Elastix Application Note #201201091:

Elastix RAID Setup Step By Step Including Recovery

Title	Elastix Raid Setup Step By Step Including Recovery
Author	Bob Fryer
Date Document Written	9th January 2012
Date of Last Revision	20th January 2012
Revision	1.1
Replaces Document	N/A
Tested on Elastix Version	2.2
Backward Compatible	Yes
Elastix Level	Beginner to Experienced
Linux Level	Intermediate to Experienced
Network Level	N/A
Latest Document Source available from	www.elastixconnection.com
Credits (See Document History)	Shauw
Licence	GNU FDL

Contents

Foreword

These application notes are intended to be a guide to implement features or extend the features of the Elastix IP PBX system.

Whilst many (but not all) guides available are basically a random collection of notes, usually while someone is implementing a feature for themselves, these guides are meant to be more definitive guide that has been tested in a lab with specific equipment, and particular versions of Elastix.

Finding information on the Internet can be haphazard due to the lack of document version control, lack of attention to software versions, and in some cases they are wrong. Then you have the cross pollination issues, where a guide has been done for another distribution, which may or may not be applicable to your Elastix system.

You will note on the front page of every Application note written in this way, will be an easy to read summary, regarding the Elastix system it was tested on, when the document was written, whether it is backward compatible, and the level of expertise needed to accomplish the implementation.

These application notes are written up and tested in a lab that has been specially setup to write these notes. This includes

- 5 x Elastix IP PBX Hardware with a mixture of SIP only, Digium, Sangoma, OpenVox Cards
- 1 x WAN Simulator (including latency, jitter, random disconnects, random packet drop)
- 8 x Consumer / Business routers, including Drayteks, Cisco 1842, Cisco 877, Linksys WRT54GL
- 2 x IBM XSeries servers running VMware with 8 images of various versions of Elastix IP PBX
- 1 x Standard Microsoft SBS Network providing DHCP and DNS and Mail system
- 2 x Linux Servers

The Elastix IP PBX systems, both hardware and Virtual based have image systems to refresh the systems to limit infection from other testing. Combined with a range of Phones, which include Aastra, Linksys, Cisco, Yealink, it provides a reasonable cross section of typical systems currently in the field.

These application notes are not just done in isolation either. Behind them is over 6-7 years of commercial implementation of IP PBX systems, utilising these methods and concepts. The Lab is just used to reconfirm the implementation in a less production like environment.

How you use these application notes is entirely up to you. However, it is highly recommended that in the first instance, that you follow the notes and configurations in their entirety (except for IP addresses) of course. If you follow it exactly, then it will be easier for others to assist you when you do have an issue.

Overview

The most common hardware failure in an Elastix system is usually one of two things. Either the hard disk fails or the power supply fails, both having a mechanical aspect to their operation, more so the hard disk.

Yes you can implement a solid state hard disk, but as many have found out, these units, depending on their application, are still prone to failure. I have personally seen about 6-7 SSD failures (within their first year) to about 2 Mechanical Hard drive failures over 6 or so years (I am referring to Asterisk based systems here). The technology is good, but just not there with the confidence factor just yet.

An Elastix system with a single drive provides no redundancy in the case of a hard disk failure. Many of the lower end systems (which are great for Elastix) provide either none or poorly supported “Fake Raid”. There are various names for it, including “ChipSet Raid”, or “Onboard Raid”.

When I say poorly supported, I mean that the drivers are hard to obtain for Linux or non-existent as they are not built into Centos. Also in the event of a drive failure, there is little information on how to recover. It is more of a case of fingers crossed, I hope this does what I think it will do, especially as each chipset does it differently.

I am sure that they are exceptions, but on the whole, it is advisable to not use the “onboard” RAID as in many cases, you are better off with no RAID than using it. If you really want to implement Hardware raid then utilise a well known and trusted RAID Controllers such as the PERC range, or 3Ware and many other RAID Controllers out there.

One of the other alternatives is to use Linux Raid or Linux Software Raid. You don’t need anything else except the Centos O/S (which is what Elastix is running on). There are plenty of documents on the Web on how to repair a broken RAID, and plenty of people who have gone before you who can assist.

This document will take you step by step, screen by screen on the setup of RAID at Elastix install time. Once you get used to the concept and method, you can setup a RAID 1 configuration in less than a few minutes, and have a little more confidence than the “ChipSet RAID” method.

What is RAID?

RAID stands for Redundant Array of Inexpensive Disks. There are various levels of RAID, and mixtures of RAID levels, each with its benefits, but there is enough documentation on RAID on the Web, and if you are interested, I recommend performing some further research as it is another subject on its own.

In this tutorial we are covering RAID 1 which is also referred to as Mirroring. In other words everything written to one drive is written to the other. This is done at the O/S level so our software e.g. Asterisk, Elastix does not have to perform any special handling.

The theory and general practice is that one drive can fail and the operating system will continue without interruption on the remaining good hard disk, allowing you to replace the faulty drive at a convenient time.

Hardware requirements for RAID

Fairly simple, two hard drives, preferably identical. One of the benefits of Linux RAID, over Hardware RAID is that the disks can be dissimilar as long as the smallest drive is what you base your partition sizes on. The Drives can be SCSI, IDE, SATA, SAS, even SSD drives.

However, whilst not as critical, it is recommended that you get into the practice of using similar drives, for any type of RAID Level.

The Concept of what we are about to commence

In a nutshell, we are going to use the partitioning tool that is part of the Elastix Installation to create three partitions on the first drive, and exactly the same partitions on the second drive.

Once we have completed this, we are then going to match each partition from both drives and bring them into a RAID Set. It is only when we bring them into the RAID set will we assign the mount points for the partitions.

Just so you can follow, this tutorial was done with two 10Gb hard drives. To perform the setup, I decided to use the following partition sizes for this system

100Mb – Boot Partition (/boot)
2000Mb – SWAP partition
8134Mb – Root Filesystem Partition (/)

That's all there is to it!

RAID Configuration – Setup of the partitions

Commence the Elastix install as you normally would.

When you get to the next screen, this is where we take a departure from the normal process that you may have followed before.

Normally you would have probably selected the “Remove all partitions on select drives......”. In this case however, we are going to select Custom Layout.

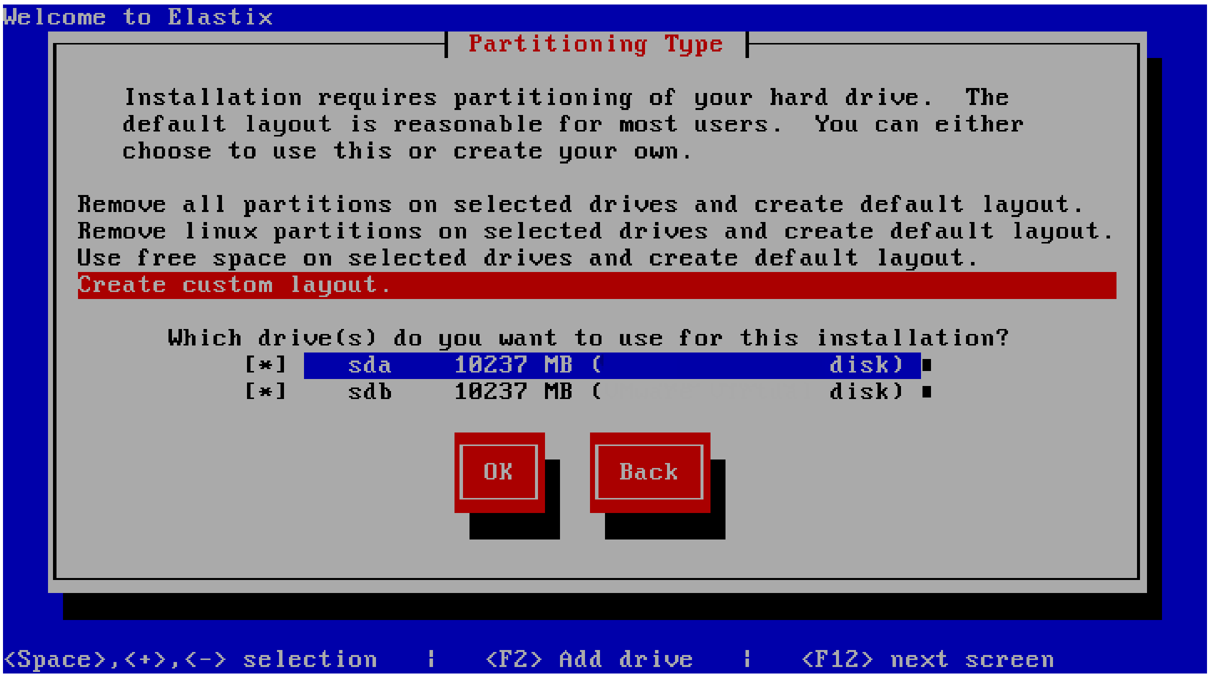

```
Welcome to Elastix
                         Partitioning Type

   Installation requires partitioning of your hard drive.  The
   default layout is reasonable for most users.  You can either
   choose to use this or create your own.

Remove all partitions on selected drives and create default layout.
Remove linux partitions on selected drives and create default layout.
Use free space on selected drives and create default layout.
Create custom layout.

      Which drive(s) do you want to use for this installation?
           [*]    sda     10237 MB (                    disk)
           [*]    sdb     10237 MB (                    disk)

                         OK          Back

<Space>,<+>,<-> selection   |   <F2> Add drive   |   <F12> next screen
```

Before we proceed, make sure that you can see two hard drives. Take note of what there device name is as it will matter as you move through this tutorial. In most systems, you will find the hard disk device names as sda and sdb, however, depending on your system, they may differ e.g. hda and hdb.. This is not an issue, but for the purpose of this tutorial, we will remain with the sda and sdb. You will just need to use the correct translation for your devices.

If you don’t see two drives, then you need to go back and correct the issue. It might be you have played with the “ChipSet” RAID and left it enabled or you have a hardware issue.

Anyhow, select CREATE CUSTOM LAYOUT and select the OK button.

Your screen should look similar except that your drive sizes will differ from this example.

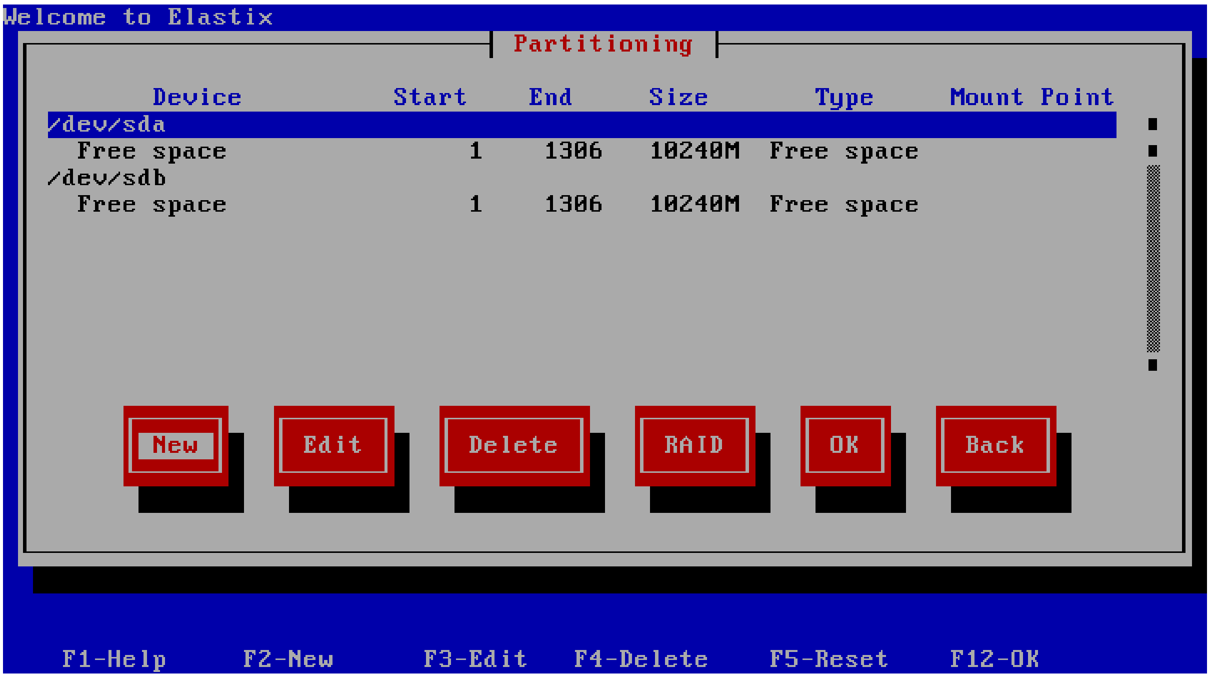

```
        Device                Start    End      Size       Type      Mount Point
/dev/sda
    Free space                   1     1306     10240M   Free space
/dev/sdb
    Free space                   1     1306     10240M   Free space
```

Click on the NEW Button and the following screen comes up

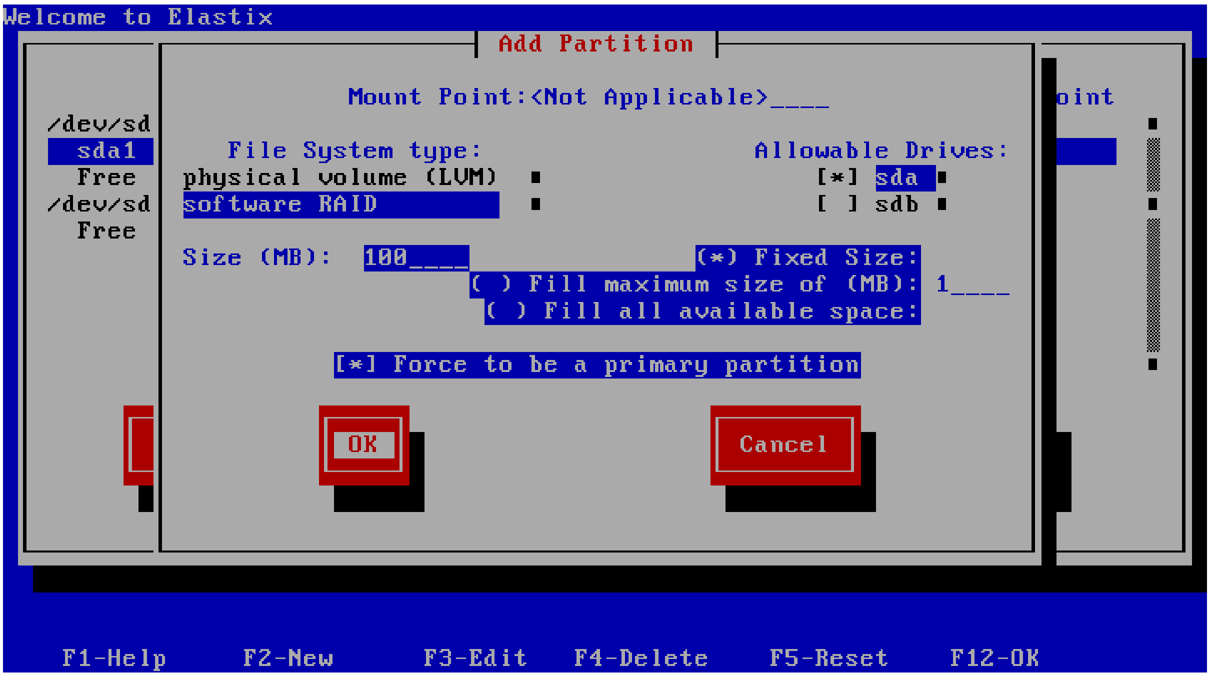

```
             Mount Point:<Not Applicable>____

    File System type:                           Allowable Drives:
 physical volume (LVM)                              [*] sda
 software RAID                                      [ ] sdb

 Size (MB):  100____                     (*) Fixed Size:
                          ( ) Fill maximum size of (MB): 1____
                           ( ) Fill all available space:

               [*] Force to be a primary partition
```

Tab your way through each of the options. There is no need to type anything into the mount point as when you select the File System type as Software RAID, it will mark it as <Not Applicable>. Make sure that you select Allowable drives as **sda** only.

The reason for this is that we are only working on the **sda** drive at the moment. If you select **sda** and **sdb**, it will place the partition on either drive, which is not what we want.

As this will be our boot partition, we only need a size of 100Mb, and it is set as a Fixed Size.

Last thing, make sure you toggle Force to be a primary Partition.

Click OK and you should now see the following screen.

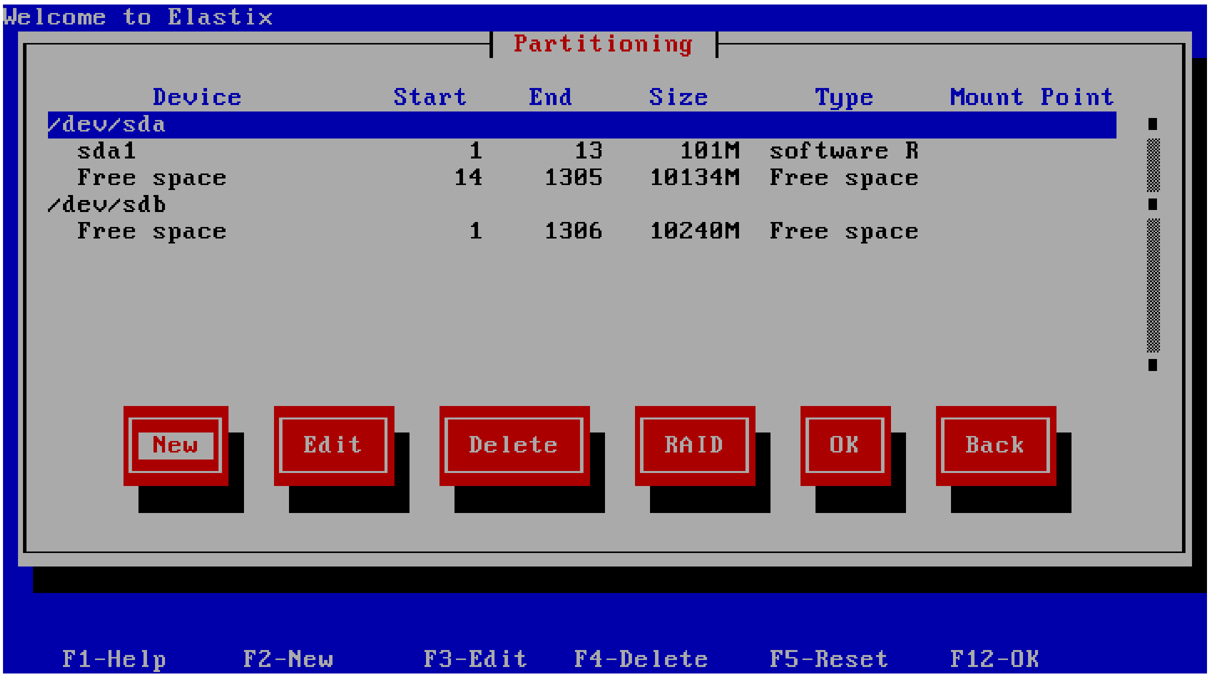

```
Welcome to Elastix
                              ┤ Partitioning ├

        Device              Start     End     Size       Type      Mount Point
/dev/sda
  sda1                          1       13     101M   software R
  Free space                   14     1305   10134M   Free space
/dev/sdb
  Free space                    1     1306   10240M   Free space

       New      Edit      Delete       RAID       OK      Back

    F1-Help     F2-New      F3-Edit    F4-Delete    F5-Reset    F12-OK
```

Nothing much to see, except you will now see that you have created a 100Mb partition on **sda**.

Click on NEW again and we are now going to setup a partition which will be the SWAP Partition.

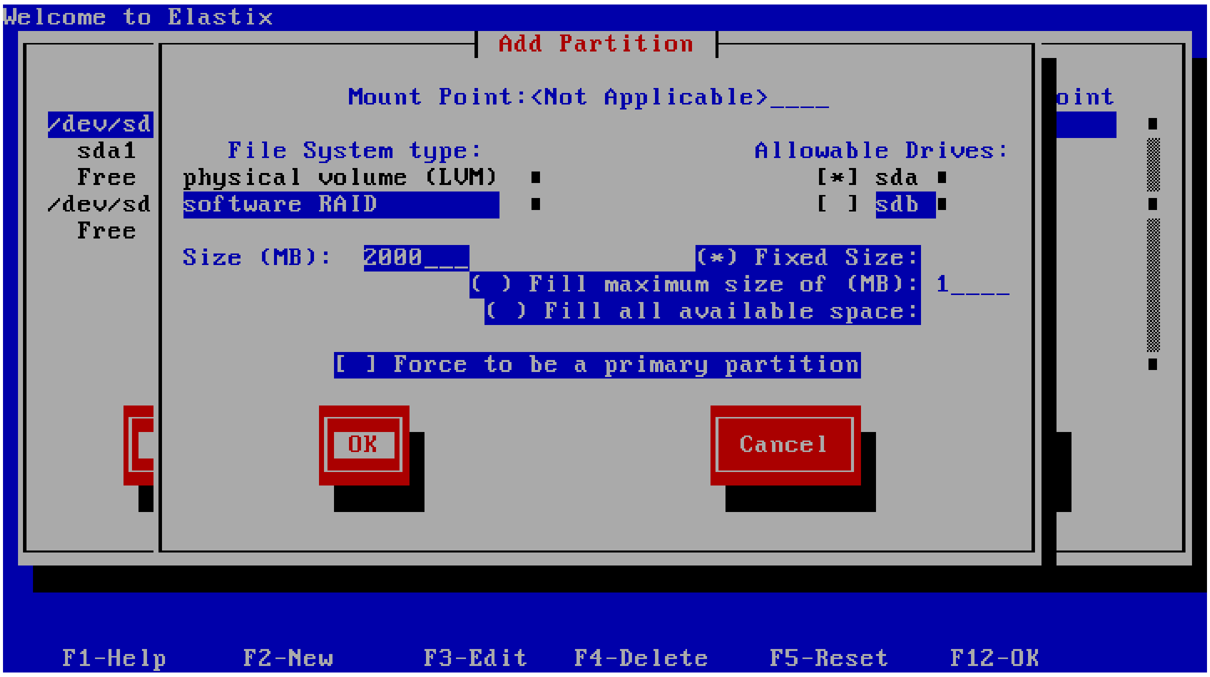

```
Welcome to Elastix
                             ┤ Add Partition ├

              Mount Point:<Not Applicable>____

        File System type:                      Allowable Drives:
   physical volume (LVM)                            [*] sda
   software RAID                                    [ ] sdb

   Size (MB):  2000___               (*) Fixed Size:
                                 ( ) Fill maximum size of (MB): 1____
                                  ( ) Fill all available space:

             [ ] Force to be a primary partition

                   OK                         Cancel

    F1-Help     F2-New      F3-Edit    F4-Delete    F5-Reset    F12-OK
```

Basically go through the same routine, except the size may vary between this tutorial and what you want to use as a swap file size. The system has 1Gb of RAM, so I have run along standard lines, which is usually the swap file partition being double that size (or thereabouts). The only other difference from the boot partition is that you do not select Force to be a Primary Partition.

Click ok and you should see the following screen. Again nothing spectacular, but it shows the two partitions that you have setup.

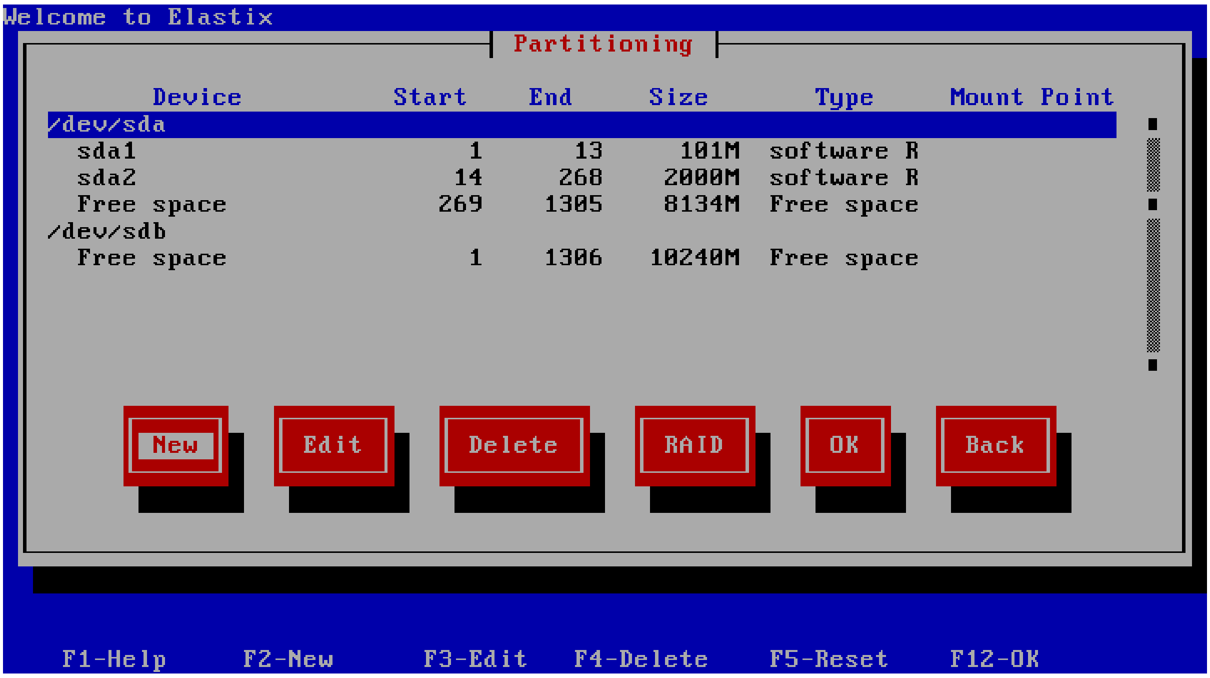

```
Welcome to Elastix
                              Partitioning

      Device             Start     End      Size        Type       Mount Point
/dev/sda
    sda1                      1      13      101M   software R
    sda2                     14     268     2000M   software R
    Free space              269    1305     8134M   Free space
/dev/sdb
    Free space                1    1306    10240M   Free space

 F1-Help     F2-New     F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Click on NEW again so that we can setup the final partition on the sda drive.

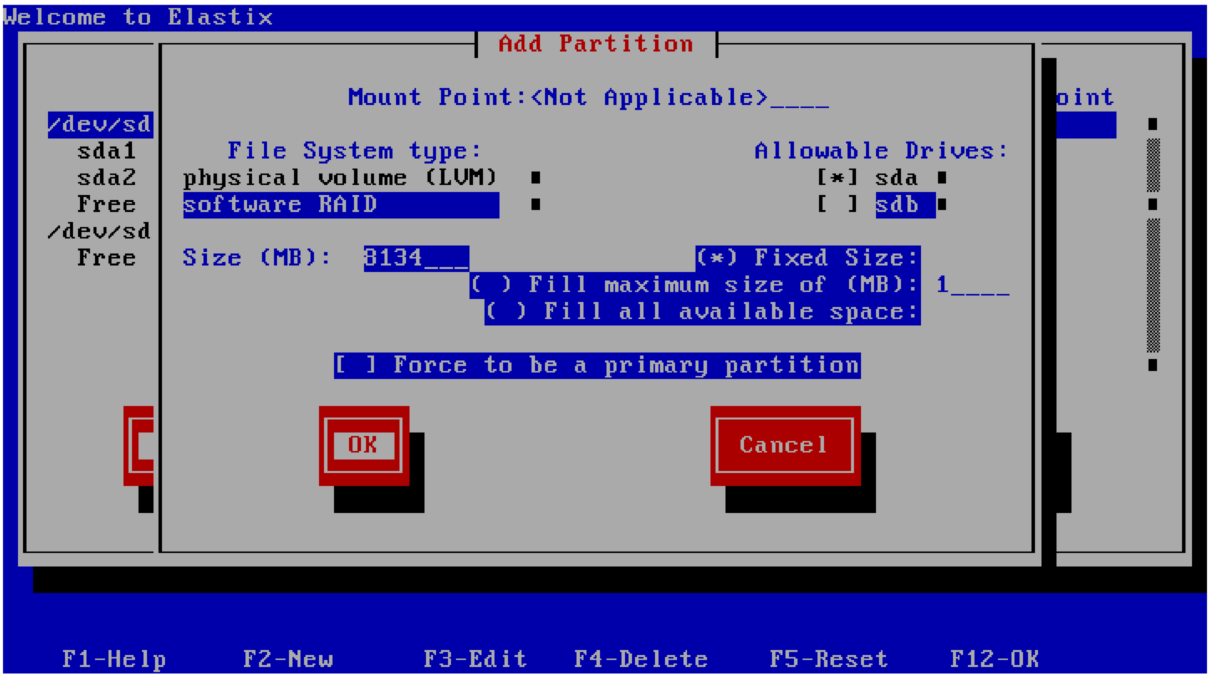

```
Welcome to Elastix
                              Add Partition

              Mount Point:<Not Applicable>____

         File System type:                       Allowable Drives:
      physical volume (LVM)                          [*] sda
      software RAID                                  [ ] sdb

      Size (MB):  8134___                  (*) Fixed Size:
                          ( ) Fill maximum size of (MB):  1____
                           ( ) Fill all available space:

                 [ ] Force to be a primary partition

 F1-Help     F2-New     F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Again same routine as the last partition we setup. This partition will be our main partition, which generally you want to make as large as possible. If you have two exact identical disks, then you could just select Fill all available Space and it will use whatever is left of the hard disk. However I prefer to set the partition size myself, and the size I used was what was shown as left on the previous screen.

Click ok and you should see the next screen.

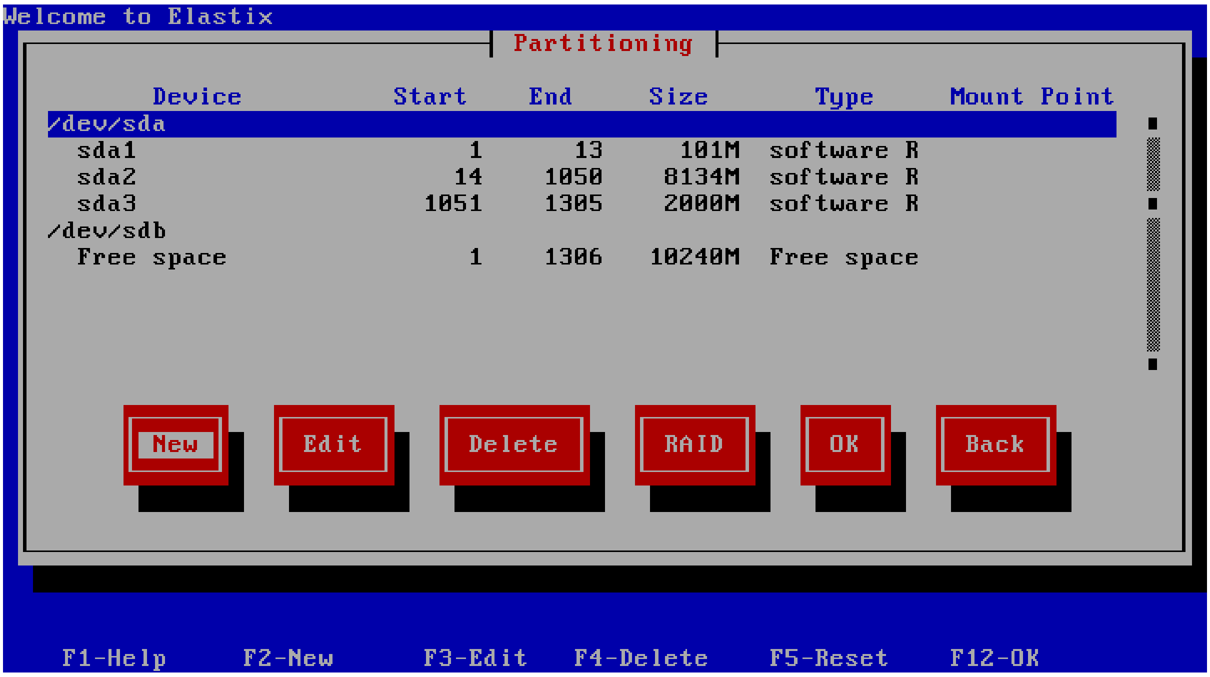

```
Welcome to Elastix
                              Partitioning

       Device              Start     End      Size       Type      Mount Point
/dev/sda
  sda1                         1      13      101M   software R
  sda2                        14    1050     8134M   software R
  sda3                      1051    1305     2000M   software R
/dev/sdb
  Free space                   1    1306    10240M   Free space

   New     Edit     Delete     RAID     OK     Back

   F1-Help     F2-New     F3-Edit    F4-Delete    F5-Reset    F12-OK
```

This screen shows you the partitions that you have setup, and in fact you have finished partitioning the **sda** hard drive.

One thing you will notice is that sda2 and sda3 have switched around, which might throw you a little bit. Centos appears to "optimise" how the partitions are laid out. Don't panic, as long as the sda and sdb partitions are the same when you are finish, you will be fine. Click on NEW again, and we run through the same thing for **sdb**

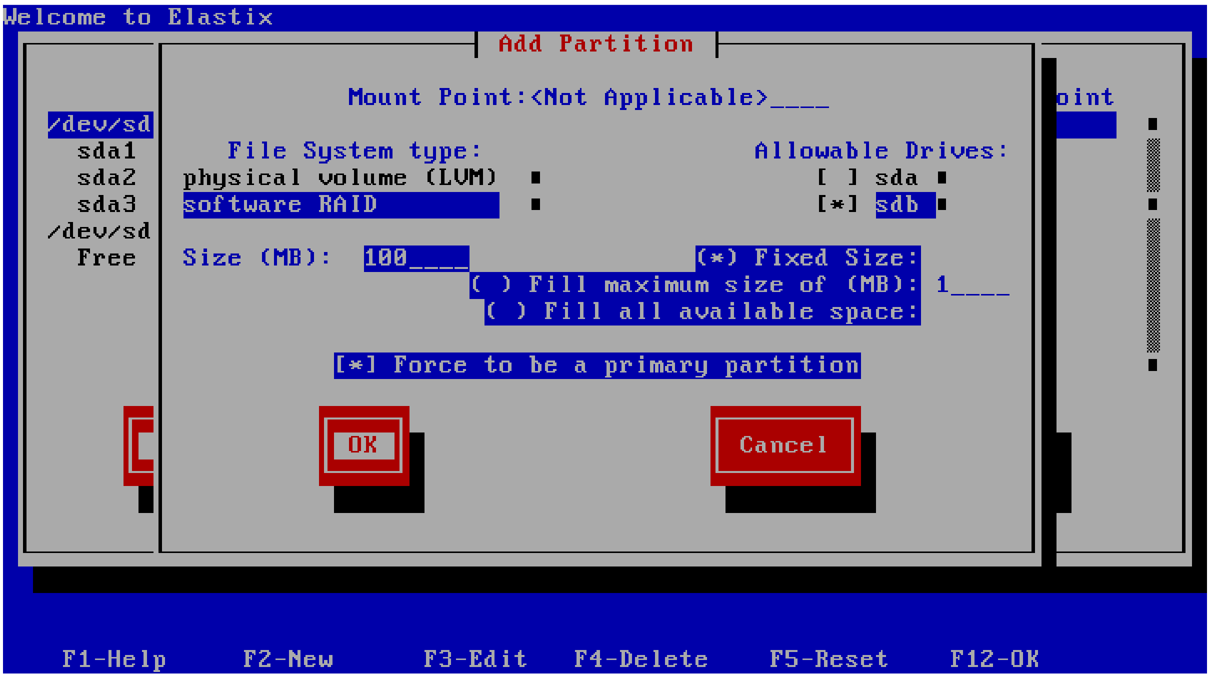

```
Welcome to Elastix
                             Add Partition

              Mount Point:<Not Applicable>____

        File System type:                     Allowable Drives:
  physical volume (LVM)                           [ ] sda
  software RAID                                   [*] sdb

  Size (MB): 100____               (*) Fixed Size:
                          ( ) Fill maximum size of (MB): 1____
                           ( ) Fill all available space:

              [*] Force to be a primary partition

                 OK                          Cancel

   F1-Help     F2-New     F3-Edit    F4-Delete    F5-Reset    F12-OK
```

You will note the same selections, except that we now only select **sdb** as opposed to **sda** in the previous three partitions. Again on this one, we mark it with Force to be a Primary Partition.

Once done click OK and you will see the following screen

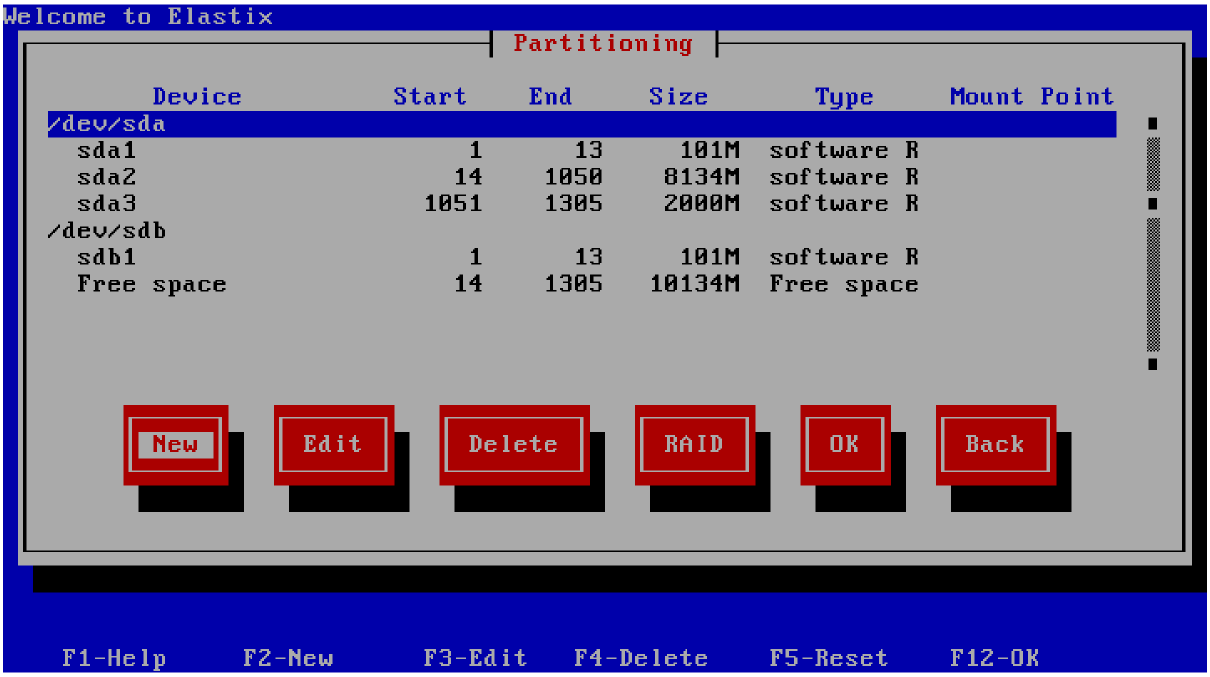

```
Welcome to Elastix
                              Partitioning

      Device               Start     End      Size       Type      Mount Point
 /dev/sda
    sda1                       1      13      101M   software R
    sda2                      14    1050     8134M   software R
    sda3                    1051    1305     2000M   software R
 /dev/sdb
    sdb1                       1      13      101M   software R
    Free space                14    1305    10134M   Free space

     New      Edit      Delete       RAID       OK       Back

    F1-Help     F2-New     F3-Edit   F4-Delete    F5-Reset    F12-OK
```

I think by now you are starting to see the idea

Same as we did on the **sda** drive, but remember to make sure that **sdb** is the only allowable drive

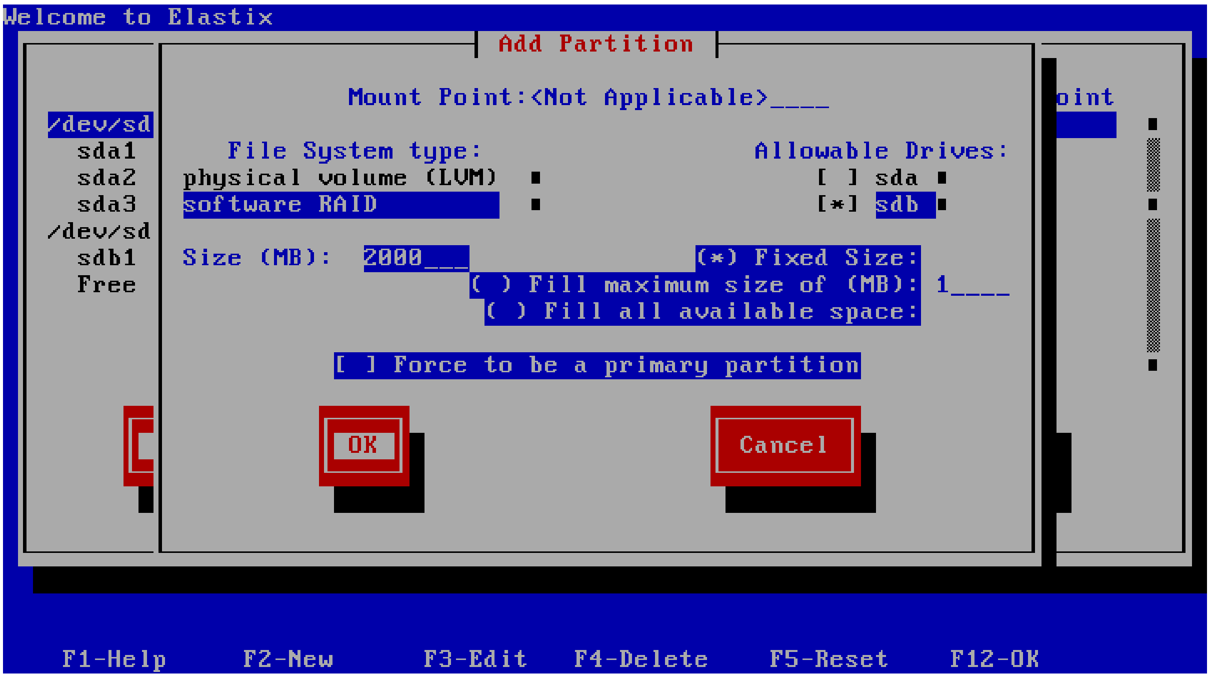

```
Welcome to Elastix
                            Add Partition

              Mount Point:<Not Applicable>____

          File System type:                       Allowable Drives:
       physical volume (LVM)                            [ ] sda
       software RAID                                    [*] sdb

       Size (MB):  2000___               (*) Fixed Size:
                                     ( ) Fill maximum size of (MB): 1____
                                      ( ) Fill all available space:

                   [ ] Force to be a primary partition

                        OK                         Cancel

    F1-Help     F2-New     F3-Edit   F4-Delete    F5-Reset    F12-OK
```

Click OK and again check that it looks correct on the partitioning table

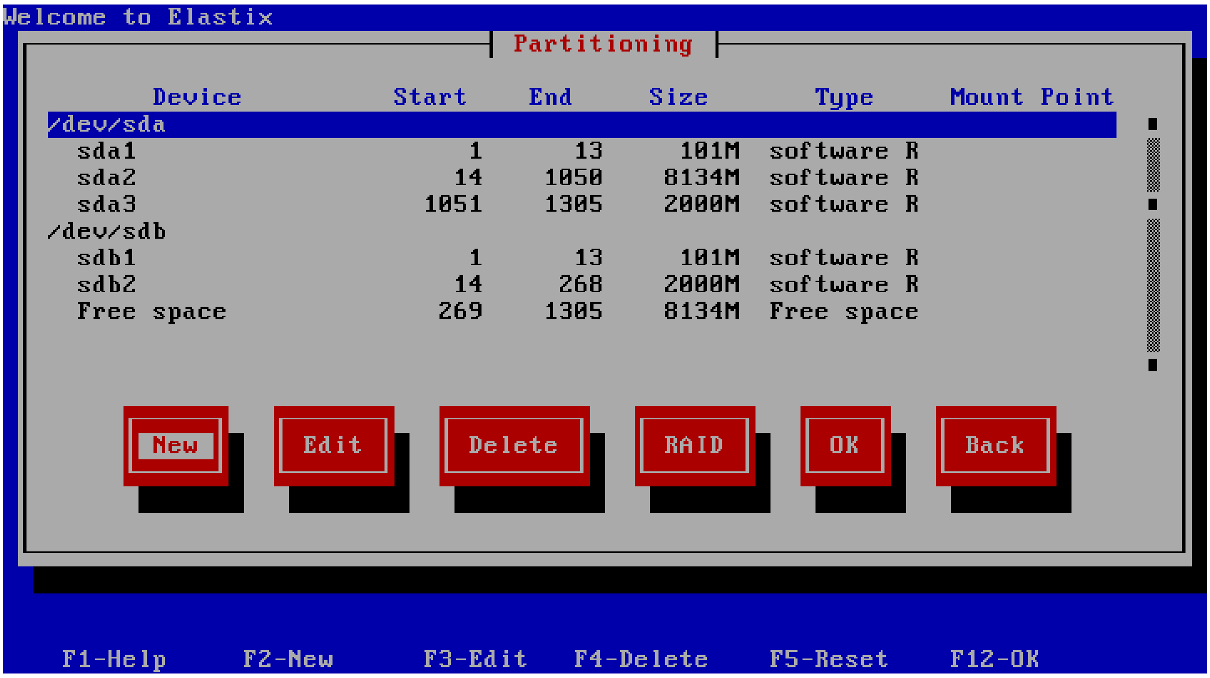

```
Welcome to Elastix
                              Partitioning

          Device                Start      End      Size        Type        Mount Point
/dev/sda
    sda1                            1       13      101M   software R
    sda2                           14     1050     8134M   software R
    sda3                         1051     1305     2000M   software R
/dev/sdb
    sdb1                            1       13      101M   software R
    sdb2                           14      268     2000M   software R
    Free space                    269     1305     8134M   Free space
```

Click on NEW and perform the final partition on the **sdb** drive.

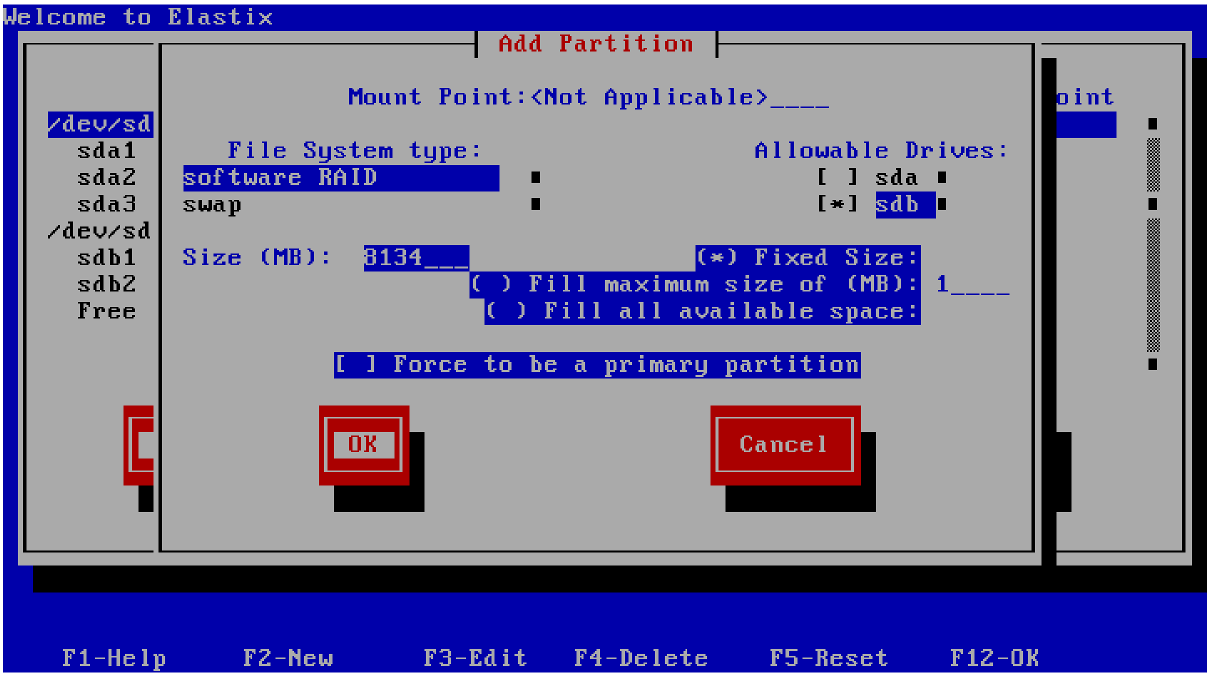

```
Welcome to Elastix
                         Add Partition

            Mount Point:<Not Applicable>____

      File System type:                    Allowable Drives:
   software RAID                                  [ ] sda
   swap                                           [*] sdb

   Size (MB):  8134___              (*) Fixed Size:
                           ( ) Fill maximum size of (MB): 1____
                            ( ) Fill all available space:

              [ ] Force to be a primary partition
```

Click on OK and you should see a screen similar to the one below

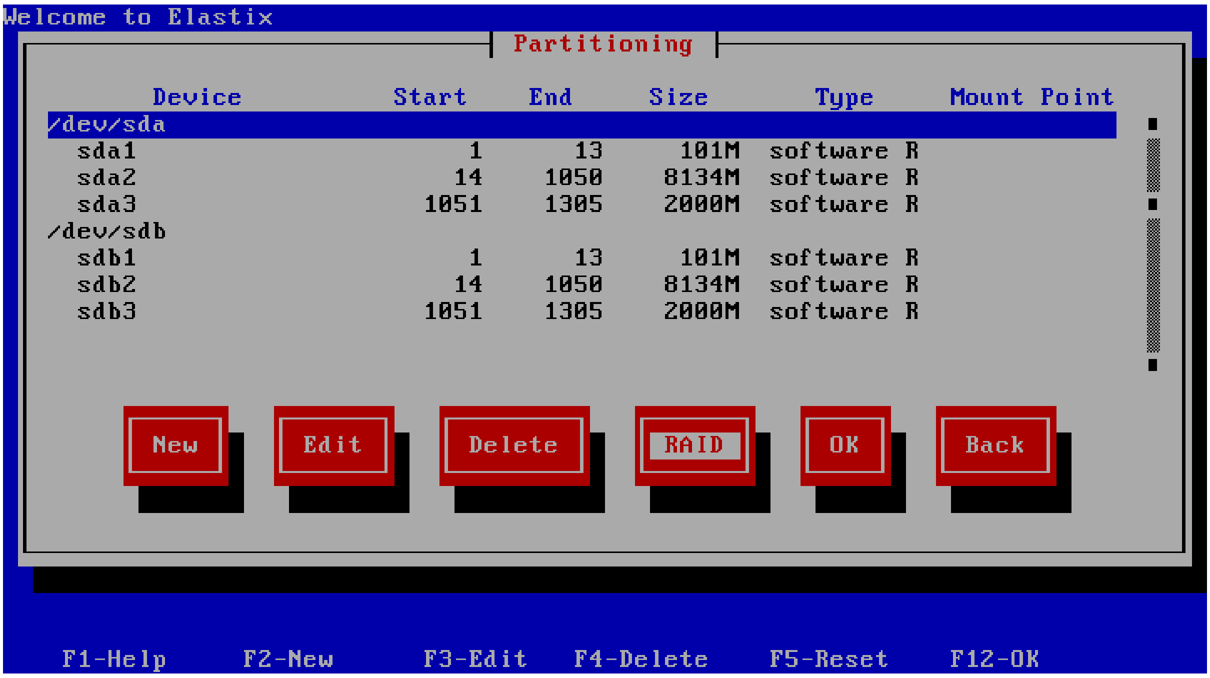

```
Welcome to Elastix
                                   Partitioning

          Device                Start    End     Size       Type      Mount Point
/dev/sda
  sda1                              1      13     101M  software R
  sda2                             14    1050    8134M  software R
  sda3                           1051    1305    2000M  software R
/dev/sdb
  sdb1                              1      13     101M  software R
  sdb2                             14    1050    8134M  software R
  sdb3                           1051    1305    2000M  software R
```

Ok we have now setup the partitions and marked them as RAID Partitions. Almost there….!!

Now click on RAID and we will move onto the next chapter.

RAID Setup – Populating the RAID Sets and setting Mount Points

Very simply, this is the section where we tie everything together.

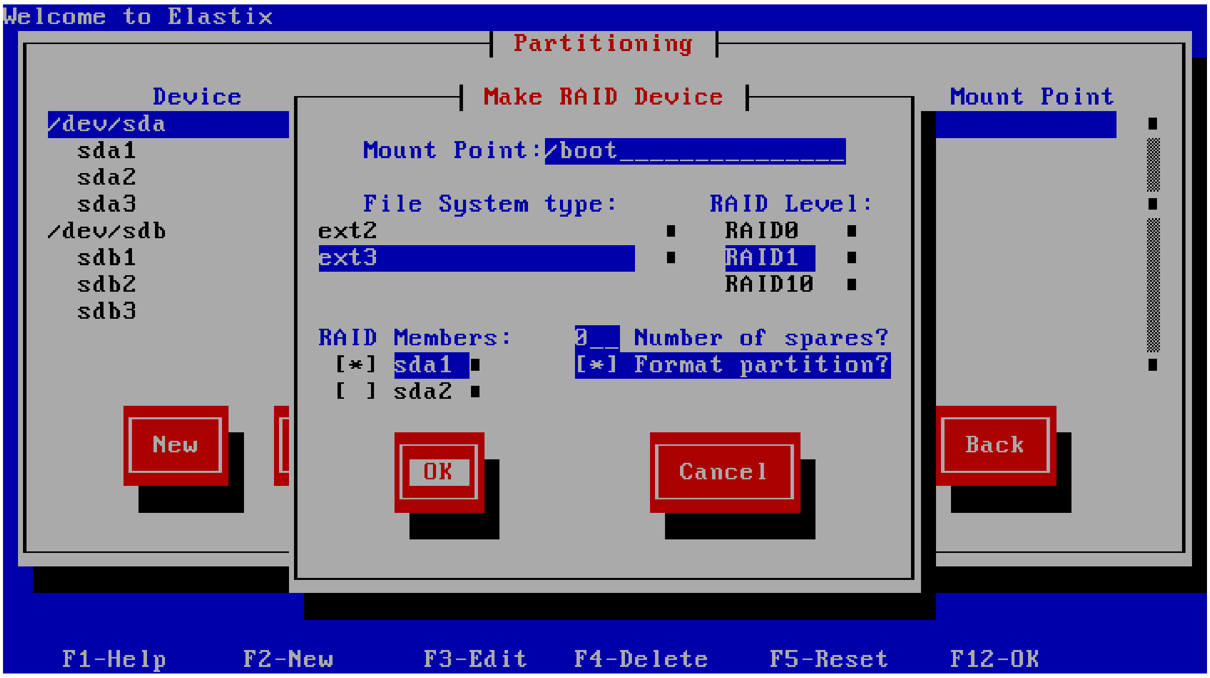

```
Welcome to Elastix
                          ┤ Partitioning ├

        Device        ┤ Make RAID Device ├           Mount Point
 /dev/sda
   sda1           Mount Point:/boot_______________
   sda2
   sda3           File System type:        RAID Level:
 /dev/sdb       ext2                    ▪    RAID0   ▪
   sdb1         ext3                    ▪    RAID1   ▪
   sdb2                                      RAID10  ▪
   sdb3
                RAID Members:         0__ Number of spares?
                 [*] sda1 ▪           [*] Format partition?
                 [ ] sda2 ▪

      New                  OK              Cancel          Back

   F1-Help    F2-New    F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Everything is done on this screen. First of all we set a mount point, which in this case will be **/boot** . It will have a file system type of **EXT3**, the RAID Level will be **RAID1**, and we need to select the partition members that will be used for this mount point, and also which partitions are the matching RAID Partitions.

So follow the settings that are on this screenshot, except the one thing you cannot see in the screen shot is the list of RAID Members. When you tab across to this option, use the down arrows and you will see that it lists

sda1
sda2
sda3
sdb1
sdb2
sdb3

The asterisk should only be beside **sda1** and **sdb1** (the two matching partitions from each hard drive).

This is where the most common mistake is for first time setup of RAID, as it is not overly intuitive, but it is the most critical part.

Once you are finished, click on ok and you will see the following list, and our first RAID partition setup.

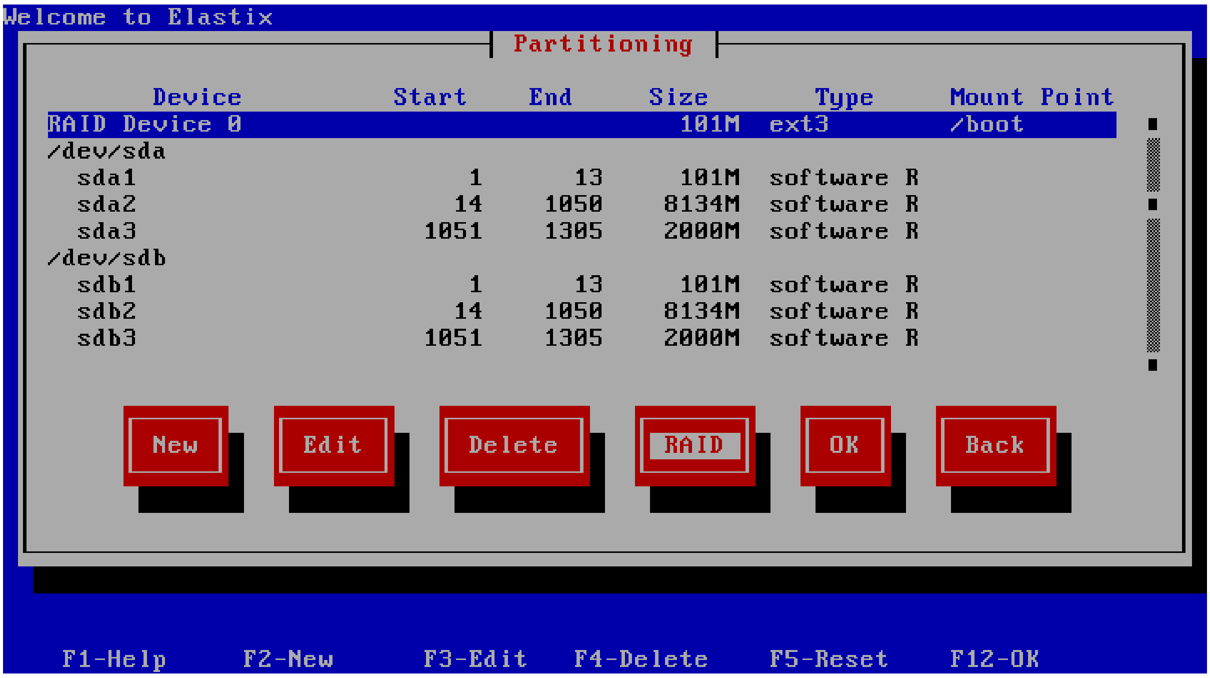

```
Welcome to Elastix
                         Partitioning

      Device             Start    End      Size      Type       Mount Point
RAID Device 0                               101M    ext3        /boot
/dev/sda
  sda1                       1       13     101M    software R
  sda2                      14     1050    8134M    software R
  sda3                    1051     1305    2000M    software R
/dev/sdb
  sdb1                       1       13     101M    software R
  sdb2                      14     1050    8134M    software R
  sdb3                    1051     1305    2000M    software R

     New     Edit     Delete     RAID     OK     Back

   F1-Help    F2-New    F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Click on RAID again and you will come to the next screen

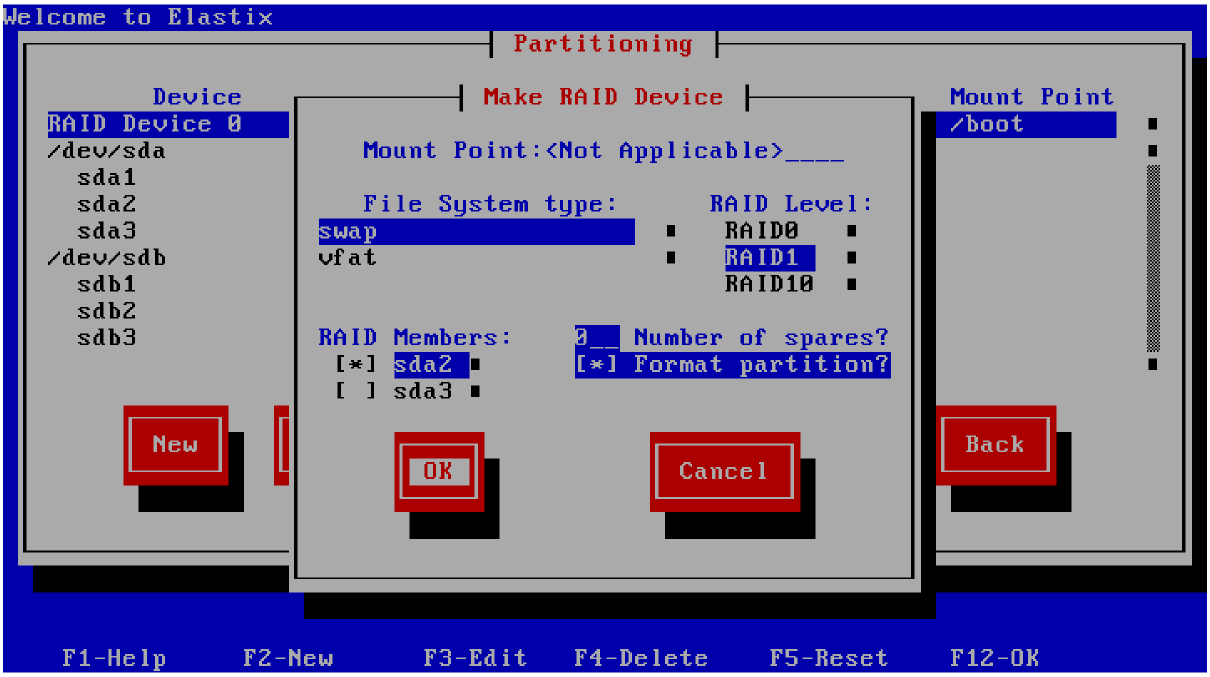

```
Welcome to Elastix
                         Partitioning
      Device          Make RAID Device               Mount Point
RAID Device 0                                         /boot
/dev/sda           Mount Point:<Not Applicable>____
  sda1
  sda2             File System type:        RAID Level:
  sda3           swap                          RAID0
/dev/sdb         vfat                          RAID1
  sdb1                                         RAID10
  sdb2
  sdb3           RAID Members:         0__ Number of spares?
                  [*] sda2             [*] Format partition?
                  [ ] sda3
     New
                        OK              Cancel                Back

   F1-Help    F2-New    F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Here we setup and assign the **SWAP** Partition. **SWAP** is not mounted, so it does not have a mount point. Leave it blank, and select the file system as **SWAP**. Again make sure the RAID level is **RAID 1** and like before, select the RAID Members.

You will find that the RAID Members has shortened and you should now only be left with

sda2
sda3
sdb2
sdb3

As you remember, we selected a 2000Mb partition as the **SWAP** partition which is **sda3** and **sdb3**, so make sure the asterisk is next to those two only, and click on OK.

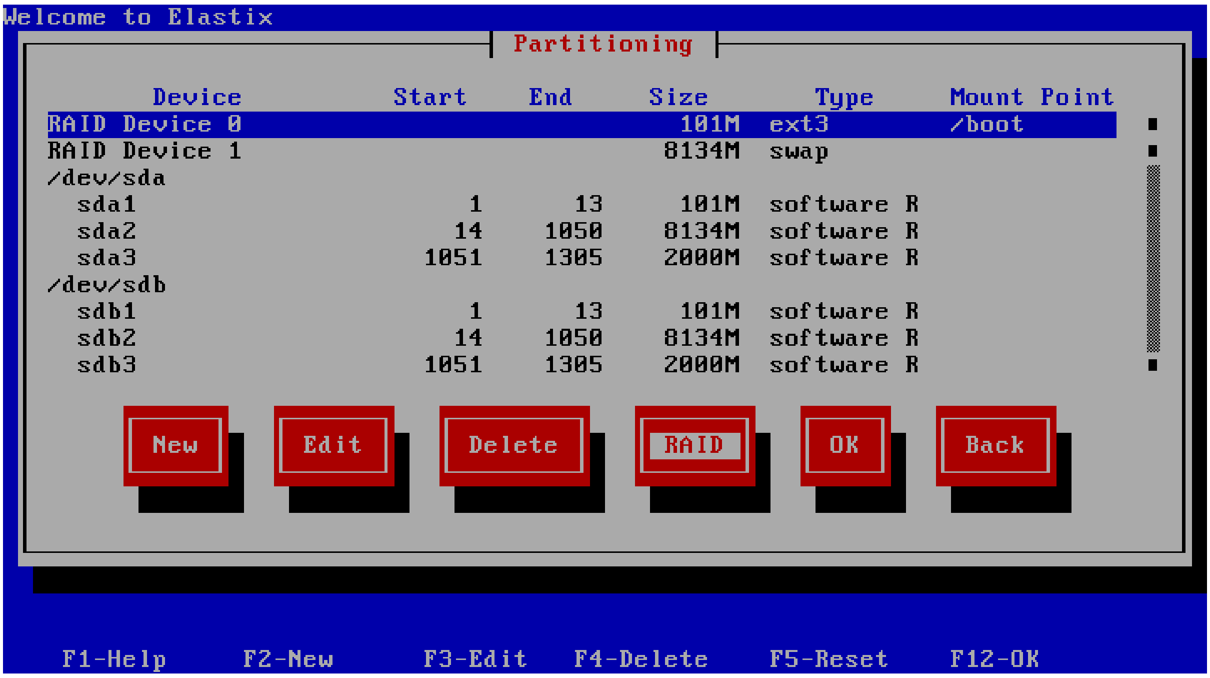

```
Welcome to Elastix
                              ┤ Partitioning ├

     Device               Start    End      Size       Type       Mount Point
  RAID Device 0                              101M    ext3        /boot
  RAID Device 1                             8134M    swap
  /dev/sda
    sda1                       1      13     101M    software R
    sda2                      14    1050    8134M    software R
    sda3                    1051    1305    2000M    software R
  /dev/sdb
    sdb1                       1      13     101M    software R
    sdb2                      14    1050    8134M    software R
    sdb3                    1051    1305    2000M    software R

     New     Edit     Delete     RAID     OK     Back

   F1-Help    F2-New    F3-Edit   F4-Delete   F5-Reset   F12-OK
```

As you can see, we are progressing and we are almost finished.

Finally we setup the Root File system partition.

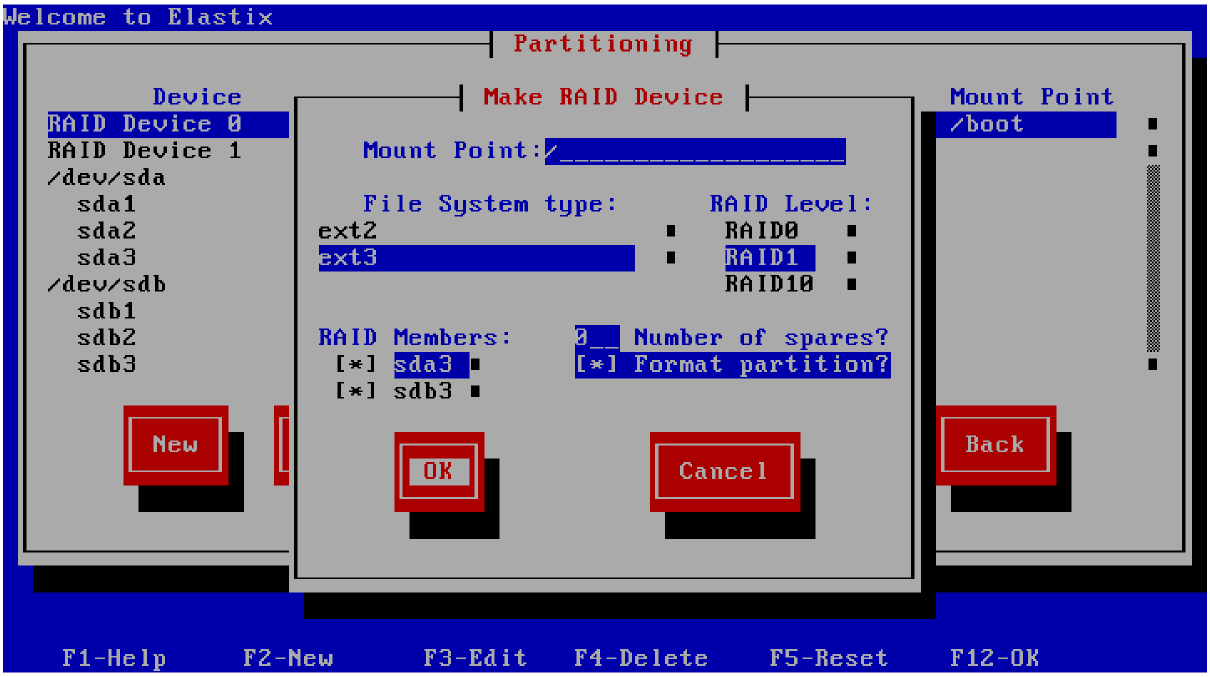

```
Welcome to Elastix
                              ┤ Partitioning ├

     Device             ┤ Make RAID Device ├                    Mount Point
  RAID Device 0                                                 /boot
  RAID Device 1           Mount Point:/____________________
  /dev/sda
    sda1                  File System type:      RAID Level:
    sda2                ext2                        RAID0
    sda3                ext3                        RAID1
  /dev/sdb                                          RAID10
    sdb1
    sdb2                RAID Members:         0__ Number of spares?
    sdb3                 [*] sda3             [*] Format partition?
                         [*] sdb3
     New                                                            Back
                             OK           Cancel

   F1-Help    F2-New    F3-Edit   F4-Delete   F5-Reset   F12-OK
```

Same routine as like the /boot Partition, instead using the **/** as the mount point but you will find only two partitions left which are

sda3
sdb3

Make sure they are selected with the asterisk, and click OK and the final screen will appear

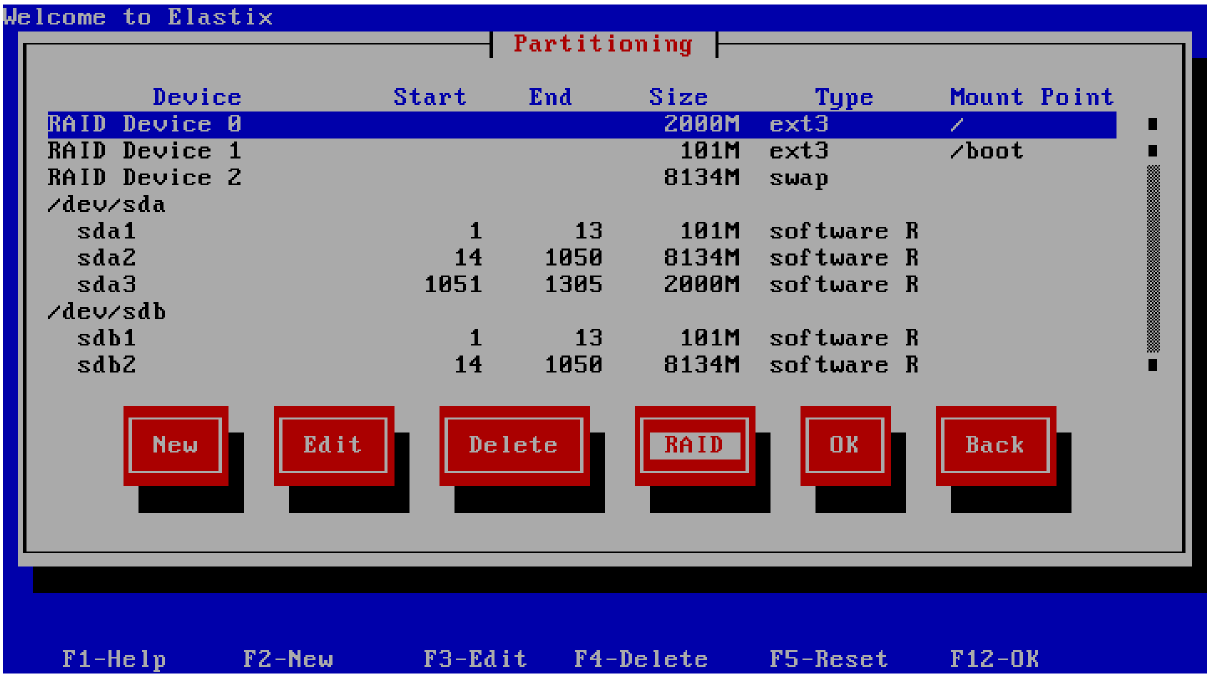

```
Welcome to Elastix
                              Partitioning

        Device              Start     End     Size       Type      Mount Point
   RAID Device 0                             2000M    ext3        /
   RAID Device 1                              101M    ext3        /boot
   RAID Device 2                             8134M    swap
   /dev/sda
     sda1                       1       13    101M    software R
     sda2                      14     1050   8134M    software R
     sda3                    1051     1305   2000M    software R
   /dev/sdb
     sdb1                       1       13    101M    software R
     sdb2                      14     1050   8134M    software R

    F1-Help      F2-New      F3-Edit    F4-Delete    F5-Reset    F12-OK
```

If you have the capability, I fully recommend taking a photo or a screen shot or copying the information down. It's not necessary for general day to day, but if you have to perform a RAID repair or replace a disk the information in that screen is invaluable, and could save you making a mistake.

Now complete the Elastix install as per normal.

One area that we need to complete once Elastix has finished its install, is that we need to install GRUB onto the second hard drive, otherwise if you need to reboot the system in the event of the first hard drive failing, it will not be able to boot.

GRUB is not replicated to the second hard drive as it is a unique item, in that it is a bootloader that normally is installed into sector 0 of the device. Very much the same thing you might have seen on dual boot systems, where you can have Windows and Linux and choose to boot onto either system.

There is also cases where GRUB does not install on some systems, so what we do is take precautions, and we install the GRUB bootloader, not just onto the second drive, but also the first drive for safety.

This is done very easily.....

As soon as the Elastix install is completed and at the login prompt, login as root and at the Linux prompt type

Grub

And the following GRUB Shell will appear

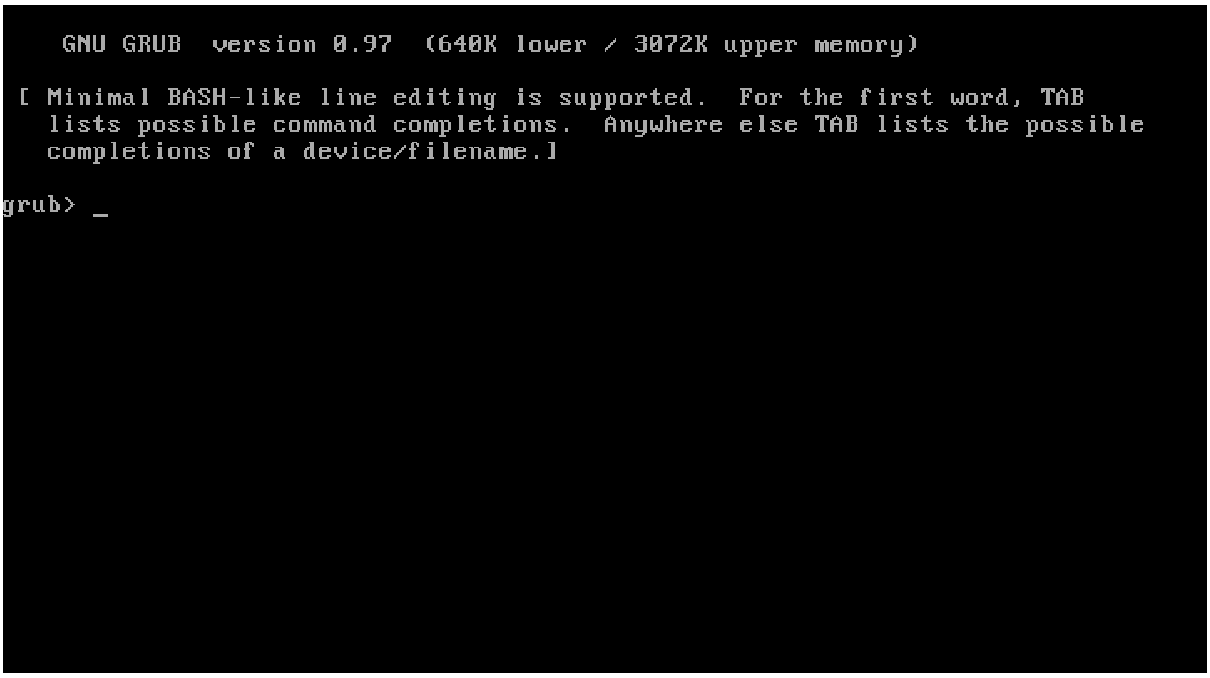

```
    GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

 [ Minimal BASH-like line editing is supported.  For the first word, TAB
   lists possible command completions.  Anywhere else TAB lists the possible
   completions of a device/filename.]

grub> _
```

Type each of this commands one by one
grub> **device (hd0) /dev/sda <enter>**
grub> **device (hd1) /dev/sdb <enter>**
grub> **root (hd0,0) <enter>**
grub> **setup (hd0) <enter>**
grub> **root (hd1,0) <enter>**
grub> **setup (hd1) <enter>**

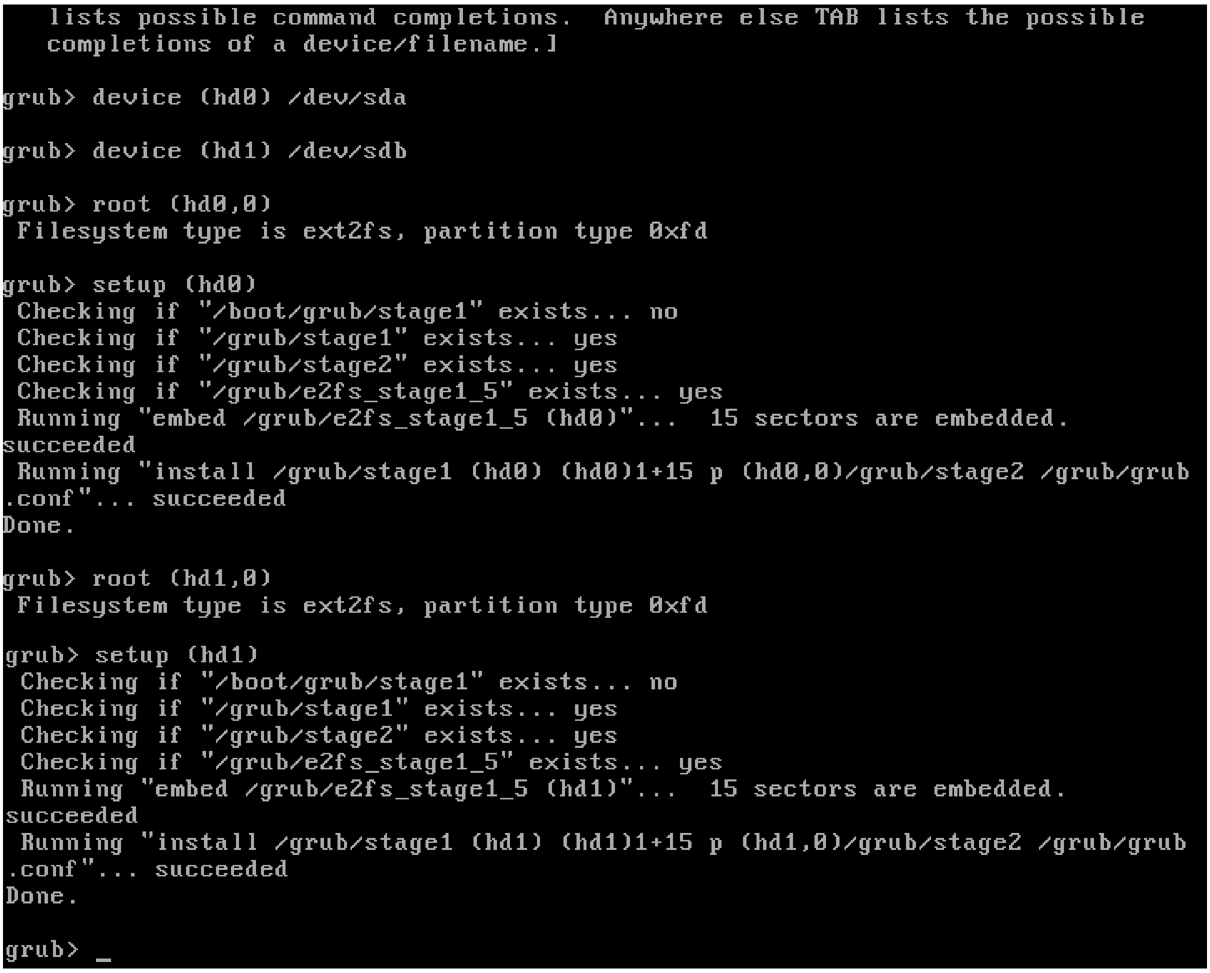

```
   lists possible command completions.  Anywhere else TAB lists the possible
   completions of a device/filename.]

grub> device (hd0) /dev/sda

grub> device (hd1) /dev/sdb

grub> root (hd0,0)
 Filesystem type is ext2fs, partition type 0xfd

grub> setup (hd0)
 Checking if "/boot/grub/stage1" exists... no
 Checking if "/grub/stage1" exists... yes
 Checking if "/grub/stage2" exists... yes
 Checking if "/grub/e2fs_stage1_5" exists... yes
 Running "embed /grub/e2fs_stage1_5 (hd0)"...  15 sectors are embedded.
succeeded
 Running "install /grub/stage1 (hd0) (hd0)1+15 p (hd0,0)/grub/stage2 /grub/grub
.conf"... succeeded
Done.

grub> root (hd1,0)
 Filesystem type is ext2fs, partition type 0xfd

grub> setup (hd1)
 Checking if "/boot/grub/stage1" exists... no
 Checking if "/grub/stage1" exists... yes
 Checking if "/grub/stage2" exists... yes
 Checking if "/grub/e2fs_stage1_5" exists... yes
 Running "embed /grub/e2fs_stage1_5 (hd1)"...  15 sectors are embedded.
succeeded
 Running "install /grub/stage1 (hd1) (hd1)1+15 p (hd1,0)/grub/stage2 /grub/grub
.conf"... succeeded
Done.

grub> _
```

grub> **quit <enter>**

And now reboot

Confirming that your RAID is working

Now to confirm that you have successfully completed your RAID Setup

At the Linux prompt type

cat /proc/mdstat <enter>

and a similar screen will appear. In this case, with the [UU] it shows that my RAID Mirror is complete and completed building.

If you are missing a U on all or any partitions, then your RAID setup is degraded.

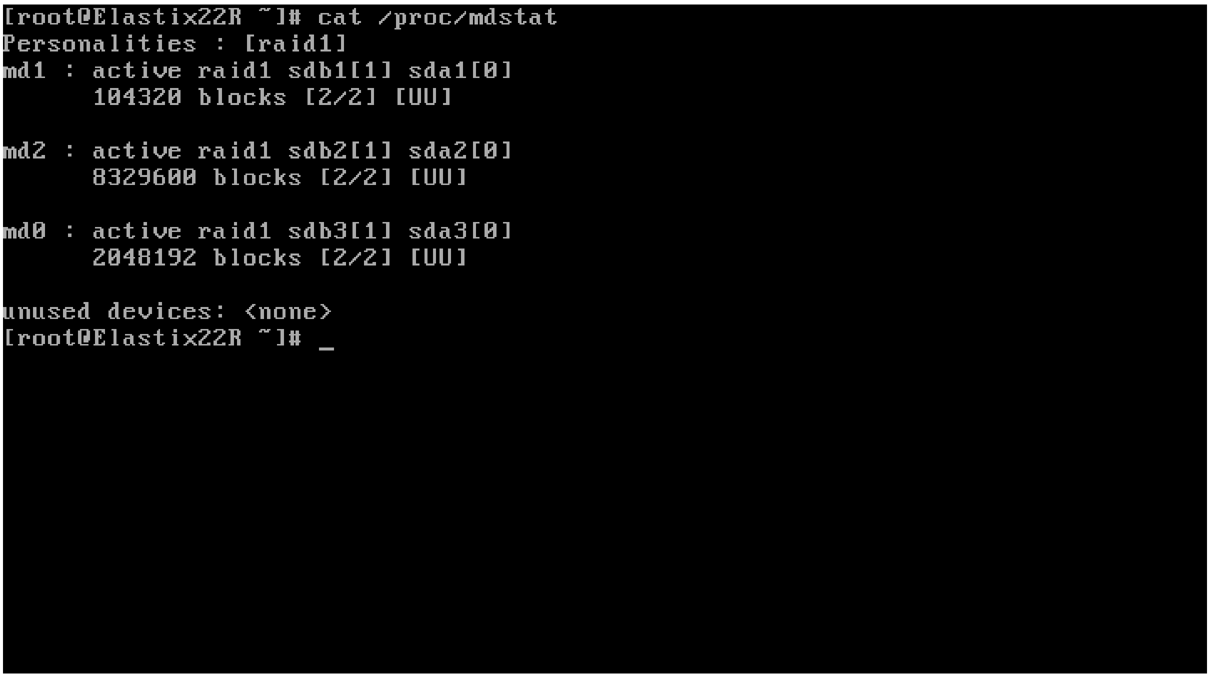

```
[root@Elastix22R ~]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb2[1] sda2[0]
      8329600 blocks [2/2] [UU]

md0 : active raid1 sdb3[1] sda3[0]
      2048192 blocks [2/2] [UU]

unused devices: <none>
[root@Elastix22R ~]# _
```

Depending on your hard disks and sizes, you may find that one or more arrays are still building and you will see its progress on the screen.

Take note that your RAID sets are now referred to as **MD0** through to **MD2** (MD stands for multiple drives). This is important when you are checking status or rebuilding the RAID set.

Management of your RAID

I mentioned at the start that the tools in some of the “Chipset RAIDs” left a lot to be desired, and that ‘s if you can trust them. Linux O/S Raid has some nice tools, and the one that is available is MDADM

For instance, I can issue the following command

mdadm --detail /dev/md0

and it will display the following screen

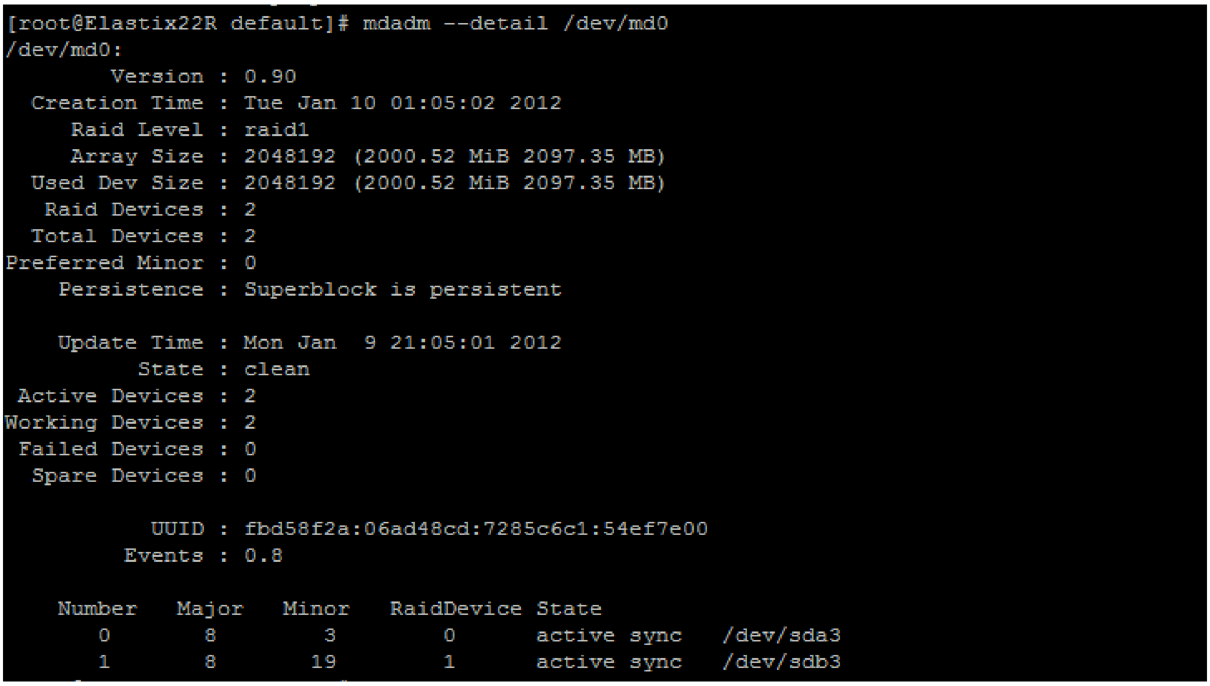

```
[root@Elastix22R default]# mdadm --detail /dev/md0
/dev/md0:
        Version : 0.90
  Creation Time : Tue Jan 10 01:05:02 2012
     Raid Level : raid1
     Array Size : 2048192 (2000.52 MiB 2097.35 MB)
  Used Dev Size : 2048192 (2000.52 MiB 2097.35 MB)
   Raid Devices : 2
  Total Devices : 2
Preferred Minor : 0
    Persistence : Superblock is persistent

    Update Time : Mon Jan  9 21:05:01 2012
          State : clean
 Active Devices : 2
Working Devices : 2
 Failed Devices : 0
  Spare Devices : 0

           UUID : fbd58f2a:06ad48cd:7285c6c1:54ef7e00
         Events : 0.8

    Number   Major   Minor   RaidDevice State
       0       8        3        0      active sync   /dev/sda3
       1       8       19        1      active sync   /dev/sdb3
```

This command tells mdadm to provide the details and status on **MD0** which if you remember from the final partition list, is our **SWAP** partition. It shows the state, shows the members of the RAID, the RAID Level.

This is one of the first commands that you might run on each of your partitions to confirm their health.

We will also use this command to reconstruct a failed RAID set, so it is worthwhile learning what it can do.

RAID Recovery

The most common thing that will occur will be a failed hard drive. You ring the vendor, they come out and replace the drive. So you now need to rebuild the RAID.

We can perform the following command from the Linux prompt

cat /proc/mdstat

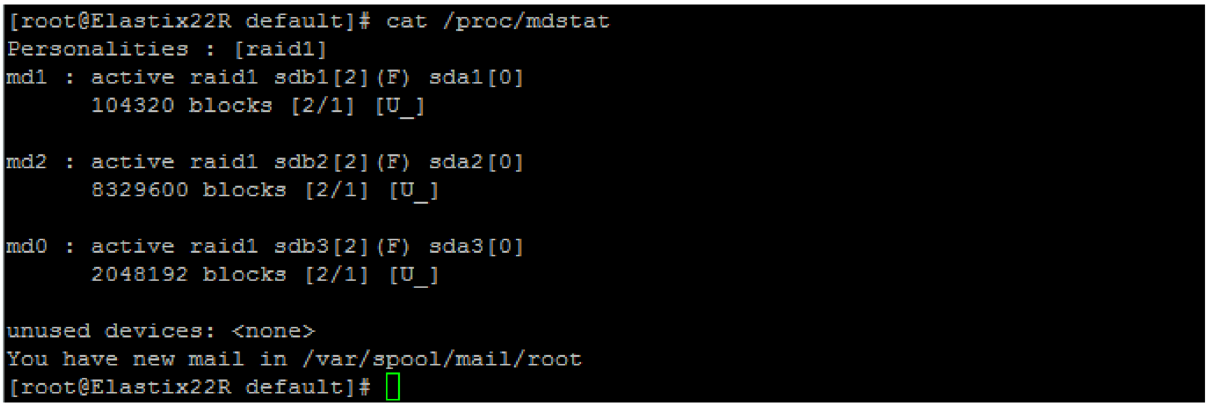

```
[root@Elastix22R default]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb1[2](F) sda1[0]
      104320 blocks [2/1] [U_]

md2 : active raid1 sdb2[2](F) sda2[0]
      8329600 blocks [2/1] [U_]

md0 : active raid1 sdb3[2](F) sda3[0]
      2048192 blocks [2/1] [U_]

unused devices: <none>
You have new mail in /var/spool/mail/root
[root@Elastix22R default]# 
```

As you can see each of the MD sets are now showing (F) for failed, and you can see what has failed (in this case all the partitions on the second drive)

If I check what the system can see, I perform the following command

fdisk –l

And the following screen shows is that the sda drive is running and has a partition table, it shows us that it can see the second drive (the replaced drive), however it complains that **md0**, **md1**, and **md2** have an invalid partition table, which is understandable.

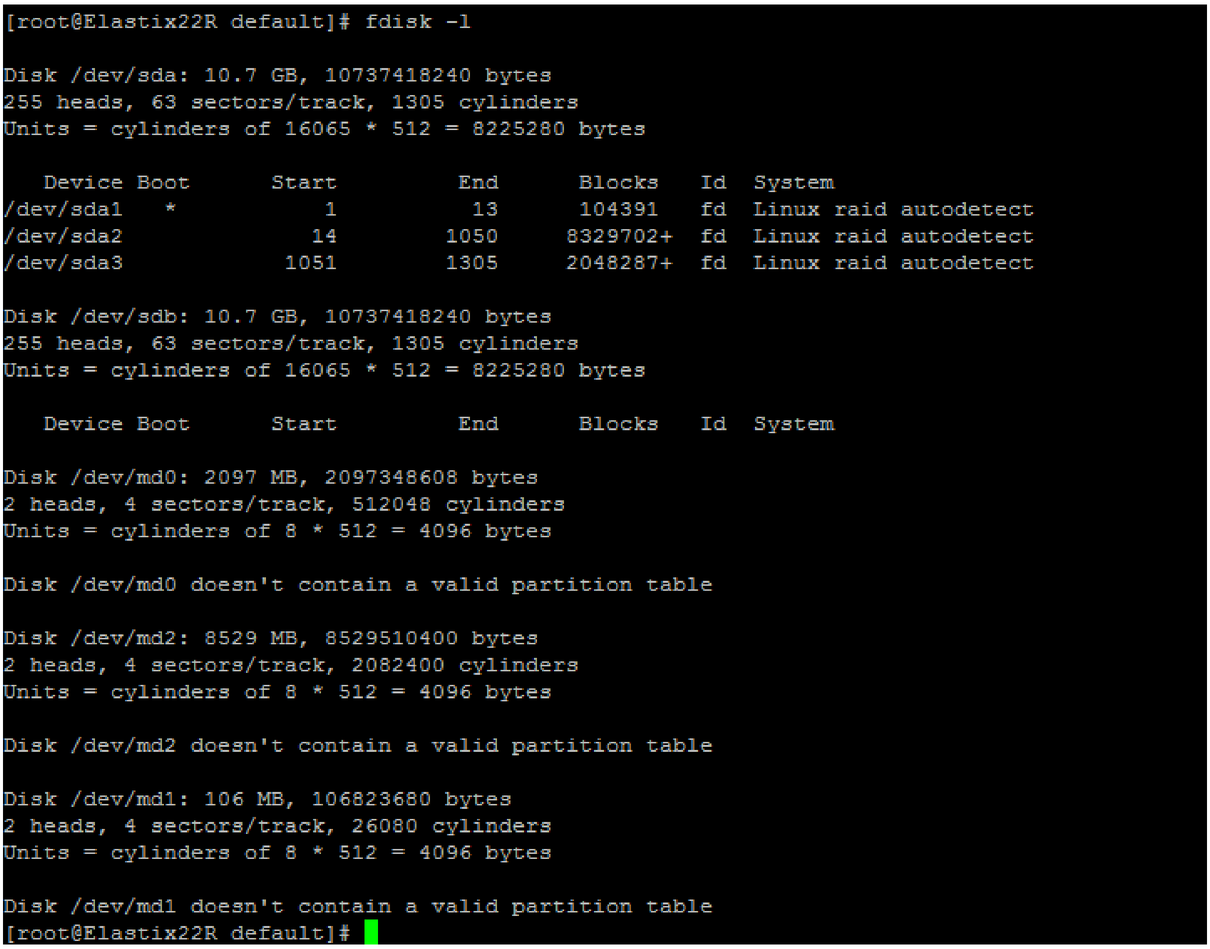

```
[root@Elastix22R default]# fdisk -l

Disk /dev/sda: 10.7 GB, 10737418240 bytes
255 heads, 63 sectors/track, 1305 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1          13      104391   fd  Linux raid autodetect
/dev/sda2              14        1050     8329702+  fd  Linux raid autodetect
/dev/sda3            1051        1305     2048287+  fd  Linux raid autodetect

Disk /dev/sdb: 10.7 GB, 10737418240 bytes
255 heads, 63 sectors/track, 1305 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System

Disk /dev/md0: 2097 MB, 2097348608 bytes
2 heads, 4 sectors/track, 512048 cylinders
Units = cylinders of 8 * 512 = 4096 bytes

Disk /dev/md0 doesn't contain a valid partition table

Disk /dev/md2: 8529 MB, 8529510400 bytes
2 heads, 4 sectors/track, 2082400 cylinders
Units = cylinders of 8 * 512 = 4096 bytes

Disk /dev/md2 doesn't contain a valid partition table

Disk /dev/md1: 106 MB, 106823680 bytes
2 heads, 4 sectors/track, 26080 cylinders
Units = cylinders of 8 * 512 = 4096 bytes

Disk /dev/md1 doesn't contain a valid partition table
[root@Elastix22R default]#
```

So the first step we need to complete is to get a copy of the partition table onto the new replacement drive.

We can do this by a very simple command...

sfdisk –d /dev/sda | sfdisk /dev/sdb

You will then see the following screen if everything is successful.

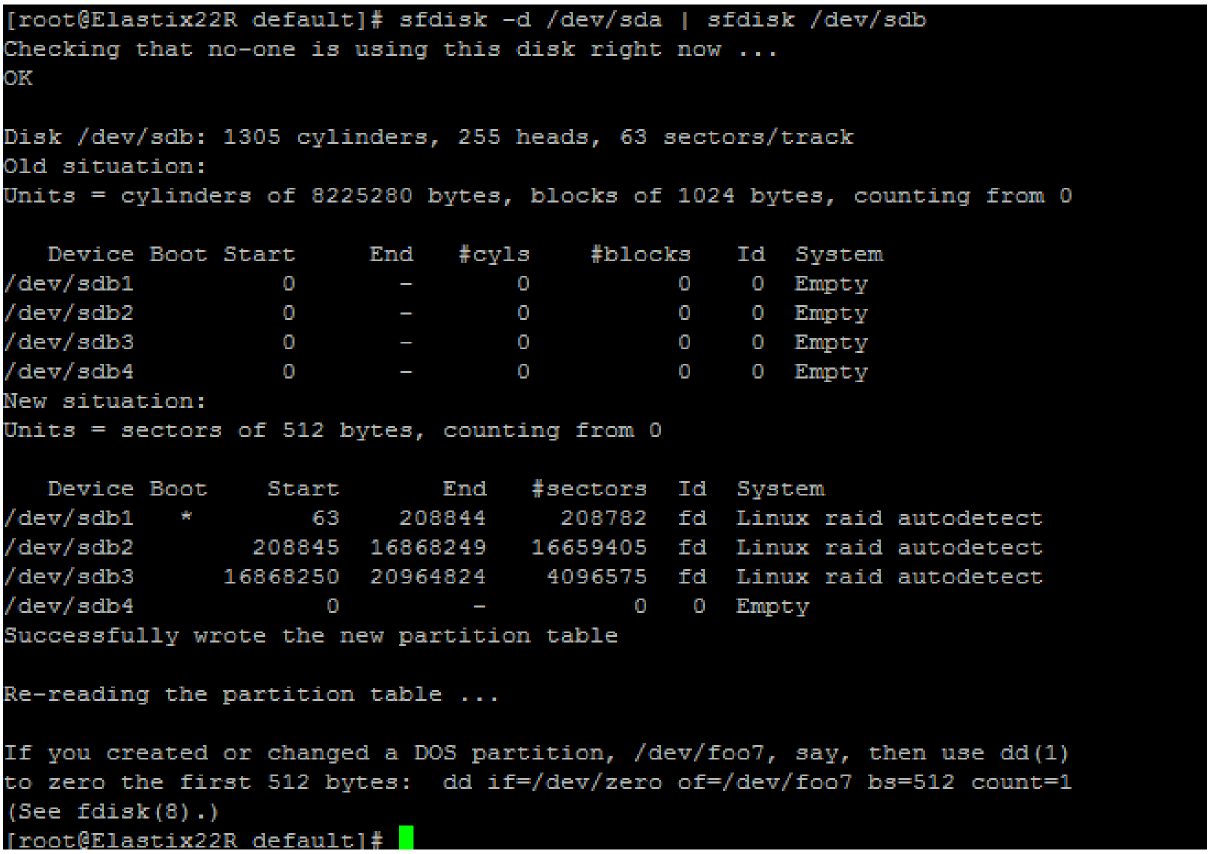

```
[root@Elastix22R default]# sfdisk -d /dev/sda | sfdisk /dev/sdb
Checking that no-one is using this disk right now ...
OK

Disk /dev/sdb: 1305 cylinders, 255 heads, 63 sectors/track
Old situation:
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0

   Device Boot Start     End   #cyls    #blocks   Id  System
/dev/sdb1          0       -       0          0    0  Empty
/dev/sdb2          0       -       0          0    0  Empty
/dev/sdb3          0       -       0          0    0  Empty
/dev/sdb4          0       -       0          0    0  Empty
New situation:
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start       End   #sectors  Id  System
/dev/sdb1   *        63    208844     208782  fd  Linux raid autodetect
/dev/sdb2        208845  16868249   16659405  fd  Linux raid autodetect
/dev/sdb3      16868250  20964824    4096575  fd  Linux raid autodetect
/dev/sdb4             0         -          0   0  Empty
Successfully wrote the new partition table

Re-reading the partition table ...

If you created or changed a DOS partition, /dev/foo7, say, then use dd(1)
to zero the first 512 bytes:  dd if=/dev/zero of=/dev/foo7 bs=512 count=1
(See fdisk(8).)
[root@Elastix22R default]# 
```

If we perform the

cat /proc/mdstat command

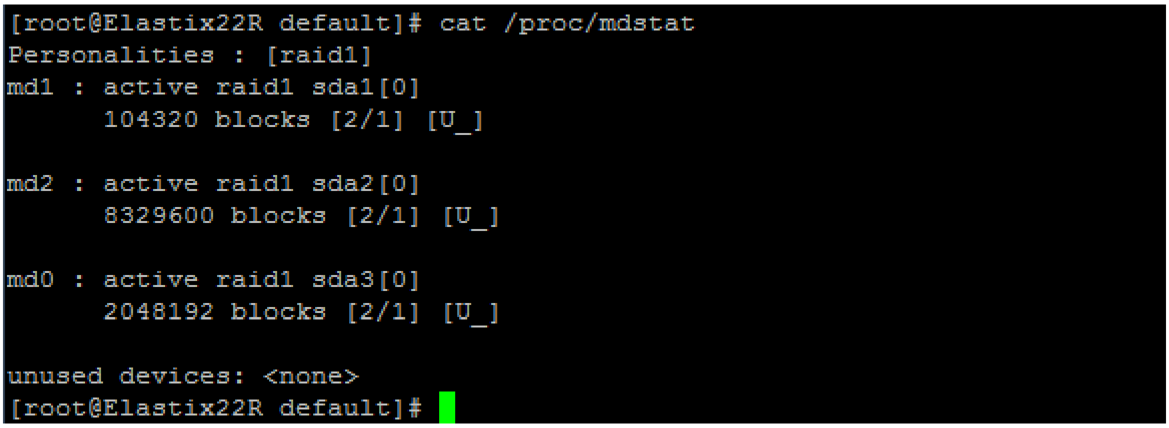

```
[root@Elastix22R default]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sda1[0]
      104320 blocks [2/1] [U_]

md2 : active raid1 sda2[0]
      8329600 blocks [2/1] [U_]

md0 : active raid1 sda3[0]
      2048192 blocks [2/1] [U_]

unused devices: <none>
[root@Elastix22R default]# 
```

We can see that the system is still degraded. Nothing is rebuilding.

We need to tell mdadm to re-add the partitions to the RAID set which we will do with the following commands.

mdadm -a /dev/md0 /dev/sdb3
mdadm -a /dev/md1 /dev/sdb1
mdadm -a /dev/md2 /dev/sdb2

Be careful with the above commands. You need to precisely map the correct (original) partitions. This is part of the reason why I recommend taking a screen shot or copying the details down when you first setup your RAID. Don't panic if you haven't as you can extract this information using fdisk and mdadm, but you need to take your time and be confident that you have them correct.

After you run them, you should see the following.

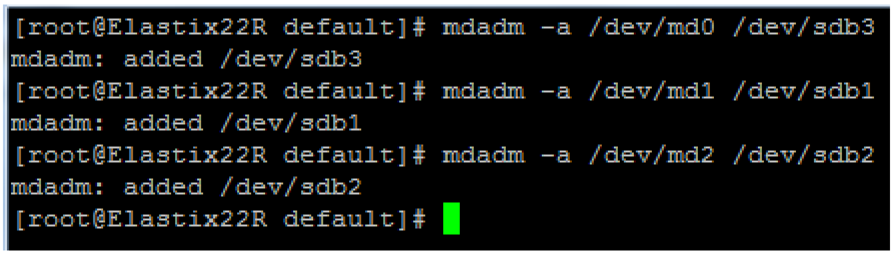

```
[root@Elastix22R default]# mdadm -a /dev/md0 /dev/sdb3
mdadm: added /dev/sdb3
[root@Elastix22R default]# mdadm -a /dev/md1 /dev/sdb1
mdadm: added /dev/sdb1
[root@Elastix22R default]# mdadm -a /dev/md2 /dev/sdb2
mdadm: added /dev/sdb2
[root@Elastix22R default]#
```

Now if you run

cat /proc/mdstat

You should see a similar screen to the one below as it commences rebuilding the RAID set.

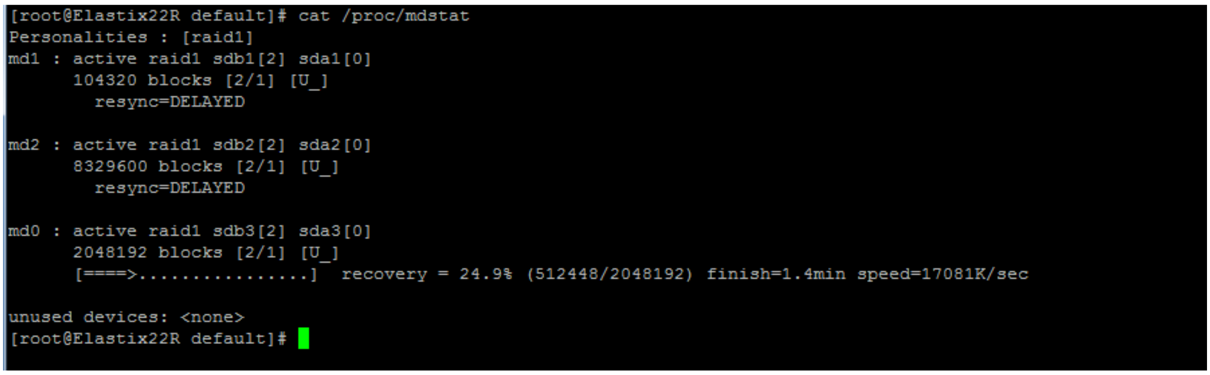

```
[root@Elastix22R default]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb1[2] sda1[0]
      104320 blocks [2/1] [U_]
        resync=DELAYED

md2 : active raid1 sdb2[2] sda2[0]
      8329600 blocks [2/1] [U_]
        resync=DELAYED

md0 : active raid1 sdb3[2] sda3[0]
      2048192 blocks [2/1] [U_]
      [====>................]  recovery = 24.9% (512448/2048192) finish=1.4min speed=17081K/sec

unused devices: <none>
[root@Elastix22R default]#
```

Depending on your Hard Drive size, this can be about 10 minutes, or many hours.

It will rebuild each partition, and once it has completed each partition, it will show the familiar [UU] next to the partition.

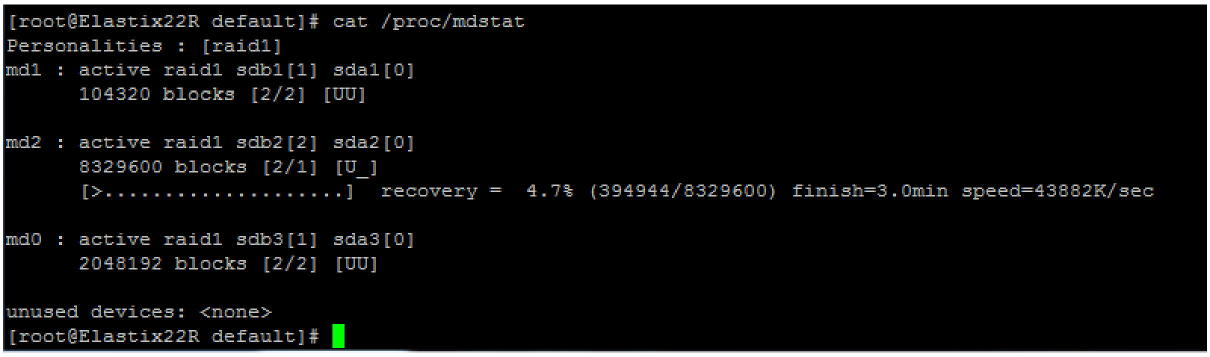

```
[root@Elastix22R default]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb2[2] sda2[0]
      8329600 blocks [2/1] [U_]
      [>....................]  recovery =  4.7% (394944/8329600) finish=3.0min speed=43882K/sec

md0 : active raid1 sdb3[1] sda3[0]
      2048192 blocks [2/2] [UU]

unused devices: <none>
[root@Elastix22R default]#
```

And finally the RAID will be back online

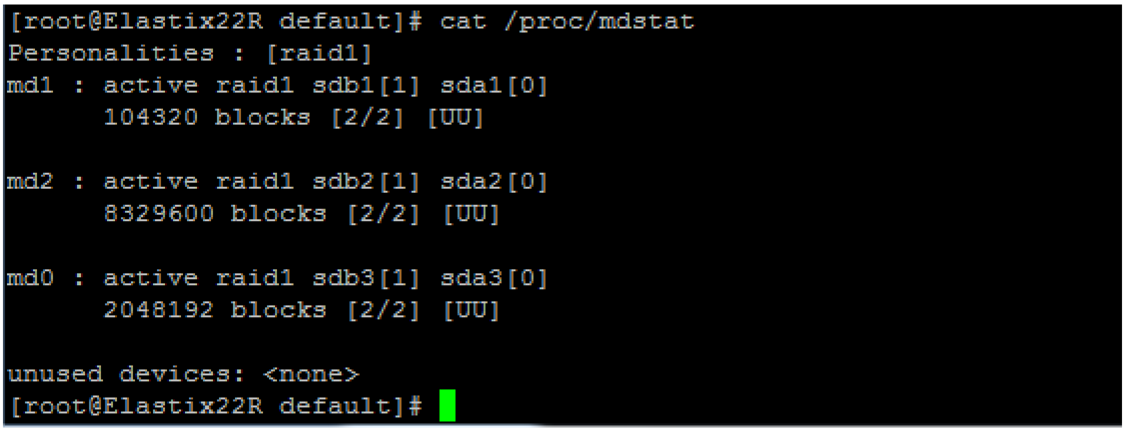

```
[root@Elastix22R default]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb2[1] sda2[0]
      8329600 blocks [2/2] [UU]

md0 : active raid1 sdb3[1] sda3[0]
      2048192 blocks [2/2] [UU]

unused devices: <none>
[root@Elastix22R default]#
```

But there is one more thing to do, especially if you have a replacement drive and that is to install GRUB on the second drive

grub> **device (hd1) /dev/sdb**
grub> **root (hd1,0)**
grub> **setup (hd1)**

We know that we already have GRUB on the first disk, so we are only running the commands to install it on the second disk.

The results should be as per the following screenshot

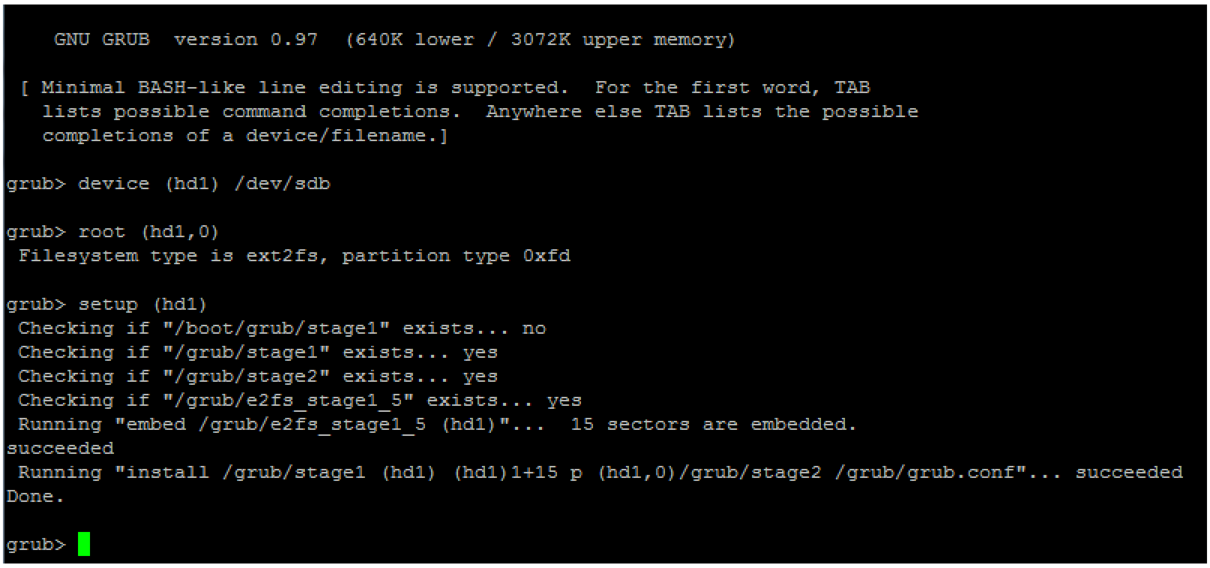

```
    GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

 [ Minimal BASH-like line editing is supported.  For the first word, TAB
   lists possible command completions.  Anywhere else TAB lists the possible
   completions of a device/filename.]

grub> device (hd1) /dev/sdb

grub> root (hd1,0)
 Filesystem type is ext2fs, partition type 0xfd

grub> setup (hd1)
 Checking if "/boot/grub/stage1" exists... no
 Checking if "/grub/stage1" exists... yes
 Checking if "/grub/stage2" exists... yes
 Checking if "/grub/e2fs_stage1_5" exists... yes
 Running "embed /grub/e2fs_stage1_5 (hd1)"...  15 sectors are embedded.
succeeded
 Running "install /grub/stage1 (hd1) (hd1)1+15 p (hd1,0)/grub/stage2 /grub/grub.conf"... succeeded
Done.

grub>
```

This chapter is not meant to be a full run down on how to manage or recover from RAID issues, but to provide a quick insight on how to at least replace a failed hard drive.

Take the time to learn mdadm and what it can do for you. Mdadm can also remove a drive from a RAID set, which actually is useful for learning how to recover. There are many great guides on how to use mdadm, so take the time, have a “play”. It’s better to do it now on a pre-production system or test system, instead of on a live system.

Trademarks

The following trademarks used in these guides and are required to be acknowledged.

Asterisk® is a *registered trademark* of DIGIUM, Inc
FreePBX® is a Registered Trademark of Bandwidth.com
Elastix® is a registered Trademark of Palosanto Solutions

Disclaimer

Your use of these application notes is subject to the following conditions:

- Your application of the information provided is entirely at your own risk
- Whilst tested in a test environment, your environment may be different and the application of these notes may be totally incorrect.
- It is up to you to test in a test environment as to the suitability of these notes.
- You will not hold myself, or any company that I am associated with, responsible for any damages arising from the use of these notes.

Document History

Version	Date	Change
1.0	9th Jan 2011	• Initial Release
1.1	20th Jan 2011	• Included table of contents • Shaunw pointed out that clarification of the partitions changing was needed – thanks. • Added Document History • Added Trademarks