Lab – MondoRescue

Lab Contents

This lab purpose is to install and use MondoRescue to produce a golden image of a server and restore it.

Lab Writer and Trainer

Bruno.Cornec@hp.com

Table of content

Overview of the MondoRescue Lab

Objectives

At the end of the Lab students should be able to install MondoRescue, use mondoarchive to create a golden image of a server and boot from it to res tore with mondorestore on a target..

Expected duration : 100 minutes

Reference documents

When dealing with the installation and configuration of MondoRescue, the first approach is to look at some reference documents:

The **HOWTO** at http://www.mondorescue.org/docs.shtml as well as **man pages**

The **Wiki** at http://trac.mondorescue.org/wiki and especially the **FAQ** at http://trac.mondorescue.org/wiki/FAQ

Most of the documentation is provided currently under a wiki format, so contributors are welcome and expected, including those finding issues in that training material. Trac also provides a bug report/enhancement request feature ;-)

A mailing-list dealing with MondoRescue is available at http://sourceforge.net/mailarchive/forum.php?forum_name=mondo-devel

Estimated time for the lab is placed in front of each part.

Environment setup

Estimated time: 10 minutes

MondoRescue installation

Mondorescue is available externaly from ftp://ftp.mondorescue.org/ and internally on the HP LAN from http://eurolinux.grenoble.hp.com/pub/linux/mondorescue/ or using ftp as well from the same place.

Version 2.2.9.1 is the current stable release. Version 2.2.10 is the development version

Here is the table of corresponding versions:

Tools	*MondoRescue 2.2.9.1*	*MondoRescue 2.2.10*
mondo	2.2.9.1	2.2.10.0
mindi	2.0.7.1	2.0.8
mindi-busybox	1.7.3	-

You may want to use yum to do the installation. For that, download the file under `rhel/5` named `mondorescue.repo` and place it under your `/etc/yum.repos.d` directory. Of course, you can use rpm and deal with the dependencies yourself ;-)

You'll also need for RHEL additional packages not provided in the base distribution, but which are brought by RPMforge (or EPEL). So create the following file:

```
# cd /etc/yum.repos.d/

# wget
ftp://eurolinux.gre.hp.com/pub/linux/mondorescue/rhel/5/mondorescue.repo

# cat > /etc/yum.repos.d/rpmforge.repo << EOF

[rpmforge]

name = Red Hat Enterprise $releasever - RPMforge.net - dag
```

```
baseurl =
ftp://eurolinux.gre.hp.com/pub/linux/rpmforge/redhat/el5/en/x86_64/dag

enabled = 1

gpgcheck = 0

EOF
```

Then use `yum` to install the application with all its dependencies and check that the correct version is installed and operational:

```
# yum install mondo.x86_64

# man mondoarchive
```

From now on, the Lab is "à la carte" ;-) Which means you cant take one of the following part at will and try to do it, once you've done at least the Disaster Recovery Set Creation. Estimated time is placed in front of each part.

I recommend that if you don't already know mondorescue you do at least the PXE as it's the simplest and quickest. The iLO is simple as well but longer. Then look at P2V, SSSTK. Those of you feeling more adventurous, or already knowing the tool, may want to help debug 2.2.10 and use it in their Lab. Bug reports welcome ;-)

Bare metal Disaster Recovery with Mondo Rescue

Estimated time: 20 minutes – Based on mondo 2.2.9.1, current version.

Disaster Recovery (DR) Set Creation

The goal of this part is to make a DR Set of media under the format of ISO images of 4GB of your server. Use for that simply the `mondoarchive` command:

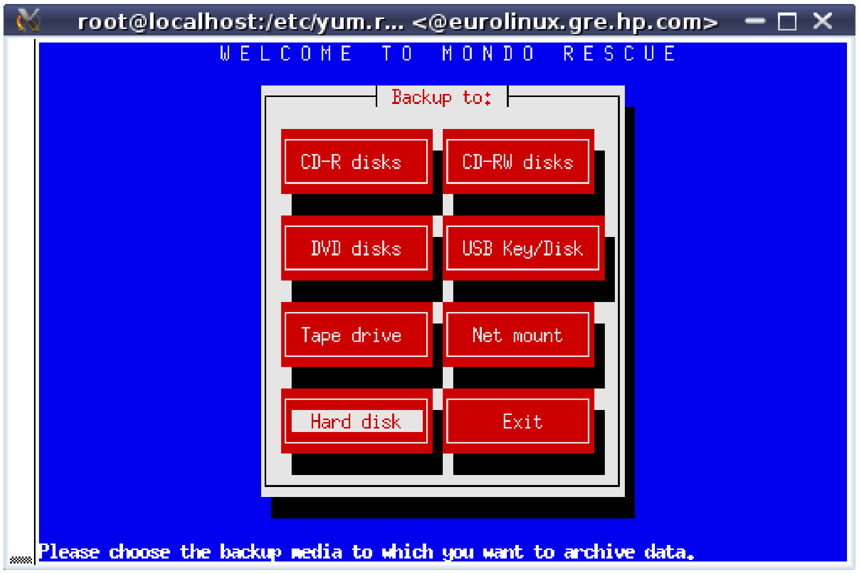

WELCOME TO MONDO RESCUE

Please choose the backup media to which you want to archive data.

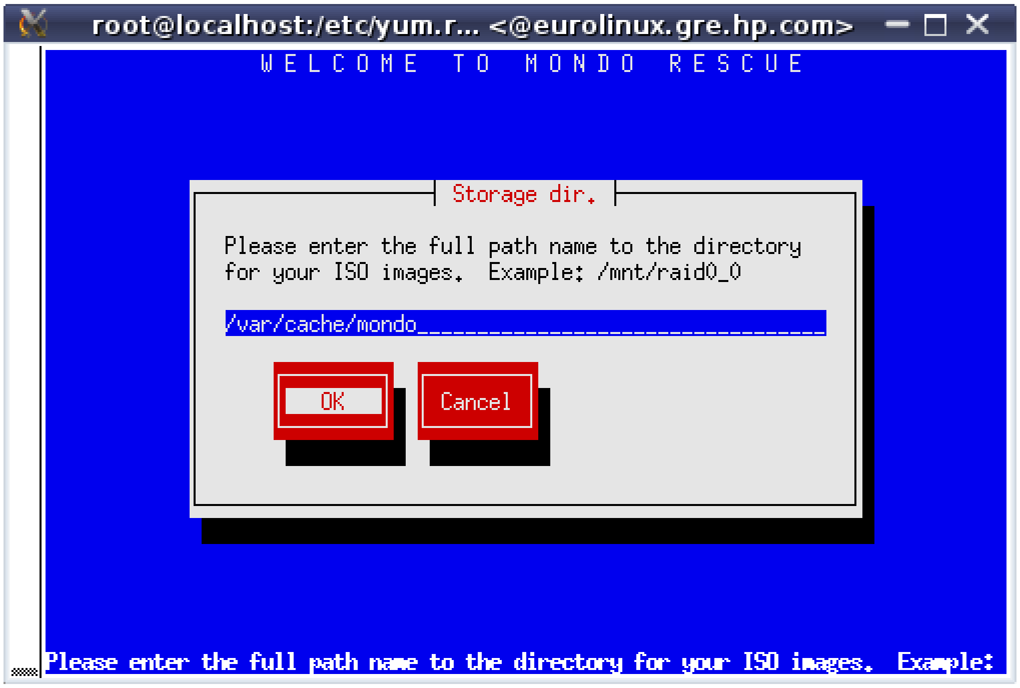

root@localhost:/etc/yum.r... <@eurolinux.gre.hp.com>

WELCOME TO MONDO RESCUE

Storage dir.

Please enter the full path name to the directory
for your ISO images. Example: /mnt/raid0_0

/var/cache/mondo

Please enter the full path name to the directory for your ISO images. Example:

That's where you ISO image will be created at the end. Note that mondo needs room to make that image, which is currently taken were there is the largest place available.

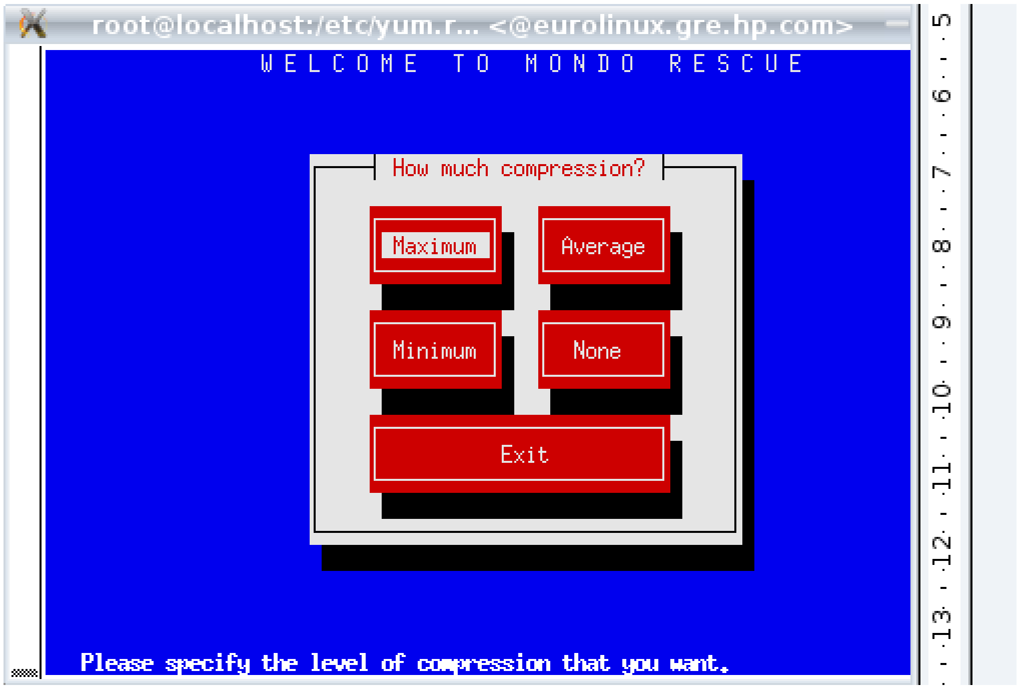

root@localhost:/etc/yum.r... <@eurolinux.gre.hp.com>

WELCOME TO MONDO RESCUE

How much compression?

Please specify the level of compression that you want.

Compression is done by default using bzip2 with that version. Another compressor can only be chosen when using the CLI and options -G (gzip) or -L (lzo) or -Y (lzma) in 2.2.10.

Specify the size of your image (could be larger than a physical media if you do not intend to really burn it).

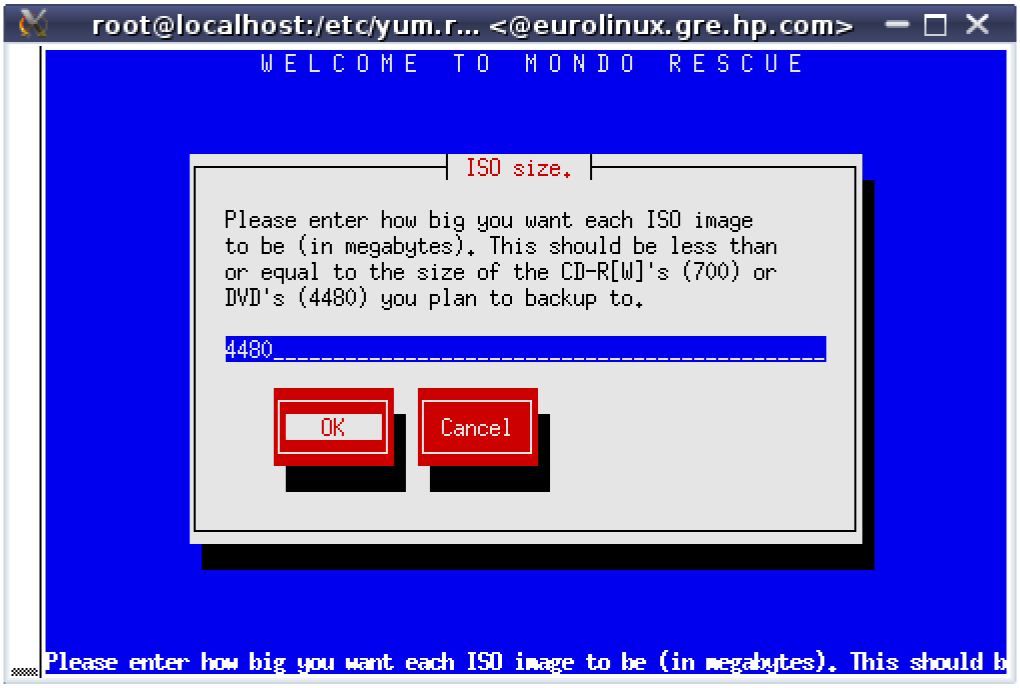

```
WELCOME TO MONDO RESCUE

ISO size.

Please enter how big you want each ISO image
to be (in megabytes). This should be less than
or equal to the size of the CD-R[W]'s (700) or
DVD's (4480) you plan to backup to.

4480

Please enter how big you want each ISO image to be (in megabytes). This should b
```

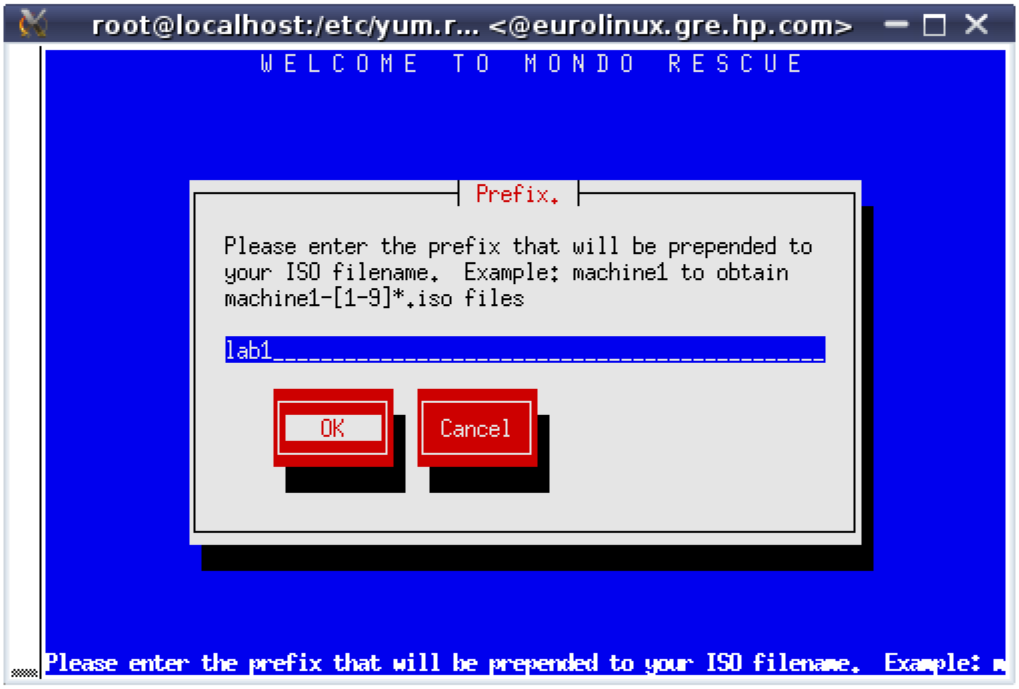

```
WELCOME TO MONDO RESCUE

Prefix.

Please enter the prefix that will be prepended to
your ISO filename.  Example: machine1 to obtain
machine1-[1-9]*.iso files

lab1

Please enter the prefix that will be prepended to your ISO filename.  Example: m
```

The prefix name allows you to keep multiple images in a central repository for example.

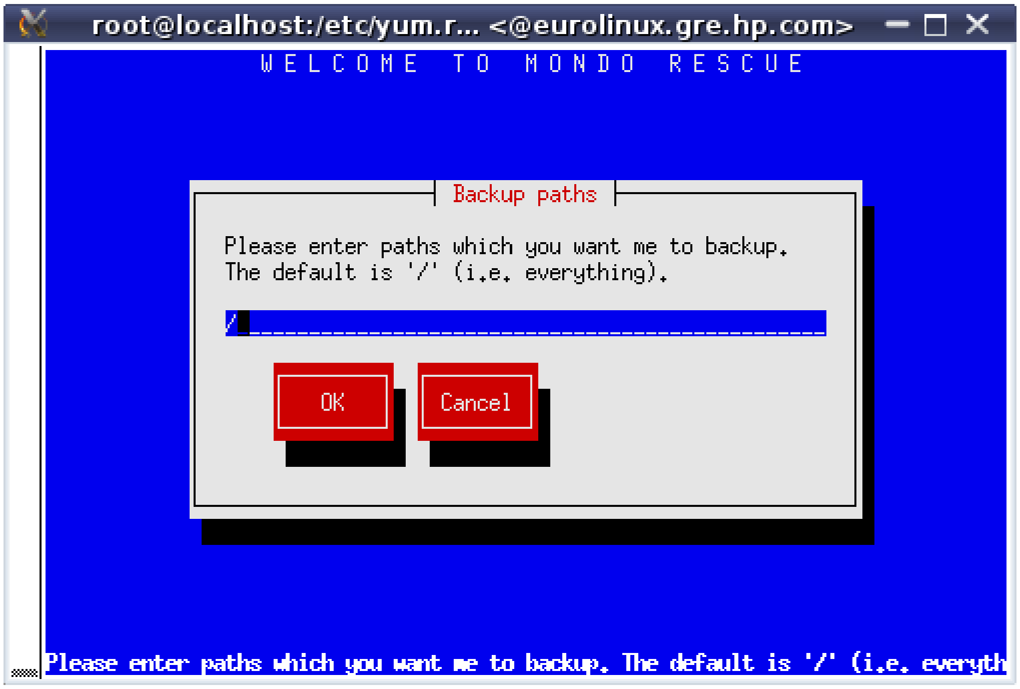

```
Please enter paths which you want me to backup.
The default is '/' (i.e. everything).

/

Please enter paths which you want me to backup. The default is '/' (i.e. everyth
```

You have to use / for a disaster recovery solution for your include dir.

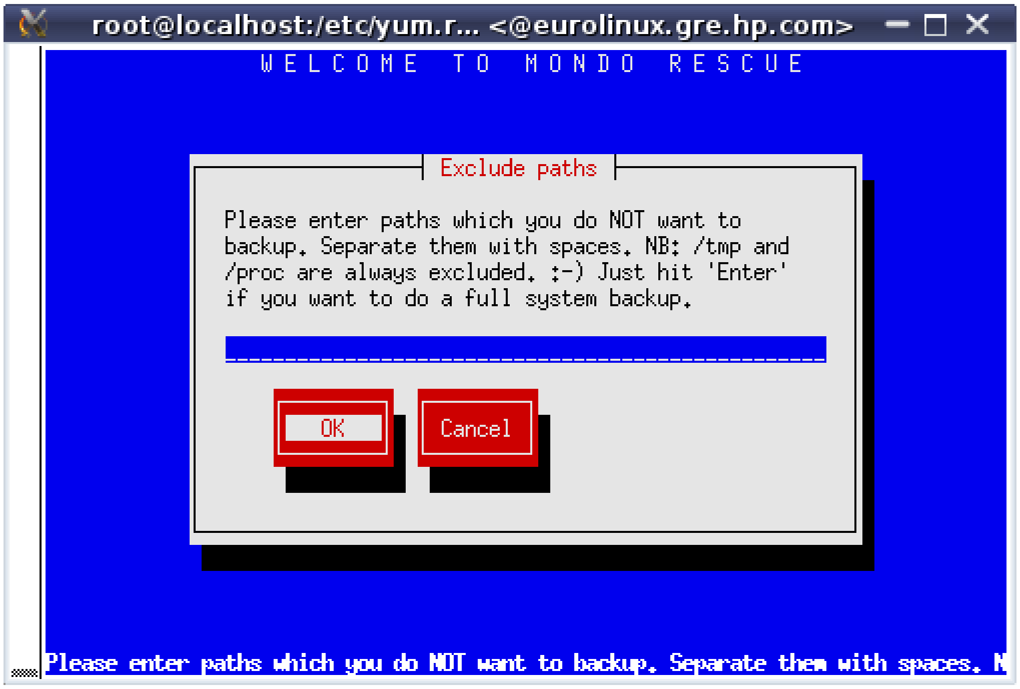

```
Please enter paths which you do NOT want to
backup. Separate them with spaces. NB: /tmp and
/proc are always excluded. :-) Just hit 'Enter'
if you want to do a full system backup.

Please enter paths which you do NOT want to backup. Separate them with spaces. N
```

In order to have a quicker backup, you may want to exclude some useless path (they may be backed up with another tool, which is often the case of data or databases). Advise here is to not exclude what would break the coherency of your package system.

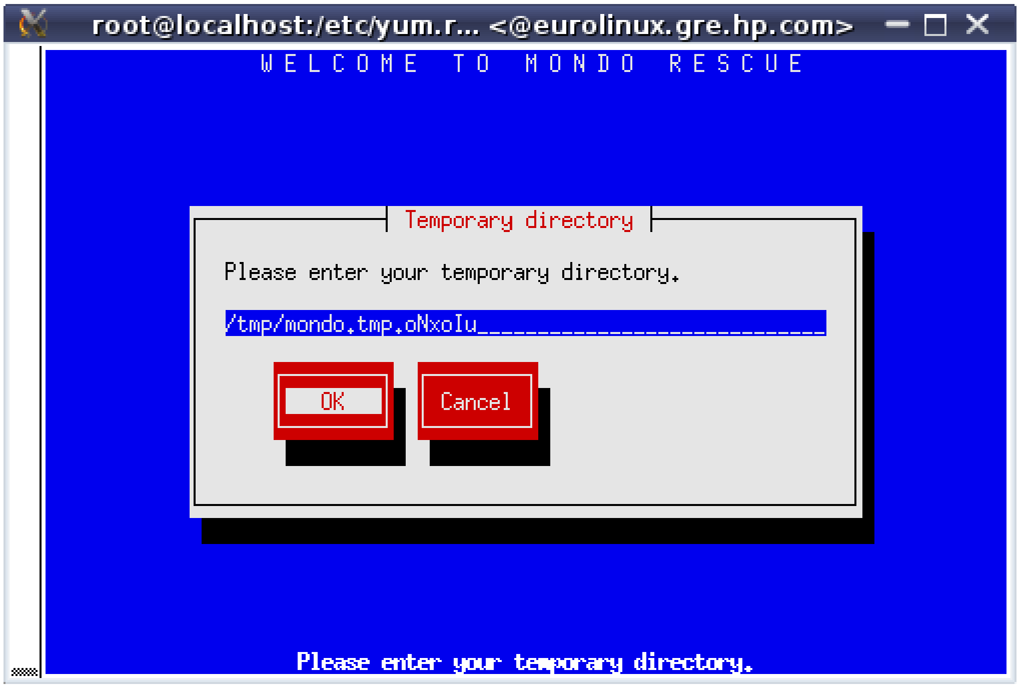

```
Please enter your temporary directory.
/tmp/mondo.tmp.oNxoIu
```

Please enter your temporary directory.

You may now precise where you want your temporary files to be generated (file system having the most place is the default) They are usually not big in size

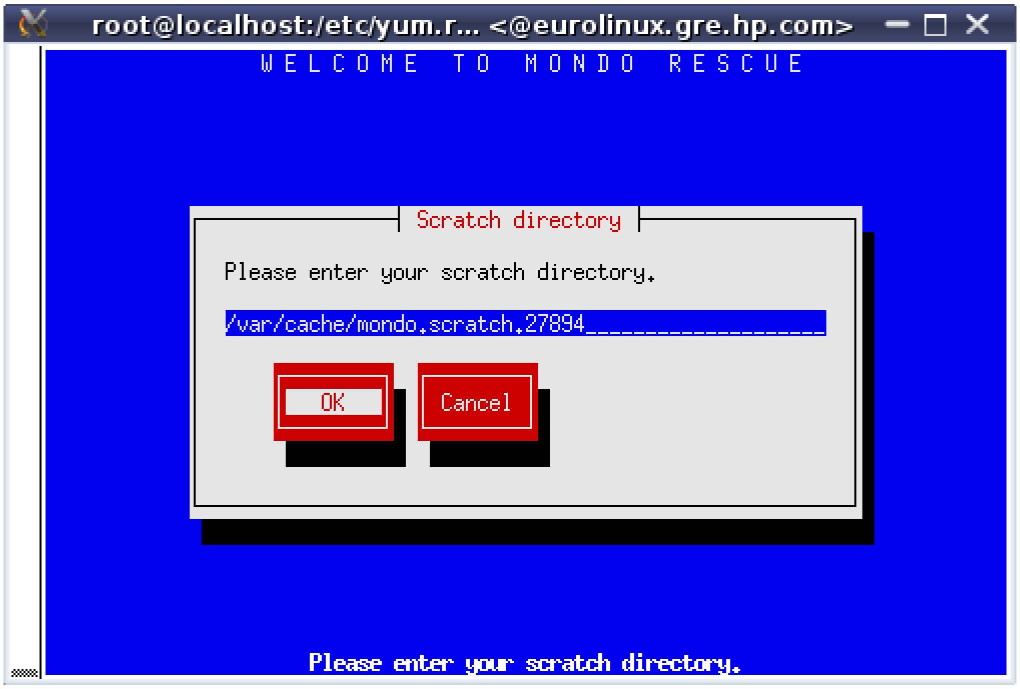

```
Please enter your scratch directory.
/var/cache/mondo.scratch.27894
```

Please enter your scratch directory.

.

You may now precise where you want your scratch files to be generated (file system having the most place is the default). They will occupy the size of a media so choose with care.

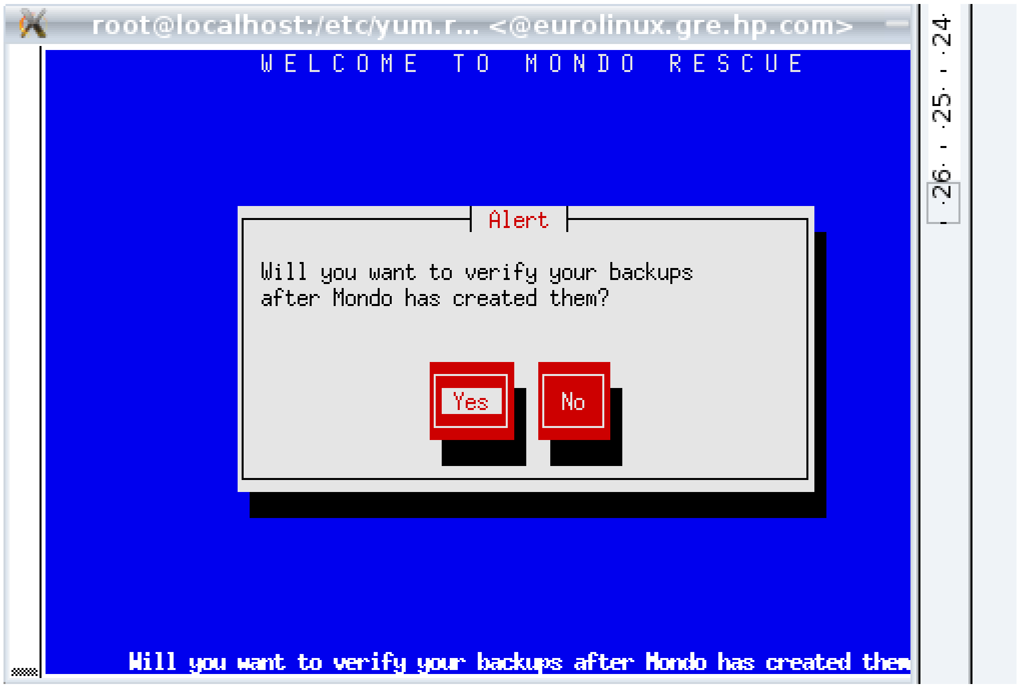

Will you want to verify your backups after Mondo has created them?

Will you want to verify your backups after Mondo has created them

At the end of the archiving run, mondo may optionaly verify all the sets created to check their consistency. Of course it takes more time.

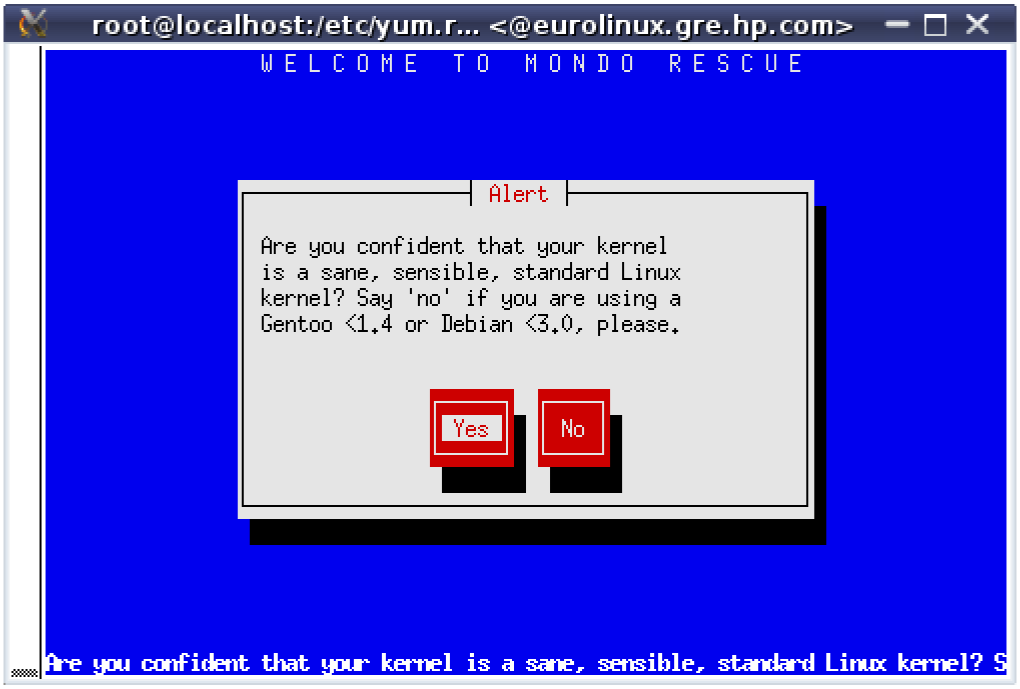

Are you confident that your kernel is a sane, sensible, standard Linux kernel? Say 'no' if you are using a Gentoo <1.4 or Debian <3.0, please.

Are you confident that your kernel is a sane, sensible, standard Linux kernel? S

With recent versions of mondo you should alwys say yes here. Even with a Gentoo or old Debian distribution. This message will disappear in 2.2.10.

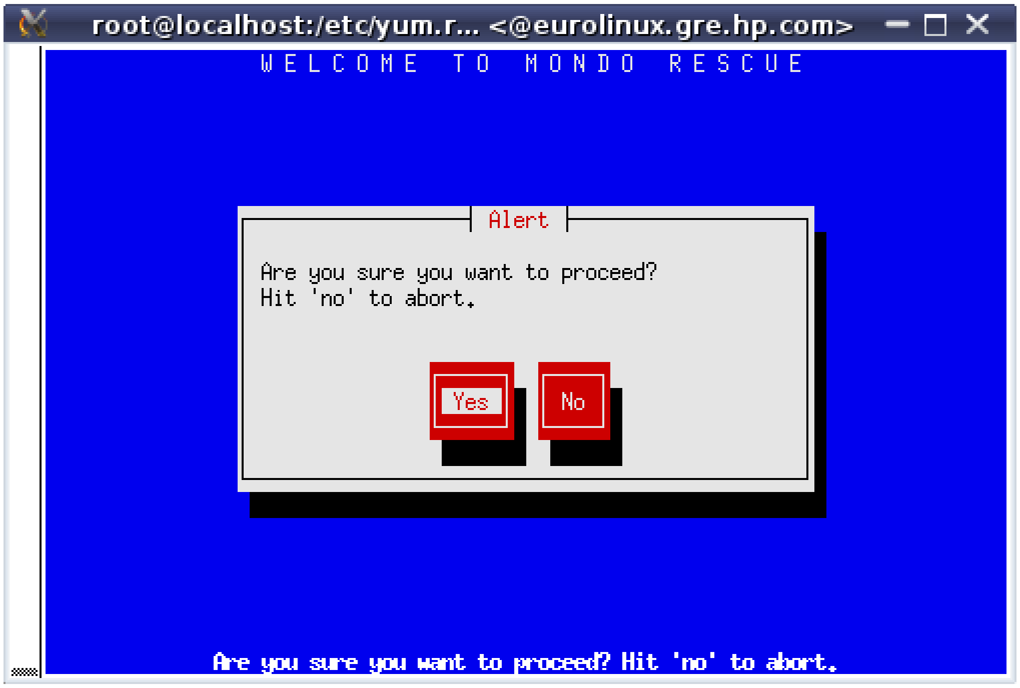

```
                 W E L C O M E   T O   M O N D O   R E S C U E

                   Are you sure you want to proceed?
                   Hit 'no' to abort.

              Are you sure you want to proceed? Hit 'no' to abort.
```

Fire !

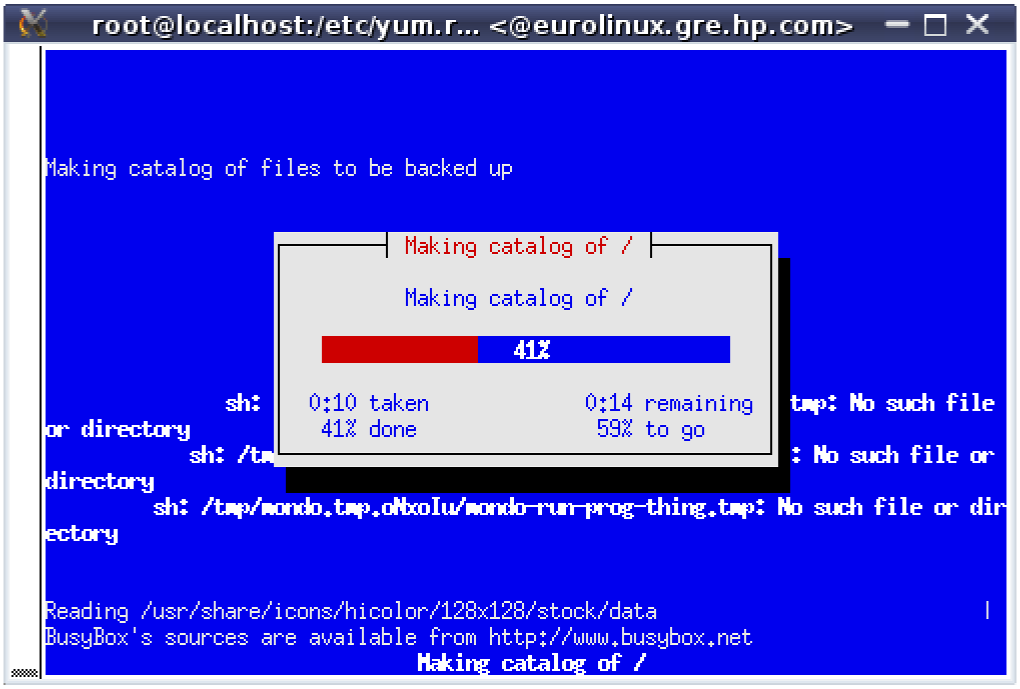

```
Making catalog of files to be backed up

                         Making catalog of /

                         Making catalog of /

            sh:   0:10 taken            0:14 remaining   tmp: No such file
or directory       41% done               59% to go
          sh: /t                                        : No such file or
directory
       sh: /tmp/mondo.tmp.oNxolu/mondo-run-prog-thing.tmp: No such file or dir
ectory

Reading /usr/share/icons/hicolor/128x128/stock/data                          |
BusyBox's sources are available from http://www.busybox.net
                         Making catalog of /
```

Mondo first creates a catalog of all files to handle.

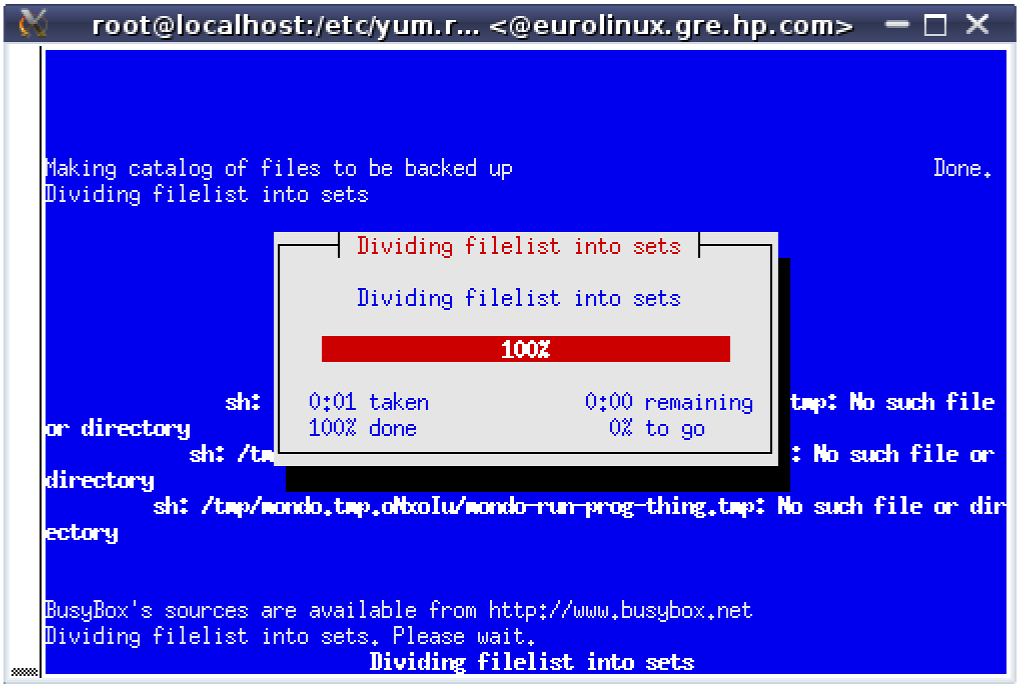

```
Making catalog of files to be backed up                                Done.
Dividing filelist into sets

                  ┤ Dividing filelist into sets ├

                    Dividing filelist into sets

                              100%

              sh:   0:01 taken            0:00 remaining  tmp: No such file
or directory        100% done               0% to go
           sh: /tm                                        : No such file or
directory
         sh: /tmp/mondo.tmp.oNxolu/mondo-run-prog-thing.tmp: No such file or dir
ectory


BusyBox's sources are available from http://www.busybox.net
Dividing filelist into sets. Please wait.
                           Dividing filelist into sets
```

Then it splits those files in various sets to be created.

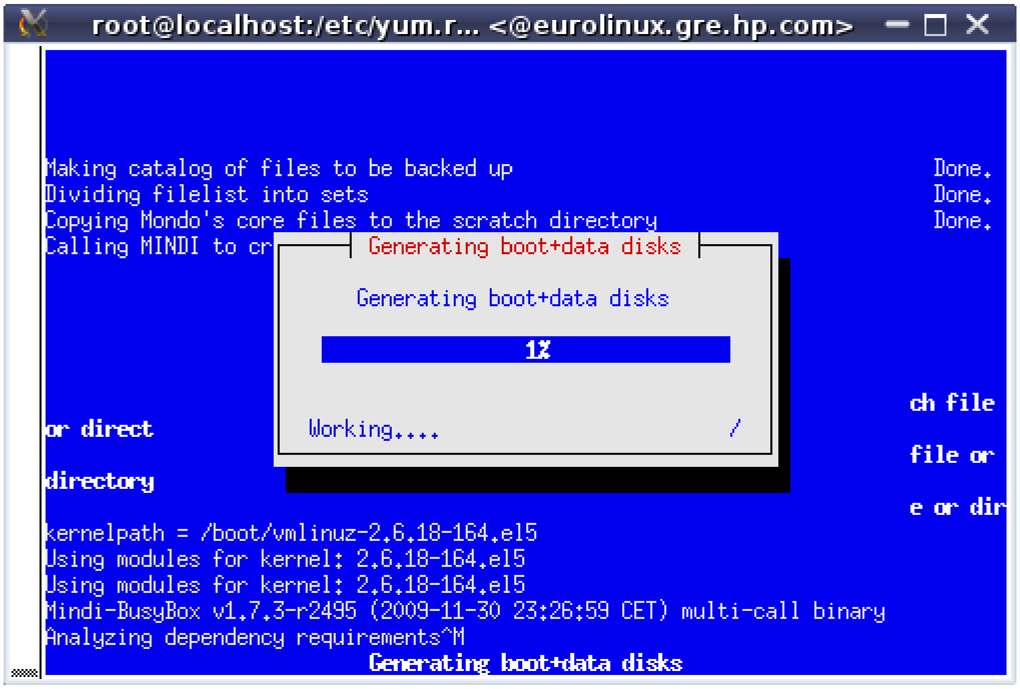

```
Making catalog of files to be backed up                                Done.
Dividing filelist into sets                                            Done.
Copying Mondo's core files to the scratch directory                    Done.
Calling MINDI to cr ┤ Generating boot+data disks ├

                      Generating boot+data disks

                                  1%

                                                                   ch file
or direct           Working....                         /
                                                                   file or
directory
                                                                   e or dir
kernelpath = /boot/vmlinuz-2.6.18-164.el5
Using modules for kernel: 2.6.18-164.el5
Using modules for kernel: 2.6.18-164.el5
Mindi-BusyBox v1.7.3-r2495 (2009-11-30 23:26:59 CET) multi-call binary
Analyzing dependency requirements^M
                           Generating boot+data disks
```

Then mondo calls mindi which generates the boot environment which will be used at restore time to prepare the machine (modules loading, device creation, partitioning, lvm management, ...)

Now it's time to create all the DR sets.

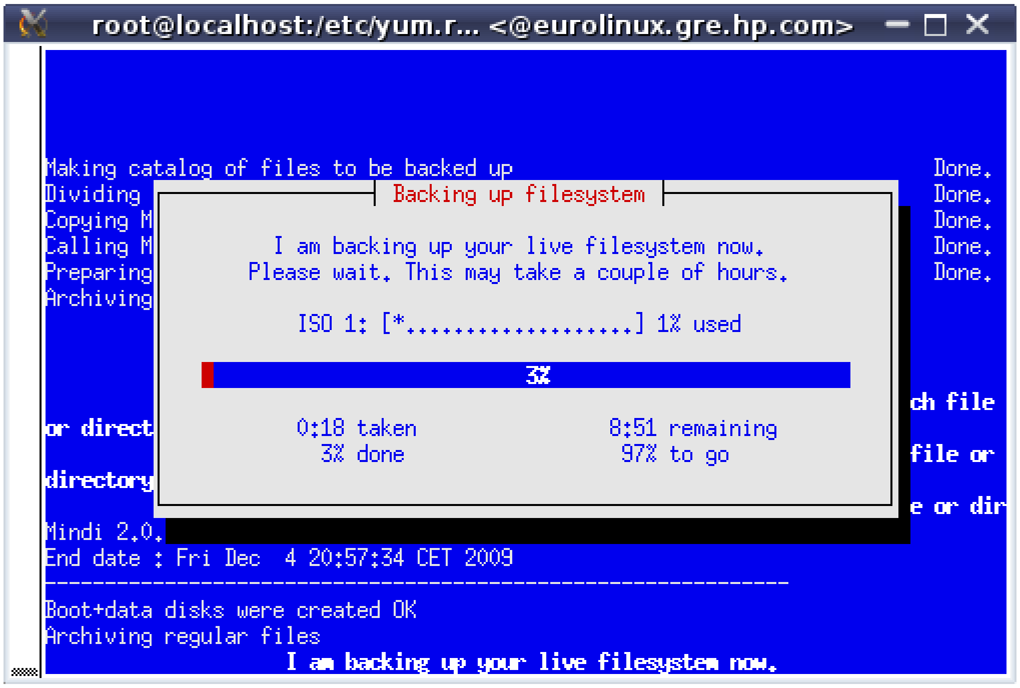

```
Making catalog of files to be backed up                          Done.
Dividing ┌───────────────┤ Backing up filesystem ├────────────────┐ Done.
Copying M│                                                        │ Done.
Calling M│          I am backing up your live filesystem now.     │ Done.
Preparing│       Please wait. This may take a couple of hours.    │ Done.
Archiving│                                                        │
         │           ISO 1: [*..................] 1% used         │
         │                                                        │
         │    ████████████████████████ 3% ███████████████████     │ch file
         │                                                        │
or direct│           0:18 taken                8:51 remaining      │file or
         │            3% done                   97% to go          │
directory│                                                        │e or dir
         └────────────────────────────────────────────────────────┘
Mindi 2.0.
End date : Fri Dec  4 20:57:34 CET 2009
--------------------------------------------------------------
Boot+data disks were created OK
Archiving regular files
                    I am backing up your live filesystem now.
```

During that phase, you may want to have a look at the documentation on tools provided with this Lab.

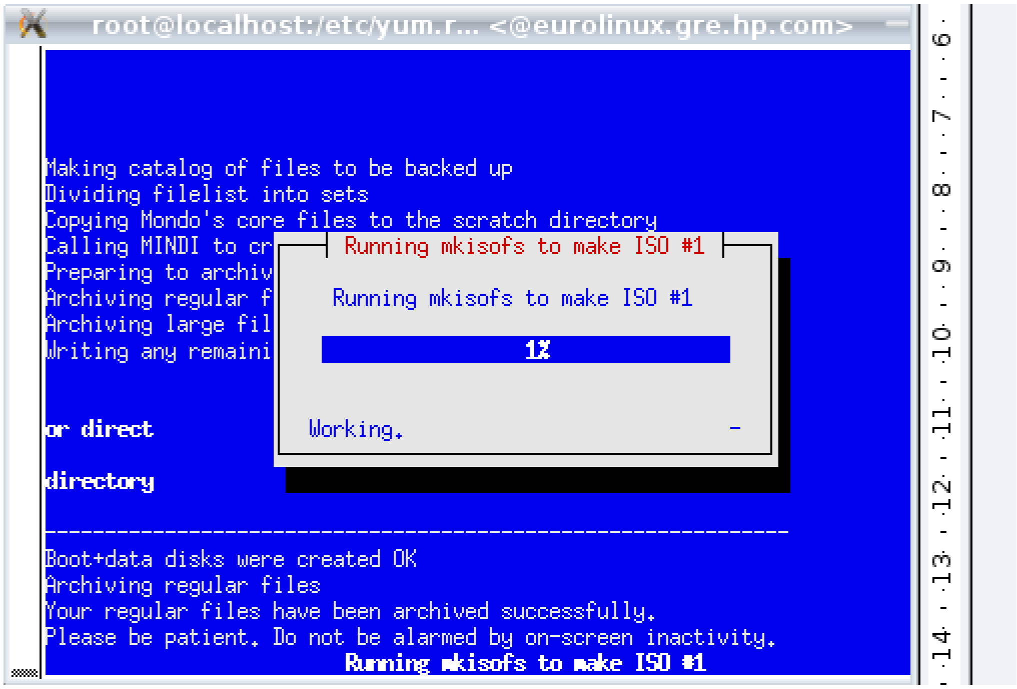

```
Making catalog of files to be backed up
Dividing filelist into sets
Copying Mondo's core files to the scratch directory
Calling MINDI to cr┌──┤ Running mkisofs to make ISO #1 ├──┐
Preparing to archiv│                                      │
Archiving regular f│   Running mkisofs to make ISO #1     │
Archiving large fil│                                      │
Writing any remaini│   ██████████████ 1% ██████████████   │
                   │                                      │
or direct          │ Working.                           - │
                   └──────────────────────────────────────┘
directory

--------------------------------------------------------------
Boot+data disks were created OK
Archiving regular files
Your regular files have been archived successfully.
Please be patient. Do not be alarmed by on-screen inactivity.
                         Running mkisofs to make ISO #1
```

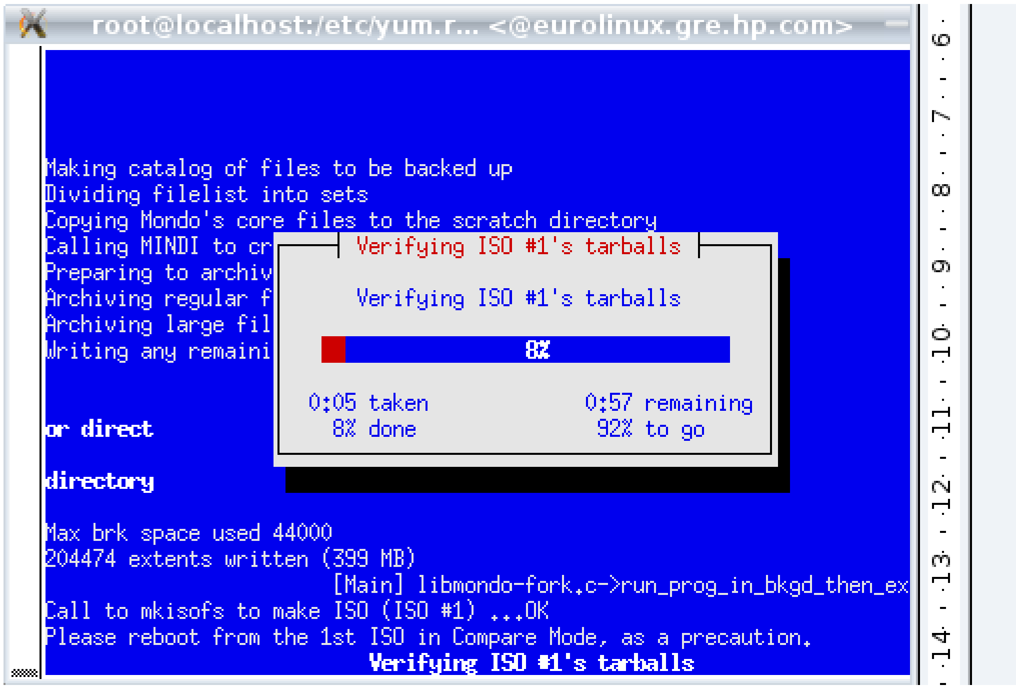

```
Making catalog of files to be backed up
Dividing filelist into sets
Copying Mondo's core files to the scratch directory
Calling MINDI to cr
Preparing to archiv
Archiving regular f
Archiving large fil
Writing any remaini

or direct

directory

Max brk space used 44000
204474 extents written (399 MB)
                        [Main] libmondo-fork.c->run_prog_in_bkgd_then_ex
Call to mkisofs to make ISO (ISO #1) ...OK
Please reboot from the 1st ISO in Compare Mode, as a precaution.
                           Verifying ISO #1's tarballs
```

```
Verifying ISO #1's tarballs

Verifying ISO #1's tarballs

8%

0:05 taken               0:57 remaining
  8% done                  92% to go
```

The verification of ISO is done upon previous request.

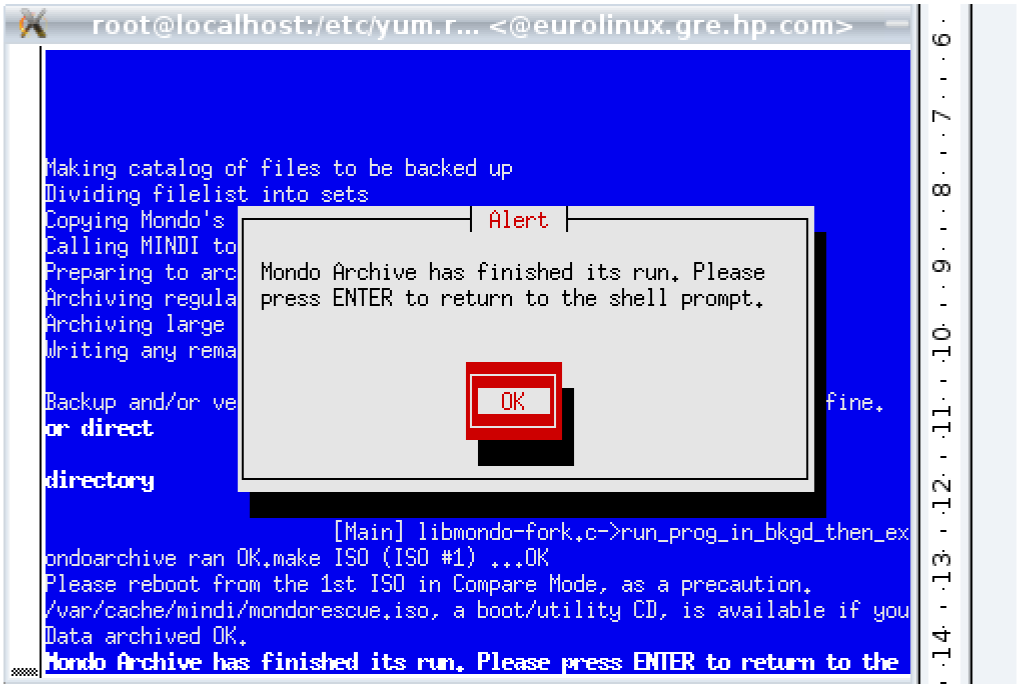

```
Making catalog of files to be backed up
Dividing filelist into sets
Copying Mondo's
Calling MINDI to
Preparing to arc
Archiving regula
Archiving large
Writing any rema

Backup and/or ve                                   fine.
or direct

directory

                        [Main] libmondo-fork.c->run_prog_in_bkgd_then_ex
ondoarchive ran OK.make ISO (ISO #1) ...OK
Please reboot from the 1st ISO in Compare Mode, as a precaution.
/var/cache/mindi/mondorescue.iso, a boot/utility CD, is available if you
Data archived OK.
Mondo Archive has finished its run. Please press ENTER to return to the
```

```
Alert

Mondo Archive has finished its run. Please
press ENTER to return to the shell prompt.

                 OK
```

That's all folks !

```
[root@localhost yum.repos.d]# ls -l /var/cache/mondo/
total 409368
```

```
-rw-r--r--  1 root root        11 Dec  4 20:55 difflevel.0
-rw-r--r--  1 root root 418762752 Dec  4 21:00 lab1-1.iso
```

Bare Metal Recovery on another system (with iLO Virtual Media)

Estimated time: 20 minutes

In order to have optimal throughput, use a browser on your system to access the OA of the target machine. Mount the produced ISO image with the virtual media of it, and restore your initial content on it.

You should have nothing special to do except accept it takes much more time than using a physical real media due to the USB emulation provided by the iLO virtual media.

You may want to try to change some parameters to see how mondorestore deals with it. You may remark that LVM modifications are not easy today, and require to edit the `/tmp/i-want-my-lvm` script. This is expected to be improved in January 2010.

Bare Metal Recovery in KVM/QEMU (aka P2V)

Estimated time: 30 minutes

For this part, in addition to the above mentioned documentation, please refer to the excellent manual provided by Lester Wade at http://www.mondorescue.org/docs/p2v.pdf.

Now you can try to see if you can re-install the image of your system just like that on a KVM Virtual Machine.

You first need some additional packages. Then you'll need to create an empty virtual machine. To do that derive the following command:

```
[root@localhost ~]# yum install kvm kvm-qemu-img
[root@localhost ~]# ln -sf /usr/libexec/qemu-kvm /usr/bin/kvm
[root@localhost ~]# modprobe kvm-intel
[root@localhost ~]# mkdir -p /home/vm
[root@localhost ~]# qemu-img create -f qcow2 /home/vm/img.qemu 130G
```

where 130 is the number of GB to allocate for the VM, which should correspond to the original disk size. The space will only increase as the VM is used.

You can now use the MondoRescue image to initiate your virtual machine with KVM. So type a command similar to the following:

```
[root@localhost ~]# kvm -cdrom /var/cache/mondo/lab1-1.iso -m 512 -boot
d /home/vm/img.qemu
```

It will launch the VM and boot on the CD which is associated with your MondoRescue image, thus beginning the installation. 512 is the size of the RAM in MB, that you may adapt at will.

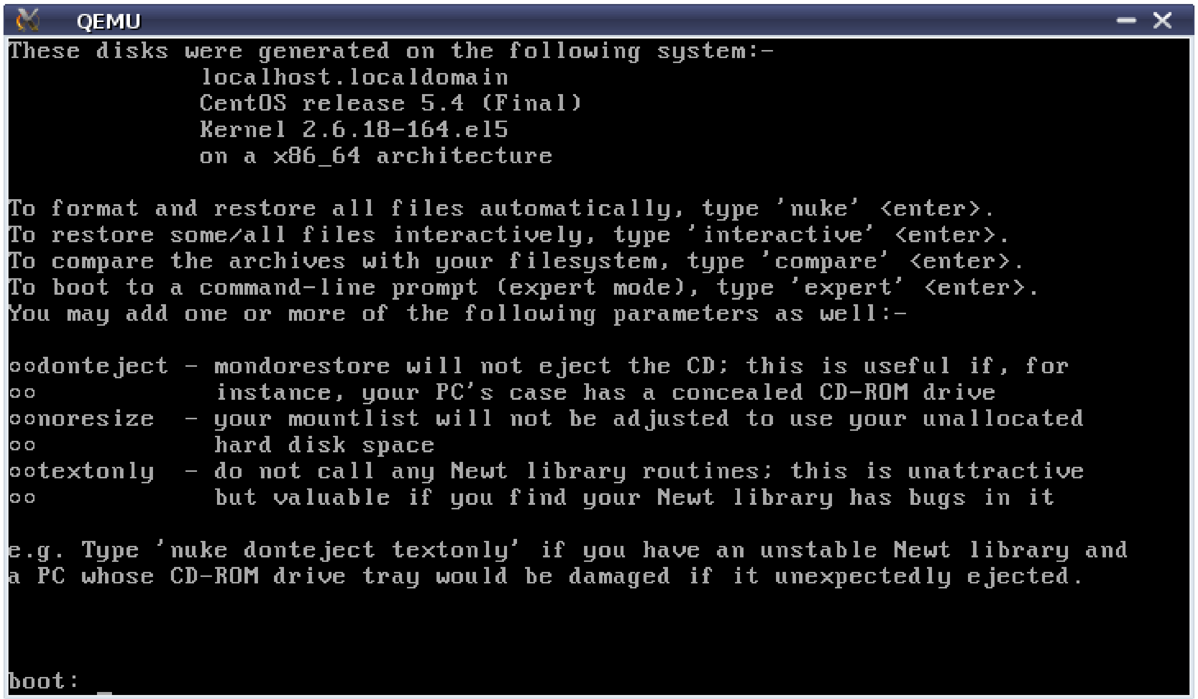

```
These disks were generated on the following system:-
             localhost.localdomain
             CentOS release 5.4 (Final)
             Kernel 2.6.18-164.el5
             on a x86_64 architecture

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: _
```

Choose `automatic` mode to restore. It should fail.

You'll have to adapt the mountlist as hard disk device names are different (`/dev/sda` => `/dev/hda`)

Then continue by trying to restore your content using that new mountlist.

However, it will still not work as you'll see, as we were using LVM and that LVM configuration uses the old device names as well, but can not be changed through the interface. So after the error message concerning the formating of the File Systems which aborts, change terminal (ALT-F2) and do the following:

:

```
# perl -pi -e 's/sda/hda/' /tmp/i-want-my-lvm

# perl -pi -e 's/sda/hda/' /tmp/mountlist.txt
```

Then relaunch `mondorestore` and this time use the `interactive` mode.

All those tasks have to be done, as we are changing fundamentally the underlying hardware of our system. However, the tool supports it, providing some adaptations. As seen earlier, the DR on the same hardware is strait forward. An enhancement request exists in trac to have a more automatic support of cloning on different hardware.

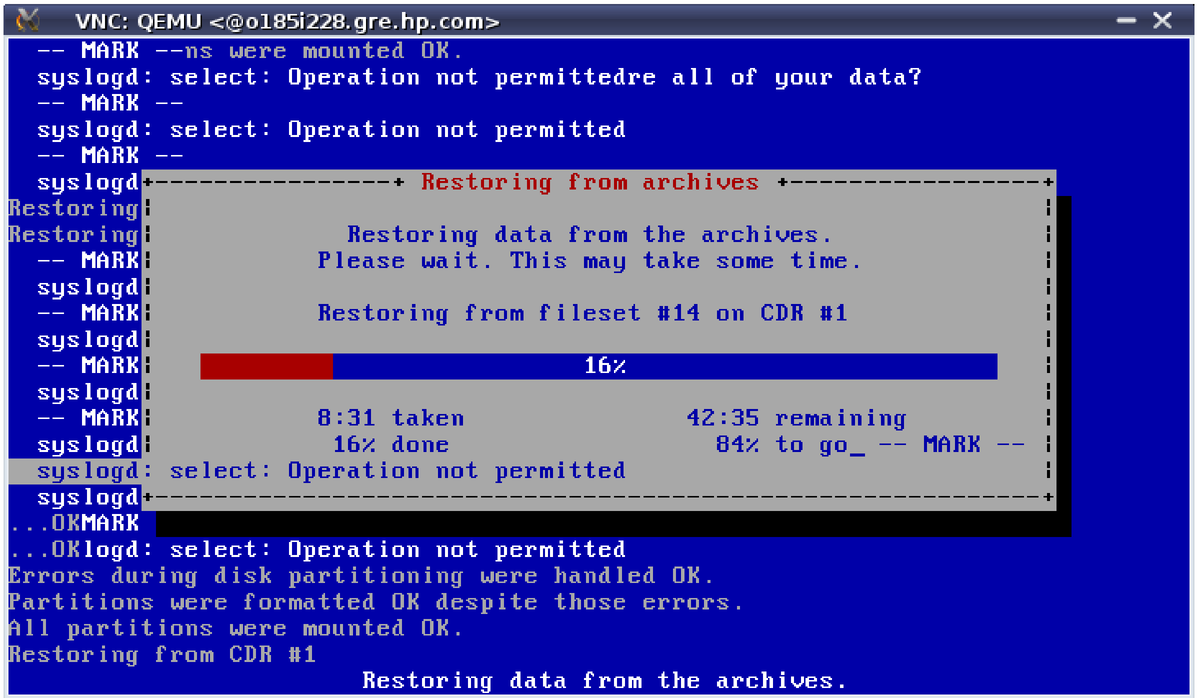

```
  -- MARK --ns were mounted OK.
  syslogd: select: Operation not permittedre all of your data?
  -- MARK --
  syslogd: select: Operation not permitted
  -- MARK --
  syslogd+----------------+ Restoring from archives +------------------+
Restoring|                                                               |
Restoring|          Restoring data from the archives.                    |
  -- MARK|        Please wait. This may take some time.                  |
  syslogd|                                                               |
  -- MARK|        Restoring from fileset #14 on CDR #1                   |
  syslogd|                                                               |
  -- MARK|                            16%                                |
  syslogd|                                                               |
  -- MARK|        8:31 taken                  42:35 remaining            |
  syslogd|         16% done                     84% to go_ -- MARK --    |
  syslogd: select: Operation not permitted                               |
  syslogd+---------------------------------------------------------------+
...OKMARK
...OKlogd: select: Operation not permitted
Errors during disk partitioning were handled OK.
Partitions were formatted OK despite those errors.
All partitions were mounted OK.
Restoring from CDR #1
                        Restoring data from the archives.
```

At the end of the restore process, you'll also have to adapt the disk used for the boot loader. Say yes when asked for a modified mountlist and review the content of the files proposed. Rebuild your initrd as proposed.

After that `mondorestore` will let you at the prompt. Review the log file `/var/log/mondorestore.log`, and check for potential issues at reboot time (especially boot loader).

Then close KVM/QEMU, and relaunch it without booting on the DVD:

```
[root@localhost ~]# kvm  -m 512  /home/vm/img.qemu
```

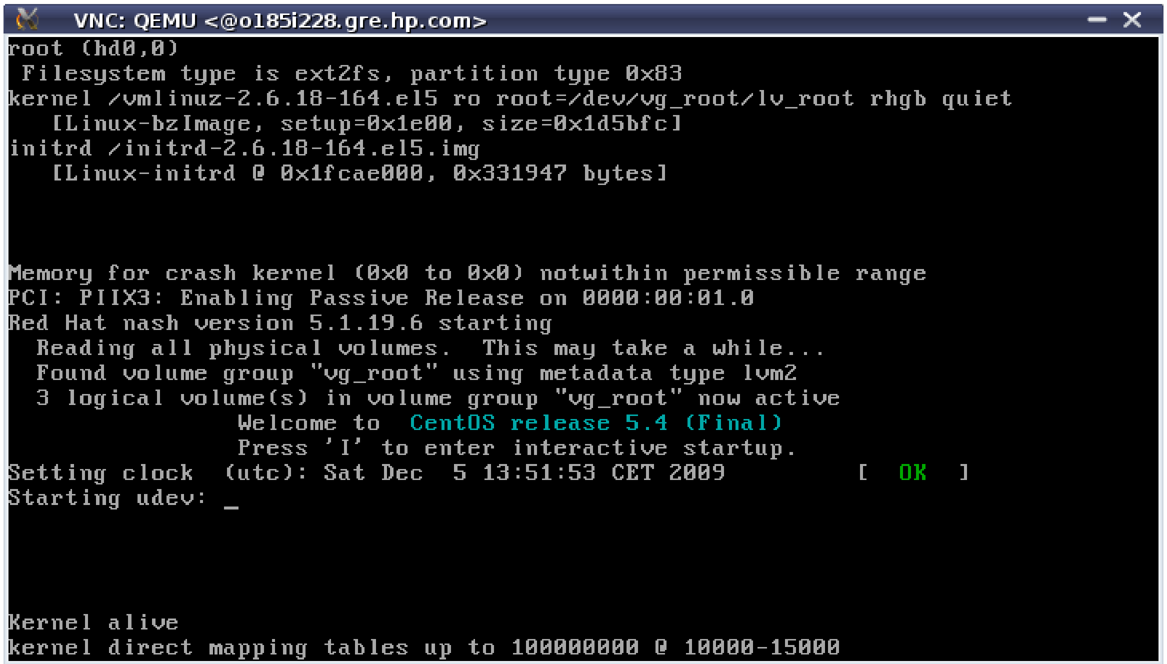

```
root (hd0,0)
 Filesystem type is ext2fs, partition type 0x83
kernel /vmlinuz-2.6.18-164.el5 ro root=/dev/vg_root/lv_root rhgb quiet
   [Linux-bzImage, setup=0x1e00, size=0x1d5bfc]
initrd /initrd-2.6.18-164.el5.img
   [Linux-initrd @ 0x1fcae000, 0x331947 bytes]



Memory for crash kernel (0x0 to 0x0) notwithin permissible range
PCI: PIIX3: Enabling Passive Release on 0000:00:01.0
Red Hat nash version 5.1.19.6 starting
  Reading all physical volumes.  This may take a while...
  Found volume group "vg_root" using metadata type lvm2
  3 logical volume(s) in volume group "vg_root" now active
                Welcome to  CentOS release 5.4 (Final)
                Press 'I' to enter interactive startup.
Setting clock  (utc): Sat Dec  5 13:51:53 CET 2009         [  OK  ]
Starting udev: _




Kernel alive
kernel direct mapping tables up to 100000000 @ 10000-15000
```

Network based Disaster Recovery

Estimated time: 30 minutes

Disaster Recovery (DR) Set Creation

You'll have to use the following script to realize your backup:

```
# mkdir -p /mondo

# cat > /usr/local/bin/postmondo << EOF

perl -pi -e 's/^HWADDR.*//' /etc/sysconfig/network-scripts/ifcfg-eth*

EOF

# mount -o rsize=32768,wsize=32768 eurolinux.gre.hp.com:/pub3/svg/mondo-
test /mondo

# mkdir -p /mondo/labn

# mondoarchive -O -d /labn -p labn -g -s 4380m -N -n
eurolinux.gre.hp.com:/pub3/svg/mondo-test -E /home/vm
```

Just replace lab*n* above by your lab number (1 to 8). And then launch the script which will do the backup through NFS on eurolinux.

Network based Bare Metal Recovery (with PXE)

In order to restore your target blade, you'll have to populate the PXE server on eurolinux. Entries like the following have already been declared:

```
label lab2

  kernel mondo/klab2

  append initrd=mondo/ilab2 nuke load_ramdisk=1 prompt_ramdisk=0
ramdisk_size=130000 rw root=/dev/ram iso_mode pxe ipconf=eth0:dhcp
prefix=lab2 nfsmount=eurolinux.gre.hp.com:/pub3/svg/mondo-test
ramdisk_blocksize=1024 post=/mnt/RESTORING/usr/local/bin/postmondo
serial=/dev/ttyS1 console=ttyS1
```

What remains to do is to create/update the kernel and initrd from your ISO image in the tftp directory. Just use the Web interface at http://eurolinux.gre.hp.com/cgi-bin/lab.cgi

then boot your target blade using F12 to activate PXE boot. At pxelinux boot prompt type lab*n* and check the status of your restore.

SSSTK usage with Mondo Rescue

Estimated time: 5 minutes – For information only – no suitable hardware available.

Disaster Recovery (DR) Set Creation

This time, before launching the backup, look at the file /usr/share/doc/mindi-2.0.7.1/README.bkphw in order to add SSSTK support to your environment. Do it only for the Smart Array Controller.

Solve the issues linked to that installation ;-) (Hint you may need to use yum search and install additional HP Insight Manager drivers. You may even need to launch some tools manually to determine the real issue)

Use `hpacucli` to change the stripe size of your source blade, so that it doesn't match the one on your destination blade.

Also check in mindi log file `/var/log/mindi.log`, that the hardware profile has been taken in account.

Bare Metal Recovery (with PXE)

You should be questioned about the possibility to restore the hardware profile of your original Blade, in the early stages of the restore process. Confirm. Then you should have nothing special to do.

During the reboot following the restore, check that the parameter of the blade has been changed back to the value of the one originally on your source blade.

There is much more to discover around MondoRescue (, …) but I leave that for an update of this document.