“Cloud & Datacenter Networking” course

- 3 CFU
- Laurea Magistrale Ingegneria Informatica @ Università degli Studi di Napoli Federico II
  - 2° year 2° semester
- Course topics:
  - Datacenter architecture
  - Datacenter networking technologies
  - Cloud Computing networking aspects
    - In Amazon AWS, OpenStack, etc…
  - Network Virtualization
  - Emerging networking paradigms (Software Defined Networking, Network Function Virtualization, …)

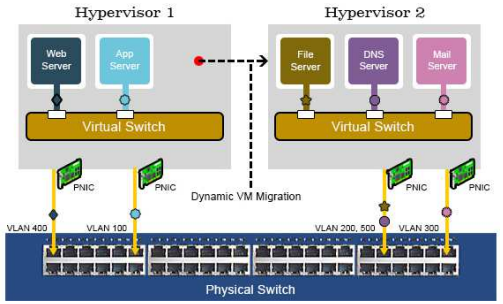

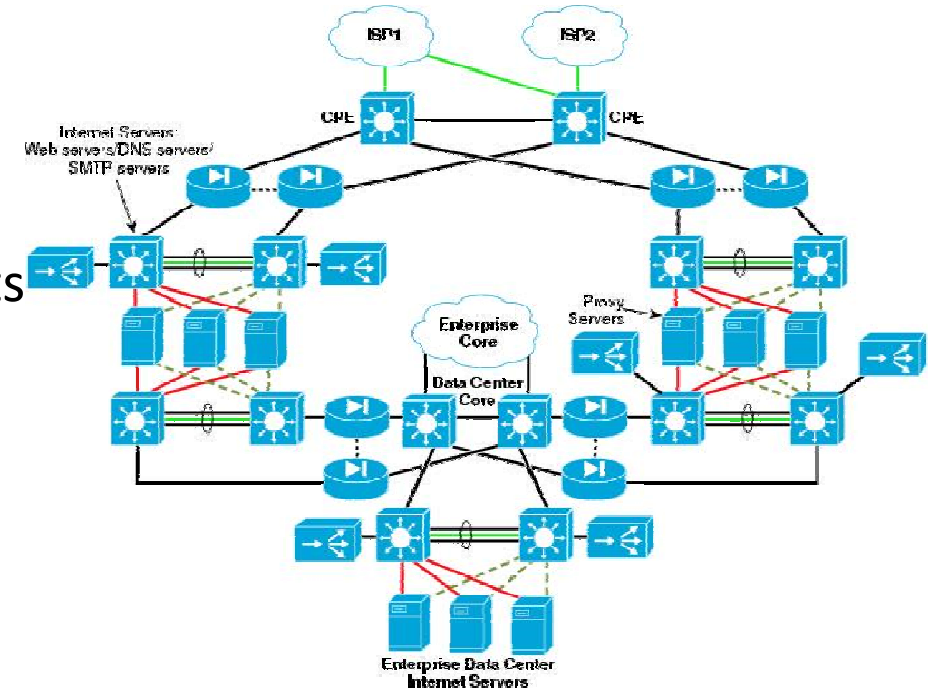

Datacenter engineering

- In a datacenter multiple infrastructures coexist:
  - IT: Computers, storage systems, switches, routers
  - Cabling
  - Servers cooling and air conditioning
  - Power supply systems
  - Physical security systems
- Design (and building) of a datacenter requires a number of engineering professionals with diverse competencies and expertise

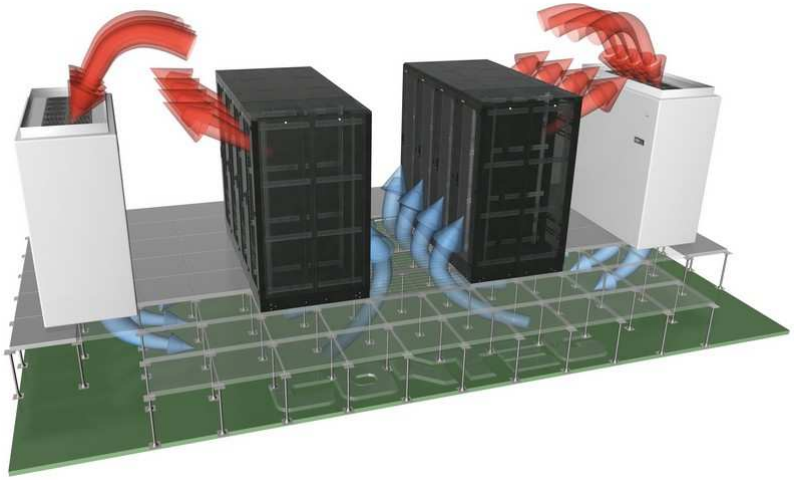

Datacenters: highly complex infrastructures

- A datacenter simply <u>cannot</u> be badly engineered
  - Service continuity requirements
    - High availability: 99,999% (five nines) $\rightarrow$ at most 5.26 minutes of outage in a whole year
  - Elasticity
    - It must be possible to reconfigure the infrastructure to adapt to a higher demand of resurces
  - Energy Efficiency
    - Power Usage Effectiveness (PUE)
    - In traditional datacenters non-IT plants (including air conditioning, uninterruptible power supply, emergency power systems, etc.) may consume as much energy as (or even more than) the IT equipments (servers, switches, …)

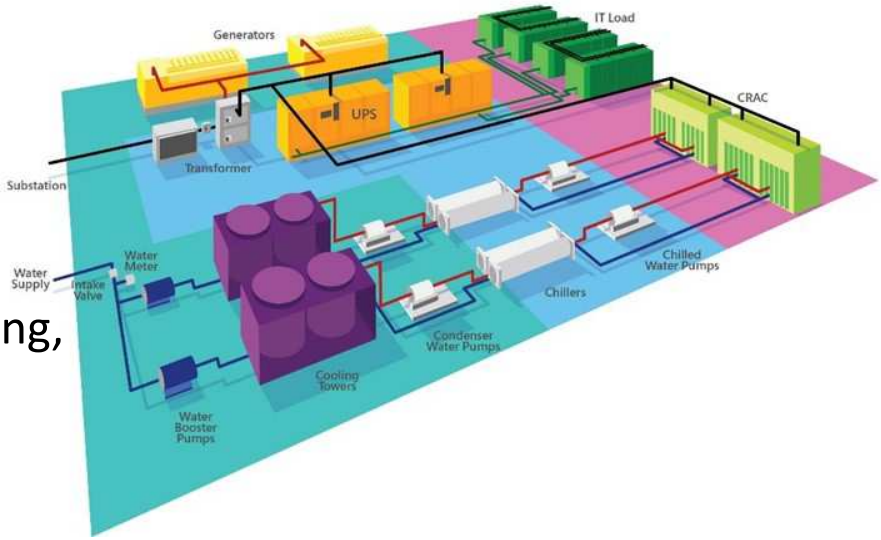

Datacenter networking challenges (1)

- Highly interactive and data-intensive applications
  - A single user interaction produce as a consequence a high number of interactions among server-side components (eg. many db queries)
    - Max response time for each server component estimated 10ms
    - The majority of the traffic remains within the datacenter
    - Traditional networking protocols exhibit problems in a DC environment
      - multipath, TCP incast, …

Datacenter networking challenges (2)

- How is it possible to make a LAN network of 10000+ servers?
  - A single big switch is unfeasible ….
  - Traditional WAN solutions are not adequate
  - A switch hierarchy!
  - Main goal:
    make “non-blocking” the interconnection system
  - A*gility*, i.e. the possibility of moving (*migrating*) the computational load (eg. a VM) from one server to any other server in the DC with no performance penalties for the application
  - Other goals:
    - avoid bottlenecks,
    - avoid loops,
    - exploit the existence of multiple paths between servers