

December 2008

# Linux P2V migrations using MondoRescue

By Hewlett Packard

Main Writer: Lester Wade

Contributions from: René Ribaud, Gallig Renaud, Bruno Cornec

[License: Creative Commons Attribution-Share Alike 3.0](#)

MondoRescue is a backup and recovery tool for Linux, it's packaged for various distributions and supports common architectures (i386, x86\_64 and ia64). It allows online and offline backups to network storage, local disks, cd/dvd and tape. It supports a large variety of filesystems (including but not limited to ext2, ext3, ReiserFS and XFS) and partition/disk layouts (software raid, hardware raid and LVM1/2). During a restore MondoRescue will also resize partitions depending upon the new disk geometry. Those coming from an HP-UX background may liken MondoRescue to HP Ignite-UX.

The methods MondoRescue uses to archive and restore a machine means it's well suited for use as a P2V/P2P (physical-2-virtual/physical-2-physical aka P2V/P2P) migration tool.

This document is aimed at users with a basic understanding of Linux, virtualization solutions and some familiarity in partitioning and configuring disks. It details the process for migrating a physical machine into a VMware virtualized environment. The process can easily be adapted to apply to other virtualization products such as Xen, Citrix XenServer, VirtualBox and KVM/QEMU as well as to other machines with different hardware components.

Migrations of HP hardware utilizing the SmartArray controller will also be mentioned. This is to ensure the reader is aware of a current bug in MondoRescue and how to work around it.

Instructions for the migration of other distributions/releases are earmarked for future inclusion in this document. Specifically SLES9, RHEL3 and Debian systems to supplement the information currently included for RHEL4, RHEL5 and SLES10.

The document is structured as follows:

## Table of Contents

Linux P2V migrations using MondoRescue.....	1
Setting the scene.....	2
Application & Key Features.....	2
How MondoRescue works.....	3
Pre-requisites and assumptions.....	3
Installing MondoRescue.....	3
Pre-migration Configuration.....	4
Archiving the server.....	5
Restoring the server.....	5

Migrating SuSE Linux Enterprise Server 10 .....	9
Notes for SLES 9 and RHEL 3.....	10
Working with the HP Smart Array controller and LVM.....	10
Troubleshooting .....	11
Summary.....	12
Acknowledgement.....	12

## Setting the scene

MondoRescue was originally developed by Hugo Rabson and a small team. Development duties have now been handed over to Bruno Cornec and Andree Leidenfrost. Bruno Cornec is the EMEA Linux Profession Lead for HP.

As the commercial virtualization solutions from vendors such as VMware, Citrix and Microsoft become more prevalent, the requirement for migration tools has increased. A company may maintain hundreds or thousands of existing physical servers which they will want move into the virtual environment rather than build from scratch. The ability to quickly and easily migrate servers is a key factor.

Effective migration tools for Linux are few and far between. Existing products which support Linux include Platespin PowerConvert and VMware Converter. However, at the time of writing none of these products support 64bit Linux migrations. Given that many companies are using 64bit Linux to take full advantage of the latest hardware, these migration tools are ineffectual.

## Application & Key Features

How can MondoRescue be used? Does it meet my requirements? Does it suit my environment?

Before progressing through this document it's worth spending time looking at the potential application of MondoRescue. The reader is probably wondering whether MondoRescue can fulfill a specific requirement and this section will outline some key features of MondoRescue to try to answer some of those questions.

Key Features:

- Network connection between target and source machine is not required. This is useful in environments where security is of key concern, a sysadmin may not want to open ports on perimeter firewalls for services such as NFS or SSH
- The archival of a system can be made onto the local disk of the source machine, DVD/CD media, USB stick or shared network volume (e.g; NFS).
- As a backup tool, MondoRescue is capable of Incremental snapshots of the source system
- The system does not need to be taken offline for archival (although it is recommended if data is volatile)
- MondoRescue will automatically resize partitions if the target machine has a different sized disk (sized by ratio new/old).

- MondoRescue is available for 32 and 64bit x86 architectures and ia64. Note that you must be migrating/restoring to the same architecture.
- MondoRescue packages are available for RedHat, Fedora, OpenSuSE, SLES, Mandriva, Debian, Ubuntu, Gentoo and Slackware. Source packages are also available.

## How MondoRescue works

A migration using MondoRescue consists of two steps. The first is the use of the **mondoarchive** utility to examine the server, group files into filesets dependant on size, archive and compress these filesets and finally to create a bootable CD/DVD containing the archive(s).

Secondly, **mondorestore** will repartition the disk(s) on the target virtual machine, setup filesystems, restore mountpoints and unpack the archived data.

## Pre-requisites and assumptions

This document presumes the following:

- You have a 64bit (not essential for following the guide) source server running Red Hat Enterprise Linux 4.
- Your virtualization environment is fully configured and a target virtual machine for the migration has been created and configured appropriately.
- There is network connectivity between the source machine and virtualization host
- You have access to the Internet for obtaining MondoRescue
- The target virtualization environment is VMware (server, workstation or ESX), although other products will follow very similar guidelines.

## Installing MondoRescue

Installing the latest version of MondoRescue is recommended. The example operating system used for this section is Red Hat Enterprise Linux 4 AS. Red Hat Enterprise Linux 5 follows the same guidelines. The RPM's can be downloaded from the MondoRescue site with the following commands (provided a direct connection to the Internet is available):

```
mkdir /tmp/mondo
cd /tmp/mondo
wget ftp://ftp.mondorescue.org/rhel/4/afio-2.4.7-1.x86\_64.rpm
ftp://ftp.mondorescue.org/rhe/4/buffer-1.19-1.x86\_64.rpm
ftp://ftp.mondorescue.org/rhel/4/mindi-2.0.4.rhel4.x86\_64.rpm
ftp://ftp.mondorescue.org/rhel/4/mindi-busybox-1.7.3-1.rhel4.x86\_64.rpm
ftp://ftp.mondorescue.org/rhel/4/
```

[mondo-2.2.7-1.rhel4.x86\\_64.rpm](http://ftp.mondorescue.org/rhel/4/mondo-doc-2.2.7-1.rhel4.x86_64.rpm)  
[ftp://ftp.mondorescue.org/rhel/4/mondo-doc-2.2.7-1.rhel4.noarch.rpm](http://ftp.mondorescue.org/rhel/4/mondo-doc-2.2.7-1.rhel4.noarch.rpm)

Installing the RPM's is a simple case of:

```
rpm -ivh *.rpm
```

## Pre-migration Configuration

Prior to the migration of the physical server, it's necessary to prepare the system for the move.

**NOTE** - For migrations in a production environment it's good practice to use a temporary IP address on the source machine. This will mean users and services can't connect to the server during the migration.

Firstly, the `initrd` (INITial Ram Disk) must be altered to include drivers (modules) for the LSI Logic SCSI controller. A physical machine's disk controller will most likely differ from that which the virtual guest (or target machine) will use. In the case of VMware (server, workstation and ESX) the LSI Logic disk controller is emulated.

The following command will build a new `initrd` containing the LSI logic controller modules:

```
mkinitrd --with=mptspi /boot/initrd-`uname -r`-virt.img `uname -r`
```

An entry in the GRUB menu should be added so that once the machine is migrated, it boots using the new `initrd`. Duplicate the default entry and alter it:

```
title Red Hat Enterprise Linux AS (2.6.9-67.ELsmp)
root (hd0,0)
kernel /vmlinuz-2.6.9-67.ELsmp ro root=/dev/VolGroup00/LogVol100
rhgb quiet
initrd /initrd-2.6.9-67.ELsmp-virt.img
```

Ensure the new entry is the default selection.

Next, a declaration in `/etc/modprobe.conf` must be added:

```
alias scsi_hostadapter qla2400
alias scsi_hostadapter1 mptsspi
alias eth0 e1000
```

You also have the possibility to do those steps after `mondorestore` has restored your system, just before rebooting the machine.

If your target system uses a SATA controller you have to do instead:

```
alias scsi_hostadapter ata_piix
```

You should also edit the `/usr/sbin/mindi` file to include the disk controller modules for the *target* machine. Look for the `FORCE_MODS` line and edit:

```
FORCE_MODS="mptsas mptspi"
```

or you can load the driver with `modprobe mptspi` so that `mindi` picks it up in the

backup.

## Archiving the server

The method in this section is to archive the system to a .iso image on the local disk. This requires plenty of free space depending upon the amount of data to be archived (MondoRescue doesn't archive free space). If you have an NFS share available, you may want to archive to that.

**NOTE** - Make sure you've assigned the system a temporary IP address prior to the archive.

Presuming that MondoRescue has been installed, run the following command to generate an .iso image:

```
mondoarchive -O -i -N -g -d /storage -T /usr/mondo/tmp -S /usr/mondo/scratch -E "/video /junk" -s 4480m
```

The options are as follows:

- O This tells mondoarchive to backup
- i The -i option specifies that ISO files are to be used as the media
- N Exclude all mounted network filesystems
- g GUI mode
- d Specifies the backup device/directory, in this case it was /storage
- T Specifies the full path of the temporary directory to be used during archival
- S Specifies the full path of the scratch directory where ISO images are built
- E Exclude these directories, /video and /junk as examples
- s The size of the ISO files, 4480m for a single sided DVD

The archival operation may take some time and depends upon how much data is to be archived. The MondoRescue GUI displays current progress.

Alternatively, running `mondoarchive` without options will run take you through the interactive GUI.

[As the archiving is done live on the system, you may want to stop all processes maintaining coherent information, such as RDBMS systems for which the operation could cause a synchronisation data issue.](#)

Once the archival is complete, copy the ISO image to the virtualization host. You could place it on a NFS share, or copy it directly to local storage:

```
scp /storage/mondorecue-1.iso user@virtualhost:/scratch/
```

## Restoring the server

On your virtual host, edit the virtual machine onto which you'd like to migrate and connect the virtual CD/DVD-ROM drive - mounting the iso image MondoRescue created and that you transferred in the previous step.

Start the virtual machine and boot from the CD/DVD-ROM drive. The bootable CD/DVD image contains a small operating environment based on mindi, which builds using the kernel and initrd from the source machine during the archive process.

On Proliant servers this can also be done using the iLO (Integrated Lights Out) interface and virtual media.

When booting from the image, the following menu will appear:

```
ISOLINUX 3.11 2005-09-02 Copyright (C) 1994-2005 H. Peter Anvin
Welcome to MINDI-LINUX v2.0.3-r2004, a cousin of Mondo Rescue

Creation date: Thu Aug 21 12:08:23 BST 2008
These disks were generated on the following system:-
  Red Hat Enterprise Linux Server release 5.2 (Tikanga)
  Kernel 2.6.18-92.el5 on an x86_64m

To format and restore all files automatically, type 'nuke' <enter>.
To restore some/all files interactively, type 'interactive' <enter>.
To compare the archives with your filesystem, type 'compare' <enter>.
To boot to a command-line prompt (expert mode), type 'expert' <enter>.
You may add one or more of the following parameters as well:-

oodonteject - mondorestore will not eject the CD; this is useful if, for
oo             instance, your PC's case has a concealed CD-ROM drive
oonoresize    - your mountlist will not be adjusted to use your unallocated
oo             hard disk space
ootextonly    - do not call any Newt library routines; this is unattractive
oo             but valuable if you find your Newt library has bugs in it

e.g. Type 'nuke donteject textonly' if you have an unstable Newt library and
a PC whose CD-ROM drive tray would be damaged if it unexpectedly ejected.
boot: _
```

Figure 1 - Main menu

Notice in figure 1 that the operating system, kernel and architecture of the system on which the disk was created is shown.

If left for 30 seconds, the MondoRescue GUI will appear prompting you to choose a function. Alternatively 'nuke' will run a non-interactive system restore and 'expert' will drop to a shell prompt.

Type 'interactive' and press enter. Mindi will then load a basic operating environment, using driver modules contained on the original host system (the one you archived).

```
Probing sr_mod...
Probing scsi_mod...
Probing scsi_transport_spi...
Probing sd_mod...
Probing floppy...
Probing paride...
Probing loop...
Probing ide-cd...
Probing ide-cs...
Probing edd...
Probing ata_piix...
Probing libata...
Probing ahci...
Probing sata_nv...
Probing pata_marvell...
Probing cdrom...
Probing usb-storage...
Probing uhci-hcd...
Probing ehci-hcd...
Probing ohci-hcd...
Probing mptbase...
Probing mptscsih...
Probing mptspi...
Probing e1000...
_
```

Figure 2 - Loading modules

Figure 2 shows the Mindi environment booting, kernel modules are being loaded. Note the mpt\* modules for LSI Logic controllers typically used by VMware virtual machines.

The first menu in interactive mode will ask you what you want to do. Select 'Interactively' to be prompted at each stage of the restore process. This will also allow you to make any necessary changes.

At the next screen, select the media from which you would like to restore. Choose 'DVD Disks'.

The next screen will prompt you to alter the mountpoints if required.

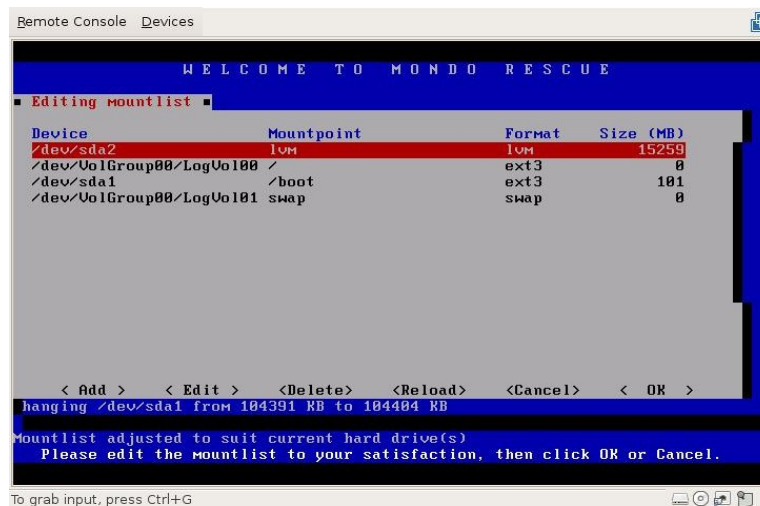


Figure 3 - Editing mountpoints

Figure 3 shows that MondoRescue has picked up the correct sizing for the partitions given the size of the new disk, it will alter any sizes by comparing the size of the new disk and the size of the old and altering proportionally. Check to ensure the mountpoints are correct and select 'OK' to proceed. Say 'Yes' to any warning messages if you know the mountlist is correct.

MondoRescue will then ask if you definitely want to erase and format the disks, select 'Yes'. Once the countdown has finished, MondoRescue will begin to

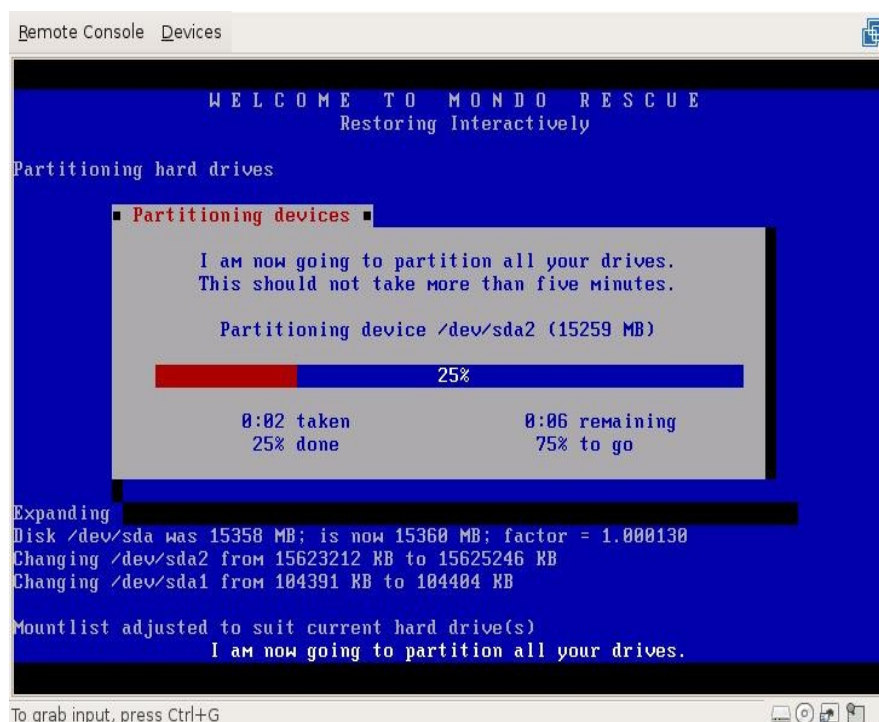


Figure 4 - Partitioning of devices



partition your drives.

Figure 4 shows MondoRescue partitioning /dev/sda2. In this guide, /dev/sda2 is an LVM partition so MondoRescue will also handle the creation of physical volumes, logical volumes and volume groups based on the layout of the source system.

Next, MondoRescue will format the partitions with the same filesystems used on the source machine.

Once the format is completed successfully (errors would be displayed), MondoRescue will ask whether you'd like all of your data restored. Select 'Yes' to begin the restoration.

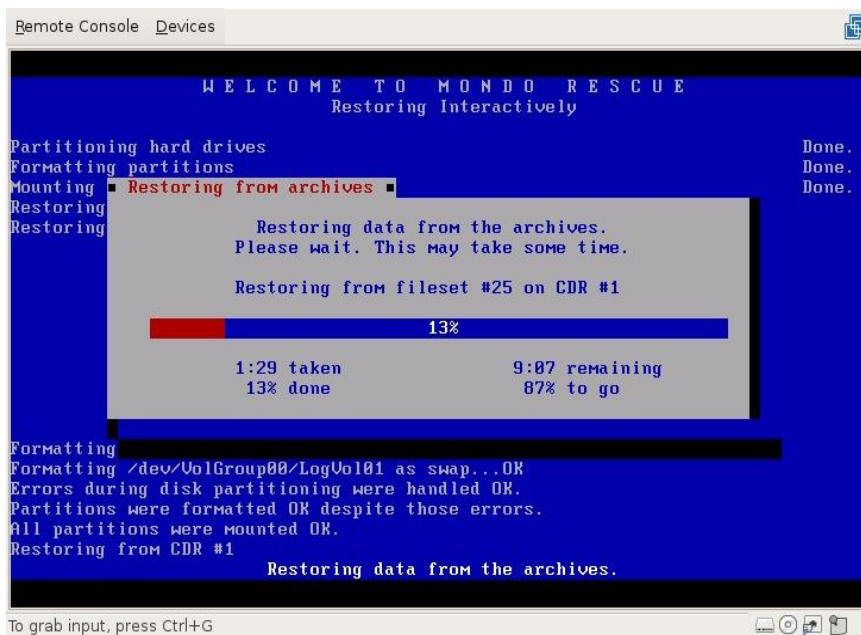
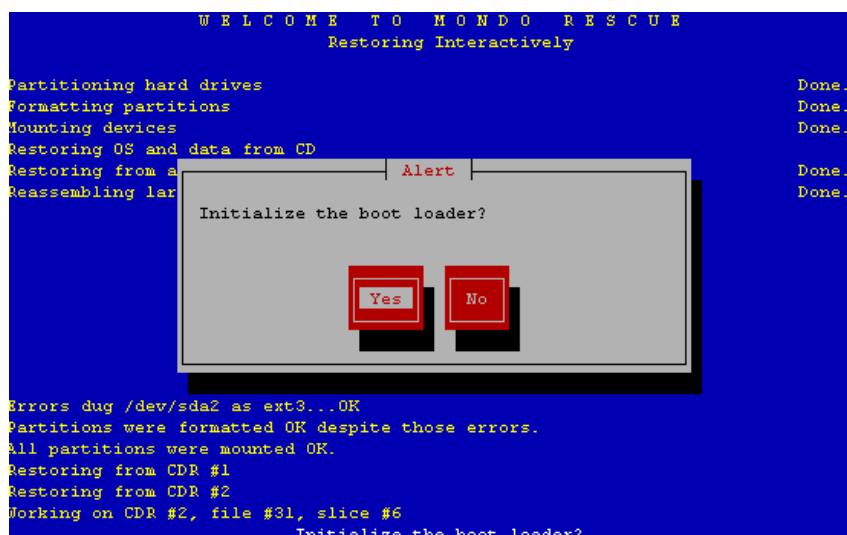


Figure 5 - Restoring data

MondoRescue will rebuild the file structure from the archived data. Depending on how much data was on the system and the performance of the virtual host, this can take a while. This is probably a good time to grab a coffee.

Once the restore is complete, MondoRescue will ask you whether you'd like to initialise the boot loader.





-Select 'Yes' and select 'No' when asked whether you changed the mountlist or not. Prior to rebooting the system, MondoRescue will ask one last question - whether you would like to label any partitions.

Typically a standard install of Red Hat will label the boot partition as '/boot'. This will then be referenced in /etc/fstab and /boot/grub/menu.lst. It's important to ensure that the partition is relabeled correctly, else referred to by its path e.g; /dev/sda1.

By selecting 'Yes' to label partitions, the restore wizard will exit and drop you to a shell prompt. From here you can use the following command to check filesystem labels:

```
e2label /dev/sda1
```

Before continuing, you may want to use your virtualization tool to take a snapshot before rebooting, in order to be able to come back to that point and fix points you may have forgotten.

Type 'exit' and hit return if everything looks ok. This will reboot the virtual machine.

The system should boot successfully. You've now migrated your physical server to a virtual server.

## Migrating SuSE Linux Enterprise Server 10

So far this document has covered the migration of Red Hat Enterprise Linux 3, 4 and 5. This section will outline specifics relating to the migration of a vanilla install of SLES 10 (SP2) 64bit.

Pre-requisites are afio, buffer, newt and newt-devel. Download these RPM's from the MondoRescue FTP site:

```
wget ftp://ftp.mondorescue.org/sles/10/afio-2.5-64.x86\_64.rpm  
ftp://ftp.mondorescue.org/sles/10/buffer-1.19-740.x86\_64.rpm  
ftp://ftp.mondorescue.org/sles/10/mindi-2.0.4-1.sles10.src.rpm  
ftp://ftp.mondorescue.org/sles/10/mindi-busybox-1.7.3-1.sles10.src.rpm  
ftp://ftp.mondorescue.org/sles/10/mondo-2.2.7-1.sles10.src.rpm  
ftp://ftp.mondorescue.org/sles/10/newt-0.52.2-1.2.x86\_64.rpm  
ftp://ftp.mondorescue.org/sles/10/newt-devel-0.52.2-1.2.x86\_64.rpm
```

mtools, gcc, gcc-c++ are also pre-requisites. Download these from Novell's SLES10 yast repositories using zypper (along with dependencies):

```
zypper -in gcc gcc-c++ mtools
```

Install the afio, buffer and newt RPM's:

```
rpm -ivh newt-0.52.2-1.2.x86_64.rpm newt-devel-0.52.2-1.2.x86_64.rpm afio-2.5-64.x86_64.rpm buffer-1.19-740.x86_64.rpm
```

Next, the source RPM's for MondoRescue, Mindi and Mindi-busybox must be built. Do the following as root:

```
rpmbuild --rebuild mondo-2.2.7-1.sles10.src.rpm mindi-2.0.4-1.sles10.src.rpm mindi-busybox-1.7.3-1.sles10.src.rpm
```

Now that MondoRescue is installed you can prepare the initrd with the LSI logic

scsi adapter modules that will be required for the system to boot under VMware. Alter the modules depending upon your virtualisation platform.

To do this, edit `/etc/sysconfig/kernel` and add the modules to the "INITRD\_MODULES" line. Run 'depmod' to create an updated list of module dependencies and then `mkinitrd` to create a new `initrd` with the modules included. Running `mkinitrd` without any options will update the existing `initrd`, rather than creating a new one.

```
depmod
mkinitrd
```

**NOTE** - Unless you specifically require the `fdomain` module (for certain Adaptec controllers), add `fdomain` to the `DENY_MODS` line in `/usr/sbin/mindi` PRIOR to archival. Failure to do so may result in your system hanging whilst booting into the restore environment.

**NOTE** - Unless required, specify the `scsi_debug` module on the `DENY_MODS` line in `/usr/sbin/mindi`. This will remove any anomalous devices when restoring within VMware and moving from `/dev/hd*` devices to `/dev/sd*` devices.

You can now begin the archival of your system using `mondoarchive`. Follow the guidance from the previous section(s) to complete an archive.

Once archived you can mount the resulting iso file on the virtual machine and begin a restore (again, see previous sections for details). After the restore is complete, you'll need to check `/etc/fstab`, `/boot/grub/device.map` and `/boot/grub/menu.lst` and alter the entries in these files if required.

**NOTE** - It is important that you edit `/boot/grub/menu.lst` (say 'Yes' to labelling partitions at the end of the restore process to reach a command prompt) rather than relying on `/etc/grub.conf`, which `mondorestore` will open for you. SLES 10 does not symlink `/etc/grub.conf` and instead relies on `/boot/grub/menu.lst`. Also, ensure that `/dev/by-id` entries for disks are removed from this file and replaced with the device names (e.g: `/dev/sdb2`). SuSE has used labelling of disk devices by id since SLES 10 SP1.

## Notes for SLES 9 and RHEL 3

On HEL 3, the following declaration in `/etc/modules.conf` must be added:

```
alias scsi_hostadapter mptbase
alias scsi_hostadapter1 mptscsih
alias eth0 pcnet32
```

On SLES 9, modify the file `/etc/sysconfig/kernel`. Change the `INITRD_MODULES` section on top:

```
INITRD_MODULES="scsi_mod sd_mod ext3 mptbase mptscsih pcnet32"
```

## **Working with the HP Smart Array controller and LVM**

If you're working with HP ProLiant's, or HP products which use the Smary Array controller (`cciss` module) then you're quite likely to come up against [bug#183](#).

Essentially this bug means that if you are cloning a system which uses LVM on Smart Array onto a different disk controller, LVM will not be restored correctly.

To work around this it is necessary to take the following steps:

When booting from the rescue iso, enter expert mode and configure the partitions prior to restoring in interactive mode. It's important that all mountpoints that existed on the source machine are serviced (unless you want to retire some of them).

- Boot from the MondoRescue ISO created during archival and enter expert mode.
- Use fdisk / parted to partition the virtual disk manually.
- Create volume groups and logical volumes on the appropriate LVM partition (you may use /tmp/i-want-my-lvm as a template)
- Once the disk is prepared enter interactive mode (type 'interactive' at the shell prompt).
- Edit the mountlist when prompted. Ensure each mountpoint is serviced by partitions / logical volumes created in step 2 and 3. Ignore any sizing discrepancies.
- Hit 'OK' and DO NOT let MondoRescue partition or format your disks.
- Select 'Yes' when asked whether you'd like to restore all your data.
- At this point you can pick up this guide from the bottom of page 9. See the troubleshooting section below if anything happened to go wrong.

## Troubleshooting

### **Problem 1**

Attempting to log onto the migrated machine causes the user to be dumped back at the login prompt.

### **Solution 1**

This may be an SELinux issue. Boot from a rescue CD, chroot to the mounted filesystem and disable SELinux with the system-config-securitylevel tool (Red Hat only).

### **Problem 2**

Migrated system fails to boot.

### **Solution 2**

Check the following; ensure filesystems are labeled as appropriate, check that the grub bootloader configuration is correct (re-install grub if required), make sure that mountpoints are correct in /etc/fstab and that the mpt\* LSI Logic modules (for VMware) are included in /etc/modprobe.conf and that they exist on the system (updatedb or find / -name mptbase.ko etc.). You may also want to check to see if Mindi included them during ISO creation (see /usr/sbin/mindi).

### **Problem 3**

Filesystems fail to mount during startup.

### **Solution 3**

Check /etc/fstab to ensure assignments are correct. Also make sure that the volumes actually exist as expected.

#### **Problem 4**

VM boots with 'operating system not found' error message

#### **Solution 4**

There is a good chance that the bootloader is not installed properly. Install GRUB by using the 'grub-install' command. You'll need to boot into a rescue environment then mount and chroot to your root partition (ensure that /proc and /sys are mounted). The command takes the option of your disk device (e.g: /dev/sda). Once GRUB is reinstalled, check the entries in /boot/grub/menu.lst and /boot/grub/device.map.

**NOTE** - You may want to use the mount-me command when booting from the MondoRescue restore CD/DVD. This will mount the require partitions for you.

## **Summary**

MondoRescue is an extremely flexible tool which lends itself to many different roles. Typically used for full and incremental system backups it works equally well for the migration of servers within both physical and virtual environments. You can find out more information about MondoRescue from [www.mondorescue.org](http://www.mondorescue.org).

## **Acknowledgement**

Many thanks to the MondoRescue team in their assistance and willingness to help when I was originally assessing this method for P2V/P2P migrations.