



OWC Digital

# Jupiter<sup>®</sup> Callisto

Quick Start Guide



# TABLE OF CONTENTS

<b>1. INTRODUCTION</b> .....	<b>1</b>
1.1 CONTENTS	
1.2 SPECIFICATIONS (8-BAY)	
1.3 SPECIFICATIONS (16-BAY)	
1.4 CONNECTIONS AND COMPATIBILITY	
1.5 REAR VIEW	
1.6 SIDE VIEW	
<b>2. ASSEMBLY &amp; REPLACEMENT</b> .....	<b>4</b>
<b>2.2 UNIT RACK INSTALLATION &amp; REMOVAL</b>	
2.2.1 Unit Installation	
2.2.2 Unit Removal	
<b>2.3 DRIVE INSTALLATION &amp; REMOVAL</b>	
2.3.1 Drive Installation	
2.3.2 Drive Removal	
<b>2.4 POWER SUPPLY FAILURE &amp; REPLACEMENT</b>	
<b>2.5 FAN FAILURE</b>	
<b>3. INITIAL SETUP</b> .....	<b>10</b>
3.1 CONNECTION	
3.2 FIRST TIME CONFIGURATION	
3.3 ACCESSING THE IPMI INTERFACE	
3.4 DEFAULT SETUP	
<b>4. STORAGE CONFIGURATION</b> .....	<b>12</b>
4.1 ZFS TERMINOLOGY	
4.2 STORAGE CONFIGURATION	
4.2.1 Creating the Pool	
4.2.2 ZFS RAID Levels	
4.2.3 Optimal RAIDZ Configurations	
4.2.4 Creating Datasets	
4.2.5 Creating a Filesystem	
4.2.6 Creating a Zvol	
4.3 DRIVE IDENTIFICATION	
4.4 DRIVE FAILURE & REPLACEMENT	
<b>5. USERS, GROUPS, &amp; PERMISSIONS</b> .....	<b>19</b>
5.1 CREATING A GROUP	
5.2 CREATING A USER	
5.3 PERMISSIONS	
<b>6. SHARES</b> .....	<b>21</b>

**6.1 CREATING SHARES AND CONFIGURING GUEST ACCESS**

- 6.1.1 Creating an AFP Share
- 6.1.2 AFP Guest Access
- 6.1.3 Creating a Windows (CIFS) Share
- 6.1.4 CIFS Guest Access
- 6.1.5 Creating an NFS Share
- 6.1.6 NFS Guest Access

**6.2 CONNECTING TO SHARES**

- 6.2.1 Connecting to AFP Shares (Mac only)
- 6.2.2 Connecting to CIFS Shares on a Mac
- 6.2.3 Connecting to CIFS Shares in Windows
- 6.2.4 Connecting to an NFS share on a Mac
- 6.2.5 Connecting to an NFS share in Windows

**7. APPENDIX ..... 26**

**7.1 ABOUT DATA BACKUP**

**7.2 ADDITIONAL SUPPORT**

**7.3 TECHNICAL SUPPORT HOURS AND CONTACT INFORMATION**

# INTRODUCTION

## 1.1 CONTENTS



OWC Digital  
Jupiter Callisto (8-bay or 16-bay)



Flash drive containing  
Jupiter manuals,  
drivers, and software



Slide rails



Eight M5 pan head  
screws



Twelve M5 cage nuts



Drives with drive trays  
attached (8 or 16)



Power cables  
(two for 8-bay, three for 16-bay)



Rail extensions



Four M5 enclosure  
securing screws

## 1.2 SPECIFICATIONS (8-BAY)

### Enclosure Specifications

<b>Front Panel Indicators:</b>	Power status, fan status	
<b>Drive Bay Indicators:</b>	Power, activity, fail	
<b>Dimensions:</b>	3.5"(H) x 19"(W) x 23"(D) 8.89cm(H) x 48.26cm(W) x 58.42cm(D)	
<b>Rack Compatibility:</b>	Square hole and round hole mounting Included rails supporting up to 19" (48.26cm) width and 30" (762mm) depth	
<b>Environmental:</b>	Operating 32°F to 104°F 0°C to 40°C 20 to 90% non-condensing	Non-Operating -4°F to 176°F -20°C to 80°C 10 to 90% non-condensing

### Power Supply Unit (PSU) Specifications

<b>Type:</b>	380W High Efficiency 1+1 Redundant hot-swappable with automatic load balancing	
<b>Operating Voltage:</b>	100~240 VAC Full Range Auto-switching with active PFC	
<b>Environmental:</b>	Operating 32°F to 104°F 0°C to 40°C 20 to 95% non-condensing	Non-Operating -4°F to 176°F -20°C to 80°C 10 to 95% non-condensing
<b>MTBF:</b>	223,678 hours	

## 1.3 SPECIFICATIONS (16-BAY)

### Enclosure Specifications

<b>Front Panel Indicators:</b>	Power status, fan status	
<b>Drive Bay Indicators:</b>	Power, activity, fail	
<b>Dimensions:</b>	3.5"(H) x 19"(W) x 23"(D) 8.89cm(H) x 48.26cm(W) x 58.42cm(D)	
<b>Rack Compatibility:</b>	Square hole and round hole mounting Included rails supporting up to 19" (48.26cm) width and 30" (762mm) depth	
<b>Environmental:</b>	Operating 32°F to 104°F 0°C to 40°C 20 to 90% non-condensing	Non-Operating -4°F to 158°F -20°C to 70°C 10 to 90% non-condensing

### Power Supply Unit (PSU) Specifications

<b>Type:</b>	650W 2+1 Redundant hot-swappable with automatic load balancing	
<b>Operating Voltage:</b>	100~240 VAC Full Range Auto-switching with active PFC	
<b>Environmental:</b>	Operating 32°F to 104°F 0°C to 40°C 20 to 95% non-condensing	Non-Operating -4°F to 158°F -20°C to 70°C 10 to 95% non-condensing
<b>MTBF:</b>	283,514 hours	

## 1.4 CONNECTIONS AND COMPATIBILITY

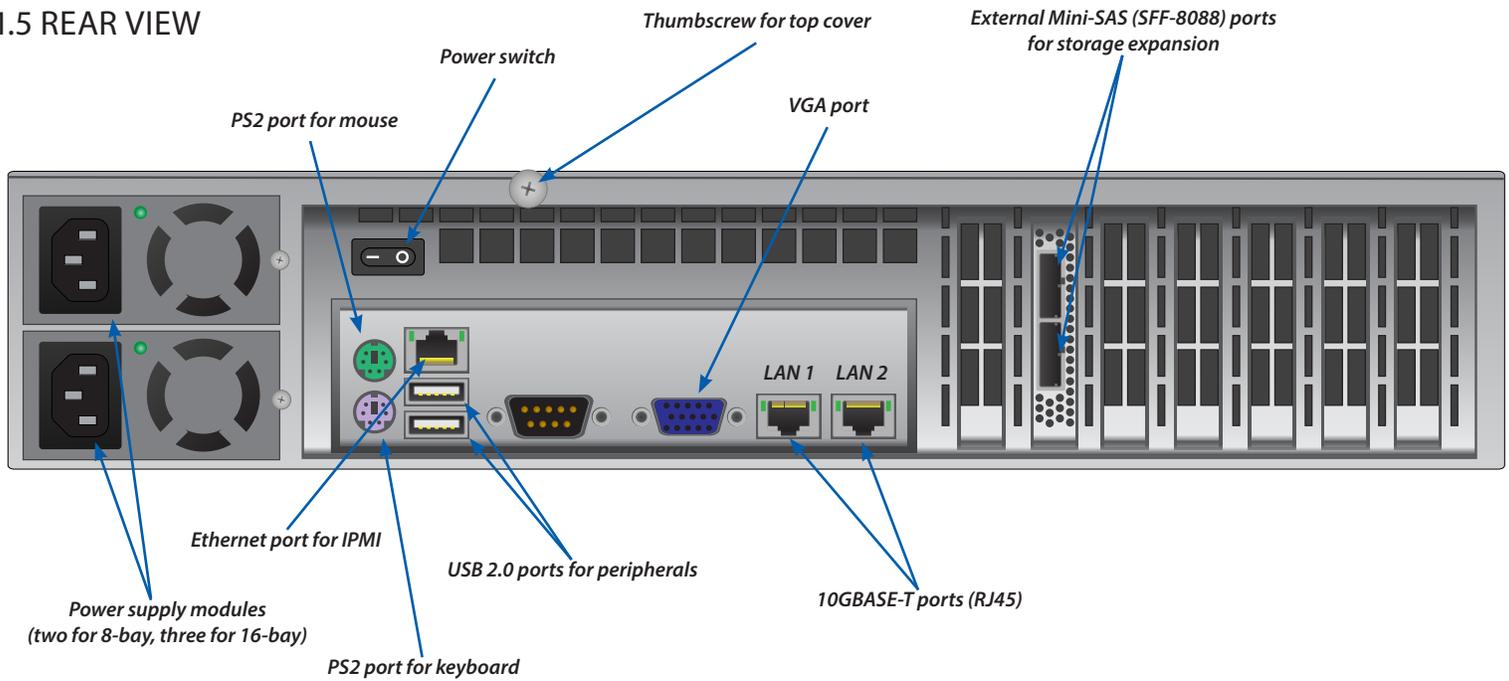
### Connection Types

<b>Internal Drive Connection:</b>	6Gb/s SAS
<b>Client Connection:</b>	10-Gigabit Ethernet (10GBASE-T), RJ45 ports x 2
<b>Storage Expansion:</b>	External Mini-SAS (SFF-8088) ports x 2
<b>IPMI:</b>	Gigabit Ethernet, RJ45 port x 1
<b>Peripheral Attachment:</b>	PS2 ports x 2 (keyboard and mouse); USB 2.0 ports x 2

### Compatibility

<b>Cable Requirements:</b>	10GbE: Cat6a or Cat7 Ethernet cable Gigabit Ethernet: Cat5 or Cat6 cable
<b>Client Requirements:</b>	RJ45 Ethernet port

## 1.5 REAR VIEW



### Ethernet Port Address Presets

Port	Mode	IP Address	Subnet Mask
IPMI:	DHCP	automatic*	automatic*
LAN 1 (ix0):	DHCP	automatic	automatic
LAN 2 (ix1):	Static	192.168.0.1	255.255.255.0

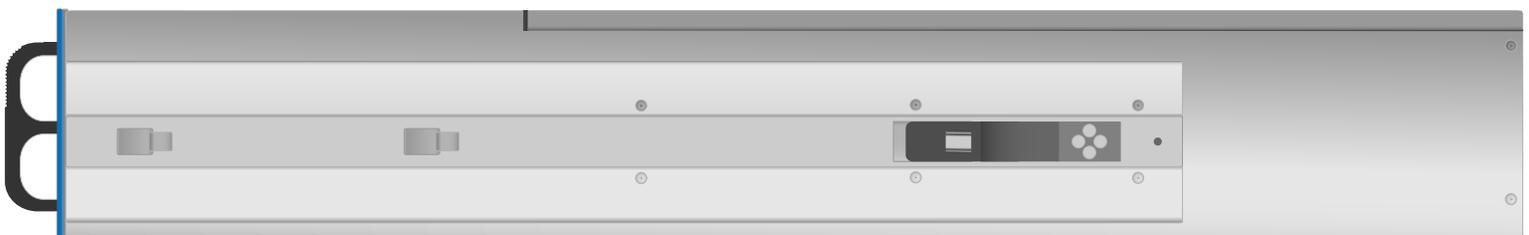
\*IPMI address can be found in Jupiter Callisto's management interface in the Network section. See section 7.1 for more details.

## 1.6 SIDE VIEW

Side view with rail attached



Side view with rail removed

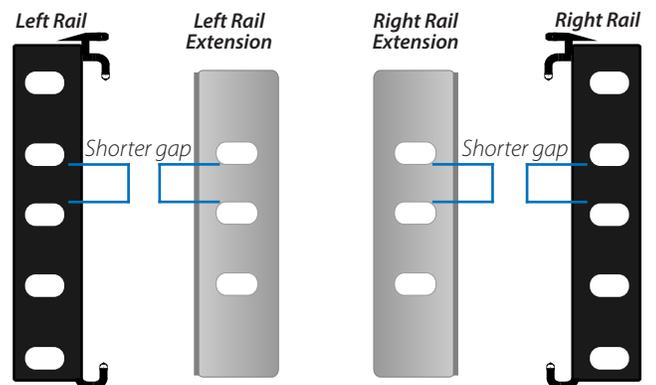
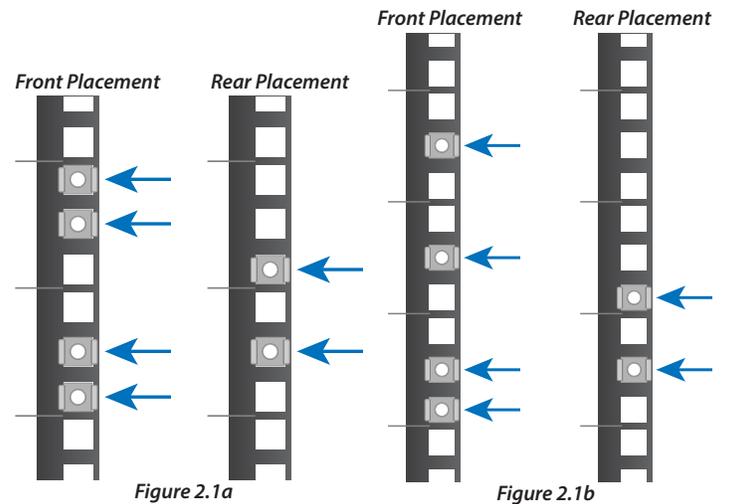


# ASSEMBLY & REPLACEMENT

## 2.1 RAIL INSTALLATION

The OWC Digital Jupiter Callisto comes with rack mount slide rails attached to the sides of the unit and the necessary screws and cage nuts for installing into 4-post square-hole rack systems. To install the slide rail brackets into your rack, do the following:

1. The Jupiter Callisto requires two (8-bay) or three (16-bay) rack units of space. Take note of which units in your rack that you will install your Jupiter Callisto into.
2. Install the included cage nuts into your rack. Use **Figure 2.1a** as a guide for cage nut placement for an 8-bay Jupiter Callisto; use **Figure 2.1b** for a 16-bay Callisto.
3. On the Jupiter Callisto, remove the securing screw from the rail on the left and right sides as seen in **Figure 2.1c**.
4. Once the securing screws have been removed, remove the rail arm by sliding it toward the rear of the unit as seen in **Figure 2.1d**. When the rail arm stops, push the release tab while continuing to slide the rail arm off. After removing the rails, set them aside taking note of the side from which they were removed.
5. Remove the rail arm extensions from the accessories box and match them to the rails. Use **Figure 2.1e** as a guide for proper matching of rails and the extension pieces.



6. Install the front rails, first making sure the bottom of the rail lines up with the bottom of the designated rack units. Using **Figure 2.1f** as a guide for an 8-bay and **Figure 2.1g** for a 16-bay unit, secure the rail using two of the included M5 pan head screws.

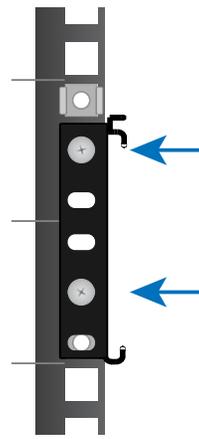


Figure 2.1f  
(8-bay)

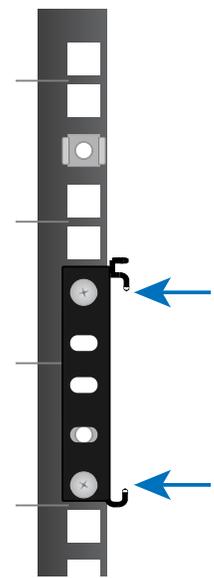


Figure 2.1g  
(16-bay)

7. Before securing the rear rail extension to the back of your rack, make sure to slide the extension into the rail as seen in **Figure 2.1h**.

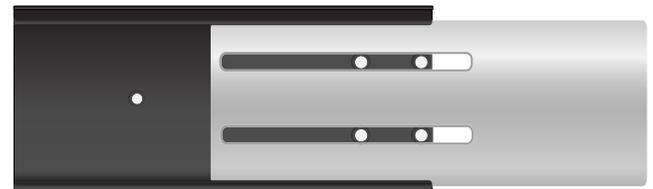


Figure 2.1h

8. Secure the rear rail extension to your rack use two of the included M5 pan head screws as seen in **Figure 2.1i**.
9. After securing the rear rail extension to your rack, secure the rail extension to the rail using two of the included M5 pan head screws as shown in **Figure 2.1j**.

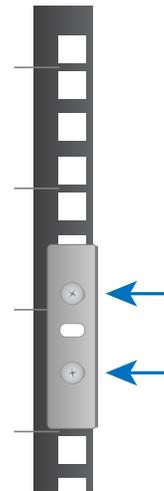


Figure 2.1i

10. After repeating the steps for the other side, you can now move on to installing the Jupiter storage system into the rack as found in **Section 2.2, Unit Rack Installation and Removal** on the next page.

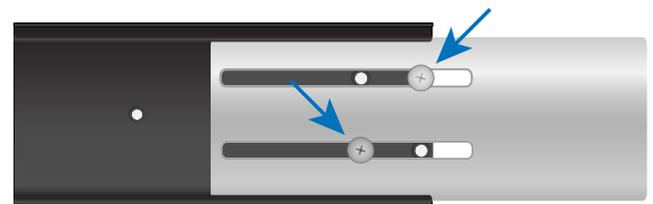


Figure 2.1j

## 2.2 UNIT RACK INSTALLATION & REMOVAL



### AVOID INJURY!

- Never lift heavy equipment by yourself
- Use proper lifting techniques



It is recommended that you remove all drives from Jupiter before installing into or removing from a rack.

After installing the rails, the Jupiter Callisto can be installed into the rack. The Jupiter Callisto is heavy even when no drives are installed. Never lift heavy equipment by yourself and always use proper lifting techniques. It is recommended that you remove all drives from the Jupiter Callisto before installing into or removing from a rack in order to avoid personal injury or damage to equipment.

### 2.2.1 Unit Installation



**Review precautionary measures from beginning of Section 2.2 before proceeding.**

1. Before lifting and installing the Jupiter Callisto into the rack, make sure all people involved in lifting know where the unit rails are attached to the Jupiter Callisto and how they will slide into the rack rails.
2. Using proper lifting techniques and keeping the Jupiter Callisto level, lift up to the rack rails where the unit will be installed into.
3. Visually align the unit rails with the rack rails.
4. Once aligned, evenly slide the unit rails into the rack rails. Do not release the unit until it has been inserted all the way or until an audible click is heard. The click indicates that the Jupiter Callisto has been inserted past the security latches as illustrated in **Figure 2.2.1a**.
5. If not already done, insert the Jupiter Callisto all the way into the rack until the rack ears touch the rack posts.
6. Install the four included rack securing screws to finish the unit installation into the rack as seen in **Figure 2.2.1b**. The 16-bay version is shown, but the screw placement is similar on the 8-bay version.

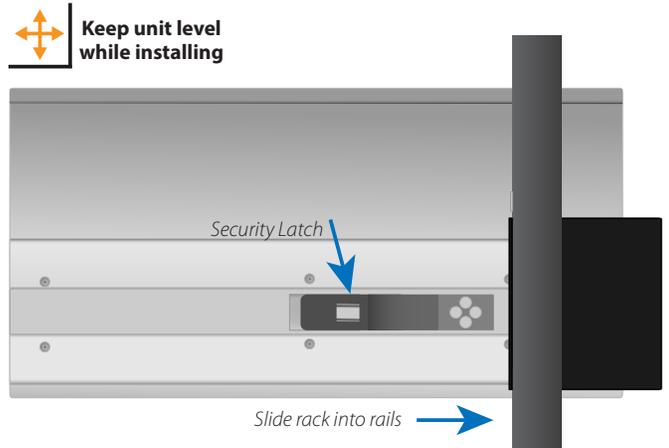


Figure 2.2.1a

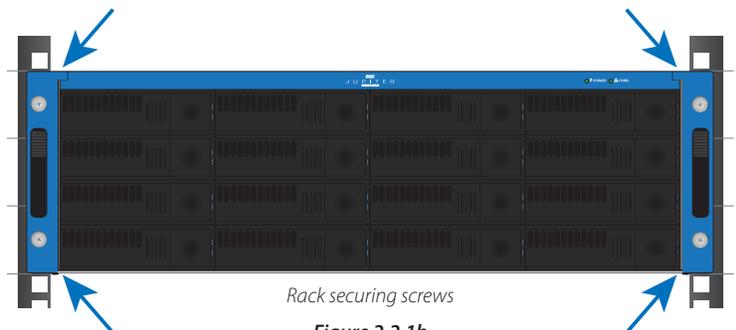


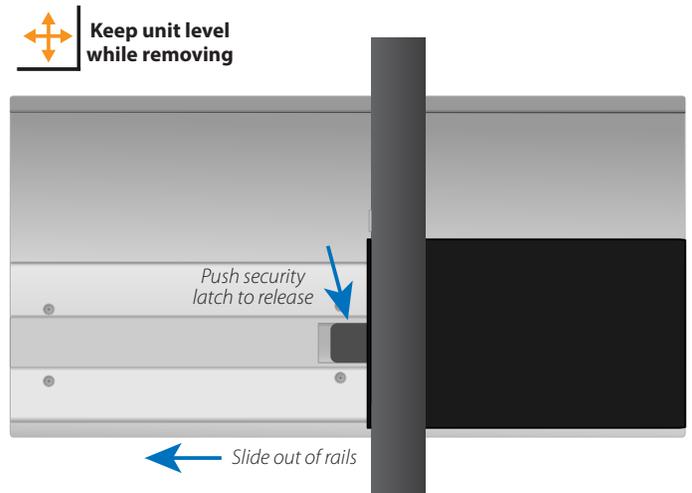
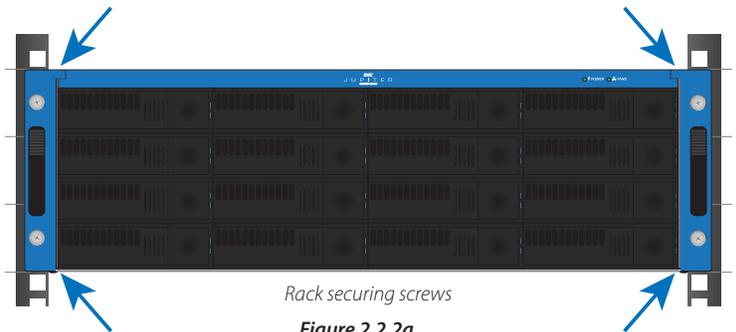
Figure 2.2.1b

## 2.2.2 Unit Removal



**Review precautionary measures from beginning of Section 2.2 before proceeding.**

1. Before removing the Jupiter Callisto from your rack, it is recommended that you remove all drives to make it easier to remove the Jupiter Callisto from the rack and to reduce the risk of personal injury and/or equipment damage.
2. Make sure all cables including data and power are disconnected from the Jupiter Callisto.
3. Remove the four rack securing screws from the front of the Jupiter Callisto as shown in **Figure 2.2.2a**.
4. Using the handles, pull the rack toward you until it is stopped by the security latches.
5. It is recommended to have a second person assist the removal at this point. Press the security latches while simultaneously sliding the unit out just past the security latch as illustrated in **Figure 2.2.2b**.
6. Make sure both people removing the unit have a firm grip on the rack front and base and are prepared to support the full width of the unit. Continue slowly pulling the rack out until it is free from the rails.



**Figure 2.2.2b**

## 2.3 DRIVE INSTALLATION & REMOVAL

The Jupiter Callisto comes with drives already installed into the individual trays to provide quicker setup time. Use the following instructions to install and remove drives from the drive bays.



### 2.3.1 Drive Installation

1. Release the drive tray handle by pushing the square eject button.
2. Insert the drive into the corresponding bay of the Jupiter Callisto, making sure the drive rests on the side guide rails within the bay. Note that each drive tray is labeled with a drive bay number. See **Figure 2.3.1a** for the drive bay numbering pattern..
3. Push the drive in using the front frame of the tray or the release button, as seen in **Figure 2.3.1b**, until it is seated into the backplane connector. DO NOT push the drive in using the handle.
4. Once the drive is seated into the backplane, push the handle in until it latches to secure the drive into the drive bay as shown in **Figure 2.3.1c**.

Figure 2.3.1a

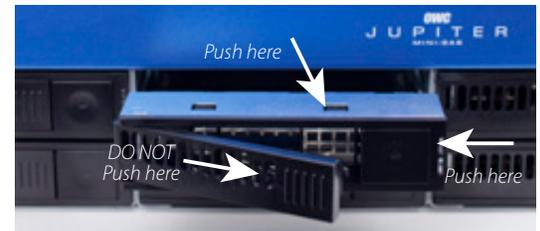


Figure 2.3.1b

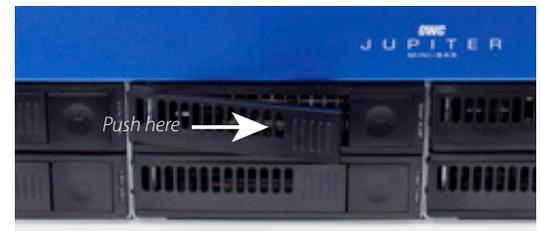


Figure 2.3.1c

### 2.3.2 Drive Removal

1. Release the handle on the drive tray by pushing the square eject button as seen in **Figure 2.3.2a**.
2. Use the handle to pull the drive out of the drive bay as shown in **Figure 2.3.2b**.



Figure 2.3.2a

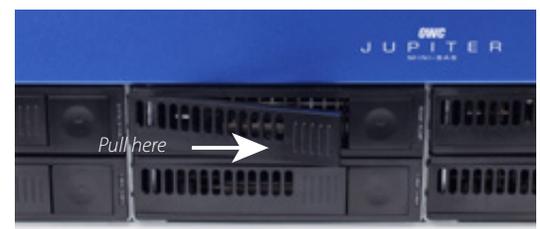


Figure 2.3.2b

## 2.4 POWER SUPPLY FAILURE & REPLACEMENT

The Jupiter Callisto has redundant, hot-swappable power supplies (two in the 8-bay, three in the 16-bay) and can withstand a failure of one power supply module. While all of the power supply modules are healthy, the power load is automatically balanced across them. In the event of a power supply failure, an alarm will sound and the 'Power' LED on the front of the unit will turn solid red. To replace a failed power supply module, do the following:

1. Identify the failed PSU module by seeing which module does not have its power LED lit.
- 2a. On the 8-bay Jupiter Callisto, each power supply module is secured with one captive thumbscrew (see **Figure 2.4a**). Loosen this thumbscrew until it spins freely. If the screw is too tight to loosen with your fingers, use a flathead screwdriver.
- 2b. On the 16-bay unit the three power supply modules are held in by a bar secured by one Phillips screw at the bottom (see **Figure 2.4b**). Remove this screw with a Phillips screwdriver, then remove the securing bar.
3. Pull the power supply module out using the black handle.
4. After removing the failed power supply module, insert the new module, making sure it is fully seated.
5. Re-install the screw. On the 8-bay, simply tighten the thumbscrew with your fingers, being careful not to over-tighten. On the 16-bay, replace the securing bar and use a Phillips screwdriver to re-install the screw.
6. After securing the new power supply module in place, plug in the power cord and turn on the Jupiter Callisto if it was powered off during the power supply replacement.

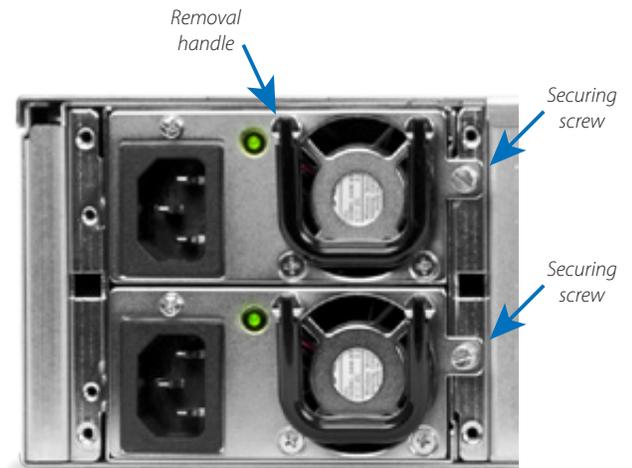


Figure 2.4a

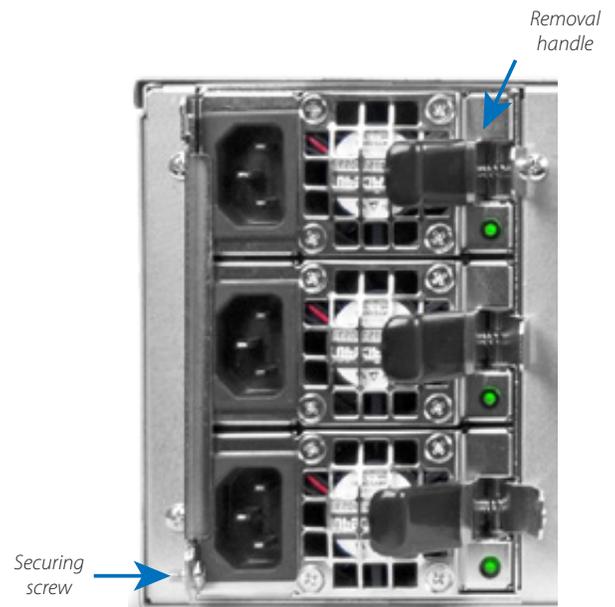


Figure 2.4b

## 2.5 FAN FAILURE

When one of the fans inside the Jupiter Callisto fails, the 'Fans' LED on the front of the unit will flash red and the remaining fans will ramp up to full speed to compensate. Contact OWC support for a replacement fan and for specific instructions on the replacement process.

# INITIAL SETUP

## 3.1 CONNECTION

Once the Jupiter Callisto is installed in the rack, the next step is to connect it to the network. Refer to **Section 1.5, Rear View** for the connectors on the Jupiter Callisto.

1. Connect the Jupiter Callisto to power. Make sure to connect all of the power supplies (two modules for the 8-bay, three for the 16-bay).
2. Connect the Jupiter Callisto to your network using the two 10GBASE-T RJ45 ports. These ports are compatible with 10Gb Ethernet as well as Gigabit Ethernet networks. These ports have been preset to the following configurations:

Port	Mode	IP Address	Subnet Mask
LAN 1 (ix0):	DHCP	automatic	automatic
LAN 2 (ix1):	Static	192.168.0.1	255.255.255.0

3. Turn on the Jupiter Callisto using the power switch.
4. The Jupiter Callisto is configured to use DHCP on LAN 1, so it will receive an IP address automatically from a network that is serving addresses via DHCP. The Jupiter Callisto is configured with a static IP address on LAN 2 to easily access the management interface in the event a network with DHCP is not available.
5. To access the Jupiter Callisto and manage it, open a web browser on any client computer connected to the network and enter the Jupiter Callisto's IP address. Instead of the IP address, you can also enter 'jupiter.local' in the web browser.

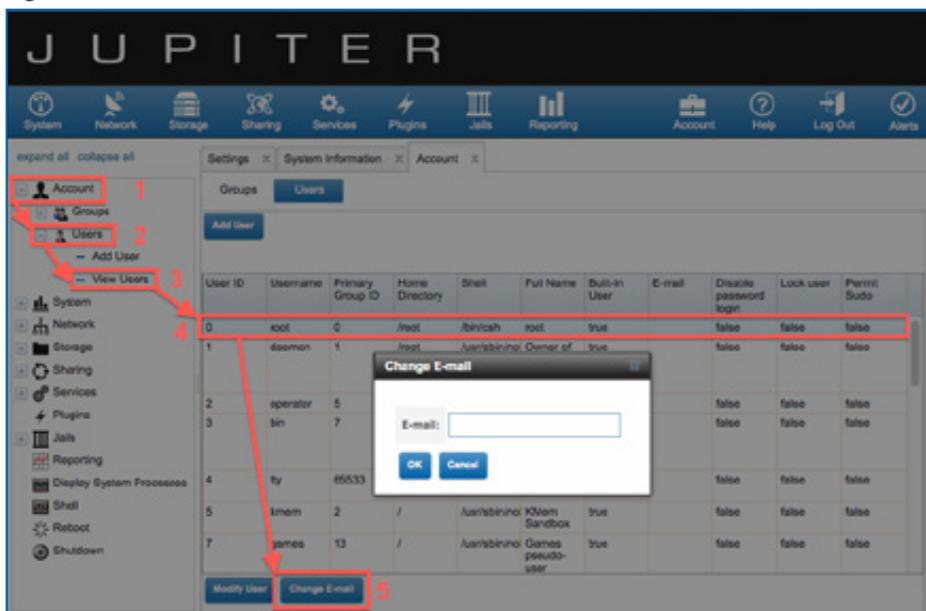
## 3.2 FIRST TIME CONFIGURATION

The first time that any user connects to the Jupiter Callisto through the web interface, they will be prompted to set a new password for the 'root' user, which is the user that has full administrator access to the Jupiter Callisto.

The next step is to set a password for the 'root' user in order to receive email alerts from the system. When the Jupiter Callisto detects an error, it will issue an alert by sending an email to the 'root' user. Additionally, any active alerts can be viewed in the upper right corner of the web interface.

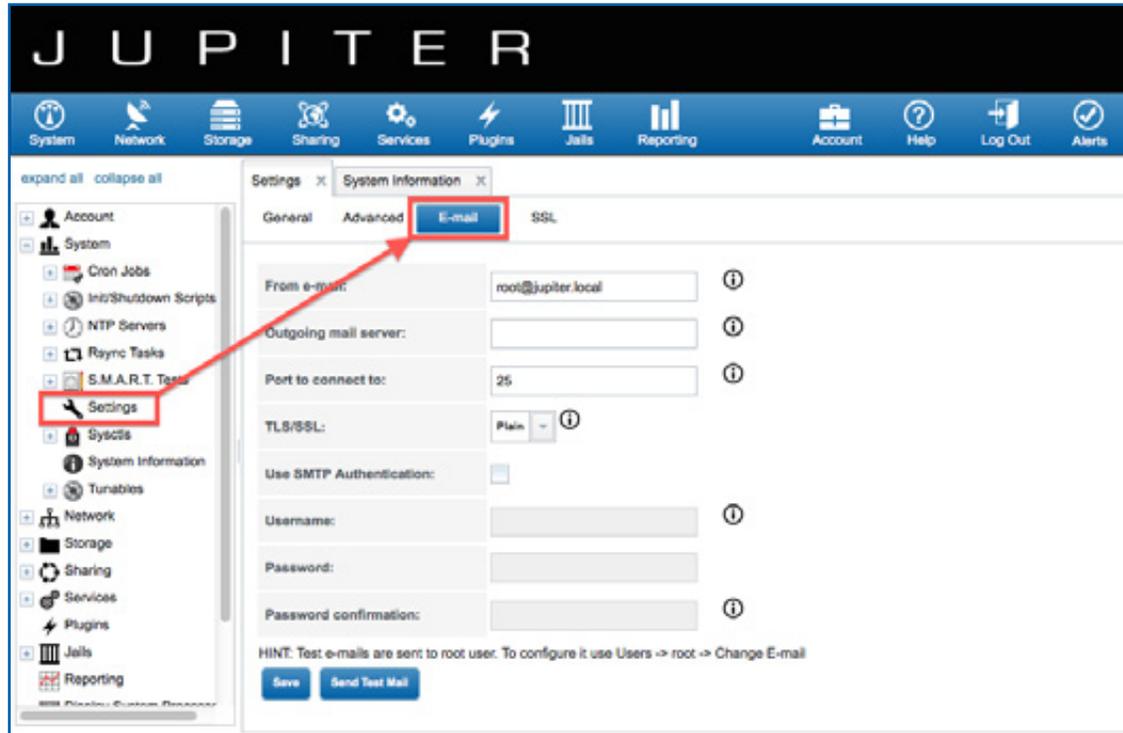
There are two steps to setting up email alerts. First, expand the 'Account' section in the sidebar, then expand 'Users' and select 'View Users'. Next, click on the 'root' user to highlight it, then click on the 'Change E-mail' button at the bottom of the screen. In the window that appears, enter the email address that you wish to receive alerts, then click 'OK' to confirm the change. See **Figure 3.2a**, below.

Figure 3.2a



Once the 'root' user email address has been set, the next step is to configure the email settings that the Jupiter Callisto will use when sending email alerts. To configure email settings, first click on 'Settings' in the sidebar, then click on the 'E-mail' tab (see **Figure 3.2b**). Configure the email settings as appropriate for your specific network.

**Figure 3.2b**



### 3.3 ACCESSING THE IPMI INTERFACE

The motherboard inside Jupiter Callisto has components that allow for management of the system independently of the power status and OS installed. The IPMI (Intelligent Platform Management Interface) provides access to the motherboard's BMC (Baseboard Management Controller), which monitors and manages system sensors, provides remote viewing KVM services, and allows for power control of the system.

To access the IPMI interface:

1. Connect to your network using the IPMI port on Jupiter Callisto (See **Section 1.5**). The IPMI port is set to DHCP in order to automatically receive a IP address from the network. These instructions assume that the network being connected to the IPMI port is serving IP addresses via DHCP.
2. To determine the IPMI address, click on the 'Network' tab at the top of the screen in the web interface, then click 'IPMI'.
3. Enter the address into a web browser to get to the IPMI interface login screen.
4. The default username is 'ADMIN' and password is 'ADMIN' (capital letters for both).

### 3.4 DEFAULT SETUP

The Jupiter Callisto system is set up by default in the following configuration.

- **Storage:** The storage pool is configured as multiple RAIDZ2 sets of four drives each. 8-bay units consist of two RAIDZ2 sets, while 16-bay units have four RAIDZ2 sets. One ZFS filesystem (labeled 'callisto') has been created from this pool (labeled 'pool'). See **Section 4, Storage Configuration**, for more information.
- **Accounts:** One user has been created, named 'callisto' with a password of 'callisto'. It belongs to the group also named 'callisto'.
- **Shares:** One CIFS (SMB) share named 'Jupiter Callisto' has been created. This share has given read/write access to the user 'callisto'. You can access the share by browsing network drives on a client computer. See **Section 6.2, Connecting to Shares**, for alternate methods to access a CIFS share.

# STORAGE CONFIGURATION

## 4.1 ZFS TERMINOLOGY

In order to use Jupiter Callisto most effectively, it is important to have an understanding of the following terms and concepts relating to the ZFS file system and volume manager.

**Pool** — Also known as 'ZFS pool' or 'zpool'. A collection (or 'pool') of all of the virtual devices (vdevs) that have been configured for use. All of the vdevs inside a pool are automatically striped together. The pool is the largest logical construction in the ZFS environment. Other ZFS components (filesystems, zvols, snapshots, and clones) can be created inside the pool. These components are called 'datasets' and they are the closest things that ZFS has to what would be a 'volume' or a 'partition' in a traditional volume manager.

**Vdev** — A vdev, or virtual device, is a building block of the pool. There are two basic types of vdev: low-level and high-level. A low-level vdev can be an entire hard drive or a partition on a hard drive. A high-level vdev is made up of multiple low-level vdevs, and can be thought of as a RAID. In ZFS, a high-level vdev can be a stripe, a mirror, or one of three levels of 'RAID-Z'.

**Dataset** — This is a blanket term for any ZFS component created out of a pool. These are the devices that will be used to store data and that can be manipulated by users. The types of datasets are: filesystems, zvols, snapshots, and clones.

**Filesystem** — This term can be confusing, because it has a slightly different meaning in ZFS than you might expect. In ZFS, a filesystem is a dataset that provides file-level access to data. It is similar to a folder, but with added capabilities. A filesystem is required in order to create a NAS.

**Zvol** — A ZFS volume (or zvol) is a dataset that provides block-level access to data. It behaves essentially like a hard drive as far as an end user is concerned, and is required in order to set up a SAN.

**Snapshot** — A read-only copy of a filesystem or zvol. Because of the copy-on-write transactional model that ZFS uses, a snapshot can be created very quickly and takes up no additional space initially, because instead of a new copy of all of the information stored on the source, a snapshot contains pointers to the original data.

**Clone** — A writable copy of a filesystem or zvol. A clone can only be created from a snapshot, and like a snapshot it can be created very quickly and initially takes up no additional space. As changes are made to the clone, the storage space that it takes up will grow.

## 4.2 STORAGE CONFIGURATION

Jupiter Callisto uses ZFS to create and manage its storage. Your Jupiter Callisto has arrived with its storage pre-configured, but you may wish to change the storage configuration to suit your specific needs.

### 4.2.1 Creating the Pool

To be able to use your Jupiter Callisto as a storage server, you must first create a pool. In simple terms, a pool is a collection of all of the drives and RAID sets attached to the Jupiter Callisto. To create a pool, go to the 'Storage' tab, then click on the 'Add/Extend ZFS Pool' button to bring up the following window (**Figure 4.2.1**).

**Figure 4.2.1**

The screenshot shows the 'Add/Extend ZFS Pool' window. It includes a 'Pool Name' input field, a 'Pool to extend' dropdown, an 'Encryption' checkbox, and a section for 'Available disks' showing '1 - 3.0 TB (16 drives, show)'. The 'Pool layout' section features a grid of 15 numbered blocks, a circle icon for resizing, and a tooltip that says 'Drag and drop this to resize'. Below the grid, it displays '0x1x0 B' and 'Capacity: 0 B'. At the bottom, there are buttons for 'Add Pool' (with a warning 'Existing data will be cleared'), 'Cancel', and 'Manual setup'.

Create a name for your pool, then move to the 'pool layout' section below that. Notice that all of the disks not currently members of a RAID are listed as 'available disks' and are grouped together by capacity. If you click on the '+' icon next to the list of available disks of one capacity, those disks will be added to the 'pool layout' section with an automatically selected RAID configuration. This configuration is simply a default and can be changed easily. Click on the dropdown menu on the left to modify the RAID level.

Each row in the chart to the right represents one RAID set. Each numbered block represents one disk. In order to create more than one RAID set out of your group of disks, simply click on the circle icon at the lower right hand of the chart and drag it down to create a new row. Dragging the circle icon from left to right will increase the number of disks assigned to each RAID set. If you create multiple RAID sets, you can modify the RAID level of each by clicking the dropdown menu to the left of each row (it is recommended to use the same RAID level for each RAID set in the pool). Each RAID set of the same RAID level will be striped together automatically when the pool is created. It is recommended to create the pool using multiple identical RAID sets rather than one large RAID. This will increase performance and reduce recovery time after a disk failure. The total usable capacity of the pool is displayed at the top of the 'pool layout' section.

### 4.2.2 ZFS RAID Levels

ZFS has its own version of RAID, including some levels that are roughly equivalent to standard RAID levels and some that are unique to ZFS.

- **Mirror** — This RAID level is equivalent to RAID 1, except that you can choose more than two disks to be members of the

RAID. Each member disk beyond the first will contain an exact copy of the information stored on the first disk.

- Stripe** — This is equivalent to RAID 0.
- RAIDZ1** — This is equivalent to RAID 5. It is a stripe with one parity disk.
- RAIDZ2** — This is equivalent to RAID 6. It is a stripe with two parity disks.
- RAIDZ3** — This is a stripe with three parity disks.

### 4.2.3 Optimal RAIDZ Configurations

Due to the interaction between the 128K stripe size of RAIDZ and the 4K block sizes of modern high-capacity hard drives, there is an optimal number of disks that each RAIDZ level (RAIDZ, RAIDZ2, RAIDZ3) should have for best performance. When creating a new RAIDZ set, use the following formula to determine the optimal number of member disks:

$$2^n + p$$

Where  $n$  is any number beginning at 1, and  $p$  is the number of parity disks in your RAIDZ level (RAIDZ=1, RAIDZ2=2, RAIDZ3=3). Based on this formula, some of the optimal numbers of disks for each RAIDZ level are:

- RAIDZ1** — 3, 5, 9
- RAIDZ2** — 4, 6, 10
- RAIDZ3** — 5, 7, 11

Additionally, for optimal performance it is recommended to create a pool using more than one RAIDZ set. Write performance especially will be negatively affected if only one RAIDZ set is used.

### 4.2.4 Creating Datasets

After the pool is created, the next step is to create devices from the pool. Think of the pool as one enormous hard disk that you will be partitioning into multiple volumes. ZFS calls these volumes 'datasets', and there are two basic types of datasets: filesystems and zvols. See **Section 4.1, ZFS Terminology** for more information on the types of datasets. If you want to set up a NAS or use iSCSI to provide file-level data storage, you will create one or more filesystems. If you want to use iSCSI to provide block-level data storage, you will create one or more zvols.

### 4.2.5 Creating a Filesystem

Filesystems are the most common datasets in the ZFS storage architecture. You can create a filesystem directly from the pool, and you can also nest filesystems inside each other. To create a filesystem, click on the pool or the filesystem that you want to use as the parent, then at the bottom of the screen click on the 'Create ZFS Filesystem' button to bring up the following window (**Figure 4.2.5**).

Figure 4.2.5

Create ZFS Filesystem

Create ZFS filesystem in pool

Filesystem Name:

Compression level:

Share type:

Enable atime:  Inherit  On  Off

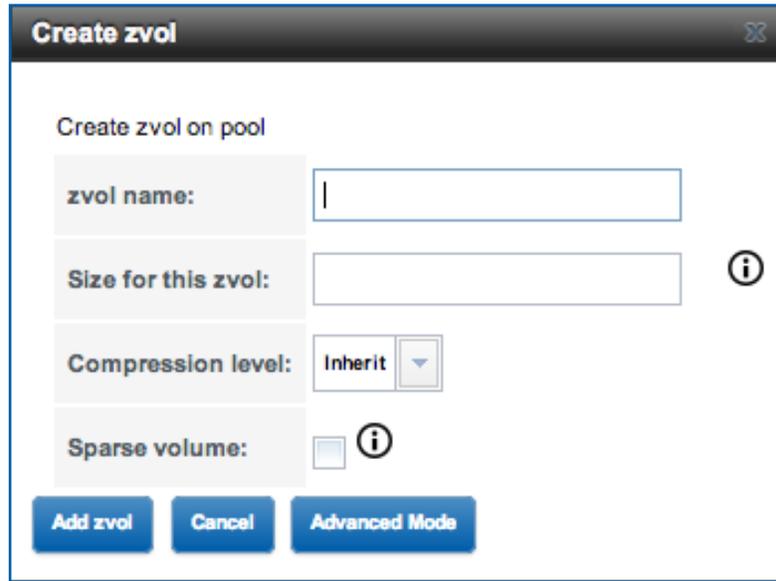
ZFS Deduplication:   
Enabling dedup may have drastic performance implications, as well as impact your ability to access your data. Consider using compression instead.

Give your filesystem a name, then create it by clicking the 'Add Filesystem' button. There are advanced settings that you can modify, including compression, deduplication, quotas, and reserved space, but if you are unfamiliar with these options, it is best to leave them at their defaults.

## 4.2.6 Creating a Zvol

Zvols can be created directly from the pool or from a filesystem. To create a zvol, select the pool or the filesystem that you want to use as the parent, then click the "Create zvol" button at the bottom of the screen to bring up the following window. (Figure 4.2.6)

Figure 4.2.6



zvol name:

Size for this zvol:  ⓘ

Compression level: Inherit ▾

Sparse volume:  ⓘ

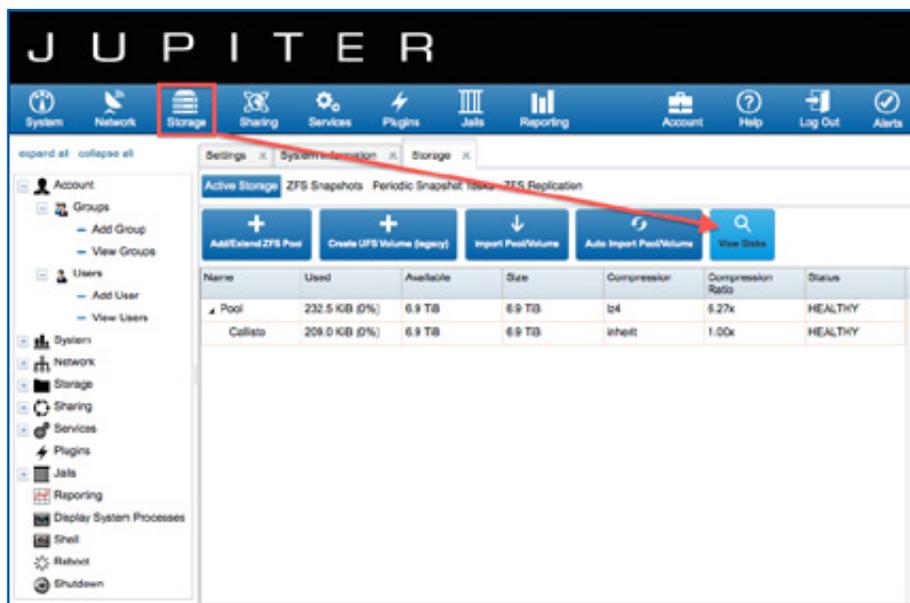
Add zvol Cancel Advanced Mode

Give the zvol a name, then click on 'Add zvol' to create the zvol. Note that unlike a filesystem, a zvol does need to have a specific size.

## 4.3 DRIVE IDENTIFICATION

In order to view the drives installed in the Callisto, click on the 'View Disks' button inside the 'Storage' tab (see **Figure 4.3a**).

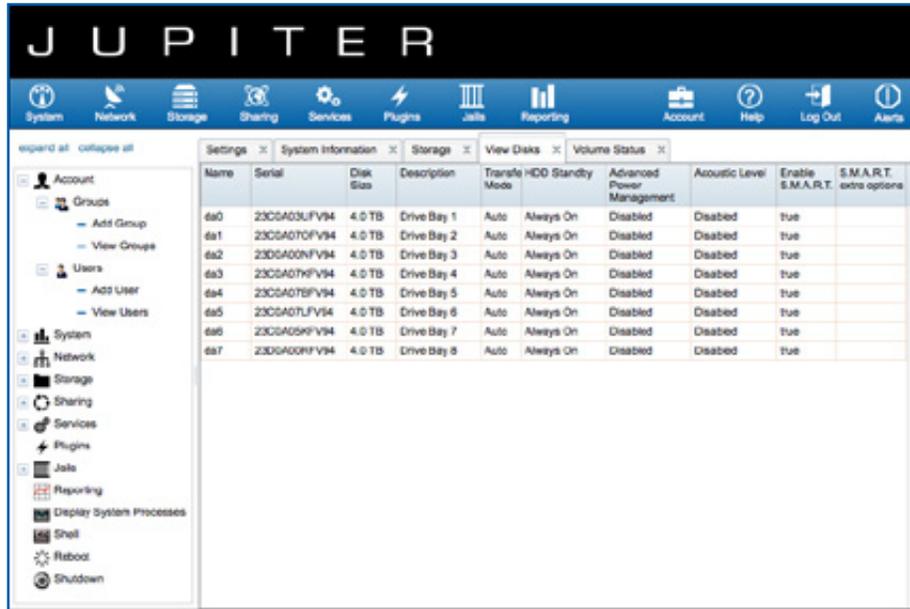
Figure 4.3a



Name	Used	Available	Size	Compression	Compression Ratio	Status
Pool	232.5 KB (0%)	6.9 TB	6.9 TB	lz4	6.27x	HEALTHY
Callisto	209.0 KB (0%)	6.9 TB	6.9 TB	inherit	1.00x	HEALTHY

Inside 'View Disks' section, the drives are identified in a number of ways (see **Figure 4.3b**). First, each drive is assigned a device ID by the Callisto operating system. This device ID is arbitrary and not based on the drive's physical location inside the Callisto. Next, each drive has a serial number. The serial number is unique to each drive and is the same as the serial number written on the drive's label. Third, the 'Description' field for each drive has been labeled with a drive bay number. This label follows the drive; the label does not stay with the drive bay. In the event that you move a drive into another drive bay or add a new drive, the label will need to be changed manually.

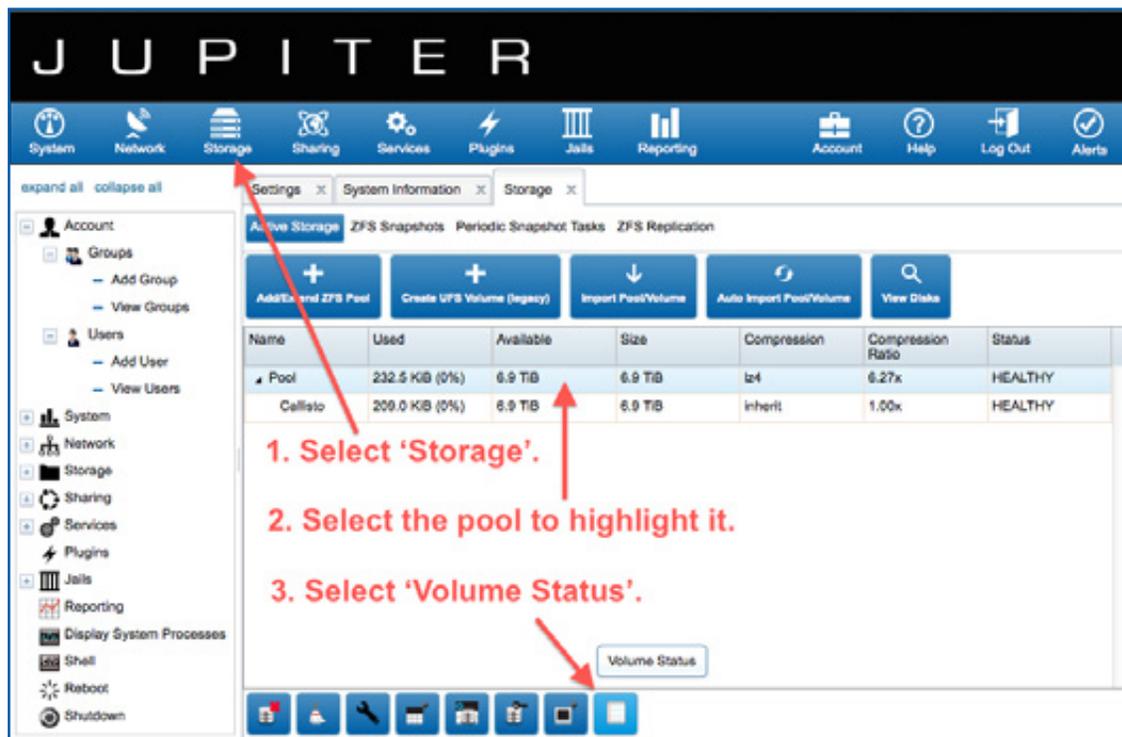
**Figure 4.3b**



#### 4.4 DRIVE FAILURE & REPLACEMENT

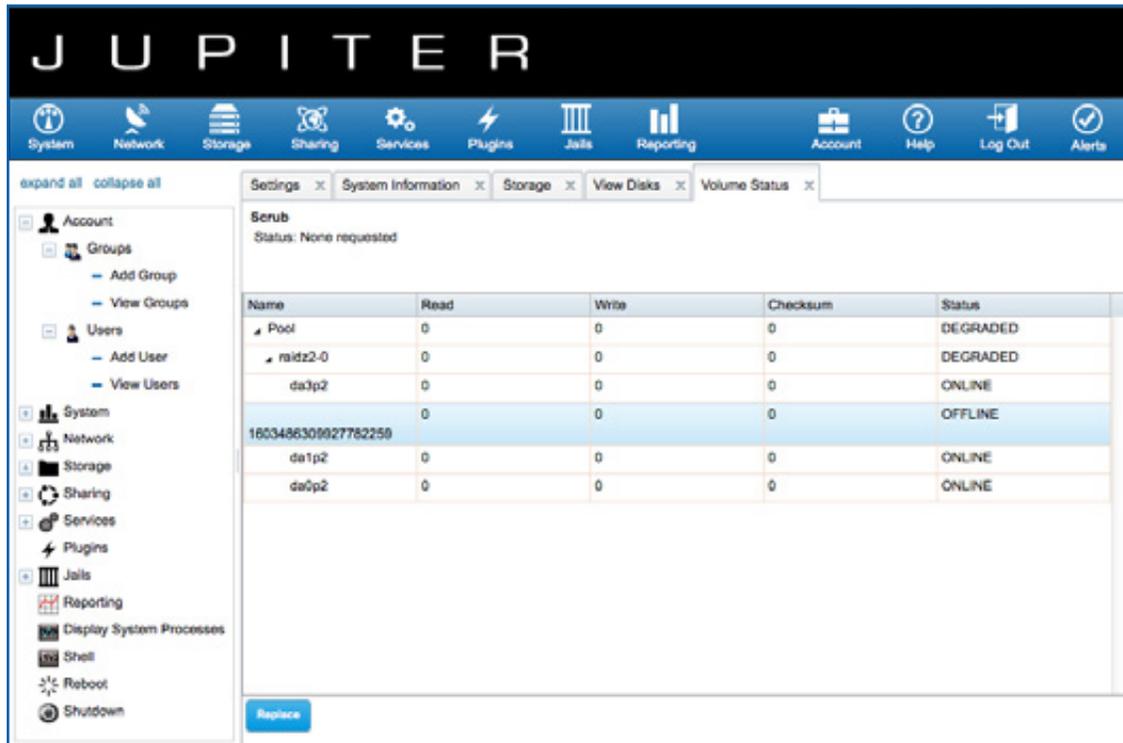
In the event of a drive failure, the Jupiter Callisto will send an alert email to the administrator (see **Section 3.2, First Time Configuration** for information on setting the administrator email address). In order to identify the failed drive, first go to 'Volume Status' in the 'Storage' menu (**Figure 4.4a**).

**Figure 4.4a**



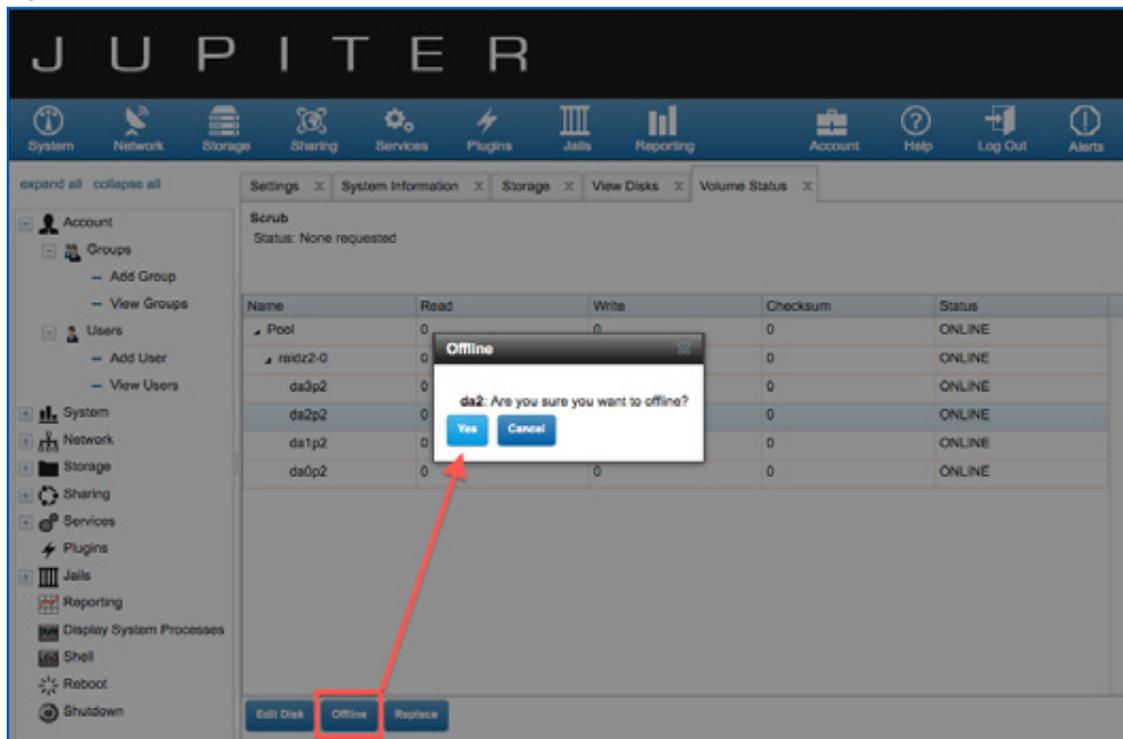
The failed drive will be marked in the 'Volume Status' screen, and the RAID set and pool of which the disk is a member will be labeled as 'DEGRADED' (**Figure 4.4b**).

**Figure 4.4b**



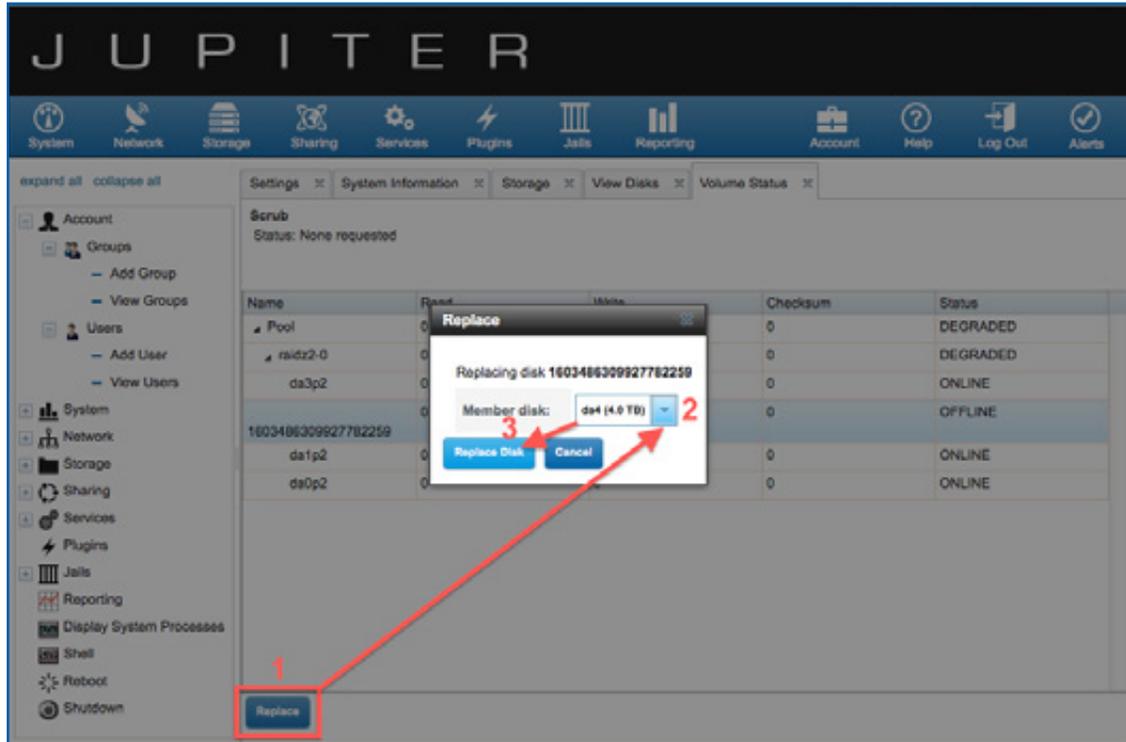
To replace the failed disk, first click on the drive to highlight it. In **Figure 4.4b** above, the failed drive is already listed as 'OFFLINE'. In some instances the disk will be listed as 'FAULTED' and will not yet be offline. In that case, you will need to click the 'Offline' button to take the disk offline (**Figure 4.4c**).

**Figure 4.4c**



Once the disk is offline, click the 'Replace' button at the bottom of the screen to bring up the replacement menu (**Figure 4.4d**). From the dropdown menu, select the disk that you would like to use as a replacement for the failed disk, then click 'Replace Disk'. For information on physical drive replacement, see **Section 2.3, Drive Installation & Removal**.

Figure 4.4d



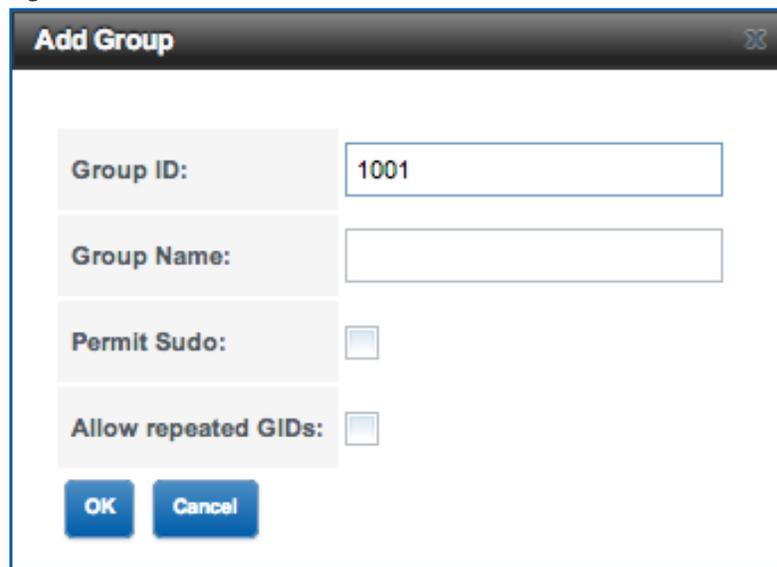
# USERS, GROUPS, & PERMISSIONS

After you have configured your storage, the next step is to set up users, groups, and permissions. You can give individual users or groups different access levels to different filesystems and zvols, or you can allow wide open access to everyone. It depends on how you prefer to administer your network. See **Section 3.3, Default Setup** for information on the pre-configured users, groups, and permissions.

## 5.1 CREATING A GROUP

To create a group, click on 'Account' in the sidebar to expand the category, then click on 'Groups' to expand it. After that, click on 'Add Group' to open the following window (**Figure 5.1**).

Figure 5.1



The screenshot shows a dialog box titled "Add Group". It has a dark header bar with the title and a close button. Below the header, there are four rows of input fields. The first row is "Group ID:" with a text box containing "1001". The second row is "Group Name:" with an empty text box. The third row is "Permit Sudo:" with an unchecked checkbox. The fourth row is "Allow repeated GIDs:" with an unchecked checkbox. At the bottom of the dialog, there are two buttons: "OK" and "Cancel".

Give the new group a name, then click 'OK' to create the group.

## 5.2 CREATING A USER

To create a user, click on 'Account' in the sidebar to expand the category, then click on 'Users' to expand it. After that, click on 'Add User' to open the following window (**Figure 5.2**).

**Figure 5.2**

The 'Add User' dialog box contains the following fields and options:

- User ID: 1001
- Username: [empty]
- Create a new primary group for the user:
- Primary Group: [dropdown]
- Home Directory: /nonexistent [Browse]
- Shell: ssh
- Full Name: [empty]
- E-mail: [empty]
- Password: [empty]
- Password confirmation: [empty]
- Disable password login:
- Lock user:
- Permit Sudo:
- SSH Public Key: [empty]
- Auxiliary groups: Available (wheel, daemon, kmem, sys) Selected

Enter a username. If the new user is going to be a member of an existing group, uncheck 'Create a new primary group for the user', then select a group from the 'Primary Group' dropdown menu. If you leave 'Create a new primary group for the user' checked, a new group will be created with the same name as the new user.

Next, choose a home directory for the user. Click 'Browse', then navigate to the filesystem that you want to use as the user's home directory. You can also select the pool as the home directory. The only other required fields are 'Full Name' and 'Password'. If you click 'Advanced Mode' at the bottom, you can modify access permissions for the new user. Click 'OK' at the bottom to create the new user.

## 5.3 PERMISSIONS

Every pool, filesystem, and network share that you create has a set of permissions associated with it. These permissions tell the Jupiter Callisto which users and groups are allowed to read, write, and execute files. These permissions are represented by an ACL (Access Control List), which looks like this (**Figure 5.3**):

**Figure 5.3**

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

The 'Owner' column sets the permissions for any user that is set as the owner of the directory, pool, or filesystem. 'Group' sets the permissions for all users that are in any of the groups that have been given specific access. 'Other' sets the permissions for all other users attempting to use the storage. As shown above, the default settings in the ACL are to give the owner read/write/execute access, while the group and any other users have read and execute access, but not write.

# SHARES

## 6.1 CREATING SHARES AND CONFIGURING GUEST ACCESS

After configuring storage, users, groups, and permissions, the next step is to set up shares. There are three types of file sharing protocols supported by Jupiter Callisto: AFP, CIFS (SMB), and NFS. AFP sharing is a good idea if all client computers are Macs® and you would like to use the Jupiter Callisto as a Time Machine® backup; CIFS sharing is preferred if all client computers are using Windows® operating systems or Macs running OS X 10.9 and later (previous OS X revisions use an older version of SMB and do not perform as well), and NFS is useful if client computers are running on a mixture of different OS versions.

### 6.1.1 Creating an AFP Share

To create an AFP share, click on the 'Sharing' tab at the top of the screen, click 'Apple (AFP)', then click 'Add Apple (AFP) Share' to bring up the following window (**Figure 6.1.1**).

Figure 6.1.1

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

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

Give the new share a name, then assign it a path by clicking the 'Browse' button and selecting the pool or the filesystem that you would like to share. Next, add any users or groups that you wish to be able to access the share by entering their names in the 'Allow List'. Separate multiple entries with commas, and add any groups by adding the '@' symbol at the beginning (e.g. @groupname).

If you wish to modify the default permissions, do so at this time. If you intend to use this share as a Time Machine backup, make sure to click the 'Time Machine' checkbox. Then, click 'OK' at the bottom to create the AFP share.

After the new AFP share is created, you will receive a message asking if you wish to activate the AFP service. Click 'Yes' to enable AFP sharing. If you wish to enable or disable AFP sharing in the future, click on the 'Services' tab at the top of the screen, then click on the 'ON/OFF' slider labeled 'AFP' to toggle the service.

## 6.1.2 AFP Guest Access

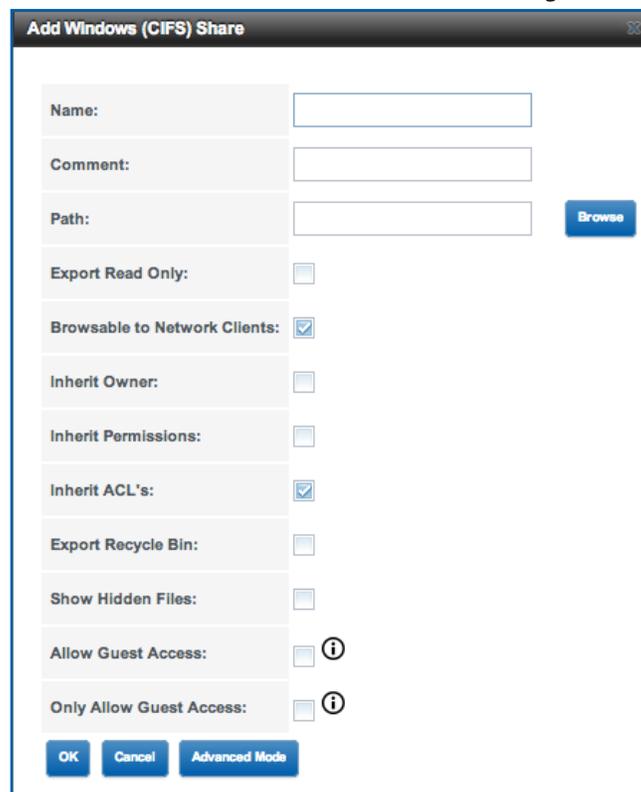
To enable guest access over AFP, there are three steps.

1. When creating the AFP share, add the user 'nobody' to the allow list. 'Nobody' is a user that is built into the system and is used for setting up guest access to shares.
2. Go to the 'Services' tab and click the wrench icon to the right of the 'ON/OFF' slider for the AFP service. This will bring up the AFP settings window. Click the 'Guest Access' checkbox, then set 'nobody' as the guest account. Click 'OK' to save these modifications.
3. Modify the permissions of the filesystem to which your AFP share is pointed. Guest access can only be configured on a filesystem because the ownership of the pool cannot be changed for security reasons. Click on the 'Storage' tab at the top of the screen, then click on the appropriate filesystem. At the bottom of the screen, click on the 'Change Permissions' icon (a stack of disks with a black key at the top), then in the window that opens select 'nobody' as 'Owner (user)' and as 'Owner (group)'. Click 'Change' to apply these modifications.

## 6.1.3 Creating a Windows (CIFS) Share

To create a CIFS share, click 'Sharing' at the top of the screen, click 'Windows (CIFS)', then click 'Add Windows (CIFS) Share' to bring up the following window (**Figure 6.1.3**).

**Figure 6.1.3**



Give the new share a name, then assign it a path by clicking the 'Browse' button and selecting the pool or the filesystem that you would like to share. Click 'OK' at the bottom to create the CIFS share.

After the new CIFS share is created, you will receive a message asking if you wish to activate the CIFS service. Click 'Yes' to enable CIFS sharing. If you wish to enable or disable CIFS sharing in the future, click on the 'Services' tab at the top of the screen, then click on the 'ON/OFF' slider labeled 'CIFS' to toggle the service.

## 6.1.4 CIFS Guest Access

To enable guest access over CIFS, there are three steps.

1. When creating the CIFS share, click the 'Allow Guest Access' checkbox.
2. Go to the 'Services' tab and click the wrench icon to the right of the 'ON/OFF' slider for the CIFS service. This will bring up the CIFS settings window. In the 'Guest account' section, select the user 'nobody' from the dropdown menu. 'Nobody' is a user that is built into the system and used for setting up guest access to shares.
3. Modify the permissions of the filesystem to which your CIFS share is pointed. Guest access can only be configured on a filesystem because the ownership of the pool cannot be changed for security reasons. Click on the 'Storage' tab at the top of the screen, then click on the appropriate filesystem. At the bottom of the screen, click on the 'Change Permissions' icon (a stack of disks with a black key at the top), then in the window that opens select 'nobody' as 'Owner (user)' and as 'Owner (group)'. Click 'Change' to apply these modifications.

## 6.1.5 Creating an NFS Share

To create an NFS share, click 'Sharing' at the top of the screen, click 'Unix (NFS)', then click 'Add Unix (NFS) Share' to bring up the following window (**Figure 6.1.5**).

**Figure 6.1.5**

The screenshot shows the 'Add Unix (NFS) Share' dialog box. It features a title bar with the text 'Add Unix (NFS) Share' and a close button. The main area contains several sections: 'Comment:' with a text input field; 'Authorized networks:' and 'Authorized IP addresses or hosts:' with text input fields and information icons; 'All Directories:', 'Read Only:', and 'Quiet:' with checkboxes and information icons; 'Maproot User:', 'Maproot Group:', 'Mapall User:', and 'Mapall Group:' with dropdown menus, all currently showing 'N/A'. At the bottom, there is a 'Path' section with a 'Path:' text input field and a 'Browse' button. Below that is an 'Add extra Path' section with 'OK' and 'Cancel' buttons.

Select a path for the NFS share and click 'OK' to create the share. Note that access to NFS shares is not user specific. You can limit access to certain IP addresses if you want, but permissions will be handled entirely on the side of the ZFS filesystem that you are pointing to in the path.

After the new NFS share is created, you will receive a message asking if you wish to activate the NFS service. Click 'Yes' to enable NFS sharing. If you wish to enable or disable NFS sharing in the future, click on the 'Services' tab at the top of the screen, then click on the 'ON/OFF' slider labeled 'NFS' to toggle the service.

## 6.1.6 NFS Guest Access

Because of the way that NFS shares operate, users do not log in with a specific username, so there is no specific 'guest access' option, but there is still a way to ensure that all users that access the NFS share have the same permissions. To do so, follow these steps:

1. When creating the NFS share, in the 'Mapall User' and 'Mapall Group' sections, select 'nobody'.
2. Modify the permissions of the filesystem to which your NFS share is pointed. Guest access can only be configured on a filesystem because the ownership of the pool cannot be changed for security reasons. Click on the 'Storage' tab at the top of the screen, then click on the appropriate filesystem. At the bottom of the screen, click on the 'Change Permissions' icon (a stack of disks with a black key at the top), then in the window that opens select 'nobody' as 'Owner (user)' and as 'Owner (group)'. Click 'Change' to apply these modifications.

## 6.2 CONNECTING TO SHARES

Once you have created the shares that you would like to use (AFP, NFS, or CIFS), the client computers connected to the network will need to access these shares. Depending on the operating system that each computer is running and the type of share, the connection process may vary.

### 6.2.1 Connecting to AFP Shares (Mac only)

Select 'Finder' either by clicking the Finder icon in your dock or by clicking on the desktop. Then, open up the 'Connect to Server' window by clicking the 'Go' menu in the menu bar at the top of the screen and then selecting 'Connect to Server'. You can also open this window using the keyboard shortcut command+K.

In the 'Connect to Server' window, enter 'afp://', followed by the IP address of the Jupiter Callisto. Click the 'Connect' button to connect to the AFP share. You will be prompted with a login screen, at which point you will either enter the appropriate user credentials or log in as a guest if you have set up guest access.

### 6.2.2 Connecting to CIFS Shares on a Mac

Select 'Finder' either by clicking the Finder icon in your dock or by clicking on the desktop. Then, open up the 'Connect to Server' window by clicking the 'Go' menu in the menu bar at the top of the screen and then selecting 'Connect to Server'. You can also open this window using the keyboard shortcut command+K.

In the 'Connect to Server' window, enter 'smb://', followed by the IP address of the Jupiter Callisto. Click the 'Connect' button to connect to the CIFS share. You will be prompted with a login screen, at which point you will either enter the appropriate user credentials or log in as a guest if you have set up guest access.

### 6.2.3 Connecting to CIFS Shares in Windows

Open an Explorer window, then right-click on 'Computer' and select 'Map network drive...'. In the next window, select a drive letter to assign to the share, then in the 'Folder' section, follow the formula: \\[IP address of Callisto]\[name of your CIFS share]

For example, if the IP address of your Callisto is 192.168.0.115 and you created a CIFS share named 'callisto', then you would enter the following: \\192.168.0.115\callisto

### 6.2.4 Connecting to an NFS share on a Mac

Select 'Finder' either by clicking the Finder icon in your dock or by clicking on the desktop. Then, open up the 'Connect to Server' window by clicking the 'Go' menu in the menu bar at the top of the screen and then selecting 'Connect to Server'. You can also open this window using the keyboard shortcut command+K.

In the 'Connect to Server' window, enter 'nfs://', followed by the complete file path of the nfs share on the Jupiter Callisto, including the IP address. For example, if the IP address of your Callisto is 192.168.0.115 and the path to your shared filesystem is /mnt/pool/callisto, then the complete address that you will enter is:

```
nfs://192.168.0.115/mnt/pool/callisto
```

Click the 'Connect' button to connect to the NFS share. The share will mount as a network folder.

## 6.2.5 Connecting to an NFS share in Windows

Depending on your Windows version, you may or may not be able to mount an NFS share without downloading a separate application to do so. These instructions are for mounting an NFS share in Windows 7 and in Windows 8 Enterprise.

First, you will need to make sure that support for mounting NFS shares is turned on. Go to 'Computer', then 'Control Panel', then 'Programs and Features'. Click the option that says, 'Turn Windows features on or off', then in the list that appears scroll down to find 'Services for NFS'. Expand this entry if possible and then click the checkbox for both items inside it. Some Windows versions will only have one entry rather than an expandable one with multiple sub-categories. Click 'OK' at the bottom of the screen to activate the services.

Next, open an Explorer window, then right-click on 'Computer' and select 'Map network drive...'. In the next window, select a drive letter to assign to the share, then in the 'Folder' section, follow the formula:

[IP address of Callisto]:[complete file path to the NFS share]

For example, if the IP address of your Callisto is 192.168.0.115 and the path that you specified when creating the NFS share is /mnt/pool/callisto, then you would enter:

192.168.0.115:/mnt/pool/callisto

Make sure to enter the colon between the IP address and the beginning of the file path. When you connect, the NFS share will mount as a network drive and be available in an Explorer window.

## 7.1 ABOUT DATA BACKUP

To ensure that your files are protected and to prevent data loss, we strongly suggest that you keep two copies of your data on separate devices. A comprehensive backup plan is essential to preventing downtime and loss of valuable data. Any data loss or corruption while using the Jupiter Callisto is the sole responsibility of the user, and under no circumstances will Other World Computing be held liable for compensation or the recovery of any lost data.

## 7.2 ADDITIONAL SUPPORT

### FreeNAS:

The Jupiter Callisto uses a modified version of the FreeNAS® operating system. Some terminology has been changed (e.g. 'pools' and 'filesystems' are called 'volumes' and 'datasets' in FreeNAS), but the online FreeNAS manual will be able to provide extended explanations of many elements that are not covered in this guide. Access the FreeNAS manual at: [doc.freenas.org](http://doc.freenas.org)

### Warranty:

For warranty information, go to: [www.owcdigital.com/jupiter/](http://www.owcdigital.com/jupiter/)

## 7.3 TECHNICAL SUPPORT HOURS AND CONTACT INFORMATION

### HOURS OF OPERATION

8AM - 8PM (CT) Monday - Friday



Telephone Support:  
866-692-7100 (North America)



Live chat is available during normal business hours, at:  
[www.owcdigital.com/livechat/](http://www.owcdigital.com/livechat/)



Or you can email us. Submit your email at:  
[www.owcdigital.com/support/](http://www.owcdigital.com/support/)

## Changes:

The material in this document is for information purposes only and subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, Other World Computing and its officers and employees assume no liability resulting from errors or omissions in this document, or from the use of the information contained herein. Other World Computing reserves the right to make changes or revisions in the product design or the product manual without reservation and without obligation to notify any person of such revisions and changes.

## FCC Statement:

### **Warning! Modifications not authorized by the manufacturer may void the user's authority to operate this device.**

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference with radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

NOTE: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference with radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

## Health And Safety Precautions:

- Use proper anti-static precautions while performing the installation of your hard drives into this drive enclosure. Failure to do so can cause damage to your drive mechanisms and/or the hard drive enclosure.
- Read this user guide carefully and follow the correct procedures when setting up the device.
- Do not attempt to disassemble or modify the device. To avoid any risk of electrical shock, fire, short-circuiting or dangerous emissions, never insert any metallic object into the device. If it appears to be malfunctioning, contact Other World Computing technical support.
- Never expose your device to rain, or use it near water or in damp or wet conditions. Never place objects containing liquids on the drive, as they may spill into its openings. Doing so increases the risk of electrical shock, short-circuiting, fire or personal injury.

## General Use Precautions:

- To avoid damage, do not expose the device to temperatures outside the range of 5° C to 40° C (41° F to 104° F).
- Always unplug the device from the electrical outlet if there is a risk of lightning or if it will be unused for an extended period of time. Otherwise, there is an increased risk of electrical shock, short-circuiting or fire.
- Do not use the device near other electrical appliances such as televisions, radios or speakers. Doing so may cause interference which will adversely affect the operation of the other products.
- Do not place the device near sources of magnetic interference, such as computer displays, televisions or speakers. Magnetic interference can affect the operation and stability of hard drives.
- Do not place heavy objects on top of the device.
- Protect your device from excessive exposure to dust during use or storage. Dust can build up inside the device, increasing the risk of damage or malfunction.
- Do not block any ventilation openings on the device. These help to keep the device cool during operation. Blocking the ventilation openings may cause damage to the device and cause an increased risk of short-circuiting or fire.
- For up-to-date product and warranty information, please visit the product webpage.

## Copyrights and Trademarks:

© 2015 OWC Digital. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of OWC Digital. Jupiter, Other World Computing, and OWC are trademarks of New Concepts Development Corporation, registered in the U.S. Patent and Trademark Office and in Canada. The OWC logo and OWC Digital are trademarks of New Concepts Development Corporation. Apple, Mac, OS X, and Time Machine are trademarks of Apple Inc., registered in the U.S. and other countries. Windows is a registered trademark of Microsoft Corporation in the U.S. and other countries. FreeNAS is a registered trademark of iXsystems. Other marks may be the trademark or registered trademark property of their owners.



Copyright 2015 OWC Digital. All Rights Reserved. Other World Computing's  
Limited Warranty is not transferable and subject to limitations.