiliary Parameters (ups.conf)*: `pollinterval = 10`. The default polling frequency is two seconds.

upsc(8) (http://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* (page 244) using this syntax:

```
upsc ups@localhost
```

The *upsc(8)* man page gives some other usage examples.

upscmd(8) (http://networkupstools.org/docs/man/upscmd.html) can be used to send commands directly to the UPS, assuming the hardware supports the command being sent. Only users with administrative rights can use this command. These users are created in the *Extra users* field.

10.18.1 Multiple Computers with One UPS

A UPS with adequate capacity can power multiple computers. One computer is connected to the UPS data port with a serial or USB cable. This *master* makes UPS status available on the network for other computers. These *slave* computers are powered by the UPS, but receive UPS status data from the master computer. See the NUT User Manual (http://networkupstools.org/docs/user-manual.chunked/index.html) and NUT User Manual Pages (http://networkupstools.org/docs/man/index.html#User_man).

10.19 WebDAV

The WebDAV service can be configured to provide a file browser over a web connection. Before starting this service, at least one WebDAV share must be created using *Sharing* → *WebDAV Shares* → *Add WebDAV Share*. Refer to *WebDAV Shares* (page 165) for instructions on how to create a share and connect to it when the service is configured and started.

Figure 10.21 shows the WebDAV configuration screen. Table 10.16 summarizes the available options.

Fig. 10.21: WebDAV Configuration Screen

Table 10.16: WebDAV Configuration Options

Setting	Value	Description
Protocol	drop-down menu	*HTTP* keeps the connection always unencrypted. *HTTPS* always encrypts the connection. *HTTP+HTTPS* allows both types of connections.
HTTP Port	string	Specify a port for unencrypted connections. Only appears if the selected *Protocol* is *HTTP* or *HTTP+HTTPS*. The default of *8080* is recommended. **Do not** reuse a port number.
HTTPS Port	string	Specify a port for encrypted connections. Only appears if the selected *Protocol* is *HTTPS* or *HTTP+HTTPS*. The default of *8081* is recommended. **Do not** reuse a port number.
Webdav SSL Certificate	drop-down menu	Select the SSL certificate to use for encrypted connections. Only appears if the selected *Protocol* is *HTTPS* or *HTTP+HTTPS*. To create a certificate, use *System* → *Certificates*.
HTTP Authentication	drop-down menu	Choices are *No Authentication*, *Basic Authentication* (unencrypted), or *Digest Authentication* (encrypted).
Webdav Password	string	Default is *davtest*. This is a known value and is recommended to be changed.

CHAPTER

ELEVEN

VCENTER PLUGIN

vCenter Server (https://www.vmware.com/products/vcenter-server.html) is server management software that uses a single console to manage a virtual infrastructure across a hybrid cloud of physical and virtual machines. The TrueNAS® vCenter Plugin makes it possible to provision and use TrueNAS® storage from within vCenter Server.

For more information, please contact iXsystems Support at `support@iXsystems.com` or by phone:

- US-only toll-free: 855-473-7449 option 2
- Local and international: 408-943-4100 option 2

CHAPTER

TWELVE

REPORTING

Reporting displays several graphs, as seen in Figure 12.1. Click the tab for a device type to see those specific graphs.

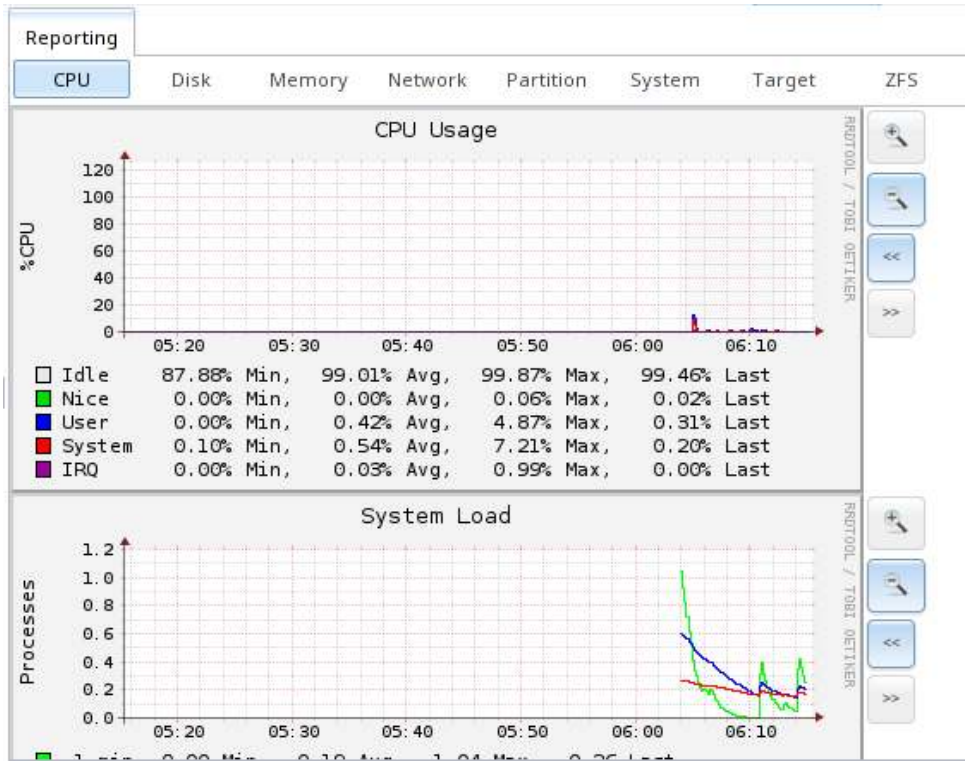

Fig. 12.1: Reporting Graphs

TrueNAS® uses collectd (https://collectd.org/) to provide reporting statistics. The resulting graphs are grouped into several tabs on the Reporting page:

- *CPU*
  - CPU (https://collectd.org/wiki/index.php/Plugin:CPU) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle.
- *Disk*
  - Disk (https://collectd.org/wiki/index.php/Plugin:Disk) shows statistics on I/O, percent busy, latency, operations per second, pending I/O requests, and disk temperature.
- *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. It is 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 compressed physical ARC size, hit ratio, demand data, demand metadata, prefetch data, and prefetch metadata.

Reporting data is saved to permit viewing and monitoring usage trends over time. This data is preserved across system upgrades and restarts.

Data files are saved in `/var/db/collectd/rrd/`.

The reporting data file recording method is controlled by the *System* → *System Dataset Reporting database* option. When deselected, data files are recorded in a temporary filesystem and copied hourly to on-disk files.

When *System* → *System Dataset Reporting database* is enabled, data files are written directly to the *System Dataset* (page 35).

Warning: Reporting data is frequently written and should not be stored on the boot pool or operating system device.

Use the magnifier buttons next to each graph to increase or decrease the displayed time increment from 10 minutes, hourly, daily, weekly, or monthly. The << and >> buttons can be used to scroll through the output.

Update on using Graphite with FreeNAS (http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html) contains instructions for sending the collected information to a Graphite (http://graphiteapp.org/) server.

CHAPTER

THIRTEEN

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.

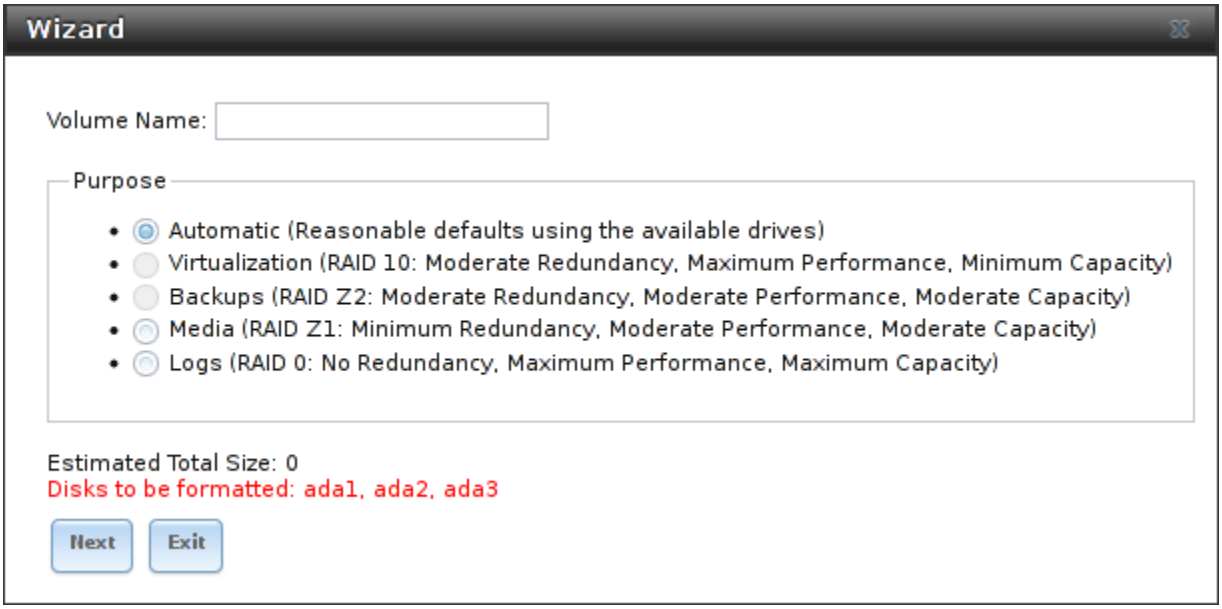

Fig. 13.2: Volume Creation Wizard

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 Volume* (page 108) 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 (https://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 250) 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 95).

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, or NIS server, click *Next* to skip to the next screen.

However, if the TrueNAS® system is on a network containing an Active Directory, LDAP, or NIS server and you wish to import the users and groups from that server, select the type of directory service in the *Directory Service* drop-down menu. The rest of the fields in this screen will vary, depending upon which directory service is selected. Available configuration options for each directory service are summarized in Tables 13.1 through 13.3.

Note: Additional configuration options are available for each directory service. The wizard can be used to set the initial values required to connect to that directory service. You can then review the other available options in *Directory Services* (page 141) to determine if additional configuration is required.

Table 13.1: Active Directory Options

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

Table 13.2: LDAP Options

Setting	Value	Description
Hostname	string	Hostname or IP address of LDAP server.

Continued on next page

Table 13.2 – continued from previous page

Setting	Value	Description
Base DN	string	Top level of the LDAP directory tree to be used when searching for resources. Example: *dc=test,dc=org*
Bind DN	string	Name of the administrative account on the LDAP server. Example: *cn=Manager,dc=test,dc=org*)
Base password	string	Password for the administrative account on the LDAP server.

Table 13.3: NIS Options

Setting	Value	Description
NIS domain	string	Name of the NIS domain.
NIS servers	string	Enter a comma-delimited list of hostnames or IP addresses.
Secure mode	checkbox	Set for ypbind(8) (https://www.freebsd.org/cgi/man.cgi?query=ypbind) to refuse to bind to any NIS server that is not running as root on a TCP port number over *1024*.
Manycast	checkbox	Set for *ypbind* to bind to the server that responds the fastest. This is useful when no local NIS server is available on the same subnet.

The next configuration screen, shown in Figure 13.5, is used to create network shares.

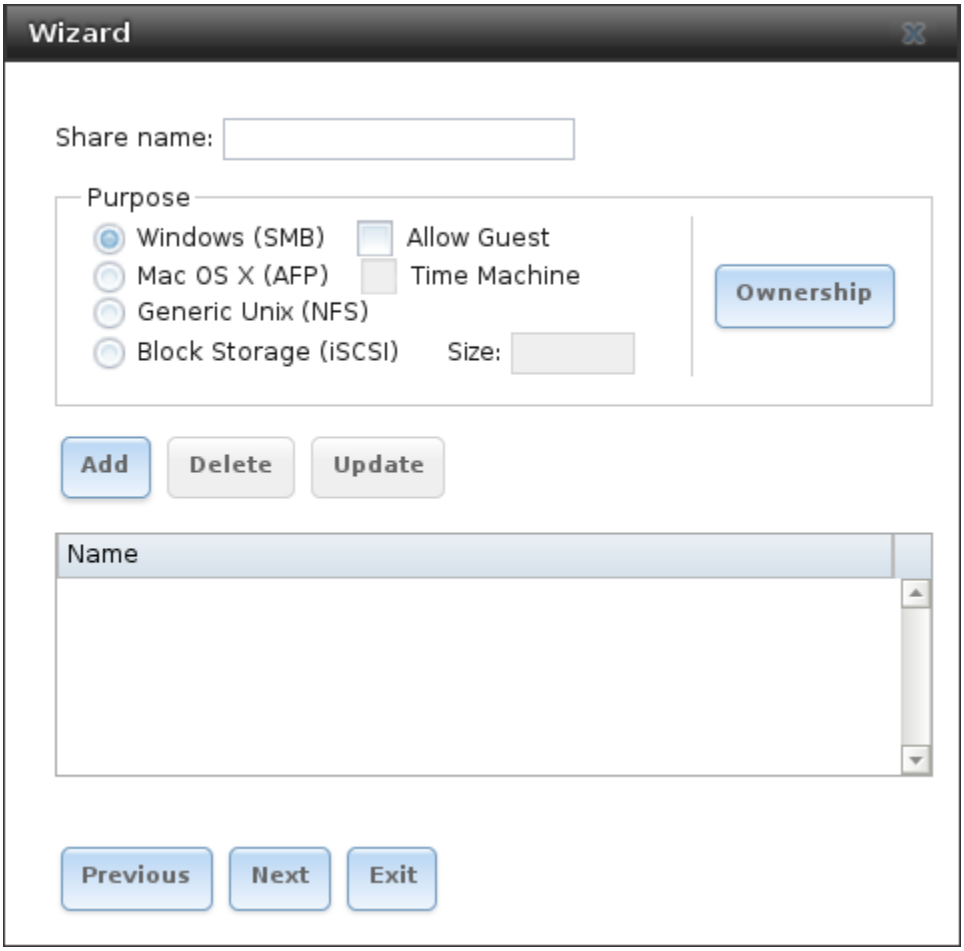

Fig. 13.5: Network Shares

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 153).

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 166).
- **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 154).
- **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 158).
- **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 GiB). iSCSI shares created with the wizard can be fine-tuned afterward with *iSCSI* (page 210).

After selecting the *Purpose*, click the *Ownership* button to see the screen shown in Figure 13.6.

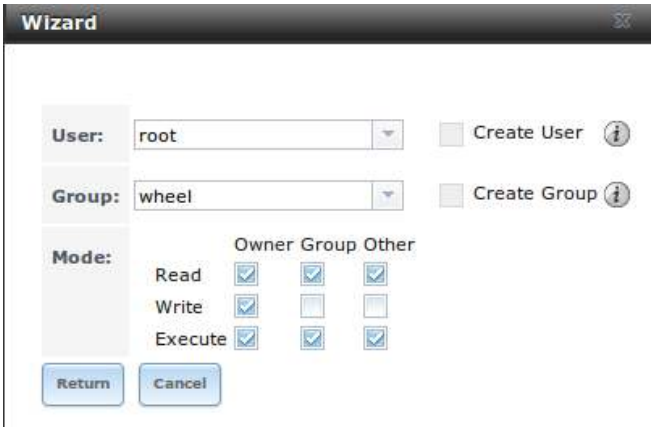

Fig. 13.6: Share Permissions

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.

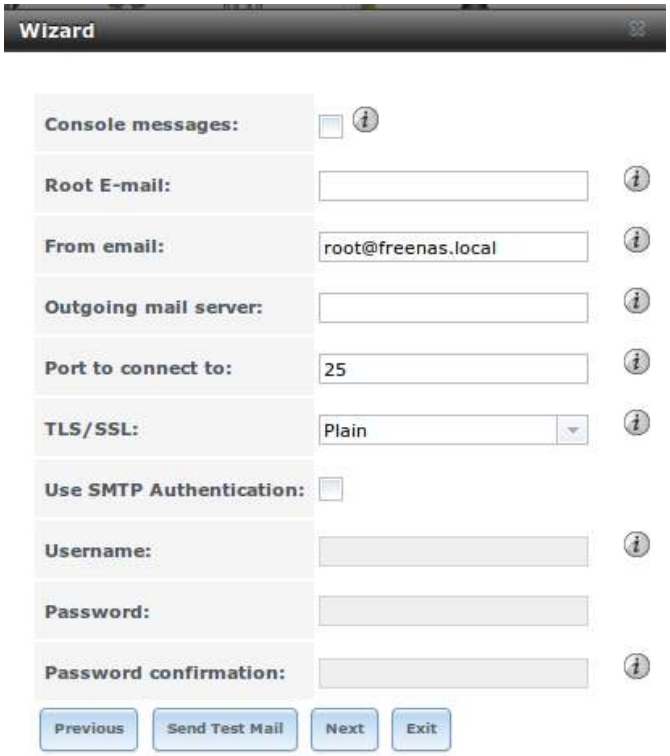

Fig. 13.7: Miscellaneous Settings

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 77), 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**.

CHAPTER

FOURTEEN

ADDITIONAL OPTIONS

This section covers the remaining miscellaneous options available from the TrueNAS® graphical administrative interface.

14.1 Display System Processes

Clicking *Display System Processes* opens a screen showing the output of top(1) (https://www.freebsd.org/cgi/man.cgi?query=top). An example is shown in Figure 14.1.

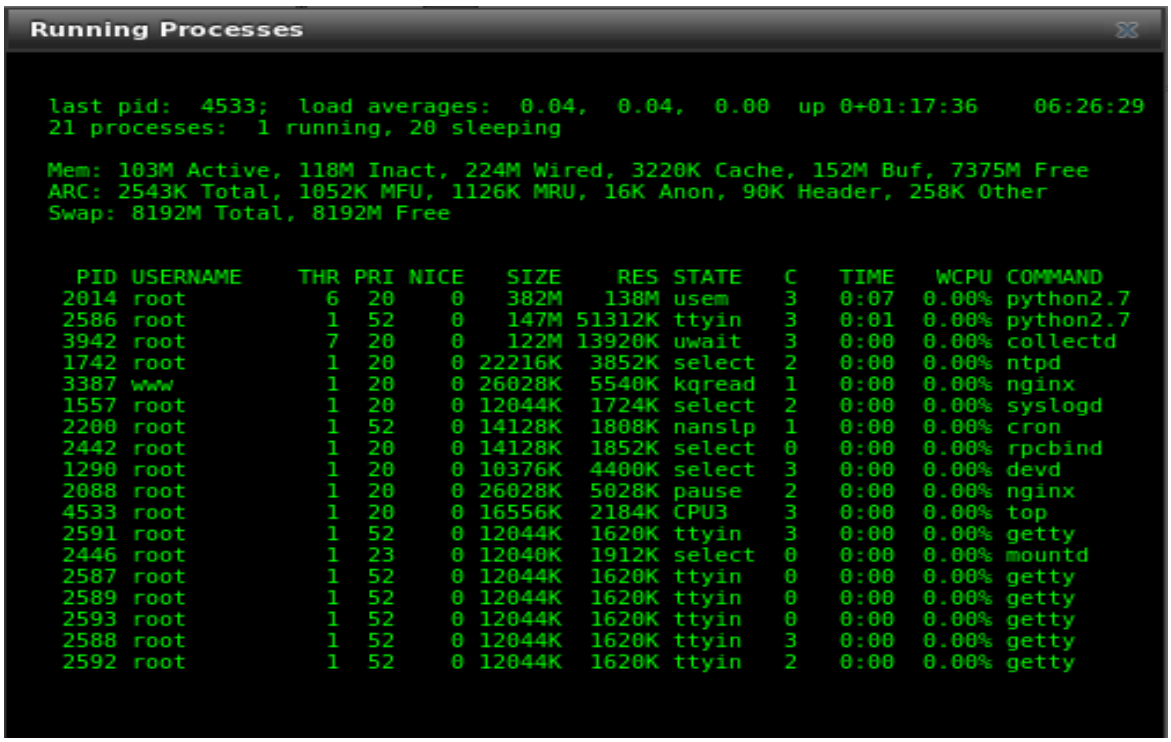

Fig. 14.1: System Processes Running on TrueNAS®

The display automatically refreshes itself. Click the *X* in the upper right corner to close the display when finished. This display is read-only, so it is not possible to give a `kill` command in it.

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.

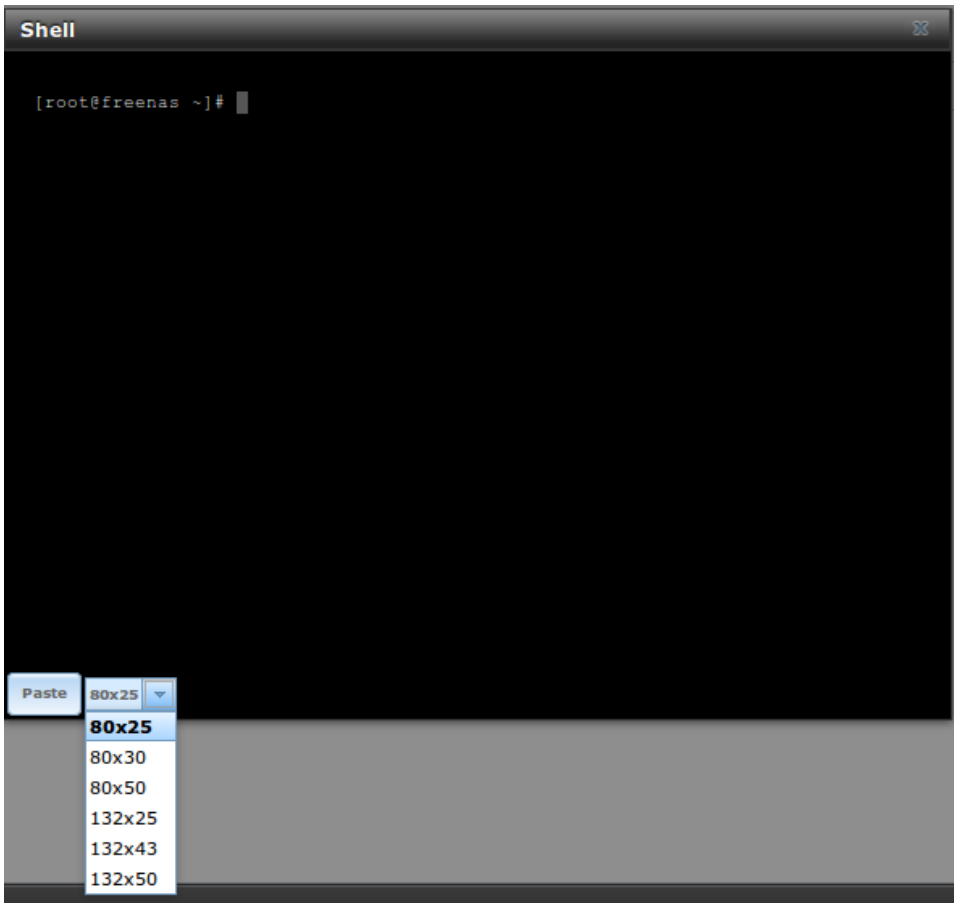

Fig. 14.2: Web Shell

The prompt indicates that the current user is *root*, the hostname is *truenas*, and the current working directory is ~ (*root*'s home directory).

To change the size of the shell, click the *80x25* drop-down menu and select a different size.

To copy text from shell, highlight the text, right-click, and select Copy from the right-click menu. To paste into the shell, click the *Paste* button, paste the text into the box that opens, and click the *OK* button to complete the paste operation.

While you are in Shell, you will not have access to any of the other GUI menus. If you need to have access to a prompt while using the GUI menus, use `tmux` instead as it supports multiple shell sessions and the detachment and reattachment of sessions.

Shell provides history (use your up arrow to see previously entered commands and press `Enter` to repeat the currently displayed command) and tab completion (type a few letters and press tab to complete a command name or filename in the current directory). When you are finished using Shell, type `exit` to leave the session.

Note: Not all of Shell's features render correctly in Chrome. Firefox is the recommended browser for using Shell.

Most FreeBSD command line utilities are available in Shell.

14.3 Log Out

Click the *Log Out* entry in the tree to log out of the TrueNAS® GUI. This causes an immediate logout. A message is displayed with a link to log back in.

14.4 Reboot

Click *Reboot* shows the warning message in Figure 14.3. The browser window background color changes to red to indicate that this option can 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 asks if you wish to proceed. In this case, it is recommended to *Cancel* the reboot request and to periodically run `zpool status` from *Shell* (page 244) until it is verified that the scrub or resilver process is complete. Once complete, the reboot request can be re-issued.

Click the *Cancel* button to cancel the reboot request. Otherwise, click the *Reboot* button to reboot the system. Rebooting the system disconnects all clients, including the web administration GUI. The URL in the web browser changes, adding `/system/reboot/` to the end of the IP address. Wait a few minutes for the system to boot, then use the browser's Back button to return to the TrueNAS® system's IP address and display the GUI login screen. If the login screen does not appear, access the system using IPMI to determine if a problem is preventing the system from resuming normal operation.

14.5 Shutdown

Clicking *Shutdown* shows the warning message in Figure 14.4. The browser window background color changes to red to indicate that this is an option that will negatively impact users of the TrueNAS® system.

Fig. 14.4: Shutdown Warning Message

If a scrub or resilver is in progress when a shutdown is requested, an additional warning will ask for confirmation to proceed. In this case, it is recommended to *Cancel* the shutdown request and to periodically run `zpool status` from *Shell* (page 244) until it is verified that the scrub or resilver process is complete. Once complete, the shutdown request can be re-issued.

On High Availability (HA) systems with *Failover* (page 58), an additional checkbox is provided to shut down the standby node.

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

14.6 Support Icon

The *Support* icon, the first icon on the right side of the menubar, provides a shortcut to *System* → *Support*. This screen can be used to verify the system license or to create a support ticket. Refer to *Support* (page 54) for detailed usage instructions.

14.7 Guide

The *Guide* icon, the second icon in the top menubar, links to the online version of the TrueNAS® User Guide (this documentation).

Fig. 14.5: User Guide Menu

14.8 Alert

The TrueNAS® alert system provides 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 14.6, the system is warning that the S.M.A.R.T. service is not running.

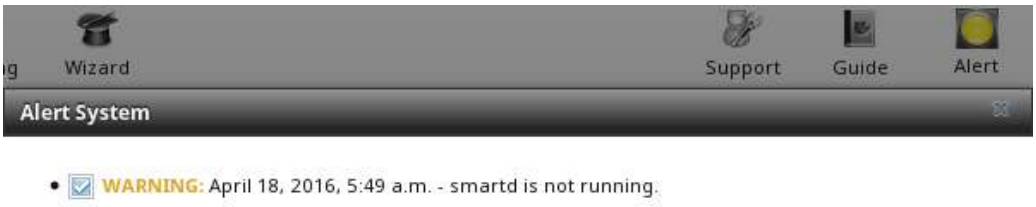

Fig. 14.6: Example Alert Message

Informational messages have a green *OK*, warning messages flash yellow, and messages requiring attention are listed as a red *CRITICAL*. CRITICAL messages are also emailed to the root user account. To remove the flashing alert for a message, deselect the option next to it.

Behind the scenes, an alert daemon checks for various alert conditions, such as volume and disk status, and writes the current conditions to `/var/tmp/alert`. The daemon retrieves the current alert status every minute and changes the solid green alert icon to flashing red when a new alert is detected.

Current alerts are viewed from the Shell option of the Console Setup Menu (Figure 2.1) or from the Web Shell (Figure 14.2) by running `alertcli.py`. Alert messages indicate which *High Availability (HA)* (page 58) node generated the alert.

Some of the conditions that trigger an alert include:

- used space on a volume, dataset, or zvol goes over 80%; the alert goes red at 95%
- new *ZFS Feature Flags* (page 253) are available for the pool; this alert can be unchecked if a pool upgrade is not desired at present
- a new update is available
- the system reboots itself
- non-optimal multipath states are detected
- ZFS pool status changes from *HEALTHY*
- a S.M.A.R.T. error occurs
- syslog-ng(8) (https://www.freebsd.org/cgi/man.cgi?query=syslog-ng) is not running
- the system is unable to bind to the *WebGUI IPv4 Address* set in *System* → *General*
- the system can not find an IP address configured on an iSCSI portal
- the NTP server cannot be contacted
- a periodic snapshot or replication task fails
- a VMware login or a *VMware-Snapshot* (page 139) task fails
- deleting a VMware snapshot fails
- a Certificate Authority or certificate is invalid or malformed
- an update failed, or the system needs to reboot to complete a successful update
- a re-key operation fails on an encrypted pool
- LDAP failed to bind to the domain
- any member interfaces of a lagg interface are not active
- a scrub is paused
- a Fibre Channel (FC) Host Bus Adapter (HBA) configured as an iSCSI target is not detected
- the interface which is set as critical for failover is not found or is not configured
- NVDIMM problems
- 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
- Storage controllers do not have the same number of connected disks
- the boot volume of the passive node is not HEALTHY
- 30 days before the license expires, and when the license expires
- the usage of a HA link goes above 10MB/s
- an IPMI query to a standby node fails, indicating the standby node is down
- *Proactive Support* (page 56) is enabled but any of the configuration fields are empty
- ticket creation fails while Proactive Support is enabled
- if VMware failed to log in (usually preceding a VMware snapshot)
- if an unlicensed expansion shelf is connected
- if a USB storage device has been attached which could prevent booting or failover
- when the passive node cannot be contacted
- when it is 180, 90, 30, or 14 days before support contract expiration

Note: If *Proactive Support* (page 56) is enabled with Silver or Gold support coverage, and there is an internet connection, alerts which can indicate a hardware issue automatically create a support ticket with iXsystems Support. These alerts include a ZFS pool status change, a multipath failure, a failed S.M.A.R.T. test, and a failed re-key operation.

CHAPTER

FIFTEEN

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) to provide continued, collaborative development of the open source version.

Here is an overview of the features provided by ZFS:

ZFS is a transactional, Copy-On-Write (COW) (https://en.wikipedia.org/wiki/ZFS#Copy-on-write_transactional_model) filesystem. For each write request, a copy is made of the associated disk blocks and all changes are made to the copy rather than to the original blocks. When the write is complete, all block pointers are changed to point to the new copy. This means that ZFS always writes to free space, most writes are sequential, and old versions of files are not unlinked until a complete new version has been written successfully. ZFS has direct access to disks and bundles multiple read and write requests into transactions. Most filesystems cannot do this, as they only have access to disk blocks. A transaction either completes or fails, meaning there will never be a write-hole (https://blogs.oracle.com/bonwick/raid-z) and a filesystem checker utility is not necessary. Because of the transactional design, as additional storage capacity is added, it becomes immediately available for writes. To rebalance the data, one can copy it to re-write the existing data across all available disks. As a 128-bit filesystem, the maximum filesystem or file size is 16 exabytes.

ZFS was designed to be a self-healing filesystem. As ZFS writes data, it creates a checksum for each disk block it writes. As ZFS reads data, it validates the checksum for each disk block it reads. Media errors or "bit rot" can cause data to change, and the checksum no longer matches. When ZFS identifies a disk block checksum error on a pool that is mirrored or uses RAIDZ, it replaces the corrupted data with the correct data. Since some disk blocks are rarely read, regular scrubs should be scheduled so that ZFS can read all of the data blocks to validate their checksums and correct any corrupted blocks. While multiple disks are required in order to provide redundancy and data correction, ZFS will still provide data corruption detection to a system with one disk. TrueNAS® automatically schedules a monthly scrub for each ZFS pool and the results of the scrub are displayed by selecting the *Volume* (page 94) and clicking *Volume Status*. Checking scrub results provides 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 95) is 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. ZFS pools created on TrueNAS® version 9.2.1 or later use the recommended LZ4 compression algorithm.

ZFS provides low-cost, instantaneous snapshots of the specified pool, dataset, or zvol. Due to COW, snapshots initially take no additional space. The size of a snapshot increases over time as changes to the files in the snapshot are written to disk. Snapshots can be used to provide a copy of data at the point in time the snapshot was created. When a file is deleted, its disk blocks are added to the free list; however, the blocks for that file in any existing

snapshots are not added to the free list until all referencing snapshots are removed. This makes snapshots a clever way to keep a history of files, useful for recovering an older copy of a file or a deleted file. For this reason, many administrators take snapshots often, store them for a period of time, and store them on another system. Such a strategy allows the administrator to roll the system back to a specific time. If there is a catastrophic loss, an off-site snapshot can restore the system up to the last snapshot interval, within 15 minutes of the data loss, for example. Snapshots are stored locally but can also be replicated to a remote ZFS pool. During replication, ZFS does not do a byte-for-byte copy but instead converts a snapshot into a stream of data. This design means that the ZFS pool on the receiving end does not need to be identical and can use a different RAIDZ level, 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 (http://www.freenas.org/blog/zfs-zil-and-slog-demystified/)). 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/)
- 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/)
- ZFS Intent Log (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 244) to determine if the system will benefit from a SLOG. See this website (http://www.richardelling.com/Home/scripts-and-programs-1/zilstat) for usage information.

ZFS currently uses 16 GiB of space for SLOG. Larger SSDs can be installed, but the extra space will not be used. SLOG devices cannot be shared between pools. Each pool requires a separate SLOG device. Bandwidth and throughput limitations require that a SLOG device must only be used for this single purpose. Do not attempt to add other caching functions on the same SSD, or performance will suffer.

In mission-critical systems, a mirrored SLOG device is highly recommended. Mirrored SLOG devices are *required* for ZFS pools at ZFS version 19 or earlier. The ZFS pool version is checked from the *Shell* (page 244) with `zpool get version poolname`. A version value of - means the ZFS pool is version 5000 (also known as *Feature Flags*) or later.

ZFS provides a read cache in RAM, known as the ARC, which reduces read latency. TrueNAS® adds ARC stats to top(1) (https://www.freebsd.org/cgi/man.cgi?query=top) and includes the `arc_summary.py` and `arcstat.py` tools for monitoring the efficiency of the ARC. If an SSD is dedicated as a cache device, it is known as an L2ARC (http://www.brendangregg.com/blog/2008-07-22/zfs-l2arc.html). Additional read data is cached here, which can increase random read performance. L2ARC does *not* reduce the need for sufficient RAM. In fact, L2ARC needs RAM to function. If there is not enough RAM for a adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of *random* reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 32 GiB of RAM, and the size of an L2ARC should not exceed ten times the amount of RAM. In some cases, it may be more efficient to have two separate pools: one on SSDs for active data, and another on hard drives for rarely used content. After adding an L2ARC device, monitor its effectiveness using tools such as `arcstat`. To increase the size of an existing L2ARC, stripe another cache device with it. The GUI will always stripe L2ARC, not mirror it, as the contents of L2ARC are recreated at boot. Failure of an individual SSD from an L2ARC pool will not affect the integrity of the pool, but may have an impact on read performance, depending

on the workload and the ratio of dataset size to cache size. Note that dedicated L2ARC devices cannot be shared between ZFS pools.

ZFS was designed to provide redundancy while addressing some of the inherent limitations of hardware RAID such as the write-hole and corrupt data written over time before the hardware controller provides an alert. ZFS provides three levels of redundancy, known as *RAIDZ*, where the number after the *RAIDZ* indicates how many disks per vdev can be lost without losing data. ZFS also supports mirrors, with no restrictions on the number of disks in the mirror. ZFS was designed for commodity disks so no RAID controller is needed. While ZFS can also be used with a RAID controller, it is recommended that the controller be put into JBOD mode so that ZFS has full control of the disks.

When determining the type of ZFS redundancy to use, consider whether the goal is to maximize disk space or performance:

- RAIDZ1 maximizes disk space and generally performs well when data is written and read in large chunks (128K or more).
- RAIDZ2 offers better data availability and significantly better mean time to data loss (MTTDL) than RAIDZ1.
- A mirror consumes more disk space but generally performs better with small random reads. For better performance, a mirror is strongly favored over any RAIDZ, particularly for large, uncacheable, random read loads.
- Using more than 12 disks per vdev is not recommended. The recommended number of disks per vdev is between 3 and 9. With more disks, use multiple vdevs.
- Some older ZFS documentation recommends that a certain number of disks is needed for each type of RAIDZ in order to achieve optimal performance. On systems using LZ4 compression, which is the default for TrueNAS® 9.2.1 and higher, this is no longer true.

These resources can also help determine the RAID configuration best suited to the specific storage requirements:

- Getting the Most out of ZFS Pools (https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/)
- A Closer Look at ZFS, Vdevs and Performance (https://constantin.glez.de/2010/06/04/a-closer-look-zfs-vdevs-and-performance/)

Warning: RAID AND DISK REDUNDANCY ARE NOT A SUBSTITUTE FOR A RELIABLE BACKUP STRATEGY. BAD THINGS HAPPEN AND A GOOD BACKUP STRATEGY IS STILL REQUIRED TO PROTECT VALUABLE DATA. See *Periodic Snapshot Tasks* (page 120) and *Replication Tasks* (page 122) 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%.
- When considering the number of disks to use per vdev, consider the size of the disks and the amount of time required for resilvering, which is the process of rebuilding the vdev. The larger the size of the vdev, the longer the resilvering time. When replacing a disk in a RAIDZ, it is possible that another disk will fail before the resilvering process completes. If the number of failed disks exceeds the number allowed per vdev for the type of RAIDZ, the data in the pool will be lost. For this reason, RAIDZ1 is not recommended for drives over 1 TiB in size.
- Using drives of equal sizes is recommended when creating a vdev. While ZFS can create a vdev using disks of differing sizes, its capacity will be limited by the size of the smallest disk.

For those new to ZFS, the Wikipedia entry on ZFS (https://en.wikipedia.org/wiki/Zfs) 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)
- ZFS Administration Guide (https://docs.oracle.com/cd/E19253-01/819-5461/index.html)
- Becoming a ZFS Ninja (video) (https://www.youtube.com/watch?v=6_K55Ira1Cs)
- Slideshow explaining VDev, zpool, ZIL and L2ARC and other newbie mistakes! (https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-for-noobs.7775/)
- A Crash Course on ZFS (http://www.bsdnow.tv/tutorials/zfs)
- ZFS: The Last Word in File Systems - Part 1 (video) (https://www.youtube.com/watch?v=uT2i2ryhCio)
- The Zettabyte Filesystem (https://www.youtube.com/watch?v=ptY6-K78McY)

15.1 ZFS Feature Flags

To differentiate itself from Oracle ZFS version numbers, OpenZFS uses feature flags. Feature flags are used to tag features with unique names to provide portability between OpenZFS implementations running on different platforms, as long as all of the feature flags enabled on the ZFS pool are supported by both platforms. TrueNAS® uses OpenZFS and each new version of TrueNAS® keeps up-to-date with the latest feature flags and OpenZFS bug fixes.

See zpool-features(7) (https://www.freebsd.org/cgi/man.cgi?query=zpool-features) for a complete listing of all OpenZFS feature flags available on FreeBSD.

CHAPTER

SIXTEEN

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.

16.1 VAAI for iSCSI

VAAI for iSCSI supports these operations:

- *Atomic Test and Set* (*ATS*) allows multiple initiators to synchronize LUN access in a fine-grained manner rather than locking the whole LUN and preventing other hosts from accessing the same LUN simultaneously.
- *Clone Blocks* (*XCOPY*) copies disk blocks on the NAS. Copies occur locally rather than over the network. The operation is similar to Microsoft ODX (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)).
- *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.

CHAPTER

SEVENTEEN

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 17.1, an online version of the API is available at api.freenas.org (http://api.freenas.org).

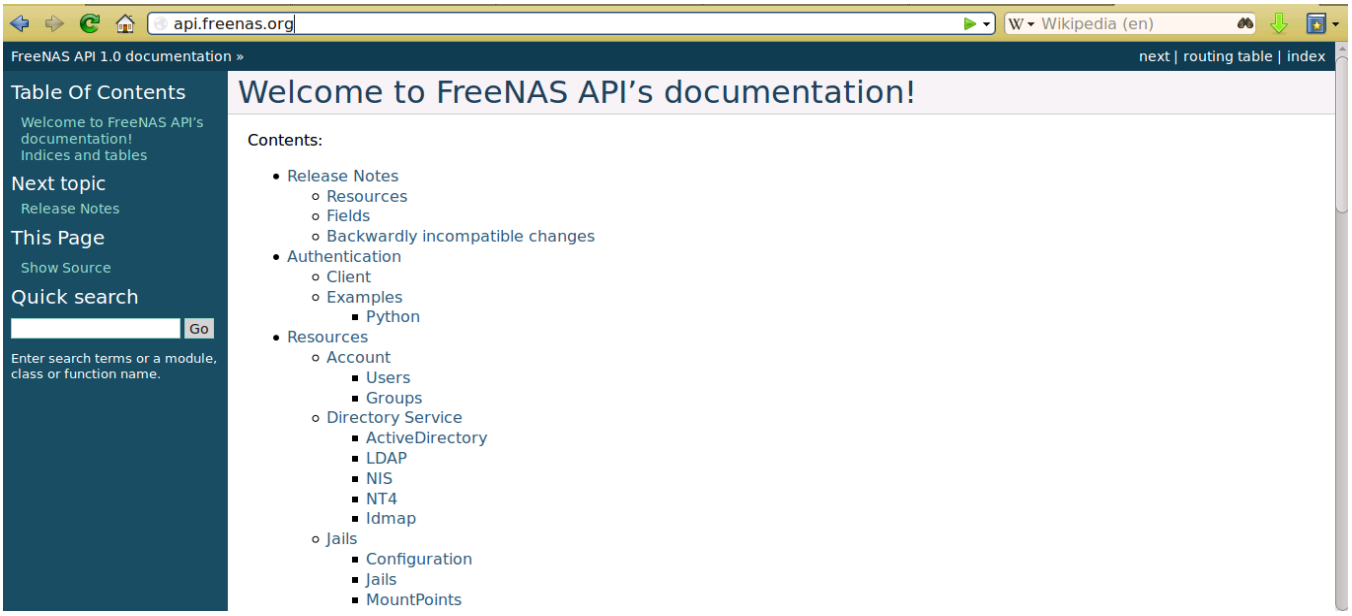

Fig. 17.1: API Documentation

17.1 APIv2

A new API was released with TrueNAS® 11.1. The previous API is still present and in use because it is feature-complete. Documentation for the new API is available on the TrueNAS® system at the */api/docs/* URL. For example, if the TrueNAS® system is at IP address 192.168.1.119, enter *http://192.168.1.119/api/docs/* in a browser to see the API documentation. Work is under way to make the new API feature-complete. The new APIv2 uses WebSockets (https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API).

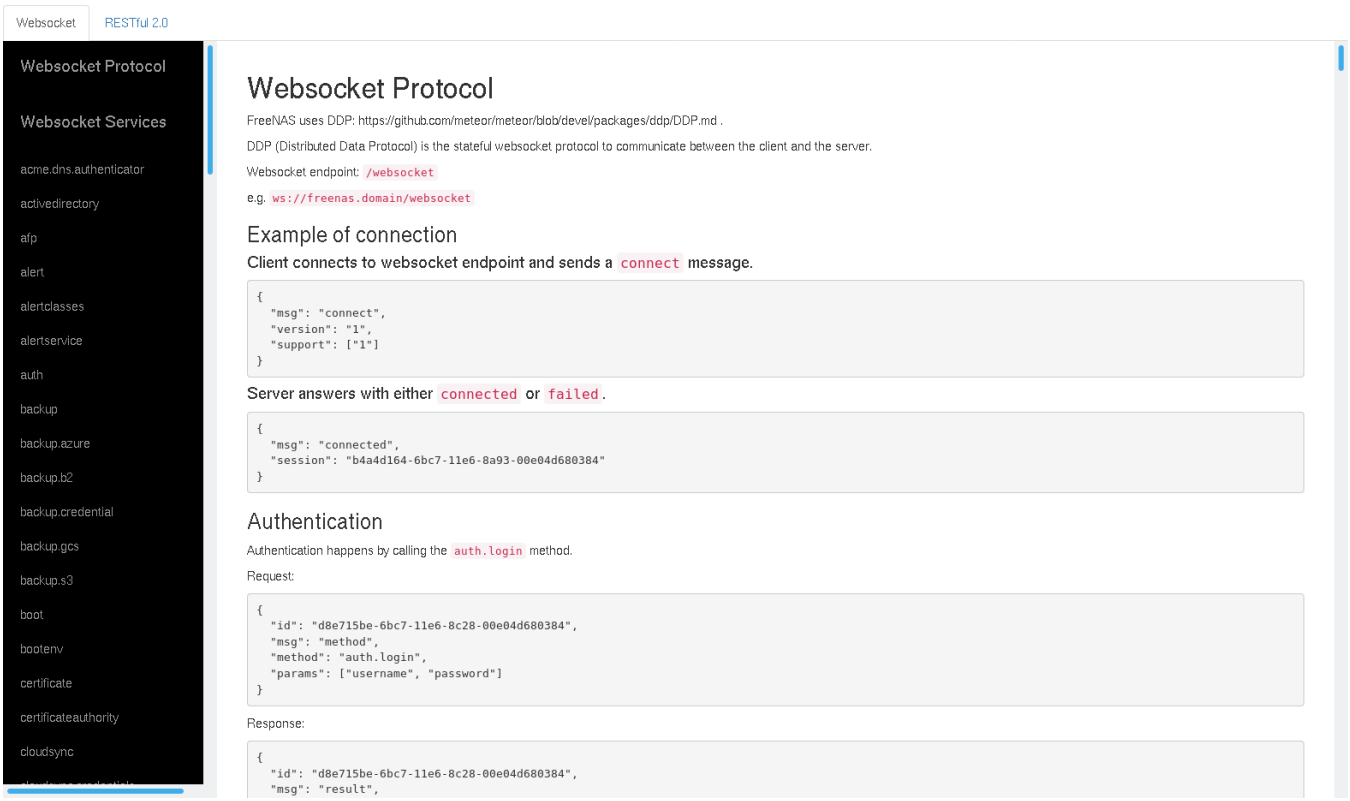

Fig. 17.2: APIv2 Documentation

This advanced technology makes it possible to open interactive communication sessions between web browsers and servers, allowing event-driven responses without the need to poll the server for a reply. When APIv2 is feature complete, the TrueNAS® documentation will include relevant examples that make use of the new API.

17.2 A Simple API Example

The api directory of the FreeNAS® github repository (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) 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 17.1 summarizes acceptable values. This resource uses JSON, so the boolean values are *True* or *False*.

Table 17.1: JSON Parameters for Users Create Resource

JSON Parameter	Type	Description
bsdusr_username	string	Enter a maximum of 32 characters. A maximum of 8 is recommended for interoperability. The username can include numerals but cannot include a space.
bsdusr_full_name	string	This field can contain spaces and uppercase characters.
bsdusr_password	string	The password 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	Specify the numeric ID of the group to create if *bsdusr_creategroup* is set to *False*.
bsdusr_creategroup	boolean	Set to *True* to create a primary group with the same numeric ID as *bsdusr_uid*.
bsdusr_mode	string	Sets default numeric UNIX permissions for the home directory of the user.
bsdusr_shell	string	Specify the full path to a UNIX shell that is installed on the system.
bsdusr_password_disabled	boolean	The user is not allowed to log in when set to *True*.
bsdusr_locked	boolean	The user is not allowed to log in when set to *True*.
bsdusr_sudo	boolean	`sudo` is enabled for the user when set to *True*.
bsdusr_sshpubkey	string	Enter the contents of the 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*.

17.3 A More Complex Example

This section provides a walk-through of a more complex example found in the `startup.py` script. Use the search-bar within the API documentation to quickly locate the JSON parameters used here. This example defines a class and several methods to create a ZFS 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:

```
class Startup(object):
  def __init__(self, hostname, user, secret):
        self._hostname = hostname
        self._user = user
        self._secret = secret
        self._ep = 'http://%s/api/v1.0' % hostname
  def request(self, resource, method='GET', data=None):
        if data is None:
            data = ''
        r = requests.request(
            method,
            '%s/%s/' % (self._ep, resource),
            data=json.dumps(data),
            headers={'Content-Type': "application/json"},
            auth=(self._user, self._secret),
        )
        if r.ok:
            try:
                return r.json()
            except:
                return r.text
        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.:

```
def _get_disks(self):
        disks = self.request('storage/disk')
        return [disk['disk_name'] for disk in disks]

def create_pool(self):
        disks = self._get_disks()
        self.request('storage/volume', method='POST', data={
            'volume_name': 'tank',
            'layout': [
                {'vdevtype': 'stripe', 'disks': disks},
            ],
})
```

The *create_dataset* method is defined which creates a dataset named `MyShare`:

```
def create_dataset(self):
        self.request('storage/volume/tank/datasets', method='POST', data={
            'name': 'MyShare',
        })
```

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.:

```
def create_cifs_share(self):
       self.request('sharing/cifs', method='POST', data={
           'cifs_name': 'My Test Share',
           'cifs_path': '/mnt/tank/MyShare',
           'cifs_guestonly': True
})
```

Finally, the *service_start* method enables the CIFS service. The *srv_enable* JSON parameter is described in the Services resource.

```
def service_start(self, name):
       self.request('services/services/%s' % name, method='PUT', data={
           'srv_enable': True,

})
```

CHAPTER

EIGHTEEN

APPENDIX A: TRUENAS SOFTWARE END USER LICENSE AGREEMENT

TrueNAS® EULA:

Important - Please Read This EULA Carefully

PLEASE CAREFULLY READ THIS END USER LICENSE AGREEMENT (EULA) BEFORE CLICKING THE AGREE BUTTON. THIS AGREEMENT SERVES AS A LEGALLY BINDING DOCUMENT BETWEEN YOU AND IXSYSTEMS, INC. BY CLICKING THE AGREE BUTTON, DOWNLOADING, INSTALLING, OR OTHERWISE USING TRUENAS SOFTWARE, YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS AGREEMENT). IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS IN THIS AGREEMENT, DO NOT USE OR INSTALL TRUENAS SOFTWARE.

This agreement is provided in accordance with the Commercial Arbitration Rules of the American Arbitration Association (the "AAA Rules") under confidential binding arbitration held in Santa Clara County, California. To the fullest extent permitted by applicable law, no arbitration under this EULA will be joined to an arbitration involving any other party subject to this EULA, whether through class arbitration proceedings or otherwise. Any litigation relating to this EULA shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California. All matters arising out of or relating to this agreement shall be governed by and construed in accordance with the internal laws of the State of California without giving effect to any choice or conflict of law provision or rule.

1.0 Definitions

1.1 "Company", **"iXsystems"** and **"iX"** means iXsystems, Inc., on behalf of themselves, subsidiaries, and affiliates under common control.

1.2 "TrueNAS Software" means the TrueNAS storage management software.

1.3 "TrueNAS Device" means the TrueNAS hardware storage appliances and peripheral equipment.

1.4 "Product" means, individually and collectively, the TrueNAS Software and the TrueNAS Device.

1.5 "Open Source Software" means various open source software components licensed under the terms of applicable open source license agreements, each of which has its own copyright and its own applicable license terms.

1.6 "Licensee", **"You"** and **"Your"** refers to the person, organization, or entity that has agreed to be bound by this EULA including any employees, affiliates, and third party contractors that provide services to You.

1.6 "Agreement" refers to this document, the TrueNAS End User License Agreement.

2.0 License

Subject to the terms set forth in this Agreement, iXsystems grants You a non-exclusive, non-transferable, revocable, limited license without the option to sublicense, to use TrueNAS Software on Your TrueNAS Device(s) in accordance with Your authorized purchase and use of a TrueNAS Device(s) for Your internal business purposes. This use includes but is not limited to using or viewing the instructions, specifications, and documentation provided with the Product.

3.0 License Restrictions

TrueNAS Software is only authorized for use with a TrueNAS Device identified by a specific serial number and manufactured by iXsystems. This license may be extended to a second TrueNAS Device if an additional TrueNAS Device was purchased for high availability data protection. The Product, including the TrueNAS Software, is protected by copyright laws and international treaties, as well as other intellectual property laws, statutes, and treaties. The Product is licensed, not sold to You the end user. You do not acquire any ownership interest in the Product, including

TrueNAS Software, or any other rights to such Product, other than to use such Product in accordance with the license granted under this Agreement, subject to all terms, conditions, and restrictions. iXsystems reserves and shall retain its entire right, title, and interest in and to the Product, and all intellectual property rights arising out of or relating to the Product, subject to the license expressly granted to You in this Agreement.

The Product, including TrueNAS Software, may contain iXsystems' trademarks, trade secrets, and proprietary collateral. iXsystems strictly prohibits the acts of decompiling, reverse engineering, or disassembly of the Product, including TrueNAS Software. You agree to use commercially reasonable efforts to safeguard the Product and iXsystems' intellectual property, trade secrets, or other proprietary information You may have access to, from infringement, misappropriation, theft, misuse, or unauthorized access. You will promptly notify iXsystems if You become aware of any infringement of the Product and cooperate with iXsystems in any legal action taken by iXsystems to enforce its intellectual property rights. By accepting this Agreement, You agree You will not disclose, copy, transfer, or publish benchmark results relating to the Product without the express written consent of iXsystems. You agree not to use, or permit others to use, the Product beyond the scope of the license granted under Section 2, unless otherwise permitted by iXsystems, or in violation of any law, regulation or rule, and you will not modify, adapt, or otherwise create derivative works or improvements of the Product. You are responsible and liable for all uses of the Product through access thereto provided by You, directly or indirectly.

4.0 General

4.1 Entire Agreement - This Agreement, together with any associated purchase order, service level agreement, and all other documents and policies referenced herein, constitutes the entire and only agreement between You and iXsystems for use of the TrueNAS Software and the TrueNAS Device and all other prior negotiations, representations, agreements, and understandings are superseded hereby. No agreements altering or supplementing the terms hereof may be made except by means of a written document signed by Your duly authorized representatives and those of iXsystems.

4.2 Waiver and Modification - No failure of either party to exercise or enforce any of its rights under this EULA will act as a waiver of those rights. This EULA may only be modified, or any rights under it waived, by a written document executed by the party against which it is asserted.

4.3 Severability - If any provision of this EULA is found illegal or unenforceable, it will be enforced to the maximum extent permissible, and the legality and enforceability of the other provisions of this EULA will not be affected.

4.4 United States Government End Users - For any Product licensed directly or indirectly on behalf of a unit or agency of the United States Government, this paragraph applies. Company's proprietary software embodied in the Product: (a) was developed at private expense and is in all respects Company's proprietary information; (b) was not developed with government funds; (c) is Company's trade secret for all purposes of the Freedom of Information Act; (d) is a commercial item and thus, pursuant to Section 12.212 of the Federal Acquisition Regulations (FAR) and DFAR Supplement Section 227.7202, Government's use, duplication or disclosure of such software is subject to the restrictions set forth by the Company and Licensee shall receive only those rights with respect to the Product as are granted to all other end users.

4.5 Foreign Corrupt Practices Act - You will comply with the requirements of the United States Foreign Corrupt Practices Act (the "FCPA") and will refrain from making, directly or indirectly, any payments to third parties which constitute a breach of the FCPA. You will notify Company immediately upon Your becoming aware that such a payment has been made. You will indemnify and hold harmless Company from any breach of this provision.

4.6. Title - iXsystems retains all rights, titles, and interest in TrueNAS Software and in TrueNAS Device(s) and all related copyrights, trade secrets, patents, trademarks, and any other intellectual and industrial property and proprietary rights, including registrations, applications, registration keys, renewals, and extensions of such rights.

4.7 Contact Information - If You have any questions about this Agreement, or if You want to contact iXsystems for any reason, please email legal@ixsystems.com.

4.8 Maintenance and Support - You may be entitled to support services from iXsystems after purchasing a TrueNAS Device or a support contract. iXsystems will provide these support services based on the length of time of the purchased support contract. This maintenance and support is only valid for the length of time that You have purchased with Your TrueNAS Device. iXsystems may from time to time and at their sole discretion vary the terms and conditions of the maintenance and support agreement based on different business environmental and personnel factors. For more information on our Maintenance and Support contract, refer to https://ixsystems.com/TrueNAS_SLA.

4.9 Force Majeure - iXsystems will not be deemed to be in default of any of the provisions of this Agreement or be

liable for any delay or failure in performance due to Force Majeure, which shall include without limitation acts of God, earthquake, weather conditions, labor disputes, changes in law, regulation or government policy, riots, war, fire, epidemics, acts or omissions of vendors or suppliers, equipment failures, transportation difficulties, malicious or criminal acts of third parties, or other occurrences which are beyond iXsystems' reasonable control.

4.10 Termination - iXsystems may terminate or suspend Your license to use the Product and cease any and all support, services, or maintenance under this Agreement without prior notice, or liability, and for any reason whatsoever, without limitation, if any of the terms and conditions of this Agreement are breached. Upon termination, rights to use the Product will immediately cease. Other provisions of this Agreement will survive termination including, without limitation, ownership provisions, warranty disclaimers, indemnity, and limitations of liability.

4.11 Open Source Software Components - iXsystems uses Open Source Software components in the development of the Product. Open Source Software components that are used in the Product are composed of separate components each having their own trademarks, copyrights, and license conditions.

4.12 Assignment - Licensee shall not assign or otherwise transfer any of its rights, or delegate or otherwise transfer any of its obligations or performance, under this Agreement, in each case whether voluntarily, involuntarily, by operation of law, or otherwise, without iXsystems' prior written consent. No delegation or other transfer will relieve Licensee of any of its obligations or performance under this Agreement. Any purported assignment, delegation, or transfer in violation of this Section is void. iXsystems may freely assign or otherwise transfer all or any of its rights, or delegate or otherwise transfer all or any of its obligations or performance, under this Agreement without Licensee's consent. This Agreement is binding upon and inures to the benefit of the parties hereto and their respective permitted successors and assigns.

5.0 Export Control Regulations

The Product may be subject to US export control laws, including the US Export Administration Act and its associated regulations. You shall not, directly or indirectly, export, re-export, or release the Product to, or make the Product accessible from, any jurisdiction or country to which export, re-export, or release is prohibited by law, rule, or regulation. You shall comply with all applicable federal laws, regulations, and rules, and complete all required undertakings (including obtaining any necessary export license or other governmental approval), prior to exporting, re-exporting, releasing, or otherwise making the Product available outside the US.

6.0 Data Collection and Privacy

TrueNAS Software may collect information relating to Your use of the Product, including information that has been provided directly or indirectly through automated means. Usage of TrueNAS Software, geolocation information, user login credentials, and device and operating system identification are allowed according to iXsystems' privacy policy (https://www.ixsystems.com/privacy-policy/). By accepting this Agreement and continuing to use the Product, you agree that iXsystems may use any information provided through direct or indirect means in accordance with our privacy policy and as permitted by applicable law, for purposes relating to management, compliance, marketing, support, security, update delivery, and product improvement.

7.0 Limitation of Liability and Disclaimer of Warranty

THE PRODUCT IS PROVIDED "AS IS" AND WITH ALL FAULTS AND DEFECTS WITHOUT WARRANTY OF ANY KIND. TO THE MAXIMUM EXTENT PERMITTED UNDER APPLICABLE LAW, IXSYSTEMS, ON ITS OWN BEHALF AND ON BEHALF OF ITS AFFILIATES AND ITS AND THEIR RESPECTIVE LICENSORS AND SERVICE PROVIDERS, EXPRESSLY DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE, WITH RESPECT TO THE PRODUCT, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NON-INFRINGEMENT, AND WARRANTIES THAT MAY ARISE OUT OF COURSE OF DEALING, COURSE OF PERFORMANCE, USAGE, OR TRADE PRACTICE. WITHOUT LIMITATION TO THE FOREGOING, IXSYSTEMS PROVIDES NO WARRANTY OR UNDERTAKING, AND MAKES NO REPRESENTATION OF ANY KIND THAT THE PRODUCT WILL MEET THE LICENSEE'S REQUIREMENTS, ACHIEVE ANY INTENDED RESULTS, BE COMPATIBLE, OR WORK WITH ANY OTHER SOFTWARE, APPLICATIONS, SYSTEMS, OR SERVICES, OPERATE WITHOUT INTERRUPTION, MEET ANY PERFORMANCE OR RELIABILITY STANDARDS OR BE ERROR FREE, OR THAT ANY ERRORS OR DEFECTS CAN OR WILL BE CORRECTED.

TO THE FULLEST EXTENT PERMITTED UNDER APPLICABLE LAW: (A) IN NO EVENT WILL IXSYSTEMS OR ITS AFFILIATES, OR ANY OF ITS OR THEIR RESPECTIVE LICENSORS OR SERVICE PROVIDERS, BE LIABLE TO LICENSEE, LICENSEE'S AFFILIATES, OR ANY THIRD PARTY FOR ANY USE, INTERRUPTION, DELAY, OR INABILITY TO USE THE PRODUCT; LOST REVENUES OR PROFITS; DELAYS, INTERRUPTION, OR LOSS OF SERVICES, BUSINESS, OR GOODWILL; LOSS OR CORRUPTION OF DATA; LOSS RESULTING FROM SYSTEM OR SYSTEM SERVICE FAILURE, MALFUNCTION, OR SHUTDOWN;

FAILURE TO ACCURATELY TRANSFER, READ, OR TRANSMIT INFORMATION; FAILURE TO UPDATE OR PROVIDE CORRECT INFORMATION; SYSTEM INCOMPATIBILITY OR PROVISION OF INCORRECT COMPATIBILITY INFORMATION; OR BREACHES IN SYSTEM SECURITY; OR FOR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY, SPECIAL, OR PUNITIVE DAMAGES, WHETHER ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, REGARDLESS OF WHETHER SUCH DAMAGES WERE FORESEEABLE AND WHETHER OR NOT IXSYSTEMS WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES; (B) IN NO EVENT WILL IXSYSTEMS' AND ITS AFFILIATES', INCLUDING ANY OF ITS OR THEIR RESPECTIVE LICENSORS' AND SERVICE PROVIDERS', COLLECTIVE AGGREGATE LIABILITY UNDER OR IN CONNECTION WITH THIS AGREEMENT OR ITS SUBJECT MATTER, UNDER ANY LEGAL OR EQUITABLE THEORY, INCLUDING BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY, AND OTHERWISE, EXCEED THE TOTAL AMOUNT PAID TO IXSYSTEMS PURSUANT TO THIS AGREEMENT FOR THE PRODUCT THAT IS THE SUBJECT OF THE CLAIM; (C) THE LIMITATIONS SET FORTH IN THIS SECTION SHALL APPLY EVEN IF THE LICENSEE'S REMEDIES UNDER THIS AGREEMENT FAIL OF THEIR ESSENTIAL PURPOSE.

You hereby acknowledge that you have read and understand this Agreement and voluntarily accept the duties and obligations set forth herein by clicking accept on this Agreement.