CLOUD COMPUTING SYSTEMS

Lecture 11

Nuno Preguiça

(nuno.preguica_at_fct.unl.pt)

OUTLINE

- Networking 101
- Virtual networks in practice
- Data center overview

OUTLINE

- **Networking 101**
- Virtual networks in practice
- Data center overview

LAN INTERCONNECTS 101

What is a LAN? (typical cases only - Ethernet)

Local Area: distance among parties is limited (100m in general, a few Km is possible).

The ubiquitous wire protocol is Ethernet; the data link layer protocol is also Ethernet.

IP is the usual network protocol.

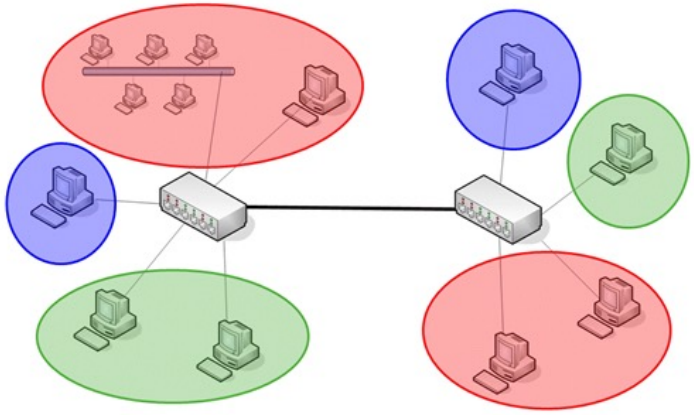

THE INTERCONNECT SPACE (INTRODUCING IB)

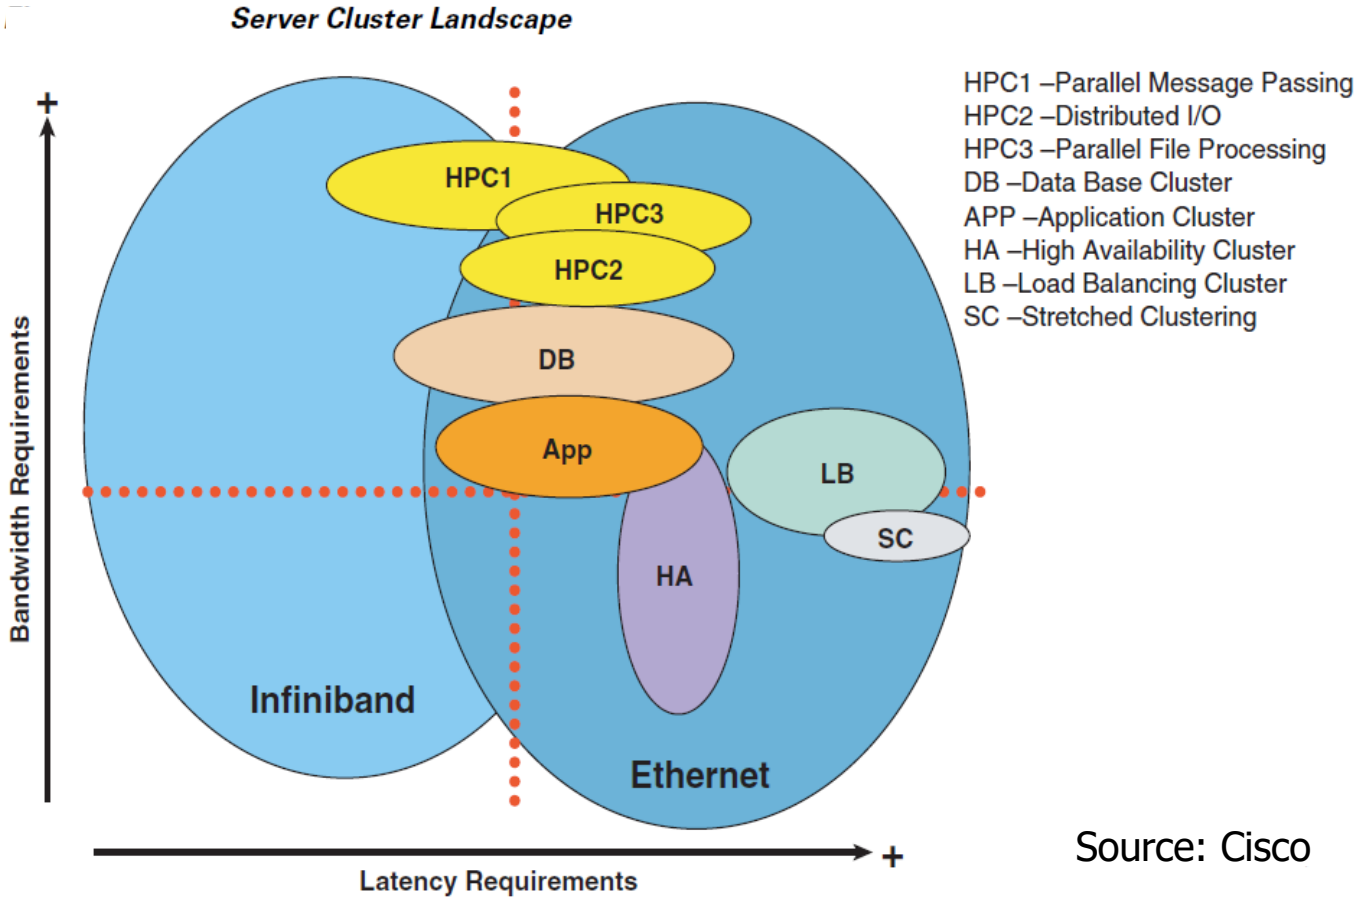

Source: Cisco

LAN INTERCONNECTS 101

How do parties communicate?

Directly connected parties exchange EM signals (could be light…). Each party has a unique address.

- Signals represent 0 & 1 bits
- Bits are packed into frames
- Frames carry IP packets
- IP packets carry TCP segments or UDP datagrams
- These usually carry "user data"

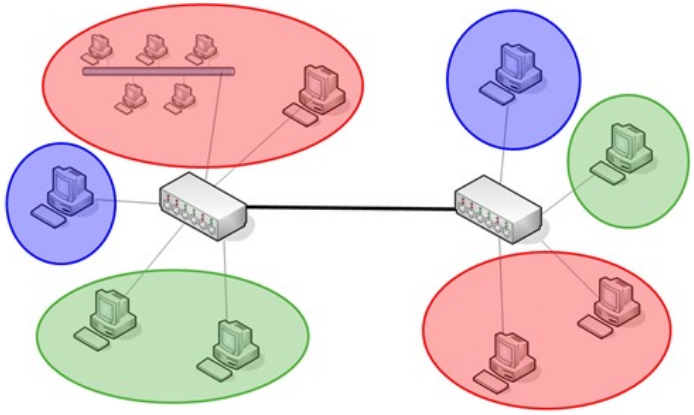

INTERNET PROTOCOL 101 (1)

How do parties communicate?

Hosts are grouped into distinct networks, each identified by an IP prefix.

In a network, each host has a different IP suffix and may communicate directly to other hosts in that network.

To communicate with hosts in a **different network**, a router device is required.

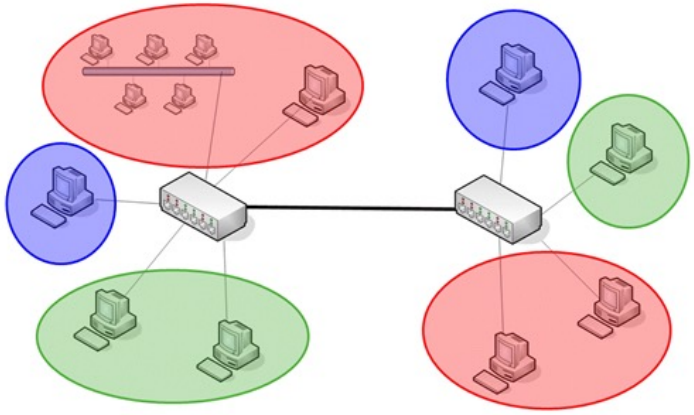

INTERNET PROTOCOL 101: IP (V4) ADDRESSES

An IPv4 address is a 32bit number, that can be split into two halves, prefix and suffix

- The usual notation is a dotted decimal followed by a slash and the prefix size, e.g., 192.168.1.23/24
- For prefixes sizes 8, 16 and 24 bits, extracting the network and host suffix is easy:

10.11.12.13/8	prefix: 10	suffix: 11.12.13
172.16.1.234/16	prefix: 172.16	suffix: 1.234
192.168.1.23/24	prefix: 192.168.1	suffix: 23

If the prefixes are equal, then the network is the same and direct communication is possible

INTERNET PROTOCOL 101 (3)

How do parties communicate?

Red net	prefix: 192.168.1
	suffixes: 1 - 8
Blue net	prefix: 192.168.2
	suffixes: 1 - 2
Green net	prefix: 172.16
	suffixes: 1.1, 1.2, 1.3

Although all servers are in the same LAN, and frames may flow unrestricted, hosts of "different colours" cannot communicate.

- However, if "someone" broadcasts a frame, every host will receive it.

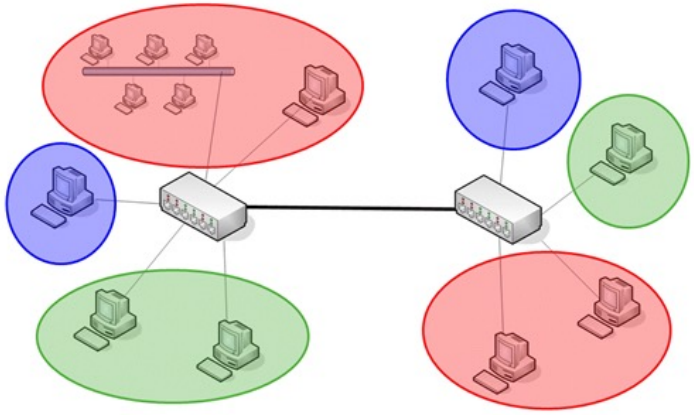

INTERNET PROTOCOL 101 (4)

How do parties communicate on a LAN?

- Hosts must communicate via IP addresses.
- But, in fact, communication "happens" at the frame level.

- How do hosts discover the destination's MAC address?

Address Resolution Protocol (ARP)

- When a host with address IP1 wants to communicate with a host with IP2, for the 1st time, it will:
  - Prepare an ARP packet with: its own IP, and the IP of the other host
  - Put the packet in a frame where the source address is its own MAC and the destination address is all ones, and broadcast the frame
  - Wait for a reply to arrive; the ARP-reply will carry the other host's MAC in the received frame "source MAC field"

ETHERNET LANS

Hub

Hub
Frames sent by CN1 to CN2 show up in all CNs…

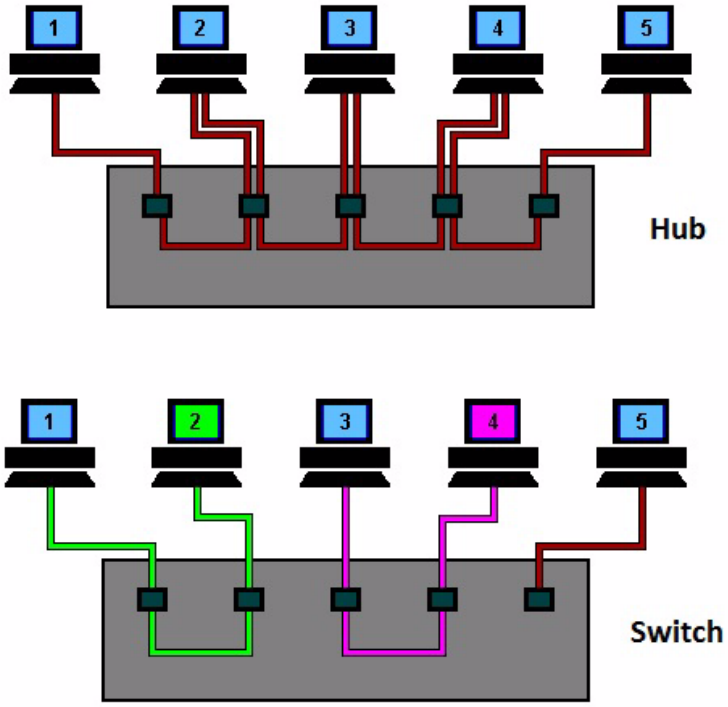

- Switch[1]

Frames sent by CN1 to CN2 only show up in CN2 – however the first frame sent by CN1 may be broad-casted to all…

[1] A switch is a.k.a. a learning bridge

ETHERNET LANS: THE SWITCH

The switch:

- Allows communications to flow among parties, and is a "passive" device – it does not initiate communication
- When a port receives a frame, the switch looks to the source and destination MACs, and
  - If this is the first frame it sees coming from that CN, the switch "memorizes" that the sender CN is reachable "trough that port";
  - If it does know what port must be used to reach the destination CN, it forwards the frame through that port;
  - If it does not know what port must be used to reach the destination CN, it broadcasts the frame to all (done! It does not have to wait for a reply)
- Step-by-step it builds a FDB (forwarding database) of ports and MACs reachable through those ports

INTERNET PROTOCOL 101 (5)

How do parties on two distinct IP networks, say N1 and N2, communicate?

- A router with a minimum of one connection for each network must exist; say, port X for N1 and Y for N2
- Each host must declare that the router (indicated by its IP) is the gateway to the other network(s)

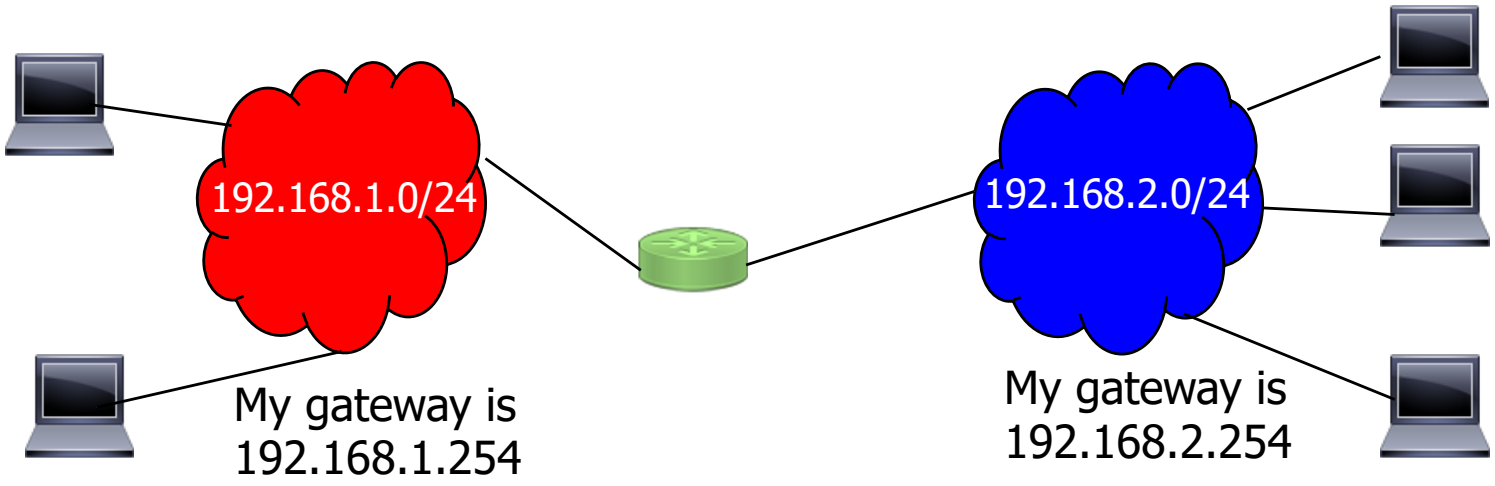

INTERNET PROTOCOL 101 (6)

How do parties on two distinct IP networks, say N1 and N2, communicate? (cont)

- When a host with address IP1 wants to communicate with a host with IP2, in a different network, it will:
  - Prepares an IP packet with data and source address IP1 and destination address IP2;
  - Sends the packet to the router (may require an ARP first);
  - The router receives on interface X the frame coming from IP1, extracts the IP packet from the frame and copies it to a new frame that it sends out via interface Y to the host IP2 (may require an ARP first).

Be careful with the terminology: switches forward frames, routers route packets!

ETHERNET LANS: VLAN

In large sites a single, undivided network, is a source of many problems:

Unsecure: traffic can be snooped, IPs changed and thus everybody may access everything.

Load: Broadcasts "steals" valuable bandwidth.

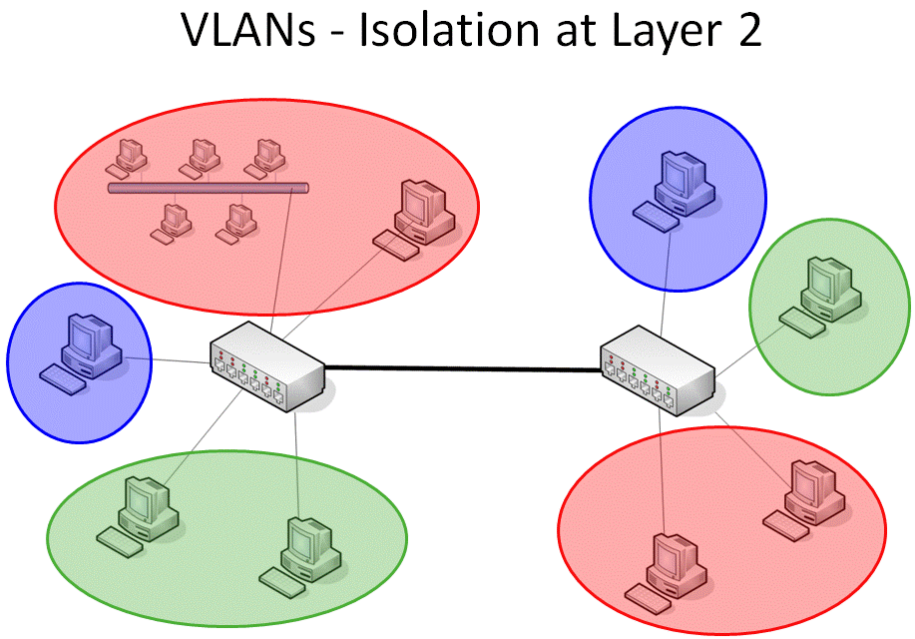

VIRTUAL LANS

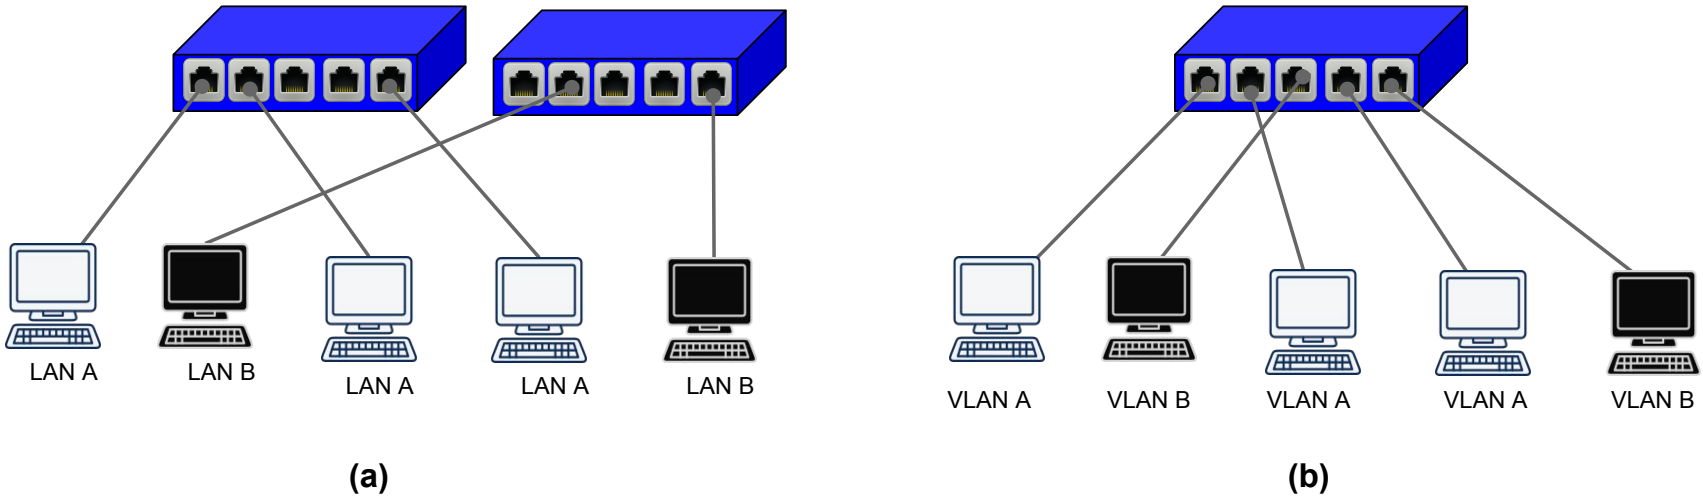

It would be possible to create different networks by connecting the computers to different switches.

Virtual networks partition the network into disjoint "areas".

Network switches separate the traffic of different virtual networks.

VIRTUAL LANS (2)

Virtual networks can encompass multiple networking devices.

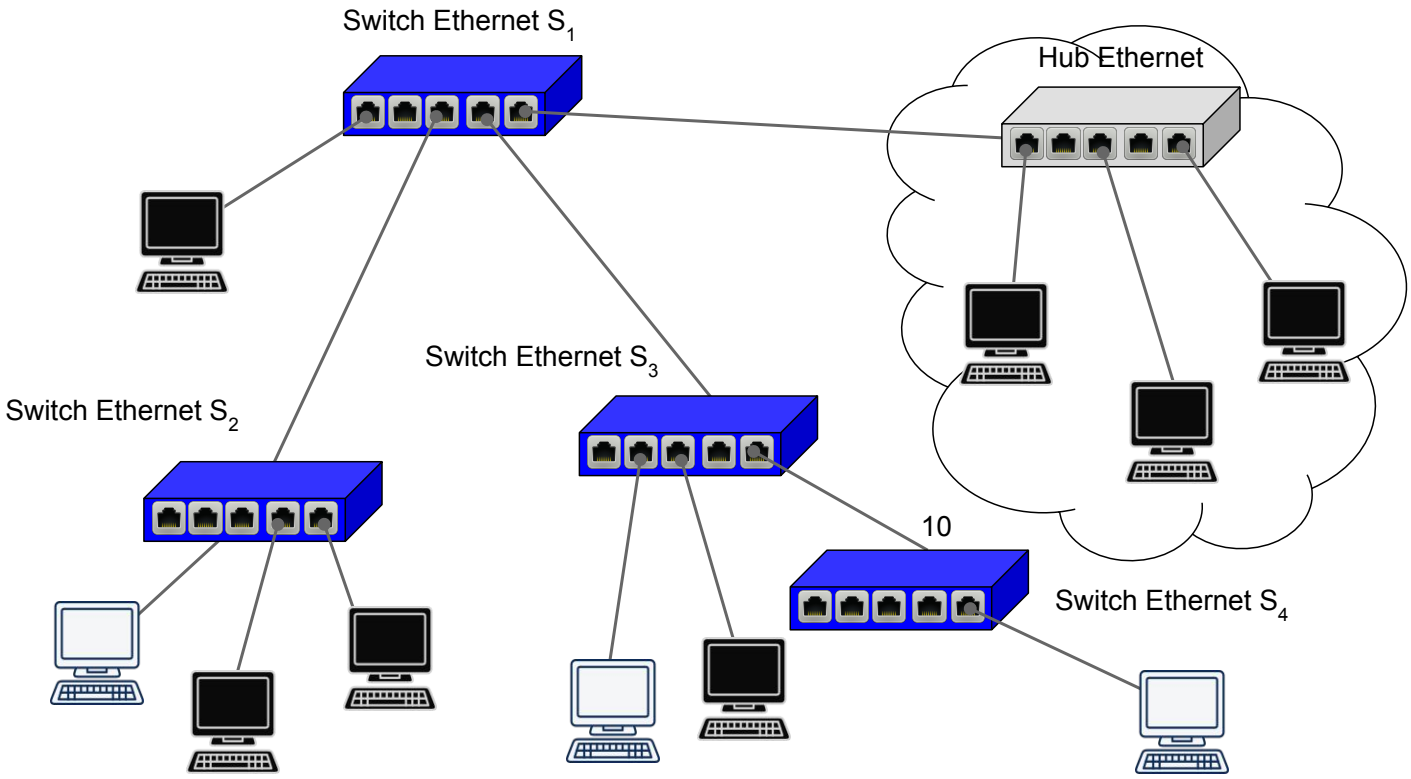

PROTOCOLS FOR VIRTUAL NETWORKS

Switches need configuration tables

- Saying which VLANs are accessible via which interfaces

Changing the Ethernet header

- Adding a field for a VLAN tag
- Implemented on the bridges/switches
- … but can still interoperate with old Ethernet cards

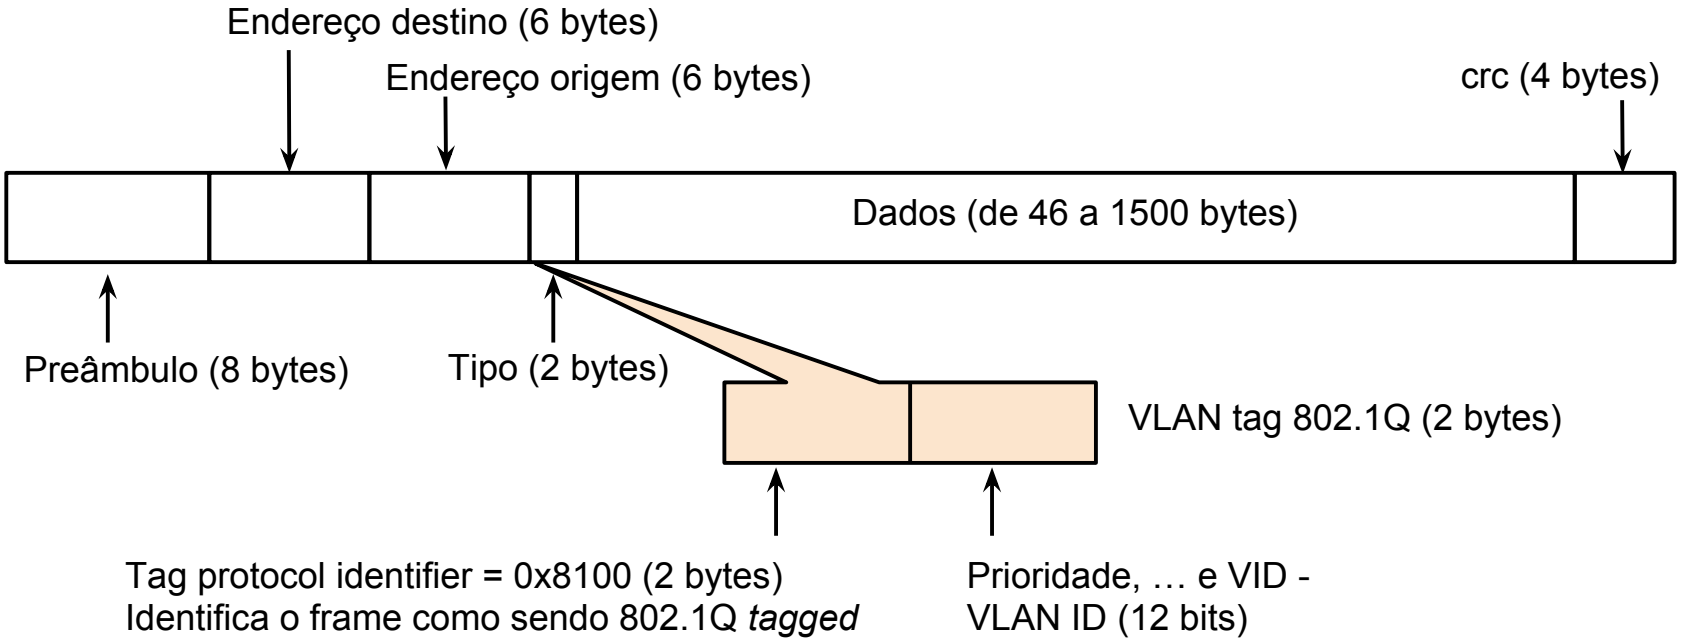

PROTOCOLS FOR VIRTUAL NETWORKS (CONT.)

The VLANs standard: 802.1Q

The Spanning Tree Protocol Standard: 802.1D

- Each VLAN has its own Spanning Tree

VLANs can be interconnected by routers

NETWORKING IN CLOUD PLATFORMS

Cloud platforms tend to internally use proprietary protocols.

Rationale: controlled environments and specific requirements are better supported by specific solutions (than by standards designed for general-purpose use).

For example, for supporting interconnectivity between Pods in Kubernetes, a number of networking solutions is supported:

https://kubernetes.io/docs/concepts/cluster-administration/networking/

[More on the "Computer Networks Architecture and Protocols" course]

OUTLINE

- Networking 101
- **Virtual networks in practice**
- Data center overview

VIRTUAL NETWORKS IN PRACTICE

Virtual networks are used for two main goals:

- Allow machines/services running in different locations to belong to the same network.
- Secure the communication between devices in a virtual network.

Azure Virtual Network

Azure Virtual Network (VNet) enables Azure resources, such as Azure Virtual Machines (VM), databases, etc., to securely communicate with each other, the internet, and on-premises networks.

VNet brings the benefits of Azure's infrastructure such as scale, availability, and isolation.

VNets are restricted to a single region.

COMMUNICATION BETWEEN AZURE RESOURCES

Azure resources communicate securely with each other in one of the following ways:

Through a virtual network: You can deploy VMs, and several other types of Azure resources to a virtual network, such as Azure App Service Environments, the Azure Kubernetes Service (AKS), and Azure Virtual Machine Scale Sets.

Through a virtual network service endpoint: Extend your virtual network private address space and the identity of your virtual network to Azure service resources, such as Azure Storage accounts and Azure SQL databases, over a direct connection. Service endpoints allow you to secure your critical Azure service resources to only a virtual network.

Communication between Azure resources (2)

Through VNet Peering: You can connect virtual networks to each other, enabling resources in either virtual network to communicate with each other, using virtual network peering. The virtual networks you connect can be in the same, or different, Azure regions.

OUTLINE

- Networking 101
- Virtual networks in practice
- **Data center overview**

The Data Centre

A safe location for: servers, storage and networking devices.

24x7 connectivity guarantees from the DC to the "outside world".

Power supply to the DC is enough to keep it operating in "any condition".

Environment parameters such as temperature, humidity and air quality are kept within specified bounds.

The Data Centre: localisation

Minimizing risks

- Ground floor or underground
- Natural risks: seismic, floods, fires, etc.
- Man-made risks: vibration, dust, EMI, etc.

Economical factors

- Real estate
- Energy, water
  - Power substation
  - Dissipation (water, atmosphere)
- Human
  - Qualified manpower, competitive wages

MARE NOSTRUM, BARCELONA

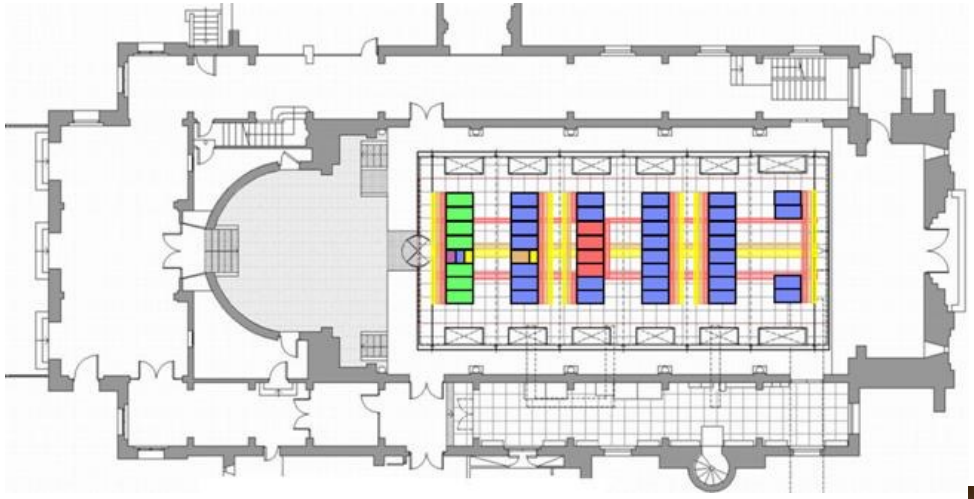

Servers +

Storage: 40 t,160 m^2

Air conditioned: 170 t

Power: 700 KW

Dissipation: 2M BTUs/h

Mare Nostrum:
Barcelona Supercomputing Center

The Data Centre: location

Counter example:

- DC in/near office buildings
  - heat dissipated in DC used to heat the building

The Data Centre: key factors

Space

- Floor, height
- Access paths

Power Supply

- Energy

Environment (HVAC: Heat, Ventilation, Air Conditioned)

- Extracting heat, injecting cold air (or…)
- New tendencies: water cooling!

THE DATA CENTRE: KEY FACTORS

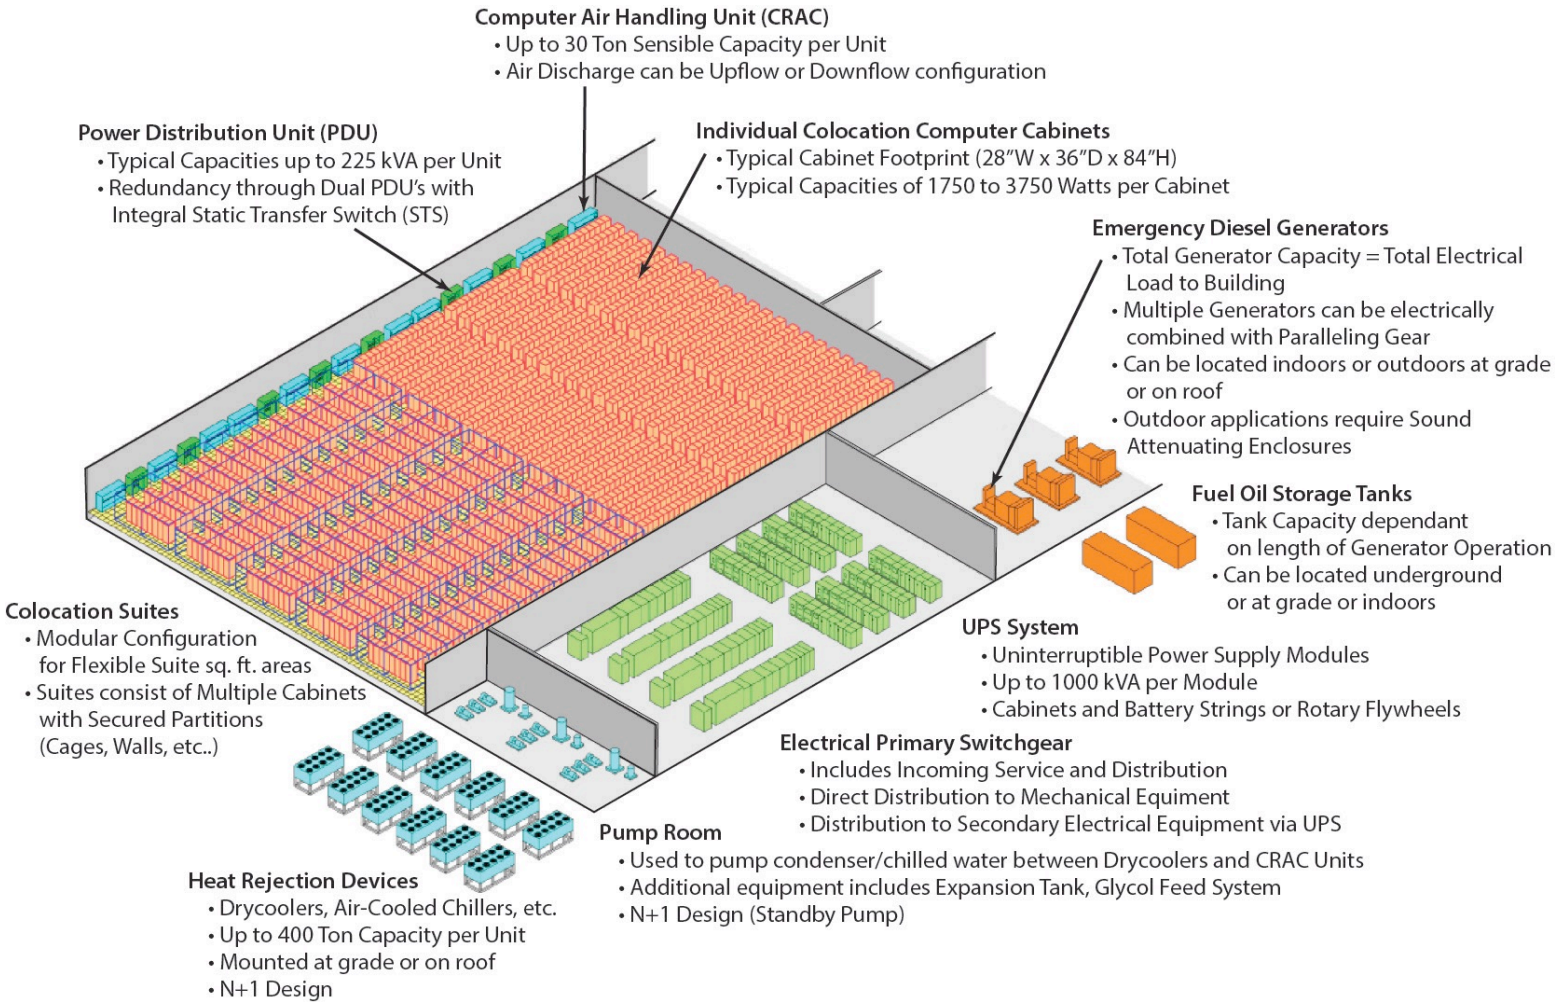

THE DATA CENTER: EQUIPMENT STORAGE (1A)

Racks

- Single "width": 19" equipment; ~28" total
- Depth: usually 30"
- Heights expressed in U units
  - 1U = 1.75"
  - Various; most used, 42 U

On the left:

Dell rack with: 1U servers (on top), 2 "Blade Center" racks of 7U each (in the middle), and 2U servers (lower part)

THE DATA CENTER: EQUIPMENT STORAGE (1B)

1U rack-mountable server

7U rack-mountable unit for blade-format servers

1U rack-mountable switch

THE DATA CENTER: EQUIPMENT STORAGE (2)

Standalone equipments:

Mainframe IBM Z800

Disk array ESS 800

The Data Center space: Floor (1)

Only 50% of the floor space should be used

- Regular distribution, but rows should not be too long
- Space for walking, specially on aisles
- Space to open rack doors to load/unload/repair equipment
- Space to separate hot/cold aisles

Beneath the raised floor

- Conveyors for power cable distribution; A/C ducts
- Conveyors for network cables
- Access zones

The Data Center space: Floor (2)

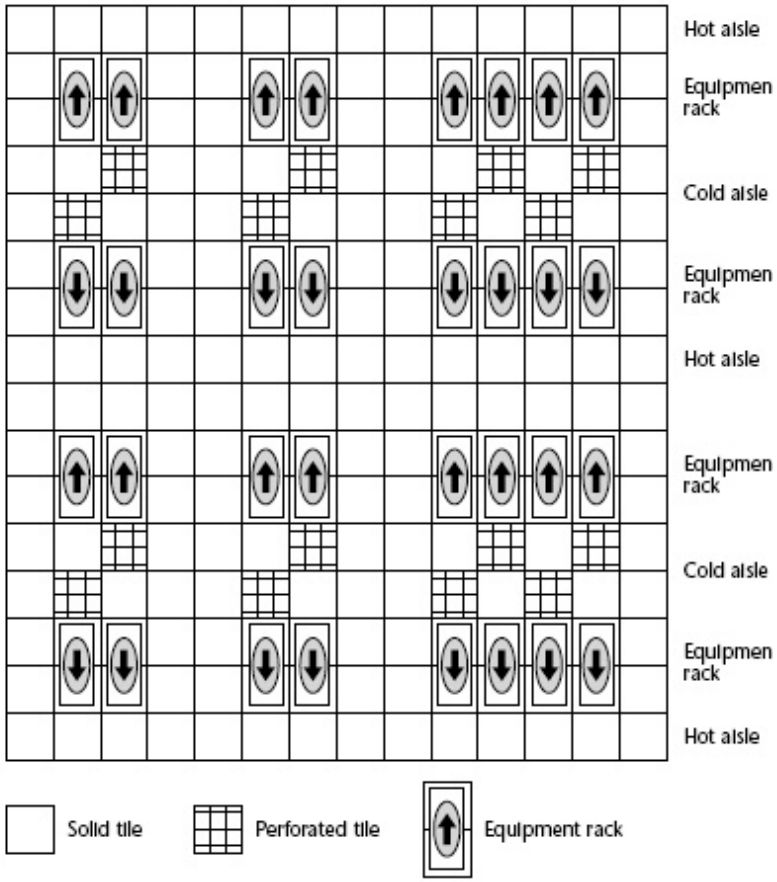

Weight/surface density

Solid and, eventually, (see HVAC slides) perforated tiles

Aligned racks (see HVAC slides)

- Cold air intake
- Hot air exhaust

The Data Center space: Ceiling

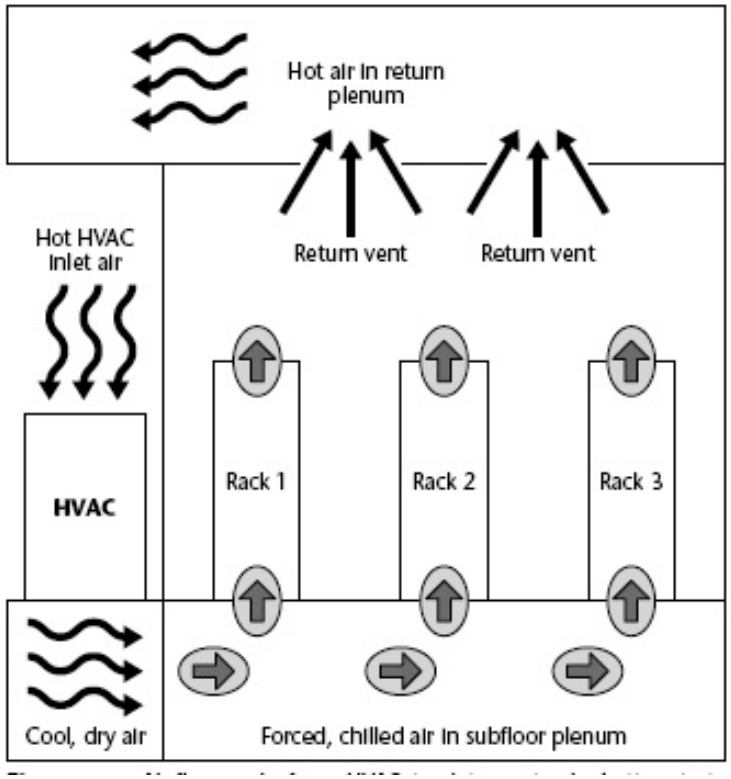

False/lowered ceiling

- Hot air extraction
- Sometimes used for cables (power or network) that drop down from the floor to the racks

Fire safety

- Detectors and sprinklers

New tendencies

- Racks with water cooled doors (ducts inside the door)

The Data Center space: Cooling

No heat extractors in the ceiling…

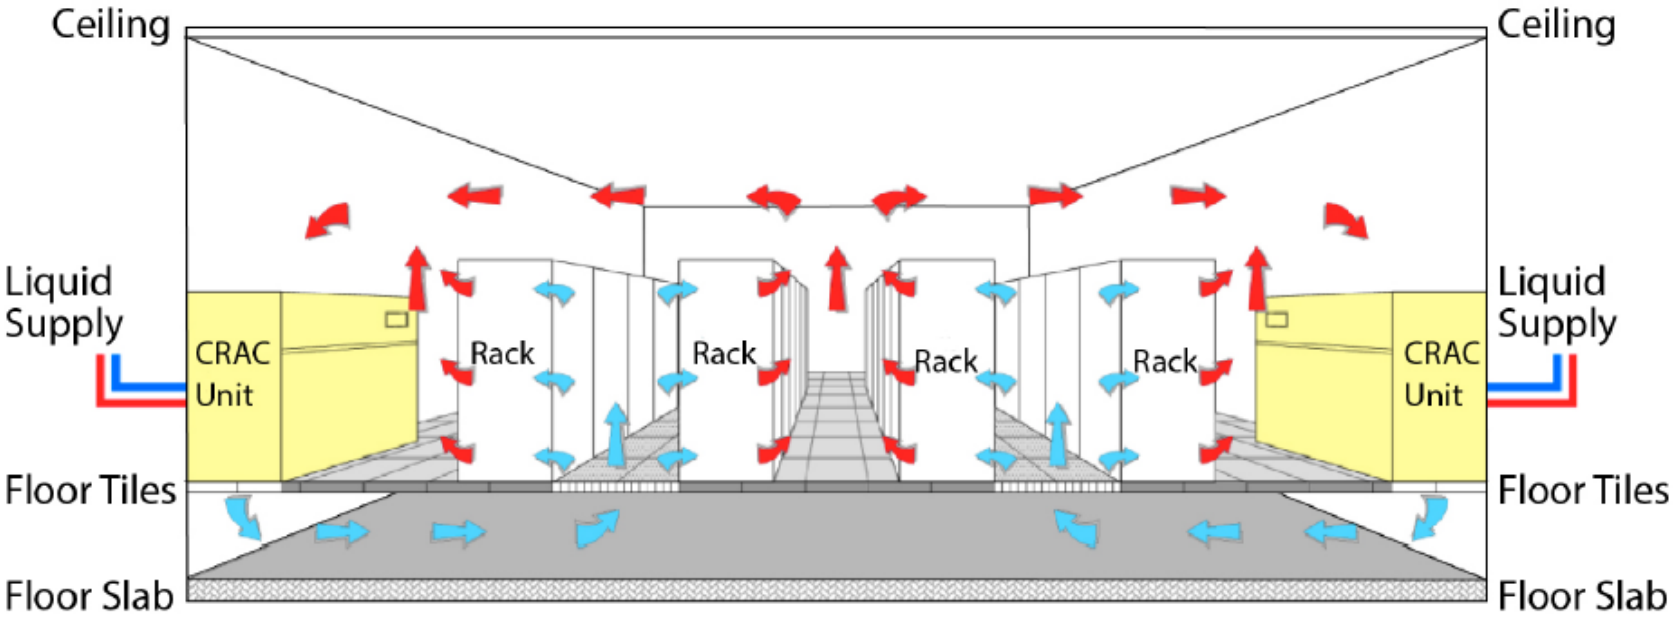

TO KNOW MORE

You will find all fundamentals (and much more) in

- In Portuguese: Legatheaux Martins, J. Fundamentos de Redes de Computadores (http://legatheaux.eu/index.php/cnfbook/)
  - 15.2 Comutação Ethernet (Ethernet switching)
  - 15.4 Virtual Local Area networks (VLANs)
  - 17. Interligação de Redes – Protocolo IP (IP Internetworking)

- In English: Introduction to Storage Area Networks (www.redbooks.ibm.com/redbooks/pdfs/sg245470.pdf)

Azure virtual networks

https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview

ACKNOWLEDGMENTS

Some slides based on a previous version by Paulo Lopes and Vitor Duarte.