ECE 590-03: Enterprise Storage
Fall 2016
Homework 3

Question 1: High Availability [15pts]

The architecture below has been deployed as a small private cloud that allows for self-service creation of virtual machines, similar to Duke OIT's VM Manage system.

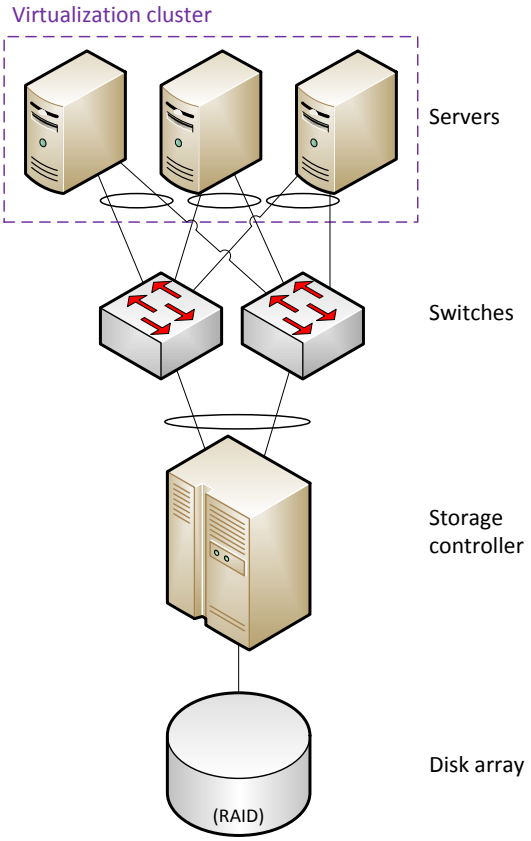

The virtualization cluster is running load-balancing and high availability clustering software, allowing automatic migration of VMs. When balanced, average utilization of each of the servers is around 75%.

On the storage side, each disk array is configured as a logical volume of RAID6 arrays.

Identify three different failure scenarios which could lead to interruption of service in this environment.

Question 2: Disaster Recovery [20pts]

Below are a number of proposals for how to do backups for disaster recovery purposes. For each, you will (1) identify the RPO and RTO, (2) determine which of "Tyler's Immutable Rules of Backup" the scenario violates and explain how each deficiency could lead to data loss.

Please try to keep answers concise. For RTO and RPO, you may respond either with an exact time interval or a description (e.g. "however long it takes to download"). Also, in computing RTO, just consider the case of simple data corruption, not a site-wide disaster.

(a) "I simply turned on automatic hourly snapshots on the storage controller." [5pts]

(b) "I have a script that mirrors everything from my NAS home directory to another NAS share every night, and it emails me if it fails. The script was easy to write, since it can login to both shares with my regular NetID credentials. Every few weeks I peek into the backup directory and verify that the timestamp is right." [5pts]

(c) "Every 7 days, a script will have the storage controller take a snapshot and write it in the background to tape. Script errors and a weekly report are communicated via email. The finished tape is then sent in an offsite vault a 1-hour drive away. To fully read the tape would theoretically take 5 hours, but we've never needed to so far." [5pts]

(d) "The production server is connected to a backup storage system in another building by a 500 meter Ethernet cable. Every 6 hours, a script on the backup server uses `rsnapshot` [as we saw on homework 2] to pull a snapshot of production data. The credentials on the backup server are unique, and the backup server can only be accessed in-person, as it has no internet uplink. Each month, I go to the server and conduct a recovery test. In a disaster, the backup server can also act as a production server." [5pts]

Question 3: Virtualization and Cloud [50 pts]

You have a startup idea: you want to make a ride-sharing app called Uber (and you're somehow unaware that this has already been done).

For the infrastructure, you know you don't want to build a datacenter, so you've narrowed the choices down to purchasing gear and running it in a colocation facility (the "BUY model"), or running on Amazon AWS (the "RENT model"). You can make the following assumptions:

- Your app will generate 2TB of stored data per year.
- Your app will need **8** dual-core VMs with 8GB RAM each.
- Only 10% of data produced needs to be hosted on "high performance" media. For the BUY model, this means SSD vs. HDD. For the RENT model, it means Amazon's Elastic Block Store (EBS) "General Purpose SSD" vs. "Cold HDD" storage.
- All purchased gear in the BUY model must be installable in a standard 19" rack to be compatible with colocation providers.
- For simplicity, you can ignore:
    - The fact that storage will get cheaper/denser over time.
    - That Amazon prices may change in the future.
    - The bandwidth usage of your app (assume it's negligible).
    - Cost of software (e.g. operating systems, hypervisor, etc.).
    - Cost of peripherals (Ethernet adapters, SAS adapters, etc.), disk enclosures, and cables.
    - Disaster recovery.
- For any other facts you need to know but are omitted, just state a reasonable assumption in your answer and go from there.
- For the servers and network switches in the BUY model, you can use these models:
    - Server: HP ProLiant DL380 G9 ($3438/ea)
    - Switch: Cisco SF500-24P ($470/ea)

Assessing the BUY model for the first year of operations:

(a) Design a fully HA virtualization environment, including storage, network, servers, and virtual machines. Draw a connectivity diagram of the environment, making sure to eliminate all single points of failure. [10pts]

(b) Research and identify the storage hardware needed (storage servers, disks). You may look at integrated vendors (e.g. Dell, HP, etc.) or component resellers (e.g. NewEgg, Amazon, CDW, etc.). Develop a Bill of Materials (BoM) for the environment. This is a spreadsheet showing each item to purchase, including the item name/model, quantity needed, unit price, total price (quantity*unit price), a link to the seller's page for the item. In your BoM, you should account for the storage gear you identified as well as the network switches and servers noted in the assumptions above. The bottom line of the BoM should show a grand total of cost. [10pts]

(c) Determine the total amount of rack space you need in Rack Units (RU). Research a colocation provider and determine the monthly or yearly cost to rent that much rack space. Ignore providers that ask you to request a quote and find one that posts prices online. [5pts]

(d) For the first year of the business, compute the total CAPEX for purchases and OPEX for rack space rental. [5pts]

Assessing the RENT model for the first year of operations:

(e) Compute the cost of using Amazon, including EC2 (for VMs) and EBS (for storage). You may wish to use the Amazon AWS calculator to help. For the first year of the business, what is the total OPEX for cloud services? [10pts]

Making a decision:

(f) For the first year of operation, which option is cheaper? [5pts]

(g) Assuming data storage continues to grow at 2TB per year, but everything else stays the same. Compute the total cost of both the BUY and RENT models over 10 years. Which is cheaper in the long term? [5pts]

Question 4: Security [15 pts]

(a) Explain the difference between "in flight" and "at rest" data encryption. What attacks can each approach mitigate? [6pts]

(b) The scenario in Question 2, part (b) of this assignment describes a backup script which likely violates a basic tenet of security. What is it? [5pts]
(HINT: what should you never do with a password?)

(c) In Linux/UNIX, give the command(s) that would set up permissions in each of the following scenarios. For each, make *only* the change prescribed; don't change permissions/settings that aren't mentioned.

i. The "code" directory should be owned by the "developer" group, and its contents be modifiable by any member of that group. [1pt]

ii. The file "secret.txt" should be readable/writable *only* by the file's owner. [1pt]

iii. The file "do_stuff" is a script, so make it executable by anyone. [1pt]

iv. You're a hacker with root access and you're setting up a "back door" to allow unprivileged users to gain root access in the future. Set permissions on the file "innocent.sh" so that it runs as the root user. [1pt]