



OSVault Tape User Manual

First Edition

OSVault Tape Users Manual, First Edition

Copyright ©2007 Timothy Scott Seson, All Rights Reserved Worldwide

1	About OSVault.....	4
1.1	Introduction	4
1.2	How OSVault Works	4
1.3	Where to Find This Software	5
1.4	About Open Source Storage Inc.	5
2	Installing OSVault.....	6
2.1	Hardware Requirements.....	6
2.2	Booting the Distribution DVD and Loading the Software.....	7
2.3	Power for the Appliance.....	9
2.4	Hooking up Cables.....	9
2.5	Plugging into Your Network	9
2.6	Accessing the Web-Based Interface	10
2.7	Accessing the Command Line Interface	10
3	Configuring OSVault.....	11
3.1	Setting the IP Address.....	11
3.2	Performing a Library Inventory.....	13
3.3	Formatting All Media.....	15
4	Initializing OSVault.....	15
5	Administering OSVault	16
5.1	Web-Based Appliance Administration	16
5.2	Command Line Interface Administration	17
5.3	Basic CLI Commands	17
6	Importing and Exporting.....	23
6.1	Individual Media.....	23
6.2	Magazines.....	Error! Bookmark not defined.
6.3	Web-Based Import/Export	24
6.4	Command Line Import/Export	25
7	Attaching Client Systems to OSVault	27
7.1	Windows 98/2000/XP/2003 Clients	27
7.2	LINUX Clients	29
7.3	Solaris Clients.....	31
7.4	HP-UX Clients.....	31
7.5	MacIntosh OS X Clients	31
7.6	AIX Clients.....	31
8	Media Duplication and Replication.....	32
9	Automatic Backups.....	32
9.1	How Automatic Backups Work.....	32
9.2	How to Restore a System Backup	33
10	Advanced Web-Based Administration Utilities	33

1 About OSVault

1.1 Introduction

This document is intended to aid the reader in installing and using an OSVault Tape appliance. We have assumed that you have purchased a tape library with an attached appliance (server) that already has the OSVault software installed.

1.2 How OSVault Works

OSVault is both a set of software to store files, automatically copy the files to tape media as configured, to reclaim hard disk space from files that have been copied to tape, and to retrieve those files when attempts are made to read them. Additionally it is a system configuration, based on LINUX, that creates a dedicated storage appliance to offer Network Attached Storage (NAS) services.

OSVault consists of the following major software packages:

- LINUX 2.6 – the appliance operating system that supports Network File System (NFS) access by UNIX and LINUX systems and allows the other various software packages to share the appliance hardware
- SAMBA 3.1 – allows Microsoft Windows systems to access files stored on the OSVault appliances using the Common Internet File System (CIFS) that Microsoft uses for file sharing
- webmin – A web (HTML) based system administration package that allows all needed configuration parameters in the appliance to be specified or queried

When fully configured, OSVault looks like a network file server such as a Network Appliance Filer, a Windows 2000 Server, or a LINUX file server. OSVault allows client computers to connect via either of two protocols, CIFS or NFS, to create, read, write, or delete files and directories on the appliance. An OSVault appliance allows direct access by the client system to the /cache directory and OSVault will move files, after a pre-determined period of time, to an available tape cartridge the tape library library.

For example, if a Windows XP user can access the OSVault appliance over a network, then that user can right-click on the My Computer icon, select Map Network Drive..., type in [\\OSVAULT\\CACHE](#), and click on the Finish button. The Windows XP system will then create a new disk drive number that maps into the OSVault appliance and will open an Explorer window showing the contents of the OSVault /cache file system. Files can then be copied into and out of the /cache file system, and application programs, such as Word or Excel, can directly save or open files in that file system.

The /cache file system is periodically scanned by OSVault, and the oldest or largest files are “migrated” from the hard disk onto an available piece of tape

media. A link is created between the /cache file name and the new location for the file on the tape media. In this way, files that are no longer resident on the hard disk inside OSVault can still be opened by the original name. When the original name is opened, OSVault will mount the required piece of tape media and copy the file back to hard disk, before the client system is returned the data it is requesting.

1.3 Where to Find This Software

The core software that makes up the OSVault Tape appliance is open-source. This means that the software source code can be freely downloaded from various sources, including web sites. The following web sites have the software components.

- <http://dvdvault.sourceforge.net> – The changes to the autofs software for an ASACA library and the migration software
- <http://fedora.redhat.com> – the LINUX operating system distribution used by OSVault
- <http://www.samba.org> – the SAMBA software for Microsoft Windows “share” access
- <http://www.webmin.com> – the web-based administration software
- <http://www.osstorage.com/OSVault> – the complete RPM software package that makes up the appliance

Other components are resident on the distribution DVD, including the KICKSTART configuration for installation and the post-LINUX-installation configuration script.

One note about open source software: open source software can be freely downloaded and changed or customized in accordance with the copyrights on the various modules. Support from Open Source Storage for installations is not free and requires the purchase of a support contract. Additionally, users of the software for commercial purposes are obligated to release their modifications to that software to the general public.

1.4 About Open Source Storage Inc.

Open Source Storage Inc. is a professional services company that develops and maintains the OSVault software. Since OSVault is open-source, OSS Inc. does not actually sell the software. Open Source Storage makes revenue to fund continuing development by assisting customers (end-users and manufacturers) with the installation, deployment and support of OSVault.

2 Installing OSVault

OSVault is intended to run as an appliance on an Intel or AMD processor system. Certain assumptions are made including:

- The appliance will not be used for other purposes or applications
- The hard disk drives on the appliance do not contain data that will be permanently preserved, other than OSVault specific data
- The hardware meets the hardware prerequisites given below

2.1 Hardware Requirements

OSVault is usually shipped to a customer site as a pre-configured appliance. However, this document also shows the steps needed to install the OSVault Library Appliance distribution on DVD onto a generic hardware platform. OSVault has been tested on various hardware platforms and is officially supported on hardware that contains the following minimum hardware components:

- Intel Pentium 4 processor (single processor is acceptable, dual CPU configurations have also been tested)
- 1GByte of RAM memory (more memory is acceptable)
- Two ATA or SATA disk drives with at least 500GBytes of storage on each disk
- Adaptec 79xx, 29160 or 39160 SCSI Host Bus Adapter (HBA) or LSI Logic SCSI Host Bus Adapter (HBA)
- A LINUX(Fedora)-supported Gigabit Ethernet Interface

Expected system performance is somewhat dependant on the hardware used in the appliance. For example, a properly configured Pentium 4 system with a single 10,000RPM SATA disk drive connected to a Gigabit Ethernet switch should be able to read and write data in cache at around 40Mbytes per second. The limiting factor is the throughput to the hard disk drive in the cache, not the network connection, the memory or the CPU. Adding or trunking a second gigabit ethernet link to the appliance will not increase performance, since the magnetic disk is still the limiting factor. Using a RAID-5 configuration on three or more disks for the cache can increase performance substantially, and the appliance will be able to move around 75Mbytes per second through a single Gigabit Ethernet link to a single client system. With multiple Gigabit Ethernet interfaces, a single OSVault appliance can move 300Mbytes per second or more with a large RAID-5 configuration.

The type of data moved over the network can also impact performance. The megabyte per second numbers given previously assumes that large files are transferred sequentially into or out of the cache. If a number of small files are

transferred, throughput will drop considerably due to the overhead involved in creating directory entries for new files in the cache. A large number of very small files all stored in a single directory can drop performance below 1Mbyte per second, even with fast magnetic disks.

The migration of data out of the cache to the tape drives imposes a performance penalty of around 45Mbytes per second per tape drive used. With a number of disk drives configured into the appliance, an OSVault appliance can drive 3 or more LTO-3 tape drives and still move data on a Gigabit Ethernet interface at full speed.

Reading files that reside on tape media can take various amounts of time to complete. The factors that impact that amount of time are:

- Load times for tape media. Some media may already be loaded into a tape drive, and the time may only include tape positioning time. Files on media that isn't loaded into a drive can take as long as two to three minutes to start retrieving
- Size of file to read. Since the file needs to be moved from tape to hard drive first, a very large file will take longer to retrieve back than a smaller file. A 50 GByte file from LTO-3 tape can take 10 to 15 minutes to retrieve after the tape media is loaded into the drive.
- Other cache based reads or writes at rates high enough to use a significant percentage of the network bandwidth -- A single client blasting data to or from the cache can use 90% of the network bandwidth for periods of time and will reduce other accesses accordingly

Although complex, the performance of the OSVault appliance is predictable and your sales representative can help you with expected system performance for your application.

2.2 Booting the Distribution DVD and Loading the Software

Quick Start Checklist

1. Unpack the tape library according to the directions. Make sure to remove the internal retaining items.
2. Load the tape library with media. You MUST put a piece of media in the first slot in the library. This piece of media is used for backups of the appliance during operation and used to determine the number and position of tape drives in the library.
3. If a separate server is being installed and the library is connected with SCSI hardware, place the server within 6 feet of the tape library to attach it to the library. If using a Fibre Channel library, connect the server fiber channel card directly to the fibre channel port on the library. For SAN (switch) configured installations, plug all Fibre Channel ports into your switch(es) and configure appropriately

4. Determine an IP (Internet Protocol) address for the server and make a note of it.

Formatted: Bullets and Numbering

5. Determine an IP address for the tape library and set the address on the tape library via the library's front panel.

6. Install CAT-5 network cables between the library and your network switch, and between the server and your network switch. Wireless (802.11a/b/g) is not supported.

7. Install the power cables to the server and the library, being careful to keep the power cables and any SCSI cables from running parallel in close proximity.

8. Attach a keyboard and monitor to the server. Some servers support serial port BIOS settings and you can use another computer with the "minicom" or HyperTerminal software on it in place of the keyboard and monitor.

9. Power up the tape library.

10. Power up the server.

11. Place the distribution DVD into the server DVD-ROM drive and reboot the server if necessary.

12. KEEP IN MIND THAT INSTALLING THE DVD DISTRIBUTION SOFTWARE WILL ERASE THE HARD DRIVES ON THE SERVER.

13. When the monitor on the server displays the "Boot:" prompt and waits for input, type in the following command:

```
linux ks=cdrom:/ks.cfg
```

14. The above command will load the OSVault software onto the server and will format the boot disk and cache disk for operation. The IP address of the server will initially be set to 10.1.1.2, which you can change after installation completes.

Formatted: Bullets and Numbering

15. You will be prompted to indicate where the distribution media resides. Select CDROM with the arrow keys. Press the TAB key until the OK button is highlighted, and then press return.

16. Software installation will take approximately 30 minutes (longer on some servers). If the distribution media is on multiple pieces of media, for example CD, you will be prompted to switch to the other pieces of media.

17. After software installation, the now-configured appliance will perform a hardware inventory on the ASACA library (15 minutes per 250 slots) and will then load every piece of media in the library and examine it. It takes about a minute for each piece of media to be formatted or checked. If the media is unformatted, OSVault will put a UDF file system on each side of each piece of media. If the media is formatted, the volume label on the media will be used to create a directory (using the volume name) in /archive.

-
18. OSVault will then reboot the server and all services on the server will be started. You can configure the appliance via a web browser (<http://10.1.1.2:10000>) or via the attached keyboard/monitor. The web interface is preferred unless you are an expert LINUX administrator.

2.3 Power for the Appliance

THIS DEPENDS ON CONFIGURATION. Check the manuals that came with the server hardware. If the OSVault has multiple power supplies using more than one power cable, then it can still run if one of the power supplies fails. Best practice is to plug each power plug into separate circuit breakers to protect against power interruptions. If you have an Uninterruptible Power Supply (UPS), you should plug one power cord into the UPS and one into the wall power.

2.4 Hooking up Cables

The following cables must be attached to the appliance:

1. SCSI cables from the VHDCI connector on the back of the appliance to the Micro-68 connector on the tape library. The exact cabling depends on the configuration of the tape library.
2. Power cables from a power source (100-240 volt) to the appliance and to the library
3. Network cables from the customer's network facility to the appliance Gigabit Ethernet port (RJ-45) and to the ASACA library 10/100BaseT port (RJ45). This cable must be at least CAT-5 compliant.

Formatted: Bullets and Numbering

2.5 Plugging into Your Network

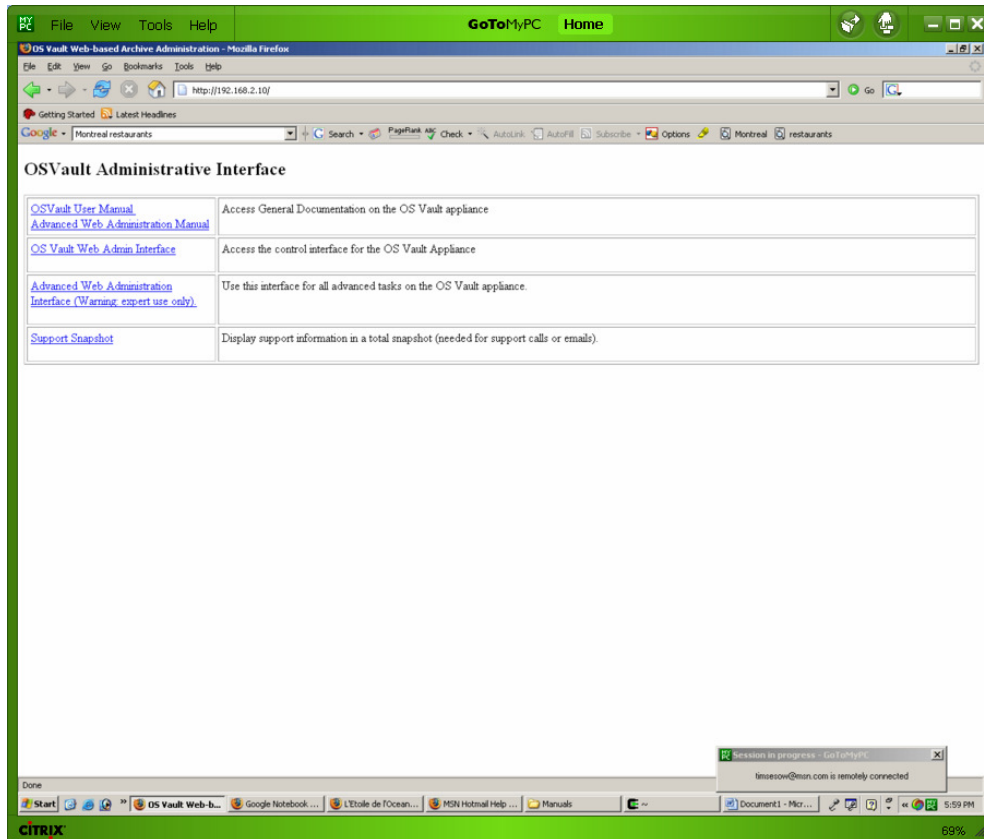
The OSVault appliance has two 10/100/1000BaseT network interfaces. During installation, the RJ-45 connector labeled "0" should be plugged into the local area network. Installing the RJ-45 connector labeled "1" is optional. The cabling from connector "0" and the premise's network switch should be a straight-thru cable, CAT-5 or better.

The network interface will automatically determine what speed to run based on the network switch capabilities, the quality of the cabling used, and the length of the cable. With a limited number of network switches, it may be necessary to lock down the duplex of the link to either half or full.

The appliance supports the ability to do 802.3ad Link Aggregation (sometimes called trunking or EtherChannel). Link Aggregation allows the appliance to be both faster and more fault tolerant, but your network switch needs to support this capability.

2.6 Accessing the Web-Based Interface

The OSVault appliance has a web portal (main page) at the default address on the appliance. When first installed, the appliance is at <http://10.1.1.2>. Any web browser on another system in the same 10.1.1.x network can type the link into the web browser and be presented with the following screen:



2.7 Accessing the Command Line Interface

Using Secure Shell

Secure shell, referred to as SSH, allows the entry of commands in a line or text mode into the appliance, with all data encrypted so that others cannot snoop on what you are doing. This is particularly important when using an external Internet connection to control the appliance, since external Internet data can be intercepted while transiting networks outside of your local office. Also, anyone within a company with physical access to your company network can, with relative ease, intercept any communications to the appliance. Therefore, SSH is recommended for all command line interface

activities and is the only interface, other than the appliance console, enabled for command line access.

SSH is available on all Windows, LINUX, Macintosh and UNIX systems. For Windows systems, a copy of PUTTY can be installed or a copy of CYGWIN can be downloaded from <http://www.cygwin.com>. See Appendix A for instructions on installing CYGWIN on a Windows system.

Installing CYGWIN with its default settings on a Microsoft Windows system will allow you to invoke SSH in a command window. Just double click on the CYGWIN desktop icon on your Microsoft Windows system and then type the following command:

```
ssh root@10.1.1.2
```

You will then be prompted for the “root” password, which by default is “password”. Your local system will then store keys used for encrypting the communications with the OSVault appliance. If you later change the IP address for the OSVault appliance, the next time you use SSH, a new set of keys will be generated.

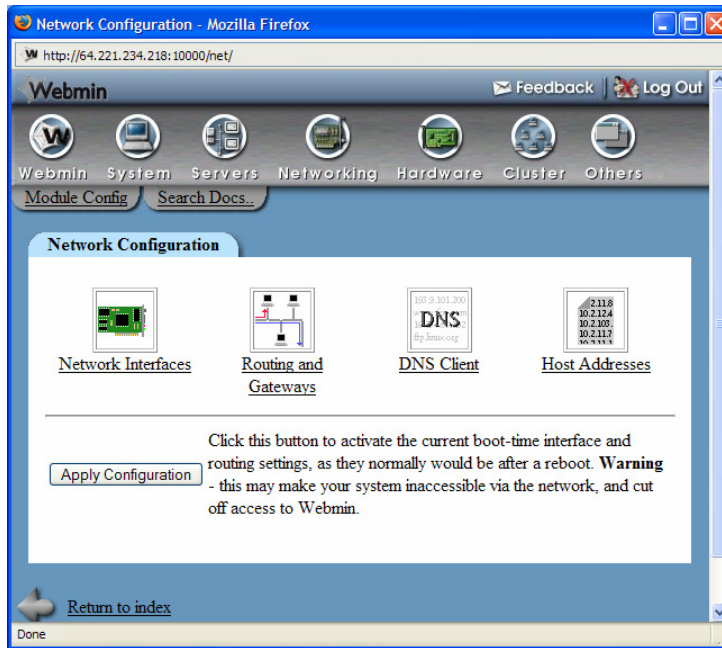
Using TELNET

TELNET is disabled by default on the OSVault appliance.

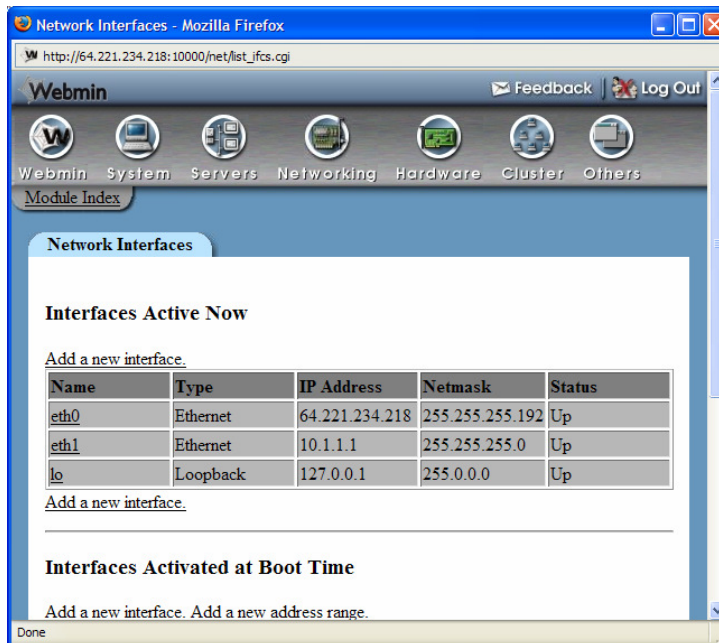
3 Configuring OSVault

3.1 Setting the IP Address

You can use the “Network Configuration” link to change the appliance network address to work in your local network. For example, if you wish to change the OSVault appliance address to 192.168.2.14, first select the “Advanced Web Administration Interface, then click on the “Network Configuration” link and you will see the following screen:

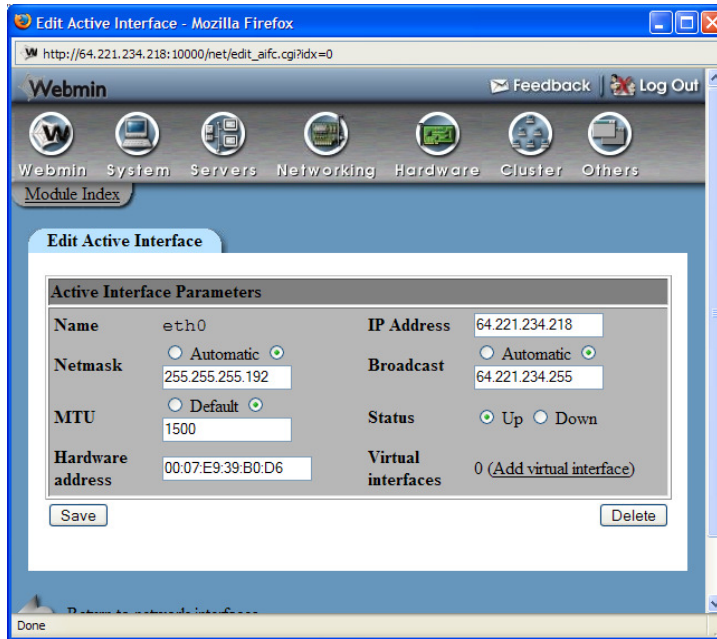


You can then click on the “Network Interfaces” icon to go to the screen that will allow you to set the IP address of the appliance, as follows:



4/16/2007

In the above screen, there are two Ethernet interfaces on the appliance, eth0 and eth1. To change eth0 to the desired IP address, click on the name “eth0”, and use the following screen to set the information:



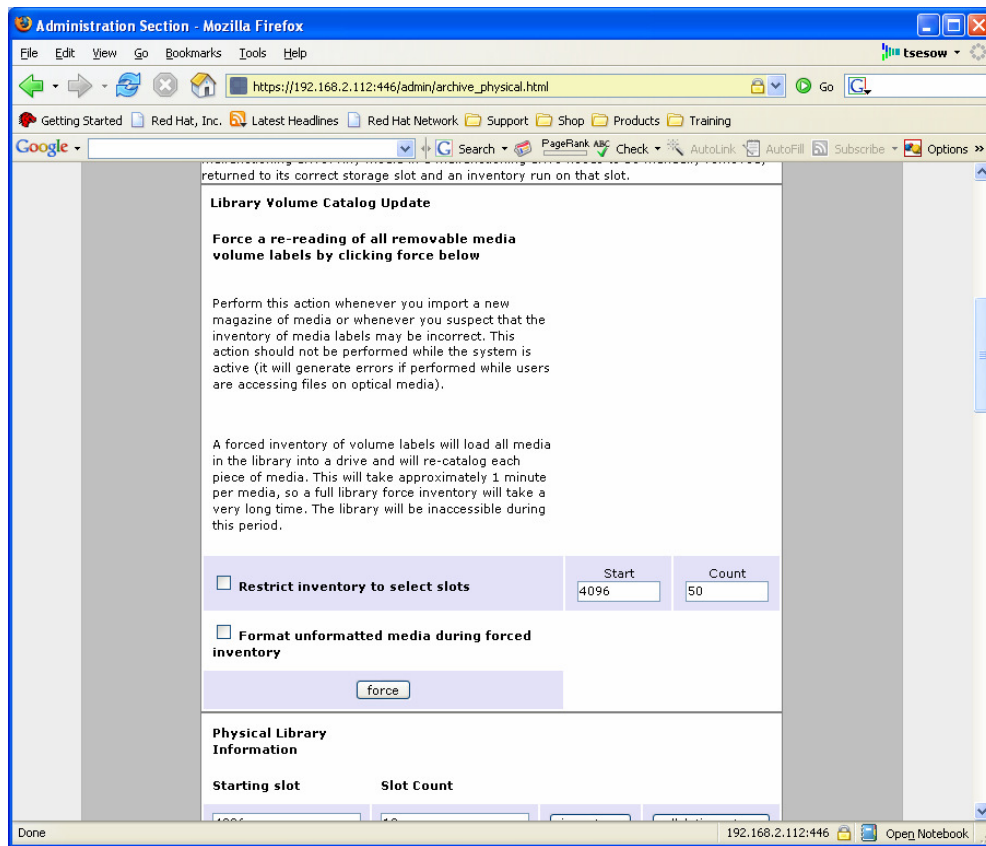
3.2 Performing a Library Inventory

OSVault, during installation, will perform a full inventory of the library. This inventory includes having the library use its internal intelligence to check for the presence of media in each slot (called an Initialize Element command), and then a load of all detected media into an optical drive to see if the media has existing data on it.

Please keep in mind that this process can take a long time. Checking for existing media on a tape drive can take 5 minutes per tape, so that a 72 slot tape library with three tape drives can take 2 hours to complete an inventory.

Web Inventory of the Library

Use a web browser to type in the IP address of the OSVault appliance. You can then select the library control by clicking on OSVault Web Administration Interface, then clicking on Archive Management, then Physical Library Management. That will present the following window:



Clicking on the “force” button will initiate a full library inventory. Each piece of media in the library will be loaded into a drive and have the label read, compared to the barcode on the tape cartridge, and stored in the appliance. The “Format unformatted media during forced inventory” checkbox has no effect when inventorying a tape library and is only used for optical media.

Command Line Inventory of the Library

```
inventory -o /dev/sg0 -f
```

The above command uses the /dev/sg0 device entry to issue commands to the robot to mount all media for data checks. You can determine the device name with the mcstat command.

4/16/2007

Inventory will mount each tape cartridge in the library. It will try to read a standard header from the tape cartridge to see if the cartridge is already written to and it will record the tape cartridge label in the database of media maintained by the OSVault database.

The inventory command will also determine the tape drive configuration of the library and store that in a database in the OSVault appliance.

3.3 Formatting All Media

In some cases, it may be necessary to format all of the media in the library, regardless of data already stored on the media. If the “format all” option is selected on the inventory command, all (REPEAT ALL!!!) data in the library will be erased.

Web-Based Method Not Available

To avoid the possibility of mistakenly formatting all media and wiping out current data, a web interface to accomplish a full format is not implemented.

Command Line Method

There is a command line interface option to the inventory command that will cause the inventory command to ignore any existing internal label on the media, and mark each media as unused and available to be completely overwritten. To indicate that “inventory” should format all media, the “-x” flag must be given twice to the “inventory” command, as shown below:

```
inventory -o /dev/sg0 -x -x
```

4 Initializing OSVault

Once the OSVault appliance software is installed, the system will prompt you to reboot the appliance. The appliance will then restart with the newly installed software and will start a series of initialization tasks, including:

Startup of the network interfaces, and setup of CIFS (Windows SHARE) and NFS services

- Startup of the web-based system control (ports 80 and 10000 by default)
- Mounting of the magnetic drives for cache space
- Scheduling of the migration task to relocate files written to cache onto tape cartridges
- (One time only) Having the library perform a hardware initialization to determine what media is stored in the library

-
- (One time only) Loading each piece of media into a drive and detecting if the tape cartridge is already written.
 - Initialization of the Relational Database that tracks file copies from cache to tape cartridge and the contents of each file moved.

Once the OSVault appliance has finished its one-time inventory and format process, any references to files migrated from the cache to tape cartridge will cause the tape library to load the appropriate media to fulfill that reference.

5 Administering OSVault

You can access all of the OSVault administration capabilities via a network interface. Additionally, if a standard PS/2 keyboard and VGA (or better) monitor is attached to the appliance, you can perform the command line utilities via that interface.

5.1 Web-Based Appliance Administration

For advanced administration tasks, OSVault uses a customized installation of the Webmin software, a web-based system administration utility. Since the underlying operating system for OSVault is LINUX, a LINUX implementation of Webmin allows for advanced system configuration beyond the defaults set by the OSVault installation.

Logging In to the Advanced Web Administration Interface

Logging into Webmin is easy. Open a web browser, like Netscape or Internet Explorer, on any machine that has network access to the server that you wish to login on. Browse to port 10000 on the IP or hostname of the server.

Webmin will then respond with either an authentication window, or an authentication web form, where you can enter the administrator user name (usually "root") and password. After successful authentication, you will be greeted with the Webmin index page. The type of login form you receive (either on a web page or in a popup window) depends on the configuration of the Webmin server. The differences between session authentication and standard HTTP authentication will be discussed later in this book. In general, all of the web interface screens required to perform system administration have been previously discussed. However, OSVault contains a full suite of system administration web tools that allow for further customization. Documentation on the general system web interfaces are contained in Section 10 of this document.

5.2 Command Line Interface Administration

The OSVault appliance allows for the use of command line administration via either Secure Shell (ssh) or TELNET. The command line interface is a full LINUX command shell and can be used to set any system configuration, to reconfigure hardware or to reinitialize the system. The command line interface should be used with care, as it is possible to destroy data stored in the appliance if the commands are used incorrectly.

5.3 Basic CLI Commands

mcstat

Mcstat is a command line program that locates the library interface in the appliance. The library interface has a name such as /dev/sg0 and mcstat will print out that name when run. Mcstat also prints out information on the library vendor and product string, as retrieved via the SCSI or Fibre Channel interface, as well as the addressing information of the library. The output from mcstat is as follows:

```
[root@OSVAULT ~]# mcstat
SCSI Media Changer Status Command
Medium changer found on /dev/sg3
Vendor ASACA
Product AM250DVD
  2 transports starting at 1
  250 storage elements starting at 4096
  3 data devices starting at 64
  1 import slots starting at 128
[root@OSVAULT ~]#
```

In the above listing, the library attached to the appliance is an ASACA 250 slot DVD-RAM library. Since the library can be virtualized or be missing magazines, the storage element count is important. In this case, there is a standard dual picker, 250 slots for media in magazines of 50 each and 3 DVD-RAM drives.

readvol

Readvol is a command line utility to report which slots are occupied in the tape library. Readvol will retrieve the hardware inventory from the library and use that information to determine what media is present. If media in the library has been manually removed or added without performing an “initelem” command, the results reported by readvol will not be correct.

You need to use the identifier returned from mcstat as the only argument to readvol, for example /dev/sg3. A sample run of readvol is shown below:

```
[root@OSVAULT ~]# readvol /dev/sg3
Querying ASACA , AM250DVD
Volume in slot 4096 (relative slot 0)
Volume in slot 4097 (relative slot 1)
Volume in slot 4098 (relative slot 2)
Volume in slot 4099 (relative slot 3)
Volume in slot 4100 (relative slot 4)
Volume in slot 4101 (relative slot 5)
Volume in slot 4102 (relative slot 6)
Volume in slot 4103 (relative slot 7)
Volume in slot 4146 (relative slot 50)
Volume in slot 4147 (relative slot 51)
Volume in slot 4148 (relative slot 52)
```

inventory

The inventory command line utility is the overall workhorse utility for OSVault. Running inventory at any time will reload all media present in the library and make sure the OSVault directory of media loaded in the library is correct.

An example of inventory output is shown below:

```
[root@OSVAULT ~]# inventory -l /dev/sg3
Querying QUALSTAR , XLS900
Drive device name not specified, autoconfiguring
Using /dev/st2 as the drive to format in
Skipping 1 slot(s) for backups
Media in slot 4097, side 0 has label TAPE0001
Media in slot 4097, side 1 has label TAPE0002
Media in slot 4098, side 0 has label TAPE0003
Media in slot 4098, side 1 has label TAPE0004
Media in slot 4099, side 0 has label TAPE0005
Media in slot 4099, side 1 has label TAPE0006
```

You can see in the above example that “inventory” can locate the order of tape drives in the library by itself. It does this by loading the first piece of media in the library into each drive and then seeing which drive reports a “media present” condition.

chkconfig

Chkconfig provides a simple command-line tool for maintaining the /etc/rc[0-6].d directory hierarchy by relieving system administrators of the task of directly manipulating the numerous symbolic links in those directories.

This implementation of chkconfig was inspired by the chkconfig command present in the IRIX operating system. Rather than maintaining configuration information outside of the /etc/rc[0-6].d hierarchy, however, this version directly manages the symlinks in /etc/rc[0-6].d. This leaves all of the configuration information regarding what services init starts in a single location.

Chkconfig has five distinct functions: adding new services for management, removing services from management, listing the current startup information for services, changing the startup information for services, and checking the startup state of a particular service.

When chkconfig is run without any options, it displays usage information. If only a service name is given, it checks to see if the service is configured to be started in the current runlevel. If it is, chkconfig returns true; otherwise it returns false. The `-level` option may be used to have chkconfig query an alternative runlevel rather than the current one.

If either `on`, `off`, or `reset` is specified after the service name, chkconfig changes the startup information for the specified service. The `on` and `off` flags cause the service to be started or stopped, respectively, in the runlevels being changed. The `reset` flag resets the startup information for the service to whatever is specified in the init script in question.

By default, the `on` and `off` options affect only runlevels 2, 3, 4, and 5, while `reset` affects all of the runlevels. The `-level` option may be used to specify which runlevels are affected.

Note that for every service, each runlevel has either a start script or a stop script. When switching runlevels, init will not re-start an already-started service, and will not re-stop a service that is not running.

putaway

Putaway is a command line utility to restore media that is located in drives, or in the import/export slot, back into a storage slot. It is actually a renamed version of “inventory”, since “inventory” will perform the same function prior to starting the loading of media into drives. When “inventory” is renamed to “putaway”, it will exit prior to starting the load of media.

The format for putaway is

```
putaway -l /dev/sg3
```

where /dev/sg3 is the name of the library device inside the appliance. An example follows:

```
[root@OSVAULT ~]# putaway -l /dev/sg3
Querying ASACA , AM250DVD
Volume in drive 64
putting cart in drive 64 in slot 4096
[root@OSVAULT ~]#
```

initelem

Initelem is a utility to tell the tape library to check each slot and see if media is present. Due to the mechanical nature of checking for media, a full hardware initialization will take a moderate amount of time. Initelem has one argument, the device name of the library interface that is returned by the mcstat command. With no arguments, initelem will inventory the entire library. Optionally, you can enter arguments to initelem to limit the number of slots checked:

- -s – Set the starting slot address to start checking at
- -n – Check “n” slots then stop

```
[root@OSVAULT ~]# initelem -n 4096 -s 50
SCSI Media Changer Status Command
Vendor ASACA
Product AM250DVD
Setting timeout to 720000
```

mm

Mm which is short for “move medium”, is a utility to move a piece of media within the tape library. The format of the “mm” command is:

mm from to invert picker device

Where “from” is the library slot to move from; “to” is the library slot, drive or import/export slot to move to; “invert” is a “1” (meaning to flip the media during the move) or “0” (do not flip); “picker” is the library picker to use to move the media (use “0” for best results); and “device” is the library device name inside the appliance (for example “/dev/sg3”).

Tape libraries do not support the flip option and that argument is ignored.

As an example, the following command will move the media from the second drive (slot 65) to the 10th storage slot (slot 4106) using the default picker

without flipping it, and then move it back again. The third command is an error where there is no media in the drive:

```
[root@OSVAULT ~]# mm 4096 65 0 0 /dev/sg3
[root@OSVAULT ~]# mm 65 4096 0 0 /dev/sg3
[root@OSVAULT ~]# mm 65 4096 0 0 /dev/sg3
Source slot is empty
Could not move medium
[root@OSVAULT ~]#
```

sourceslot

Sourceslot is a utility to display the slot that a piece of media belongs in. The one argument required is a drive number (starting at 0, not 1) that has a piece of media loaded into it. Running “sourceslot /dev/sg3 0” will return the slot number the library pulled the media from. The following is an example of three commands you can run to empty the media from a drive, back to the slot it came from.

```
[root@OSVAULT mclib]# mcstat
SCSI Media Changer Status Command
Medium changer found on /dev/sg3
Vendor ASACA
Product AM250DVD
  2 transports starting at 1
 250 storage elements starting at 4096
   3 data devices starting at 64
   1 import slots starting at 128
[root@OSVAULT mclib]# mm 4096 64 0 0 /dev/sg3
[root@OSVAULT mclib]# sourceslot /dev/sg3 0
Querying ASACA , AM250DVD
Volume in drive 64 came from 4096
[root@OSVAULT mclib]# mm 64 4096 0 0 /dev/sg3
```

tape_migration

The “tape_migration” command is used to copy files from the “/cache” file system to an appropriate tape cartridge based on rules given as arguments. Usually, the “tape_migration” command is invoked from the automatic job scheduler, “cron”, so that files are periodically copied to removable media.

The “migration” command has several arguments:

- -t# – Don't migrate files that are not at least “#” minutes old

- `--source-name` – The name of the file system root directory to search for candidate files to relocate to optical media
- `--force-mig` (or `-f`) – Migrate all files from the `/cache` file system to optical media without checking if the files are old enough;
- `--debug-level level` – Print out different levels of information; valid values for “level” are “critical”, “serious”, “error”, “warning”, “debug”, and “everything” (Use “`--debug-level everything`” to show the most information.)

The default entry in the scheduler (cron) for migration is:

```
0 * * * migration -t10 --source-name /cache -f
```

purge

The “purge” command is used to truncate files from the “/cache” file system if the file is already copied to tape cartridge and if space is needed on the /cache file system. Usually, the “purge” command is invoked from the automatic job scheduler, “cron”, so that the “/cache” file system is emptied frequently enough so that it does not fill completely.

The “purge” command has several arguments:

- `-t#` – Don't purge files that are not at least “#” minutes old
- `--source-name` – The name of the file system root directory to search for candidate files to relocate to optical media
- `--low-water #` – Don't start purging files from the “source-name” until the file system that contains “source-name” has more than “#” percent used.
- `--high-water` – Stop purging files from “source-name” when the file system that contains “source-name” is less than “#” percent used
- `--debug-level level` – Print out different levels of information; valid values for “level” are “critical”, “serious”, “error”, “warning”, “debug”, and “everything” (Use “`--debug-level everything`” to show the most information.)

The default entry in the scheduler (cron) for migration is:

```
0-55/5 * * * purge -t10 --source-name /cache --high-water 30 --low-water 20
```

tape_stage

The “tape_stage” command called automatically by the operating system to restore files from tape cartridge back to hard disk cache when a file is opened for reading.

Tape_stage can also be invoked manually from the command line to restore files back into hard disk cache. You can restore files by file name (as they would appear in the /cache directory) or by inode number. The inode number is an internally generated, unique number that uniquely identifies the file, even when renamed. The system uses this version of stage to restore files; users generally use the file name to restore a file.

The syntax of a stage by filename might be:

```
tape_stage --filename /cache/userdata/music/recording.mp3
```

The syntax of a tape stage by inode number might be:

```
tape_stage --inode 32
```

6 Importing and Exporting

6.1 Individual Media

Import

Media to be imported must first be placed into the Import/Export slot (I/E slot) in the tape library. Use the front panel on the library to insert the media into the I/E slot. Refer to the tape library User Manual on how to accomplish this. The media must be inserted into the library BEFORE telling OSVault to import it.

When importing media, OSVault will perform the following operations:

- Check to see if media is present in the I/E slot; if not an error is returned and processing of the import operation stops
- Moves the piece of media to an unoccupied slot in the tape library; if no empty slots are available in the library an error is returned and processing of the import operation stops
- Update the database to show which media label is in the slot
- Indicate success of the operation when complete.

Export

Exporting media from the library makes the files on that media unavailable for I/O in the OSVault appliance, so requests for files on that media (or to files

that were migrated from cache to that media) will result in an error being returned to the requesting application.

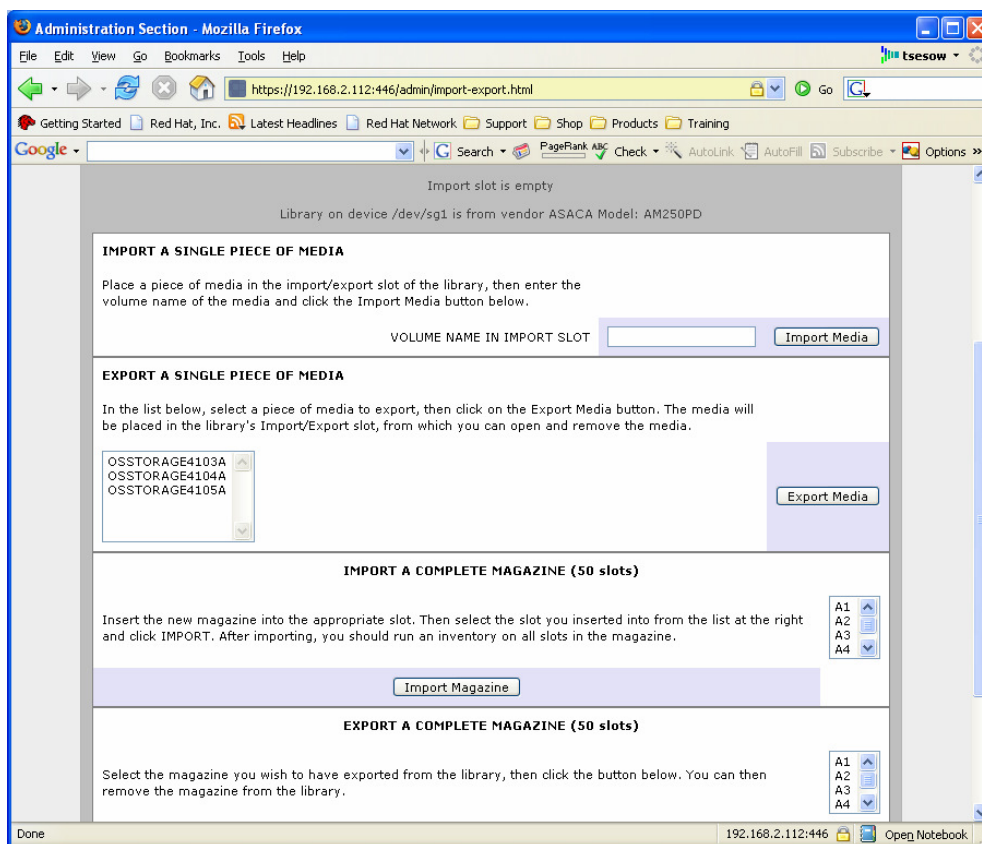
Attempting to export media that is currently in use, i.e. mounted in an optical drive, will probably fail with an error such as “Unmount failed, device in use”. If you receive this error, wait until the media is no longer in use and retry the export operation.

To removed the media from the I/E slot after exporting, refer to the Tape Library Users Manual for directions on use of the library front panel.

Formatted: Bullets and Numbering

6.2 Web-Based Import/Export

If you open a web browser at <http://10.1.1.2> (the default OSVault IP address), then select “OSVault Web Admin Interface”, then select “Archive Mgt”, then select “Import/Export”:



This screen can be used to import single pieces of media through the import/export slot of the library. It is assumed that media being imported already has a barcode label installed with the same Volume Name specified in the screen.

4/16/2007

Import Media

Once the media is inserted into the I/E slot (see above), click on the IMPORT button on the web browser. This will move the piece of media out of the import/export slot, place it in an empty slot in the library and update the inventory of media.

Export Media

Using the Library/Media Management screen in your web browser, select the media name that you wish to export. Clicking on the export button will then start the media export process, which performs the following operations:

- The I/E slot in the library is checked to verify that it is currently empty, if it is occupied the export operation is terminated with an error and the media selected for export remains in the library
- The media is moved from its storage slot to the I/E slot
- The database is updated to show that the media is no longer present in the library

To remove the media from the I/E slot after exporting, refer to the Tape Library Users Manual for directions on use of the library front panel.

Formatted: Bullets and Numbering

6.3 Command Line Import/Export

Import Media

Use the “importmedia” command to import media into the tape library. Media to be imported must first be placed into the Import/Export slot (I/E slot) in the tape library.

For example, to import a tape cartridge and put it into slot 4100, use the following command after inserting the cartridge into the import/export slot of the tape library:

```
importmedia -store_slot 4100
```

This will move the piece of media out of the import/export slot in the library, place it in slot 4100 and update the inventory of media to show that the media is now back in the library. The media name is retrieved from the barcode on the tape cartridge. Tape cartridges without barcodes or with unreadable barcodes will not be imported, will be left in the I/E slot and an error will be returned from the importmedia command.

Export Media

To export media from the library, use the “exportmedia” command. The “exportmedia” command takes one argument which is the media label on the media, e.g. TAPE0007. To removed the media from the I/E slot after exporting, refer to the Tape Library Users Manual for directions on use of the library front panel.

For example, to export the tape cartridge labeled TAPE0007 from the library, use the command:

```
exportmedia TAPE0007A
```

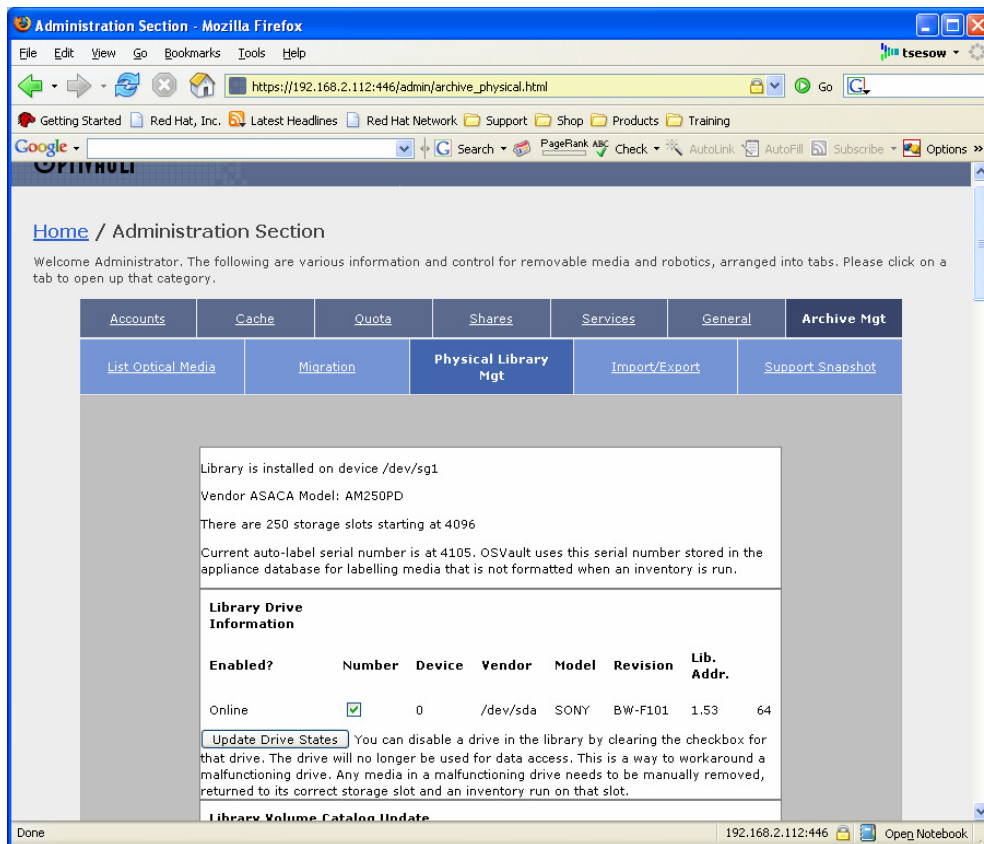
This will move the media from its storage slot to the import/export slot in the tape library and will update the OSVault inventory to show that TAPE0007 is absent from the library.

6.4 Tape Drive Management

OSVault will display the drive status of all drives in the tape library in the “Archive Mgt”/“Physical Library Mgt” screen of the web-based admin interface. If the checkbox for a particular drive is checked (“enabled”) then that drive will be used by OSVault as an available drive to read or write optical media.

If a drive were to become disabled (hardware failure of some type), this will interfere with the operation of OSVault. Until the drive can be repaired, you can un-check the checkbox on the faulty drive and select “Update Drive States” and the drive will no longer be used for I/O requests. After the drive is repaired and placed back into service, simply check the checkbox and “Update Drive States” to start using the drive again. Keep in mind that a piece of media may become jammed in the faulty drive and be unavailable for use by the system.

In general, it is best to repair hardware problems as soon as practical to avoid interference with normal system operations.



7 Attaching Client Systems to OSVault

7.1 Windows 98/2000/XP/2003 Clients

OSVault appears on your network as a Windows 2000 server. The “cache” share is automatically created when the appliance starts up. Writing to the “cache” share will always result in the files being created on hard disk first. After a period of time, older files will be moved to optical media and all further I/O operations will work directly against the optical media.

If your Windows machine is on the same subnet as your OSVault appliance, you should now be able to connect to the OSVault server by clicking on Start / Run and typing:

\\OSVAULT\pub

Mangling Method

Name mangling is a method where Windows allows long filenames to retain a short filename equivalent. For example, C:\Program files can also be referred to

4/16/2007

as C:\PROGRA~1. There are two algorithms available, hash and hash2. If you store a lot of files on your server, you can help to avoid name collisions by adjusting your global parameter section with mangling method = hash2. This is set in the OSVault Advanced Web Interface.

Oplocks

Opportunistic locking essentially means that the client is allowed to download and cache the file on their hard drive while making changes; if a second client wants to access the file, the first client receives a break and must sync the file back to the server. This can give significant performance gains in some cases; in others, some programs insist on syncing back the contents of the entire file for a single change. Level1 Oplocks (aka just plain "oplocks") is another term for opportunistic locking.

Level2 Oplocks is a fancy way of saying that you are providing opportunistic locking for a file that will be treated as "read-only". Typically this is used on files that are read-only or on files that the client has no intention to write to (at least, not initially).

Kernel Oplocks are essentially a method that allows the Linux kernel to co-exist with Samba's oplocked files, although this is simplifying things a bit. SGI IRIX and Linux are the only two UNIX's that are oplock aware at the moment. Unless your system supports kernel oplocks, you should disable oplocks if you are accessing the same files from both Unix/Linux and Smb clients.

Regardless, oplocks should always be disabled if you are sharing a database file (e.g., Microsoft Access) between multiple clients, as any break the first client receives will result in the entire file needing to be synced (not just the single record), which will result in a noticeable performance delay and, more likely, problems accessing the database in the first place. Notably, Microsoft Outlook's personal folders (*.pst) react very badly to oplocks. If in doubt, disable oplocks and tune your system from that point. If client-side caching is desirable and reliable on your network, you will benefit from turning on oplocks. If your network is slow and/or unreliable, or you are sharing your files among other file sharing mechanisms (e.g., NFS) or across a WAN, or multiple people will be accessing the same files frequently, you probably will not benefit from the overhead of your client sending oplock breaks and will instead want to disable oplocks for the share. Another factor to consider is the perceived performance of file access. If oplocks provide no measurable speed benefit on your network, it might not be worth the hassle of dealing with them.

Windows 98/ME

Fortunately, both Windows 98 and Windows ME use encrypted passwords by default. All that's left is to configure your Network properties in the Control Panel. Make sure you have Client for Microsoft Networks and TCP/IP installed along with your network adapter. If you're not using a DHCP server, you'll probably need to configure the TCP/IP settings as well.

Under the properties for Client for Microsoft Networks, check the box for Log on to Windows NT domain and specify the domain name that your Samba server is using. Finally, under the second tab labeled Identification, give your computer a unique name and specify the name of the workgroup (which should match the name of the domain, although not strictly necessary).

Windows XP/2000

You'll need to enable TCP/IP on your system, as well as configure a WINS server, from the control panel. The IP address of the WINS server should be your OSVault appliance server (where the `nmbd` daemon is running). To map a drive using the Windows GUI, open My Computer. On the toolbar are many options; look for one that provides a list that includes the phrase Map Network Drive (under Windows XP this option is under the Tools menu). Once selected, a new box will open up. In the Drive box, click a drive letter that you wish to use. In the Folder box, type the path for the server and path that you wish to connect to.

Alternately, from a Command Prompt, you can type:

```
NET USE F: \\MYSERVER\\PUB /YES
```

7.2 LINUX Clients

Before beginning, you should double-check to make sure your mount program is new enough (version 2.10m if you want to use Version 3 NFS), and that the client machine supports NFS mounting, though most standard distributions do. If you are using a 2.2 or later kernel with the `/proc` filesystem you can check the latter by reading the file `/proc/filesystems` and making sure there is a line containing `nfs`. If not, typing `insmod nfs` may make it magically appear if NFS has been compiled as a module; otherwise, you will need to build (or download) a kernel that has NFS support built in. In general, kernels that do not have NFS compiled in will give a very specific error when the `mount` command below is run.

To begin using a machine as an NFS client, you will need the portmapper running on that machine; and to use NFS file locking, you will also need `rpc.statd` and `rpc.lockd` running on both the client and the server. Most recent distributions start those services by default at boot time.

With `portmap`, `lockd`, and `statd` running, you should now be able to mount the remote directory from your server just the way you mount a local hard drive, with the mount command. Suppose our server is called *OSVault*, and we want to mount the `/cache` directory on *slave1.foo.com*. Then all we have to do from the root prompt on *slave1.foo.com* is type:

```
mount OSVault:/cache /mnt/home
```

and the directory `/cache` on master will appear as the directory `/mnt/home` on *slave1*. (Note that this assumes we have created the directory `/mnt/home` as an empty mount point beforehand.)

You can get rid of the file system just like you would for a local file system,
by typing:

```
umount /mnt/home
```

Getting NFS Servers Mounted at Boot Time

NFS file systems can be added to your `/etc/fstab` file the same way local file systems can, so that they mount when your system starts up. The only difference is that the file system type will be set to **nfs** and the dump and fsck order (the last two entries) will have to be set to zero. So for our example above, the entry in `/etc/fstab` would look like:

```
none 10.1.1.2:/cache /OSVAULT_CACHE nfs defaults 0 0
```

See the manual pages (using the *man* command) for *fstab* if you are unfamiliar with the syntax of this file. If you are using an automounter such as *amd* or *autofs*, the options in the corresponding fields of your mount listings should look very similar if not identical.

At this point you should have NFS working, though a few tweaks may still be necessary to get it to work well.

Soft vs. Hard Mounting

There are some options you should consider adding at once. They govern the way the NFS client handles a server crash or network outage. One of the cool things about NFS is that it can handle this gracefully, if you set up the clients correctly. There are two distinct failure modes:

soft

If a file request fails, the NFS client will report an error to the process on the client machine requesting the file access. Some programs can handle this with composure, most won't. We do not recommend using this setting; it is a recipe for corrupted files and lost data. You should especially not use this for mail disks, if you value your mail.

hard

The program accessing a file on an NFS mounted file system will hang when the server crashes. The process cannot be interrupted or killed (except by a "sure kill") unless you also specify **intr**. When the NFS server is back online the program will continue undisturbed from where it was. We recommend using **hard,intr** on all NFS mounted file systems.

Picking up from the previous example, the `fstab` entry would now look like:

```
# device mountpoint fs-type options dump fsckord
```

```
...  
master.foo.com:/home /mnt/home nfs rw,hard,intr 0 0  
...
```

Setting Block Size to Optimize Transfer Speeds

The **rsize** and **wsize** mount options specify the size of the blocks of data that the client and server pass back and forth.

The defaults may be too big or too small; there is no size that works well on all, or most, setups. On the one hand, some combinations of Linux kernels and network cards (largely on older machines) cannot handle blocks greater than 128Kbytes. On the other hand, if they can handle larger blocks, a bigger size might be faster.

Getting the block size right is an important factor in performance and is a must if you are planning to use the NFS server in a production environment.

7.3 Solaris Clients

Not yet tested.

7.4 HP-UX Clients

Not yet tested.

7.5 Macintosh OS X Clients

OS X supports both the Network File System and the CIFS (Windows) Share.

If you wish to access the OSVault appliance via the CIFS protocol, first make sure that you've enabled SMB support in the Macintosh Directory Access Utility (in the Applications/Utility folder). Select SMB and then click Configure. You can now enter in your preferred workgroup and WINS server information. Next, activate the Finder (e.g., click on the Finder in the Dock). Select Go, then Connect to Server. In the Server Address field, type in your server address. For example:

```
smb://OSVAULT/CACHE
```

After clicking Connect, you may be prompted for your authentication.

7.6 AIX Clients

Not yet tested.

8 Media Duplication and Replication

The OSVAULT appliance supports a configuration option to make the library automatically copy data on tape cartridges to a second cartridge. This is usually done to either export a copy of your data to an external location or to guard against tape media failure. Keep in mind that configuring the system to make copies of all tapes in a library will reduce your overall online tape storage space in half.

Automatic tape copies in an OSVAULT appliance are enabled by removing the comment header “#” from the command line in the cron file. This can be done from a command line with the “crontab -e” command, editing the files and saving it back out again. The editor that “crontab -e” invokes is, by default, the “vim” editor, sometimes also called “vi”.

For example, when first installed, the cron file contains the following line:

```
# 0 23 * * * /usr/sbin/tape_copy -auto
```

If you remove the “#” from this line, the tape_copy program will run every night at 11pm and will make a copy of any media that changed during the previous day. So if you wrote two tape cartridges during the day, the tape_copy program will find two unused cartridges in the library and copy the data to the copy tapes. The database will be updated to show that these tapes are copies and not available for future migration jobs.

The list of tapes that are marked as copies of other tapes can be listed from the command line with the “tape_copy_list” command.

9 Automatic Backups

9.1 How Automatic Backups Work

According to the schedule set in the “cron” schedule, a backup will run periodically (at 1am by default) where the cache directory will be backed up to the first piece of tape media in the library. Also backed up at this time is certain system configuration information.

The purpose of the backup is to enable the restoration of operation after a hard disk failure without needing to full read all media in the tape library. Even if the backup media is not used, it is possible to restore the full system operation after a hard disk failure by reinstalling the system software and performing a full inventory.

The automatic backup does not make a copy of the /cache file system data, so migration must run frequently enough to ensure that data is stored on the tape media. If migration is running less frequently, or if a higher level of reliability is required, then the /cache file system should be placed on a RAID hard drive system with RAID-1 or RAID-5 level configured. OSVault has been tested with LINUX software RAID (MD driver), the 3Ware 9xxx series RAID controllers,

Engenio RAID systems (marketed by STK, SGI, and IBM), the Xyratex RAID systems (with Chaparral RAID controllers) and the Nstor RAID systems.

9.2 How to Restore a System Backup

The backup on the first optical media is a “tar” file that needs to be restored via the command line interface with the following commands:

```
mcstat
putaway 0
mm 4096 64 0 0 /dev/sg3
tar xvf /dev/st2
```

The above steps assume that /dev/st2 is the name of the first drive (topmost) in the library. /dev/sg3 is assumed to be the device name of the library interface.

10 Advanced Web-Based Administration Utilities

For more detail on the more advanced GUI interface screens, see the OSVault Advanced Web Administration Manual.