



---

TitaniumZ and ZenOS

# Setup and Administration Guide

ZenOS 6.0.0  
Guide Version 3.0.1  
March, 16 2018



---

# TitaniumZ and ZenOS

## Setup and Administration Guide

The goal of this manual is to provide a starting point for you to install and use your TitaniumZ. The first chapter is for initial setup. Chapter 2 covers WebGui access. Chapter 3 is detailed configuration of the disks and volumes. Chapter 4 covers adding additional users and groups who access the TitaniumZ, including Active and Open Directory support. Chapter 5 is managing network connections. Chapter 6 shows diagnostics, log and troubleshooting info.

Small Tree  
Oakdale, MN  
[support@small-tree.com](mailto:support@small-tree.com)  
twitter @smalltreecomm

# Table of Contents

<b>Chapter 1: Introduction and Basic Setup .....</b>	<b>4</b>
Introduction.....	4
Assemble the TitaniumZ-8 and TitaniumZ-16 .....	7
Connecting a Client Workstation .....	12
Connecting Additional Network Client Workstations.....	15
<b>Chapter 2: Accessing and Using the WebGUI.....</b>	<b>17</b>
Default Login Name and Password .....	18
WebGUI Layout.....	18
Required Setup.....	20
Additional Setup.....	21
Setting up email .....	22
<b>Chapter 3: Storage Management .....</b>	<b>24</b>
ZFS Management.....	24
ZFS Terminology.....	24
Creating Additional Shared Datasets .....	27
Deleting a Dataset .....	30
Creating and Destroying Zpools .....	36
Creating an iSCSI Target .....	41
Restoring Data From a Snapshot .....	50
View Current Configuration .....	52
Your Data and Rsync .....	53
<b>Chapter 4: Users and Groups.....</b>	<b>53</b>
Managing Groups .....	54
Managing Users .....	55
Configuring Directory Support .....	58
Active Directory Integration .....	58
Open Directory Integration .....	61
LDAP Integration.....	64
Access Control Lists .....	64
<b>Chapter 5: Advanced Networking.....</b>	<b>65</b>
Changing the Configuration of the LAN Interface.....	69
<b>Chapter 6: System Status and Troubleshooting .....</b>	<b>73</b>
Troubleshooting .....	76

# Chapter 1: Introduction and Basic Setup

## Introduction

### Chapter

# 1

The TitaniumZ is designed to be simple to administer and easy to use for creative professionals. This first chapter provides the basics to get your TitaniumZ up and running, get client workstations connected and allow access to the TitaniumZ shared storage. Advanced features are found in chapters 2-6.

This guide covers the following products:

TitaniumZ-5  
TitaniumZ-8  
TitaniumZ-16  
CAS

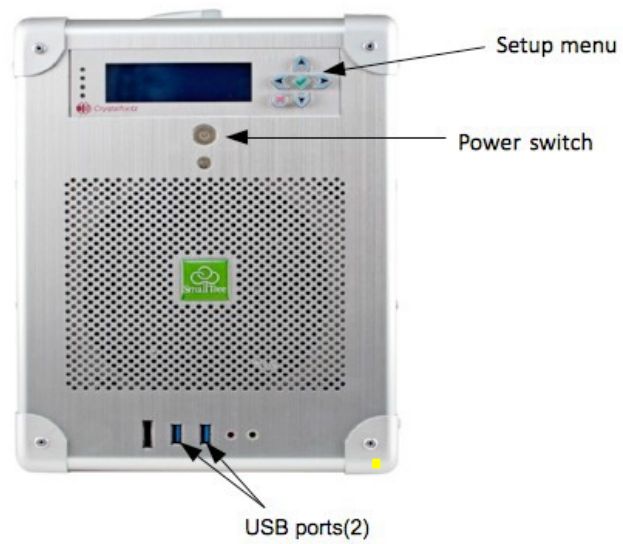
### **TitaniumZ-5 what's in the box:**

- TitaniumZ server with disk drives installed
- ZenOS preloaded and configured
- Power cable
- Quickstart Guide

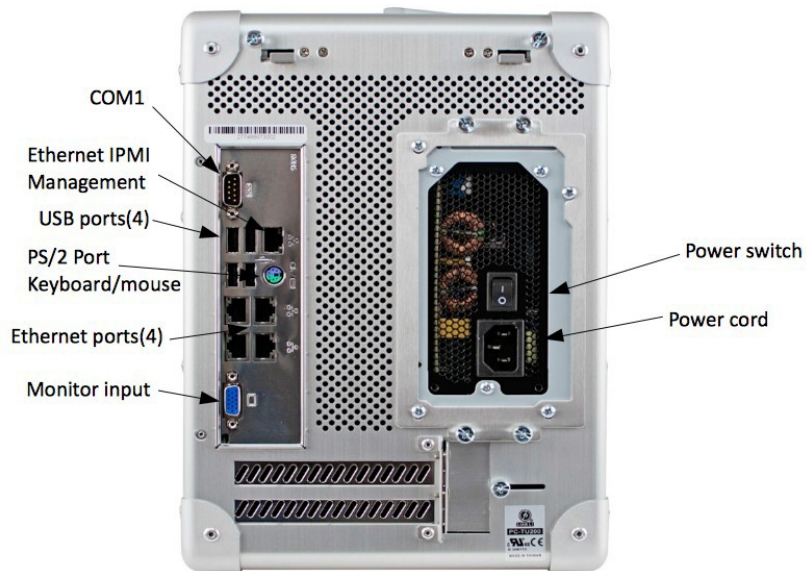
### **Optional Items (if purchased):**

- Ancillary cables
- Ethernet network adapters installed in the TitaniumZ

TitaniumZ-5 Front



TitaniumZ-5 Back view Figure 1



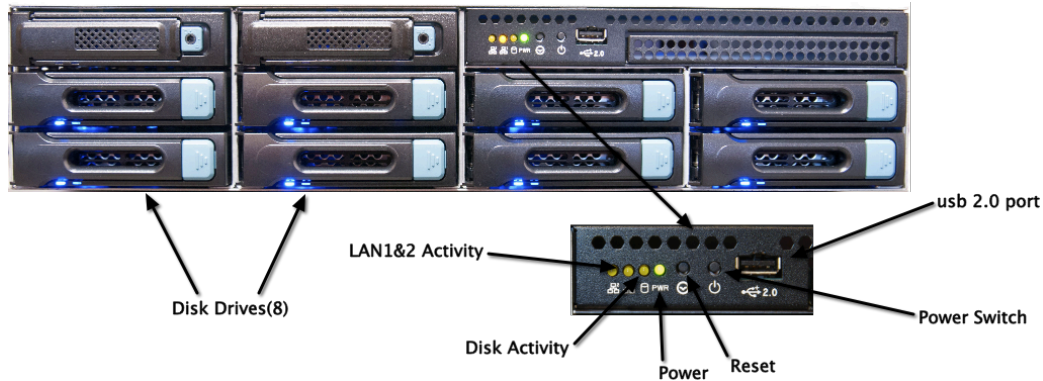
**TitaniumZ-8 and TitaniumZ-16 what's in the boxes:**

- TitaniumZ server
- ZenOS preloaded and configured
- Disk drives (not installed)
- Rack kit
- Two power cables
- Quickstart Guide

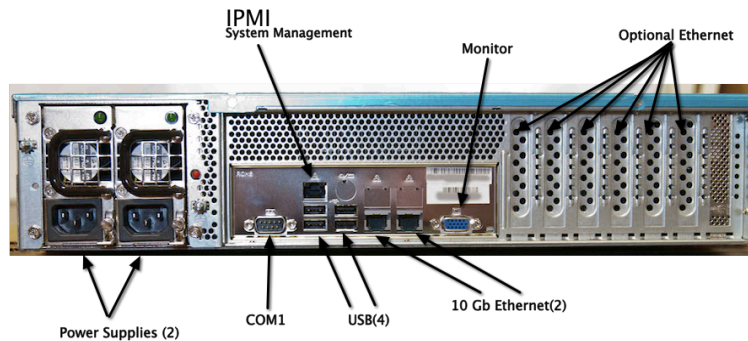
**Optional Items (if purchased):**

- Ancillary cables
- Ethernet network adapters installed in the TitaniumZ
- Mini-SAS cable for optional Expansion Chassis

TitaniumZ-8 Front view



TitaniumZ-8 Back view Figure 2



Note that the TZ-16 is the same as the TZ-8, except that it is 3U versus the 2U TZ-8.

Please don't hesitate to contact Small Tree support for assistance at any time.

Email: [support@small-tree.com](mailto:support@small-tree.com)

North American Callers: (866)-782-4622 Option 2

Callers Outside North America: +1 651-209-6509 Option 2

## Assemble the TitaniumZ-8 and TitaniumZ-16

*NOTE: The TZ-5 is shipped fully assembled. If you are setting up a TZ- 5, skip this step.*

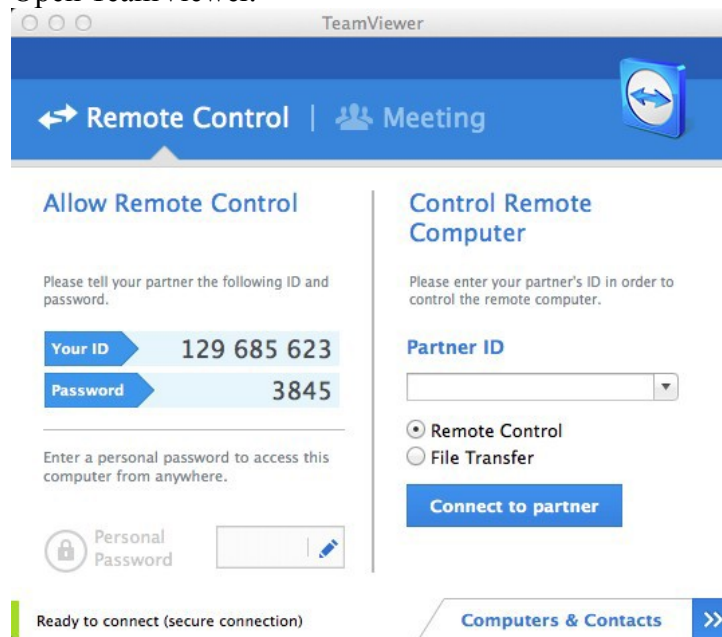
1. Install the disk drives in the TitaniumZ chassis. Drives are installed left to right. Drive 5 should be below drive 1 in both TZ 8 and TZ 16. Drive 2 should be directly right of Drive 1 in a TZ 8.

Drive 1	Drive 2	Drive 3	Drive 4
Drive 5	Drive 6	Drive 7	Drive 8

2. Connect the power cord(s) to the TitaniumZ and the UPS.

***NOTE: Do NOT power on the system at this point***

3. Any Small Tree network adapters purchased for the TitaniumZ are installed and configured by Small Tree prior to shipment. No action on your part is required to configure them.
4. Contact Small Tree support to schedule a time to complete the setup.
5. Before the scheduled time:
  - a. Download TeamViewer at this link:  
<https://www.small-tree.com/support/downloads/>
  - b. Open TeamViewer.



- c. Note Your ID and password.
- d. Send email to Small Tree support: [support@small-tree.com](mailto:support@small-tree.com). Please include in your email:

- a.i. Your ID from TeamViewer.
- a.ii. A phone number and/or Skype Name at which you can be reached.

If you are comfortable with completing the setup yourself, follow the instructions below. If you complete the setup on your own, it is a good idea to have a TeamViewer session with Small Tree to verify the setup and answer questions you might have that are not covered in this manual or the Small Tree Knowledge base found here:

[http://www.small-tree.com/mac\\_networking\\_shared\\_storage\\_faq\\_a/204.htm](http://www.small-tree.com/mac_networking_shared_storage_faq_a/204.htm)

If you don't plan to make any changes to the TitaniumZ settings, skip the next section and start at "Connecting workstations to the TitaniumZ network drive share1:"

## Login to the WebGUI

1. On your client workstation you first need to configure a wired Ethernet port:

### Mac:

Open System Preferences. Click Network. Select a wired Ethernet interface on the left (Ethernet 1 for example). In the Configure IPv4 pulldown, select Manually. Enter settings listed below.

### Windows:

Open Control Panel. Click "Network and Internet." Under "Network and Sharing Center," click "View network status and tasks". Under "View your active networks"; you should see a "Local Area Connection <x>". Click on that. In the "Local Area Connection Status" window click the "Properties" button. In the "This connection uses the following items:" list, select "Internet Protocol Version 4 (TCP/IPv4)" and click the "Properties" button below the list. Click the "Use the following IP address" radio button. Enter settings listed below.

### Settings for Mac and Windows:

IP Address: 192.168.150.10  
Subnet Mask: 255.255.255.0  
Router/Default Gateway: leave blank  
DNS Server: leave blank

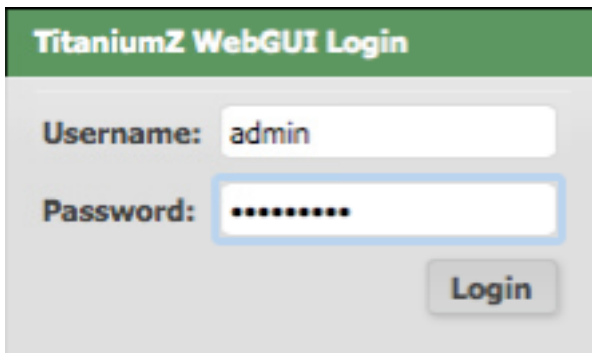
2. Connect the Ethernet port on the Mac or Windows system configured above to:
  - a. TitaniumZ-16 and TitaniumZ-8: The left port of the two RJ45 ports on the back of the TitaniumZ. This is the RJ45 port between the USB and VGA

ports. See Figure 2.

b. Titanium Z-5: The upper left of the four ports grouped together on the back of the TZ-5. See Figure 1

3. Power on the TitaniumZ. Be patient as it takes on average up to 3 minutes for the TitaniumZ to power up and be accessible, depending on your configurations.
4. Open a web browser and enter the address <http://192.168.150.220>. This should bring you to this login page:


**NOTE: If you want to change the IP address of this port, or any other ports, chapter 5**



Default User name: admin

Default Password: smalltree

After clicking on “Login” you will see this screen:



SYSTEM > INFORMATION

adamstz8.local

System

Network

Disks

Services

Access

Status

Diagnostics

Advanced

Help

System Information

Hostname	adamstz8.local
Version	2.1.3 - Redwood (revision 4)
Build date	Thu Mar 26 16:19:34 UTC 2015
Platform OS	FreeBSD 10.0-RELEASE (rdate 1000510)
Platform	x64-full on Intel(R) Xeon(R) CPU E5-2609 0 @ 2.40GHz
System	Supermicro X9DRH-7TF/7F/TF/IF
System bios	American Megatrends Inc. version: 3.0a 12/27/2013
System time	Wed Apr 15 10:51:27 CDT 2015
System uptime	8 day(s) 19 hour(s) 41 minute(s) 41 second(s)
Last config change	Mon Apr 13 13:57:06 CDT 2015
CPU frequency	2400MHz
CPU usage	0%
Memory usage	77% of 32710MiB
Load averages	0.00, 0.04, 0.03 <a href="#">[Show process information]</a>
Disk space usage	<div>ssdShare</div> <div>Total: 72G   Used: 32M   Free: 66G</div> <div>tank0</div> <div>Total: 12.6T   Used: 2.85T   Free: 9.58T   State: ONLINE</div>
Disk group status	<div>Controller: 0   EID:Slot: 252:0   DID: 11   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:1   DID: 10   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:2   DID: 9   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:3   DID: 8   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:4   DID: 14   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:5   DID: 15   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:6   DID: 12   Size: 1.818 TB   State: Onlin</div> <div>Controller: 0   EID:Slot: 252:7   DID: 13   Size: 1.818 TB   State: Onlin</div>
UPS status	<div>Status</div> <div>UPS On Line, Charging Battery <a href="#">[Show ups information]</a></div> <div>Load</div> <div>Battery Charge</div> <div>Battery Remain Time</div> <div>4 minutes 56 seconds</div>

Note: Small Tree strongly recommends that the default username and password be changed prior to putting the TitaniumZ into production.

## System > General

Enter a new user name.

## System > General > Password

Enter a new password.

Be sure to record the new user name and password. If you forget/lose the user name and password you will not have access to the WebGUI.

You do not need to make any other modifications at this point. If you want to make modifications, see the following chapters.

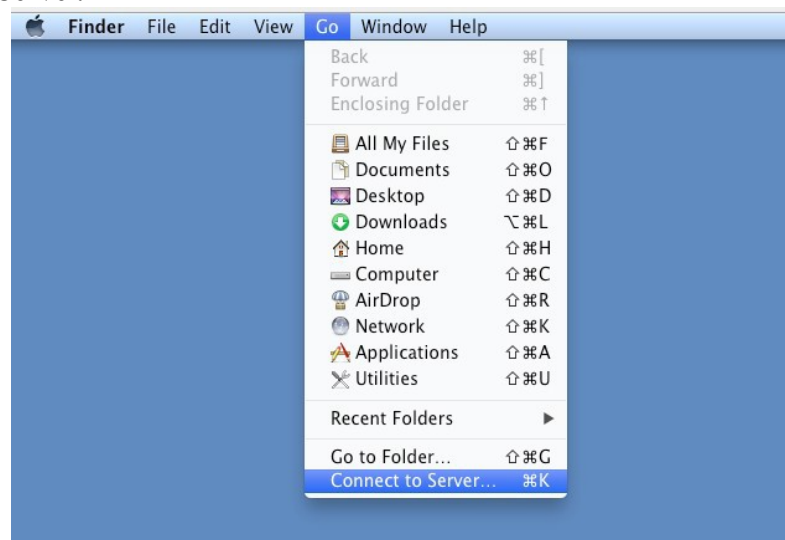
Now we will move on to providing users access to the TitaniumZ shared storage.

## Connecting a Client Workstation

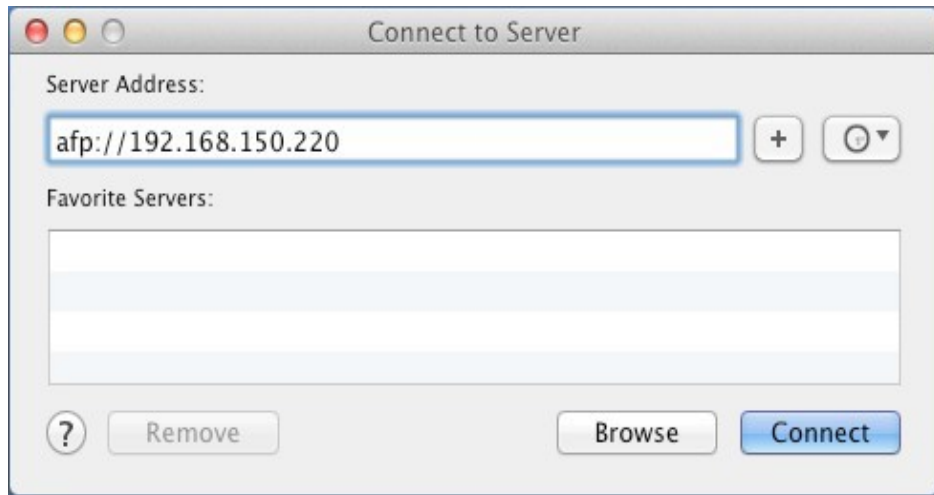
Note: The following instructions apply if you are connecting client workstations directly to Ethernet ports on the TitaniumZ. If you are using a switch to connect client workstations, please contact Small Tree support for assistance.

### Mac:

1. In the Finder (click on your Desktop), under the “Go” menu, select "Connect to Server."



Which will take you to this screen:



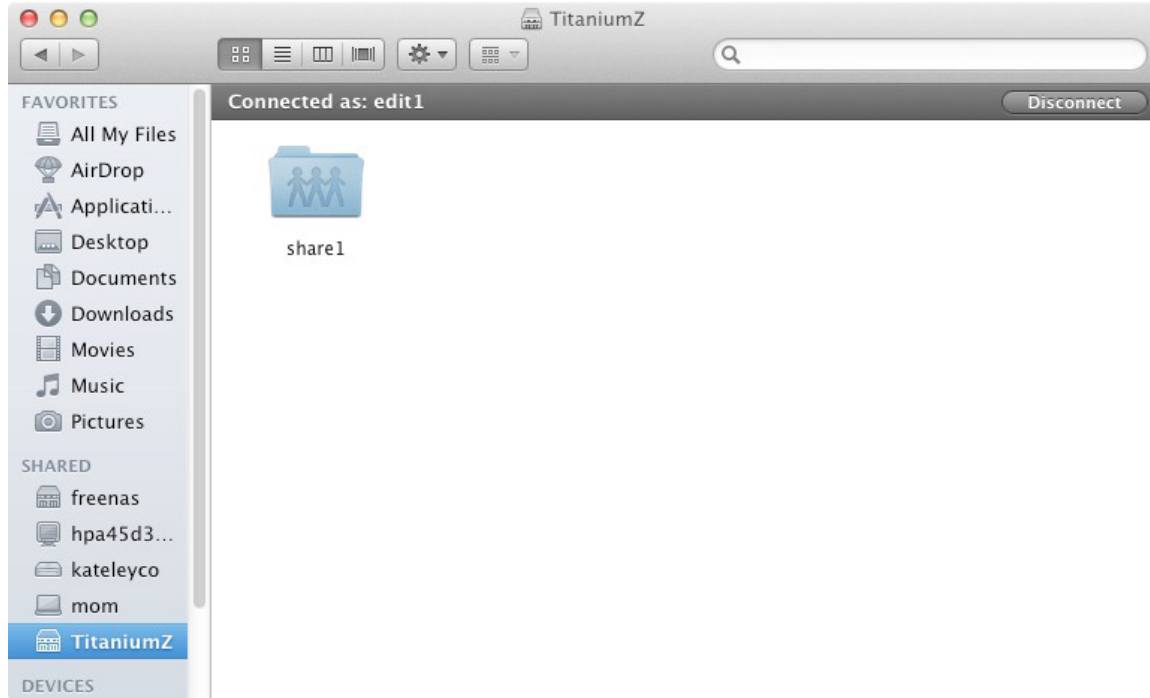
2. In the “Server Address” box enter:  
For AFP shares enter afp://192.168.150.220  
For CIFS/SMB shares enter smb://192.168.150.220  
For NFS shares enter nfs://192.168.150.220/mnt/tank0/share1/

Click “connect”. The login screen will display:



3. Login with:  
Default User name: edit1  
Default Password: edit1

Click “Connect”. You will see this screen:



4. The default network drive is called share1. Double click on share1; it will show up on your desktop as another drive. Use share1 as you would any drive.

**To mount share1 automatically each time you start up (optional):**

1. Go to System Preferences.
2. Click on “Accounts” or “Users and Groups”, depending on your version of OS X.
3. Click on “Login Items”.
4. Click on the + button to add another login item.
5. Locate share1 and click “Add”.
6. Quit System Preferences.

**Windows:**

1. In the “Start Menu”, click on “Computer”. Right click on “Computer” in the

- left column of the Computer window and then select “Map Network Drive”.
2. In the “Map Network Drive” window, click the “Browse” button.
  3. In the “Browse for Folder” window, select “TITANIUM-SMB”.
  4. Login with:  
    Default User name: edit1  
    Default Password: edit1
  5. In the Browse for Folder window, below TITANIUM-SMB, select share1 and then click the “OK” button at the bottom.
  6. Back in the “Map Network Drive” window, click the “Reconnect at login” check box then click on the “Finish” button at the bottom.

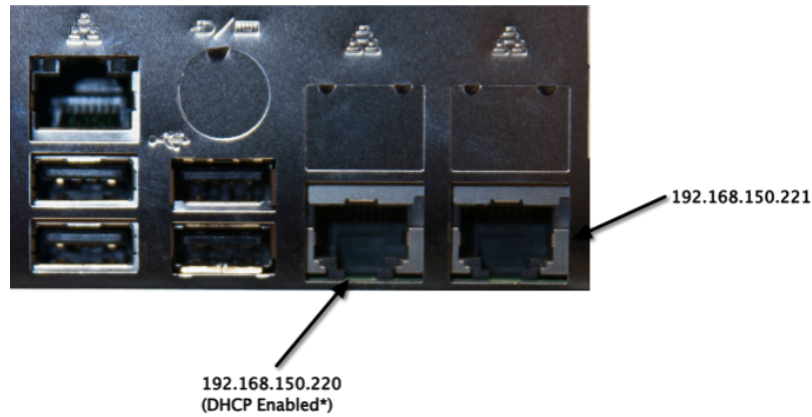
## **Connecting Additional Network Client Workstations**

You can add additional client workstations to the additional Ethernet ports by using the following IP addresses.

### **Titanium Z-8 and Titanium Z-16 Ethernet Port Assignments**

Additional Ethernet ports for connections to client workstations on the TitaniumZ-8 and Titanium Z-16 are assigned as follows:

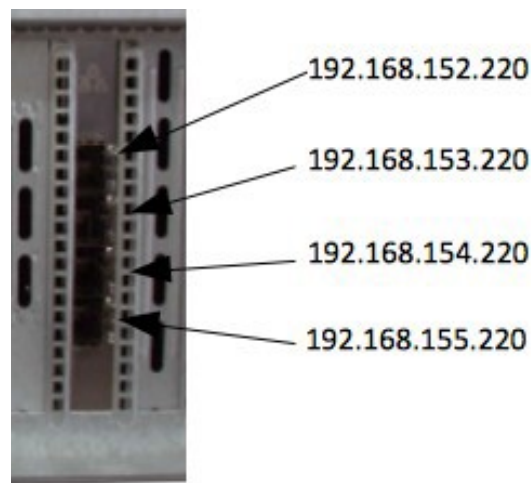
The default RJ45 port, between the USB and VGA ports, is 192.168.150.220.



**NOTE: DHCP Enabled port will automatically find an IP address.**

Note: The third Ethernet port built-in to the TitaniumZ-8 and TitaniumZ-16 is an IPMI port. This is a special management port, which cannot be used to connect client workstations.

Ports on vertical Ethernet adapters start at 192.168.152.220 for the top leftmost port. The network number is incremented for each port going down and right. For a single 4-port adapter the default IP addresses are:



For other configurations contact Small Tree technical support.

## Titanium Z-5 Ethernet Port Assignments

Additional Ethernet ports for connections to workstations on the TitaniumZ-5 are assigned as follows.



Note: The fifth Ethernet port built-in to the TitaniumZ-5 is an IPMI port. This is a special management port, which cannot be used to connect client workstations.

Ports on horizontal Ethernet adapters start at 192.168.154.220 for the left most port. The third number (154 for the left most port) increases by 1 progressing left to right. For a single 4-port adapter the default IP addresses are:

192.168.154.220	Left most port
192.168.155.220	Middle left port
192.168.156.220	Middle right port
192.168.157.220	Right most port

## Connecting the Client Workstations

Follow instructions for Mac and Windows earlier in this chapter, but change IP address to the network you are connecting to.

## Chapter 2: Accessing and Using the WebGUI

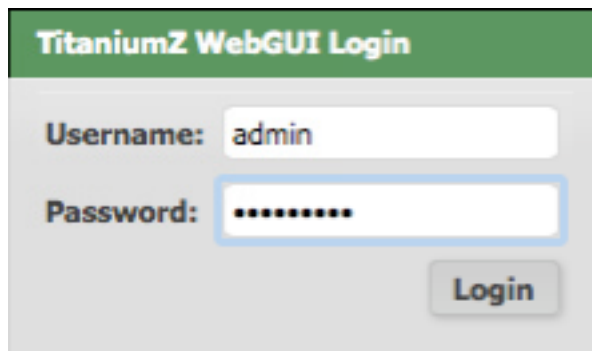
The TitaniumZ administrative console (WebGUI) is accessible over the network using a web browser. This chapter discusses some of the setup options configurations for optimal operations.

## Chapter

# 2

### Default Login Name and Password

To access and administer the TitaniumZ , open a web browser on a system connected to the TitaniumZ, either directly (Ethernet port must be properly configured per Chapter 1) or through a properly configured network. In the address bar of the web browser, type the IP address of the TitaniumZ in as the URL. The default IP address is <http://192.168.150.220>. You should then see the following graphic:



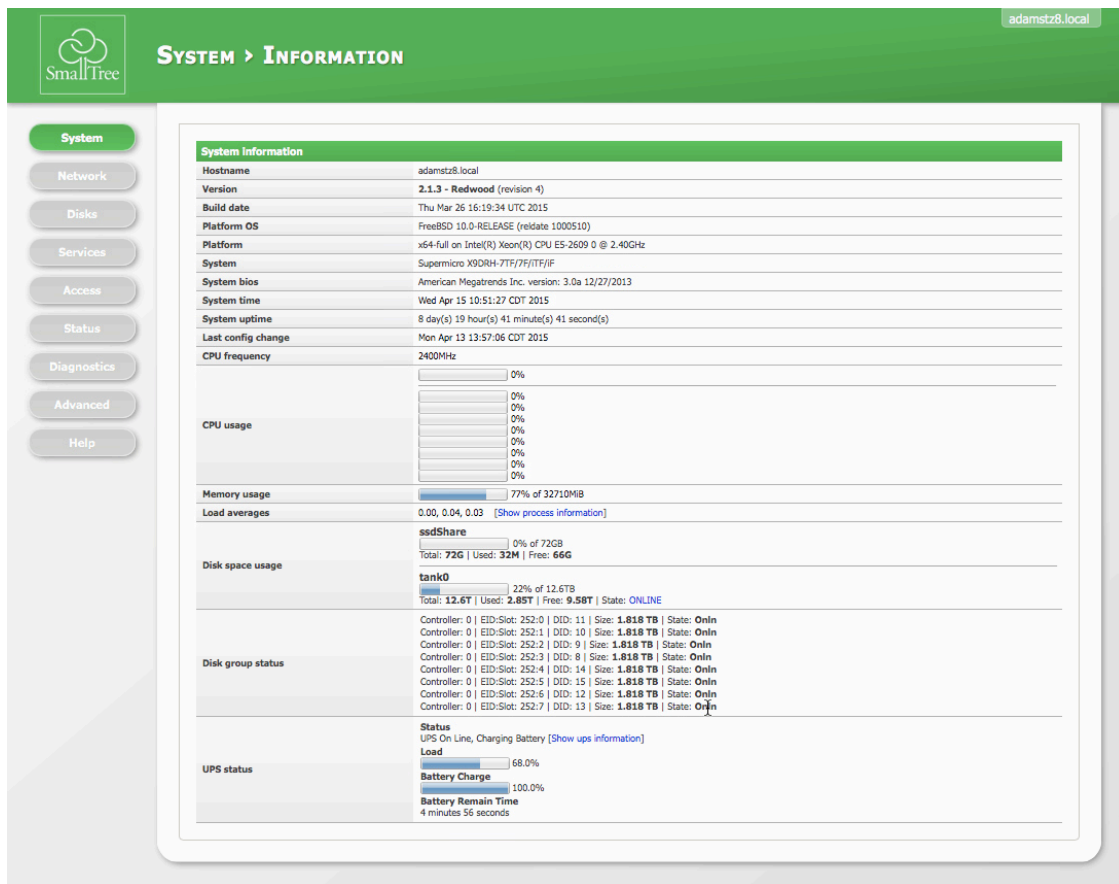
To login, enter:

1. Default username: admin
2. Default password: smalltree
3. Click “Login”

*NOTE: Small Tree strongly recommends that the default username and password be changed before putting the Titanium Z into production. Instructions are in Chapter 5.*

### WebGUI Layout

After login, you should see the System Information page.



The menu items at the top of this page allow admin access in pulldown menus for specific management functions.

All admin functions can be done from the following menus:

**SYSTEM:** Contains system configuration and status information.

**NETWORK:** Allows admin to view, modify and configure network interfaces. (See Chapter 5.)

**DISKS:** Allows admin to modify disks, pools, volumes and datasets. (See Chapter 3.)

**SERVICES:** Allows admin to enable and configure network services. For example, turn on ftp access, setup iSCSI target.




**ACCESS:** Allows admin to add users and groups either manually or setup directory services.

**DIAGNOSTICS:** Allows admin to view error logs and run network diagnostics scripts.

**ADVANCED:** Allows admin to edit files or run commands on the TitaniumZ.

**HELP:** Find additional help on topics.

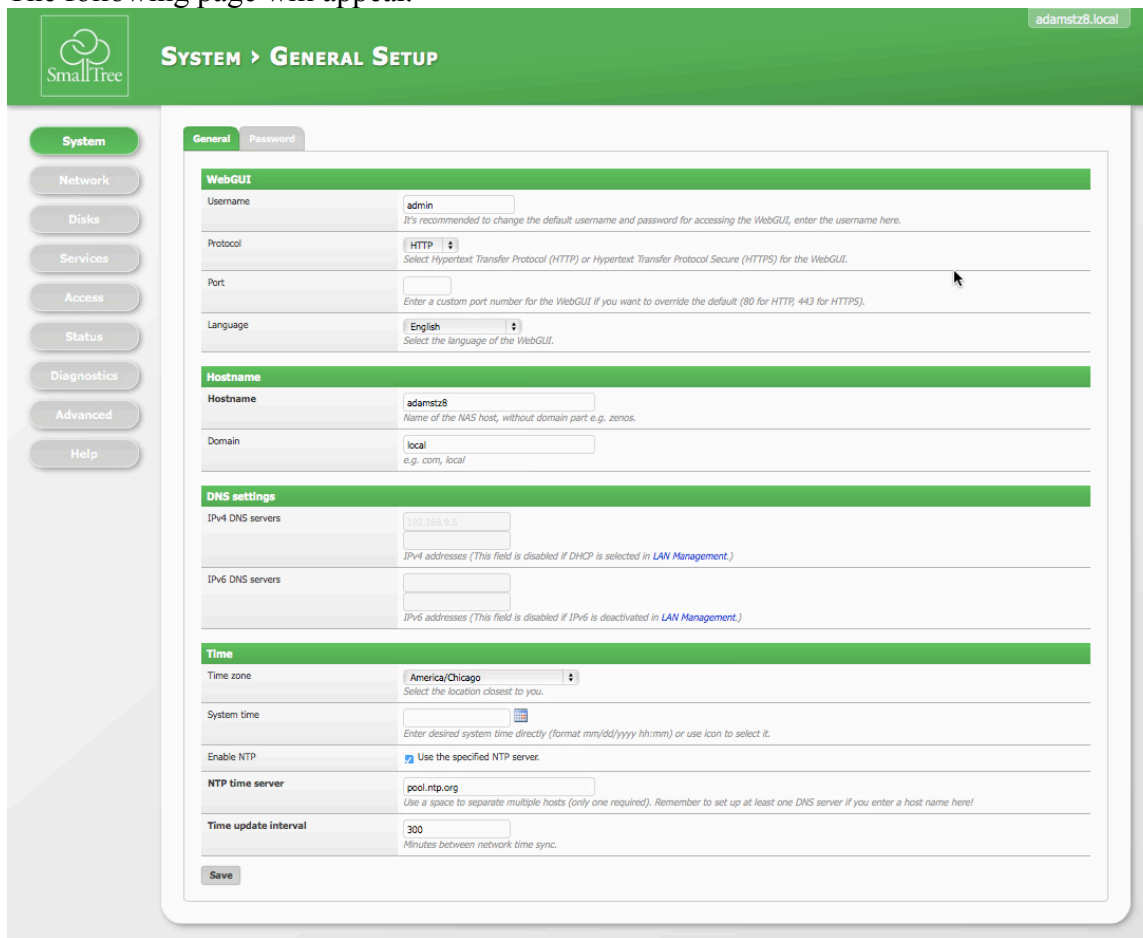
Some of WebGUI pages use these icons:

- ☐  - Allows adding of item.
- ☐  - Controls deleting or removing an item.
- ☐  - Permits the admin to edit the item.

## Required Setup

### System->General

The following page will appear.



The screenshot shows the 'SYSTEM > GENERAL SETUP' page in the Small Tree WebGUI. The page is divided into a sidebar and a main content area. The sidebar contains buttons for 'System', 'Network', 'Disks', 'Services', 'Access', 'Status', 'Diagnostics', 'Advanced', and 'Help'. The main content area is titled 'General' and contains several sections for configuration:

- WebGUI**: Includes fields for 'Username' (default: admin), 'Protocol' (default: HTTP), 'Port' (default: 80), and 'Language' (default: English).
- Hostname**: Includes fields for 'Hostname' (default: adamstz8) and 'Domain' (default: local).
- DNS settings**: Includes fields for 'IPv4 DNS servers' (default: 192.168.0.5) and 'IPv6 DNS servers'.
- Time**: Includes fields for 'Time zone' (default: America/Chicago), 'System time' (with a clock icon), 'Enable NTP' (checked), 'NTP time server' (default: pool.ntp.org), and 'Time update interval' (default: 300 minutes).

A 'Save' button is located at the bottom of the main content area.

There are a number of items that may need to be set on this page. Most of these items have been pre-configured by Small Tree.

### WebGUI

**Username:** The WebGUI user name. Default is admin.

**Protocol:** HTTP or HTTPS. TitaniumZ will self-create a certificate for HTTPS.

**Port:** WebGUI port (default is empty).

**Language:** Default is English.

**Hostname:** Default is “titaniumz.” Host name may be changed to a preferred name. All lower case, no special characters.

**Domain:** DNS domain name.

**DNS Settings:**

IPv4 DNS servers: (Required if using static IP address).

IPv6 DNS Servers: Default is empty.

You will need to set time zone and system time in order for the TitaniumZ to work properly. It is recommended to set a timeserver, but you can also set time manually.

Time	
Time zone	<div>America/Chicago</div> <div>Select the location closest to you.</div>
System time	<div></div> <div>Enter desired system time directly (format mm/dd/yyyy hh:mm) or use icon to select it.</div>
Enable NTP	<input checked="" type="checkbox"/> Use the specified NTP server.
NTP time server	<div>pool.ntp.org</div> <div>Use a space to separate multiple hosts (only one required). Remember to set up at least one DNS server if you enter a host name here!</div>
Time update interval	<div>300</div> <div>Minutes between network time sync.</div>

Save

**Time zone:** Set by selecting on the dropdown menu.

**System Time:** Can be set manually by entering time. OR

**Enable NTP:** NTP will automatically regularly update the system time by requesting the current information from a timeserver.

**NTP timeserver:** If you have a timeserver on your network you can add the name or IP address of that server, or you can use the ntp.org servers on the Internet, pool.ntp.org.

**Time update interval:** Setup the minutes in between requests to the timeserver. The default is 300 minutes.

## Additional Setup

The following includes additional systems administration settings that should be configured.

## Changing the WebGUI Admin Username and Password

Small Tree highly recommends changing the default WebGUI user name and password. If you have not already done so, do it now.

## System -> General -> General

WebGUI	
Username	<input type="text" value="admin"/> <small>It's recommended to change the default username and password for accessing the WebGUI, enter the username here.</small>
Protocol	<input type="text" value="HTTP"/> <small>Select Hypertext Transfer Protocol (HTTP) or Hypertext Transfer Protocol Secure (HTTPS) for the WebGUI.</small>
Port	<input type="text"/> <small>Enter a custom port number for the WebGUI if you want to override the default (80 for HTTP, 443 for HTTPS).</small>
Language	<input type="text" value="English"/> <small>Select the language of the WebGUI.</small>

Enter the new WebGUI user name.

Click “Save” at the bottom of the page.

## System -> General -> Password

The screenshot shows the 'Password' tab in the 'General' section of the SmallTree WebGUI. The 'Current password' field is empty. The 'New password' field is empty, and the 'Confirmation' field is also empty. A 'Save' button is located below the password fields. A note at the bottom states: 'The new password is also the default root password of the system!'.

Enter the new password name.

Click “Save” at the bottom of the page.

**NOTE: Remember that changing both of these but only clicking save will only remember the latest change you made.**

## Setting up email

To setup and modify the Email settings, go to:

**System -> Advanced -> Email**

The screenshot shows the SmallTree TitaniumZ web interface. The top navigation bar is green with the SmallTree logo on the left and the text 'SYSTEM > ADVANCED > EMAIL' in the center. On the far right of the top bar, the text 'adamstz8.local' is visible. A left sidebar contains a vertical list of buttons: 'System' (highlighted in green), 'Network', 'Disks', 'Services', 'Access', 'Status', 'Diagnostics', 'Advanced', and 'Help'. The main content area has a tabbed interface with tabs for 'Advanced', 'Email' (selected), 'Proxy', 'Swap', 'Command scripts', 'Cron', 'rc.conf', and 'sysctl.conf'. The 'Email' tab is active, showing a form with the following fields: 'From email' (with a text input and a hint 'The 'From:' email address on mails sent from this system.'), 'Outgoing mail server' (with a text input and a hint 'Outgoing SMTP mail server address, e.g. smtp.mycorp.com.'), 'Port' (with a text input and a hint 'The default SMTP mail server port, e.g. 25 or 587.'), 'Security' (a dropdown menu currently set to 'None'), 'Authentication' (a checkbox labeled 'Enable SMTP authentication.'), and 'Send test email' (with a text input and a hint 'Enter an email address to receive the test email; multiple addresses should be separated by a semicolon'). Below the 'Send test email' field is a 'Send test email' button. At the bottom left of the form is a 'Save' button.

The TitaniumZ is designed to automatically email problems and diagnostics. You will need to setup the basic email info for this to work properly. You will need to set:

**From email:** Enter a company email address. [titaniumz@mycompany.com](mailto:titaniumz@mycompany.com) is popular.

**Outgoing mail server:** SMTP outgoing server name.

**Port:** Use default, unless your company uses a different port.

**Security:** Setup security protocol that your mail server requires.

## Systems -> Advanced -> Proxy

This would be where you would setup proxy information if required by company policy. Default is empty.

## Other settings under System-> Advanced

Ignore the other tabs in System | Advanced; any changes to these settings should be made by Small Tree support. If you have questions about these settings, contact Small Tree support.

## Chapter 3: Storage Management

### ZFS Management

TitaniumZ storage servers are powered by ZenOS, Small Tree's powerful operating system. ZFS is the component of ZenOS used for most storage management tasks. ZFS is open source filesystem that provides higher performance, scalability and a number of enterprise features available only with ZFS.

ZFS, being a combined file system and logical volume manager, uses some unique terminology. This terminology is explained in the following paragraphs.

### ZFS Terminology

#### Virtual Device(vdev)

vdevs are individual disk drives. Each of the disk drives in the Titanium-5 series are vdevs. They are also what ZFS "sees" as a disk drive.

## Chapter

# 3

TitaniumZ-8 and TitaniumZ-16: A vdev is a hardware RAID controller card with eight disk drives connected to it. The hardware RAID controller card and eight connected hard drives are configured by default as RAID5. With the TitaniumZ-8 (no expansion chassis) there is only one vdev. With the TitaniumZ-16 (no expansion chassis) there are two vdevs.

The default configuration from Small Tree has the vdev(s) created and configured.

**WARNING:** Do not add or modify a vdev without Small Tree support involvement. Incorrectly configuring vdevs will likely lead to complete data loss.

**NOTE:** For optimal performance reasons, the capacity of all vdevs in a pool should be identical. This does not mean that vdevs are limited to be the same size however.

Adding a RAID controller card and an Expansion Chassis with eight disk drives creates a new vdev for the TitaniumZ-8 and TitaniumZ-16.

Vdevs are used to create pools.

## Pools

Pools are the basic storage units of ZenOS with ZFS. It is similar to what other storage management software calls a Drive Group or Volume. A pool is composed of one or more vdevs. When there are multiple vdevs, Small Tree creates a RAID50 by striping multiple vdevs.

The default configuration from Small Tree has a pool named “tank0”.

**WARNING: Do not add or modify a Pool without Small Tree support involvement. Incorrectly configuring Pools can lead to complete data loss.**

Datasets and volumes, described below, share the available storage from the pool.

## Datasets

Think of a dataset like a partition on a hard drive that may be shared using file-sharing protocols such as AFP, CIFS/SAMBA and NFS. It is a logical way to organize folders and files, except; you never need to resize a dataset. The capacity consumed from the pool is only as much space as you use populating it with files. Datasets have no size until files are added. The number of datasets on the TitaniumZ is virtually limitless.

The TitaniumZ has been pre-configured with a dataset named “share1.”

You may change the name of a dataset to something more meaningful to you. If you do change the dataset name, you should also change the name and path of any services (AFP, FTP, etc.) by going to the Services tab and selecting the appropriate service.

The storage space in the pool is shared by all datasets. You may limit the size of a dataset by setting a quota. You may also guarantee a minimum amount of space available to the dataset by reserving space.

You may want to create additional datasets to create a finer separation of files at the dataset level. You may also create datasets within datasets. For example, you could create share1/art department, share1/news, share1/promotions.

The logical separation of datasets allows the administrator to grant specific users and groups appropriate permissions to specific datasets.

To allow network user access to a dataset it must be shared using AFP, CIFS/SMB or NFS. A dataset may be simultaneously shared using any and all of the three file sharing protocols.

The TitaniumZ has been pre-configured with “share1” accessible utilizing AFP, CIFS/SMB

and NFS.

OS X systems see shared datasets as network volumes.

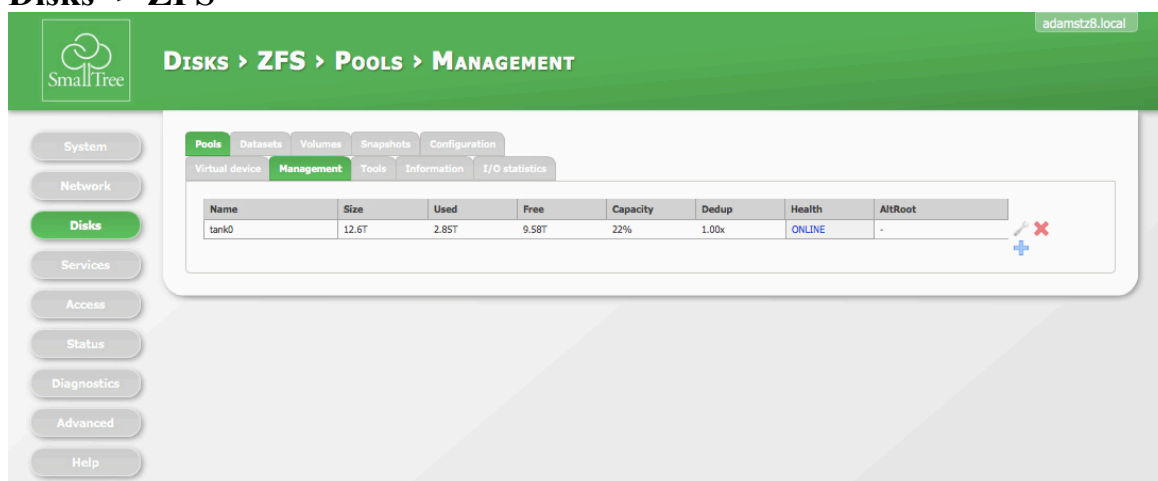
Windows systems see shared datasets as network drives.

## Creating Additional Shared Datasets

Your TitaniumZ will come from Small Tree with a default dataset named “share1”. By default, it is shared via AFP, CIFS/SMB and NFS and is ready to use. Clients connect using the IP address/datasetname on the TitaniumZ as outlined in Chapter 1.

## Pool Management

### Disks -> ZFS



On this page modifications can be made to:

**Pools** - A grouping of disks. (Only Small Tree support makes modifications to Pools).

**Datasets** – Logical directory/folder that will be exported as a network share.

**Volumes** – Space reserved for block level storage. Typically iSCSI. **Snapshots** – Point in time copy of a dataset.

### Disks -> ZFS -> Pools

On this page you can see the available pools on the system. It will show detailed information about each pool including:

**Name:** Name of pools on the system.

**Size:** Total space in the pool.

**Used:** Storage space used in the pool.

**Free:** Currently available space.

**Capacity:** Shows as 0%. Disregard this entry.

**Dedup:** Ignore this. Deduplication is not recommended for media files.

**Health:** Should show ONLINE. If not, contact Small Tree support.

## Creating a New Dataset

### Disks -> ZFS -> Datasets > Dataset

This page will show all of the datasets that are available on the TitaniumZ. By default only share1 will be listed.

Select the “+” icon to the far right to create an additional dataset. This page will be displayed.

Small Tree's suggested settings are:

**Name:** You may select any name you'd like. No spaces or special characters. If you want to create a subset of share1 type as "share1/name".

**Pool:** Select pool from which the new dataset will consume capacity. Default is tank0.

**Compression:** Leave compression off (compression does not work well for video data).

**Dedup:** Leave dedup off for video data.

**Sync:** Disabled.

**Access time:** Set this to "off". This will improve performance for most applications. If

viewing the last time accessed is required set to “on.”

**NOTE: When set to “off” created and modified information is still recorded and displayed.**

**Canmount:** Must be checked.

**Readonly:** Leave off.

**Extended Attributes:** Should be checked.

**Snapshot visibility:** Do not check. If checked will create very annoying finder problems.

**Reservation:** By default, all space in a pool is shared with all datasets. You can reserve space in the pool for a dataset to ensure it has “at least” this much space.

**Quota:** By default, all space in the pool is shared with all datasets. You can set a limit of space that can be used by a dataset.

**Description:** Allows admin to assign descriptive information about this dataset.

At the bottom of the page you can set:

Access Restrictions				
Owner	root			
Group	wheel			
Mode		Read	Write	Execute
	Owner	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Group	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Others	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Add Cancel

**Owner:** User who will have owner permissions on the dataset. Default is root.

**Group:** Group owner. Default is wheel.

**Mode:** Default permissions for this dataset. These can be selected for Owner, Group and Others. Default is to check all.

**Read:** View the contents of files.

**Write:** Create new files, or save files into the dataset.

**Execute:** Run applications from this dataset.

When all the settings are configured: Click “Add”.

Click “Apply Changes”.

The new dataset will not be created unless both steps are completed.

## Deleting a Dataset

### Disks -> ZFS -> Datasets->Dataset

To delete a dataset, navigate to the dataset management page and look for the dataset you would like to delete. Once you've located this, simply click the red "X" next to the dataset. A prompt will ask if you really want to delete the dataset. If you are sure, then simply click "OK" and then click "Apply Changes."

**NOTE: When deleting a Dataset, the dataset is simply "forgotten" by the webUI, and is not deleted. This is similar to the pools and vDevs, but after a dataset is deleted it cannot be restored through the sync menu. In order to delete a dataset completely, it is suggested that you call Small-Tree Support for help.**

## Sharing a Dataset

After the new dataset has been created, the dataset must be shared to allow user access.

**NOTE: When you share a new dataset either with AFP, CIFS/SMB or NFS, the sharing service is restarted. Any users actively working with data will be disconnected and will have to reconnect. Only add new shares after you've warned your users that there will be an interruption of the sharing service.**


Information on connecting to shared datasets using AFP, CIFS/SAMBA and NFS is found in Chapter 1.

## Sharing a Dataset Using AFP

### Services -> AFP > Shares

This page shows all the existing AFP shares. The default is "share1".

Select the "+" icon at the lower right of the list. You will be presented with the AFP share setup page.



SERVICES > AFP > SHARE > ADD

adamatz8.local

System  
Network  
Disks  
**Services**  
Access  
Status  
Diagnostics  
Advanced  
Help

Settings Shares

Name

Comment

Path

Path to be shared.

Share Password

Set share password.

This controls the access to this share with an access password.

Share Character Set

UTF8

Specifies the share character set. For example UTF8, UTF8-MAC, ISO-8859-15, etc.

Allow

This option allows the users and groups that access a share to be specified. Users and groups are specified, delimited by commas. Groups are designated by a @ prefix.

Deny

The deny option specifies users and groups who are not allowed access to the share. It follows the same format as the allow option.

Read Only Access

Allows certain users and groups to have read-only access to a share. This follows the allow option format.

Read/Write Access

Allows certain users and groups to have read/write access to a share. This follows the allow option format.

Automatic disk discovery

Enable automatic disk discovery.

Automatic disk discovery mode

Default

Note! Selecting 'Time Machine' on multiple shares will may cause unpredictable behavior in MacOS.

Advanced Options

Case Folding

No case folding

This controls how the case of filenames are viewed and stored.

dbpath

Sets the database information to be stored in path. You have to specify a writable location, even if the volume is read only.

cnidscheme

Default

Set the CNID backend to be used for the volume.

cacheid

If set afpd uses the ID information stored in AppleDouble V2 header files to reduce database load. Don't set this option if the volume is modified by non AFP clients (NFS/SMB/local).

crif

Enables crif translation for TEXT files, automatically converting macintosh line breaks into Unix ones. Use of this option might be dangerous since some older programs store binary data files as type 'TEXT' when saving and switch the filetype in a second step. Afpd will potentially destroy such files when 'erroneously' changing bytes in order to do line break translation.

mswindows

This forces filenames to be restricted to the character set used by Windows. This is not recommended for shares used principally by Mac computers.

noadouble

This controls whether the .AppleDouble directory gets created unless absolutely needed. This option should not be used if files are access mostly by Mac computers.

nodev

Always use 0 for device number, helps when the device number is not constant across a reboot, cluster, ...

noffield

Don't advertise createfield, resolveid, deleteid calls.

nohex

Disables :hex translations for anything except dot files. This option makes the '/' character illegal.

prodos

Provides compatibility with Apple II clients.

nostat

Don't stat volume path when enumerating volumes list, useful for automounting or volumes created by a preexec script.

upriv

Use AFP3 unix privileges.

Add Cancel

Small Tree's suggested settings are:

**Name:** The name of the share (this must match the dataset name).

**Comment:** You must add a comment. You can use this space for a description of this share.

**Path:** Select the path to the data set. Select the <...> button and it will let you browse and select from the available datasets. The path should look like /mnt/tank0/datasetname.

**Share password:** Blank.

**Share Character Set:** UTF8 is the default but you may want to select other character sets if your environment has special character localization requirements. **Allow:** If left blank all users and groups are allowed access to the share. This can be set to the users and groups that you'd like to have access to the share. Groups are added with @groupname. Users are added by user name, comma separated. User name and password will still be required to connect.

**NOTE: Users and Groups must first be added in the Access->Users and Groups panel. Users and Groups are discussed in detail in Chapter 3.**

**Deny:** Can be left blank unless there's a user in the office you don't want to have access.

**NOTE: Once you define a set of users and groups in allow, everyone else is denied.**

**Read Only Access:** if left blank, all users will have read/write privilege. Any groups or users added will only be able to read data in the dataset, without being able to write to it. Groups are added with @groupname. Users are added by user name, comma separated.

**Read/Write Access:** This follows the "Read Only Access" rules, but allows users to read and write. Any users or groups that are specified in allow, but not specified in "Read Only Access" or "Read/Write Access" will be given "Read Only Access" by default.

**Automatic disk discovery:** If you enable this, you will have the option of making this dataset a time machine compatible volume. Only do this on one dataset on the server to avoid conflicts.

Leave all of the advanced options blank **except** upriv. This should be checked to use AFP3 privileges. Once done, select "Add" and then select "Apply Changes" so this new share will be made available and remain after a reboot. You must click "Save" AND "Apply Changes" for your changes to be persistent. If not, your entries will be temporary, and will disappear upon reboot.

**WARNING: When you select "Apply Changes", all sharing services (AFP, CIFS/SMB and NFS) will restart and any users with mounted shared datasets will need to reconnect. The connected user will get a message saying the service has stopped.**

## Sharing a Dataset Using SMB/CIFS

Datasets may be shared using multiple protocols. Settings for SMB/ CIFS shares are very similar to AFP settings in the prior section.

### Services -> SMB/CIFS -> Shares

This page shows all the existing SMB/CIFS shares. The default is “share1”. Select the “+” icon at the lower right of the list. You will be presented with the SMB/CIFS share setup page.

Small Tree

SERVICES > CIFS/SMB > SHARE > ADD

adamstz8.local

System  
Network  
Disks  
**Services**  
Access  
Status  
Diagnostics  
Advanced  
Help

Settings Shares

Name

Comment

Path

Path to be shared.

Read only ☐ **Set read only**  
If this parameter is set, then users may not create or modify files in the share.

Browseable ☒ **Set browseable**  
This controls whether this share is seen in the list of available shares in a net view and in the browse list.

Users allow

This option allows the users and groups that access a share to be specified. Users and groups are specified, delimited by commas. Groups are designated by a @ prefix.

Users deny

The deny option specifies users and groups who are NOT allowed access to the share. It follows the same format as the allow option.

Guest ☒ **Enable guest access**  
This controls whether this share is accessible by guest account.

Inherit permissions ☒ **Enable permission inheritance**  
The permissions on new files and directories are normally governed by create mask and directory mask but the inherit permissions parameter overrides this. This can be particularly useful on systems with many users to allow a single share to be used flexibly by each user.

Recycle bin ☐ **Enable recycle bin**  
This will create a recycle bin on the share.

Hide dot files ☒ **Enable shadow copy**  
This parameter controls whether files starting with a dot appear as hidden files.

Shadow Copy ☒ **Enable shadow copy**  
This will provide shadow copy created by auto snapshot. (ZFS only)

Shadow Copy format

The custom format of the snapshot for shadow copy service can be specified. The default format is auto-%Y%m%d-%H%M%S used for ZFS auto snapshot.

ZFS ACL ☒ **Enable ZFS ACL**  
This will provide ZFS ACL support. (ZFS only)

Hosts allow

This option is a comma, space, or tab delimited set of hosts which are permitted to access this share. You can specify the hosts by name or IP number. Leave this field empty to use default settings.

Hosts deny

This option is a comma, space, or tab delimited set of host which are NOT permitted to access this share. Where the lists conflict, the allow list takes precedence. In the event that it is necessary to deny all by default, use the keyword ALL (or the netmask 0.0.0.0/0) and then explicitly specify to the hosts allow parameter those hosts that should be permitted access. Leave this field empty to use default settings.

Enable Avid Support ☒ **Enable Avid Support**  
Allow users to share Avid media files on this share.

Auxiliary parameters

These parameters are added to [Share] section of smb.conf. Please check the [documentation](#).

Small Tree's suggested settings are:

**Name:** The name of the share. This must match the name of the dataset.

**Comment:** Cannot be blank. Fill in a useful description.

**Path:** Name of /mnt/tank0/datasetname. Select <...> to browse choices and select.

**Read Only:** Off.

**Browseable:** On.

**Guest:** Optional.

**Inherit Permissions:** This will allow all new files and directories created in this share to have the same default permissions as the share.

**Recycle bin:** Off.

**Hide dot files:** On. This allows Windows systems to use tools for determining hidden files.

**Shadow Copy:** On. This only allows ZFS to take a snapshot for a shadow copy. Shadow copy must be configured on the Windows client.

**Shadow Copy Format:** Do not modify default.

**ZFS ACL:** Off. There are rare cases where this may need to be turned on. Check with Small Tree support before enabling.

**Hosts allow/deny:** This can typically be left blank. Comma separated list of hosts that will have or not have access. Once hosts are defined as allowed, no undefined hosts will have access. Use cautiously.

**Auxiliary Parameters:** Expert level configuration; do not set unless instructed by Small Tree support.

When all settings are configured:

Click “Add”.

Click “Apply Changes”.

The new share will not be available after reboot unless both steps are completed.

## Sharing a Dataset Using NFS

### Services -> NFS -> Shares

This page shows all the existing NFS shares. The default is “share1”.

Select the “+” icon at the lower right of the list. You will be presented with the NFS share setup page.

The screenshot shows the 'Add' page for NFS shares in the Small Tree interface. The top navigation bar is green with the Small Tree logo and the text 'SERVICES > NFS > ADD'. On the left, there is a sidebar with buttons for 'System', 'Network', 'Disks', 'Services' (highlighted in green), 'Access', 'Status', 'Diagnostics', 'Advanced', and 'Help'. The main content area has a 'Settings' tab and a 'Shares' tab. The 'Shares' tab is active, showing a form with the following fields: 'Path' (with a browse button), 'Map all users to root' (set to 'Yes'), 'Authorised network' (set to '24'), 'Comment', 'All dirs', 'Read only', and 'Quiet'. Each field has a small text description below it. At the bottom of the form are 'Add' and 'Cancel' buttons.

Small Tree’s suggested settings are:

**Path:** Select <...> to browse datasets to share.

**Map all users to root:** Yes.

**NOTE: NFS ID mapping is not currently supported.**

**Authorized network:** Can set an IP network to restrict access to only clients on that network. Must be filled in

**Comment:** Description of this share.

**All dirs:** Off. Will share all of the subdirectories of this dataset.

**Read only:** Off. Makes all files in this dataset available to see but not change or edit.

**Quiet:** Off. If checked on, limited diagnostics info is provided.

Connect clients with `nfs://mnt/tank0/datasetname`. Or access in a terminal by `/net/IPADDRESS/mnt/tank0/share1`.

## Creating and Destroying Zpools

When unboxing your TZ unit, there will already be a pool on it. As mentioned earlier, pools are a basic unit of storage for ZFS. The default pool that is on your TZ unit is named “tank0”.

The creation of pools requires Virtual Devices (Vdevs). This means that a pool can only be created if a vdev already exists. If trying to create a pool without a vdev to support it, you will be asked to make a new vdev. Creation and destruction of vdevs will be discussed in its own section, after this one.

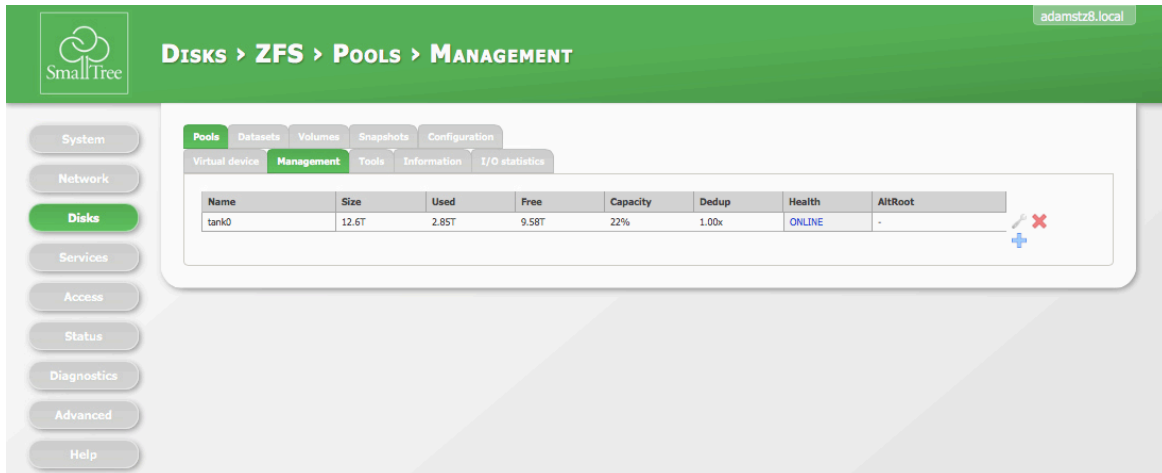
**NOTE: While Small-Tree specifically states not to add or modify a pool without contacting us, this guide will handle how to destroy pre-existing pools, and create pools or sync pools using the WebUI.**

One important thing to remember is that destroying a pool will also destroy all of the dependent data that exists with it. Because datasets all share a zpool, destroying the pool will destroy any and all data sets that are related. This also means that any snapshots of the pool or datasets will be obsolete and will not help to restore any data. It is because of this fact that Small-Tree highly suggests that all data is backed up before attempting this without support.

## Creating a Zpool

### Disks->ZFS->Pools->Management

This screen will display a list of pools on the system. In the picture, there is only the default pool “tank0”



Clicking the blue “+” button will then bring you to this screen:

The settings that are configurable are:

**Name:** The name of your created pool.

**Virtual Devices:** This will display all vdevs available to create a pool. If no virtual devices are visible here, there are none available to create a new pool, and one will need to be created, or a pool destroyed.

**Root:** Allows for the creation of a pool with an alternate route. Default is to leave blank.

**Mount Point:** Set an alternate mount point for root data. Default is /mnt

**Force Use:** Forces the use of a vdev, even if they appear in use or specifically different sizes.

**NOTE: “Force Use” is not recommended without the help of Small Tree**

personnel, as deleting a pool through the WebUI does not delete the pool, but rather cause the system to “forget” about the pool. The Vdev will still remember the pool and it is recoverable UNLESS you select “Force Use”. This will overwrite the pool with all data lost and all snapshots of the pool will not be viable for recovery.

**Description:** A unique description for the pool for your own reference.

When the settings are all configured, click “Add”  
Then click “Apply Changes”

## Destroying a Zpool

### **Disks->ZFS->Pools->Management**

The instructions for deleting a pool are similar to adding one, as you navigate to the same page. To “delete” the pool, simply click the red “X” next to the pool you would like to delete. Note that earlier it was stated that deleting a pool through the WebUI does not actually wipe away the pool, but rather causes the system to simply forget that it is there. In order to delete it permanently through the WebUI, a new pool must be created with the force button selected.

# Recovering a Zpool

## Disks->ZFS->Configuration->Synchronize

This step is most useful for recovering pools and vdevs. The screen below will display:

SmallTree

DISKS > ZFS > CONFIGURATION > SYNCHRONIZE

adamstz8.local

System  
Network  
**Disks**  
Services  
Access  
Status  
Diagnostics  
Advanced  
Help

Pools Datasets Volumes Snapshots **Configuration**

Current Detected **Synchronize**

**Pools (1)**

	Name	Size	Used	Free	Dedup	Health	Mount point	AltRoot
<input checked="" type="checkbox"/>	tank0	12.6T	2.85T (22%)	9.58T	1.00x	ONLINE	-	-

**Virtual devices (1)**

	Name	Type	Pool	Devices
<input checked="" type="checkbox"/>	tank0_disk_0	disk	tank0	/dev/mfid0

**Datasets (6)**

	Name	Pool	Compression	Dedup	Sync	Canmount	Quota	Extended attributes	Readonly	Snapshot Visibility
<input checked="" type="checkbox"/>	ds0	tank0	off	off	standard	on	none	off	off	hidden
<input checked="" type="checkbox"/>	ds54k	tank0	off	off	standard	on	none	off	off	hidden
<input checked="" type="checkbox"/>	dsINT	tank0	off	off	disabled	on	none	off	off	hidden
<input checked="" type="checkbox"/>	rev4-svn	tank0	off	off	standard	on	none	off	off	hidden
<input checked="" type="checkbox"/>	test	tank0	off	off	disabled	on	none	off	off	hidden
<input checked="" type="checkbox"/>	vol0	tank0	off	off	standard	on	none	off	off	hidden

**Volumes (0)**

Name	Pool	Size	Blocksize	Sparse	Compression	Dedup	Sync
------	------	------	-----------	--------	-------------	-------	------

**Options**

Leave auto snapshot configuration ☒ Leave already configured auto snapshots.

Import disks ☒ Import disks used in configuration.

Overwrite disks configuration ☐ Overwrite already configured disks (only affects filesystem value).

Synchronize

Simply clicking the “Synchronize” button at the bottom of the page is sufficient to restore any deleted pools or vdevs. It is also useful for synching any data if moving the drives to a new head unit.

## Creating a Virtual Device

### Disks->ZFS->Pools->Virtual Device

Similar to the pools, when unboxing your TZ unit, there will be a vdev or several, depending on which model you own. Vdevs are devices made up of block devices that serve as the actual storage for your TZ product.

The creation of vdevs requires memory. Small Tree makes their vdevs from their raid devices, which means your device will be able to support only as many raid cards in the device. This

means that a TZ-5 or TZ-8 can support 1, while a TZ-16 can support 2, etc. If there are no more raid cards to support a new vdev, one cannot be created.

**NOTE: While Small-Tree specifically states not to add or modify vdevs without contacting them, this guide will handle how to destroy pre-existing vdevs, and create vdevs or sync vdevs using the WebUI.**

One important thing to remember is that destroying a vdev will also destroy all of the dependent data that exists with it. Because datasets all share a zpool, and zpools are built from vdevs, destroying the vdev will destroy any and all pools, and the datasets in them. This also means that any snapshots of the pool or datasets will be obsolete and will not help to restore any data. It is because of this fact that Small-Tree highly suggests that all data is backed up before attempting this without support.

To start the process, click on the blue “+” button and you will be redirected to the following screen:

The settings that are configurable on this page are as follows:

**Name:** The display name of the new vdev

**Type:** Specify the type of vdev. In order to support a pool, it must be set to the type of “data”

**Devices:** All available devices will be displayed in this window. If there are no devices

available, then no vdev can be created.

**Advanced Format:** Checking this will allow

**Description:** A unique description for the pool for your own reference.

## Creating an iSCSI Target

There are several steps involved in creating an iSCSI target. You need to:

1. Create a volume.
2. Setup at least one portal group.
3. Create an initiator group.
4. Associate an extent to the volume.
5. Add a target.
6. Enable the iSCSI service.


## Creating a Volume

### Disks -> ZFS -> Volume

This screen will display a list of volumes defined on the system. A volume is a quantity of reserved space in a pool to be used as a block device.



Select the “+” icon at the lower right of the list. You will be presented with the volume creation page.


adamstz8.local

## DISKS > ZFS > VOLUMES > VOLUME > ADD

System

Network

**Disks**

Services

Access

Status

Diagnostics

Advanced

Help

Pools Datasets **Volumes** Snapshots Configuration

Volume Information

**Name**

**Pool** tank0: 12.6T ▾

**Size**   
ZFS volume size. To specify the size use the following human-readable suffixes (for example, 'K', 'KB', 'M', 'Gb', etc.).

**Compression** Off ▾  
Controls the compression algorithm used for this volume. The 'lzjb' compression algorithm is optimized for performance while providing decent data compression. Setting compression to 'On' uses the 'lzjb' compression algorithm. You can specify the 'gzip' level by using the value 'gzip-N', where N is an integer from 1 (fastest) to 9 (best compression ratio). Currently, 'gzip' is equivalent to 'gzip-6'.

**Dedup** Off ▾  
Controls the dedup method.  
**NOTE/WARNING:** See [ZFS volumes & deduplication](#) wiki article BEFORE using this feature.

**Sync** standard ▾  
Controls the behavior of synchronous requests.

**Sparse Volume** ☐ Use as sparse volume. (thin provisioning)

**Description**   
You may enter a description here for your reference.

Add Cancel

Small Tree's suggested settings are:

**Name:** Can use any name to identify this volume.

**Pool:** There is a pulldown menu to select an available pool.

**Size:** Required space for iscsi target in kilobytes-K, megabytes-M, or gigabytes-G.

**Compression:** Off.

**Dedup:** Off.

**Sync:** Disabled.

**Sparse Volume:** Can check to not use storage until needed. Be careful not to over allocate storage.

**Description:** Add any description for this volume.

When all the settings are configured:

Click “Add”.

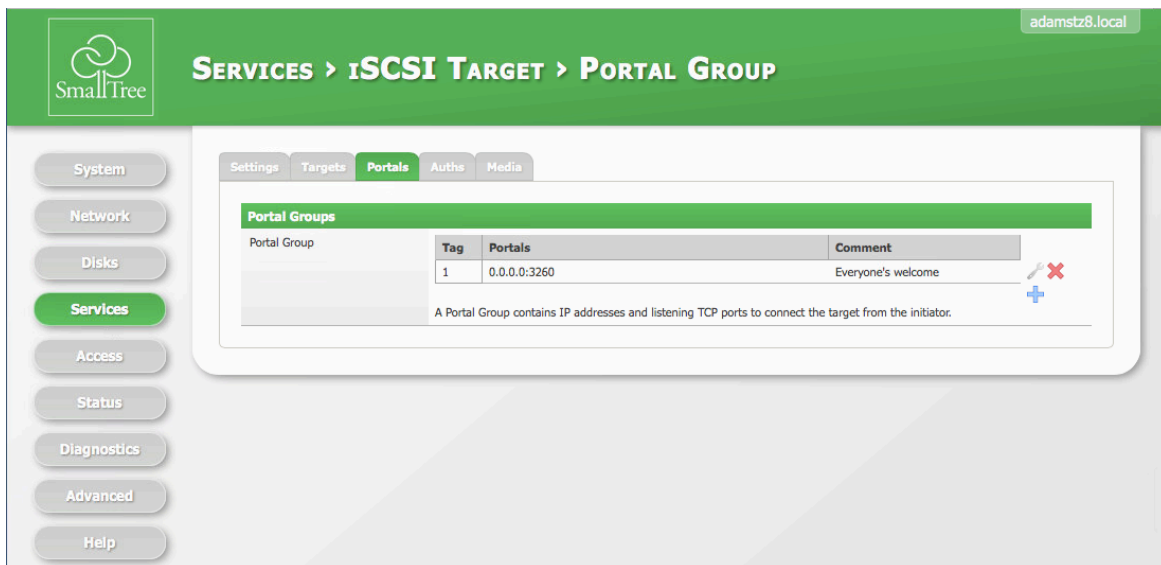
Click “Apply Changes”.

The new volume will not be created unless both steps are completed.

## Adding a New Portal Group

Services -> iSCSI Target -> Portals

You should see this screen:



The screenshot shows the SmallTree iSCSI Target configuration interface. The top navigation bar is green and contains the SmallTree logo, the breadcrumb "SERVICES > iSCSI TARGET > PORTAL GROUP", and the username "adamstz8.local". On the left is a sidebar with buttons for System, Network, Disks, Services (highlighted in green), Access, Status, Diagnostics, Advanced, and Help. The main content area has tabs for Settings, Targets, Portals (selected), Auths, and Media. Below the tabs is a "Portal Groups" section with a table. The table has columns for Portal Group, Tag, Portals, and Comment. There is one entry with Tag "1", Portals "0.0.0.0:3260", and Comment "Everyone's welcome". To the right of the table are edit and delete icons. Below the table is a blue "+" icon and a red "X" icon. A note at the bottom states: "A Portal Group contains IP addresses and listening TCP ports to connect the target from the initiator."

Portal Group	Tag	Portals	Comment
	1	0.0.0.0:3260	Everyone's welcome

Click “+” to add a portal group. You should see this screen.

Small Tree's suggested settings:

**Tag Number:** You can have many portal groups. Each will need a unique number.

**Portals:** Use default. This is the IP address and port number that the iSCSI service will listen on. O's in the IP address will use all. 3260 is the standard port for the service.

**Comment:** Optional comment. When all settings are configured: Click "Add".

Click "Apply Changes".

The new volume will not be created unless both steps are completed.

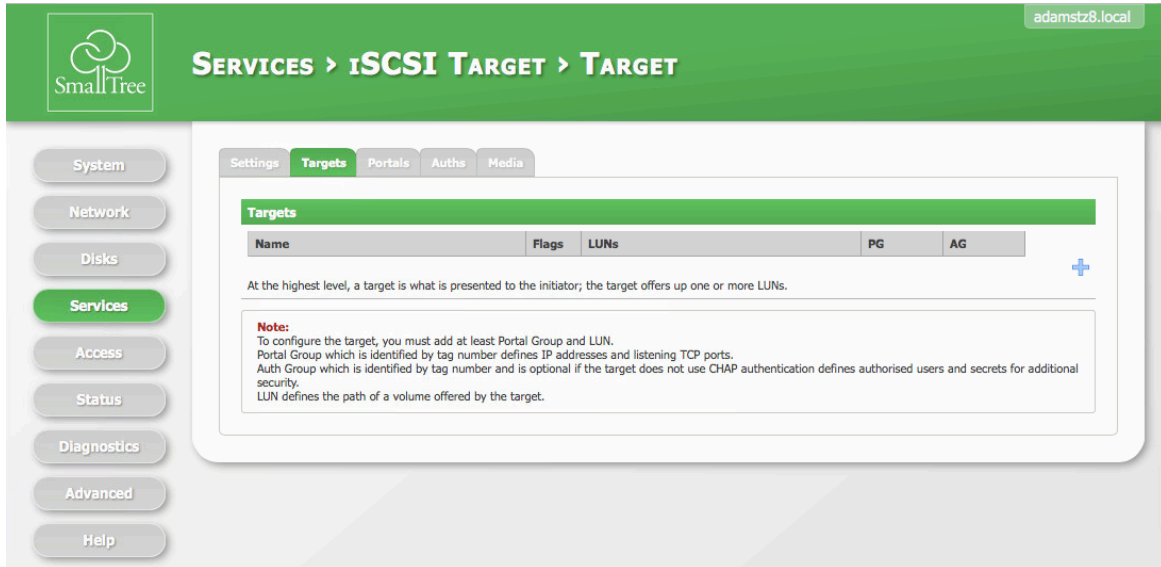
## Adding an Initiator Group

An initiator group is a group of IP addresses that can initiate connections to the iSCSI target.

## Creating a Target

**Services -> iSCSI Target -> Target**

You should then see this page:



## Add a Target

From the Target screen, click “+” to add a target. You should see this screen:

Small Tree

SERVICES > iSCSI TARGET > TARGET > ADD

adamst28.local

System  
Network  
Disks  
**Services**  
Access  
Status  
Diagnostics  
Advanced  
Help

Settings Targets Portals Auths Media

**iSCSI Target** Enable

Target Name   
Base Name will be appended automatically when starting without 'iqn.'

Target Alias   
Optional user-friendly string of the target.

Type   
Logical Unit Type mapped to LUN.

Flags

Portal Group   
The initiator can connect to the portals in specific Portal Group.

Comment   
You may enter a description here for your reference.

**LUN0**

Storage   
The storage area mapped to LUN0.

**LUN1** Enable

**LUN2** Enable

**LUN3** Enable

**Advanced settings**

Auth Method   
The method can be accepted by the target. Auto means both none and authentication.

Auth Group   
The initiator can access to the target with correct user and secret in specific Auth Group.

Initial Digest   
The initial digest mode negotiated with the initiator.

Queue Depth   
0=disabled, 1-255=enabled command queuing with specified depth. The recommended queue depth is 32.

Write Cache ☒ Enable write cache mode.  
It can be changed from the client side by standard SCSI mode page at any time.

Inquiry Vendor   
You may specify as SCSI INQUIRY data. Empty as default. (up to 8 ASCII chars)

Inquiry Product   
You may specify as SCSI INQUIRY data. Empty as default. (up to 16 ASCII chars)

Inquiry Revision   
You may specify as SCSI INQUIRY data. Empty as default. (up to 4 ASCII chars)

Inquiry Serial   
You may specify as SCSI INQUIRY data. Empty as default. (up to 16 ASCII chars)

Logical Block Length   
You may specify logical block length (512 by default). The recommended length for compatibility is 512.

Add Cancel

Small Tree's suggested settings:

**Target Name:** Use default. **Target**

**Alias:** Optional name.

**Type:** Use pulldown menu to select Disk.

**Flags:** Read/Write. You can also make target read only.

**Portal Group:** Use default Tag1. Contact Small Tree support if you need to configure multiple groups.

**Comment:** Optional comment.

**Storage:** Use pulldown menu to select the extent you want to use for this target.

On the bottom half of this screen there are advanced settings. You can set up multiple volumes/extends on each target and add specific authentication methods. Contact Small Tree support for assistance with these configurations.

Make sure the enable service button is checked. When all settings are configured: Click "Add".

Click “Apply Changes”.

The iSCSI Target is then ready to be used by a client with an iSCSI initiator.

## Snapshots

Snapshots are datasets that are point in time, read-only copies of a dataset or volume. Snapshots can be created almost instantly, and they initially consume no disk space from the pool.

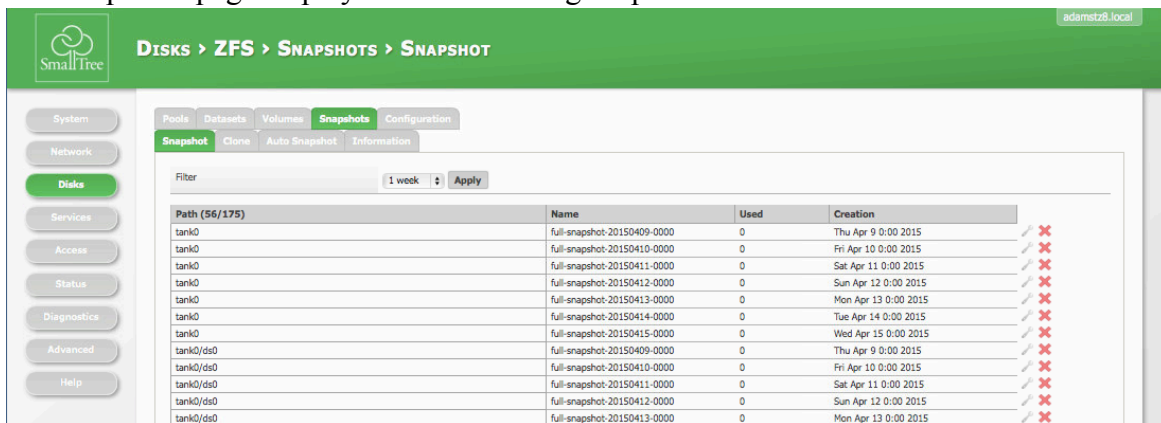
However, as data within the active dataset or volume changes, the snapshot consumes disk space by continuing to reference the old data, preventing the disk space from being freed. This is referred to as “copy on write.”

Snapshots may be created one at a time or automatically at defined time intervals. The only limit to the number of snapshots is available disk space in the pool.

## Creating a Snapshot

### Disks -> ZFS -> Snapshots -> Snapshot

The snapshots page displays a list of existing snapshots.

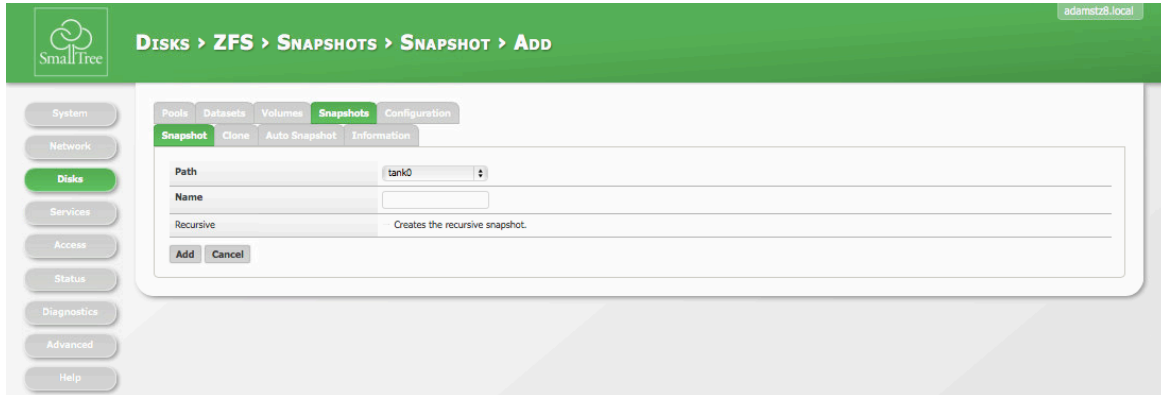


The screenshot shows the 'Snapshots' page in the Small Tree management interface. The breadcrumb path is 'DISKS > ZFS > SNAPSHOTS > SNAPSHOT'. The left sidebar contains navigation links for System, Network, Disks (selected), Services, Access, Status, Diagnostics, Advanced, and Help. The main content area has tabs for Pools, Datasets, Volumes, Snapshots (selected), and Configuration. Under the Snapshots tab, there are sub-tabs for Snapshot, Clone, Auto Snapshot, and Information. A filter dropdown is set to '1 week' with an 'Apply' button. Below the filter is a table of snapshots.

Path (56/175)	Name	Used	Creation		
tank0	full-snapshot-20150409-0000	0	Thu Apr 9 0:00 2015	✓	✗
tank0	full-snapshot-20150410-0000	0	Fri Apr 10 0:00 2015	✓	✗
tank0	full-snapshot-20150411-0000	0	Sat Apr 11 0:00 2015	✓	✗
tank0	full-snapshot-20150412-0000	0	Sun Apr 12 0:00 2015	✓	✗
tank0	full-snapshot-20150413-0000	0	Mon Apr 13 0:00 2015	✓	✗
tank0	full-snapshot-20150414-0000	0	Tue Apr 14 0:00 2015	✓	✗
tank0	full-snapshot-20150415-0000	0	Wed Apr 15 0:00 2015	✓	✗
tank0/ds0	full-snapshot-20150409-0000	0	Thu Apr 9 0:00 2015	✓	✗
tank0/ds0	full-snapshot-20150410-0000	0	Fri Apr 10 0:00 2015	✓	✗
tank0/ds0	full-snapshot-20150411-0000	0	Sat Apr 11 0:00 2015	✓	✗
tank0/ds0	full-snapshot-20150412-0000	0	Sun Apr 12 0:00 2015	✓	✗
tank0/ds0	full-snapshot-20150413-0000	0	Mon Apr 13 0:00 2015	✓	✗

In the filter, there is a pulldown menu to select viewable snapshots. This will show snapshots created manually and by Auto Snapshot (see below). You can select a time range of snapshots to view. You can also select “All” to see all of the snapshots on the system. After selecting range, click “Apply” to see an updated list. All snapshots should have a unique name.

To create a new snapshot, select the “+” icon on the right side of the page. The list of snapshots can be quite long. You may need to scroll down to see the “+” icon. This screen should appear.



**Path:** Select the pool, dataset or volume to snapshot. The list should show all pools, datasets and volumes available to snapshot.

**Name:** Each snapshot must have a unique name. You can use any name that helps you identify this copy. In the example, typing in the date and time helps keep the name unique.

**Recursive:** Check on, if there are nested datasets below the dataset you have selected and you want all of them to get snapped at the same time. For example, taking a snapshot of tank0 with recursive set will also snapshot share1 and any other datasets in the pool tank0.

Click add and the snapshot will automatically be created

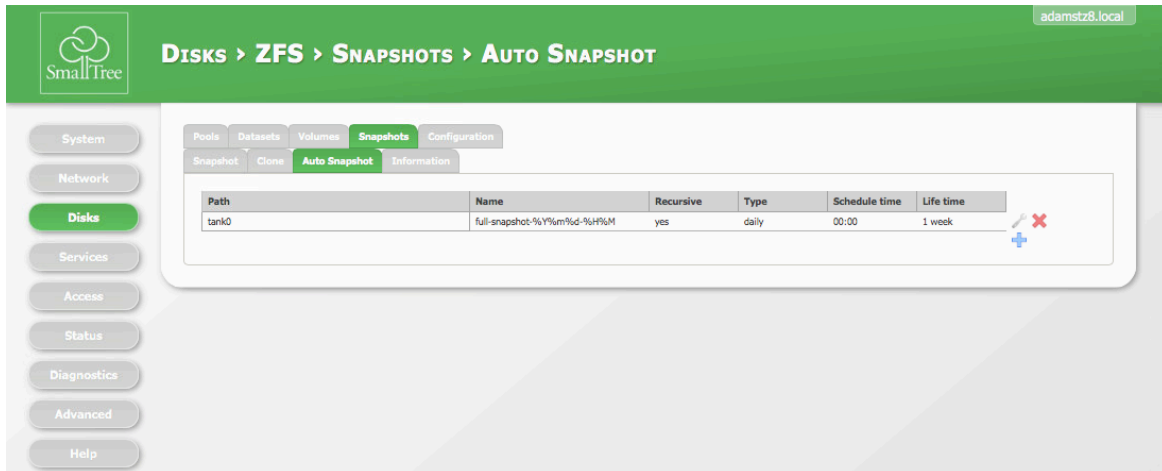
**NOTE:** Snapshots are set-up by default to take one every day at midnight, and set to expire after a week. This means that anything from over a week ago will not be obtainable through the default snapshot setting.

## Creating Automatic Snapshots

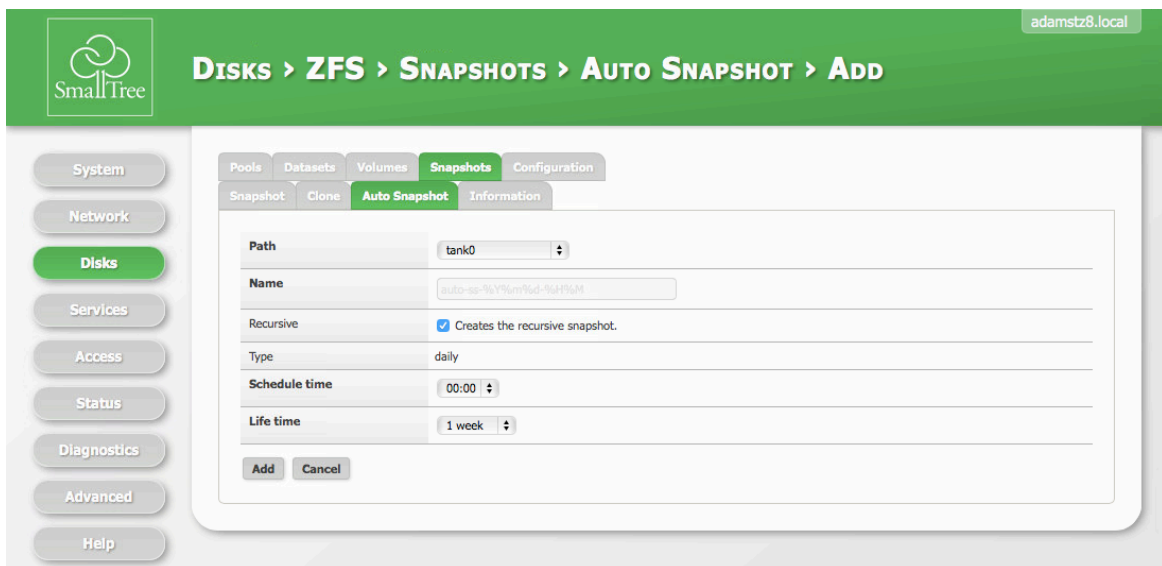
You may setup the TitaniumZ to automatically take snapshots on a scheduled basis. To do this you add an Auto Snapshot policy.

### Disks -> ZFS -> Snapshots -> Auto Snapshot

This page displays a list of existing Auto Snapshot policies.



Click on “+” to configure a new Automatic Snapshot. The following screen should appear.



Small Tree’s suggested settings:

**Path:** From the pulldown menu select the pool, volume or dataset to snapshot.

**Name:** Small Tree recommends you use the default. It creates a snapshot with Year, Month, Day and hour.

**Recursive:** Check on.

**Type:** Daily.

**Scheduled time:** Select a time at which the snapshot should be taken. It is recommended to take at a time with the least storage access.

**Lifetime:** 30 days. Allows admin to define how long snapshots will be saved. You have a choice between 1 week and infinity. The TitaniumZ will automatically remove snapshots that are older than the chosen lifetime.

## Restoring Data From a Snapshot

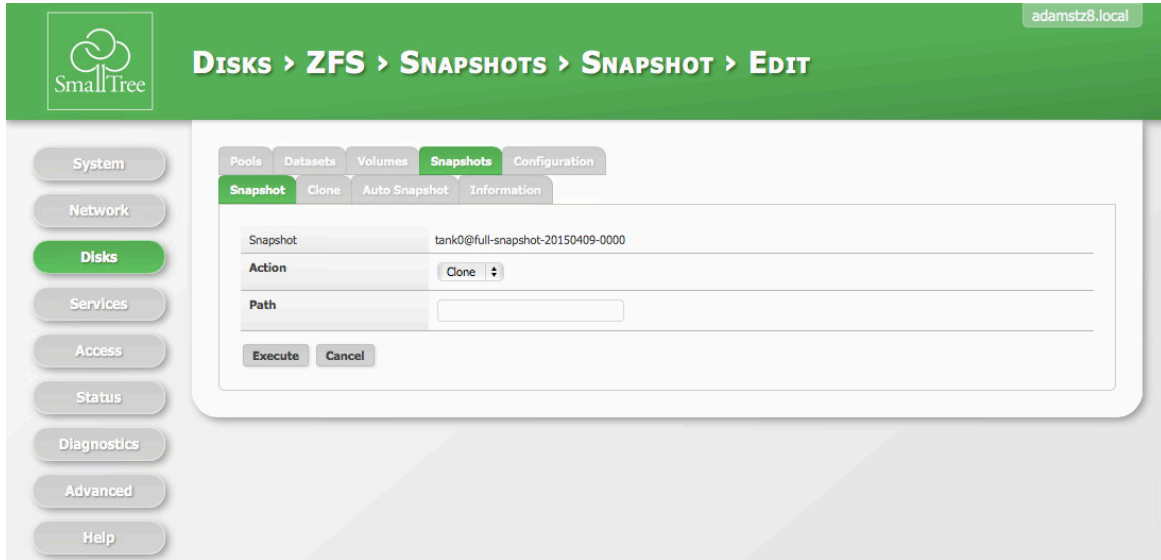
To allow users access to the data in a snapshot it must be cloned and shared.

### Cloning a Snapshot

The first step to provide users read-write access to the files contained in a snapshot is to “clone” the snapshot. This creates special type of dataset (parallel to the existing dataset) and makes the contents accessible as a new dataset. That new dataset will be immediately available and have the identical copies of the files as the snapshot.

### Disks -> ZFS -> Snapshot

Select the wrench icon to edit the snapshot you wish to clone. You will be presented with the snapshot panel.



**Snapshot:** The name of the Snapshot being cloned.

**Action:** Select “Clone”.

**Path:** The location you want the clone to be stored. This is something like putting it into the root directory of a share if you’d like to access it like you do with all other data you’ve added to your pool. This will create a folder with the snapshot clone in it. If your share is at tank0/share1, then pointing the path to something like tank0/share1/clone-copy/ would be a good idea. This way you can see it when you mount your share.

Click “Execute”. It will then navigate you back to the Snapshot tab. If you navigate back to the “Clone” tab, you should then be able to see the clone by clicking the clone tab. The second step is to share the new dataset created using the sharing services such as AFP or SMB. (See pages 27-29 if not familiar with this) If you followed the example in the “Path” section above, you can access the share simply by mounting your share as usual. This means that if you want to access the files, simply connect to the server like usual, and look for the new folder you created. The other approaches to this are to place the share in a unique place or somewhere that you give specific access to. In this case, it may be worth it to go into share creation and set permissions and things so that only you can access the clone. This would be a good idea if dealing with a situation where you want to have exclusive privilege to the snapshot, like in an administrative position.

Once you are done with the snapshot clone that you have been taking data from, and you went the route of creating its own share to access it you will need to also remove it and its contents in order to ensure you don’t waste any space on your storage device. You will start this process by un-mounting the share. You’ll then want to remove the share in the WebGUI. If you didn’t make a new share, then you can skip that step. Finally, you’ll want to delete the clone. This can be done under “ZFS->Snapshots->Clone”. Deleting the clone will ensure that all the storage space

is given back to the system and that the snapshot does not sit idly taking up space on the disk.

## View Current Configuration

### Disks -> ZFS -> Configuration -> Current

This screen displays a summary of all of the current ZFS components:

1. Pools
2. Virtual devices (vdevs)
3. Datasets
4. Volumes

**SmallTree** adamstz8.local

**DISKS > ZFS > CONFIGURATION > CURRENT**

System  
Network  
**Disks**  
Services  
Access  
Status  
Diagnostics  
Advanced  
Help

Pools Datasets Volumes Snapshots **Configuration**

**Current** Detected Synchronize

**Pools (1)**

Name	Size	Used	Free	Dedup	Health	Mount point	AltRoot
tank0	12.6T	2.85T (22%)	9.58T	1.00x	ONLINE	-	-

**Virtual devices (1)**

Name	Type	Pool	Devices
tank0_disk_0	disk	tank0	/dev/mfid0

**Datasets (4)**

Name	Pool	Compression	Dedup	Sync	Canmount	Quota	Extended attributes	Readonly	Snapshot Visibility
ds0	tank0	off	off	disabled	on	none	off	off	hidden
ds64k	tank0	off	off	standard	on	none	off	off	hidden
dsINT	tank0	off	off	disabled	on	none	off	off	hidden
test	tank0	off	off	disabled	on	none	on	off	hidden

**Volumes (0)**

Name	Pool	Size	Blocksize	Sparse	Compression	Dedup	Sync
------	------	------	-----------	--------	-------------	-------	------

**Note:**  
This page reflects the configuration that has been created with the WebGUI.

**NOTE:** The other items in the Disks menu should be ignored. They are not supported with ZenOS.

## Your Data and Rsync

### Services->Rsync

Included in your Titanium Z is the file-copying tool, Rsync. Rsync is a program that is widely used for backups and mirrors. It is used to either copy files to or from a remote host, or locally on the current host. There are two different options on how to connect to carry out an Rsync transfer. The following sections will explore both ways on how to connect using your Titanium unit, and what some of the Rsync terminology stands for.

## Chapter 4: Users and Groups

The purpose of creating multiple users and groups on the TitaniumZ is to control access to the system's shares, volumes, directories and files. If you have no requirement to control access, you do not need to create additional users and groups.

ZenOS offers simple access control mechanisms. For more advanced options, see “Configuring Directory Support” later in this section.

A *user* is anyone who uses a computer. In this case, we are describing the names that represent those users. It may be Mary or Bill, and they may use the names Dragonlady or Pirate in place of their real names. If you want to control access by user, a user account must be set up for every user accessing the TitaniumZ.

**NOTE: The TitaniumZ default configuration has one user created, “edit1.” For ease of management in environments where access control is deemed unnecessary, you may have all users access the TitaniumZ as edit1.**

A *group* is set of users in which all members have the same access controls applied. Groups may be as simple as Employees and Guests, or Edit Department, Graphics Department, etc. or Show1, Show2, or Small Tree, Contractors, etc.

**NOTE: All users must have a primary group. Edit1 is set up with “Staff” as its primary group. You may use this as your primary group for all users or create new groups to create logical groupings of users.**

Users and groups on TitaniumZ can be added, modified and removed using the Access menu.

### Chapter

# 4

# Managing Groups

## Add a Group

Access -> Users and Groups -> Groups

The following screen will be displayed, listing all the existing groups. There are a large number of Predefined system groups. You may not add to, delete from or modify the predefined systems groups.

Click on the “+” under Local user groups to add a group. This screen should appear:

SmallTree

ACCESS > GROUPS > ADD

adamstz8.local

System  
Network  
Disks  
Services  
Access  
Status  
Diagnostics  
Advanced  
Help

Users Groups

Name   
Group name.

Group ID   
Group numeric id.

Description   
Group description.

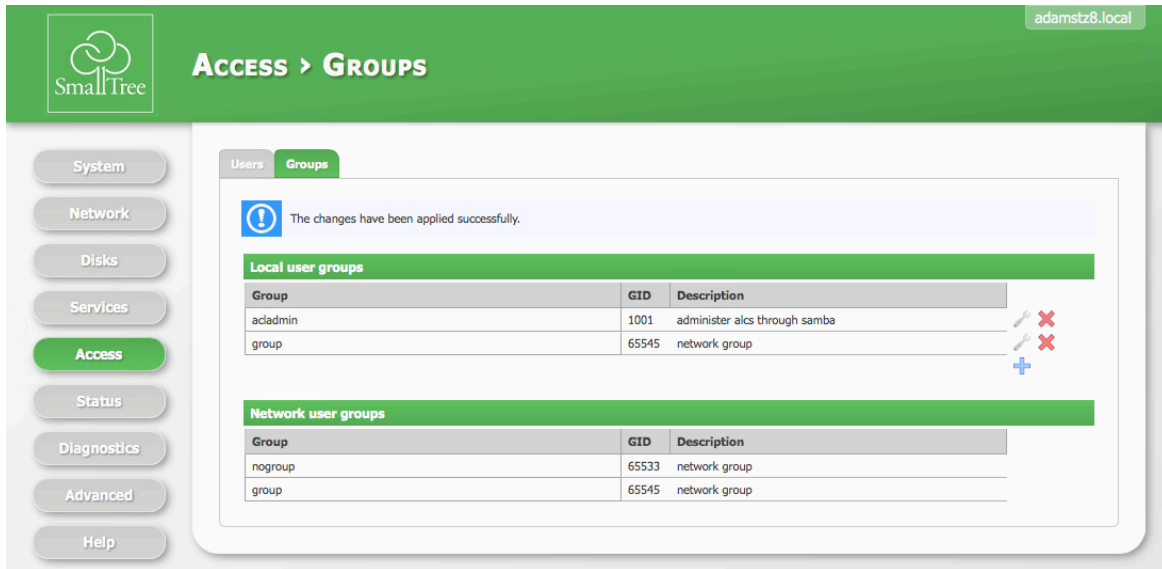
Add Cancel

Add the following information:

**Name:** A short, unique name for the Group. No spaces.

**Group ID:** The TitaniumZ will automatically assign the next available group ID.

**Description:** A general description of this group.



Click “Add.”

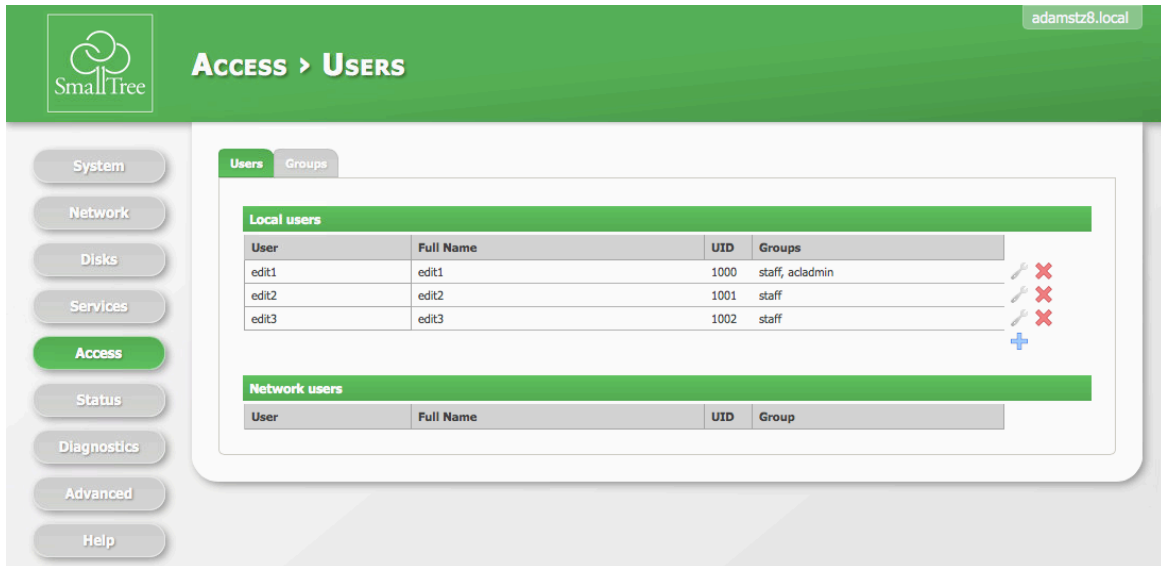
Click “Apply Changes” to complete.

## Managing Users

### Add a User

**Access -> Users and Groups -> Users**

You should see this screen, which displays all current users.



Click on the “+” on the bottom right of the list.

This screen will be displayed:

Each user account consists of:

**Name:** 8-32 Characters, 8 characters recommended.

**Full name:** User's real name (optional).

**Password:** Add a password.

**User ID:** Numeric user ID (or UID). The TitaniumZ will assign the next available User ID.

**Shell:** Nologin.

**Primary group:** Select a primary group from pulldown menu. The default group is "staff".

**Additional group(s):** Optional. Users may belong to up to 16 groups. To select multiple groups use Ctrl-click (or command-click on the Mac). Ctrl-click (or command-click on the Mac) is also used to remove groups.

**Home directory:** Leave blank.

**User Portal:** Allows user access to the WebGUI.

Click “Add” and then “Apply”.

**NOTE: The Titanium Z follows the Unix/POSIX model of permissions, not the windows model of “rights”. This is the same mechanism used by Mac OS X for tracking permissions and ownership. All files are given a read-write-exec permissions for user, group, all.**

To view the primary group of all users, go to Access -> Users and Groups ->Users. To determine all of the members of a group you must view each user and note the group(s) the user is a member of.

## Configuring Directory Support

Directory services allow a group of systems to have all the users and groups defined in a central location for use by many hosts in a workgroup or organization. The TitaniumZ supports multiple directory services.

## Active Directory Integration

TitaniumZ supports Active Directory on Windows 2003r2, Windows 2008r2 and Windows2012r2. Before you bind the Titanium to your Active Directory Server, there are a couple steps that have to be taken.

### Before AD Setup

#### System->General

There are a couple Steps that need to be taken before and after you bind your Active Directory to the Titanium Z.


From here we will look at an example and follow along. For this example, our Active Directory server will be located at 192.168.9.218. Along with the IP address of the Server, we will also want to know its Domain. This is set during the setup of the Active Directory server, and will be unique to you. For this example, we have set the Domain to “stc.domain.com.” Now it is time to set the Hostname, DNS, and Time Settings on the Titanium Z. Example images are shown below.

Hostname	
Hostname	<input type="text" value="AD_Binding_Example"/> <i>Name of the NAS host, without domain part e.g. zenos.</i>
Domain	<input type="text" value="stc.domain.com"/> <i>e.g. com, local</i>

“Hostname” is related to the TitaniumZ and Domain is in reference to the Active Directory Server.

DNS settings	
IPv4 DNS servers	<input type="text" value="192.168.9.218"/> <input type="text"/> <small>IPv4 addresses (This field is disabled if DHCP is selected in LAN Management.)</small>
IPv6 DNS servers	<input type="text"/> <input type="text"/> <small>IPv6 addresses (This field is disabled if IPv6 is deactivated in LAN Management.)</small>

Set the DNS Servers to point at your Active Directory Server.

Time	
Time zone	<div>America/Chicago</div> <div>Select the location closest to you.</div>
System time	<div><input type="text"/></div> <div>Enter desired system time directly (format mm/dd/yyyy hh:mm) or use icon to select it.</div>
Enable NTP	<input checked="" type="checkbox"/> Use the specified NTP server.
NTP time server	<div><input type="text" value="192.168.9.218"/></div> <div>Use a space to separate multiple hosts (only one required). Remember to set up at least one DNS server if you enter a host name here!</div>
Time update interval	<div><input type="text" value="30"/></div> <div>Minutes between network time sync.</div>

Set the “NTP time Server” to point to your Active Directory. This is because if the two are more than 5 minutes off, you will not be able to authenticate.

## Setup AD Integration

### Access -> Active Directory

Now we are ready to add in our information for Active Directory. When you’ve navigated to the page you should see the following:

**Domain Controller Name:** This is where you will put the name of your Active Directory Server.

**Domain Name (DNS/Realm-Name):** This is the same name we entered above when we were setting up the Hostname.

**Domain Name (NetBIOS-Name):** This would have also been set up on your server when installing Active Directory. This will be in all capital letters.

**Administrator Name:** This is NOT an account that has admin privilege, but is the actual login for the administration account on the Active Directory Server. It is the account that can change and access password information and is not a user in the Active Directory database. This is also set up during installation.

**Administration Password:** This is the password for the Administrator Account.

To fill in the required fields for connecting to your Active Directory server, make sure the enable button in the upper right corner is checked and proceed. After clicking “Save”, Users and Groups will then be queried from the Active Directory server.

## After AD Setup

### Services->SMB/CIFS->Settings

In order to use the Active Directory Server however, you will need to select when it gets queried. Because Active Directory is used with Samba, we will need to navigate to the Samba page and the first four fields on this page will look like the following:

Common Internet File System	
Authentication	<div>Active Directory</div> <div><i>If the Domain Controller is enabled, then this field will be disabled.</i></div>
Max Protocol	<div>SMB2</div> <div><i>SMB2 is for recent OS like Windows 7 and Vista. NT1 is for legacy OS like XP.</i></div>
NetBIOS name	<input type="text"/>
Workgroup	<div><input type="text"/></div> <div>The workgroup in which the server will appear when queried by Windows or SMB clients (maximum 15 characters).</div>

**Authentication:** Make sure this has been switched to say “Active Directory.”

**Max Protocol:** This should use the highest possible protocol at the time as default. During the creation of this, it was what you see above.

**NetBIOS name:** This is typically what is set for “hostname” from the “Systems->General” page, but can be changed as it is unique to the TitaniumZ.

**Workgroup:** This is related to the Active Directory lookup and must match a work group on that server.


When all of these are formatted and saved, you should now be ready to connect and authenticate with your Active Directory Server. The TitaniumZ will query the Active Directory Server every 15 minutes for new users and groups.

## Open Directory Integration

TitaniumZ integrates with Open Directory, the directory typically used with Mac OS X server systems. You can query your users and groups from the Open Directory Server.

### Access -> Open Directory

This screen should appear.


adamstz8.local

**ACCESS > OPEN DIRECTORY**

System  
Network  
Disks  
Services  
**Access**  
Status  
Diagnostics  
Advanced  
Help

Open Directory (LDAP) Enable

URI  
  
The space-separated list of URIs for the LDAP server, e.g. ldap://hostname.server

Base DN  
  
The default base distinguished name (DN) to use for searches, e.g. dc=hostname,dc=local

Bind DN  
  
The DN to bind to the Open Directory server. NOTE: directory admin user required, e.g. uid=diradmin,cn=users,dc=hostname,dc=local

Bind password  
  
 (Confirmation)  
The cleartext credentials with which to bind.

Save

Fill in the appropriate information to access your OD server. The required fields are as follows:

**URI:** This is how you will connect to the server. For Open Directory, it will be by typing:

*ldap://<IP ADDR>* where the “<IP ADDR>” is the ip of the machine, or

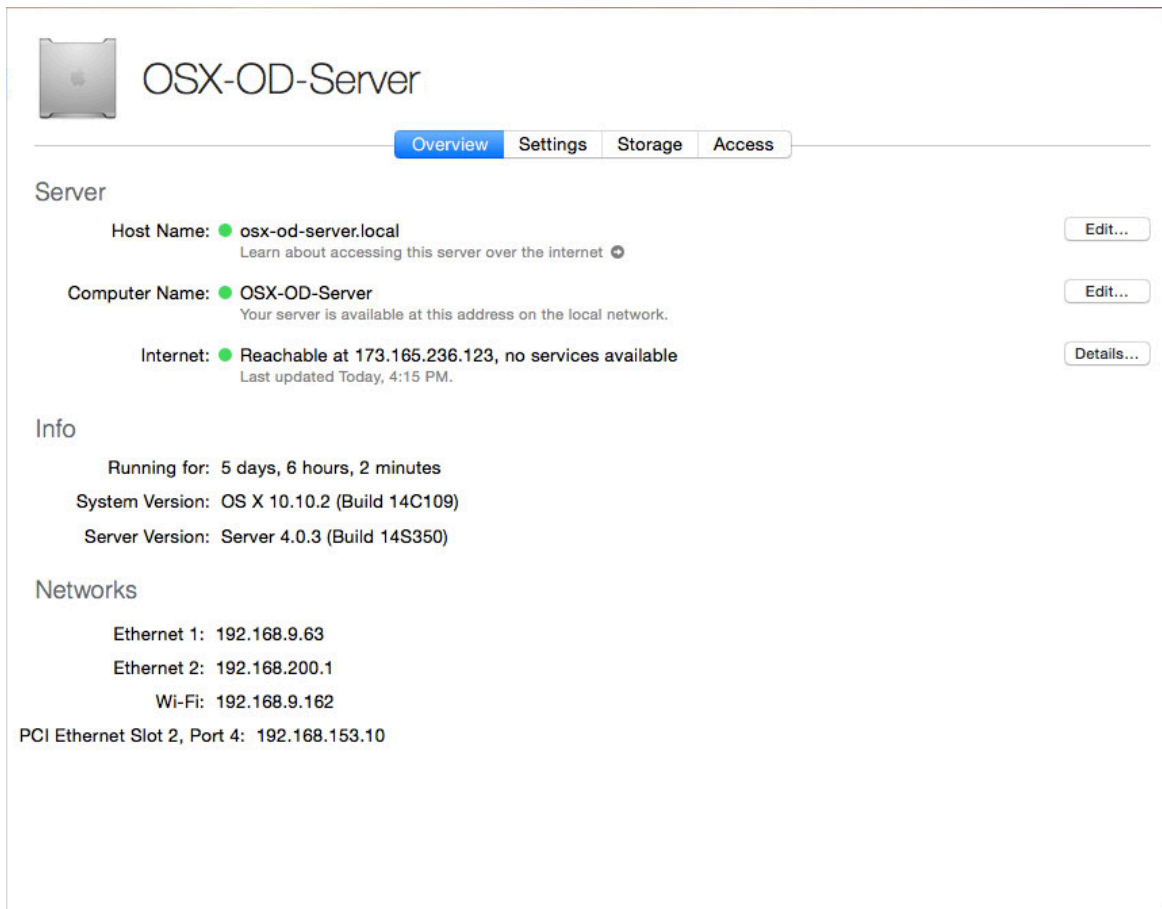
*ldap://<DN>* where “<DN>” is the domain name. In the following example, we will use the “<DN>” form, as often this is easier to read and remember.

**Base DN:** Base DN is domain name. This will be shown in an example below.

**Bind DN:** This is where it is binding. Note that it is required that it bind to the directory admin. This does not mean a group with admin privileges, but the directory admin. If you have not changed your default settings, this will be “diradmin.”

**Bind Password:** The password of the directory admin, so the TitaniumZ can login and authenticate as the directory admin. This is required so that all users and groups can be queried and added to the TitaniumZ’s network users and groups.

Now we will include a picture of a sample Open Directory server’s information based on the view you would get by opening the “server” app:



This page gives us everything we need for the first two fields. The computer name is “OSX-OD-Server”. This computer name is what we will use for our URI, which will then say:

*ldap://OSX-OD-Server*

For the Base DN, we take the host name, and separate the name by where the periods are located. In this example, we write:

*dc=osx-od-server,dc=local*

For the Bind DN, can copy paste everything in the Base DN field, and in front of that info we add in the extra information for this field. Because we are using the default settings from the server, the entire field will be populated as such:

*uid=diradmin,cn=users,cn=groups,dc=osx-od-server,dc=local*

For the password, you would simply put the diradmin password in twice to verify.

After all this information is added, simply click the “Save” button and test mounting a share over AFP with user credentials from a member of the Open Directory Server.

# LDAP Integration

TitaniumZ integrates with LDAP. You can query your users and groups from the LDAP server.

## Access -> LDAP

The following screen should appear.

SmallTree

ACCESS > LDAP

adamstz8.local

System

Network

Disks

Services

Access

Status

Diagnostics

Advanced

Help

**Lightweight Directory Access Protocol** Enable

**URI**  
The space-separated list of URIs for the LDAP server.

**Base DN**  
The default base distinguished name (DN) to use for searches, e.g. dc=test,dc=org

**Anonymous bind**  
☒ Enable anonymous bind.

**Root bind DN**  
The distinguished name with which to bind to the directory server, e.g. cn=admin,dc=test,dc=org

**Root bind password**  
(Confirmation)  
The credentials with which to bind.

**Password encryption**  
clear   
The password encryption protocol to use.

**User suffix**  
This parameter specifies the suffix that is used for users when these are added to the LDAP directory, e.g. ou=Users

**Group suffix**  
This parameter specifies the suffix that is used for groups when these are added to the LDAP directory, e.g. ou=Groups

**Password suffix**  
This parameter specifies the suffix that is used for passwords when these are added to the LDAP directory, e.g. ou=Users

**Machine suffix**  
This parameter specifies the suffix that is used for machines when these are added to the LDAP directory, e.g. ou=Computers

**Auxiliary parameters**  
ldap\_version 3  
timelimit 30  
bind\_timelimit 30  
bind\_policy soft  
perm\_ldap\_attribute uid  
These parameters are added to ldap.conf.

Save

Fill in the appropriate information to access your LDAP server. Click “Save”.

## Access Control Lists

### ACL Support

The Titanium Z supports ACLs on all datasets and volumes. Currently this can only be managed

via the command line on the TitaniumZ.

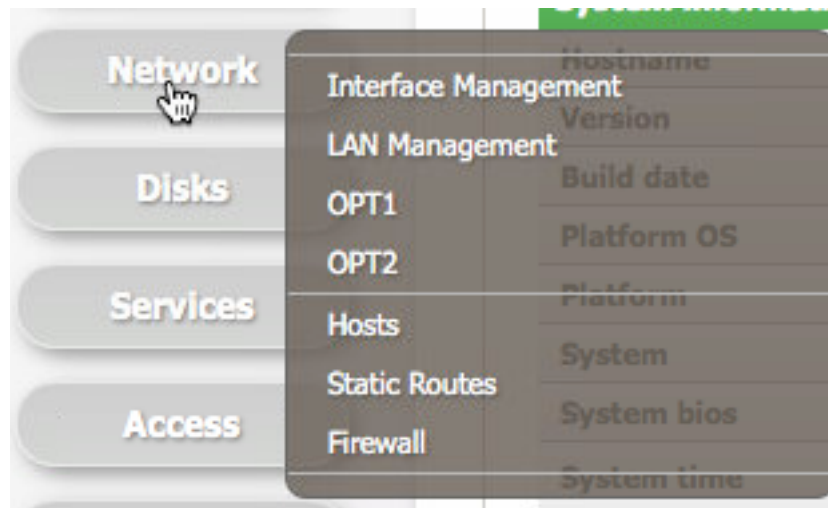
Contact Small Tree support for assistance.

## Chapter 5: Advanced Networking

### Chapter

# 5

You may find it necessary to configure and use multiple network interfaces to provide separation and distribute workload for many clients. The TitaniumZ is equipped with multiple network interfaces. There are a number of ways a network can be configured. Small Tree doesn't recommend any specific network topology. We try to provide the tools necessary to configure the TitaniumZ to fit into your company's network architecture. In the Network menu there are a number of things you can modify.



**Interface Management:** Allows you configure network utilities like vlan tags or link aggregation.

**LAN Management:** Allows you to manage IP configurations.

**OPT1:** Is the setting of the primary interface.

**Hosts:** Allows you to configure access information per host, domain or network for each service.

**Static routes:** Allows you to setup one or more routers per interface. Required for static IP's.

**Firewall:** Allows you to setup firewall rules.

## Interface Naming/Management

Before you get started with modification, it is good to get an understanding of how the physical interfaces are mapped to names inside the Titanium Z.

TitaniumZ models support from 2-20 network interfaces. Gigabit Ethernet (GbE) and 10 gigabit Ethernet (10GbE) interfaces may both be installed in the TitaniumZ, in many combinations. It is important to understand which physical port is configured to each IP address/network. The following assumes a basic understanding of IP networking. Inside of the TitaniumZ each interface has a unique name.

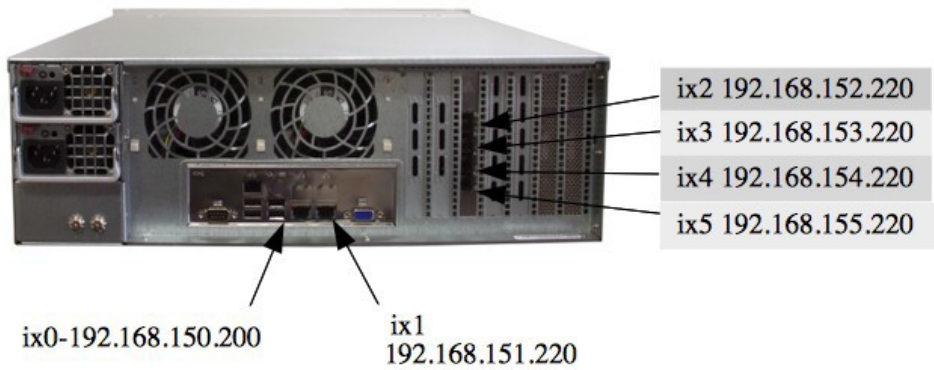
The 10GbE interfaces starts with “ix,” followed by a number starting with 0 and incrementing by 1, to uniquely identify it. The first interface, typically the default interface, will be ix0.

The GbE interfaces start with “em” followed by a number, starting with 0 and incrementing by 1. The first GbE interface will be em0.

**NOTE: If your system has different naming, contact Small Tree technical support.**

The default TitaniumZ configuration uses the first detected Ethernet interface (ix0) and it is set to an IP address of 192.168.150.220. The remaining interfaces are each assigned an address on a different subnet. The first subnet address is 192.168.150.220. The subnet is incremented by 1 for additional interfaces from top to bottom and left to right with an address of this form: 192.168.15X.220. The X is replaced with the next integer in the sequence: 0, 1, 2 etc. The next port will be 192.168.151.220. The next port to the right (at the top of any installed PCIE card) would be 192.168.152.220. For port numbers higher than 9, we would use 192.168.16X.220 and continue to increment the network number.

See example below:



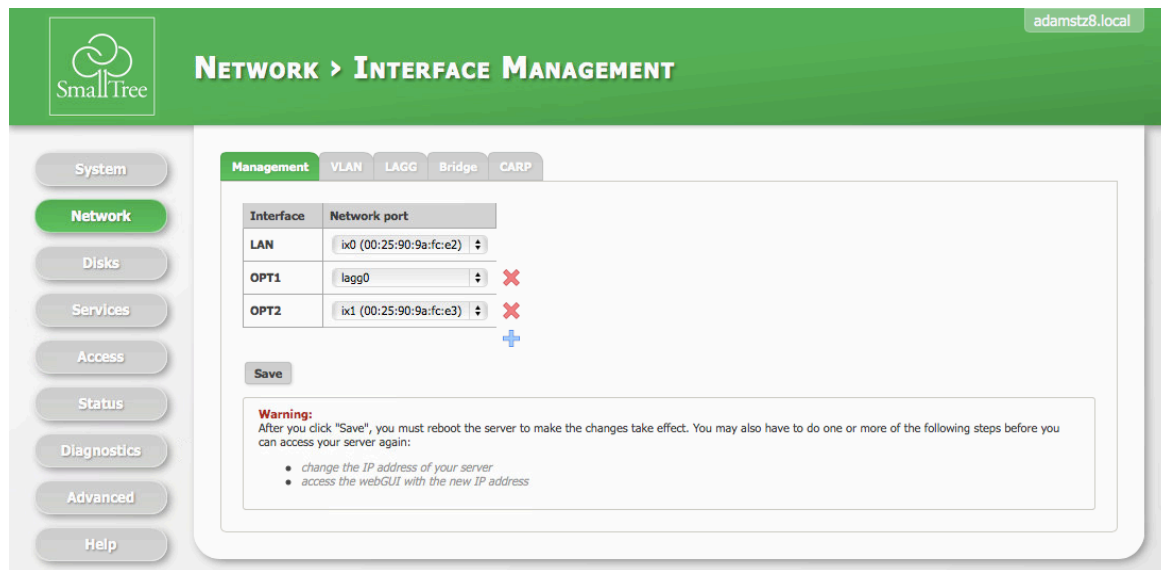
The network adapters are configured by Small Tree and assigned as LAN, OPT1, OPT2, etc. The LAN port is the left-most integrated port (as viewed from the back of the TitaniumZ) and is set to an address of 192.168.150.220. em0 or ix0 will map to LAN, em1 or ix0 to OPT1 and so on.

## Changing the Primary (LAN) Interface

The primary interface is the interface that is assigned the IP address to access the WebGUI and the “first” client workstations to access the shared storage. ZenOS designates this as the LAN interface. To change the primary LAN interface navigate to:

### Network -> Interface Management

This screen will appear:



The interface marked LAN is the primary. To change, use the pull-down menu to select a different port. Click “Save”. A system reboot must occur for the changes to take effect. Remember that the WebGUI will be accessible at the new IP address, and check that there are no users accessing the TitaniumZ prior to rebooting.

**NOTE: Incorrect settings in the network or the TitaniumZ may adversely affect performance of other network devices/services or prevent them from working at all, so make sure you understand your network prior to making changes.**

## Other Network Interface Management Settings

There are additional options in Interface Management including, VLAN, LAGG, Bridge and CARP. It is recommended that you contact Small Tree support to configure these features.

# Changing the Configuration of the LAN Interface

## Network -> LAN Management

The screenshot shows the 'NETWORK > LAN MANAGEMENT' page in the SmallTree web interface. The left sidebar contains navigation buttons: System, Network (highlighted), Disks, Services, Access, Status, Diagnostics, Advanced, and Help. The main content area is divided into three sections: IPv4 Configuration, IPv6 Configuration, and Advanced Configuration. The IPv4 section shows 'Assigned interface' as 'ix0', 'Type' as 'DHCP', 'IP address' as '192.168.9.38 / 24', and 'Gateway' as '192.168.9.5'. The IPv6 section has 'Type' as 'Auto' and an 'Activate' button. The Advanced Configuration section includes 'MTU' (with a text box), 'Media' as 'Autoselect', and 'Extra options' (with a text box). A 'Save' button is at the bottom. A warning message states: 'Warning: After you click "Save", you may also have to do one or more of the following steps before you can access this server again: • change the IP address of your computer • access the webGUI with the new IP address'.

You need to change the IP address of the LAN port, ask your network administrator for IP LAN configuration information to make the TitaniumZ accessible on your network. Otherwise, use the default setting of 192.168.150.220.

You may set these values:

IPv4 Configuration:

**Type:** Static.

**IP Address:** Default is 192.168.150.220.

**Gateway:** Set address of Router for Static, leave blank for DHCP.

**IPv6 Configuration:** Typically not necessary to modify.

Contact Small Tree support if you have questions.

Advanced Configuration:

**MTU:** Typically set to 9216 for video editing.

**Media:** Autoselect is default.

**Extra Options:** Leave blank.

## Creating a Domain Controller

### Services->Domain Controller

#### Creating the Domain Controller

Before starting, make sure that you have an NTP server enabled. If you do not, you will see a warning asking you to set one up when navigating to the page. When you do have one set up and move to the page, you will see it looks like this:

The screenshot shows the SmallTree web interface for configuring a Samba Domain Controller. The 'Enable' checkbox is checked. The 'Realm name' is 'zenos.zen', the 'Domain name' is 'zenos', and the 'Administrator password' field is present with a confirmation field. A note specifies password requirements: 8 characters, 1 capital, and 1 number, with an example 'ChangeM3'. A 'Save and Restart' button is at the bottom.

In the upper right corner, there is an enable box that needs to be checked. Once that is checked, the parameters will be customizable. The default settings that Small-Tree has auto-fill the parameters should then fill in. These are the parameters.

**Realm Name:** zenos.zen  
**Domain Name:** zenos  
**Administrator Password:** (Put your own password here)  
(Confirm your password in the box below)

Remember also that the password you choose should contain one capital letter and one number. The example given for this is “ChangeM3.”

When satisfied with your modifications, Click “Save and Restart.” This will cause the page to reload and should give a blue dialogue box declaring that the changes have been applied successfully.

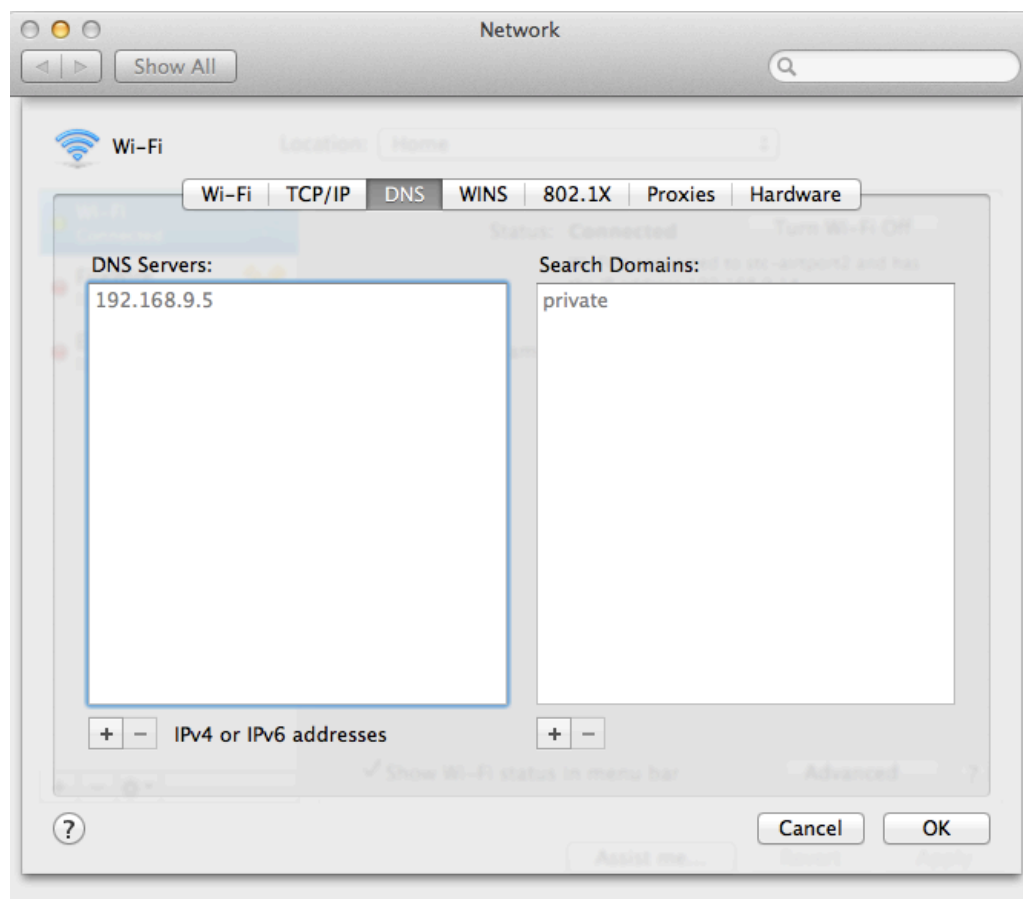
Afterwards, it is suggested that you reboot the system to avoid any possible errors, but it is not required, as the changes you have applied should already be changed.

NOTE: Any and all groups/users that already exist are not added automatically. In order to make them part of the DC server, you will need to make new groups and users, or save them again.

## Binding your Mac

In order to make sure we can now access the server using the groups and users we have, we have to bind our mac to it. To start, navigate to your “Network” page in the “System Preferences.”

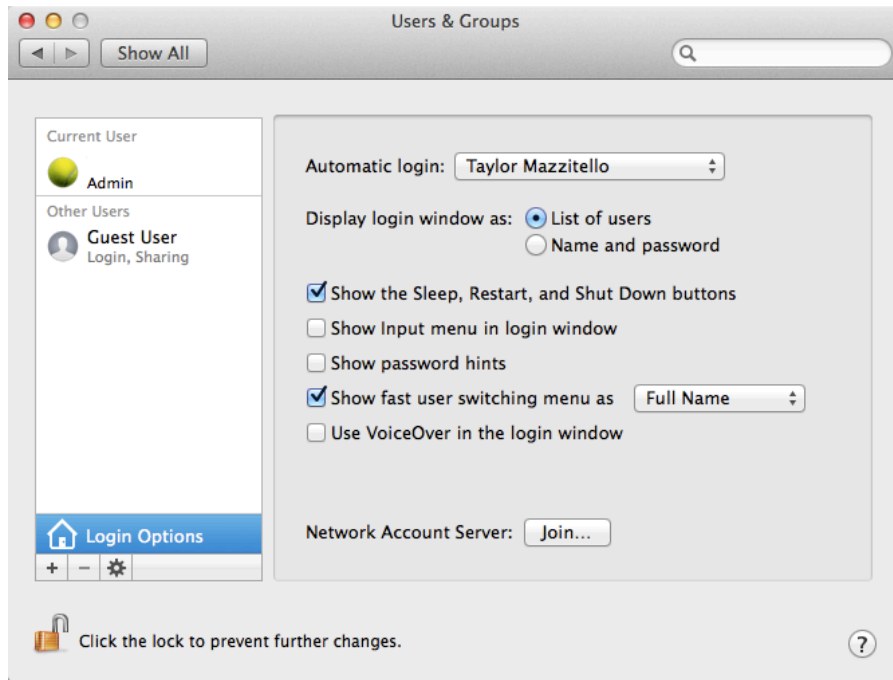
Click on the “Advanced...” settings button. On this new box, click on the DNS tab on the top. It will look like the following:



In the Example above, the IP address of our DNS server is 192.168.9.5. This should be the IP address of the Titanium Z here, so whatever the IP address of your TitaniumZ is, input

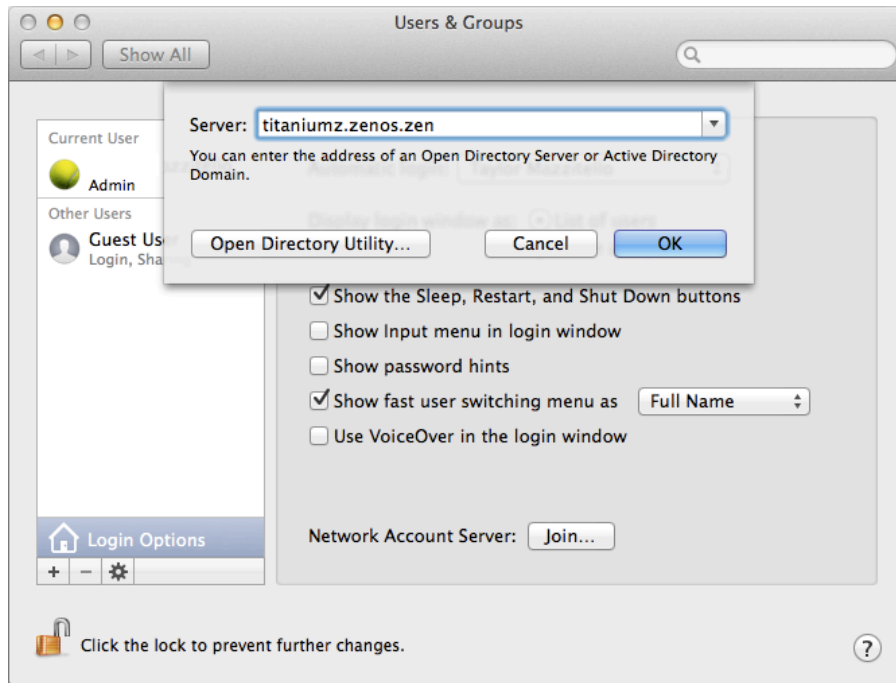
it here with the “+” button below the box.

Next, navigate to the “Users and Groups” page in the “System Preferences” and you will be greeted with the user page. Click “Login Options” under the current user list. Here is a picture for reference:



The button that should be picked is highlighted in blue, and in order to access it, you must unlock the changes lock located directly below it.

From here, click “Join” in order to bring up a screen that will allow you to enter a server name. This is “titaniumz.<Realm name>” a visual example is below that uses the example server name from before.



When clicking ok, it will ask you for “AD Admin User” and for “AD Admin Password.” For AD Admin User, input “Administrator” and for the password, it is whatever you put in on the WebUI when starting this process.

## Chapter 6: System Status and Troubleshooting

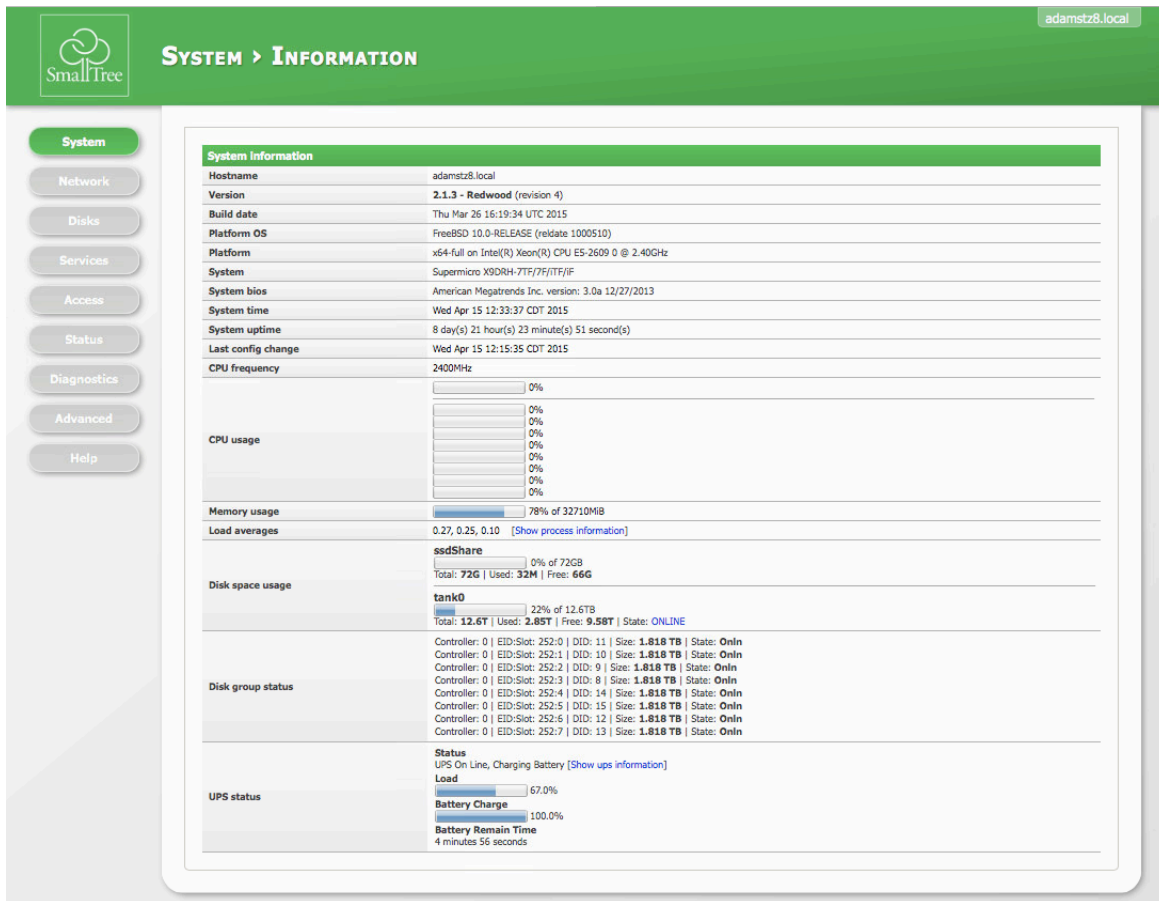
### Chapter

# 6

### Accessing Status Information

#### Status -> System

This screen provides a general system overview. This is the same screen displayed at login.



## Status Information

The TitaniumZ maintains massive amounts of status information. You may be asked by Small Tree support to access this information and show or email this information to Small Tree support.

Additional status information available in the status menu:

**Processes Status:** Shows the processes running on the system in order of runtime.

**Service Status:** Displays running and configured services. A green check in “enabled” means the service should be running. A green check in “status” means the service is configured.

**Interfaces Status:** Displays information about your network adapters. **Disk Status:** Displays information about your TZ-5 configured disks. **Graph:** **Traffic Graph:** Activity of an interface over time. **CPU Load:** CPU Load over time. **Email Report:** Settings to email reports to a specified email address at specified times.

**NOTE: You must have first setup Email under System->Advanced->Email.**

**STATUS > EMAIL REPORT**

**Email Report** Enable

**To email**  
Destination email address. Separate email addresses by semi-colon.

**Subject**  
The subject of the email. You can use the following parameters for substitution:  
• %id - Date  
• %h - Hostname

**Reports**

- ☒ System info
- ☒ System message buffer
- ☒ System log
- ☒ FTP log
- ☒ RSYNC log
- ☒ SSHD log
- ☒ S.M.A.R.T. log
- ☒ Daemon log
- ☐ Custom script

**Polling time**

Minutes					Hours					Days					Months					Week days				
Selected ..					Selected ..					Selected ..					Selected ..					Selected ..				
0	12	24	36	48	0	12				1	13	25			January			Sunday						
1	13	25	37	49	1	13				2	14	26			February			Monday						
2	14	26	38	50	2	14				3	15	27			March			Tuesday						
3	15	27	39	51	3	15				4	16	28			April			Wednesday						
4	16	28	40	52	4	16				5	17	29			May			Thursday						
5	17	29	41	53	5	17				6	18	30			June			Friday						
6	18	30	42	54	6	18				7	19	31			July			Saturday						
7	19	31	43	55	7	19				8	20				August									
8	20	32	44	56	8	20				9	21				September									
9	21	33	45	57	9	21				10	22				October									
10	22	34	46	58	10	22				11	23				November									
11	23	35	47	59	11	23				12	24				December									

Note: Ctrl-click (or command-click on the Mac) to select and de-select minutes, hours, days and months.

**Save and Restart** **Cancel**

## Accessing Diagnostics

### Diagnostics -> Log

The TitaniumZ by default will maintain log information for the following services:

- System
- SSH
- Daemon
- FTP (and FTP transfers)

- S.M.A.R.T
- Rsync(client, server, and local)
- Netatalk3
- UPnP
- DAAP

To access a specific log, select from the pull down menu located above “Date & Time.”

## **Diagnostics -> Log -> Settings**

Allows for the setting of log parameters and configuring a sys log server.

## **Network Diagnostic Tools**

There are three tools used to test network connectivity.

## **Diagnostics ->Ping/Traceroute**

**Ping:** Allows you to test connect between the TitaniumZ and a single IP address or host on the network. This is typically the first test to determine if there is network connectivity.

**Traceroute:** Allows you to see the networks or router which are used for a packet to go from the TitaniumZ to another host. It allows you to see the route or path that is taken. It also allows you to see if the communication stops at a router.

## **Diagnostics ->ARP tables**

**ARP tables:** Maps IP addresses of connected computers to their physical (MAC) addresses.

# **Troubleshooting**

## **Can’t Access a Share**

### **Step1: Check Network Connection**

First, check to make sure the cables are all attached, the TitaniumZ is powered up and connected.

Next, check to see if you can reach the IP address. On a Mac, open a terminal window and ping IP.

```
Jacobs-Mac-mini:~ lkateley$ ping 192.168.150.220
PING 192.168.150.220 (192.168.150.220): 56 data bytes
Request timeout for icmp_seq 0
Request timeout for icmp_seq 1
Request timeout for icmp_seq 2
Request timeout for icmp_seq 3
Request timeout for icmp_seq 4
Request timeout for icmp_seq 5
Request timeout for icmp_seq 6
Request timeout for icmp_seq 7
^C
--- 192.168.150.220 ping statistics ---
9 packets transmitted, 0 packets received, 100.0% packet loss
Jacobs-Mac-mini:~ lkateley$
```

If you see “Request timeout”, there is an addressing problem.

Check the TitaniumZ to make sure it is set to the IP address you are testing.

Check if TitaniumZ has access to a router/gateway on your network.

Check that your system can access other things on the network.

If all of these check okay, contact Small Tree support.

If your ping is successful you should see:

```
Jacobs-Mac-mini:~ lkateley$ ping 192.168.0.61
PING 192.168.0.61 (192.168.0.61): 56 data bytes
64 bytes from 192.168.0.61: icmp_seq=0 ttl=64 time=5.039 ms
64 bytes from 192.168.0.61: icmp_seq=1 ttl=64 time=52.376 ms
64 bytes from 192.168.0.61: icmp_seq=2 ttl=64 time=17.927 ms
64 bytes from 192.168.0.61: icmp_seq=3 ttl=64 time=4.483 ms
64 bytes from 192.168.0.61: icmp_seq=4 ttl=64 time=5.285 ms
^C
--- 192.168.0.61 ping statistics ---
6 packets transmitted, 5 packets received, 16.7% packet loss
round-trip min/avg/max/stddev = 4.483/17.022/52.376/18.381 ms
Jacobs-Mac-mini:~ lkateley$
```

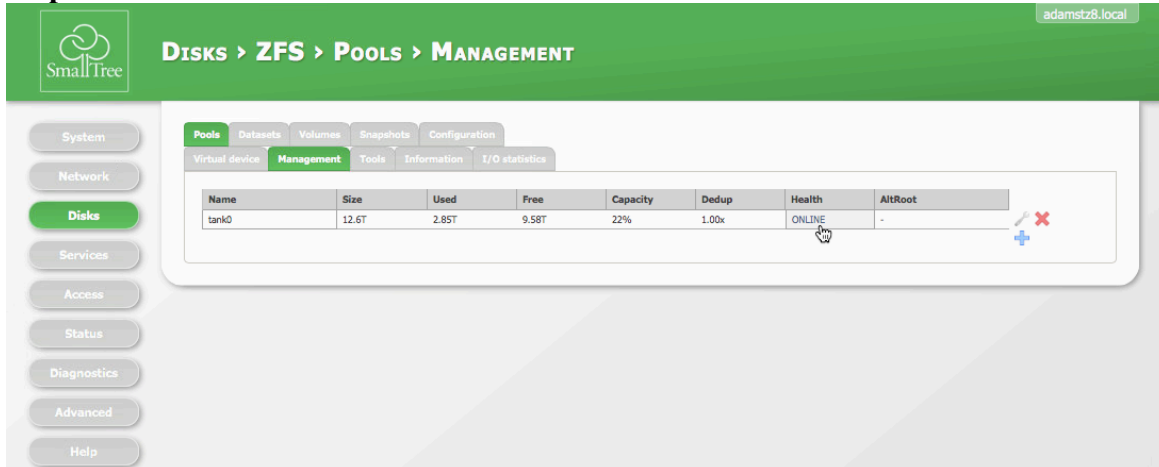
Bytes received from the IP address. If you see this you have connection on the network.

## Step 2: Check Share

First, check to see that the share is being shared and you are sharing with the same protocol you are trying to connect with, eg afp, cifs, nfs.

Second, make sure you are connecting to the server using the right protocol and IP address.

## Step 3: Check Users/Permissions



Name	Size	Used	Free	Capacity	Dedup	Health	AltRoot
tank0	12.6T	2.85T	9.58T	22%	1.00x	ONLINE	-

First, check that you are trying to access the TitaniumZ with a user and password that are defined on the server.

Second, check that permissions and ownership of the share are set correctly for the user you are testing.

## Replacing a Disk

To check the status of your storage navigate to:

### Disks -> ZFS -> Pools -> Management

This page shows health status. Health should always say ONLINE. If you see anything but ONLINE, for example, degraded or offline, you may have a failed disk. Your storage may be inaccessible until the disk is replaced.

To determine which disk is causing the pool not to be ONLINE navigate to:

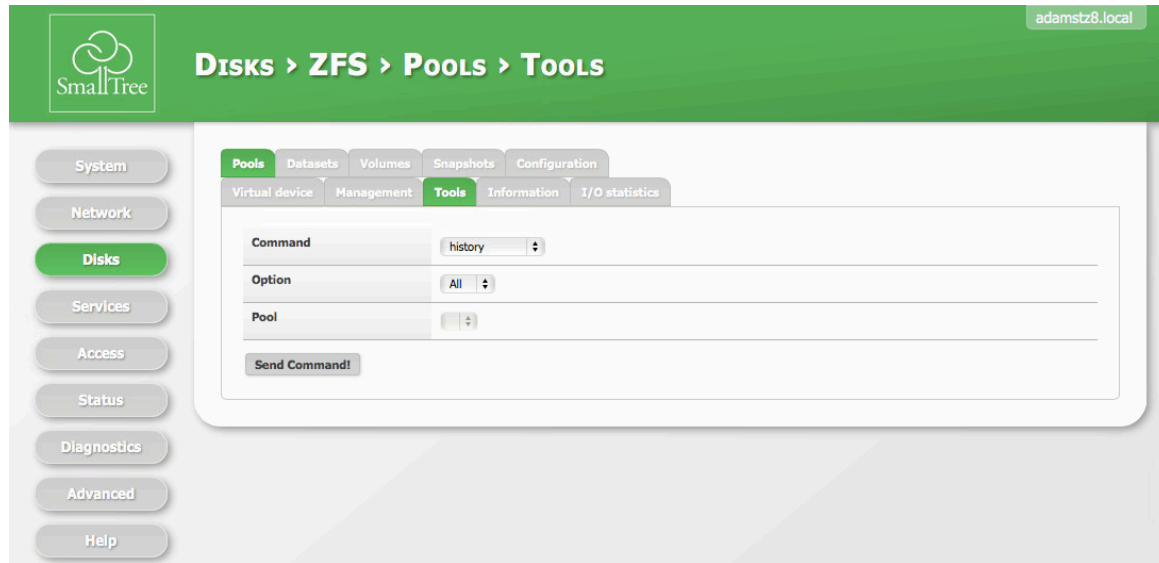
### Status -> Disks

This enumerates all the disks and will show you which disk has a status other than ONLINE. It is important to replace a problem drive as soon as possible to prevent data loss. If you have a TZ8 or a TZ16, the drive that is bad can be found by referencing the physical front panel. A red LED on the chassis will indicate which drive needs to be replaced. Replacing the drive will allow the

RAID card to rebuild the drive. If you are using a TZ5, then you will have to navigate to:

**Disks -> ZFS -> Pools -> Tools**

The following panel will be displayed:



Follow these steps to remove and replace the drive:

1. Locate the drive you wish to remove.  
Find drive serial number under **Disks > Management** and remember the disk number that corresponds to the serial number.
2. Navigate to **Disks > ZFS > Pools > Tools**
3. Set the drive you are removing to offline by using the following parameters:  
**Command:** offline  
**Option:** Device  
**Pool:** Tank  
**Devices:** Check the box next to the drive you want to replace.

Now click **Send Command**.

4. Navigate to **System > Shutdown > Now** and click **Yes** to shutdown the TZ.
5. After the TZ is off, find the drive you wish to replace. Remove this drive from inside the TZ.
6. Make sure that the new replacement drive does not contain any data on it. If it is completely clear, place the new drive in the slot.
7. Power up the TZ
8. You will likely see that it says the pool is degraded. Ignore this for now and navigate to **Disks > ZFS > Pools > Tools**.
9. Set the new drive to replace the old by using the following parameters:

**Command:** replace

**Option:** Device

**Pool:** Tank

**Devices:** Check the box next to the drive you have replaced.

**New Device:** Select the same device that you chose for **Devices**.

Now click **Send Command**.

10. Go to **Disks > Management** and click button **Clear config and Import disks**.

***For all other troubleshooting contact Small Tree Support.***

# Acknowledgements:

Copyright (c) 2006-2008 Volker Theile (votdev@gmx.de)  
All rights reserved.

part of FreeNAS (<http://www.freenas.org>)  
Copyright (C) 2005-2008 Olivier Cochard-Labbé <olivier@freenas.org>.  
All rights reserved.

NAS4Free Copyright © 2012-2017 The NAS4Free Project.  
([info@nas4free.org](mailto:info@nas4free.org)). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.