Analyzing networks availability

High Availability (HA): HA means providing services with maximum uptime by avoiding unplanned downtime.
HA is the concept or goal of ensuring your critical systems are always functioning.
In practice, this means creating and managing the ability to automatically "failover" to a secondary system if the primary system goes down for any reason as well as eliminating **all single points of failure** from your infrastructure. So the system must be **Fault Tolerant.**

FT - Fault Tolerant:

- Fault Tolerance describes a computer system or technology infrastructure that is designed that when one component fails (be it hardware or software), a backup component takes over operations immediately so that there is no loss of service.
- Fault tolerance is the property that enables a system to **continue operating properly in the event of the failure** of some of its components.
- Something in the network will fail - A router power supply might fail, or a cable might break, or a switch might lose power.
- **Networks need redundant links to improve the availability of the network.**

Network engineers refer to any one component that, if it fails, brings down that part of the network as a single point of failure. (SPOF)

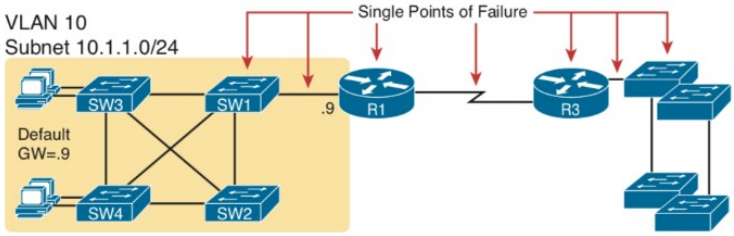

Redundancy is used to create systems with high levels of availability.

Networks need redundant links to improve the availability of the network.

Defining a high availability plan usually starts with a **Service Level Agreement (SLA)**

- **SLA** is created for an IT department or service provider. It defines the services and metrics that must be met for the availability and performance of an application or service. How much time during the month the service (a mailboxes for example) need to be available to end users.
- Achieving a network uptime of 99.999% (commonly referred to as "five nines") may be your organization's goal.

When determining what goes into an SLA, two other factors need to be considered (typically from a disaster recovery context). These factors are:

- **The recovery point objective (RPO)**: An **RPO** is essentially the amount of data that must be restored in the event of a failure. For example, in a single server or component failure, the **RPO** would be 0, but in a site failure, the RPO might allow for up to 20 minutes of lost data.
- **The recovery time objective (RTO):** An **RTO** is the length of time an application can be unavailable before service must be restored to meet the SLA. For example, a single component failure would have an **RTO** of less than five minutes, and a full-site failure might have an **RTO** of three hours.

Two important factors that affect an SLA are the:

- **Mean time between failure (MTBF)**
- **Mean time to recovery (MTTR):** Time taken to recover after an outage of service.

Flat vs Hierarchical Network Design

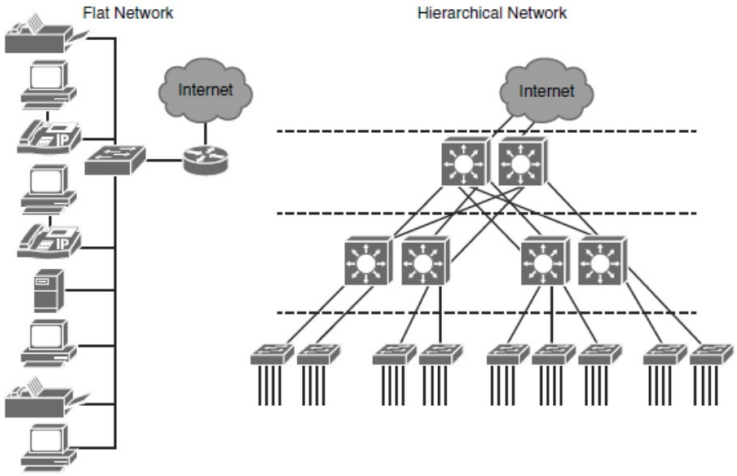

Access layer is used to grant the user access to network. E.g.workstations, IP phones, access points, and printers.
Distribution layer aggregates the access layer switches wiring closets, floors, or other physical domain by leveraging module or Layer 3 switches.
Core layer (also referred to as the backbone) is a high-speed backbone, which is designed to switch packets as fast as possible. In most campus networks, the core layer has **routing capabilities**.

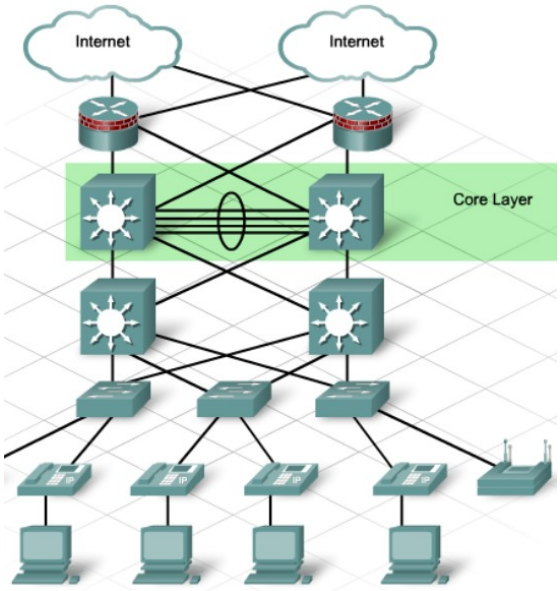

Virtualization

Virtualization is a method for abstracting physical resources (the hardware) from the operating system (it separates the software from the underlying hardware)
This way, you are able to move the operating system between different physical systems.
Let us assume that you are using a virtual server for your business process. What it means is that you are using a small part of a physical server. The server that you access performs like a physical server with dedicated resources, but in a virtualized environment, multiple users can share the same physical server simultaneously. This process is called **server virtualization**. Some other forms of virtualization include **network**, **data**, **desktop**, and **storage**.

Why organizations use virtualization today?

- **Automated Scalability**

Virtualization, when integrated with the cloud environment, anable automated scalibility. Although the cloud encompasses multiple technologies, virtualization is the most important among them that make possible automated scalibility of a system.

- **Cost Reduction (Reduced Hardware Costs):**
Before virtualization, the small and medium businesses had to buy or rent the entire physical server to host their process. They had to pay for the total resources even if they did not utilize all the resources.
Moreover, if the server was owned, the maintenance of the server added some additional costs such as cooling and power equipment, cabling, etc.
With Virtualization, a single physical server can be shared among many business processes. Hence, they have to pay as per usage and not a penny more.

- **Space Optimization:**
Virtualization helps businesses as well as hosting providers to optimize their physical space. As the hosting providers can divide a server into multiple Virtual Machines (VM), with each VM being used by individual businesses, the number of physical servers required reduces considerably.
For instance, if two businesses are utilizing only half of their resources on two separate physical servers, because of virtualization, both can share a single physical server and still function properly reducing the number of servers by half.

- **Flexible Operation**
Virtualization offers the businesses a more flexible operation. As it separates the software from the underlying hardware, the servers can host multiple operating systems and applications. You can host your business process on VMs, and each VM can have its own operating system and applications.

- **Disaster Recovery - Local:** Many companies now use Virtualization as a standard. Virtualization has provided many capabilities not least the ability to copy machines while they are running. This ability is often used to implement disaster recovery on site.

  How Does Virtualization Simplify Disaster Recovery?
  https://www.gowhiteowl.com/blog/virtualization-benefits-disaster-recovery/
  A virtual machine is effectively a single file that contains everything, including your operating systems, programs, settings, and files. Virtualization greatly simplifies disaster recovery, since it does not require rebuilding a physical server environment. Instead, you can move your virtual machines over to another system and access them as normal.

- **Easy Provisioning (Faster Server Provisioning and Deployment)**
The VMs can be provisioned easily as compared to the physical servers. Due to the virtual nature, a VM can be added or deleted on the server as per the requirement. Moreover, the operating system and applications can be installed almost instantly.

- **Less Hardware, Fewer Problems**
- **Significant Energy Cost Savings**

Disadvantages:

- **Not all hardware or software can be virtualized**
The drawback, however, is that not all servers and applications are virtualization-friendly, Livesay said. "Typically, the main reason you may not virtualize a server or application is only because the application vendor may not support it yet, or recommend it," he said.

- **Shared Resources**
VMs share the hypervisor, the host OS, and the same shared resources as with multicore processors: processor-internal resources (L3 cache, system bus, memory controller, I/O controllers, and interconnects) as well as processor-external resources (main memory, I/O devices, and networks). These shared resources imply the existence of single points of failure, that two applications running on the same VM can interfere with each other, and that software running on one VM can impact software running on another VM (i.e., interference can violate spatial and temporal isolation).

What is a hypervisor?
A hypervisor (or virtual machine monitor, VMM) is a computer software, firmware or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine, and each virtual machine is called a guest machine.
A hypervisor is software with only one purpose - abstracting and isolating operating systems from hardware. This abstraction allows you to host and operate multiple operating systems inside virtual machines on one physical host machine.

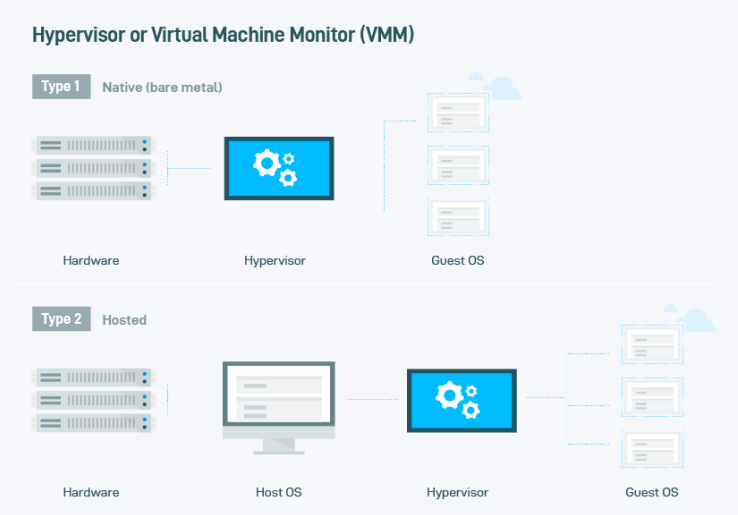

Cloud

Service Models

- **Software as a Service (SaaS):** The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage etc.

- **Platform as a Service (PaaS):** Platform as a Service (PaaS) or Application Platform as a Service (aPaaS) or platform-based service is a category of cloud computing services that provides a platform allowing customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure typically associated with developing and launching an app. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

- **Infrastructure as a Service (IaaS):** The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models

- **Private cloud:** The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g.,

business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

- **Community cloud:** The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

- **Public cloud:** The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

- **Hybrid cloud:** Hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud and third-party, public cloud services with orchestration between the two platforms. By allowing workloads to move between private and public clouds as computing needs and costs change, hybrid cloud gives businesses greater flexibility and more data deployment options.

AWS Global Infrastructure: AWS Global infrastructure can be brooked down into three elements:

- **Regions:**
  An AWS Region is a geographical area. Each Region is made up of two or more Availability Zones.
  - AWS has 18 Regions worldwide.
  - You enable and control data replication across Regions.
  - Communication between Regions uses AWS backbone network connections infrastructure.

- **Availability zones:**
  - Each Availability Zone is:
    - Made up of **one or more data centers.**
    - Designed for **fault isolation.**
    - Interconnected with other Availability Zones using **high-speed private links.**
  - You choose your availability zones. AWS recommends replicating across Availability Zones for resiliency.

- **Edge Locations:**
  - An Edge Location is where users access AWS services.
  - It is a global network of 114 points of presence (103 Edge Locations and 11 regional Edge Caches) in 56 cities across 24 countries.
  - Specifically used with **Amazon CloudFront**, a Global Content Delivery Network (CDN), to deliver content to end users with reduced latency.
  - **Regional edge caches** used for content with infrequent access.

VPC's

A virtual private cloud (VPC) is a secure, isolated private cloud hosted within a public cloud.

Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can use both IPv4 and IPv6 in your VPC for secure and easy access to resources and applications.

Benefits:

- **Security:** Information passed through a VPC stays within a customer's control without crossing the internet. Amazon VPC provides advanced security features, such as security groups and network access control lists, to enable inbound and outbound filtering at the instance and subnet level.
- **Easy integration:** A VPC can be integrated with other VPCs, the public cloud, or an on-premise infrastructure.
- **Savings:** Because VPCs are within a public cloud, customers still benefit from economies of scale, sharing costs with other organizations without compromising the aforementioned security.

Business continuity planning (BCP) an Disaster Recovery

file:///home/adelo/1-system/1-disco_local/1-mis_archivos/.stockage/desktop-dis/it_cct/3-Network_Management_and_High_Availability/9-2-NMHA%202019%20BCP,%20DR%20and%20DRaaS%20Lecture.pdf

Disaster Recovery: Disaster Recovery involves a set of policies, tools and procedures to enable the recovery or continuation of vital technology infrastructure and systems following a natural or human-induced disaster. Disaster recovery focuses on the IT or technology systems supporting critical business functions, as opposed to **business continuity**, which involves keeping all essential aspects of a business functioning despite significant disruptive events. Disaster recovery can therefore be considered as a subset of business continuity.

Business continuity planning (BCP):

- Business continuity planning (or business continuity and resiliency planning) is the process of creating systems of prevention and recovery to deal with potential threats to a company.
- A BCP is a set of contingencies to minimize potential harm to businesses during adverse scenarios.

Virtual LANs (VLANs)

A virtual LAN is a LOGICAL subnetwork that groups a collection of devices within the same or different physical LANs. VLANs allow network administrators to group hosts together even if the hosts are not directly connected to the same network switch.

Another definition can be: A VLAN is a logical grouping of nodes (clients and servers) residing in a common broadcast domain.

- The broadcast domain has been artificially created within a LAN switch by a LAN manager
- By default, traffic from one VLAN cannot pass to another VLAN
- 1 VLAN = 1 Broadcast domain = 1 IP subnet
  - Each of the users in a VLAN would also be in the same IP subnet

Reasons for use of VLANs:

- To create more flexible designs that group users by department, instead of by physical location (same network switch).
- To enforce better security by keeping hosts that work with sensitive data on a separate VLAN (broadcast domain) e.g. group users by department.
- To separate traffic sent by an IP phone from traffic sent by PC's connected to the phones
- To segment devices into smaller LAN (broadcast domains) to reduce overhead caused to each host in the VLAN

Software Defined Networking (SDN)

Software-Defined Networking (SDN) is a network architecture approach that enables the network to be centrally controlled, or 'programmed,' using software applications. This helps operators manage the entire network regardless of the underlying network technology.

SDN attempts to centralize network intelligence in one network component by disassociating the forwarding process of network packets

(data plane) from the routing process (control plane) «Separate Control plane and Data plane entities»

Network programmability: it refers to the ability of controling the network behavior by the software that resides beyond the networking devices that provide physical connectivity. As a result, network operators can adapt the behavior of their networks to support new services, and even individual customers.

Why separate control plane from the data plane?
By decoupling the hardware from the software, operators can introduce innovative, differentiated new services rapidly, regardless of the proprietary platforms.

The Data, Control, and Management Planes
Everything that networking devices do can be categorized as being in a particular plane:

- **The Data Plane:**
  The term data plane refers to the tasks that a networking device does to forward a message. Anything to do with receiving data, processing it, and forwarding that same data - whether you call the data a frame, packet, or, more generically, a message - is part of the data plane.The data plane is often called the forwarding plane.

- **Control plane:**
  The term control plane refers to any action that controls the data plane. Most of these actions have to do with creating the tables used by the data plane, tables like the IP routing table, an IP ARP table, a switch MAC address table, and so on.
  The information supplied to the data plane controls what the data plane does. For instance, a router with no routes in the routing table cannot forward packets. What controls the contents of the routing table? Various control plane processes.

- **The Management Plane:**
  The management plane includes protocols that allow network engineers to manage the devices.
  Telnet and SSH are two of the most obvious management plane protocols.
  To emphasize the difference with control plane protocols, think about two routers: one configured to allow Telnet and SSH into the router, and one that does not. Both could still be running a routing protocol and routing packets, whether or not they support Telnet and SSH.

Spanning Tree Protocol
The main purpose of STP is to ensure that you do not create loops when you have redundant paths in your network. STP ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop.

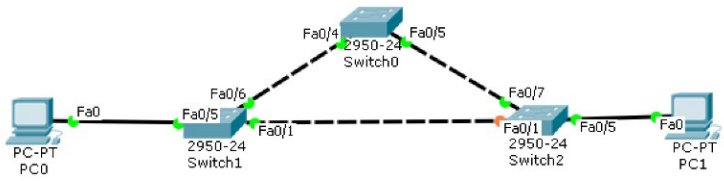

EtherChannel
With two redundant links, spanning-tree will block on one port to prevent loops. This way, using only one link utilizes only half of the available bandwidth.

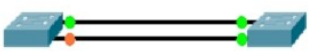

EtherChannel combines multiple parallel segments of equal speed (up to eight) between the same pair of switches, bundled into an EtherChannel.

EhterChannel allows spanning-tree to treat the two physical links as **one logical port** and thus both ports can operate in full forwarding mode.

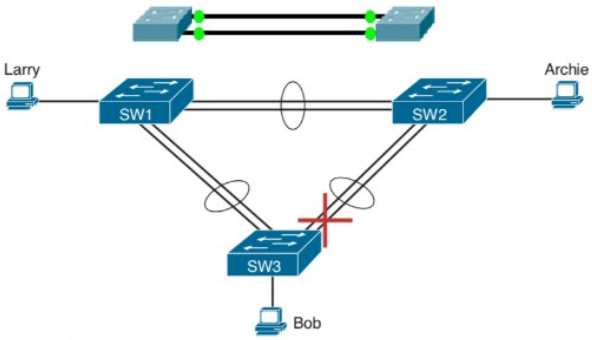

Figure 2-12 *Two-Segment EtherChannels Between Switches*

Benefits:

- Reduces the number of times STP must converge, which in turn makes the **network more available.**
- A LAN design that uses EtherChannelsmakes **much better use of the available bandwidth**

Two EtherChannel thecnologies:

- Port Aggregation Protocol (PAgP), a Cisco EtherChannel protocol
- Link Aggregation Control Protocol (LACP), an IEEE802.3adopen standardversion of EtherChannel.

First Hop Redundancy Protocols (FHRPs)
FHRPs provides redundant default gateways for end devices with no end-user configuration necessary. Hot Standby Routing Protocol (HSRP) is a FHRP.

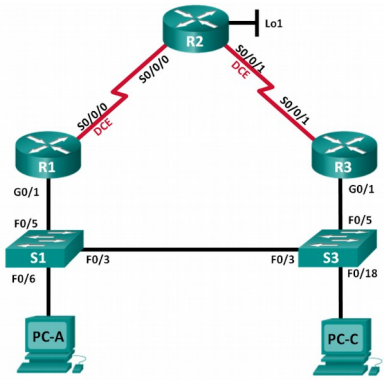

Why do many organizations take advantage of using snapshots? Can you explain this in detail?
Snapshots are a very common feature in today's storage systems that allow for users to create a point in time image of a volume, without actually creating a full, stand-alone copy. Snapshots set the blocks associated with that data to read only, then, as changes are made to the original, they are branched off and stored separately. This allows for an active image of that data to be maintained as well as the point in time copy.

What about using clustering...
What is a cluster? A cluster is a set of computing nodes that work together and can be loosely viewed as a single system to provide high availability of services for clients.
Why cluster…
Organization wants to achieve higher availability. This is a good step toward improved availability, but the return on the investment (ROI) of a cluster doesn't always add up.
You have to balance the improvement over the increased hardware cost, complexity, and level of training required for administrators.

Network performance issues and outages:

Employees have often complained there can be a problem getting access to the Ethernet network as in some areas of the building there are no available wall jacksto connect to the network:

- Incorporate a Wireless network by adding at least 2 Access point in the Access Layer.

One of the first hoprouters (router 0 in network diagram below) went down (for 3 hours) and it effectively split the network in half not allowing some of DAC's staff access the Internet or route trafficto the server farmand this also impacted customers requesting data.

- To solve this problem, we have to elimitate all **single point of failure** by addding redundancy. In the network design proposed in Figure 1, we have elimitated all SPOF by using a Hierarchical design. This way, if one of the devices or links fail, the connectivity is guareted using another path.
- It is also important to configure a First Hop Redundancy Protocol. Spanning tree provides loop-free redundancy between switches within a LAN. However, it does not provide redundant default gateways for end-user devices within the network if one of the routers fails.
  First Hop Redundancy Protocols (FHRPs) provide redundant default gateways for end devices with no end-user configuration necessary.
  Hot Standby Routing Protocol (HSRP), a First Hop Redundancy Protocol (FHRP).

Utilization of servers:

Some servers have very low utilization while others are crashing due to over-provisioning and lack of virtualization/VM migration capabilities.

- When servers are underutilized, a large amount of hardware, space, power, and management cost of these server is wasted. Server virtualization enables smaller resource allocation. It is possible to create several VMs on a physical server and allocate only the necessary resources depending on the process' demand. This way, less physical servers are needed.
  Using Virtual Machine Migration, its possible to migrate operating systems and applications from older servers to newer servers easily and without disrupting the services.
  Live migration refers to the process of moving a running virtual machine or application between different physical machines without disconnecting the client or application.

DAC data scientists perform a vast amount of testing on the physical servers whereby they need to perform an activity and then delete all configuration. At present, they have no way of automating this and they spend a large amount of time on installing/re-installing operating systems.

- This procedure can be automated cloning virtual machines.
  For example, a master images of a VM can be configured with all the parameters needed. As a new VM is needed you would create a clone of that master. This way, new VM with all the configurations needed can be easily deployed in a short time.
  http://www.storage-switzerland.com/Articles/Entries/2011/2/4_Using_Clones_To_Manage_VMware_Storage_Growth.html

Server downtime and loss of customer data:

A recent ransomware out-break (WannaCry attack-October 2017) infected many Microsoft servers (most of their servers are Microsoft Server 2012) across the data centre as there was no network segmentationand patching procedures. Due to this outbreak, customers data was lostas backups were also on the same network where the malware attacked. Numerousrestore procedures also failed as backups were found not to be working correctly on some servers.

Only one slow Internet connection:

Current DSL line not meeting the needs of the business. This also went down for 2hours recently as DAC's ISP (Four Ireland) experienced an outage and this impacted the Dublin area. As a direct result of this, one customer couldn't get data analysis resultsfrom DAC's server farm. Since this outage, the customer has now left DAC and have taken their business to a competitor.

Lack of manageability:

The IT manager needs to physically visit the data centre to carry out works if there is a problem. No remote access is currently available so the IT team need to walk between head office and the data centre to resolve issues. Also, issues are sometimes only discovered on visit to the data centre e.g. link continually flashing on a device but no monitoring in place.

What network management changes would you recommend making to help the IT manager and IT staff to manage equipment remotely securely? Is there any remote accessprotocols you would recommend DAC to avoid and why? What VLAN design would you recommend for the network? Provide a sample network addressing scheme with your answer.

To resolve this issue, a secure remote access can be configure. Two of the most common remote management protocols are telnet and SSH (Secure Shell). The key different between these 2 protocols is that SSH encrypts the data while Telnet sends data in plain text. So, SSH is a secure protocolo while telnet is not recommended on insecure networks.

From the information given, explain and provide a high level network design strategyindicating a best practice design moving forward DAC should implement based on access, distribution and core layers. Ensure to provide a labelled diagramof the proposed design.

Discuss TWO technologies in detail you would introduce to improve availability and performanceto help eliminate system downtime and bottlenecks. Use diagrams to help explain your technologies

- EtherChannel
- First Hop Redundancy Protocol (FHRP)
  - Hot Standby Routing Protocol (HSRP)

What is Network Functions Virtualization (NFV)? Provide a brief description of any ONE benefit introducing NFV into THM Ltd networks.

Network Function Virtualization, or NFV, is a way to reduce cost and accelerate service deployment for network operators by decoupling functions like a firewall or encryption from dedicated hardware and moving them to virtual servers.

This collapses multiple functions into a single physical server, reducing costs and minimizing truck rolls.

Introducing NFV into THM Ltd networks will benefit:

- Reduced space needed for network hardware
- Reduce network power consumption
- Reduced network maintenance costs
- Easier network upgrades
- Longer life cycles for network hardware
- Reduced maintenance and hardware costs

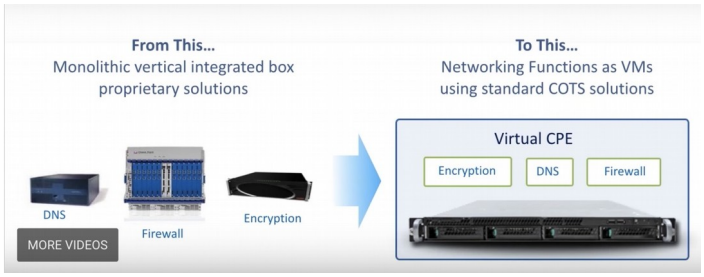

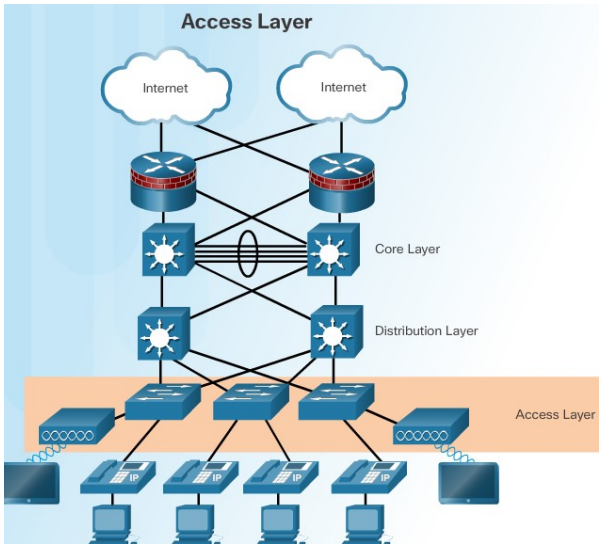

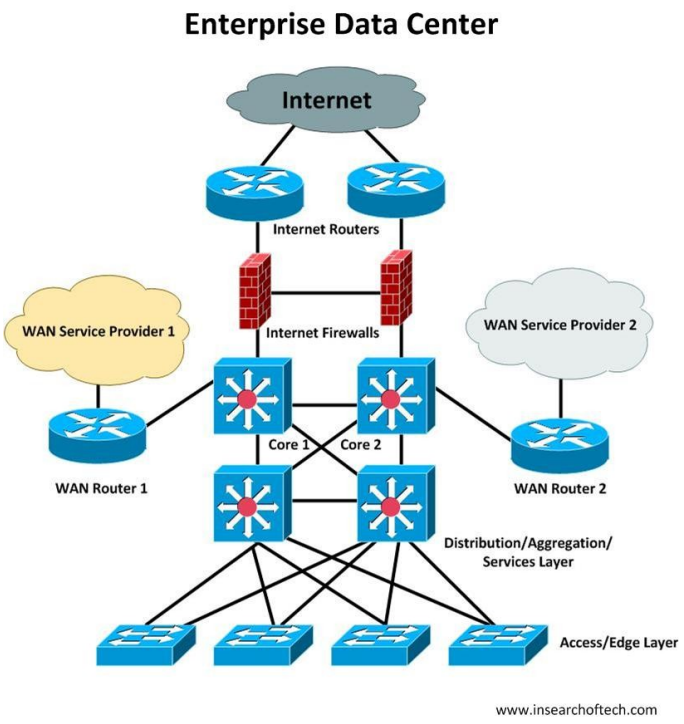

www.insearchoftech.com