



TECHNICAL HowTo

Imaging Linux systems with hardware changes using Mondo Rescue

ABSTRACT

This document describes the process to create and deploy system images from HP Proliant equiped with a Smart Array Controller to some other hardware controller using mondorescue software.

TABLE OF CONTENTS

	2
General remarks	4
Copyright and confidentiality	4
Version history	4
References	4
Prerequisites	5
Requirements	5
Backup the source server	5
Backup onto the local server	5
Backup onto NFS share	6
Deployment on target server	8
Restoring from Media	8
Restoring from network	8
Mondorestore part	9
Manual part	12

General remarks

Copyright and confidentiality

The information contained in this document is subject to change without notice. Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Use, duplication, or disclosure is subject to restrictions as set forth in contract subdivision (c)(1)(ii) of the Rights in Technical Data and Computer Software clause 52.227-FAR14.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304, USA

© Copyright 2003 Hewlett-Packard Company.

Version history

Version	Date	Comments
1.0	2005.07.06	First draft / FX. Horel
2.01	2005.10.05	Review, add mkinitrd process / FX. Horel
2.02	2005.10.07	Review / B. Cornec
2.1	2007.05.21	Addition for ATCA / E. Montaut
2.2	2007.06.08	Complete review + adding P2V / G. Renaud
2.3	2007.07.23	Review + new tips / A. Polaert

References

Persons or documents	Description

Prerequisites

Requirements

The following steps must be done on the source system (HP Proliant with a SCSI Raid controller – cciss driver) :

- Mondorescue must be installed (mondo + mindi RPMs/debs) (Latest version available at <http://www.mondorescue.org>). This has been tested with mondo 2.2.3/mindi 1.2.3.
- You may need to remove the previous version as rpm version are not well handled when those lines are written. Use the command 'rpm -e mindi mondo mindi-busybox'
- You may require additional package such cdrecord,afio and/or buffer
- Stop all the processes on the server for whose the online backup could cause problems (for example: Databases)

Backup the source server

Backup onto the local server

Enter the following lines in a new file called mondo-bck (with execute rights):

```
# cat > /usr/local/bin/mondo-bck << EOF
# Our data are on a separated XFS FS
umount -at xfs
# You need room under /usr/mondo
rm -fr /usr/mondo/*
mkdir -p /usr/mondo/images /usr/mondo/tmp /usr/mondo/scratch
/usr/sbin/mondoarchive -O -i -H -N -g -d /usr/mondo/images -T
/usr/mondo/tmp -S /usr/mondo/scratch -E /video -s 4300m
EOF
```

Add the following line to /etc/modules.conf file (if your target system uses e.g. a SATA controller managed by the ata_piix module) :

```
# echo 'alias scsi_hostadapter ata_piix' >> /etc/modules.conf
```

Or use the method below with the FORCE_MODS variable in /usr/sbin/mindi script file.

Note: this option is not available before v1.2.3!

Ensure that you'll have enough space to store the mondo images and temporary files generated and launch the script:

```
# /usr/local/bin/mondo-bck
```

Then, when mondo has finished the backup, burn the images on DVD (in the script above, the image size is configured for 4.3GB).

Backup onto NFS share

If you want to do a P2V (Physical to virtual migration), Virtual being a VMWare Virtual Machine, you will need to modify `mindy` (located in `/usr/sbin/mindy`) to specify the modules that will be required by the new hardware on boot. Change the `FORCE_MODS` line to match:

```
FORCE_MODS="diskdump lib diskdump mptbase mptscsi mptfc mptspi mptsas mptscsih ata_piix crc32
mii pnet32 sunrpc nfs nfs_acl lockd loop"
```

Modify the path and run the script `mondoscript.sh` (you can find it in the folder `/d2/apps/i386/mondo/scripts` on deploy server):

```
#!/bin/sh

# better to got those locally
TEMP_PATH=/tmp/backupmondo/temp
SCRATCH_PATH=/tmp/backupmondo/scratch

NFS_HOST=10.3.252.25
NFS_MOUNT=/dploy/mondo
MOUNT_PATH=/mnt/backupmondo

# Relative path that will be used locally and remotely on the NFS server
BCKUP_PATH=images/ims

if (! test `mount | grep $MOUNT_PATH | wc -l` -eq 0);
  then umount $MOUNT_PATH;
fi

if (! test -d $TEMP_PATH);
  then mkdir -p $TEMP_PATH;
fi
echo "Temp directory $TEMP_PATH created"

if (! test -d $SCRATCH_PATH);
  then mkdir -p $SCRATCH_PATH;
fi
echo "Scratch directory $SCRATCH_PATH created"

if (! test -d $MOUNT_PATH);
  then mkdir -p $MOUNT_PATH;
fi
echo "Mount directory $MOUNT_PATH created"

echo "Try to mount NFS point $NFS_HOST:$NFS_MOUNT on $MOUNT_PATH..."

mount 10.3.252.25:$NFS_MOUNT $MOUNT_PATH
echo "NFS point $NFS_HOST:$NFS_MOUNT mounted on $MOUNT_PATH"

if (! test -d $MOUNT_PATH/$BCKUP_PATH);
  then mkdir -p $MOUNT_PATH/$BCKUP_PATH;
fi
echo "Backup directory $MOUNT_PATH/$BCKUP_PATH created"

BCKUP_NAME=`hostname -s`-"`date +%Y%m%d%H%M`
echo "Backup name: "$BCKUP_NAME

CMD_LINE="mondoarchive -O -n $NFS_HOST:$NFS_MOUNT -p $BCKUP_NAME -N -d $BCKUP_PATH -s 4380m
-F -S $SCRATCH_PATH -T $TEMP_PATH -E /tmp/rhel40u1"
echo "Will run command: $CMD_LINE"

$CMD_LINE
echo "Mondorescue terminated"

umount $MOUNT_PATH
echo "Backup directory $MOUNT_PATH unmounted"

rm -rf $TEMP_PATH
echo "Temp directory $TEMP_PATH removed"

rm -rf $SCRATCH_PATH
```

```
echo "Scratch directory $SCRATCH_PATH removed"  
echo "Backup over"
```

Copy the initrd and the kernel needed to boot using PXE onto deploy server in the correct folder:

```
# mount -o loop /var/cache/mindi/mindi.iso /mnt/cdrom  
# scp /mnt/cdrom/isolinux/initrd.img 10.3.252.25:/d2/v1/mondo/i-3pf.img  
# scp /mnt/cdrom/isolinux/vmlinuz 10.3.252.25:/d2/v1/mondo/k-3pf  
# umount /mnt/cdrom
```

Deployment on target server

Restoring from Media

Insert the first media created by mondo rescue and start the server. It will boot automatically on the linux image (-H option of mondoarchive).

Restoring from network

To restore from the network, you need to use some services provided by the deployment server (10.a.b.c).

Update the DHCP config file (for instance /etc/dhcpd.d/10.x.y) by adding some line corresponding to your hardware:

```
host yourhostname {
    hardware ethernet 00:50:56:AD:11:11 # your server mac address
    fixed-address 10.x.y.z             # IP address for your server
}
```

Then, reload the DHCP configuration files:

```
# service dhcpd reload
```

You now have to modify the PXE configuration. In the folder /d2/v1/pxelinux.cfg, create a file using the following naming:

01-<your mac address in lower case dash-separated>
for instance: 01-00-50-56-ad-11-11

In that file, copy the following:

```
default local
prompt 20
timeout 600
label local
    localboot 0
label ims-3pf
    kernel mondo/k-3pf
    append initrd=mondo/i-3pf.img load_rmdisk=1 prompt_ramdisk=0
    ramdisk_size=36864 rw root=/dev/ram iso_mode nuke devfs=nomount
    exec-shield=0 pxe prefix=3pf-2007060610
    nfsmount=10.a.b.c:/dploy/mondo/images/ims ipconf=eth0:dhcp
```

IMPORTANT:

If you are restoring the image on an ATCA you are likely needed to add 'console=ttyS0,115200' option on the "append initrd" line (check the speed you want! 115200bds is a sample).

Note:

Do NOT include the '-x.iso' (with x a number) in the 'prefix=' option

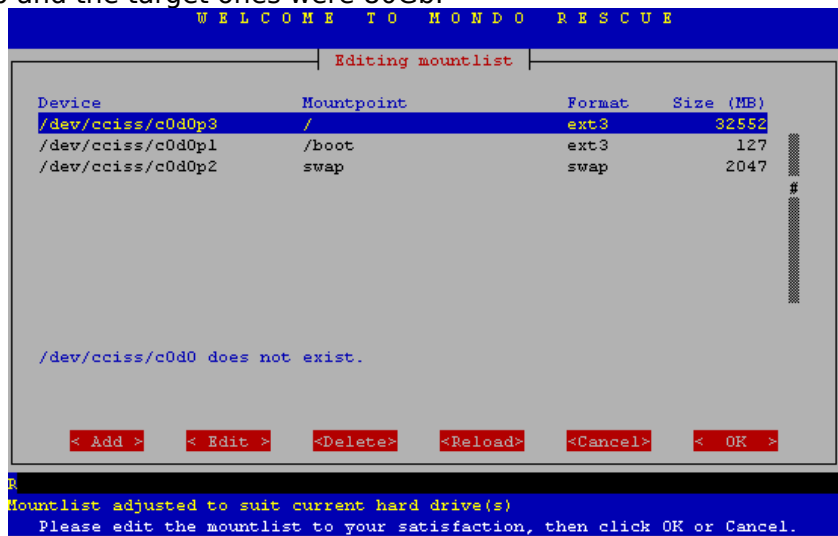
Boot your target system on the network (press F12 at boot or change boot order in Bios).

Mondorestore part

After few minutes, the software detects that the SCSI hardware is not present and that it can not mount the correct devices. It proposes to switch to interactive mode (answer Yes).



Edit the mountlist and change the devices to reflect the current hardware (going from /dev/cciss/c0d0p to /dev/sda or /dev/sdb for the SATA controller). In this example, the source disks were 36Gb and the target ones were 80Gb.



NOTE for LVM:

If ever you restore LVM volumes on a different HW (from smartarray /dev/cciss/c0d0p to /dev/sda for example), please check status and/or workaround of the bug related to the i-want-my-lvm file ([bug #183](#)).

Press OK and answer Yes twice. It will repartitioned the disks and format them.

```
WELCOME TO MONDO RESCUE
Restoring Interactively

Partitioning hard drives Done.
Formatting partitions
  Formatting partitions
    I am now formatting your hard disk partitions.
    This may take up to five minutes.
    Formatting /dev/sdal as ext3 _
    1% done          99% to go
  I'm thinking...
  Recovered mondo-restore.cfg
  Mountlist adjusted to suit current hard drive(s)
  I'm thinking...
  Recovered mondo-restore.cfg
  Mountlist adjusted to suit current hard drive(s)
  I am now formatting your hard disk partitions.
RC4 Secured (128 Bit)
```

Mondo rescue will ask you for restoring the data on the new system. Answer Yes. At the end, it asks you to initialize the boot loader:

```
WELCOME TO MONDO RESCUE
Restoring Interactively

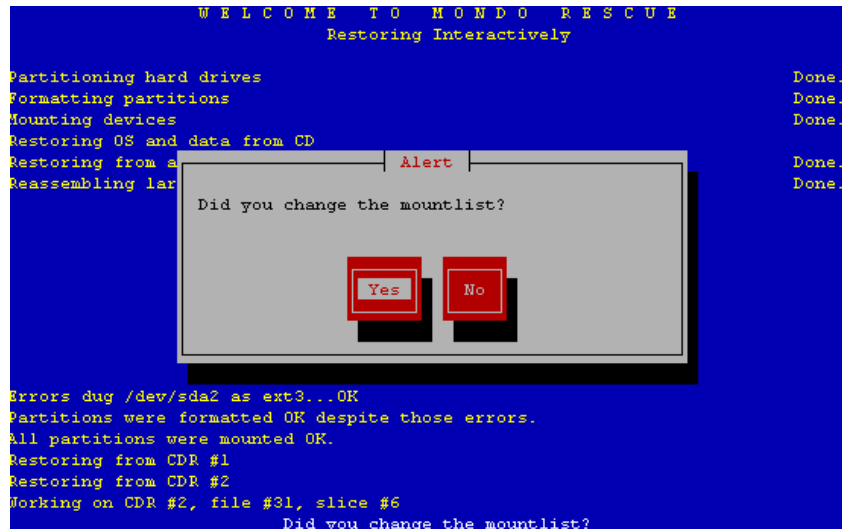
Partitioning hard drives Done.
Formatting partitions Done.
Mounting devices Done.
Restoring OS and data from CD Done.
Restoring from a CD Done.
Reassembling lar Done.

Alert
Initialize the boot loader?
  Yes No

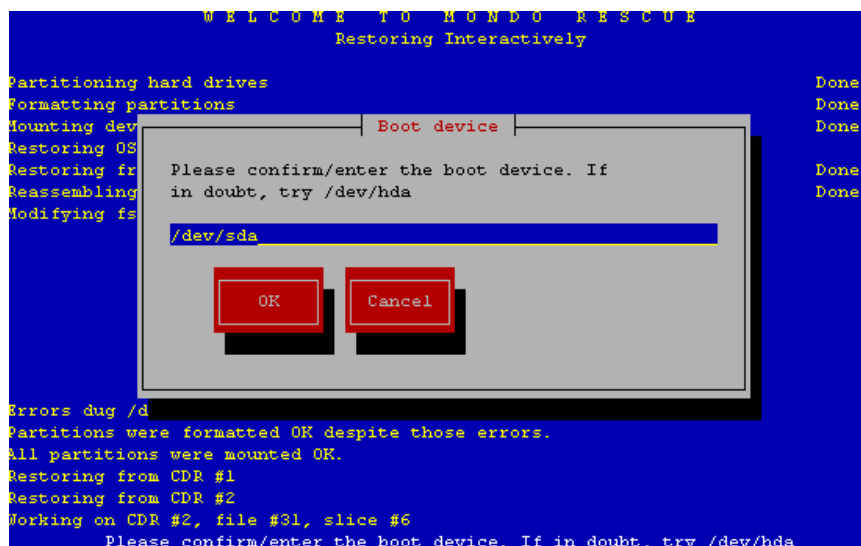
Errors dug /dev/sda2 as ext3...OK
Partitions were formatted OK despite those errors.
All partitions were mounted OK.
Restoring from CDR #1
Restoring from CDR #2
Working on CDR #2, file #31, slice #6
Initialize the boot loader?
```

Answer Yes.

Then you'll have to change the mountlist. Answer again Yes:



Change the mountlist and point it to /dev/sda (in case you have restored the image on /dev/sda):



Then it will ask to review the /etc/fstab, /etc/grub.conf files before rebooting the server. For fstab, you should change the file to reflect the change of hard drive reference (in our case from /dev/cciss/c0d0p? To /dev/sda?) as in the lines below. If you are using Label on drives, you can stay with it, Mondo will handle it properly.

IMPORTANT: If you're using LVM, check that your fstab contains no LABEL= line for Logical Volumes partitions (Cf [bug #185](#)). Change them for device name instead if there are some.

You have to modify the `/etc/inittab` file and add the following line at the end:

```
#S0:2345:respawn:/sbin/mingetty ttyS0
```

Then, in the `/etc/securetty` file, add at the end of the file a line with `ttyS0`. Finally, enter the following command :

```
# /sbin/mingetty ttyS0
```

You can now exit the chroot environment, umount the volumes and reboot:

```
# umount /proc
# umount /sys
# exit (CTRL+D)
# umount /mnt/sysimage/boot
(optional) # umount /mnt/sysimage/usr
# umount /mnt/sysimage
# reboot
```

After the reboot of the server, you might have to go into some reconfiguration process (such as `kudzu` on redhat).