ECE566 Enterprise Storage Architecture Lab #0: Your Server

A major focus of this course is hands-on experience with real hardware. As such, you have been provided with a server recently retired from service here at Duke. You will explore and rejuvenate this server, giving it a new life under your command.

Directions:

- This assignment will be completed in your groups. However, **every member of the group must be fully aware of every part of the assignment**. Further, while you can discuss concepts with other groups, actual steps and answers should not be shared between groups.
- Format:
    - **Physical work:** Much of this assignment will be setting up and deploying your server; this will be graded by looking at the server when you're done.
    - **Data recording:** Some of the assignment will ask you to collect data and update a Google Sheet. This will include fields in green, yellow, or magenta.
    - **Written deliverables:** Some of the assignment will ask for short answers or screenshots; this material should be collected in a PDF file submitted via GradeScope. Anything you need to include in this document is highlighted in cyan.

1 Introduction to server management

Servers are managed very differently than laptops that you're used to. Before we get started, let's discuss some key differences.

1.1 Management

People don't work physically near servers except during initial installation or maintenance, so servers do not have a screen or keyboard attached except for initial setup or if things have gone VERY wrong. Because of this, most servers have a mechanism to control them without physically being there. This is called **out-of-band management:** an embedded microcontroller totally separate from the main CPU that has its own network port and the ability to connect a virtual keyboard, display, and bootable storage over this network port. You can even turn the server on and off. In Dell servers, this is called DRAC ("Dell Remote Access Console"). This system works even if the server OS crashes or the server is powered down. This is in contrast to **in-band management**, which is when you use SSH, remote desktop, or other protocols to connect to the actual running system. We have out-of-band management to recover when things go wrong without physically visiting the data center.

You'll be configuring the DRAC cards in your servers to allow for this remote access.

1.2 Networking

Another difference is network management. For your laptops, you just need *some* IP address to access the network, and the laptop is generally a client that consumes network resources. The way the laptop gets an **IP address** is to request one from the network upon first connecting via **Dynamic Host Configuration Protocol (DHCP)**. The DHCP server on the network picks a free one and assigns it. Further, you generally connect to wireless networks, which are fine for consumer devices.

Servers are the *providers* of resources, so a server usually needs to have a *specific, known* IP address. There are two main ways this is done. First, a server may be assigned a static IP address, which is stored on the server itself. This is simple, but requires you to manually configure each server. As an alternative, you can use DHCP like your laptop uses. The difference is that the DHCP server in these environments doesn't just hand out a random IP to every random device that gets attached. Rather, the DHCP system is informed of what machines it needs to provide IPs to by their **MAC address**: a unique 6-byte number generally burned into the ROM of your network card. We'll be doing this method, so this guide will have you make note of your MAC addresses (two 'normal' Ethernet interfaces plus the DRAC interface).

Related to this, we also need to ensure that **DNS** works. DNS (the "Domain Name System") is the network protocol that translates human-friendly names ("google.com") to IP addresses the system can connect to (64.233.177.138). Without getting into too much detail, it will be important for your server to have the correct **hostname** (the leftmost token of a domain name). The way the Duke network works, we'll be registering your names and MAC addresses together and getting Duke to assign IP addresses accordingly using the asset tracking spreadsheet described below.

1.3 Meet your server

Each group has been issued a Dell storage server (various models). In production IT environments, all equipment is managed in some form of **asset tracking** database. Your first step is to figure out what you have and update such a database with the basic physical info.

Our asset tracking database will be this simple Google Sheet. If you do not have access to modify this sheet, request it.

Note: Some of the servers were used last year and have already been assigned names. As some network configs are already set up for these servers, we want to assign the same name as last year. To this end, servers are pre-populated in the sheet up to the serial number (also known as service tag). Check to see if your server's serial number matches the ones pre-populated, and if so, use that row of the spreadsheet – you'll populate the rest of the cells per normal. If your server is new, add a row for it and record all fields.

The name of your server is "esa<NUM>" where <NUM> is either pre-labeled on the server or, for new servers, a number provided by the instructor.

By examining the server, identify the columns highlighted in green. Manufacturer and Model should be easy (see front bezel). For height, you have to understand how servers are measured in Rack Units

(simply written as 'U'), see this article for details. For the serial number, note that on Dell servers this is also known as the service tag, and is a seven-character alphanumeric string[1].

2 Poke around inside

Let's see the guts. Open the chassis and examine the interior. Take one or more photographs for PDF submission. Annotate the photograph(s) to indicate the location of the CPU heatsink (under which sits the CPU), the RAM, the storage controller card (if separate from the mainboard), the power supply(s). *Don't remove the CPU heatsink* (as replacing it may require needlessly cleaning and re-applying thermal paste). You may need to temporarily remove plastic cowlings or covers to see everything. The chassis may contain a map to help identify parts, and the instructor can help guide you.

Let's check how much RAM we have. Ground yourself by touching the chassis, and then eject a single stick of RAM (push on the side-levers around it). Examine the label on it to determine the following for your write-up:

(a) The type of RAM (DDR1, DDR2, DDR3, etc.; this may be indicated by the prefix "PC1-", "PC2-", etc.; see here for details),
(b) The speed of the RAM (e.g. "PC2-3200", "PC3-6400", etc.; the latter gives the peak transfer rate in MB/s).
(c) The size of the RAM stick (e.g. 1GB, 2GB, etc.)
(d) The number of RAM sticks installed

Compute the total RAM installed by multiplying the size of the stick by the number of sticks, and document this in your write-up. (It is possible that you may have a mix of sizes installed, but this is fairly uncommon.) Reinstall the stick in the same slot it came from, ensuring the orientation is correct (to install, just slot it in and push down until it clicks).

Replace any cowlings or covers you removed and close the chassis.

3 Boot your server

In this phase, we'll be recording some hardware details: the yellow columns in our asset database spreadsheet.

Hook up a monitor, keyboard, and power cable to the server (you only need one power cable even if you have two power supplies). If the lab you're working in has workstations, you may borrow a monitor/keyboard from a workstation, ***but you must restore the keyboard/monitor to the workstation before you leave!***

[1] One thing that Dell does well is track the hardware configuration of things they sell. If you like, you can put that service tag into support.dell.com and it will tell you exactly how the system was shipped when Duke originally purchased it. Of course, it may have been upgraded or modified since then, but this resource can provide helpful clues regarding a system's history.

Once connected, even though the system has no hard drives so far, we can still boot into the firmware and check it out. Unlike a laptop or desktop machine, a server has many different modules, each with firmware that initializes during boot, including:

- The **BIOS (Basic Input Output System)**, the basic firmware of the system
- The **storage controller card(s)**, which typically provides RAID functionality (Redundant Array of Independent Disks), which we'll learn a lot about in class.
- The **network card(s)**, which are equipped with PXE (Preboot Execution Environment), the ability to boot an operating system over the network as opposed to from a local storage device.
- A **remote access card** (also known as a RAC, a management card, a Baseboard Management Controller (BMC), or some vendor-specific name, such as the Dell Remote Access Card (DRAC)), which is a separate processor that allows network-based control of the server in a manner almost equivalent to physical access; this is to allow for "lights-out" datacenters where all non-physical server maintenance is done over the network, even when a server is severely misconfigured or malfunctioning. (HINT: you'll want to become familiar with your server's remote access capability to avoid in-person trips to the datacenter once you rack your server).

Each piece of hardware will appear in turn during boot, most flashing up their own prompt to allow for user setup (e.g. "Press F2 for menu"). At any point, you can restart the computer by pressing Ctrl+Alt+Delete. If you curious, you could visit each boot-time configuration menu, but don't make any changes. If you do so, the storage controller may complain that all its drives are gone; press whatever keys needed to ignore this for now.

You'll need to specifically visit the main BIOS menu and the remote access card setup menu to make note of the MAC address (the physical Ethernet identifier) for each Ethernet port. Specifically, your server probably has four general onboard Gigabit Ethernet ports and a management Ethernet port (marked with a wrench icon) which is managed by the remote access card. We'll just be using the first two normal ports plus the management port; you can ignore the latter two normal ports as well as any additional add-on Ethernet cards your server may have. We need the MAC addresses of the ports to register them on the Duke network (for use during this lab), and later we'll turn them over to Duke IT so our servers can be registered in the Datacenter network.

While you're in the BIOS (pictured below), you can also confirm the amount of installed RAM and the type and speed of CPU. **Note the MAC addresses and hardware stats in the yellow tabs of the asset database.**

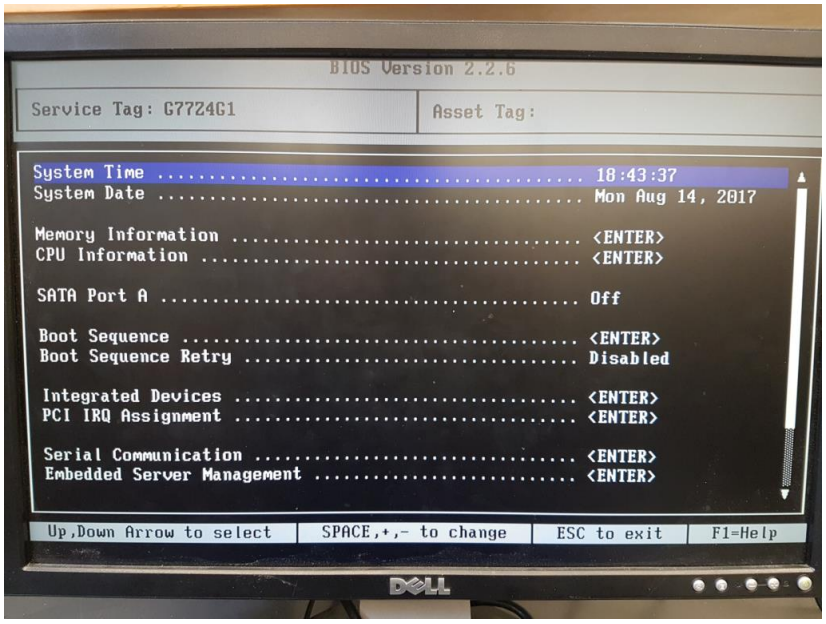

A typical (if old) server BIOS interface

4 Set up remote access

Recall that the remote access card will allow us to manage the server over the network, once it's configured. The first step is the network configuration.

Connect an Ethernet cable between an internet-connected switch and the management port (marked with a wrench). If the lab you're working in has workstations, you may borrow the Ethernet cable from a workstation ***but you must put it back when you're done***.

Reboot (Ctrl+Alt+Del) and get into the remote access card setup menu when prompted. An example prompt is show below.

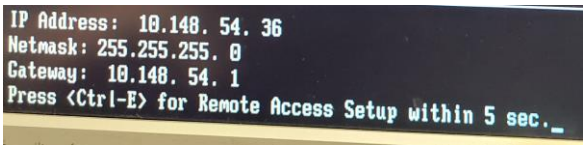

In this server model, the key to press is Ctrl+E

In this interface, do the following:

- Enable DHCP so the card automatically gets an IP address from the Duke network (e.g. set the **IP Address Source** to **DHCP**, though this procedure may differ depending on the exact hardware/firmware you have).
- Set the DHCP hostname to esa<NUM>-oob (e.g. "esa00-oob"); we'll adopt a naming standard where the hostname of the out-of-band management interface is the same as the host, but with '-oob' appended (it stands for 'out of band').

- Configure with username “root” and the password provided by the instructor.
- Disable IPMI as follows[2]:
    - Set “IPMI Over LAN” to “Off”
    - Within “LAN User Configuration”, Set “IPMI LAN Privilege” to “No Access”

What does this mean? DHCP is a protocol that allows a machine to get an IP address from the network itself when it boots. The DHCP server that gives out such addresses typically tracks clients via their MAC address, which is associated with the physical interface (i.e. burned into the chip). These requests may also contain a requested hostname, which, if available, is associated with whatever IP address the server is assigned by DNS servers, which translate names (e.g. “esa00-oob.egr.duke.edu”) to IP addresses (e.g. 10.22.34.55). Note: it may take up to an hour DNS to sync up, so plan to use the bare IP address in the near term.

At Duke, the network environment is set up so that hosts must be *registered* on the Duke network, which we’ll do next. Visit https://dukereg.duke.edu and login with your Duke NetID. Paste in the MAC address for the remote access card, set description to “esa<NUM>-oob”, set type to “other devices”, and click Register.

While we’re here, go ahead and paste in the MAC address for the first Ethernet port, set description to “esa<NUM>”, set type to “Fedora/Ubuntu/Debian”, and click Register.

Now reboot your server, and enter the remote access card menu again. You should now see a valid IP address for remote access, e.g. 10.22.34.45 – Note this in your write-up and keep it handy for later testing!

At this point, you should be able to navigate your web browser to the IP address and/or hostname of the remote access card and see the remote access interface. When this is working, do the following:

- Use the power control to power the server off and on.
- Access the remote console. Take a screenshot and include this in your write-up.

NOTE: If the remote access console doesn’t work and you’re on an R710 or similar server with an iDRAC6, follow the “iDRAC6 setup” guide linked from the course site. If you are unable to get the remote console working, consult the instructor, and if we can’t get it working with your laptop, you can skip this, but I recommend getting it working at some point, as it will save you a trip to the datacenter if things go wrong.

---

[2] In the time since these servers were used, multiple security advisories have been published about IPMI (Intelligent Platform Management Interface), an open standard for server management, including disclosure of password hashes and even an authentication bypass. These servers are old enough that a patch is not available. Having to compensate for a security vulnerability when no patch is available is a common situation on IT. The good news is that the issues only affect the IPMI protocol, not Dell’s proprietary web interface, which is what we use. This is why we disable IPMI support in this way.

Once remote access is working, note this in the magenta "Remote access setup?" column in the asset management spreadsheet. Also, confirm that IPMI is disabled in the "IPMI disabled?" column.

5 Install drives

5.1 Physical installation

At present, your server has *no* storage devices. You will install drives from a pool of spares acquired for this course. Each student server will end up with:

- Four traditional spinning-platter hard disk drives (HDDs): SAS protocol, either 73 or 146 GB, either 10k or 15k RPM. *Note: These hard drives are pretty old, so drive failures are anticipated. This is a good thing, as tolerating drive failure is a big part of the course!*
- One solid state drive (SSD): Kingston SATA 120GB or similar.

See the instructor to obtain your drives. For HDDs, be sure to end up with a matched set, where each drive has the same capacity and speed.

If the drives aren't already installed into Dell drive trays, you'll need to do that before installing them.

Power off your server, and install the drives into the front of your server, with the SSD going first (top left), then the HDDs. Remove any blanking panels as needed.

Note the quantity, size, and rotational speed of your SSDs and hard drives in the yellow "SSD/HDDs" column of the asset tracking sheet, e.g. "1x 120GB SSD, 4x 72GB 10krpm HDD".

5.2 Storage card setup

ABOUT RAID AND STORAGE ADAPTERS: Your server is equipped with a high-performance storage adapter card, which is what connects the front-mounted drives to the system. This card typically handles drive redundancy via RAID (Redundant Array of Independent Drives) in *hardware*. This is a good thing, but for this course, we want to do the RAID stuff ourselves, so we can experiment with it. Therefore, we will be handling the drives in *software* instead, bypassing the RAID features of the storage adapter card.

To make this possible, we need to tell the card to do nothing but expose the drives directly to the operating system. Reboot and enter the storage card's configuration menu by pressing the appropriate key (e.g. Ctrl+R) when prompted (it may vary depending on your exact hardware). Note: for Dell servers, the storage card is often called a PERC (PowerEdge Expandable RAID Card). The PERC 6 interface is shown below.

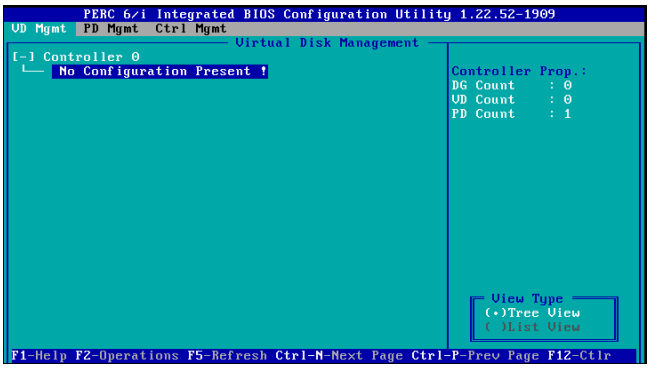

The PERC 6 interface

Once in the configuration tool, create a single virtual device (VD) for each physical drive (PD). For example, the shot below shows the creation of a virtual device with just the attached SSD (as indicated by the size of ~120GB):

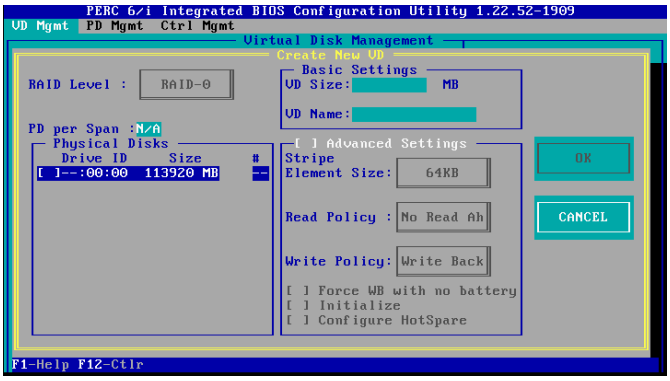

The PERC 6 interface when creating a single-drive VD

Proceed to make 5 total VDs: one for the SSD, and one for each of the four HDDs.

NOTE: Some disks may appear as “Foreign”, meaning they were part of a RAID array on a Dell server before we got to them. In this case, go to the “Foreign View” tab, and in Controller 0, use F2 to see operations, and select Foreign Config -> Clear. Then you can make VDs out of them as described above.

When done, take a remote screenshot or local photograph of the PERC configuration interface showing your VDs and include it in your write-up.

6 Operating System Installation

We’re going to install Linux as our operating system, specifically Ubuntu Linux Server. There are many trade-offs between operating systems in general and Linux distributions specifically, but Ubuntu Linux Server should work well for us as it combines decent ease of use with a rich software library.

We’ll be installing **Ubuntu server 22.04**.

Note that the install process for a server OS includes more steps than a traditional desktop OS (including even Ubuntu Desktop), as server admins (like you!) need more fine-grain control of the installation.

Before proceeding, ensure that the first regular Ethernet port is connected (if you have only one Ethernet uplink cable, move it from the management port).

6.1 Basic installation steps

Obtain an Ubuntu Server 22.04 USB stick, and boot from it. Proceed to install Ubuntu Server when prompted. Navigate through the installation prompts per usual (defaults are fine for most of it; please keep the language settings set to English so the instructor can assist with the server if needed).

If prompted to update the installer, feel free to do so.

The type of install should be left as “Ubuntu Server”, and you can leave “third party drivers” disabled.

Select the first network interface when prompted, and because you registered this interface’s MAC address in DukeReg back in section 4, it should get an IP address successfully – Note this in your write-up and keep it handy for later testing!

6.2 Partitioning

ABOUT PARTITIONING: Hard drives are commonly divided by into very large chunks called partitions to logically divide the storage. For example, desktop Linux systems commonly have a partition for the root directory (e.g. byte locations 0 through 110GB) and a separate smaller partition for swap space (e.g. byte locations 110GB through 120GB). At this point in the installation, we’ll be partitioning drives as we see fit, and selecting which partitions will hold the operating system.

When prompted to pick a partitioning method, choose **Custom storage layout**, as shown below:

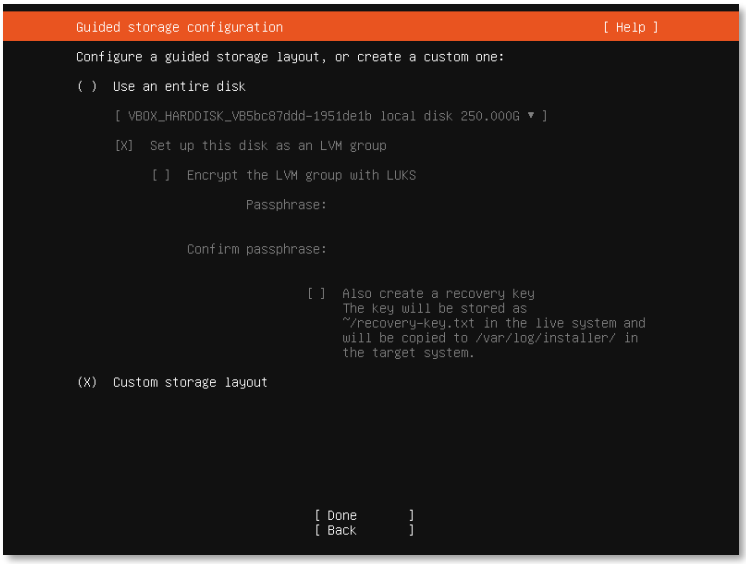

On a production server, you would always have the OS on redundant storage (e.g. a pair of hard drives configured for *mirroring*), either in hardware (using the storage adapter's hardware RAID functionality) or in software (by setting up OS-managed RAID at install time).

However, our goal is to learn about storage systems, so we want raw block-level access to our hard drives, so we're going to deviate from the standard practice, and install the OS to a partition on the single SSD (while reserving some of the SSD space for experimentation as well).

Therefore, when presented with the partitioner interface, just select the SSD. Note: the screenshot below was prepared on a VM, so the names will be different than what you see. If the devices are unclear, use the sizes to determine which is the SSD.

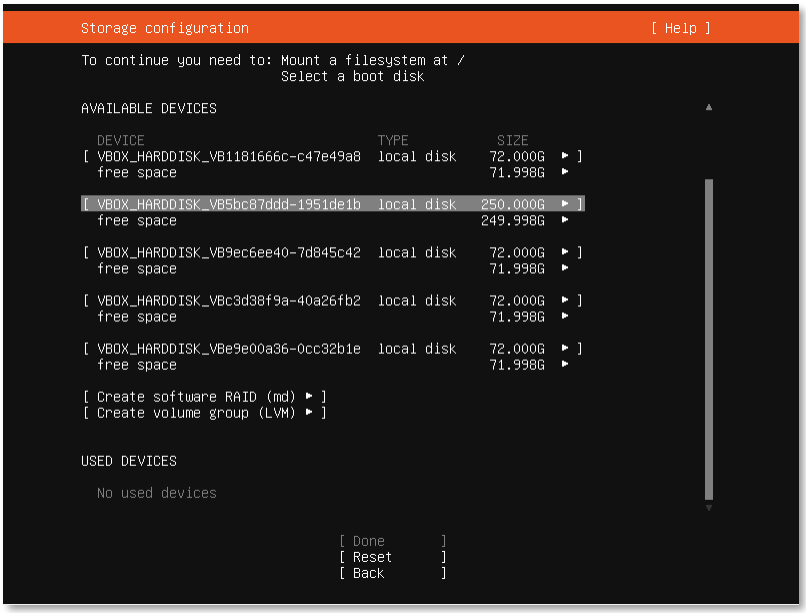

Let’s make a partition for the OS root directory. Add a partition by selecting the free space under the SSD, pressing Enter, and choosing “Add GPT Partition” as shown:

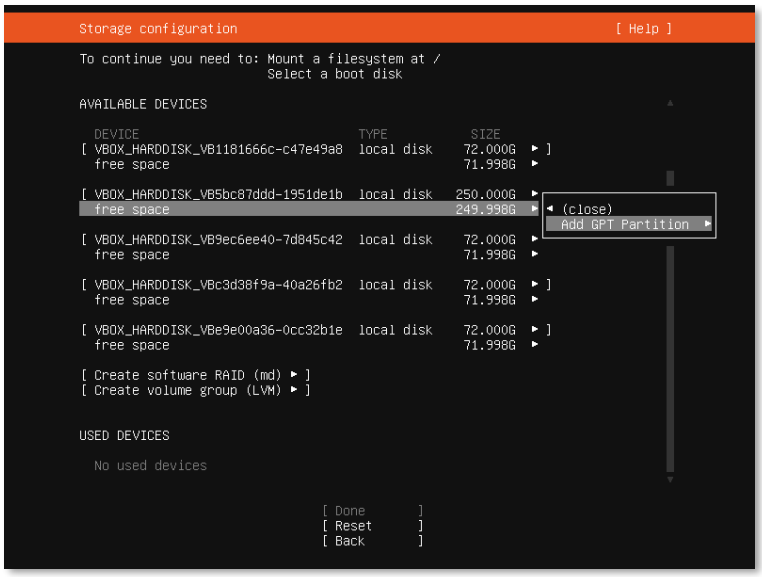

Make the new partition **40 GB**, formatted with the **ext4** filesystem[3], mounted as **/** (root):

[3] Filesystems are what translate raw block locations into the concept of files and directories we’re familiar with. Ext4 is a particular filesystem commonly used in Linux. Filesystems will be covered in great detail later in the course.

Now we'll set up some swap space, as even though these servers have tons of RAM, the Linux kernel still prefers to have some. Since we don't anticipate needing virtual memory swapping, it need not be large; just 2GB will suffice. Select the SSD's free space again and make a **2GB** partition whose format is set to "**swap**" instead of ext4:

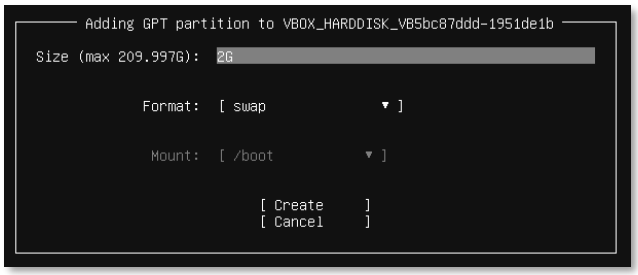

At this point, we're done partitioning, and you should see something like this:

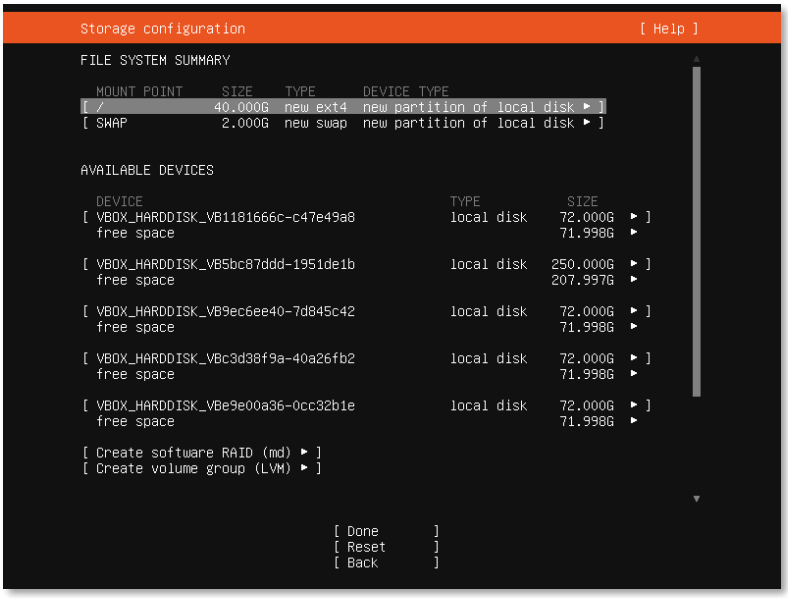

Note that we did not touch any of the HDDs: we'll worry about them much later, after the OS is installed. We could have created additional partitions to divide the operating system (such as making a separate area for /home, where user home directories live) or create partitions not in the normal UNIX naming scheme (such as /movies or /music). However, a simple single root is good enough for us. If our SSD is 120GB, this setup leaves us 120-42 = 78GB of free raw SSD space for later experimentation. Ensure your partitioner view looks like above, take a screenshot or photograph for your write-up, then and select "**Done**", and confirm when prompted.

6.3 OS configuration

For the “profile setup” step:

- Set “your name” to “**localadmin**”.
- Set “your server’s name” to “**esa<NUM>**”.
- Set “username” to “**localadmin**”.
- Set the password to the one provided by the instructor.

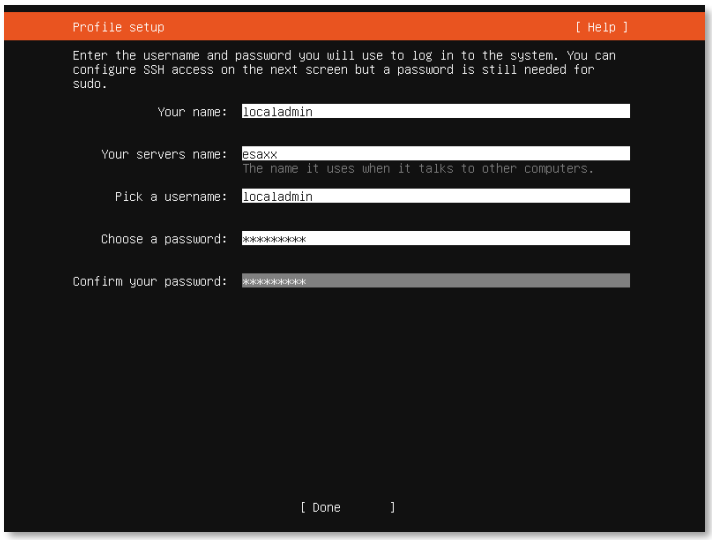

If prompted, skip “Ubuntu Pro”.

Enable “Install OpenSSH server”.

Do not install any “Snaps” when prompted.

6.4 Operating system data is copied; updates applied

Time passes.

6.5 Reboot

Once the install and updates are complete, your display will look like this and you can choose **Reboot**!

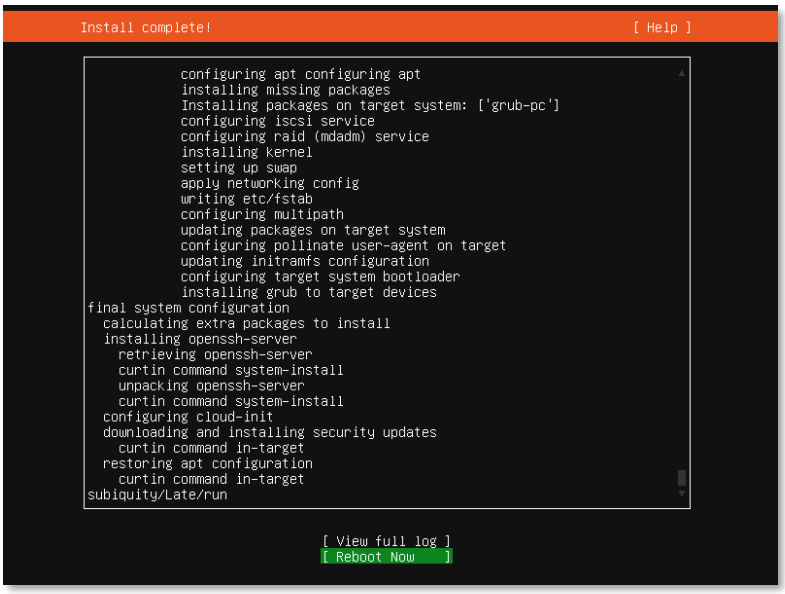

Remove the USB key when prompted.

7 OS configuration and checkup

7.1 OS checkup

Ensure that you can login to your management interface (esa##-oob) via a web browser, and that remote console works. Take a screenshot for your write-up.

Ensure that you can SSH to your primary interface (esa##) with the login and password you created – use the IP address you noted earlier.

Look at the "`/proc/meminfo`" and "`/proc/cpuinfo`" files and/or run the "`top`" command to verify your CPU/memory stats.

Run "`sudo fdisk -l /dev/sd?`" to look at the drive info and partition tables of all our drives. Use the "`lsblk`" command to view this info in abbreviated form. Take a screenshot for your write-up.

Once all this is done, note this in the magenta "OS install ok?" column in the asset management spreadsheet.

7.2 Installation of endpoint security software

Duke requires all university-owned systems to have endpoint security software: CrowdStrike and Planisphere Report. To make deployment easy, the instructor has prepared an all-in-one package for you to deploy these.

Via SSH into the server, run:

```
$ wget https://people.duke.edu/~tkb13/courses/ece566/resources/tyler-rapid-deploy.tgz
$ tar xzf tyler-rapid-deploy.tgz
$ cd tyler-rapid-deploy/
$ ./go
```

This is the first time we're using this particular automation, so let the instructor know if something goes wrong.

Once everything seems to be in order, it's time to send this server to its new home: the FitzWest Datacenter.

8 Rackmount installation

Servers such as these are meant to live in racks. As discussed in section 1, the width and height of the server is standardized so it can be packed with total density into racks with dozens of other systems from various vendors. Servers are mounted in with rails.

8.1 The FitzWest Datacenter

At some point, after a tour and some training, you will be granted access to the FitzWest Datacenter.

DO NOT TOUCH ANYTHING THAT'S NOT YOUR SERVER IN THERE!

A rack has been set aside for use by this course; you'll be installing your server in there. Rolling carts are available to help transport your server to there, and a day will be arranged where most if not all groups will be expected to do the move with instructor support.

8.2 Rails?

A selection of Dell-style rails has been procured, but unfortunately, there is not a label indicating which rails go with which servers, and not all servers have rails. So our first step is to look up the physical mounting guide for these servers and compare it to the rails on hand.

Take the service tag you identified in section 1 and put it into support.dell.com. Under "Manuals and Documents", find the "Rack Installation Instructions" for sliding rails. Find which rails we have, if any, match your server. Indicate the answer in the magenta "Rails ok?" column in the asset management spreadsheet.

If you found good rails, you'll use them (see section 8.3). If not, see section 8.4.

8.3 If you have matching rails

When you move your server to FitzWest, install the rails and your server per the manual. Servers should be installed at the bottom-most free spot in the rack!

8.4 If you do not have matching rails

If your server is small enough (1U or 2U), we can just have it sit on top of a server that does have rails. This is not ideal, and would not be accepted in a production datacenter, but it won't affect us for this course.

If your server is big (4U+) or if there are not enough rail-having servers available, you can use universal rails. See the instructor if this is the case; additional universal rails may have to be ordered, so see if you need them ASAP.

8.5 Cabling

Run Ethernet cables from the network switch "esa-sw1" to your management Ethernet port and your first regular Ethernet port. The exact switch port you should use is documented in the "Cabling" tab of the asset management spreadsheet; *only make the connections noted in black, not the red one*. (We'll add network connection redundancy later in the course.)

Run a power cable from the Power Distribution Unit (PDU) on the left of the rack to the left power supply of your server. Similarly, run a power cable from the right PDU to your server's right power supply. Each PDU is fed from a different transformer, and therefore this provides power redundancy in the event of a power failure.

Note that this is a different network from the one you initially set up the server in, so the IP addresses you had before will not work. Some servers are already registered in this datacenter, but some aren't. Temporarily connect a keyboard and monitor to your server and collect the management IP address (shown at boot time) and the system IP address (available via the `ip addr` command after logging in), noting that either or both may fail to populate. If they are shown, ensure you can access these addresses from another system (such as your laptop).

In the magenta columns "IP" and "Mgmt-IP", put the system IP address and management IP address, respectively. If an interface didn't get an IP address, put "UNASSIGNED", and also **let the instructor know**. I'll be putting in a bulk IP allocation request for all servers based on this spreadsheet later on.

Remove the keyboard and monitor: this is a lights-off server, so once the network gets set up, you should be able to do everything over the network. If you do everything right, you should never have to come back to the datacenter except for cabling-related issues or hardware failures.

Once system is racked and working, update the magenta "Location" column in the asset management spreadsheet to indicate the FitzWest datacenter and U position (written on the rack itself) of the bottom of the server.

8.6 Post install network check

Ensure you can connect to your server both via SSH to the main esa## hostname as well as via the browser to the DRAC interface on the esa##-oob hostname. If not, work with the instructor to check if DHCP is working. If it's not, then your server may not be configured for use in the FitzWest network yet, in which case the instructor will need to request IP address allocation – be sure to request this!

Sometime later, once the instructor confirms that IP address have been allocated, try to connect both via SSH to the main esa## hostname as well as via the browser to the DRAC interface on the esa##-oob hostname. If this doesn't work, you may need to visit and reboot the server. If it still doesn't work, contact the instructor for help.

9 You did it!

Congratulations, you now have an enterprise-class server online, logged in our asset tracker, and ready for work!

Take a photograph of your group with your server and include it in your write-up! ☺