Linux: Getting Started HowTo Guide

A technical howto document presented to H3ABioNet

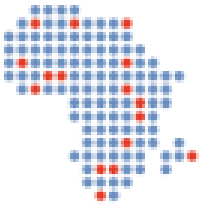

H3ABioNet

Pan African Bioinformatics Network for H3Africa

Created by

The System Administrator Task-force

Prepared for

The greater H3ABioNet and H3Africa Consortium community

Document Control

Date	Author	Authorization By	Version	Description
17 February 2014	Suresh Maslamoney	System Administrator Task-force	1.0	First draft – Dell C6145 Ubuntu 12.04 and RAID installation and configuration
19 June 2014	Suresh Maslamoney	System Administrator Task-force	2.0	Combined all three OS's into one document

Task-Force Members

Last Name	First Name	Institution	Country
Alibi	Mohamed	Pasteur Institute of Tunis (IPT)	Tunisia
Brown	David	Rhodes University (RU)	South Africa
Indome	David	Noguchi Memorial Institute for Medical Research (NMIMR)	Ghana
Scheepers	Inus	Centre for High Performance Computing (CHPC)	South Africa
Maslamoney	Suresh	Computational Biology Group – UCT (CBIO)	South Africa
Panji	Sumir	Computational Biology Group – UCT (CBIO)	South Africa
Van Heusden	Peter	South African National Bioinformatics (SANBI)	South Africa

Reviewers

Last Name	First Name	Institution	Country
Marcello	Lucio		Burkina Faso

Acronyms and Abbreviations

Acronym and Abbreviations	Description
BIOS	Basic Input Output System
I/O	Input / Output
OS	Operating System
POST	Power On Self-Test is performed once the server is switched on. The POST tests the hardware before booting into the OS
RAID	(Redundant Array of Independent Disks) is often implemented to introduce redundancy across the server's internal physical disk drives.

Table of Contents

1. Introduction

This howto guide provides step by step instruction for installing the Scientific Linux 6.4, Ubuntu 12.4 and Debian 7 server operating systems onto your Dell C6145 server. It includes detailed instruction for installing and configuring a software RAID volume using the default RAID utility bundled with Debian and Ubuntu as well as creating hardware RAID volume using the internal PERC H700 RAID controller shipped with the server chassis. The document includes three additional sections which discuss how to navigate the Linux file system, using command line (CLI) text editors and some useful commands to get you started.

The Dell C6145 is referred to as a headless server meaning that it does not come equipped with a built-in DVD ROM. Your installation options therefore are: (a) use an external USB DVD ROM, (b) a USB memory stick or (c) a network based installation (PXE).

The simplest method of installing your OS is via an external USB DVD ROM, hence, this howto guide follows a USB DVD ROM installation.

In this howto guide we will provide instruction for the following:

- Configure the server BIOS to boot from a USB DVD device
- Configure a RAID volume (hardware or software based), and
- Installation of one of the three supported OS's (Ubuntu, Debian or Scientific Linux)
- How to navigate a Linux file system
- How to edit files from the command line
- Useful commands
- To learn more about configuring and securing your server -refer to document two in this series- "Linux: Configuring and securing your server howto guide".

2. Support Contact Information

Table 1 below lists all the support contact details for the C6145 server. Both groups of support personnel will provide both hardware and software support to H3ABioNet consortium members. The H3ABioNet helpdesk will however provide additional bioinformatics support.

Table 1

Vendor	Contact Number	Contact Person	Description
Dell Support	011 709 7729	Call Centre	To access the official dell helpdesk, you would need to supply the server's service tag. This is often made up of seven characters and located on the face of the server
H3ABioNet Helpdesk	helpdesk@h3abionet.org	Helpdesk	Log all calls via the H3AbioNet helpdesk and a support specialist will be assigned to your call

3. Hardware Settings

Table 2 below has been provided to allow you to record your hardware settings:

Table 2

Asset Tag:		**Hostname:**		**DNS name:**	
Operating System:				**Version:**	
RAID Configuration:					
Networking					
Eth0					
IP Address:			**Static or DHCP Assigned:**		
Subnet Mask:		**Default Gateway:**			
Primary DNS:		**Secondary DNS:**			
MAC Address:		**Card Speed:**		**Switch Port:**	
Additional Notes:					
Eth1					
IP Address:			**Static or DHCP Assigned:**		
Subnet Mask:		**Default Gateway:**			
Primary DNS:		**Secondary DNS:**			
MAC Address:		**Card Speed:**		**Switch Port:**	
Additional Notes:					

4. BIOS Configuration

Before you can install your OS, you need to configure your server to use your preferred installation method. The Dell C6145 server does not have an internal DVD ROM, this therefore leaves you with three options of installing your OS; (a) use a USB DVD ROM, (b) use a USB memory stick or (c) install your OS via PXE (network installation). In this step by step guide, we used an external USB DVD ROM for installing the OS.

4.1. First we need to configure your server to boot from USB > switch on the C6145 server > when the Dell logo is displayed, depress the "**F2**" key to invoke the BIOS configuration

4.2. From the BIOS configuration screen > using your keyboard's arrow keys, navigate to the "**Boot**" menu option and then select the "**Boot Device Priority**" option

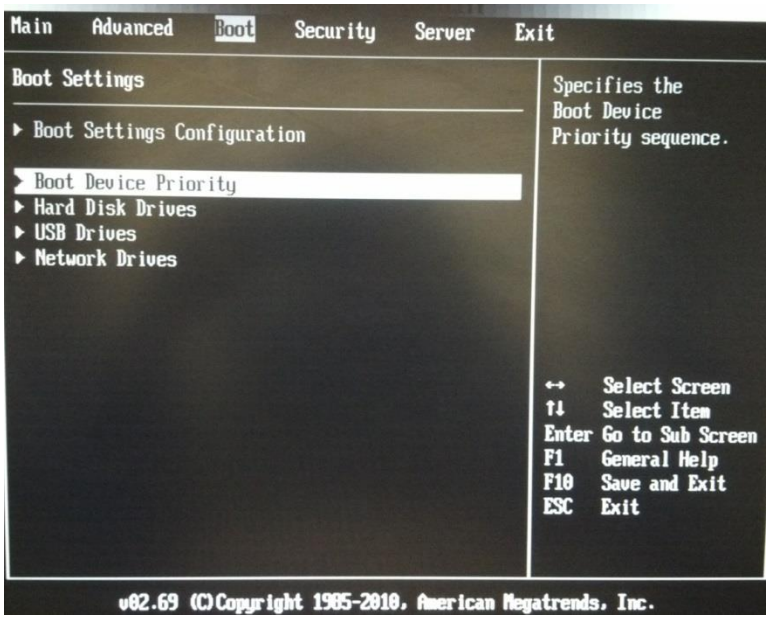

4.3. With the "**Boot Device Priority**" option highlighted, depress the "**enter"** key to access this option > from the pop up screen, scroll to the desired option and depress the "**enter"** key to select your boot device option. For the purposes of this step by step guide, we will be using the external USB DVD ROM method of installation.

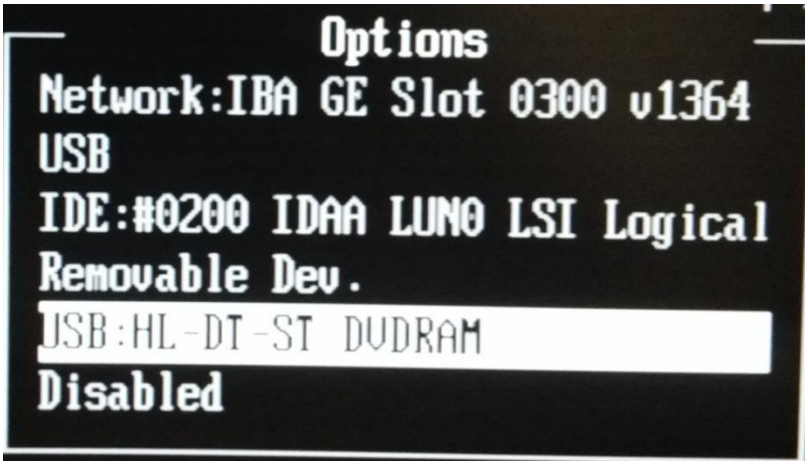

4.4. Now exit and save the configuration by using your arrow keys on your keyboard to navigate to the "**Exit**" menu option > scroll down and select the "**Save Changes and Exit**" option – alternatively you can depress the "**F10**" key on your keyboard to save and exit the BIOS menu.

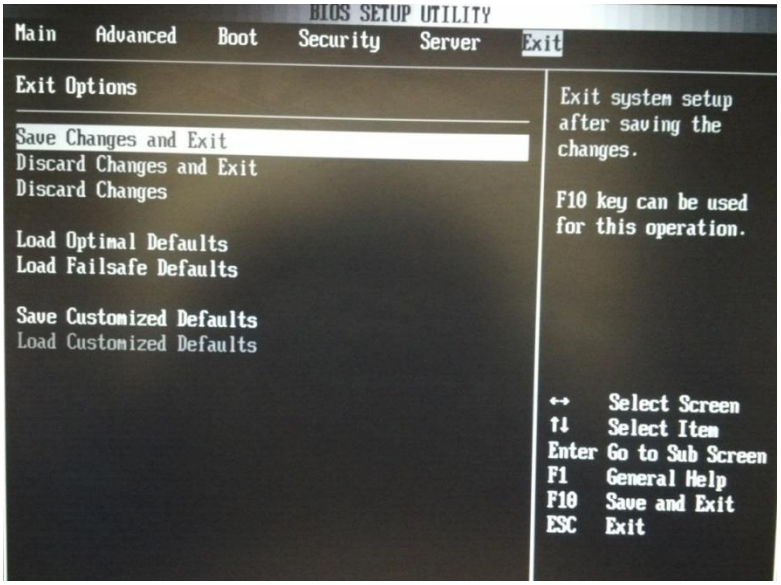

4.5. Excellent, you have just configured your server to boot from your external USB DVD ROM device. The next step is to create a RAID volume on which to install your OS. There are two methods of creating a RAID volume presented in this howto guide. The first method is a software RAID which was used on the Ubuntu and Debian installations where the software RAID utility bundled with the respective Linux OS was used. The second is a hardware RAID configuration which was used in the Scientific Linux installation utilizing the Dell PERC H700 hardware RAID controller supplied with the C6145 server.

5. Redundant Array of Independent Disks (RAID)

5.1. RAID Overview

The acronym "RAID" stands for "Redundant Array of Independent Disks" which ultimately is the process of striping data across multiple physical disks to either provide the OS access to a larger disk volume or to introduce data protection via various levels of fault tolerance. See table 3 below for RAID level descriptions for some of the more commonly used RAID configurations.

Most RAID controllers / utilities make allowances for using a physical disk/s as "hot spares". A hot spare is a physical disk that is not part of any RAID configuration but is used by the RAID volume to automatically replace a failed physical disk drive in the event of a physical disk failure within the specified RAID volume.

Note:
Once you have a physical disk failure, it is of utmost importance that you order your replacement disk drive from Dell and replace any faulty disk drive/s as soon as possible. If you continue to operate your server with the failed disk drive -you lose all disk drive fault tolerance and are at great risk of losing your data contained on the RAID volume.

5.2. RAID Levels

There are multiple RAID levels. The more commonly used RAID configurations are listed in the below table with a brief description for each level

Table 3

RAID Level	Description
RAID0	This RAID level does not offer any fault tolerance. Instead it serves as a way of grouping all of your physical disk drives together and presents a single larger volume to the OS as opposed to individual disk drives. This level offers the fastest read / writes speeds compared to other RAID levels as there is no fault tolerance to be written back to the disks. In the event of a single disk failing, your RAID will become inaccessible and you could potentially lose all data on this RAID. This RAID level requires a minimum of two physical disk drives
RAID1	Referred to as a "mirror" RAID. This configuration requires two physical hard disks and effectively mirrors each other. In the event of one physical disk failing, the mirrored disk drive takes over operation without any disruption to production. The failed disk should be replaced and the mirror rebuilt to restore the fault tolerance provided by this RAID level.
RAID5	RAID5 is configured across a minimum of three physical disk drives and is often referred to as a "striped with parity RAID". This level allows for the failure of one physical disk drive in the RAID without disrupting productivity. Should a second disk fail –the RAID will become inaccessible and all data would potentially be lost.
RAID6	This RAID level is often referred to as "striped with dual parity" and works much same as RAID5 above but allows for a maximum of two disk failures without disrupting productivity.
RAID10	RAID10 is a combination of RAID1 and RAID0. Effectively, the physical disk drives are looped together in a RAID0 configuration and then mirrored to introduce a level of redundancy. This option is often reserved for devices housing many disks and would not prove useful for your C6145 server.

5.3. Hardware RAID vs software RAID

There are pro's and con's to both systems. Hardware RAID controllers manage the RAID volume/s which reduces the I/O burden on the OS and is considered to be faster. The disadvantage however is that in the event of a RAID controller failure, you would need to find the same type of RAID controller to replace the faulty one as other makes or models will often be incompatible. This really only becomes a problem for older servers where the RAID controller is so old, it has been discontinued.

A software RAID is as its name suggests –is built into the OS layer. There is a high I/O burden as the OS has to manage the RAID I/O and is often shied away from. The advantage here though is that if you ever need to replace your RAID, it is easily achieved as the RAID is software based and can be reinstalled at any time.

On the C6145, if you would like to use a RAID5 or RAID6 volume –you will have to use a software RAID as the built-in hardware RAID controller does not accommodate RAID level 5

It is recommended that you use a RAID 1 configuration for your OS installation and RAID 5 or 6 for your data storage area. As you are limited by the number of physical disk drives your server chassis can accommodate, RAID5 would probably be better suited as it will provide sufficient fault tolerance while maximizing your disk storage space.

The major difference between RAID5 and RAID6 is the amount of disks you lose to parity. In RAID5, you lose one disk to parity while in a RAID6 you will lose two disks to parity. The benefits however is that you can survive a two disk failure in a RAID6 and only a single disk failure in RAID5.

6. Hardware RAID Setup and Configuration

The PERC H700 internal RAID controller supplied with your C6145 server can accommodate a RAID0, 1 or 1E/10 configuration. The RAID controller is capable of hosting two separate configurations simultaneously. For example: configuring a RAID 1 for your OS and a RAID 5 for your data storage. The below instruction will guide you through configuring a RAID volume. Simply repeat the steps for each RAID configuration you require.

To successfully create a working hardware RAID volume you would need to first invoke the RAID configuration utility > Choose your RAID option > Select the physical disks that will form part of the RAID > define a hot spare should you prefer having one > Create the RAID > Initialize your newly created RAID > Ensure the RAID you intend booting with is selected as the boot device > Redo to configure a second RAID option or finish to begin installing your server OS.

6.1. Your Dell C6145 server ships with the physical hard disks already installed. Your server should have 5 x 2TB 3.5" Near Line SAS disk drives > Switch on your Dell C6145 server > during the boot process, look for the text displayed in the below image > depress the "**Ctrl**" and "**R**" keys simultaneously to invoke the built-in RAID configuration utility.

6.2. From the virtual disk management screen > navigate down to the "**RAID Properties**" option and depress your "**enter**" key

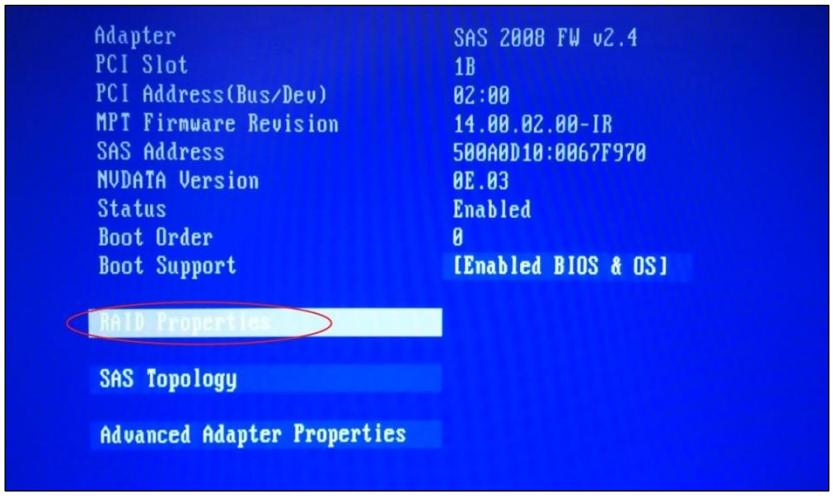

6.3. The next screen will allow you to create a RAID volume or view an existing volume's properties: to view the properties of an existing volume, navigate to the "**View Existing Volume**" option and depress your "**enter"** key > to create a RAID volume, navigate to the desired RAID level and depress the enter key.

Note:
The option to view an existing RAID volume is only displayed if there is an existing RAID volume present. The RAID controller which comes bundled with the C6145 can only accommodate a RAID0, 1 or 1E/10

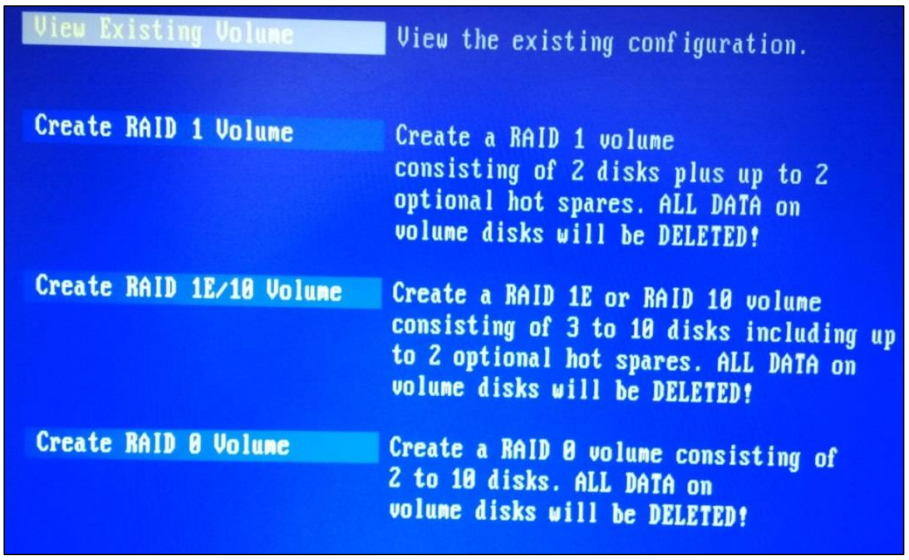

6.4. The following screen allows you to select the physical disks you wish to include in your RAID configuration. Use your arrow or tab key to navigate to the "**RAID Disk**" column and use your space bar or +/- keys to toggle between "**No**" or "**Yes**". Changing this setting to "**Yes**" includes the physical disk into RAID configuration, similarly, selecting "**No**" removes the physical disk from the volume > Once you have chosen all disks you which to include in the RAID, depress the "**C**" character on your keyboard to create the RAID volume.

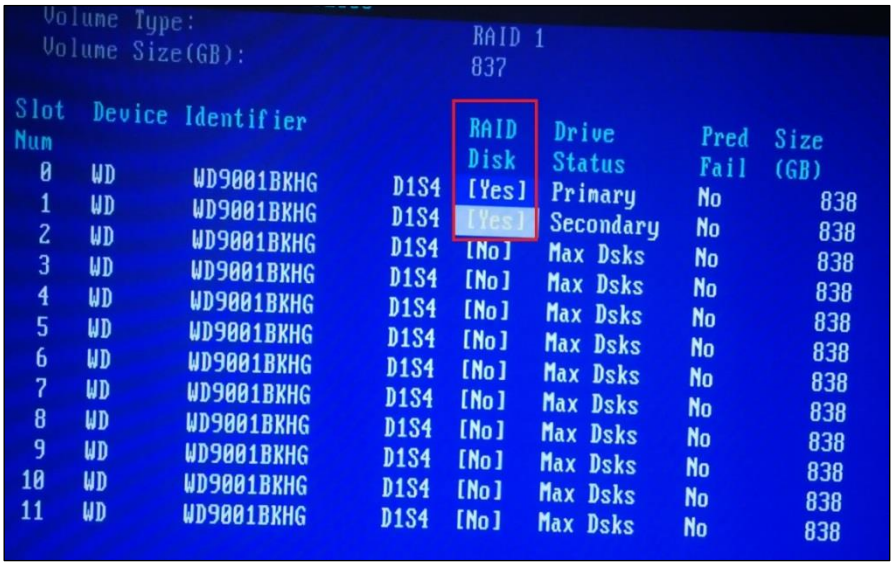

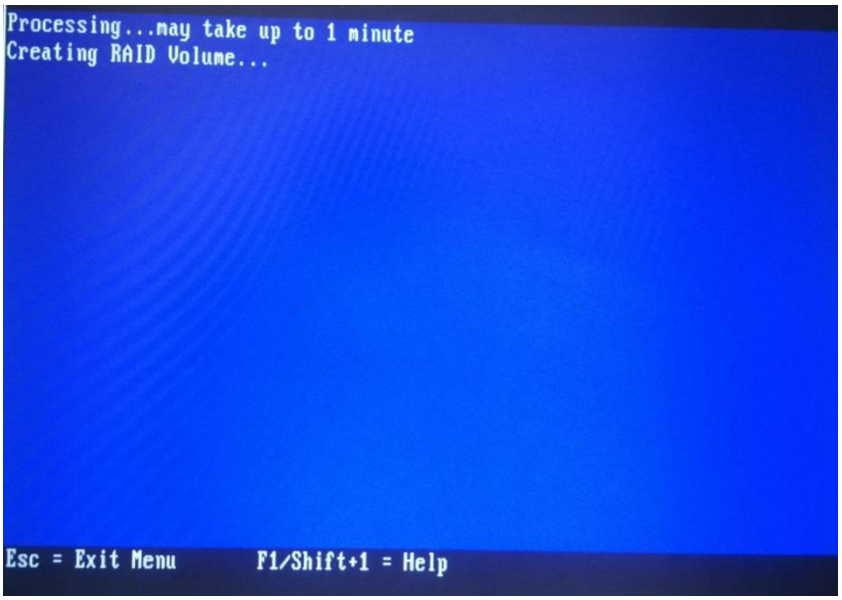

6.5. Once the RAID volume has been built, you will be redirected to the RAID utility's home screen. From this point you can chose to configure another RAID volume or view the properties of an existing volume. Select "**View existing volume**" -the next screen will present a summary of the RAID volume you just created. Information supplied will show the RAID level and the disks included in the RAID. Note the RAID is listed as 0% initialized. You need to first initialize the RAID volume before it can be made available to the server and OS.

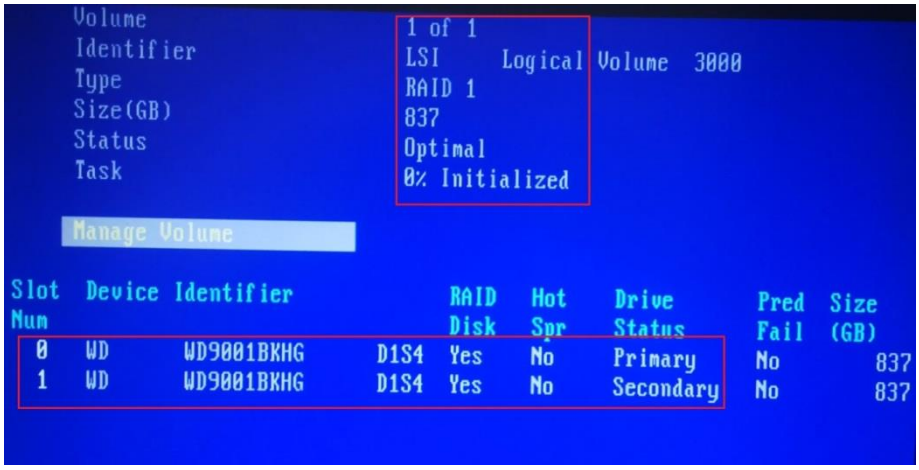

6.6. Navigate to "**Manage Volume**" and depress your enter key. From the manage volume screen you could delete, activate the RAID volume or create and manage a hot spare should you wish to use any.

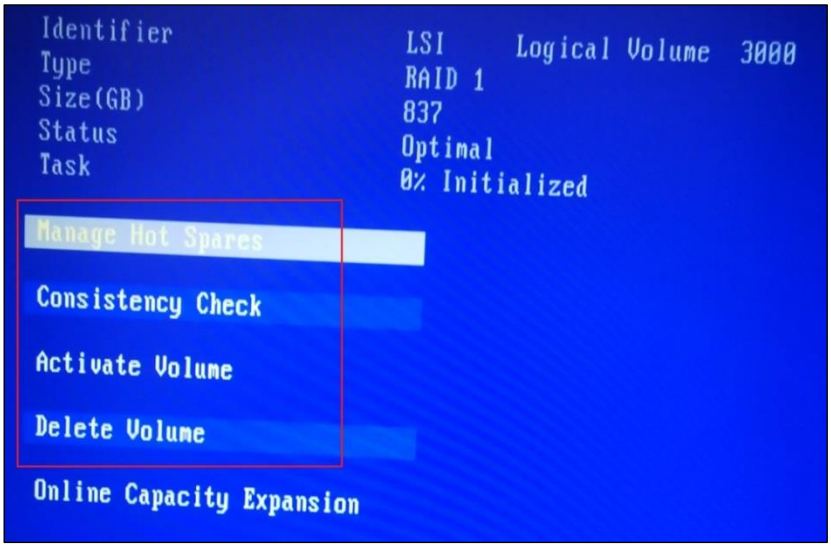

- To define a hot spare > choose the "**Manage Hot Spares**" option from the below screen > navigate to the physical disk you wish to dedicate as the hot spare > tab or navigate via the arrow keys to the "**Hot Spr**" column and depress your space bar or use the +/- keys to toggle between the "No" and "Yes" options. Selecting the "**Yes**" option will configure the selected disk as the nominated hot spare > depress the "**C**" character to commit changes > depress the escape key to return to the previous menu.

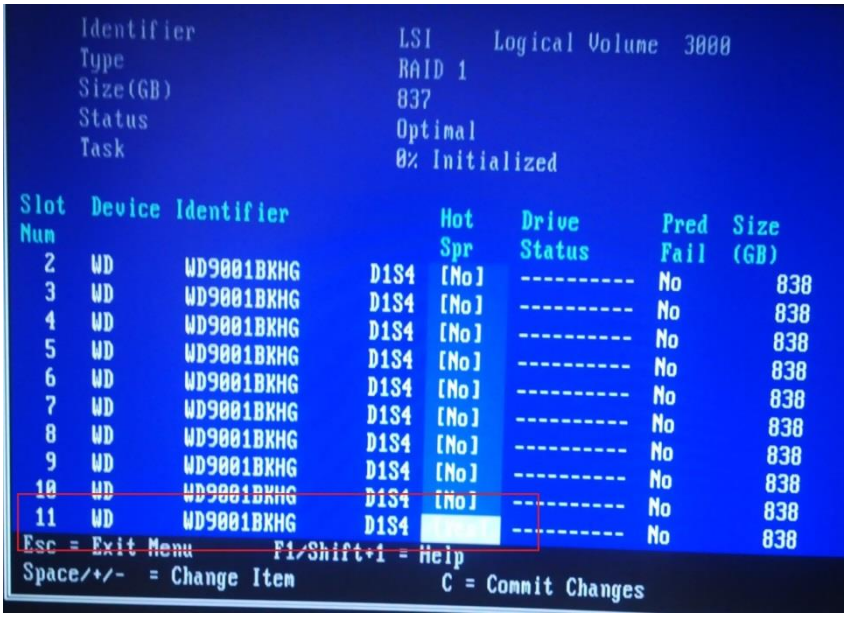

- From the manage volume window choose to active the volume. Once initialized, the newly created RAID volumes will display during the boot process.

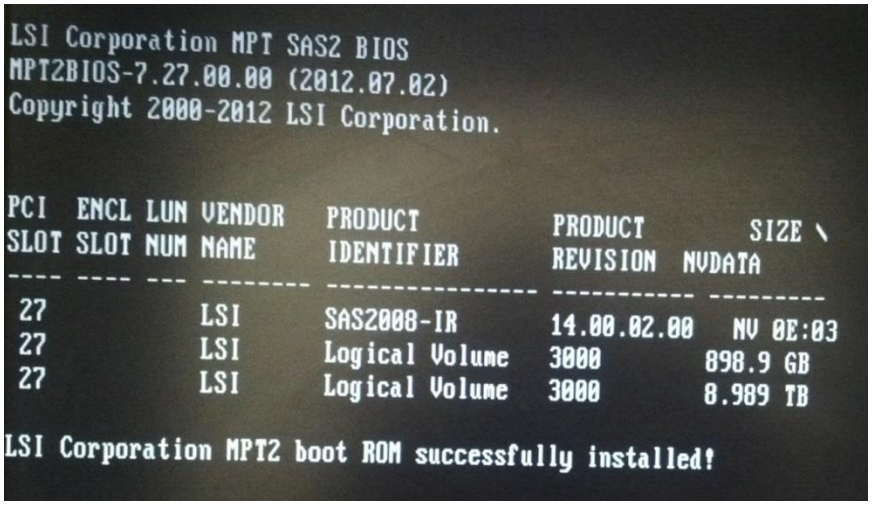

- If you see the above screen during the boot process, you have successfully created a hardware RAID volume. You can now proceed onto the next step - installing your server OS.

Note:
If you prefer to us a software RAID, instructions to setup a software RAID are part of the Ubuntu and Debian installations. Should you wish to configure a hardware RAID, simply ignore the software RAID instructions in the Ubuntu and Debian installation and use the above hardware RAID instructions

7. Operating System Installation

An Operating System (OS) at its most basic definition is a piece of software which installed on a computer system's hard drive. Its purpose is to manage the computer system's hardware and act as a go between the user applications and the hardware present in the computer. An OS is an integral part of any computer system, all applications developed is dependent on an underlining OS to function.

7.1. Choosing a Linux distribution

Linux has many distributions and each distribution have multiple releases of the OS. Each one has its own set of advantages and disadvantages. To simplify support via the H3ABioNet helpdesk, the system administrator task force has identified the following community based Linux distributions: Ubuntu, Scientific Linux and Debian as the officially supported OS's. This decision by no means excludes other Linux distributions; it does however mean that official documentation will only be generated for the above distributions. With that in mind, one of the biggest questions asked when deciding to install Linux is "Which Linux distribution is best for me?" The metrics you use to compare between distributions largely depend on what you want to do with your Linux server. Below is a short list of the more commonly used metrics to determine which distribution is better suited for your need and that of the organization.

Ease of use

This has to be the number one question on everyone's mind when deciding on using Linux as the OS. The term "ease of use" is often loosely used when talking about applications or operating systems but what does it really mean to you? A novice to Linux might find any Linux distribution intimidating and not user friendly while an expert would easily navigate around the file system. The desktop components of an OS are generally much easier to use when compared to their server counterparts, this is largely due to the fact that a desktop system has a graphical user interface (GUI) whereas a server does not. Unlike Microsoft Windows, all interaction with a Linux server is via the command line which can be intimidating at first.

The Debian and Scientific Linux (SL) Linux distributions are not considered user friendly for the novice Linux user as they come shipped with minimal software and generally provide a slower software release cycle which could prove problematic for hardware compatibility. Ubuntu however is a bit more feature rich and maintains a shorter software release cycle. Ubuntu was specifically designed for the Linux newbie while Debian and SL is gear more towards the seasoned Linux user. Out of the box, Ubuntu has a lot more applications and driver support while the Debian and SL adopts a more minimalistic approach to their OS builds. From an ease of use perspective, Ubuntu seems to win this round as you will find an Ubuntu installation a lot easier with all the hardware often detected and supported - this is not always the case with the Debian and SL distributions.

Hardware Compatibility

The next question on everyone's lips after ease of use is often the hardware compatibility of an OS. Different distributions package a range of drivers and useful applications while other package not so useful applications making the OS application heavy. The first step in deciding which distribution to use would be to research and identify if the distribution you intend installing on your server has driver support for all your hardware, next look at what default features it comes bundled with such as system administration, network management applications etc. Please bear in mind that if your

distribution does not have a driver or feature that you need; you could easily acquire the required software via the repository or the manufactures website. The important part is that your OS supports the hardware.

Lucky for us, the Dell C6145 server hardware is compatible with all three distributions we are officially supporting; hence, you can choose either OS. In this instance, all three distributions are winners.

Application and OS stability

Another hot topic is the matter of OS updates and release life cycles. Different Linux distributions have different release and update cycles. Debian for instance does not have pre-set release dates. They do however have at any one point, three versions of a distribution running: Unstable, testing and stable. Unstable is the initial version which is then polished and moved into a testing release and once all the release bugs have been resolved, it is moved into a stable release. This robust testing of the OS before it is released is considered by most the makings of a very stable OS, the disadvantage however is felt in potential hardware compatibility issues.

Ubuntu which is based on the Debian distribution provides releases every 6 months which is supported for 9 months while the long term support (LTS) releases are produced every 2 years and supported for 5 years. While the OS is considered stable due to the Debian base, many consider the OS to be less desirable as a production server. This is largely due to the frequency of releases and the additional applications bundled with the OS.

When choosing Ubuntu as the preferred OS, it is recommended to use the Ubuntu LTS release. When an Ubuntu release reaches its end of life cycle, it received no further update. To review the status of your Ubuntu release, from the command line type the following command:

:: Ubuntu-support-status

Scientific Linux (SL) is based on the commercially available Red Hat Enterprise Linux (RHEL) distribution and much like Debian is considered to be a stable OS. SL closely mimic's the release cycle of RHEL.

With Debian, all previous stable releases are generally supported for 1 year after it has been superseded by another stable release. In short, of the three distributions, Ubuntu offers the longest support life cycle but as it provides releases at 6 month intervals, it's considered less stable compared to its counterparts.

Debian and SL seem to win this round.

Security

Security is another big deciding factor. It is argued that Ubuntu is less secure compared to the Debian and SL distributions – the reason for this seems to centre around the Ubuntu short OS release cycle and the additional applications bundled with their installations to make the release more user-friendly.

Side Note: it is good practice to regularly monitor any open ports and manage the local accounts specifically the ones with sudo rights on your server. This will go a long way in securing your server reducing its attack surface.

Conclusion

In conclusion, the hardware bundled with the Dell C6145 servers is compatible with all three Linux distributions so from a novice point of view and in the context of bioinformatics. It does not make too much of a difference which OS you decide on as none of these OS's come with pre-bundled applications for bioinformatics use.

The following three sections look into the actual installation process of the three officially supported Linux distributions.

7.2. Ubuntu 12.04 OS Installation

This tutorial was completed using the Virtual Box application. For the purposes of this guide, a virtual machine was created with 3 virtual hard disk drives and 4GB RAM. Should you require the latest version of Ubuntu server, the iso can be obtained from the following website >> http://www.ubuntu.com/download/server.

Note:
This tutorial doesn't deal with the issues that arise when working with GPT partition tables

Let's get started…

- Make sure your external USB DVD ROM is connected to your server's USB port > insert your Ubuntu 12.04 server installation DVD and switch on your server.
- Following a POST test, the installation wizard might prompt you to choose a language for the wizard installation > select your preferred language for the OS and installation.

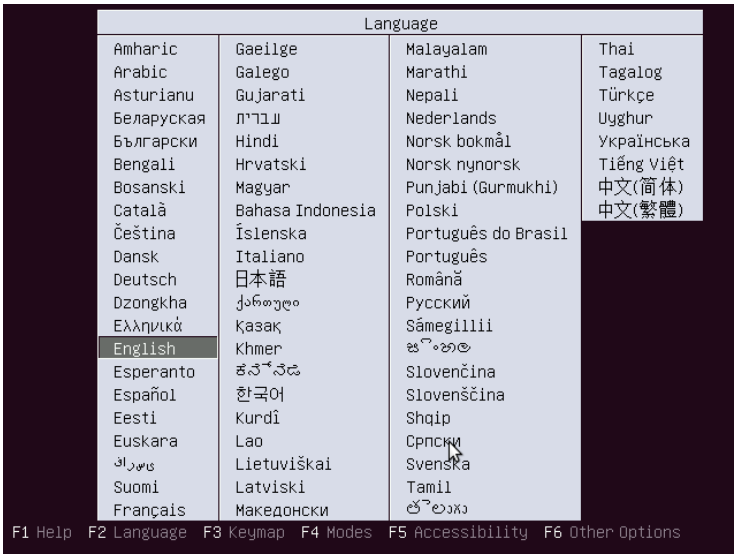

- The next screen will present some options to you > Select "**Install Ubuntu Server**" from the list of options to begin your OS installation.

- The wizard will prompt you once again to select your language of choice > choose your preferred language to continue.

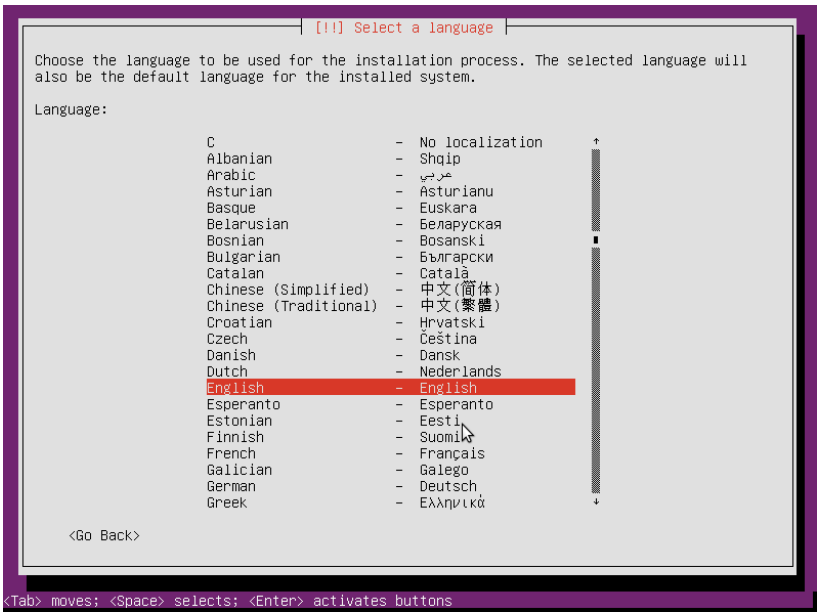

- Next, select the country you reside in or where the server will be physically located.

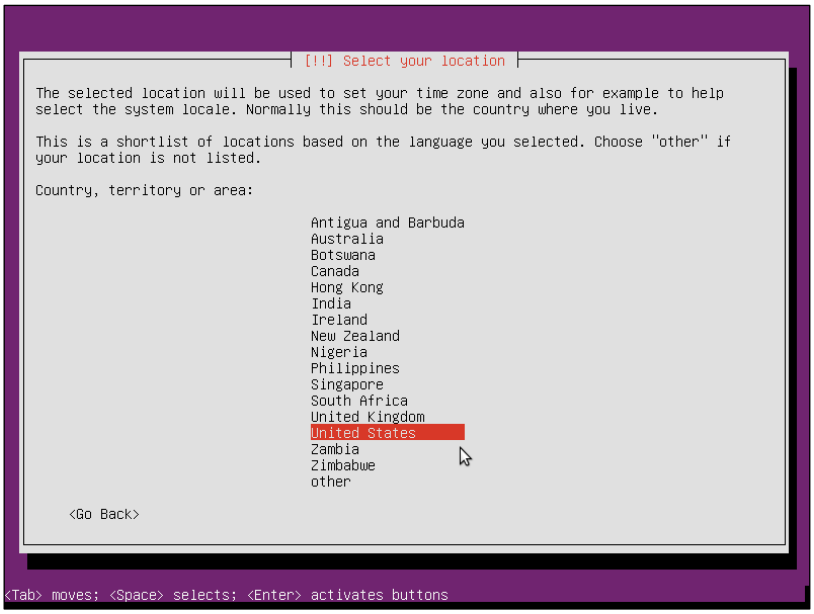

- Next you are prompted to configure your keyboard. You could choose "**Yes**" to allow the installation wizard to detect your keyboard, should you choose "**Yes**" –you will be presented with some additional wizard screens to determine your keyboard type. I often select the "**No**" option.

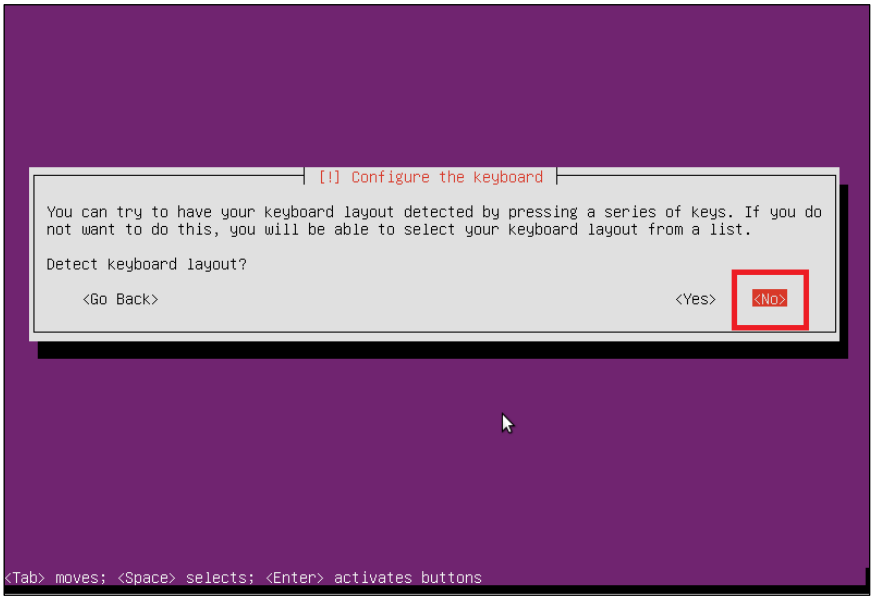

- Choosing "No" will prompt you to select a keyboard layout > select the keyboard layout of your choice to continue with the wizard

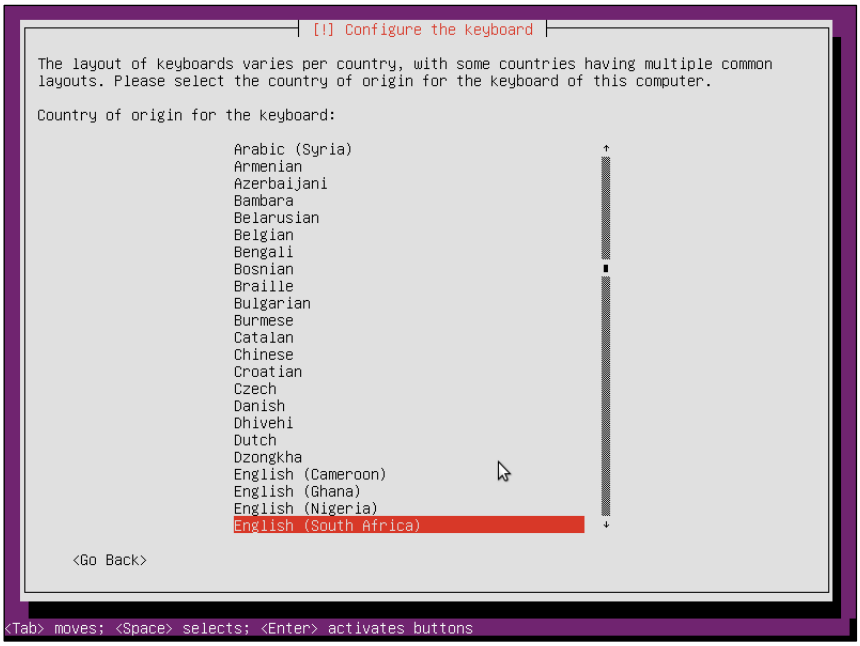

- Based on your selections thus far, the wizard will begin detecting your hardware and loading the necessary drivers.

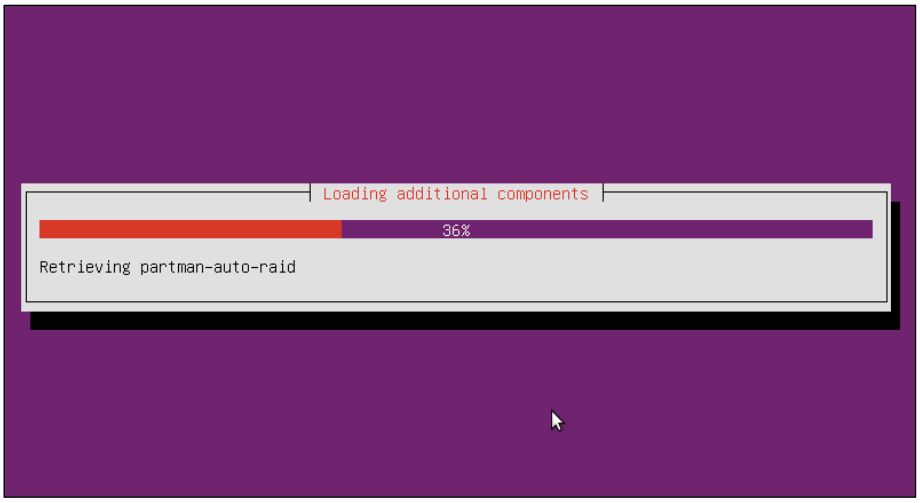

- Once the drivers have been loaded, you will be prompted to give your new server a name. Chose a hostname or alternatively leave the default of “Ubuntu” for now. This can be edited later post the OS installation.

TIP: *It’s less hassle to simply set the hostname now. Try and choose a short descriptive name for the server instance.*

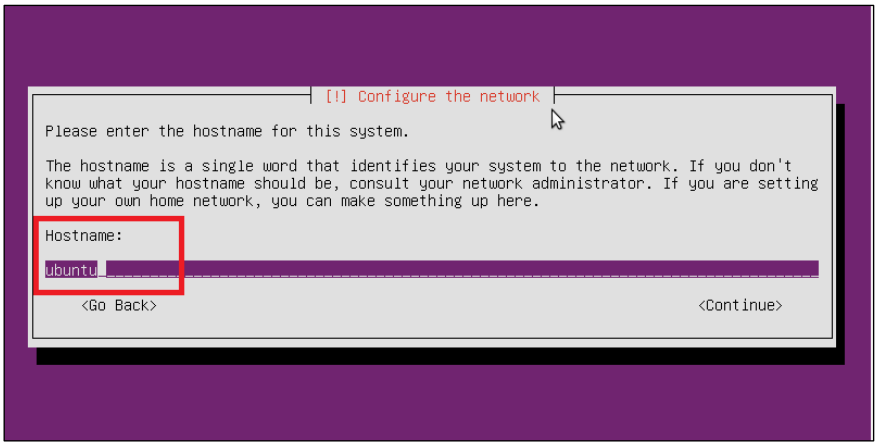

- Next you will be prompted to create a user account to perform non-administrative activities. The first screen prompts you for the full name of the user while the second screen prompts you to choose a username. The wizard chooses the first name supplied in the previous box as the username. You could leave this as is or choose a specific username for the account.

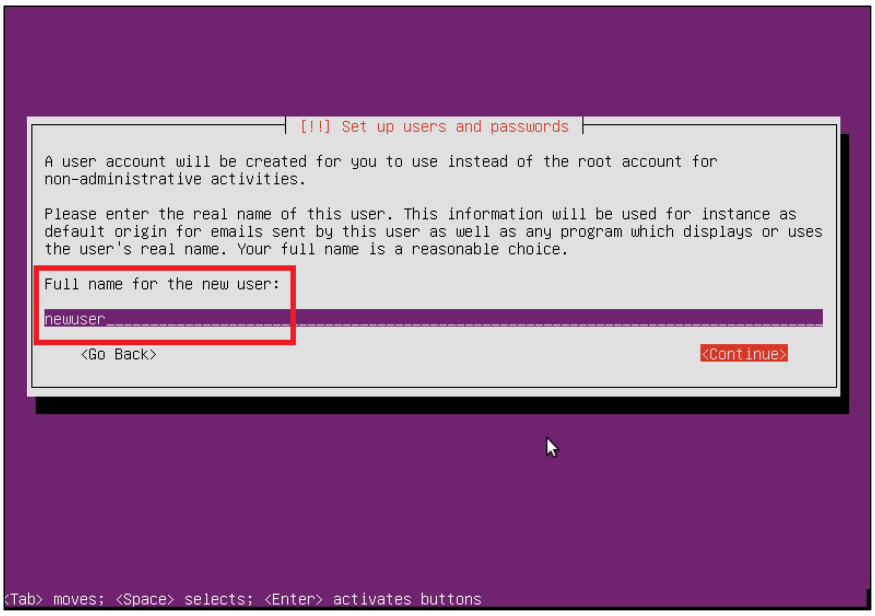

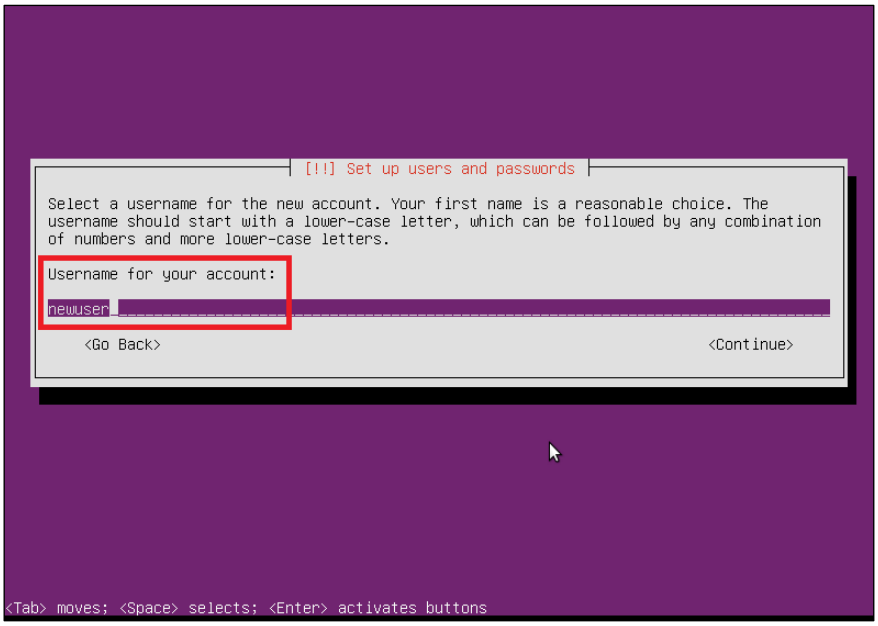

- Choose a password for the user account and click "**Continue**" to proceed. The wizard will prompt you to confirm the password by retyping it > retype the password and click "**Continue**" to proceed.

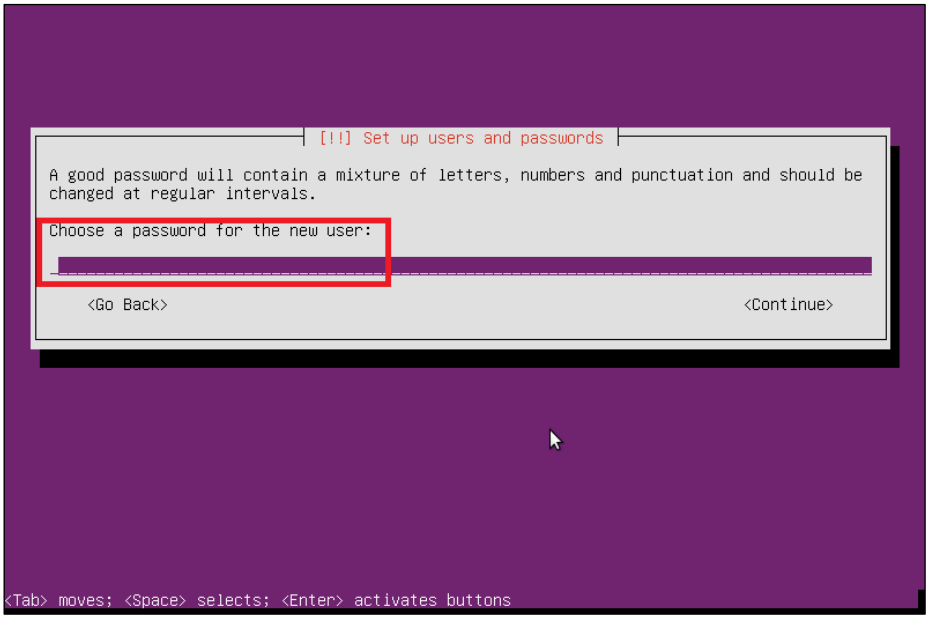

- The next screen will offer to encrypt the home directory of the newly created user account. I always select the "**NO**" option.

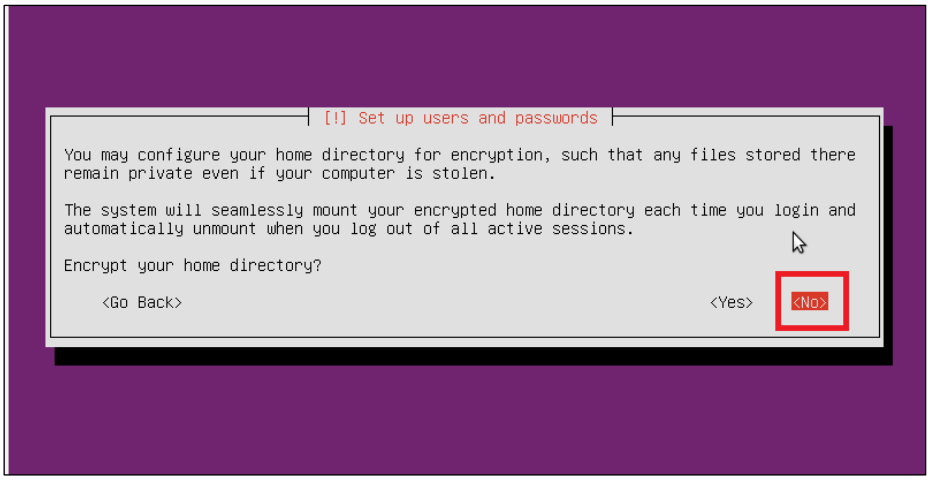

- You will then be allowed to select your time zone (Ubuntu may detect this automatically if connected to the Internet).
- You will then be presented with the "**Partition disks**" window.

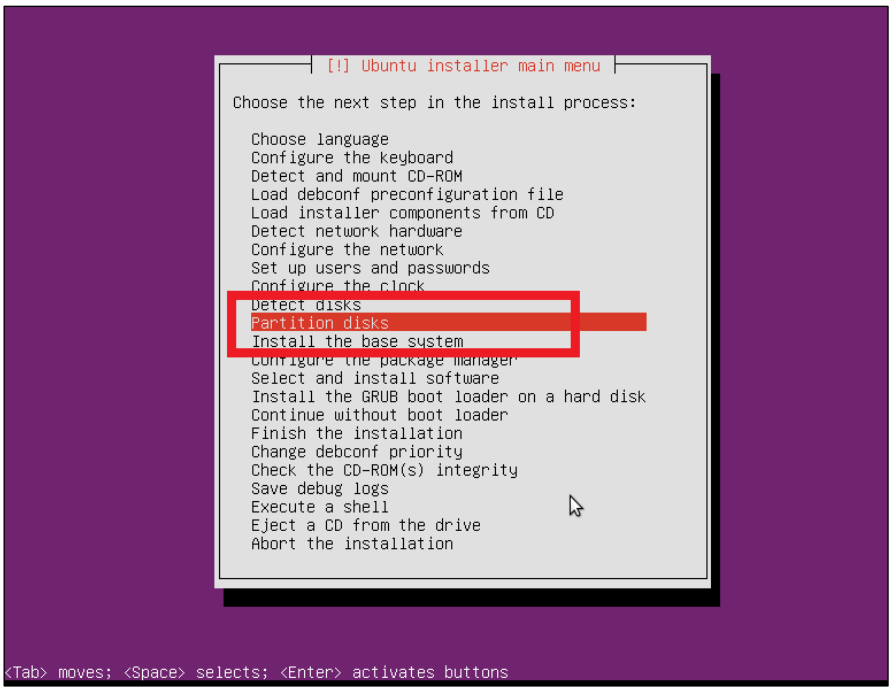

- The next screen presents the "**Partition Disks**" screen. This is the stage where we configure a RAID volume. Should you prefer to install the OS on the hard disk without implementing any RAID fault tolerance, simply highlight the first disk drive usually represented as sda. The wizard will guide you through a few screens where you will be allowed to partition the disk or use the entire disk for the installation. The purpose of this howto guide is to demonstrate the RAID volume setup so we will select the "**Manual**" option as we would like to partition our hard drives ourselves.

TIP: In Linux, y*our hard disk drives are usually represented as sda, sdb, sdc,….sdn*

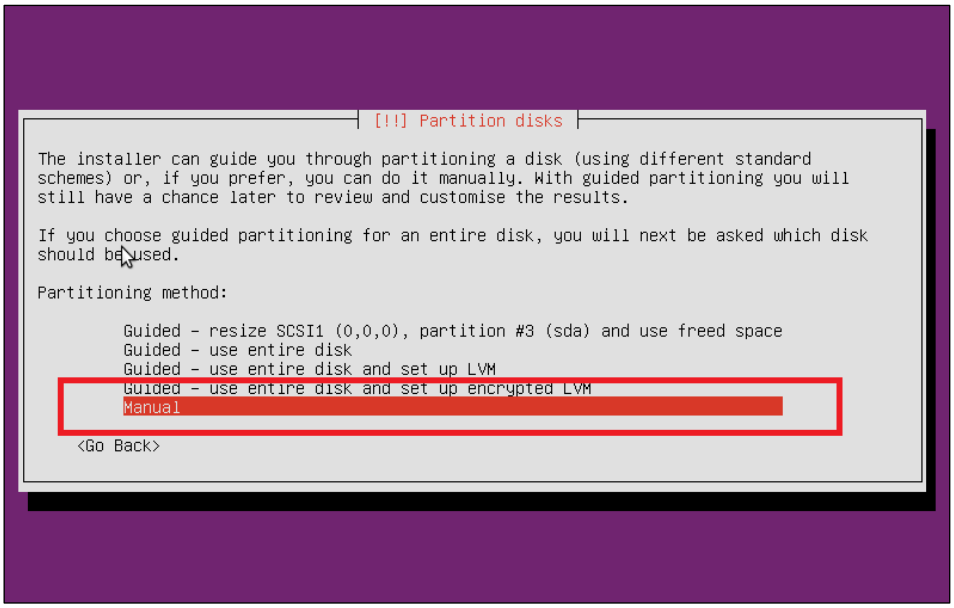

- The wizard prompts you with a warning that all data will be lost if you repartition the disk drive. As this is a new installation, there shouldn't be any existing partitions or data existing on the hard disk. Choose the "**Yes**" option to continue.

Note:
If you are partitioning a disk drive that has an existing partition or data residing on the disk. It is highly recommended that you first backup all data on the disk before continuing. Failure to do so will result in all your data being deleted

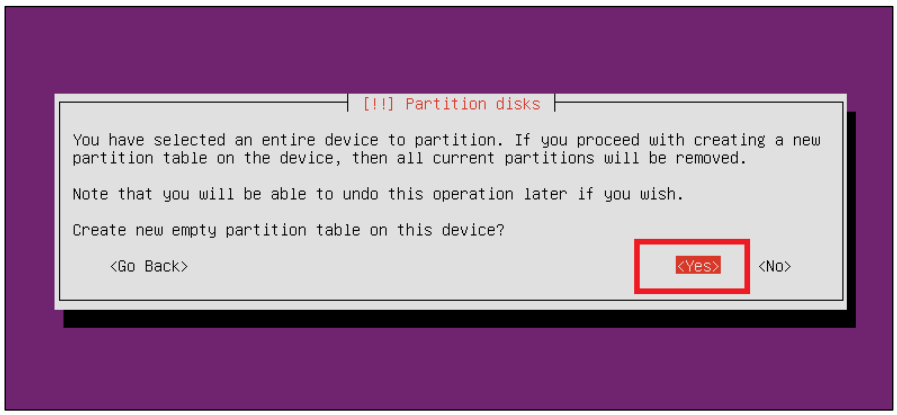

- You will be presented with a screen that shows you your hard drives with the partitions currently on them. If you are installing an OS for the first time, your hard drives will not have any partitions on them. Go down to your first hard drive (it should be labelled something along the lines of "**SCSI (0,0,0) (sda)**". Push enter and you will be asked if you would like to create a new partition table on that hard drive. Select "**YES**".

- Repeat the above steps for all your hard drives. Your partitions should now look similar to the below image (The amount of disk drives will vary based on your system)

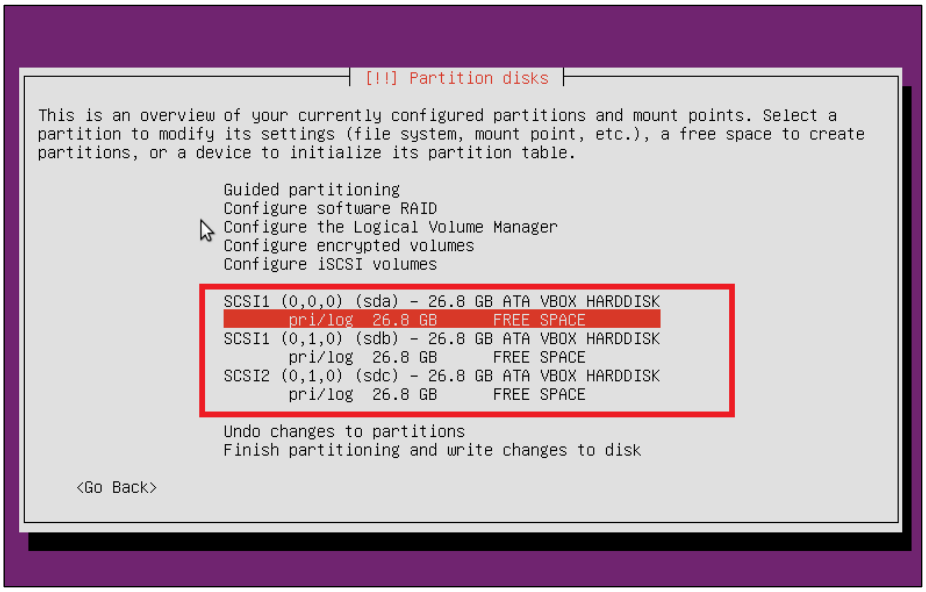

- Now that we have partition tables on each drive, we need to carve up the partitions. We are going to create two partitions on each drive. Both partitions will be RAID'd. One RAID partition will then be set up as the swap space and the other as the root file system.

Note:
This tutorial was done on virtual box and I was unable to recreate the problems I had doing the installation on the actual C6145. There, problems arose due to the partition tables being GPT partition tables (I'm assuming Ubuntu automatically sets this when using big HDDs). This is, because the MSDOS tables cannot handle partitions greater than 2TB. When using GPT tables you need to create a small 50MB partition at the beginning of each drive as a "Reserved BIOS boot area". GRUB must then be installed to this partition

- Select the free space below your first hard drive and push enter > you will be presented with the below screen > select "**Create a new partition**".

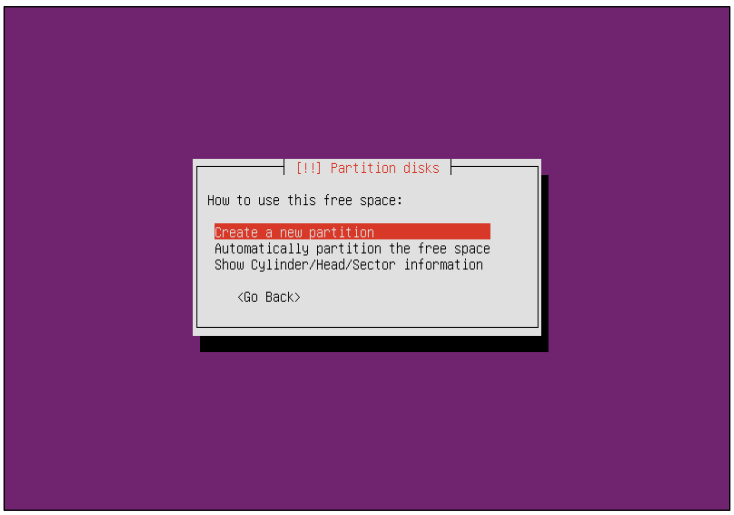

- I set this partition to 4GB as I want a partition of 8GB for swap space after I've RAIDed the

partitions. You may set this size to whatever suits your needs.

- Select “**Primary**” for the partition type.

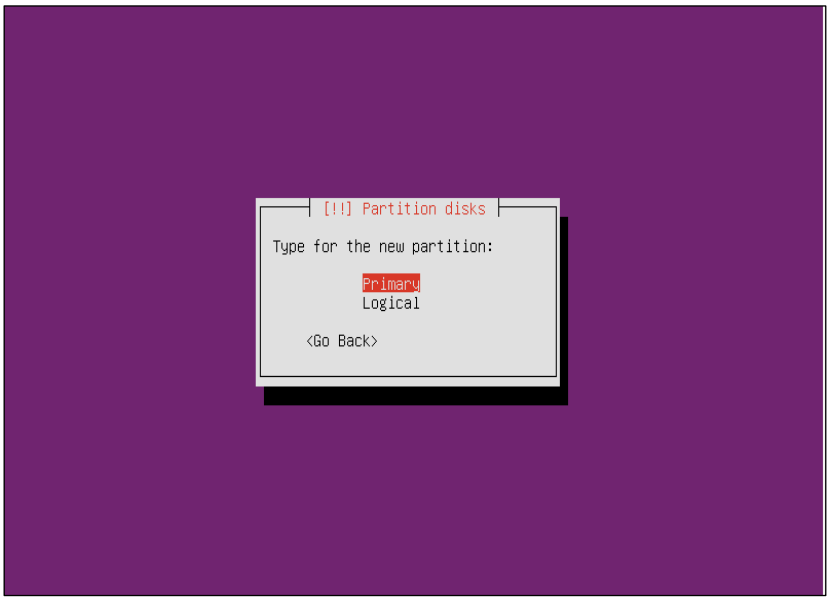

- Select “**Beginning**” for the location of the partition

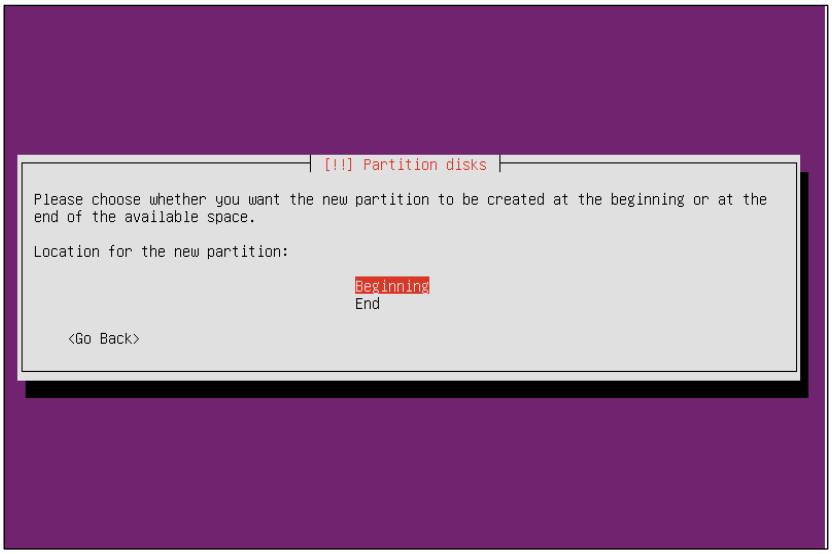

- On the next screen, when prompted how to use this new partition, select the "**Physical volume for RAID**" option and then select "**Done setting up the partition**"

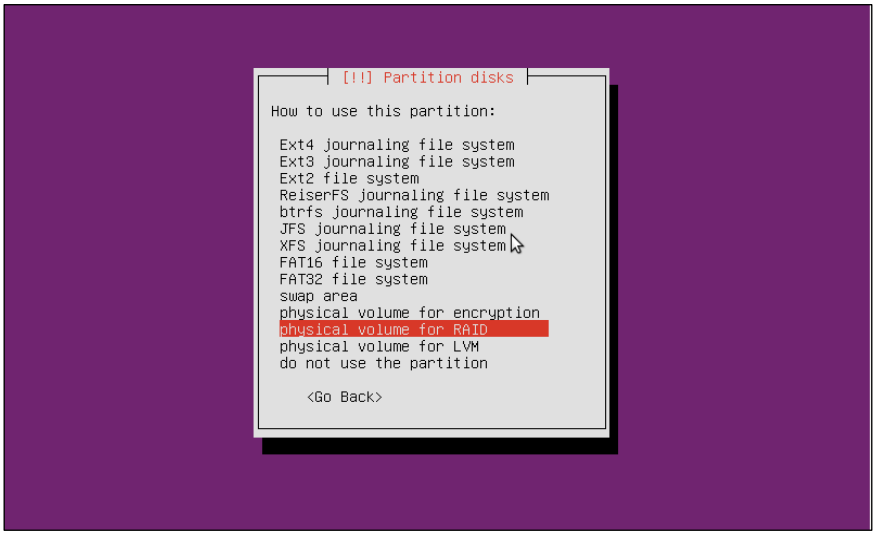

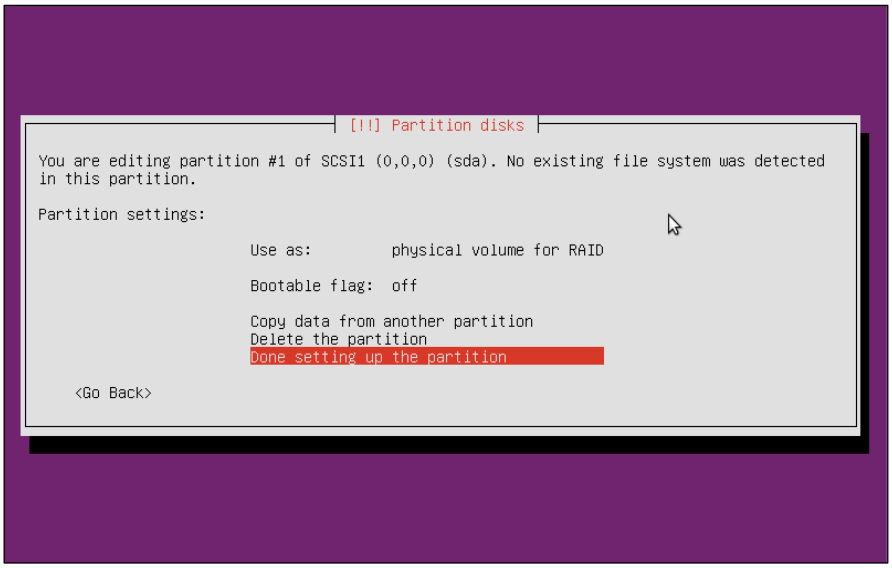

- Now repeat steps from 7.1.20 through to 7.1..24 to create a second partition on the first hard drive, but this time -allow this partition to use the remainder of the space on the drive. Once again, select the "**Physical volume for RAID**" option when prompted
- Your first hard drive should now have two partitions, each set to "**Physical volume for RAID**".
- Repeat steps 7.1.20 through to 7.1.24 for each hard drive. Your partitions should look similar to the below image (The amount of disk drives will vary based on your system)

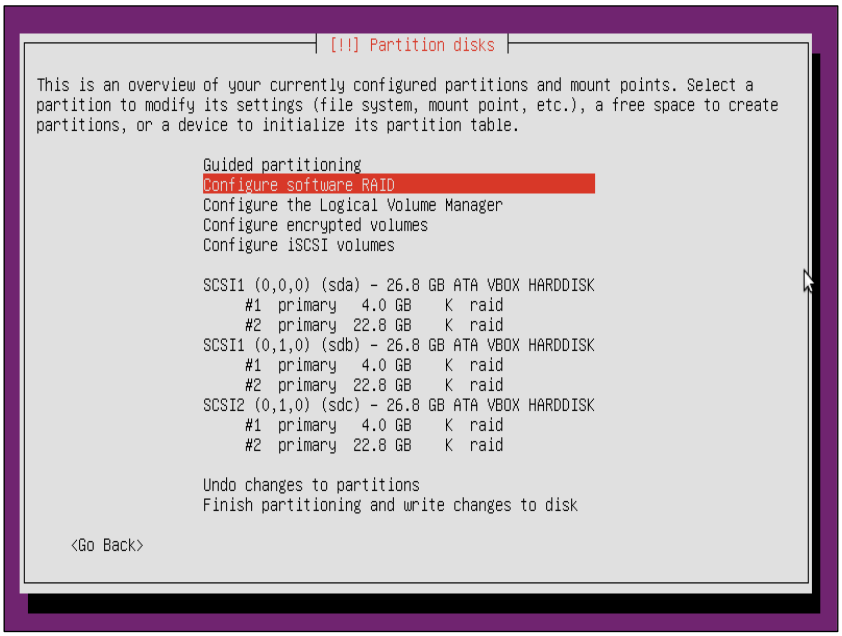

- You have just prepared all your physical disk drives for your RAID setup. It is now time to configure our RAID. Select "**Configure Software RAID**" from the options in the above > select "**Yes**" when asked if you would like to write the changes to the storage devices.
- From the list option on the following screens > select the "Create MD device" option.

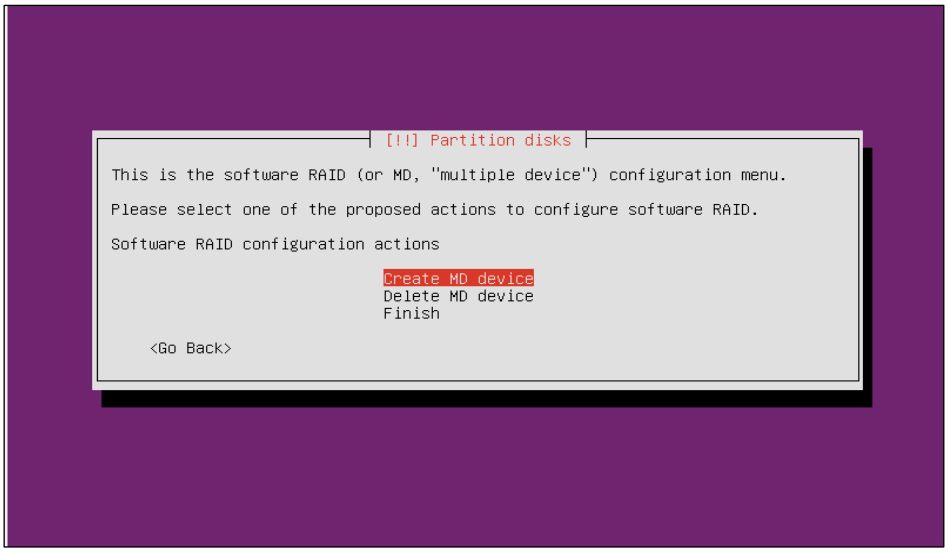

- Select “RAID 5” from the list of options.

Note:
You may want to select RAID 1 here if you are worried about data security. RAID 1 will mirror all your data across drives. This means you will have half the amount of usable space as drives i.e. if you have 4 x 2TB HDDs, you will only have 4TBs worth of usable space instead of 8TB. This data will be mirrored on the other drives, so you have backup should any drive fail

- Select the number of active devices for the RAID 5 array. The minimum number of physical disks for a RAID 5 volume is 3. If you have more than 3 drives, you may wish to leave one out at this point to use as a spare in the next screen.

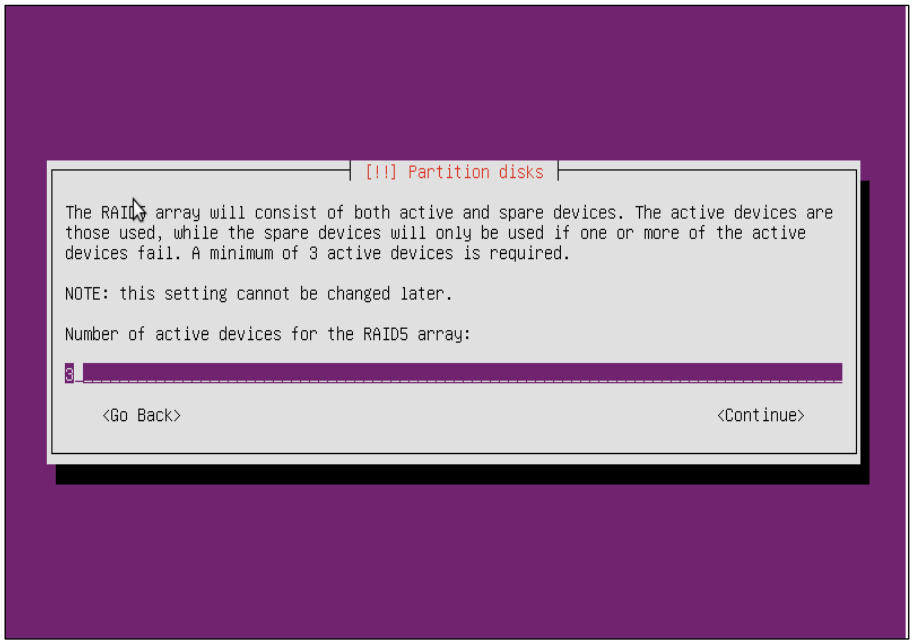

- Select the number of drives you want to use as spares. If you used all your drives as active devices in the last screen, leave this as 0.

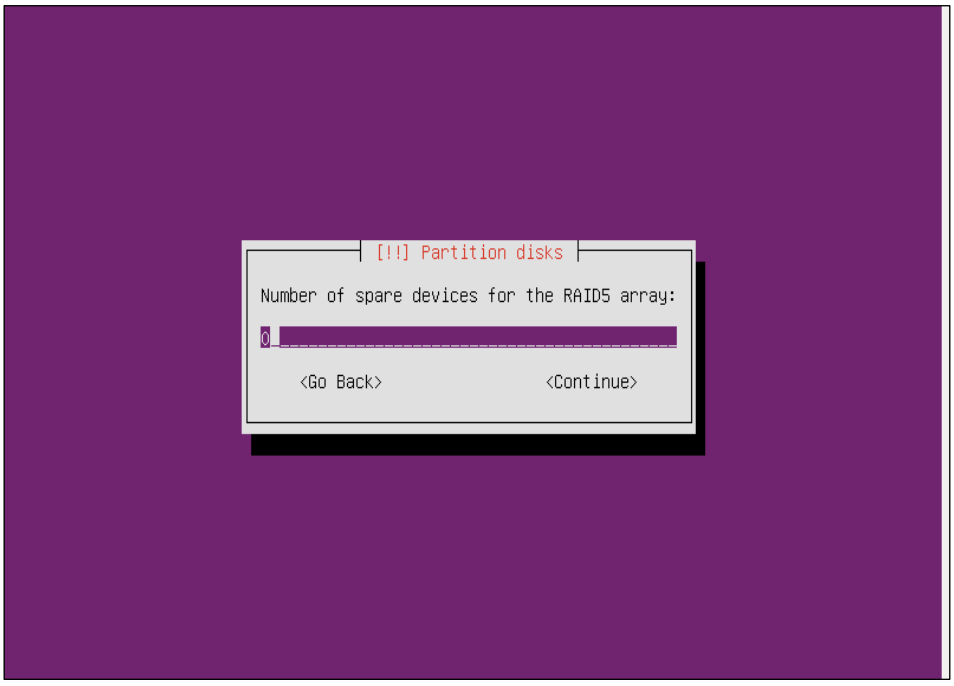

- Select the partitions you set out as the swap partitions. These will most likely be /dev/sda1, /dev/sdb1, /dev/sdc1, …, /dev/sd<n>1, and continue.

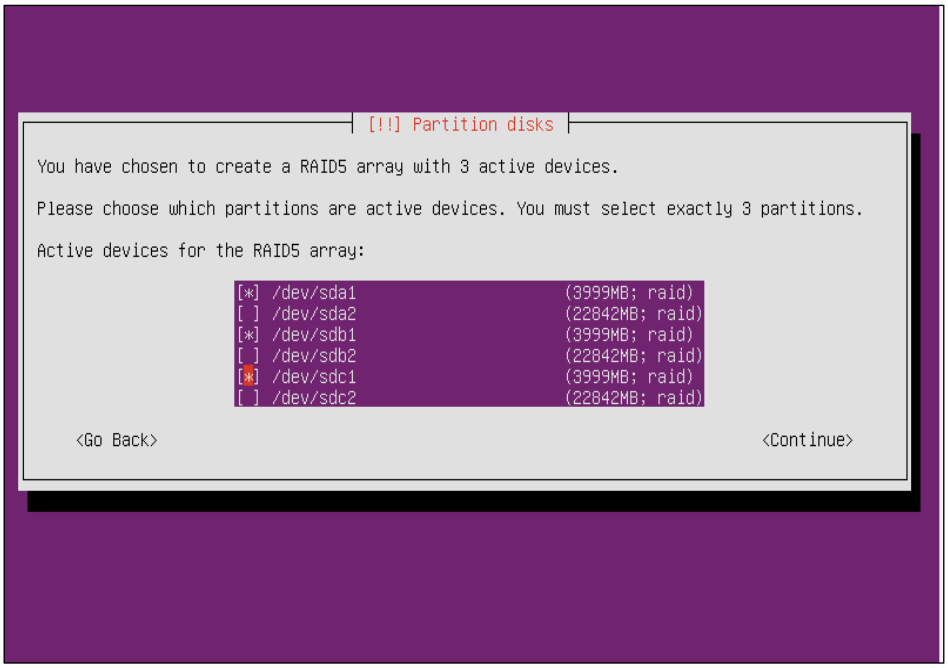

- Repeat steps from 7.1.28 through to 7.1.33 to set up your RAID 5 partition that will be used as the "/" root of the file system. This RAID will use the remaining partitions.
- Select "**Finish**" to go back to the partitions screen

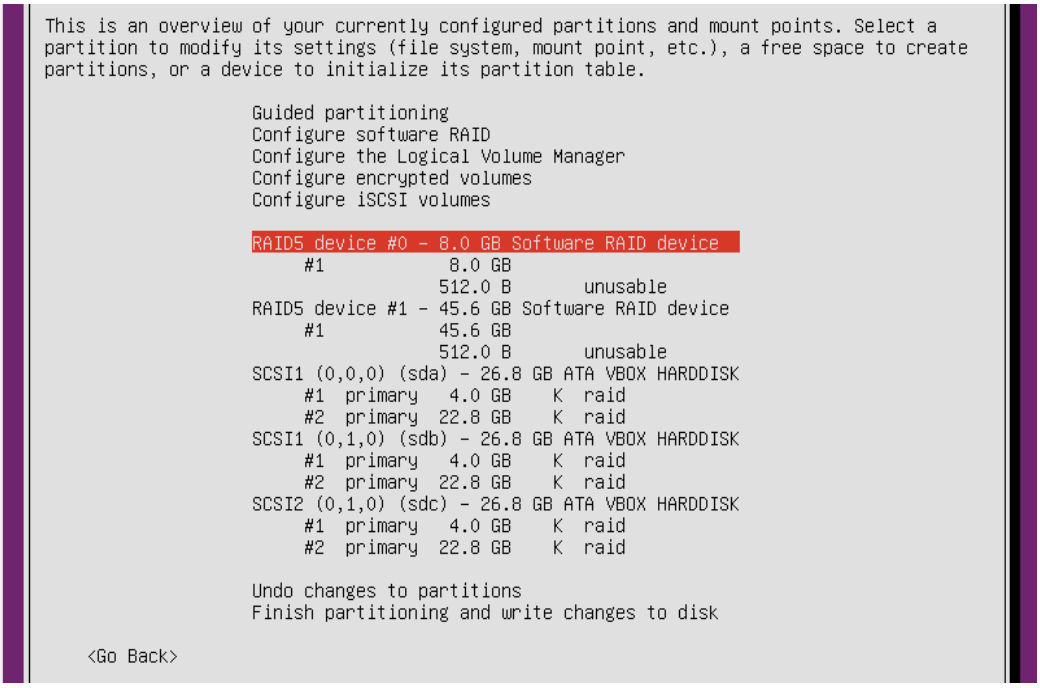

- You should now see two RAID 5 devices above your hard drives similar to the above image. Select the partition on the first RAID device. You will see the below screen:

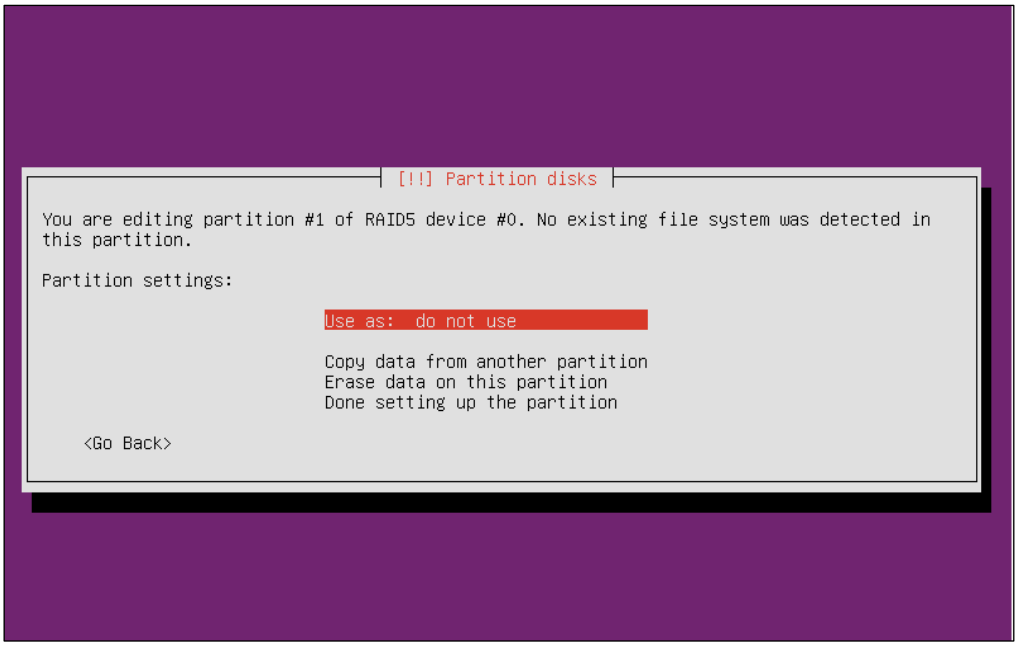

- Select the “swap area” from the list of options when asked how to use this newly created partition.

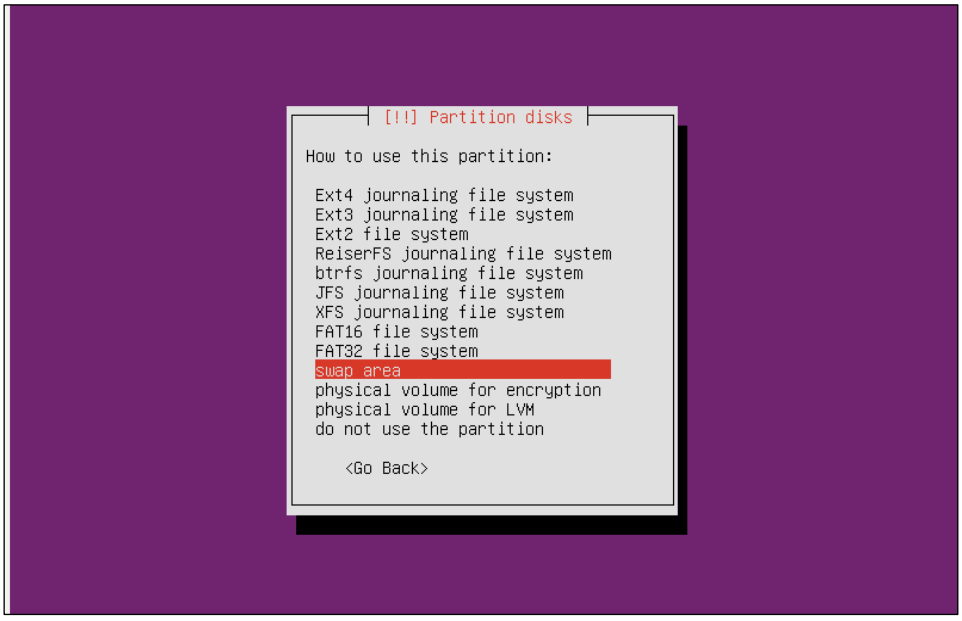

- Go back and select the partition on the second RAID device > select the “**Ext 4 journaling file system**” option and set the “**Mount point**” to “**/”** - root of the file system. The partition should look like the below image

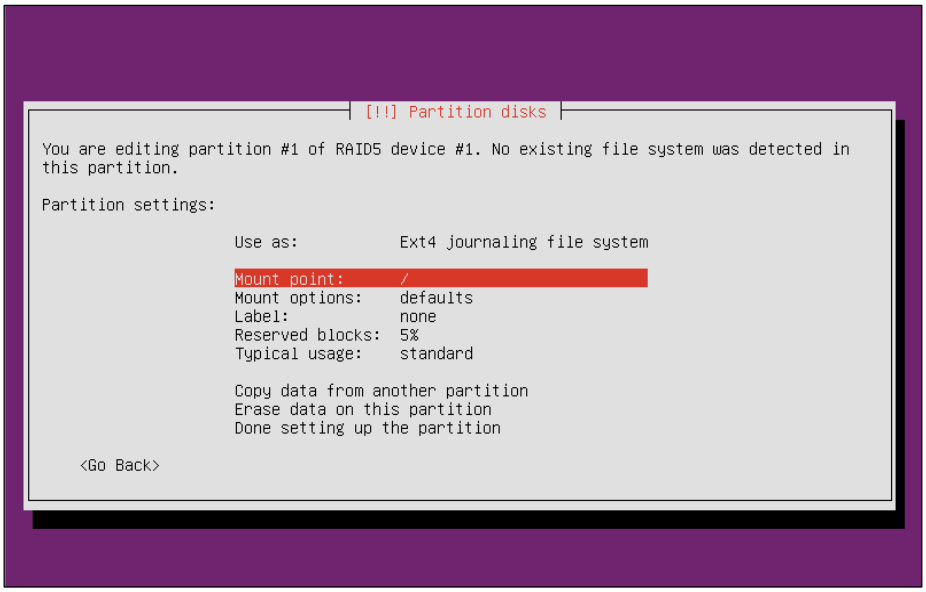

- You partitions should look similar to the below image.

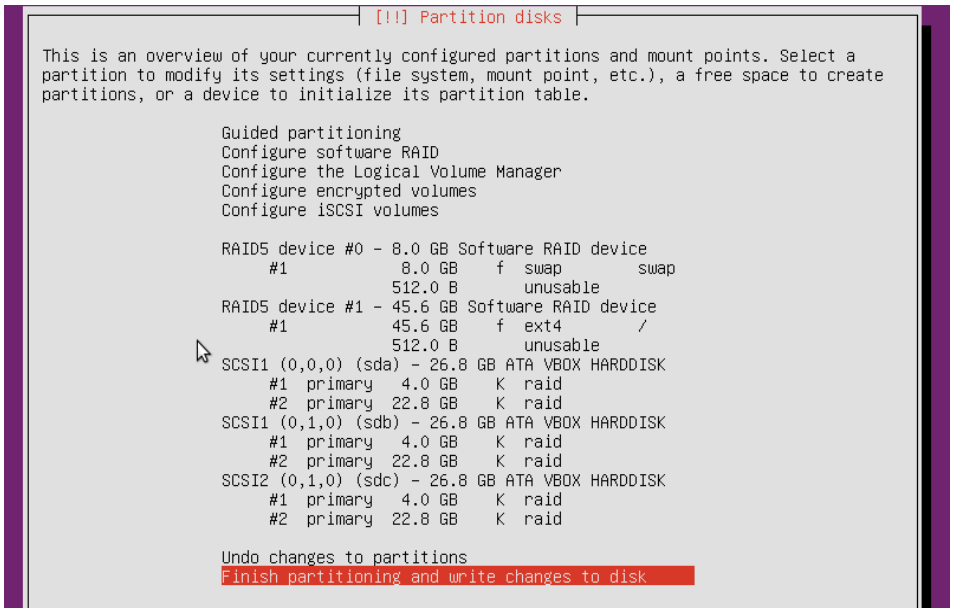

- You can now carry on with your OS installation > select the "**Finish partitioning and write changes to disk**" option and select "**Yes**" if prompted. The installation of the system will then begin.
- At some point you will be asked to enter your HTTP proxy information. If you are not behind a proxy, you can leave this blank.

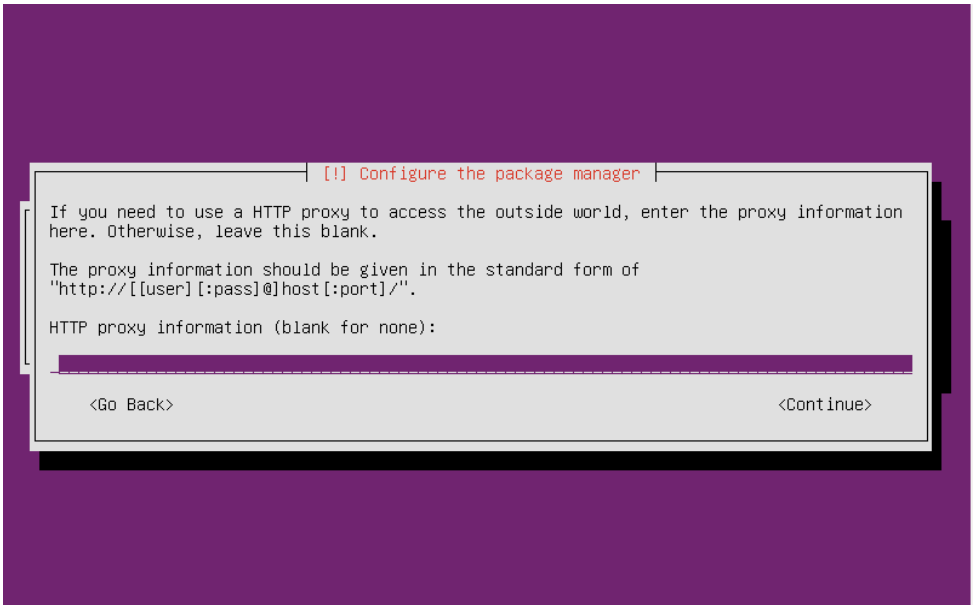

- The installation will proceed until you are asked how you want to handle updates to the system. I told it to install security updates automatically.

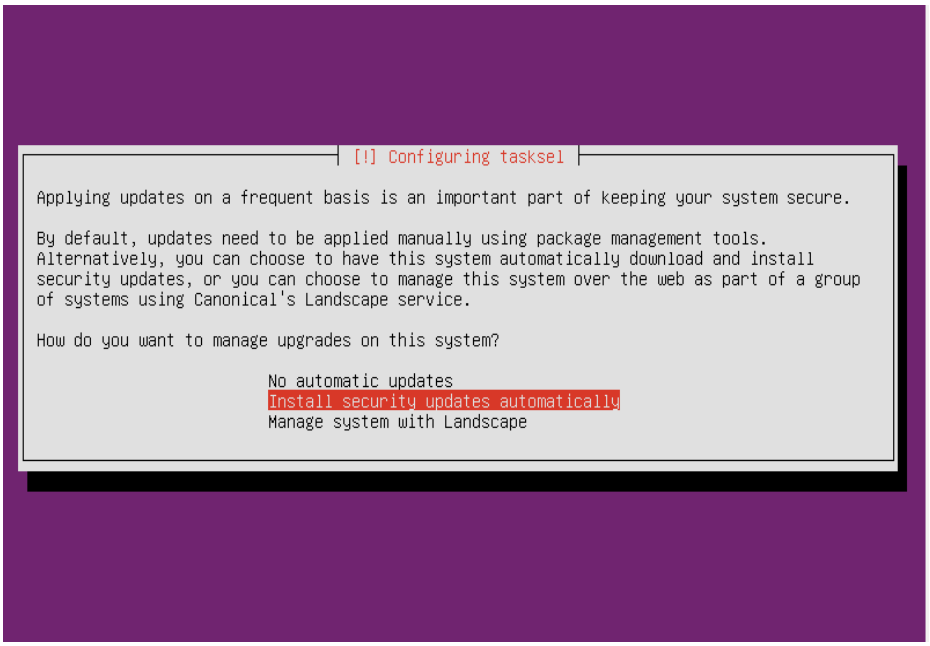

- The installation will proceed again until you get asked what software you would like to install. This will depend on what you plan on using the server for. I just selected “OpenSSH Server” for now. You can, of course, install any of this software at a later stage

TIP: *As a minimum install OpenSSH Server as this will allow you to access this server remotely once configured*

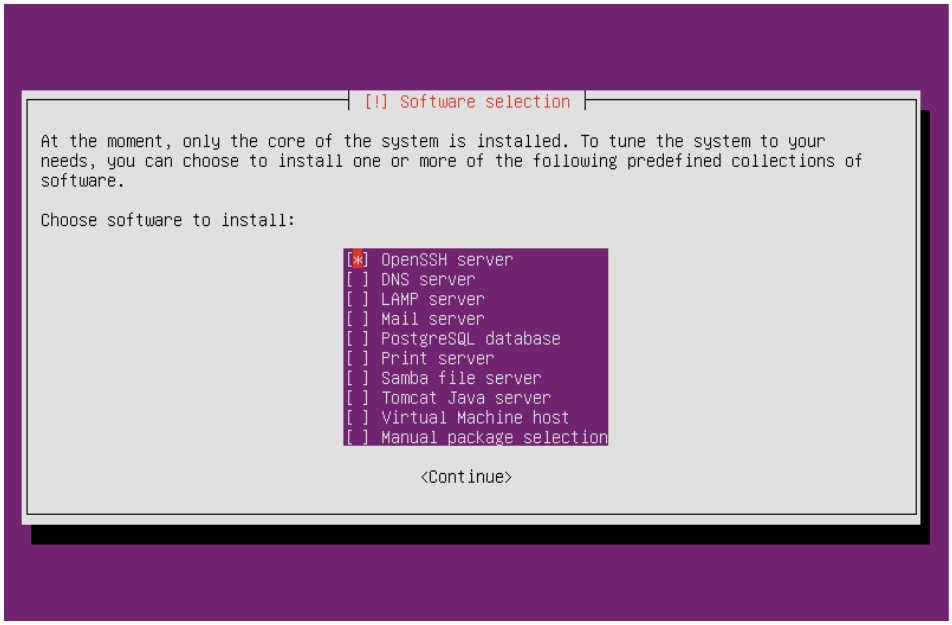

- When asked if you would like to install grub to your master boot record, select “Yes”.

Note:
When using GPT tables, you would select NO here and then select the partition you created as “Reserved BIOS boot area” to install grub

- Once grub is installed, the installation completes and you may reboot. You should boot to a login screen that looks similar to the below image.

:: Instruction Complete ::

7.3. Scientific Linux OS Installation

The following installation was performed on the actual C6145 server. Should you require a copy of the latest Scientific Linux distribution, a copy of the iso can be obtained from either of the following websites >> http://ftp.scientificlinux.org/linux/scientific/6.4/x86_64/iso/ or https://www.scientificlinux.org/ . Choose the "install-dvd.iso" option.

- Once you have booted up using your preferred installation option, you will be presented with the below welcome screen > select your preferred option to begin the installation.

Note:
For a new installation, I would recommend using option 1 "Install or upgrade an existing system"

- The next screen will offer you the option to test your installation media before you begin the installation. I often choose the "**Skip**" option

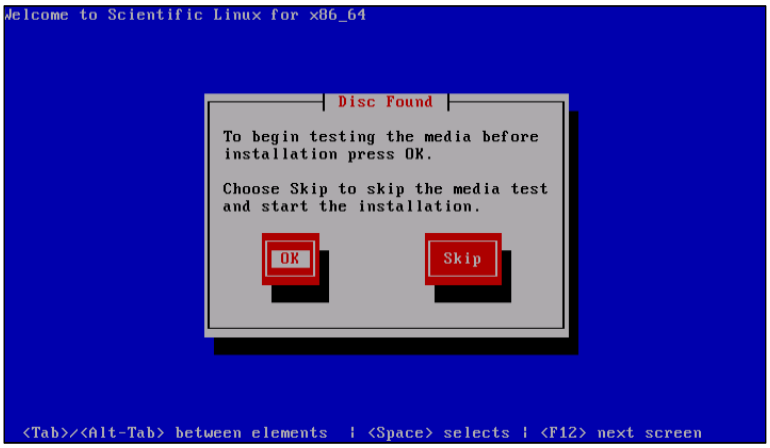

- If you chose to test your media, a quick test will be performed. If you opted to skip the media test > you will be presented with the below screen > click "**next"** to continue

- The next screen will prompt you to choose a language > choose your language of choice and click "**next**" to continue

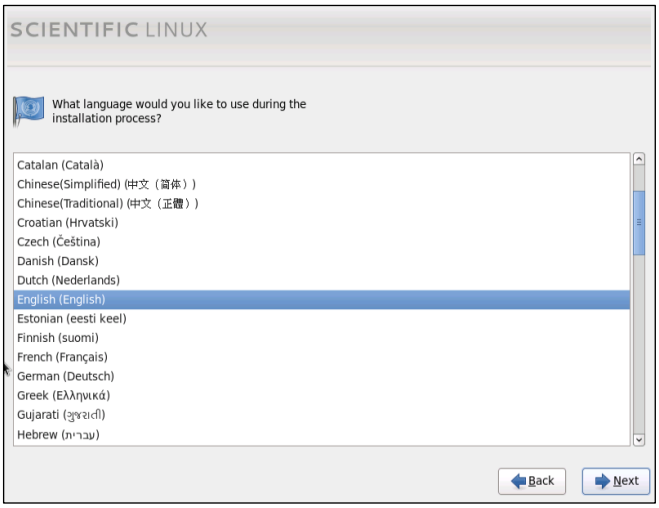

- Next you will be prompted to choose your keyboard layout > make your preferred choice and click "**next**" to continue

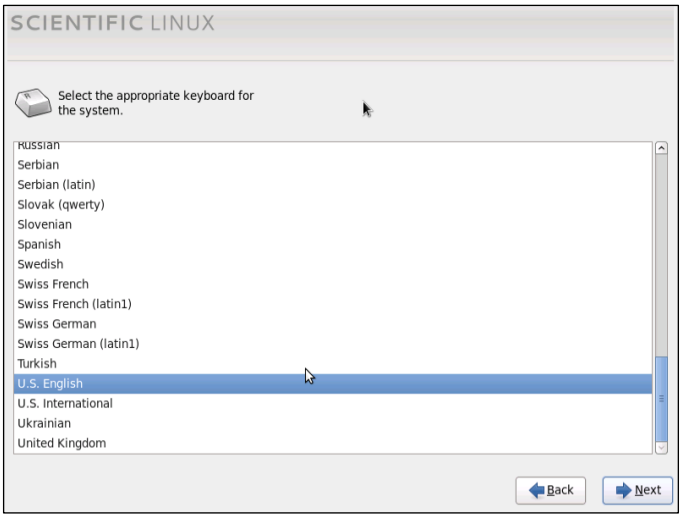

- The following screen will prompt you to choose the type of installation for your OS > option (a) "**Basic Storage Devices**" is the default option and will allow you to install your OS onto a local hard disk located inside your server > option (b) "**Specialized Storage Devices**" will allow you to install your OS on a Storage Area Network (SAN) or select iSCSI disks > choose your preferred option and click "**next**" to continue

Note:
If you purchased one of the H3ABioNet recommended server options and will be installing your OS on your local server. The default option A should suffice

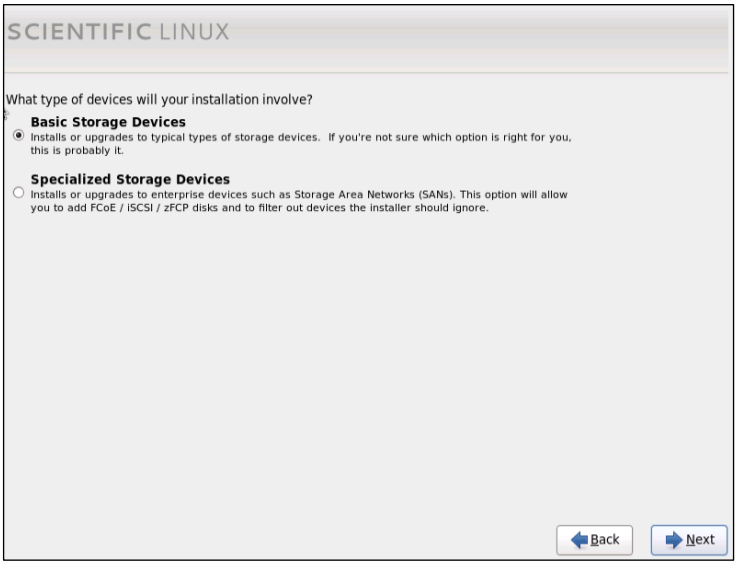

- Choosing option A will display a warning screen notifying you that should you continue, all data on the selected disk will be lost. Click your preferred option and then click "**next**" to continue

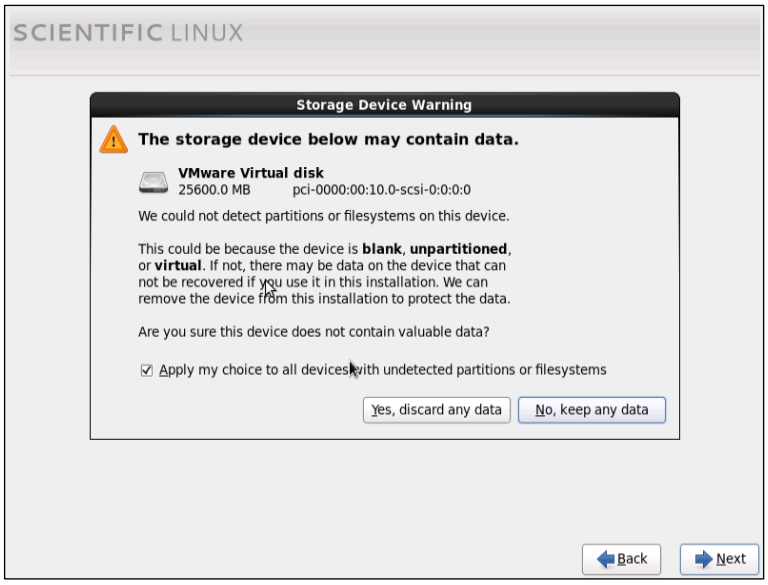

- The following screen will prompt you to type in a unique name that will identify this server on your network > type in a hostname for this server and click "**next**" to continue

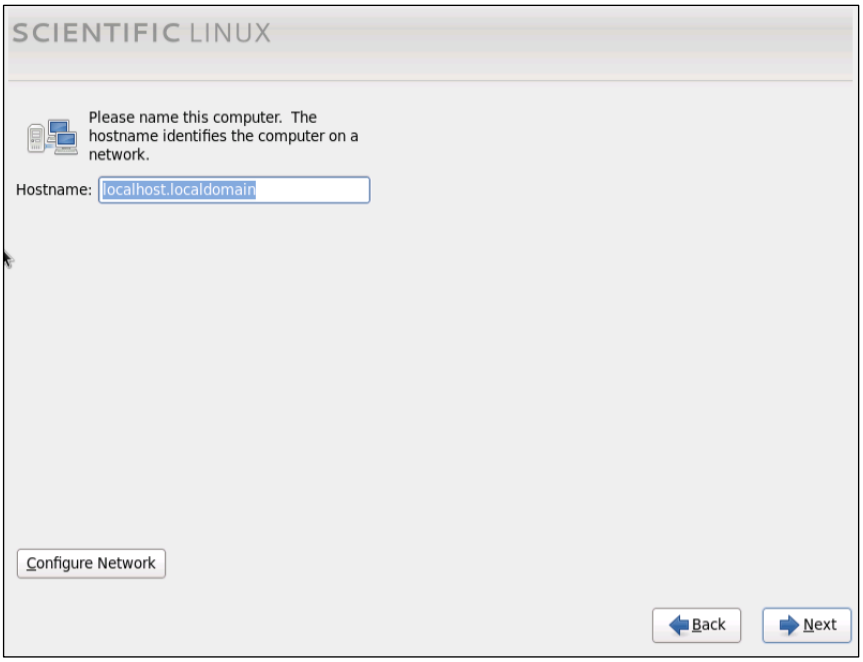

- In the following screen, choose a time zone which your server would use for its data and time stamps > deselect the "**System clock uses UTC**" > click "**next**" to continue

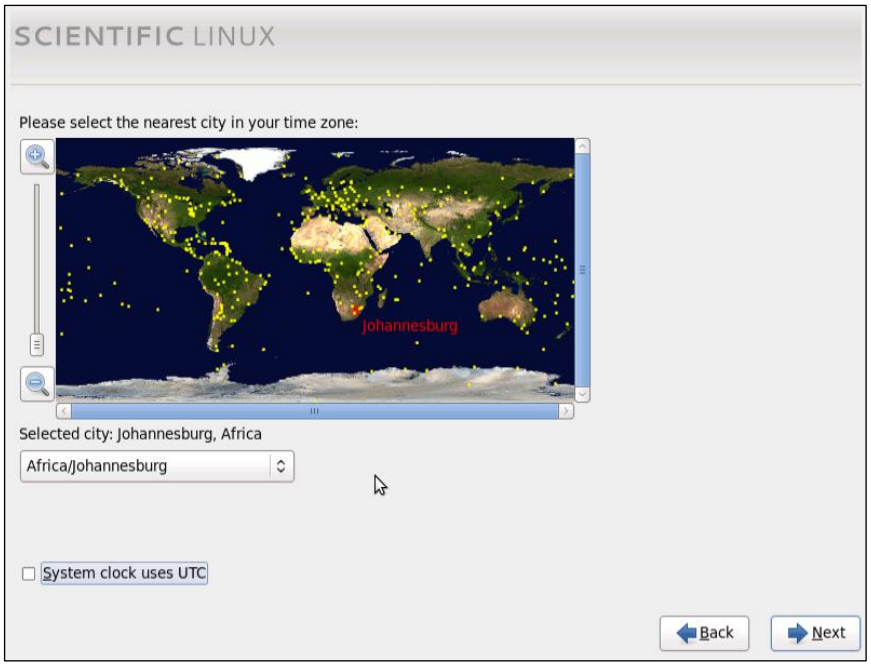

- The following screen will prompt you to type in a root password for this server. This root password will serve as the default root account on the server

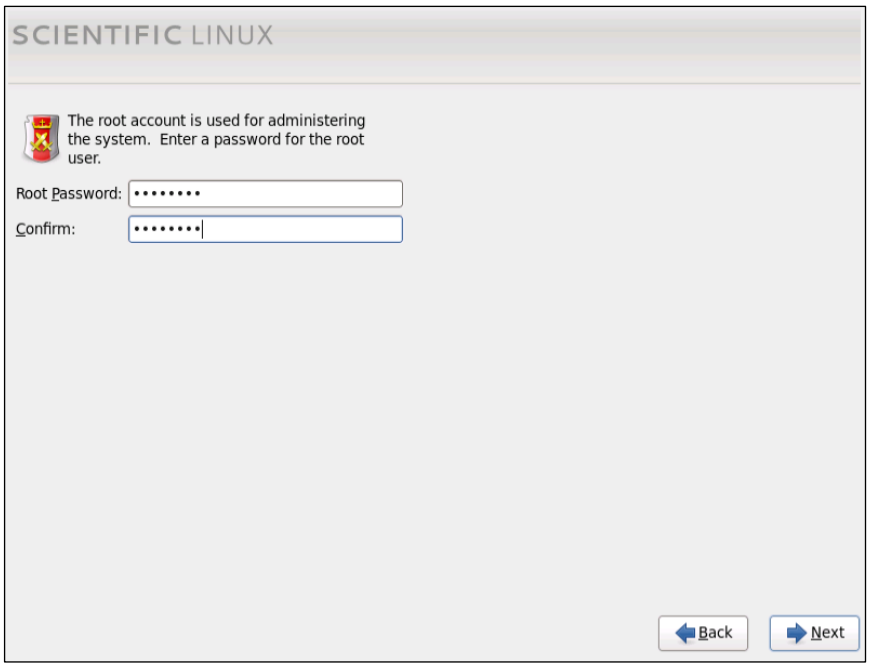

- The following screen will prompt you to choose the type of installation for your OS > review the various options and decide which best suits your environment > click "**next**" to continue

Note:
The default option "**Use All Space**" would suffice if you are following a basic default installation. This option will erase all data on the selected hard disk / partition and install the SL 6.4 OS in this location. I normally leave the "**Encrypt system**" and "**Review and modify partitioning layout**" option deselected

- The next screen will prompt you to choose a disk or partition to install your OS on > double click or highlight your option on the left hand side of the screen > if using the keyboard only, tab across to the right hand arrow and depress your space bar to move the highlighted option to the right hand side of the screen > in the right hand side of the screen, select your disk / partition where the boot loader will be installed on > click "**next**" to continue

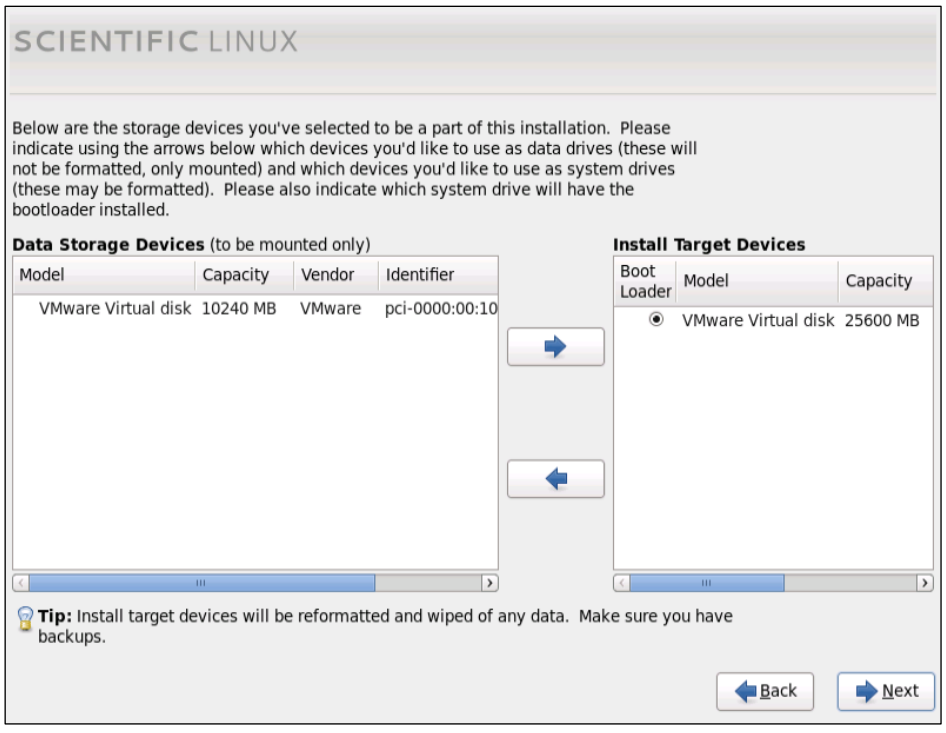

- Clicking next on the above screen will display a warning screen > if you are happy to proceed, click "**Write changes to disk**" to continue

- Clicking "**Write changes to disk**" will begin formatting the disk / partition selected in point 7.11 above

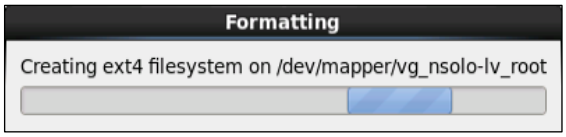

- Once the disk / partition has been formatted, the next screen will prompt you to choose a type of OS installation displayed in figure A below > choose the "**Basic Server**" option > click "**next**" to continue

Note:
Should you wish to add additional repositories at this stage > click the "**Add additional software repositories**", doing this will present the screen displayed in figure B below. To successfully setup your new repositories at this stage, you would need a valid IP with gateway to the internet. You will be able to setup new repositories once the OS has been installed. Instruction to do this is discussed in point 7 below

Figure A

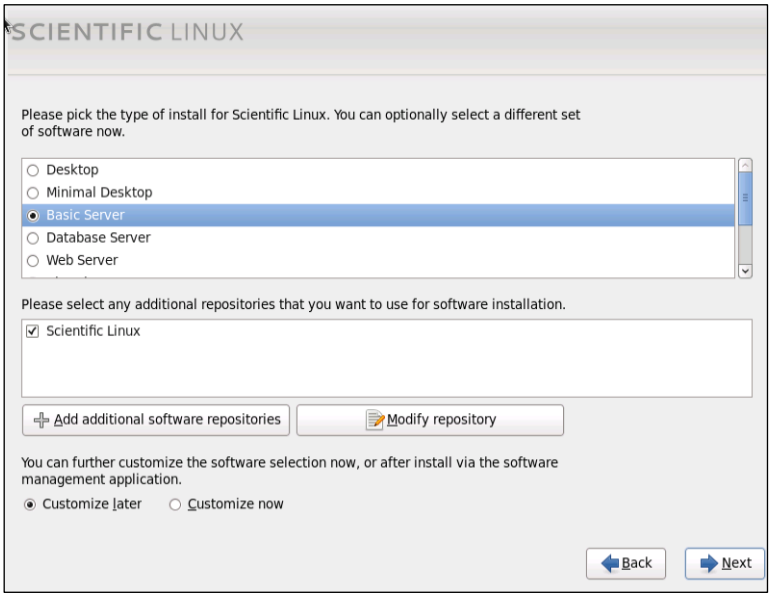

Figure B

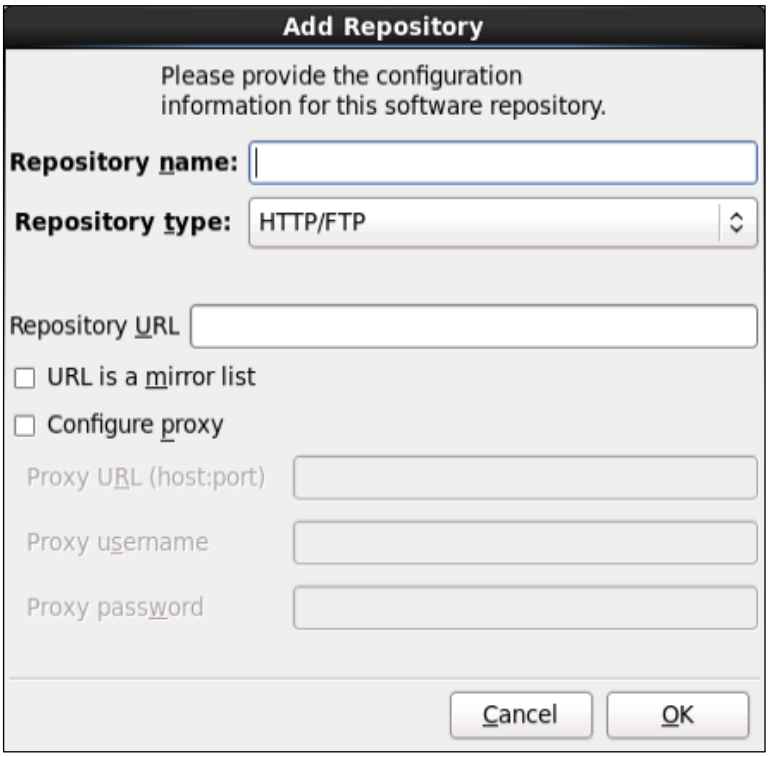

- Once you have made your choice and clicked next, the wizard will look for any dependencies that are needed

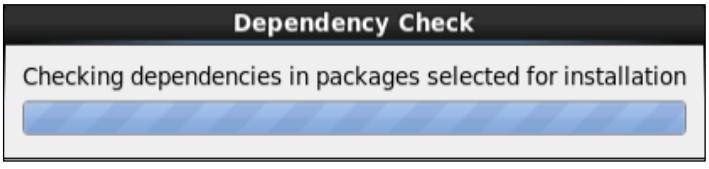

- If the wizard finds all the dependencies, it will kick off the OS installation

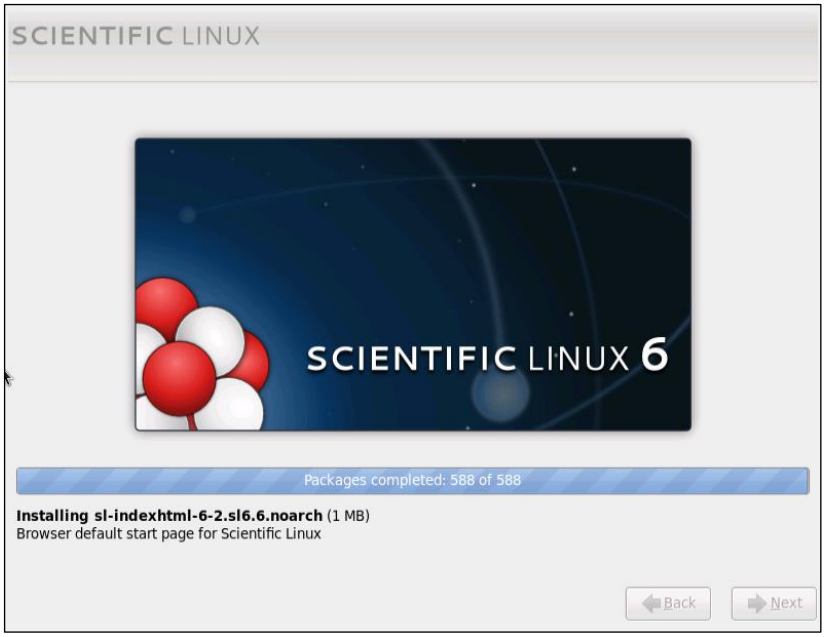

- On completion of the installation, you will be presented with the "**installation is complete**" screen > remove all install media and click the "**Reboot**" button to login to your new SL 6.4 installation.

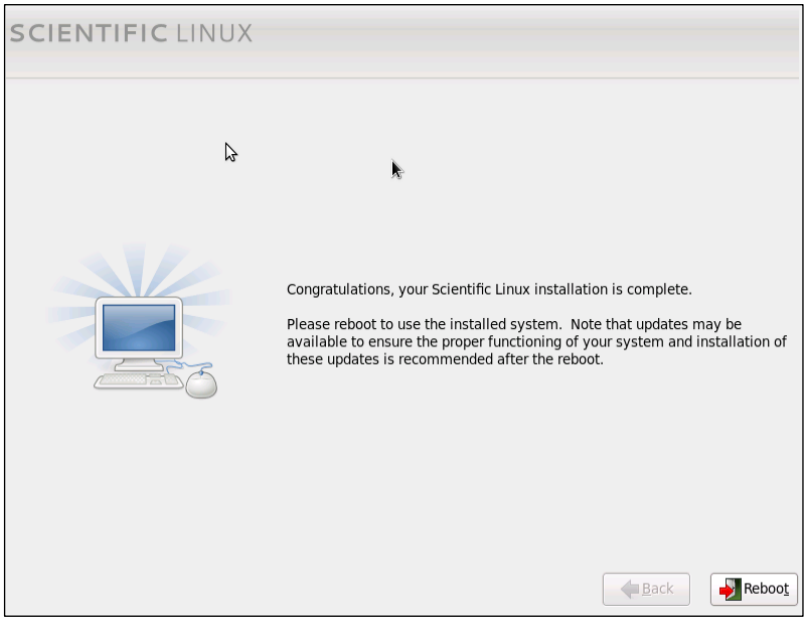

:: Instruction Complete ::

7.4. Debian OS Installation

This howto guide was configured using a virtual server created in the VMWare workstation application. The steps outlined are relevant for both virtual and physical server installations. The virtual hardware used for this server is listed below. To get a copy of the latest Debian release, please visit the following website >> http://www.debian.org/

- **Server Type:** VMware Workstation 10
- **Memory:** 1GB
- **Processor:** 1,53GHZ 2 core
- **Hard disk:** 2 SCSI Disk with 20GB each

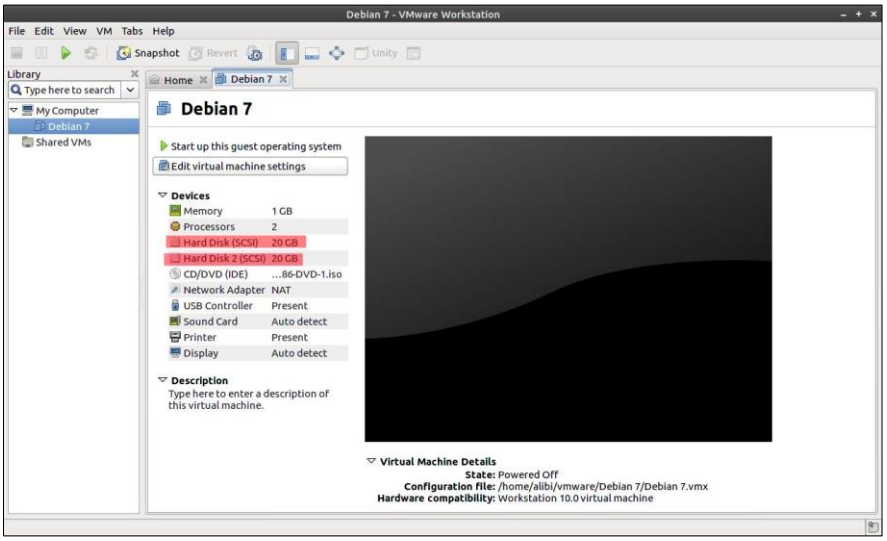

- Start the machine (Server or PC) and select Graphical install

- Choose your language

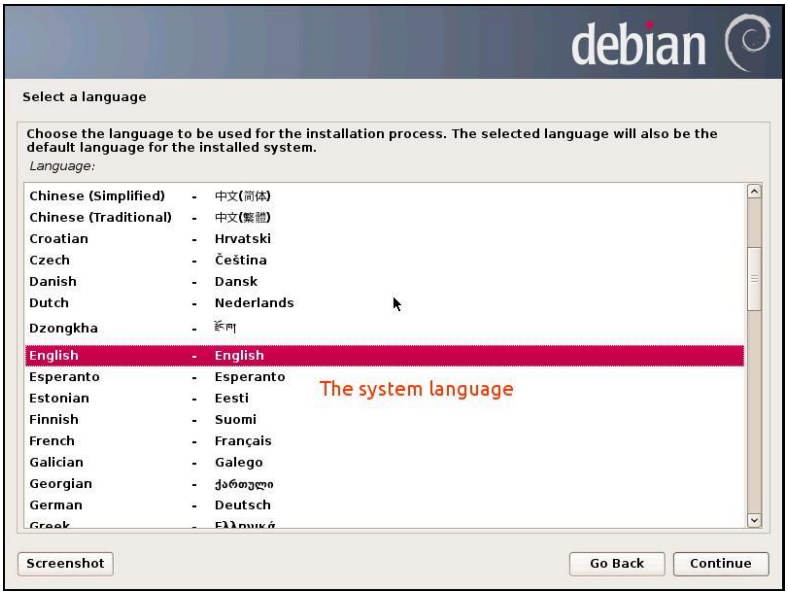

- Choose the keyboard layout

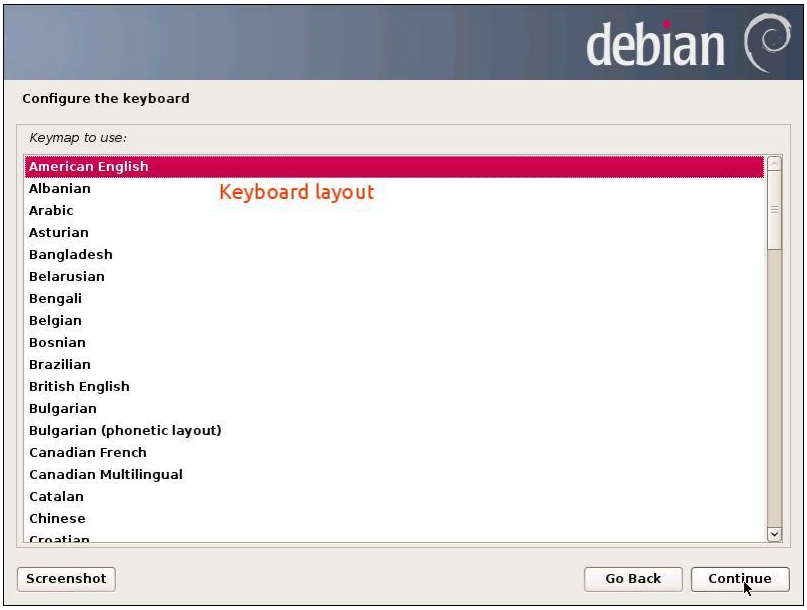

- Configure the basic network setup by typing the machine name

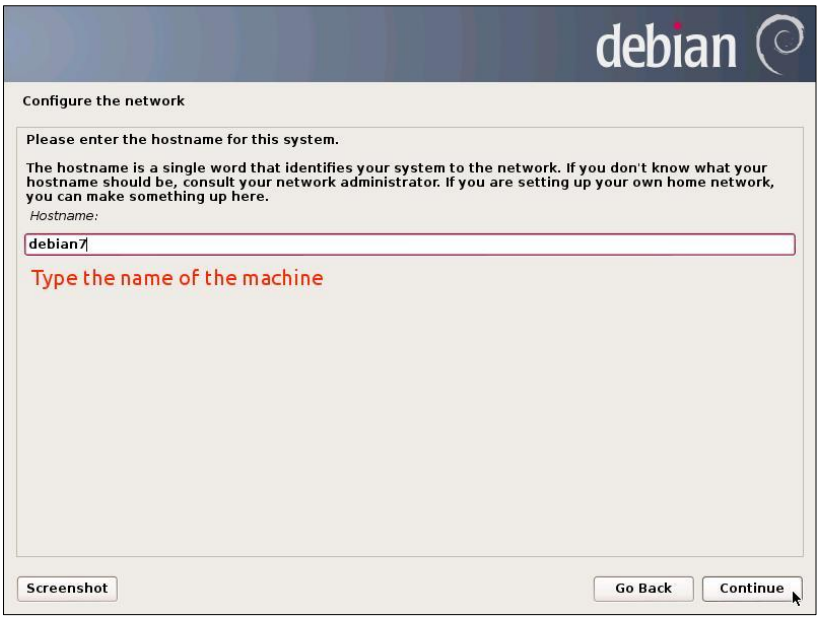

- Choose a root password for the system

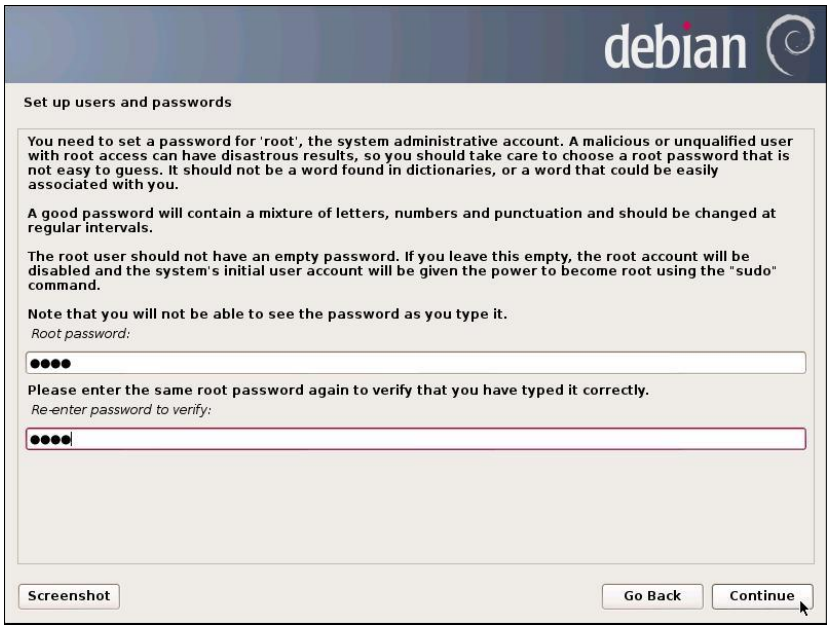

- Create a default user account by typing a name in the following box.
  You will use this default user account to log into the server once the installation is complete.

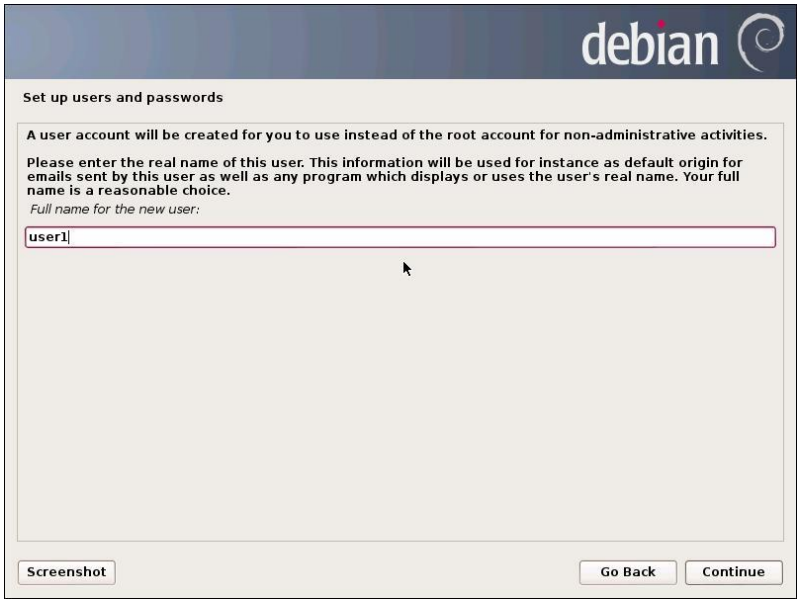

- Choose a secure password for this user account. It is common practice to choose a password which conforms to the following criteria: a minim of 8 characters is made up of numbers, special characters and both upper and lower case letters.

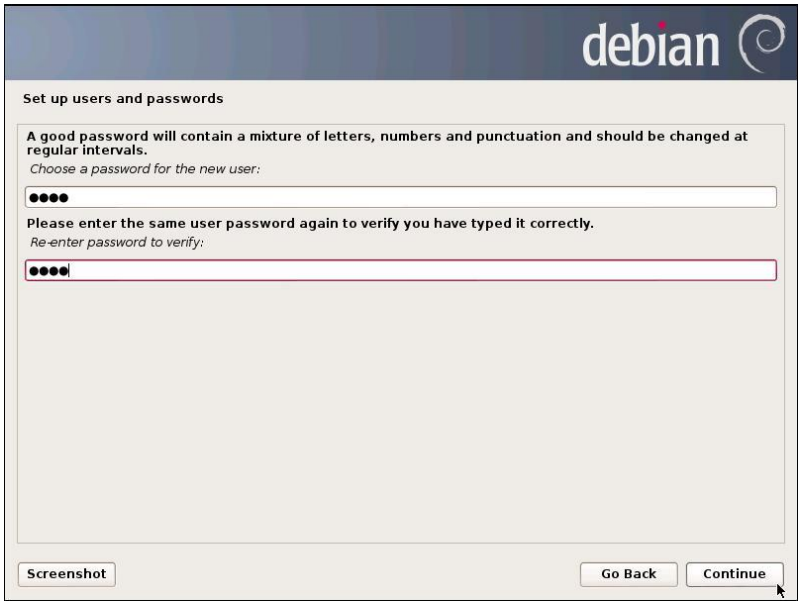

- In the following box you will define a domain.

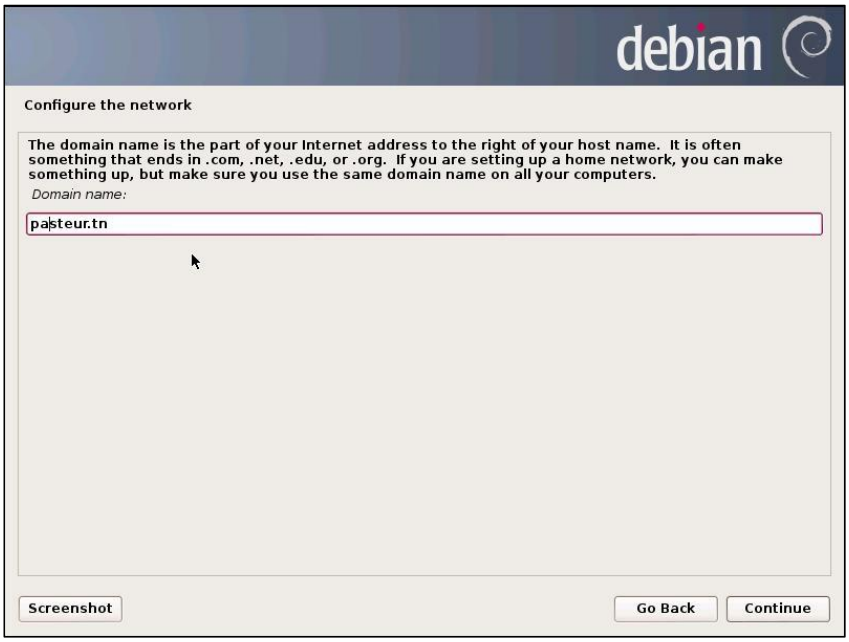

- The next screen prompts you to choose a disk partitioning method. As we will be creating a RAID configuration > choose the option to manually partition your disks.

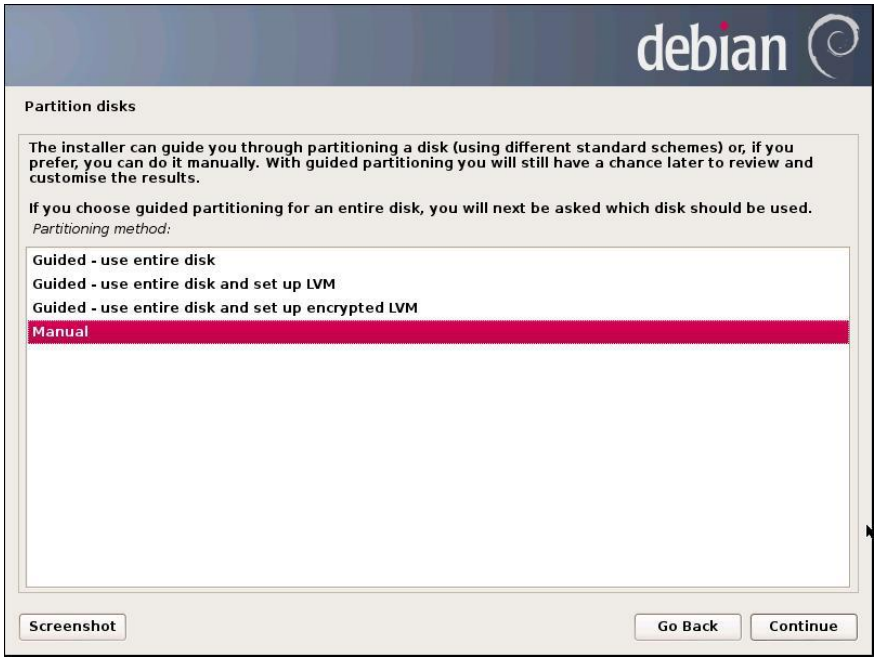

- You will now be asked to choose a disk to partition. Choose the first disk which is usually represented as "**sda**" in Linux based installations.

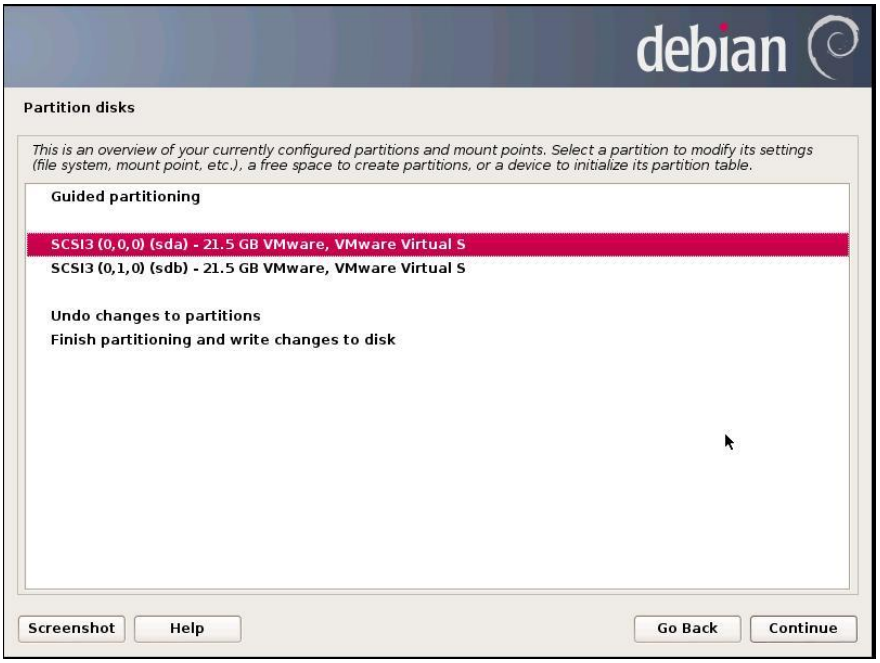

- Click the continue button to begin partitioning this disk. The wizard will prompt you to confirm that you want proceed with the partition. As there is no data on this hard disk > choose "Yes" to continue. If you are performing this operation on a hard disk which has data, it is strongly recommended that you first backup the data before proceeding as all data will be destroyed.

- Start partitioning by selecting the first partition.

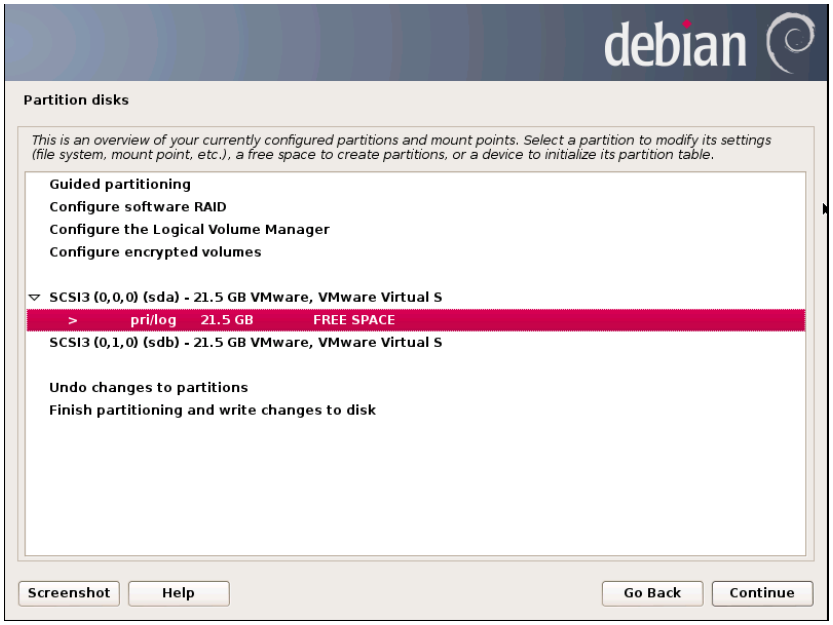

- Choose the size for this first partition.

- Choose the file system and the mount point then confirm.

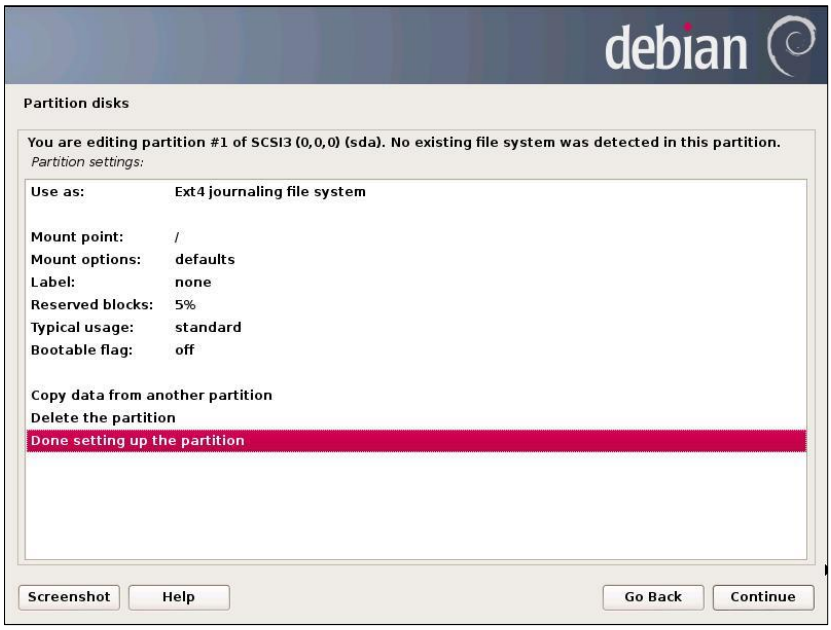

- Select the free space to setup the other partition.

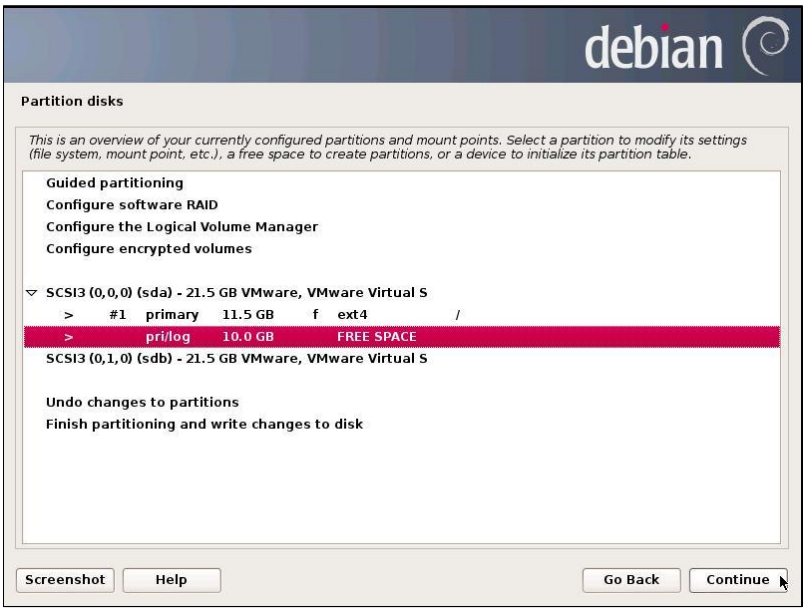

- Choose the second partition disk space and assign it a partition size.

- Setup the file system, mount point and confirm.

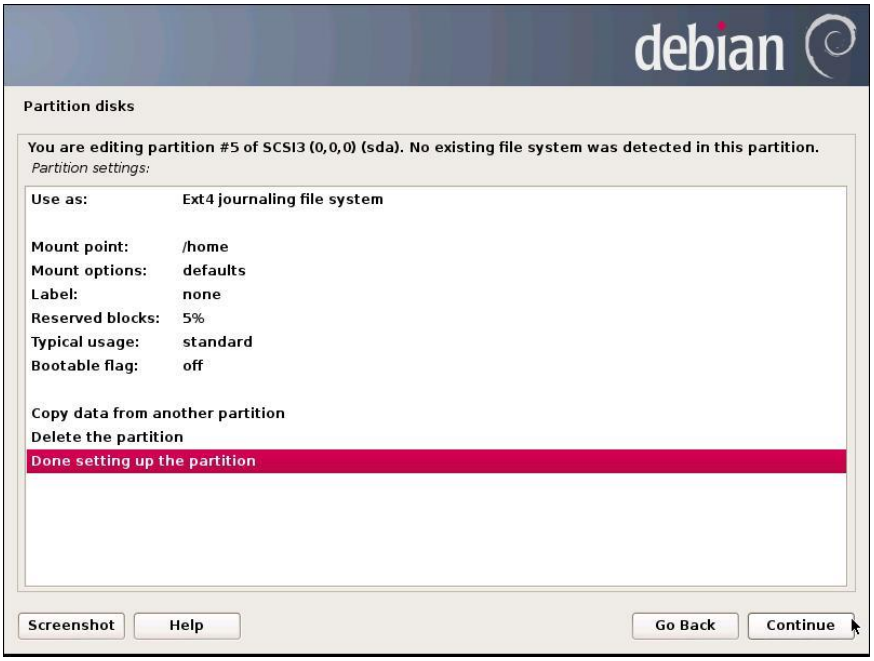

- Select the second disk for partitioning the same as the first to make the software RAID on it.

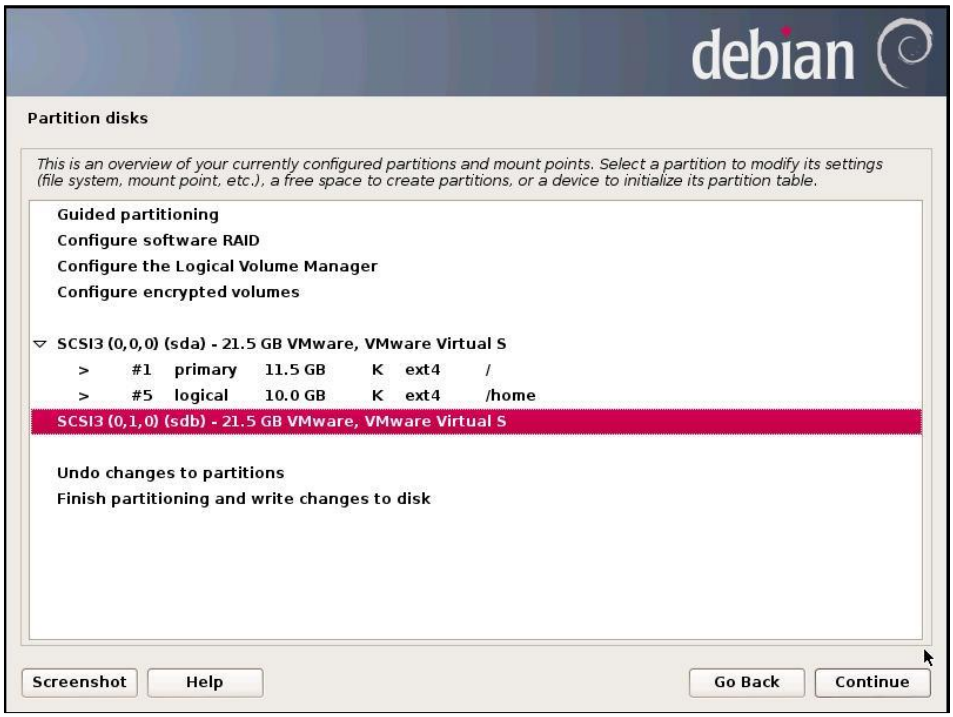

- Create the partition table and start partitioning.

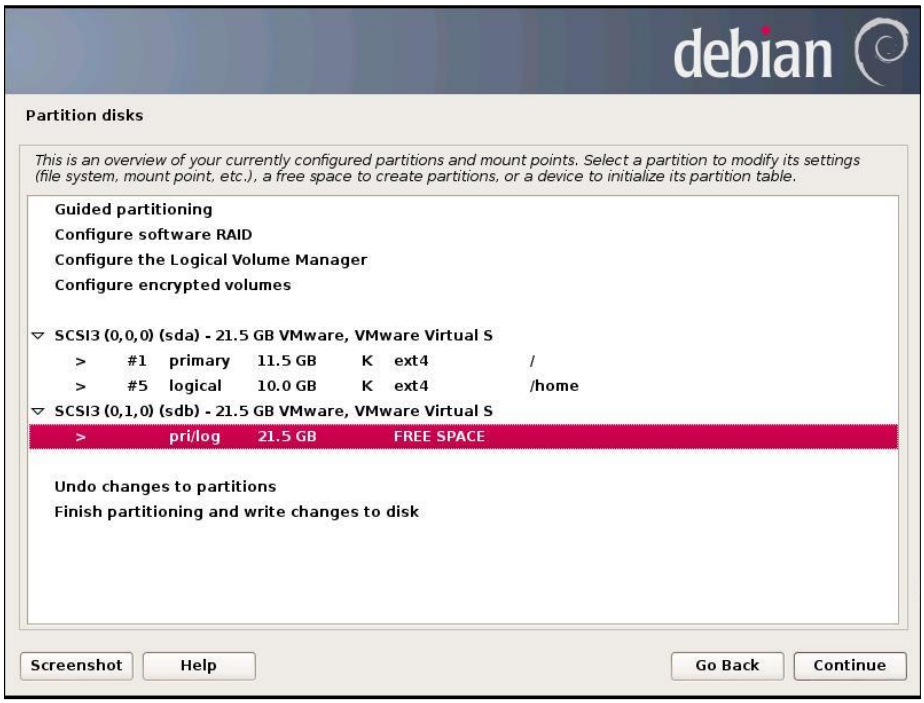

- Create the first partition with the same space as the one in the other disk.

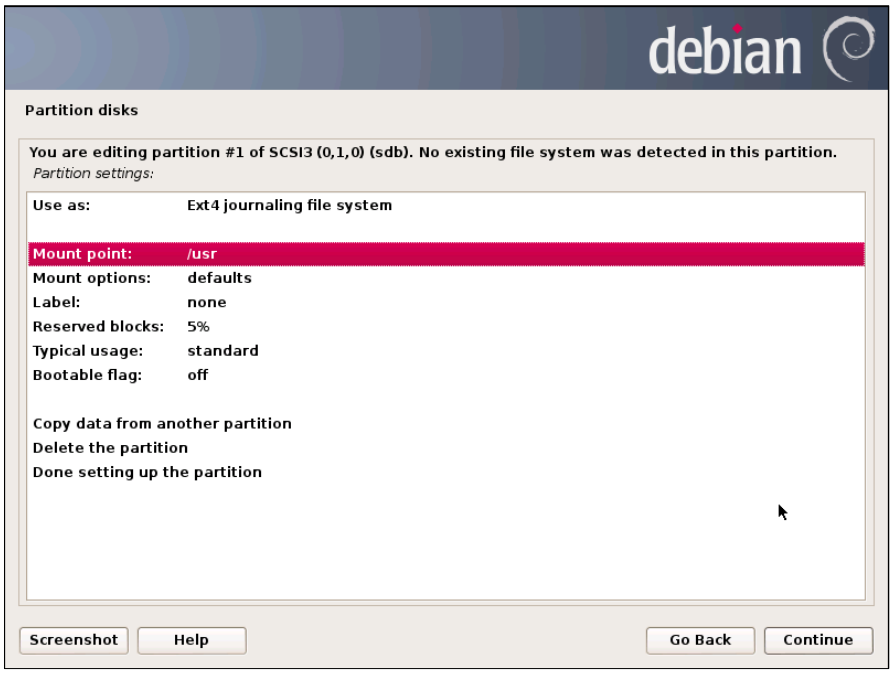

- Choose the “Do not mount it” option.

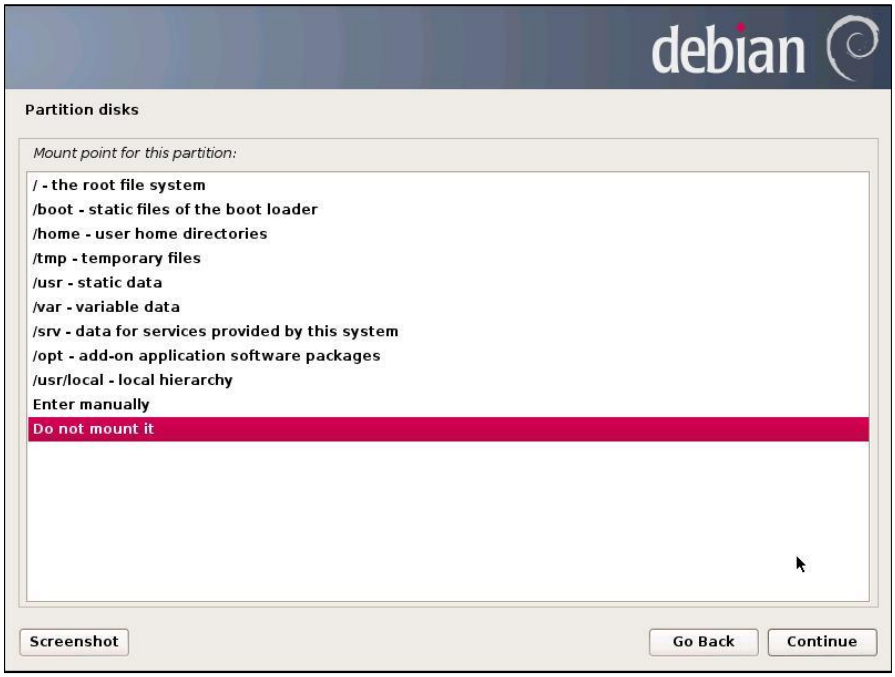

- Confirm the modification of the second partition.

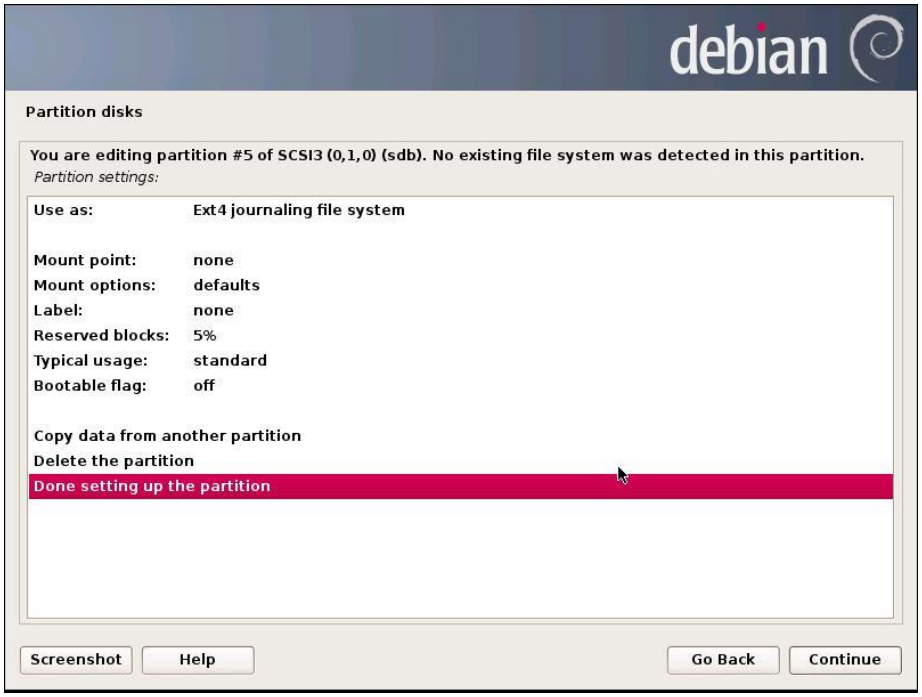

- Start the RAID configuration by choosing the "**Configure software RAID**" option.

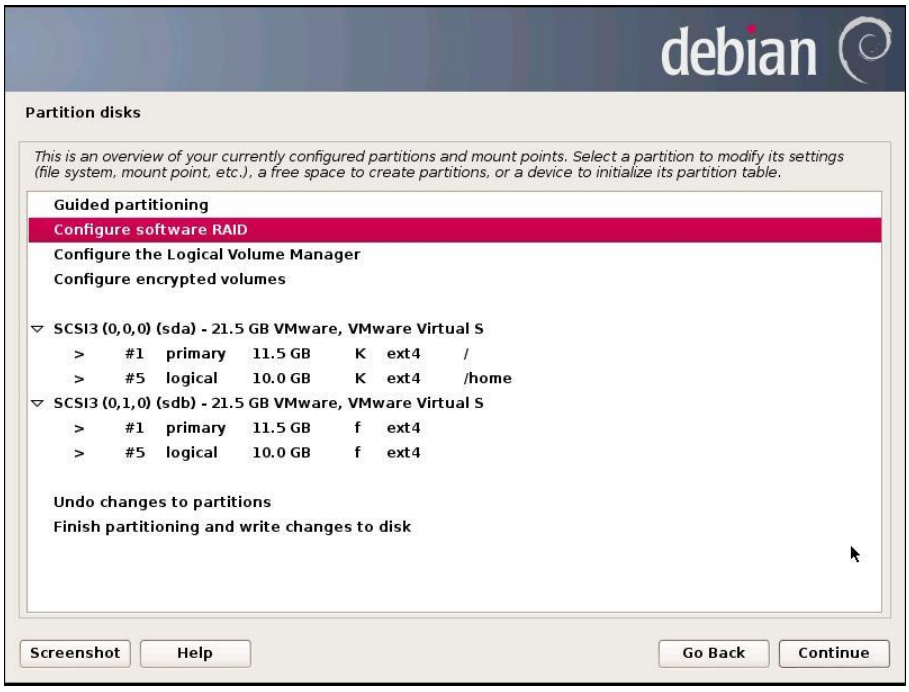

- Choose the "**Create MD device**".

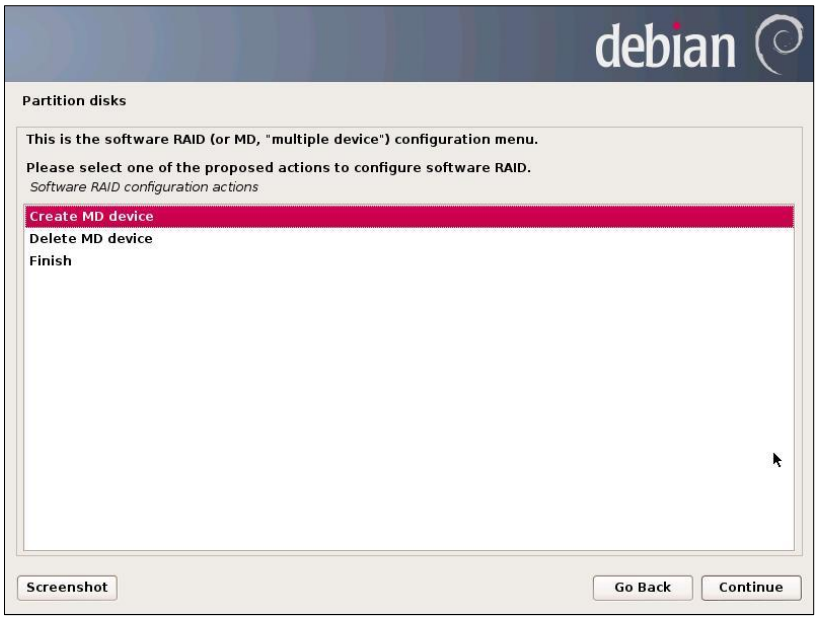

- Choose the “RAID1” option.

- Type the number 2 for the number of disks to be used.

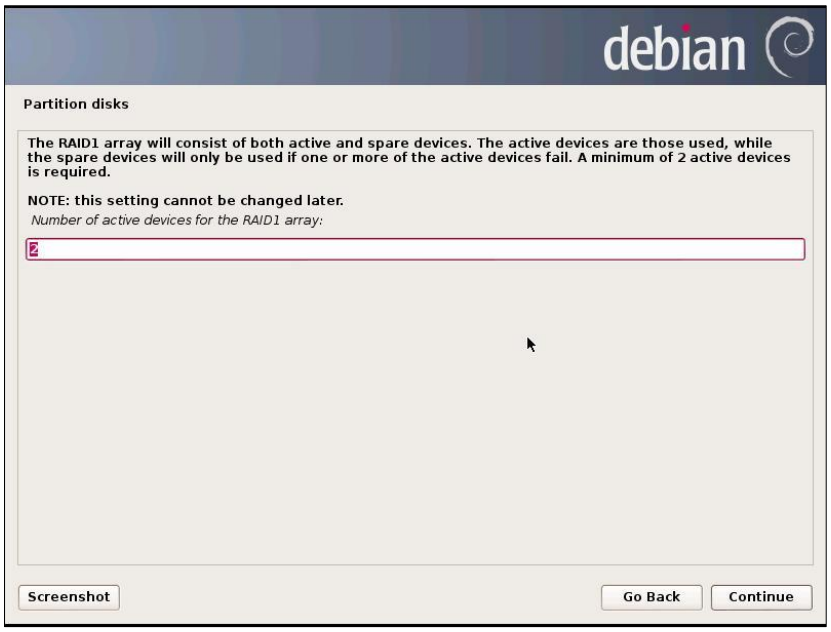

- Type the number 0 when asked to specify any spare devices. We will not be using hot spares for this demonstration.

- Select the two primary partitions to be the active devices.

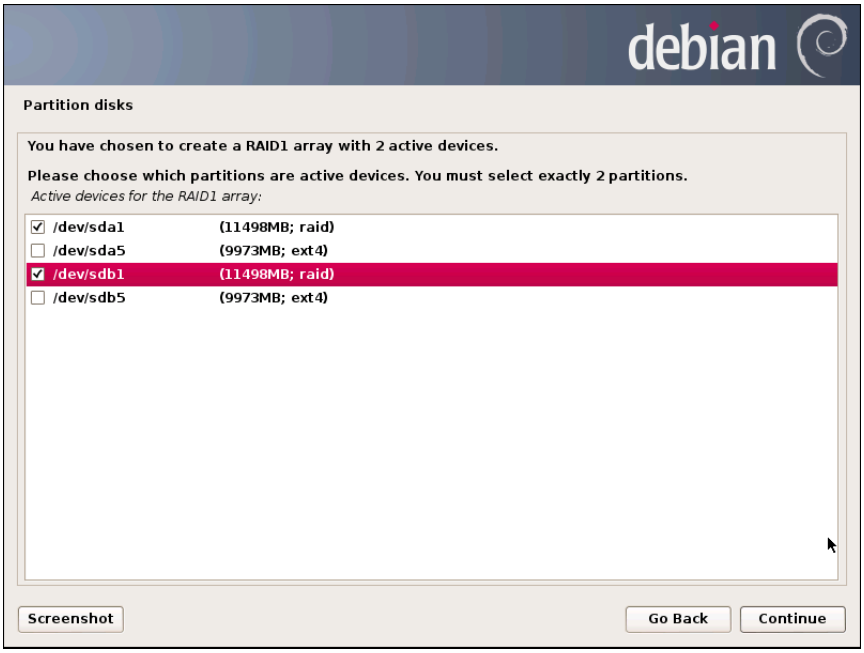

- The first partition is done. Redo the RAID configurations on both disks to create the second partition.

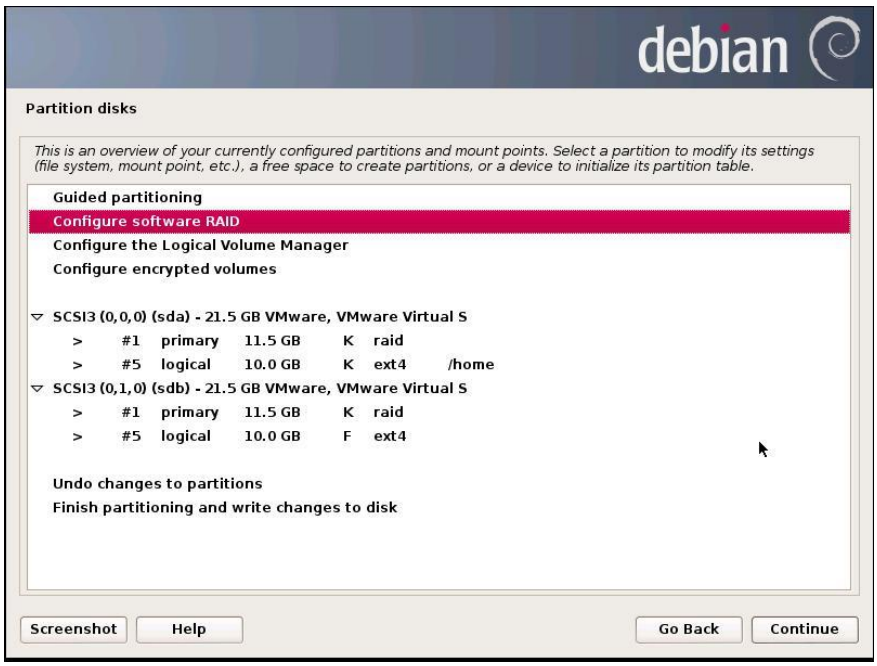

- Select the “Create new MD device” option.

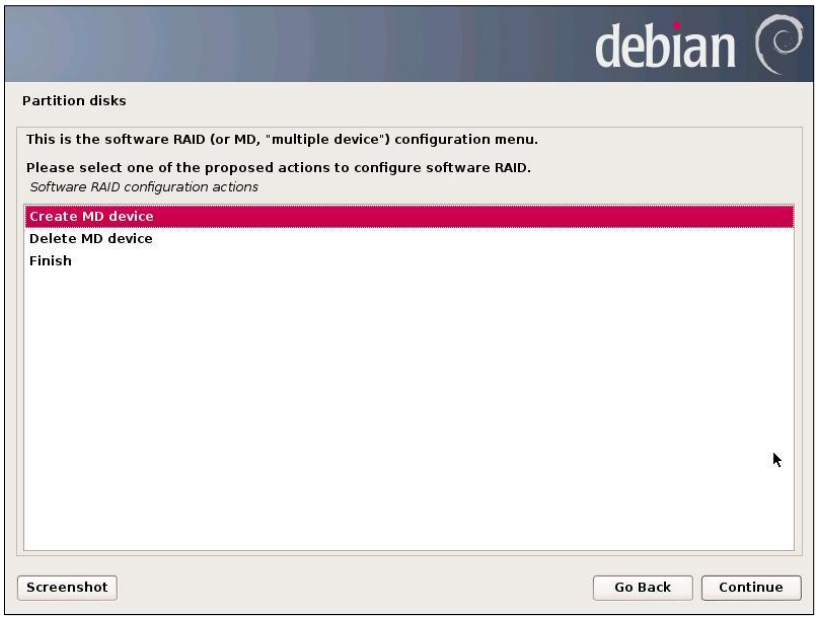

- Choose the “RAID1” option.

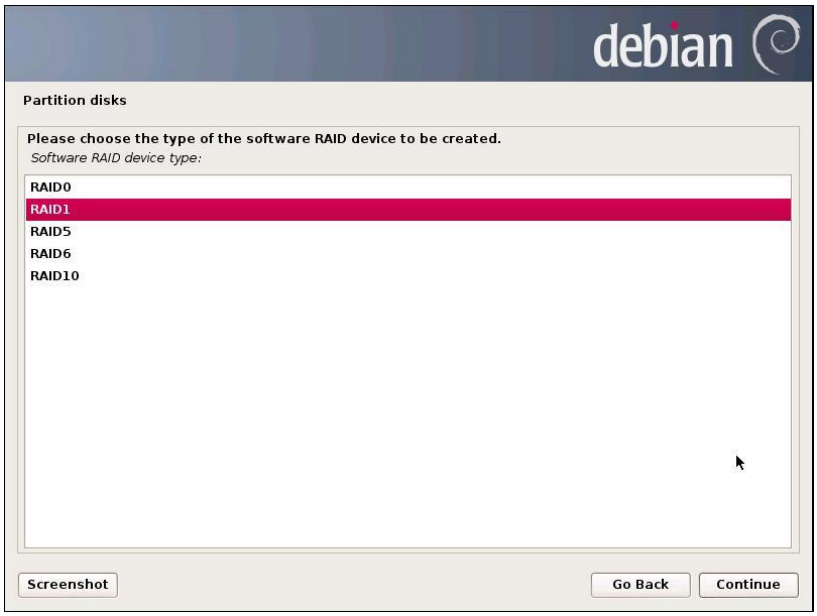

- Select the second two partitions created for the active devices.

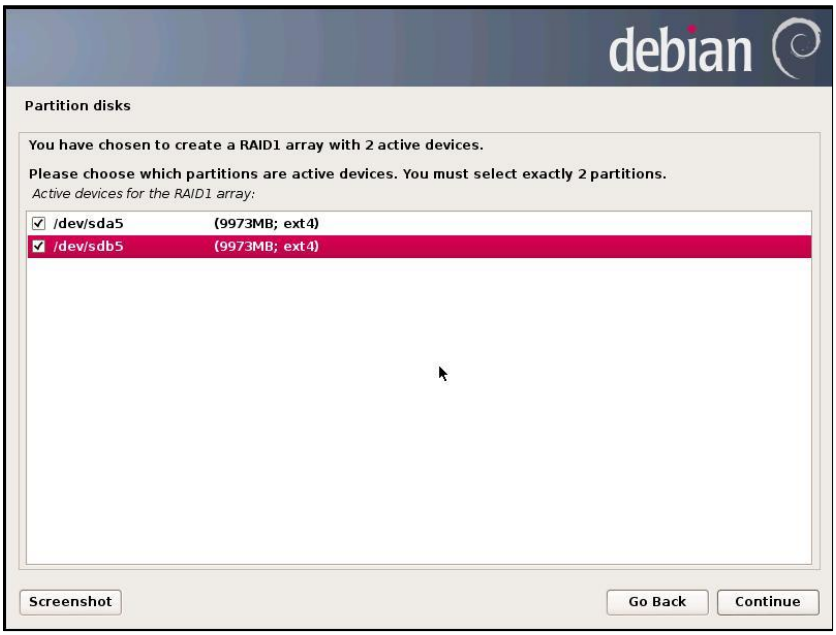

- Confirm the RAID configuration.

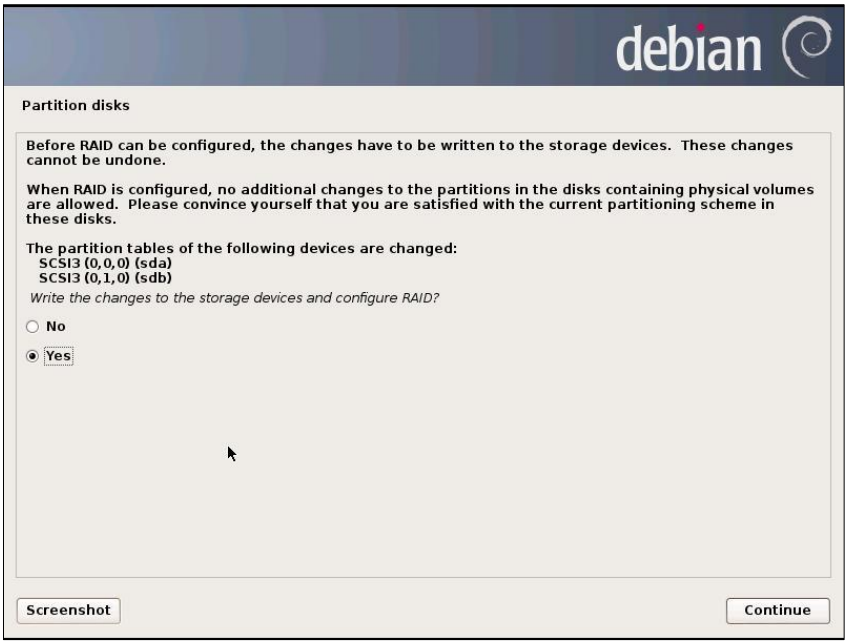

- Choose to finish the RAID configuration.

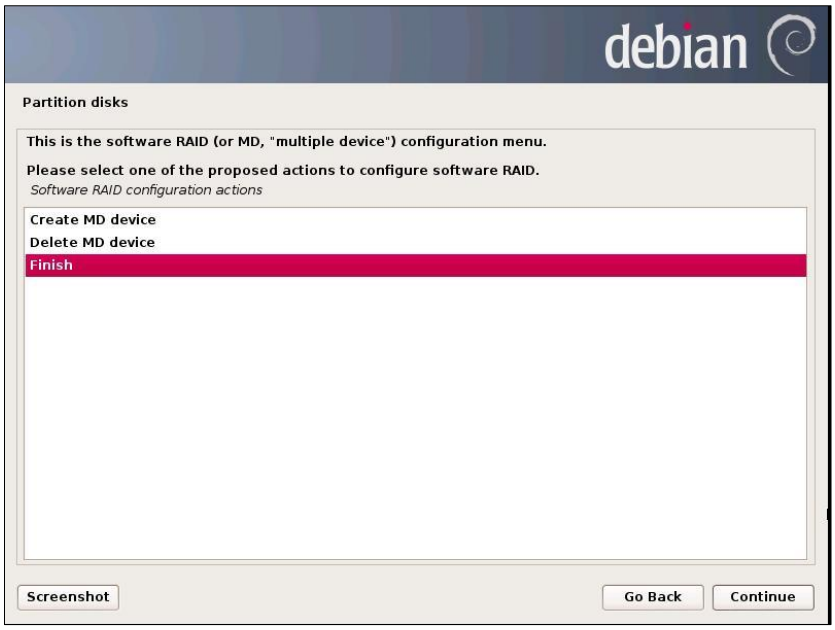

- Start the RAID disk partitioning by defining the file system and the mount point.

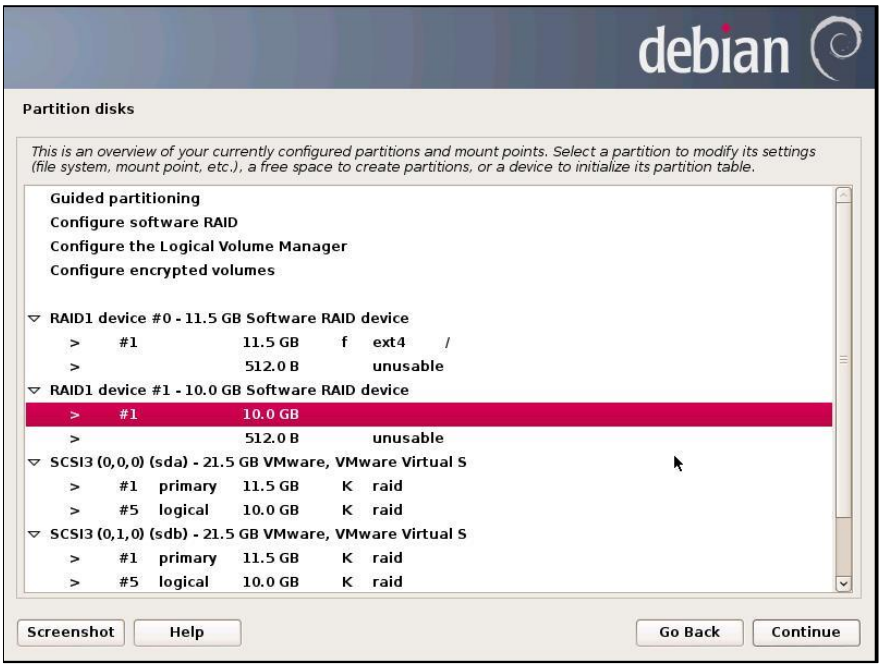

- Start configuring the partition.

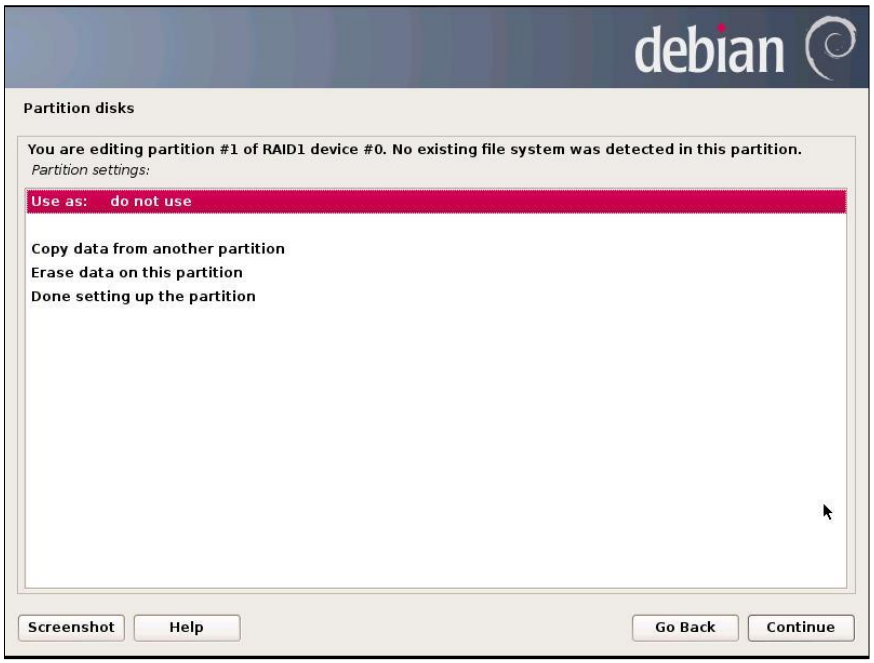

- Choose the file system as Ext4 and mount point as / and confirm.

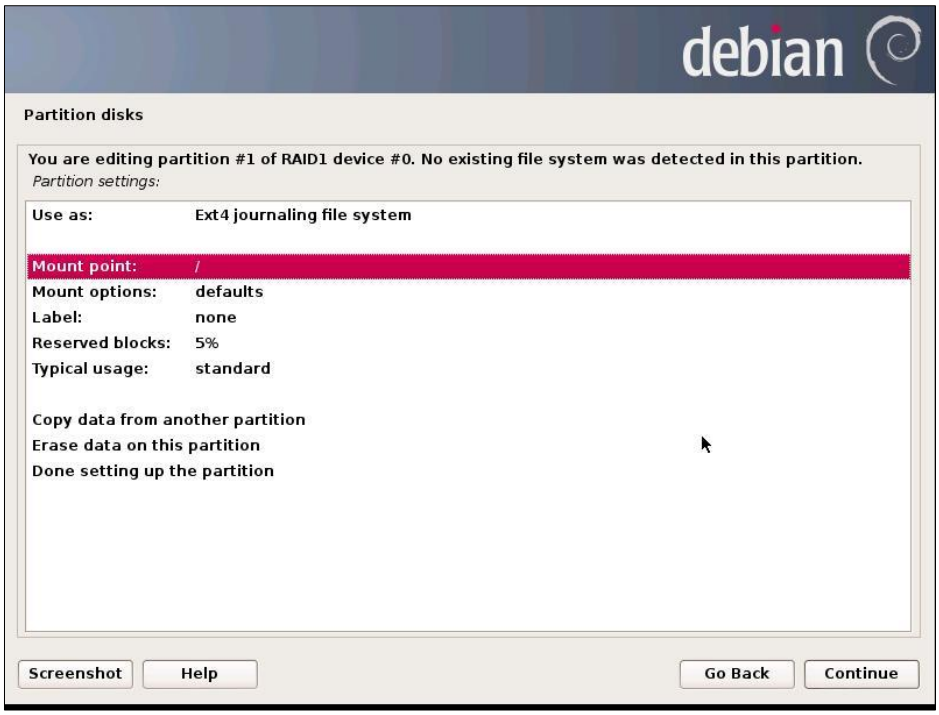

- Choose the second RAID partition to start configuring it.

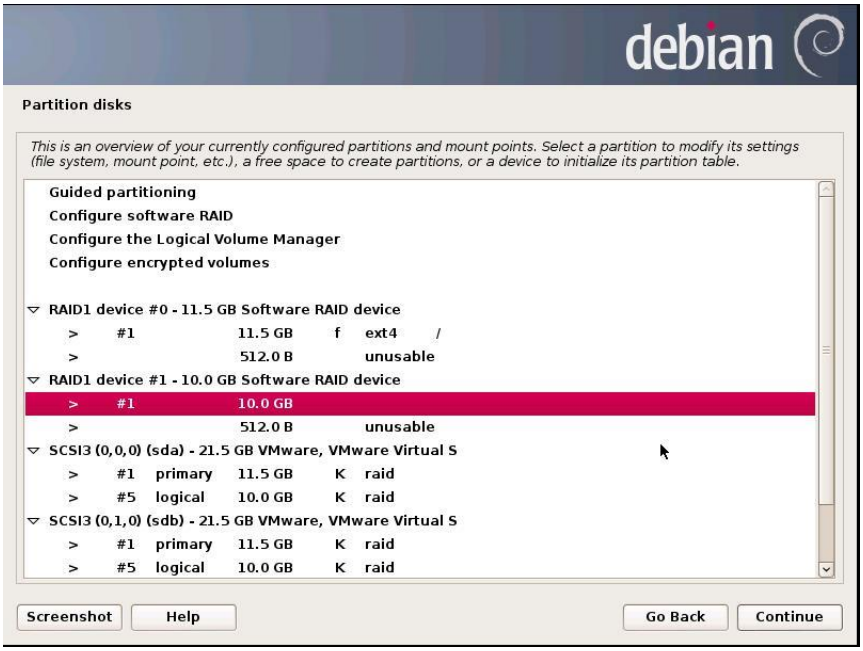

- Setup the file system as Ext4 and the mount point as /home and confirm.

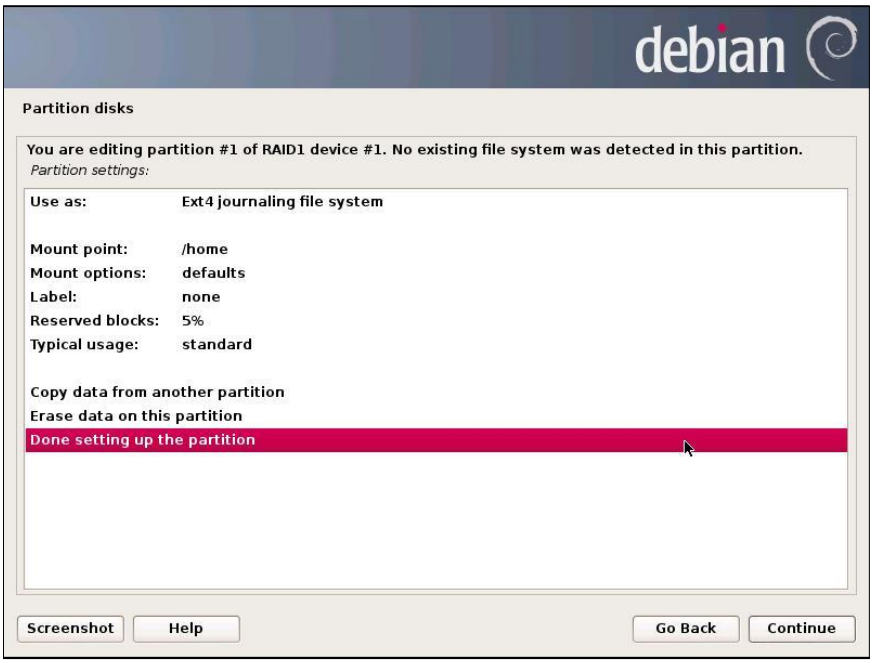

- Choose finish partitioning.

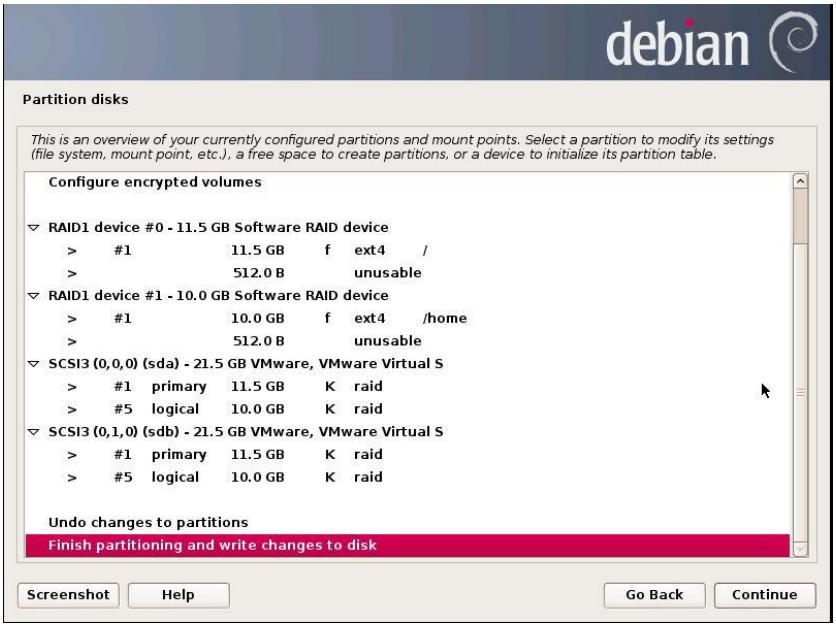

- Confirm the application of the changes on the disk.

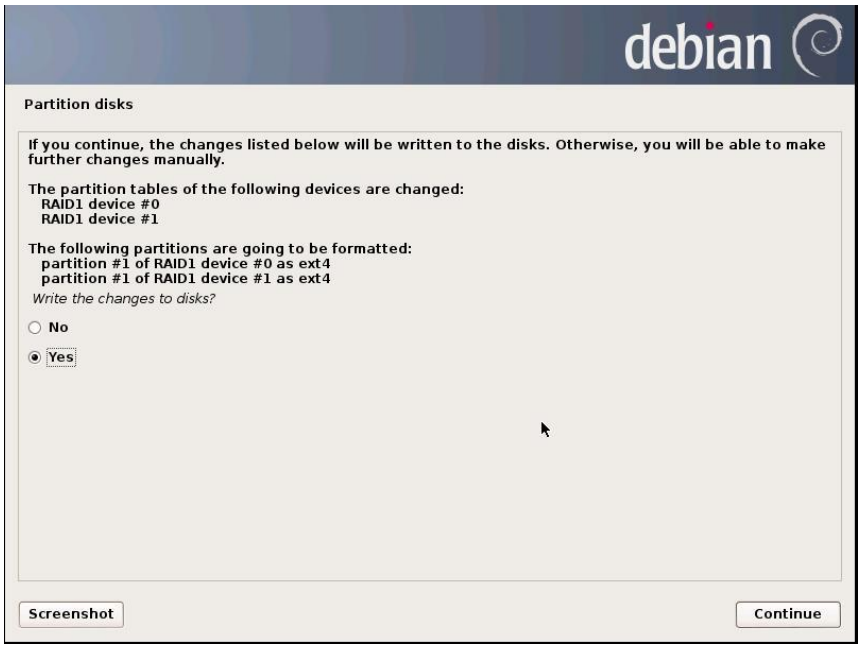

- Wait for the basic system program to be installed.

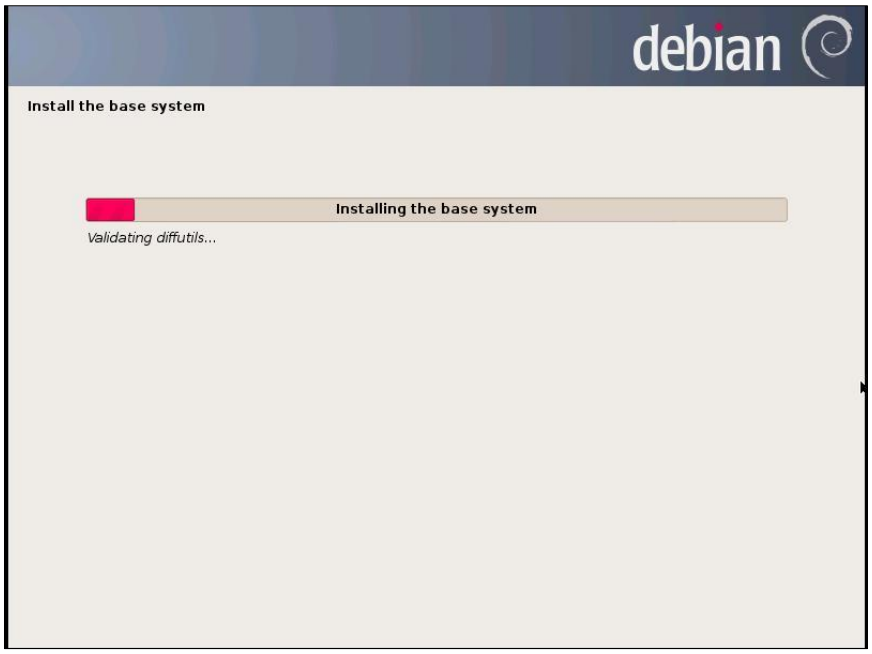

- Select “No” we will be using the first Debian DVD only.

- Choose “Yes” to setup a mirror to get programs and updates from.

- Choose a country to use its mirror server.

- Choose the mirror server link.

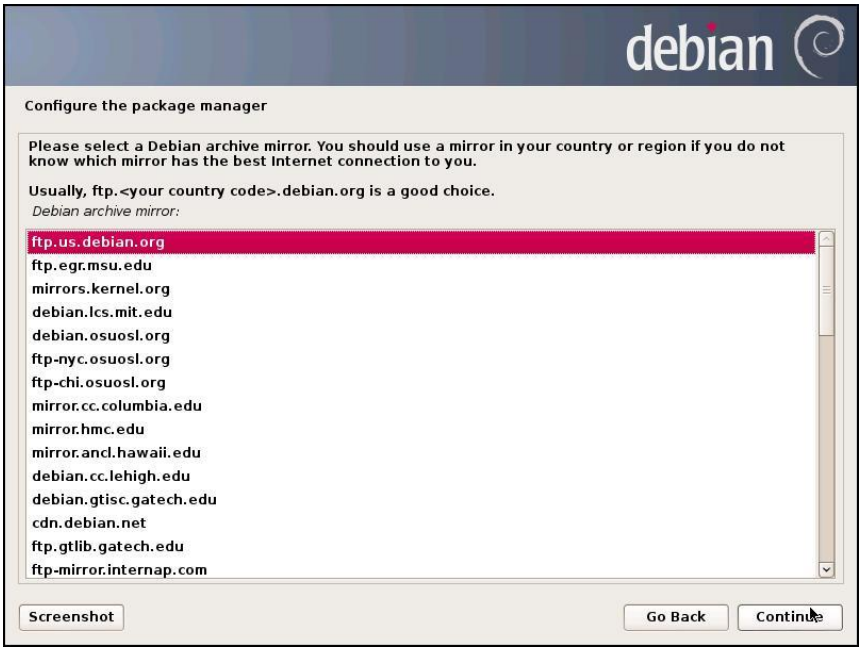

- Leave the proxy configuration blank if you have no proxy configuration at the network.

- Choose some environment or software if needed for the working of the machine.

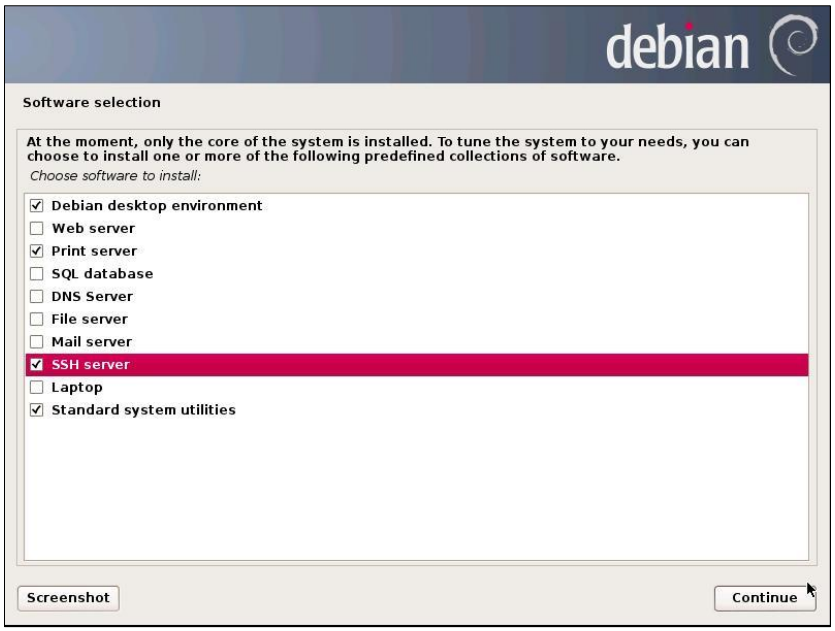

- Choose "**Yes**" to setup the boot loader menu at the startup of the machine.

- The installation is complete, tap Continue to restart the server.

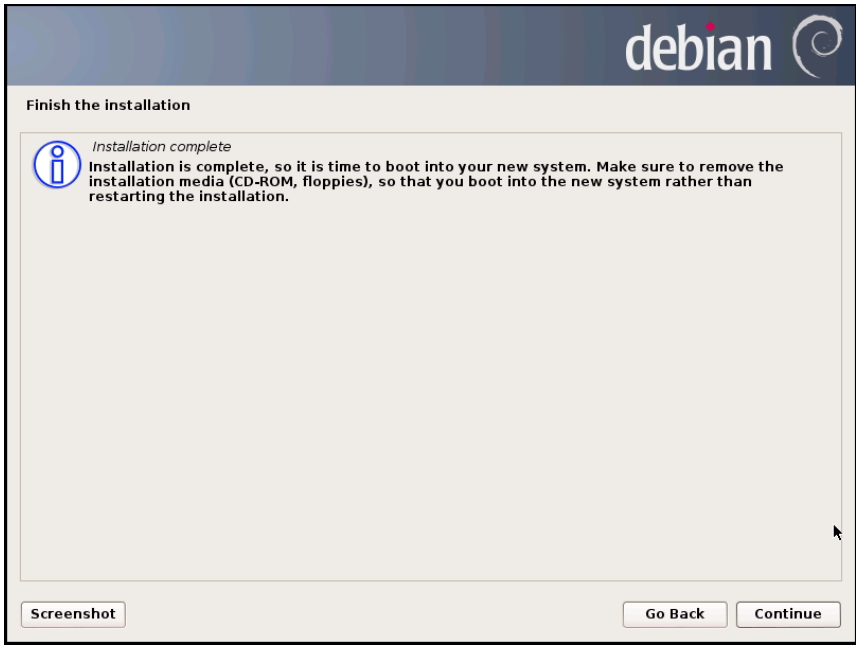

- The boot menu loads up, choose the first option to start your system in normal mode.

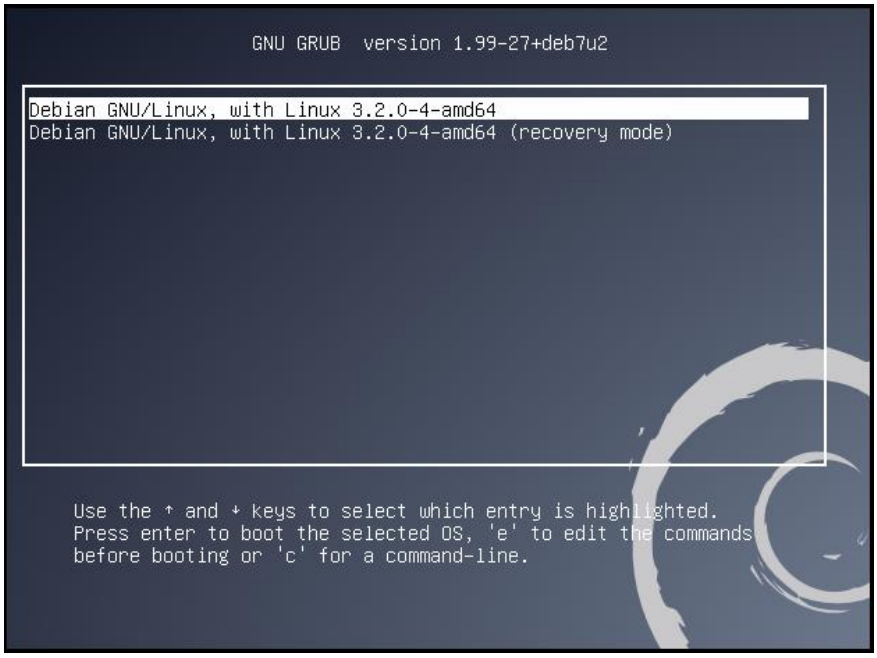

- Click on the user account you created during the installation and type in the password you set.

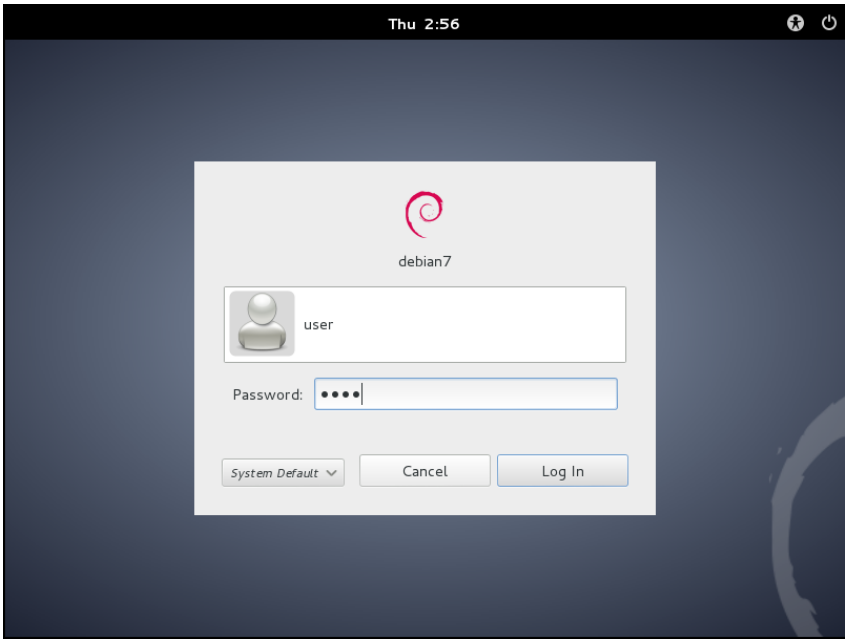

- The desktop GUI or server interface will load.

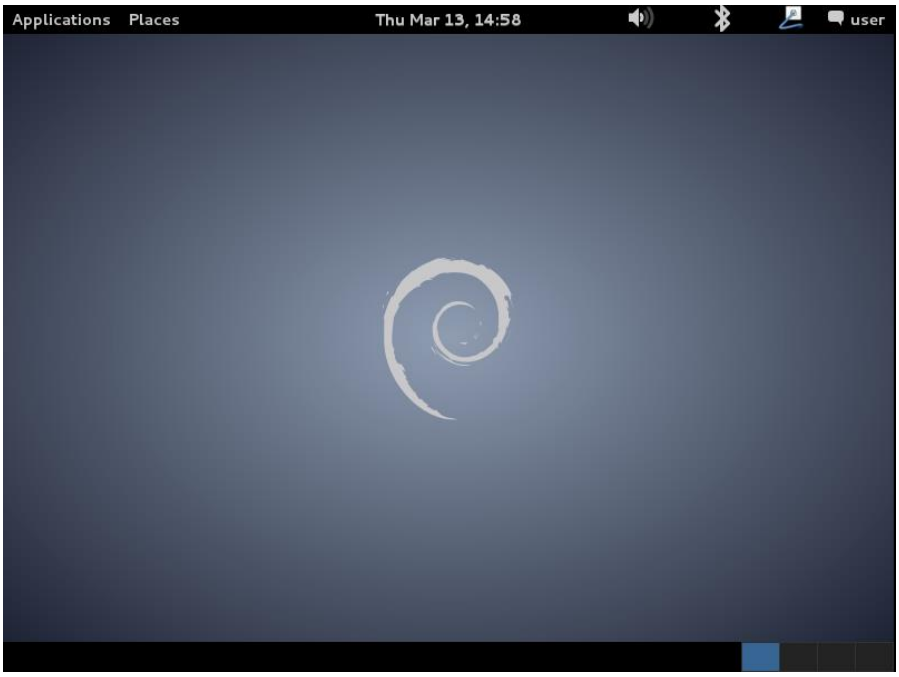

PS: the Basic install (Non graphical install) have the same menu options without the mouse and the screen-shots option.

:: Instruction Complete ::

8. Navigating the Linux File System

Linux has a different directory structure to Windows. Linux servers typically do not have a Graphical User Interface (GUI), it is therefore vitally important that as a Linux system administrator you are familiar with the Linux directory structure and knows how to navigate and manipulate the files and the file system. In this section, we try to give you an overview of the Linux directory structure.

Linux directory structure

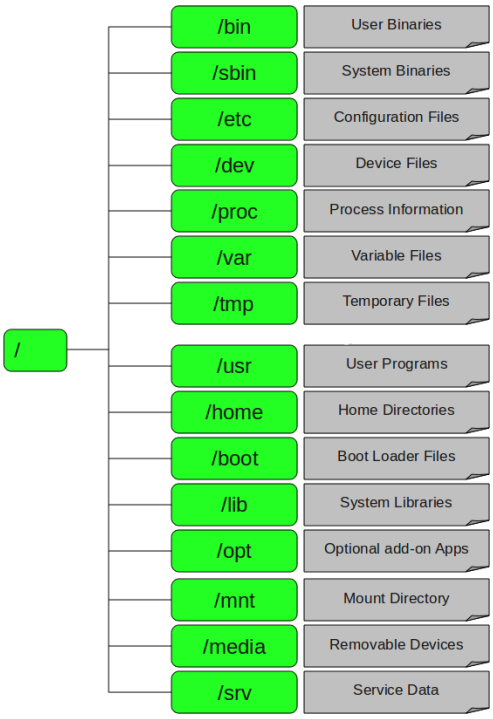

(Source: http://training.h3abionet.org/qiime_hands-on_workshop_2014/?page_id=101)

The below table describes the components listed in the above graphic.

Structure	Description
/	The forward slash is referred to as root and signifies the starting point of the Linux file system. This will be the C:\ equivalent in Microsoft Windows.
/bin	The bin folder contains the user usable binary files that are essential for general operation of your computer. These executable files include commands such as the ls, ping, mv type commands and allows you to navigate and interact with the file system.
/sbin	/sbin is similar to the /bin folder with the exception that these binaries are reserved for the root user or when you run commands with elevated rights such as with the sudo command. Examples of typical binaries would include the mke2fs, ifconfig and fdisk commands.
/etc	In Linux, all devices are configured using text files. These configuration files are usually stored in the /etc folder. Typical files would be your network interface cards, software repository files, etcetera.
/dev	The /dev folder contains the devices attached or mounted to the system. Unlike

	Windows, all Linux devices are represented as directories.
/proc	The /proc folder contains all the run time system information such as the system memory, processor, all mounted devices. In essence, it provides information on all locally installed hardware accessible to the server.
/var	The /var folder contains all the server log files and email messages.
/tmp	/tmp is the area used for temporary storage. Running programs or programs which are being installed often make use of location and is similar to the temp folder in Windows. The contents of this folder is often delete after a system reboot.
/usr	The /usr folder is similar to the bin and /sbin folders but differs in that it contains all the user specific binaries for the user installed applications.
/home	The home folder contains a folder for each user who logs into the server and is referred to the user's home folder. This subfolder is used to house all user specific documents and is similar to the "document and settings" folder in Windows.
/boot	This folder contains all the files that are necessary for the Linux system to boot correctly.
/lib	The lib folder contains all the system library binaries and is equivalent to the system32 folder in Windows which houses the .dll files. The Linux library files are represented with a .so extension.
/opt	The /opt folder is reserved for additional software you install. This is similar to the "program files" in the Microsoft Windows machines
/mnt	In Linux, any device needs to be mounted to the system before it can be accessed. Typically you would mount a device to a folder and then access the folder to access the content of the device. The /mnt folder typically contains mount points for internal devices such as the Linux root file system which is mounted to /. It also contains mounts for internal devices such as the internal hard drives which in Linux are referred to as sda1. The number 1 increments by one for each new hard drive device mounted to the system.
/media	The media folder is similar to the /mnt folder but would typically contain mount points for removable media such as CD ROM's, memory sticks, and removable hard drives, etcetera. When the user temporarily mounts a device to the system, they should mount it in this folder.
/srv	The /srv folder contains protocol specific data such as the ftp, rsync, www protocols.

9. Command line text editors

Unlike Windows based server operating systems, Linux servers do not have a graphical user interface. All interaction with the server is via the command line. Below is a short section about editing files using the nano command line text editor. There are a multitude of text editors available to be used in the CLI. In this section we discuss the nano command line text editor

In this tutorial most of the files that we have been editing are only root editable. So you will be using sudo every time you use nano. To use nano you should type: *sudo nano <<file_location/file_name>>*

Example:

sudo nano /etc/ntp.config

To move around in the file you should use the cursor keys

To save and exit nano: Type *Ctrl + x* then type *Y* then *Enter*

To exit without saving: Type *Ctrl + x* then type *N* then *Enter*

To search for a line: Type *Ctrl + w* then type *the work you looking for* then *Enter*

To check to line number: Type *Ctrl + c*

To delete a line: Go to the line then type *Ctrl + k*

To copy something: Select the line or the word then type *Ctrl + Shift + c*

To past something: Select the position then type *Ctrl + Shift + v*

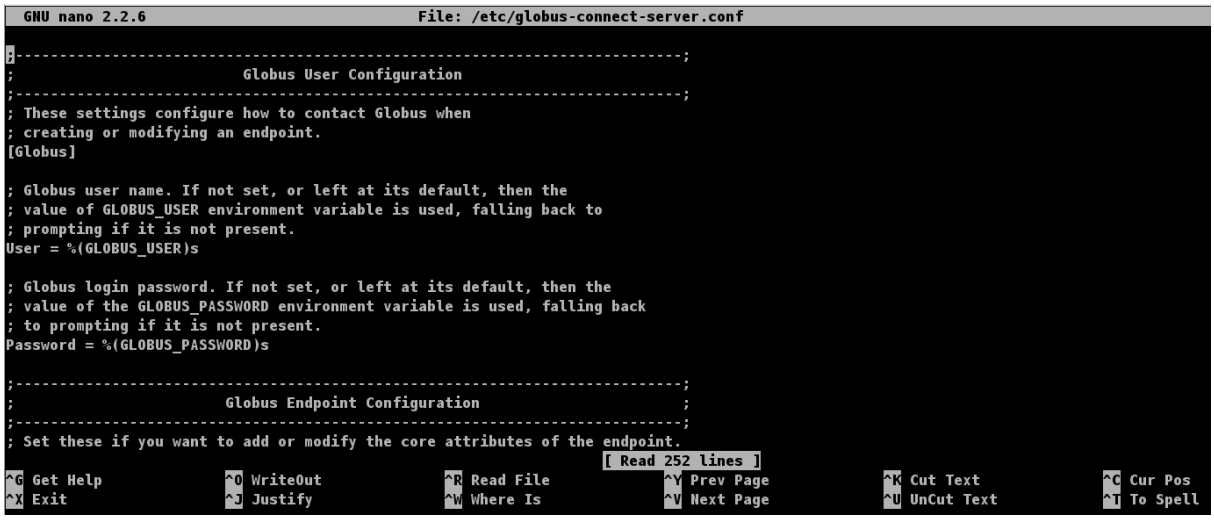

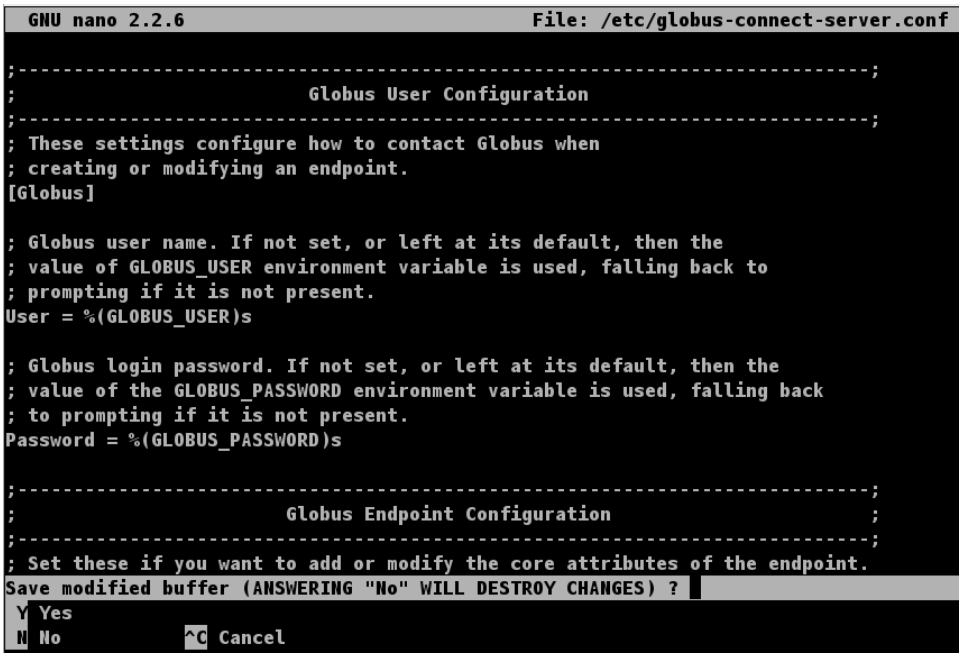

10. Network File Sharing (NFS)

The Network File System (NFS) is a distributed file system protocol originally developed by Sun Microsystems in 1984. NFS allows a user on a client computer to access files over a network much like a user on a Windows based machine will access a shared drive or folder over the network.

Terminology:
Server: the machine that will be exporting the folder i.e. the machine on which the folder is locally based. In this tutorial, the server's address will be serverA.co.za

Client: the machine that the folder will be mounted on i.e. the machine on which the folder will be remotely accessed. In this tutorial, the client's address will be clientA.co.za

Setting up NFS would require you to have root privileges. As opposed to appending the “sudo” option to your command in order to run the installation with elevated rights. One could change into “sudo mode” by typing “su” at the CLI. The main difference between using “sudo” and “su” is that with “sudo”, the single command is run with administrative rights and is usefull when installing a single application. The sudo password is held in cache for a few minutes but thereafter you would be required to retype the administrator password for each subsequent command requiring elevated rights. When you will be running a series of commands which require elevated rights, it is easier to change into sudo mode. This will run all commands typed after the “su” command until the user exists’s sudo mode.

To enter into sudo mode in all there OS’s, at the CLI type:

sudo su

You will be prompted for your administrator password. Type this in to complete the change to sudo mode.

10.1 Server side setup

Install the NFS server-side software on the server (e.g. serverA.co.za) using either the “sudo” command or switching to sudo mode.

Ubuntu or Debian

apt-get install nfs-kernel-server

SL 6.4

yum –y install nfs-kernel-server

Create the directory you want to export (if it does not already exist):

mkdir -p /path/to/directory

Set permissions (required if you are not using LDAP authentication):

chmod -R 777 /path/to/directory

To export the folder to an IP address, add the following line to /etc/exports

/path/to/directory client.co.za(rw,fsid=0,insecure,no_subtree_check,async)

Restart the NFS service

service nfs-kernel-server restart

10.2. Client side setup

Install the NFS client software on the client machine (e.g. clientA.co.za). You could either use the "sudo" command to elevate the logged on user's rights or switch to sudo mode as with the server installation.

Ubuntu / Debian

apt-get install nfs-common

SL 6.4

yum –y install nfs-common

Create the directory you would like the filesystem to be mounted to (if it does not already exist)

mkdir -p /path/to/directory

Edit the /etc/fstab file to add the filesystem you would like to mount:

serverA.co.za:/path/to/directory /path/to/directory nfs auto,noatime,nolock,bg,nfsvers=3,intr,tcp,actimeo=1800 0 0

Now from the CLI mount the filesystem

mount –a

The mount –a switch will remount all the mount points found in the fstab configuration file

11. Task scheduling

Another useful utility is the ability to schedule repetitive tasks to automatically run at a predetermined point in time. This ability is managed and configured by the "Cron" software utility. Cron is similar to the Windows based task scheduler program. It is designed to schedule jobs

(commands or shell scripts) to run periodically at fixed times, dates, or intervals. It typically automates system maintenance or administration—it is also useful for general-purpose scheduling for things like connecting to the Internet and downloading email at regular intervals.

To configure Cron to schedule a task > log into the machine as the user you would like to run the task as. If the task needs to be run with root privileges, for example, you must be logged in as root. To open a crontab, from the command prompt type;

crontab -e

Jobs are added to this configuration file as a single line consisting of 6 columns in the following order:

- Minute
- Hour
- day_of_month
- month
- day_of_week (0 is Sunday, 1 is Monday...)
- command_to_be_run

As an example, to schedule a script to run every Monday at 4:30pm, you would enter the following:

*30 4 * * 1 script.sh arg1 arg2* ## runs the command at 04h30am

*30 16 ** 1 script.sh arg1 arg2* ## runs the command at 16h30pm

If you have your server set up to send emails (see "Set up postfix to send emails from your server"), you can set cron to email you when it has notifications for you using the MAILTO directive. At the top of the crontab, add:

MAILTO=email_address

12. Useful commands

Linux provides several powerful administrative tools and utilities which will help you to manage your systems effectively. If you don't know what these tools are and how to use them, you could be spending lot of time trying to perform even the basic administrative tasks. The focus of this section is to help you understand the basic system administration tools, which will help you to become an effective Linux system administrator.

Command	Description
man	The man command stands for "manual" and is similar to the "help" or "?" command in Windows. To review the syntax for a particular command, at the CLI, type *man command_name*
ls	ls stands or list and is similar to the "dir" command in Windows. ls will display the contents of a folder. Using the ls -l switch will display the file date stamp and ownership
pwd	pwd stands for "print working directory". When navigating around the file system, you might get to a point where you don't know where you are in the file system.

	Running the pwd command will display where you are in the file system.
cp	Is the copy command and is used to copy files to different locations on the local machine. The cp syntax is *cp file new_location*
scp	scp is the secure copy command and is used when copying data across the network. The syntax for coping a file from the local serve to a remote host is *scp file user_@remote_host:location* when coping a file from a remote server to the local server *scp username_@remote_host:location location_on_local_server*
mv	The mv command stands for "move" and has two functions. The first is as the name suggests, it moves a file from one location to another. The mv command is also used to rename a file. To rename a file, use the mv command but instruct the command to move the file into the same location as the original adding only the new file name. This will rename the file. **Examples** Moving a file *mv file destination_location* Renaming a file *mv filename new_filename* if the file you renaming is in a remote location, add the path to the above command
rm	The rm command is used to remove of delete a file. When removing a folder, you need to add the -r option. *rm file_name* or, *rm -r folder_name*
mkdir	The make directory command is as the name suggests used to make a directory. *mkdir new_directory_name*
cd	The cd command is used to change your working directory. The syntax for this would be: *cd /path_to_directory* *cd..* will revert the cd command by one folder at a time
Viewing file content	There are a few ways to achieve this. One method is to use a text editor; instruction for this is in point 9 above. To view the content of a file from the CLI, the following commands can be used *cat* – will display the contents of the file on the screen *more* – using the more option will fill display the file contents one screen at a time *less* – works the same as the more command *tail* – will give you the first 10 lines of the file *head* – will give you the last 10 lines of the file

13. References

13.1. RAID

- http://www.thegeekstuff.com/2010/08/raid-levels-tutorial/

13.2. Linux hard disk naming convention

- http://www.tldp.org/HOWTO/Partition-Mass-Storage-Definitions-Naming-HOWTO/x99.html

13.3. Nano

- http://en.wikipedia.org/wiki/Vim_(text_editor) – vim background
- http://www.openvim.com/tutorial.html - vim tutorial
- http://www.fprintf.net/vimCheatSheet.html - vim cheat sheet

13.4. VIM

- http://www.howtogeek.com/howto/42980/the-beginners-guide-to-nano-the-linux-command-line-text-editor/ - nano howto
- https://wiki.gentoo.org/wiki/Nano/Basics_Guide
- https://help.ubuntu.com/community/Nano - how to install nano
- http://mintaka.sdsu.edu/reu/nano.html

13.5. Basic Linux commands

- http://www.ee.surrey.ac.uk/Teaching/Unix/

Please forward any queries, comments or complaints you may have about howto to the H3ABioNet System Administrator Task-force: sys_admin_tf@lists.h3abionet.org