

Cloud and Datacenter Networking

Università degli Studi di Napoli Federico II

Dipartimento di Ingegneria Elettrica e delle Tecnologie dell'Informazione DIETI

Laurea Magistrale in Ingegneria Informatica

Prof. Roberto Canonico

Datacenters: architecture and organization



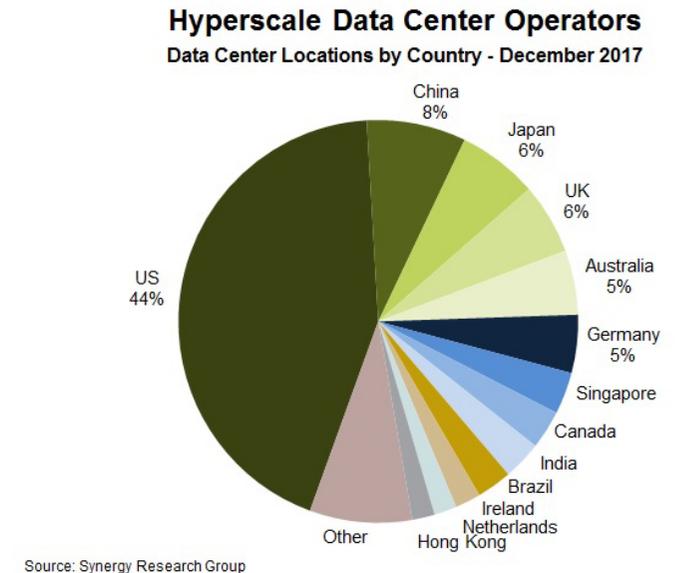
- ▶ **Datacenter: functional organization**
- ▶ **Datacenter: IT infrastructure and other plants**
- ▶ **Server racks organization**
- ▶ **The TIA-942 standard**
- ▶ **Cabling**
- ▶ **Datacenter cooling plant and CRAC units**
- ▶ **UPS plant**
- ▶ **Datacenter energy efficiency and PUE**
- ▶ **Datacenter tiers**
- ▶ **Modular datacenters**
- ▶ **Some datacenter examples**

- ▶ A datacenter is an engineering infrastructure designed and built to host a large-size computing system as well as all the associated auxiliary plants, such as telecommunication systems, power, air conditioning and cooling, etc.
- ▶ A datacenter (DC) comprises:
 - ▶ IT devices: servers, network devices, data storage devices
 - ▶ A cabling system
 - ▶ Uninterruptible Power Supply and Emergency Power Systems
 - ▶ Computer Room Air Conditioning (CRAC) and cooling systems
 - ▶ An infrastructure for physical security management: access control, fire detection and suppression, etc.
 - ▶ A Network Operations Center (NOC)

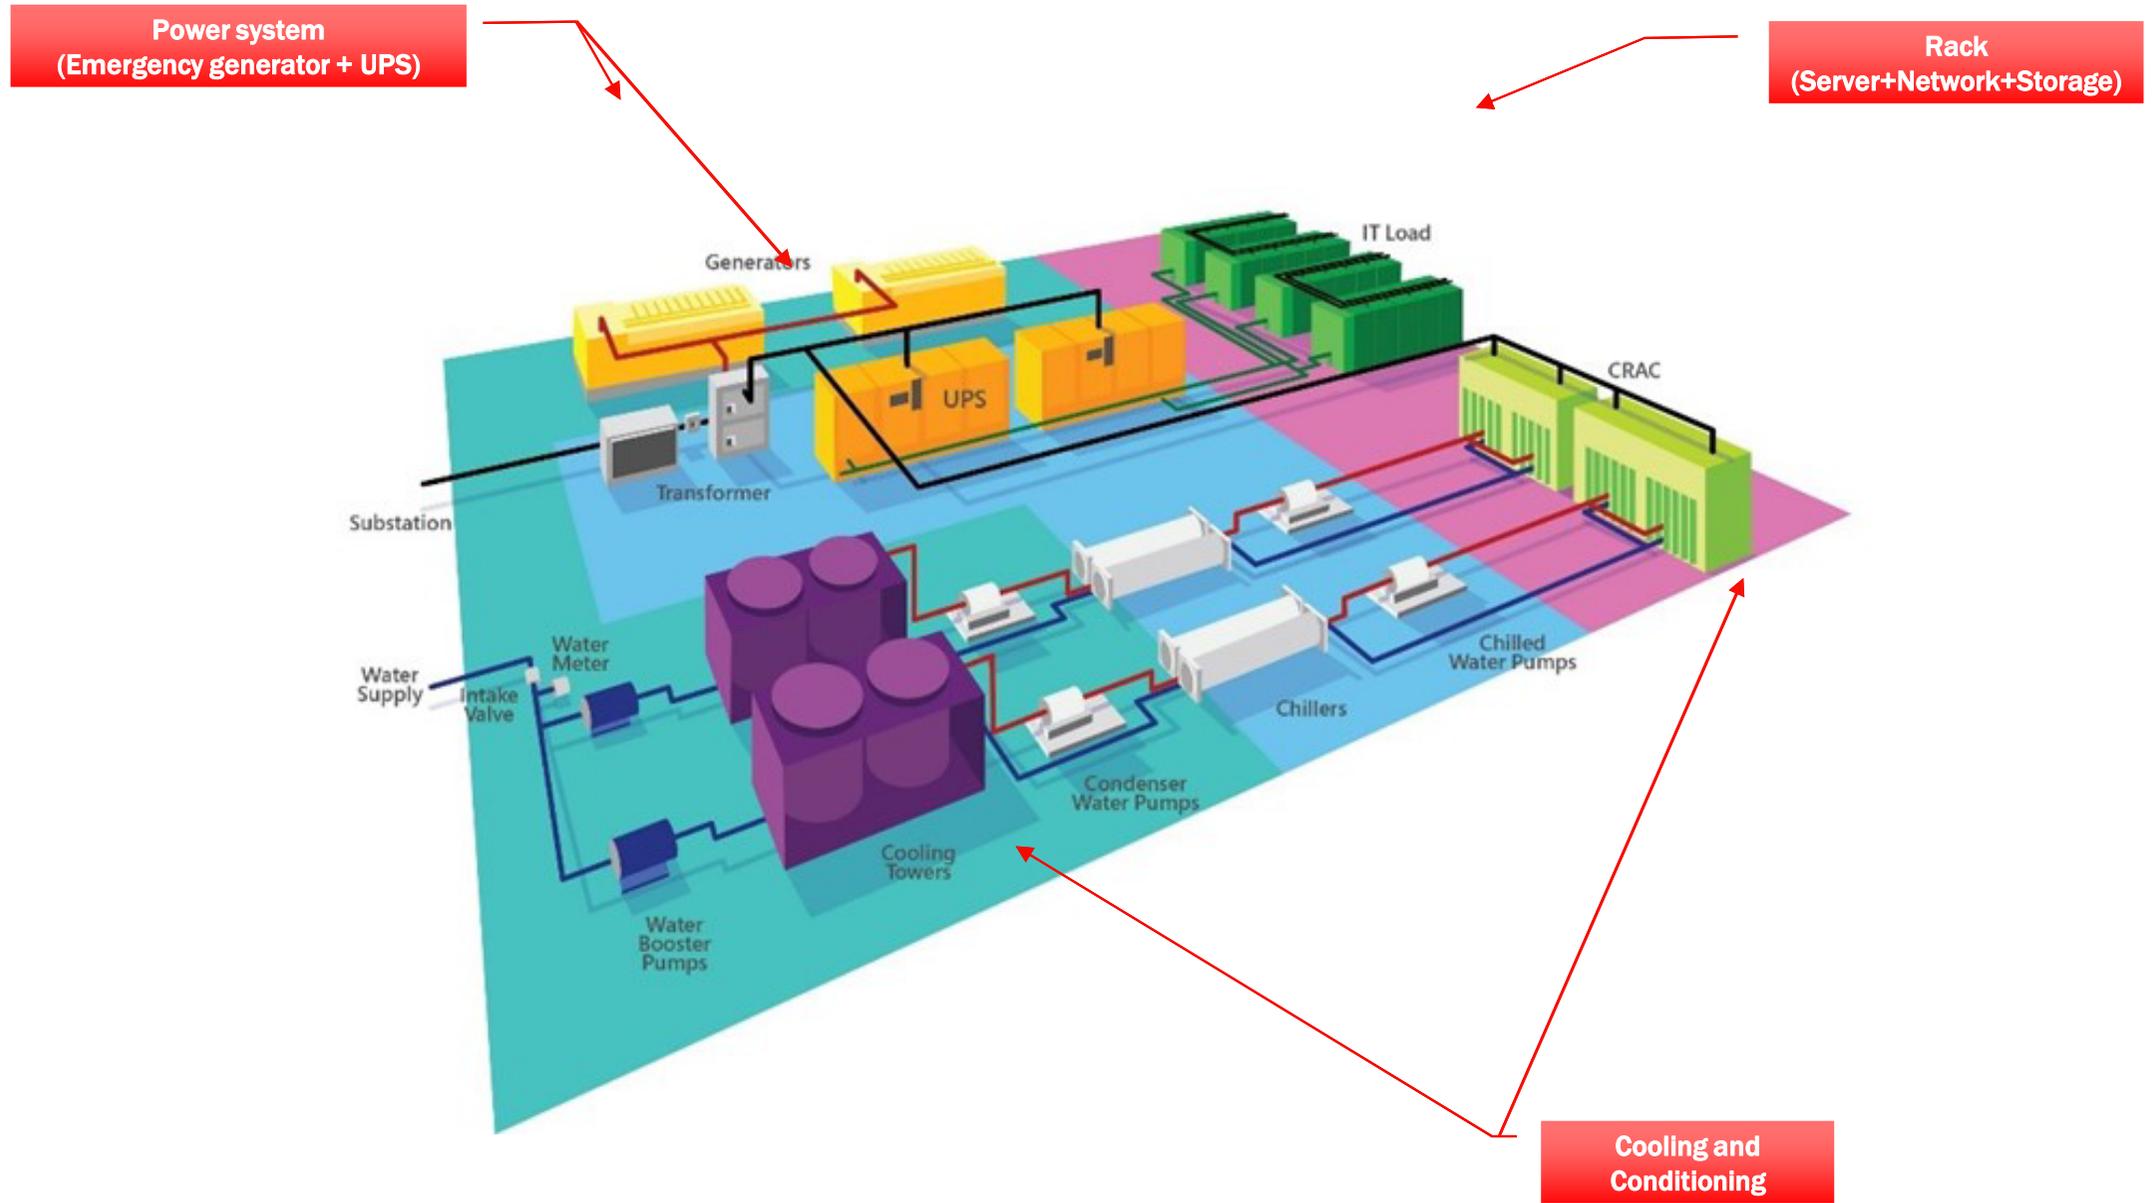
Large scale datacenters



- ▶ Number of large scale datacenters worldwide:
 - ▶ > 300 at the end of 2016
 - ▶ > 390 at the end of 2017
- ▶ Amazon, Apple, Microsoft, IBM and Google own at least 45 datacenters each worldwide
- ▶ Companies like Apple, Twitter, Salesforce, Facebook, eBay, LinkedIn, Yahoo have their DCs mainly in US
- ▶ Tencent and Baidu have large scale DCs in China
- ▶ Hyperscale datacenters:
 - ▶ *at least 5000 servers and 10,000 square feet of available space, but often much larger (IDC,2016)*
- ▶ Location of hyperscale datacenters:
 - ▶ US (44% of the total)
 - ▶ China (8%)
 - ▶ Japan and UK (6%)
 - ▶ Australia and Germany (5%)



Datacenter: IT infrastructure and auxiliary plants

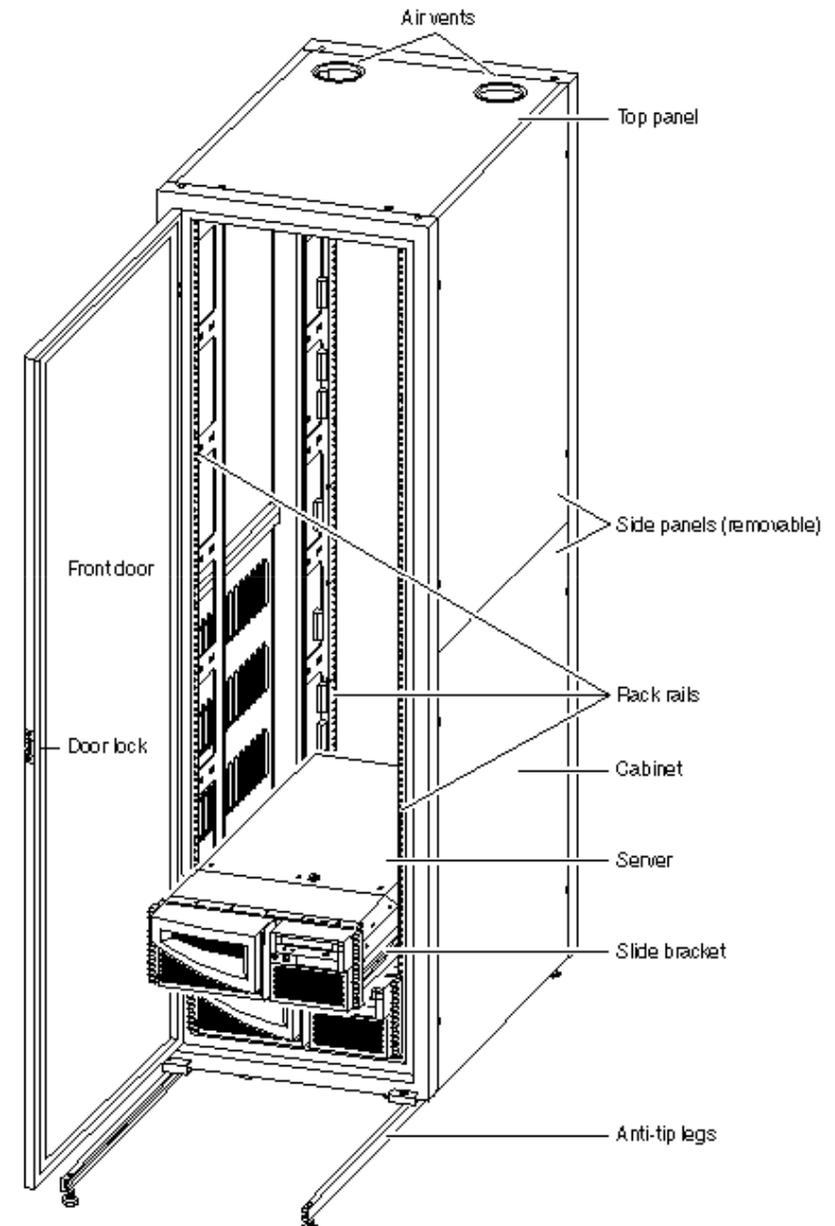


Datacenter: rack and cabinets



- ▶ A cabinet is the enclosure that includes the *rack* hosting servers, other IT devices and network switches
- ▶ Rack and cabinet size specified by EIA 310D standard
- ▶ Vertical space within a cabinet is measured in *modular units* (U), also referred to as *rack units* (RU)
- ▶ Typical rack height: 24U/25U and 42U

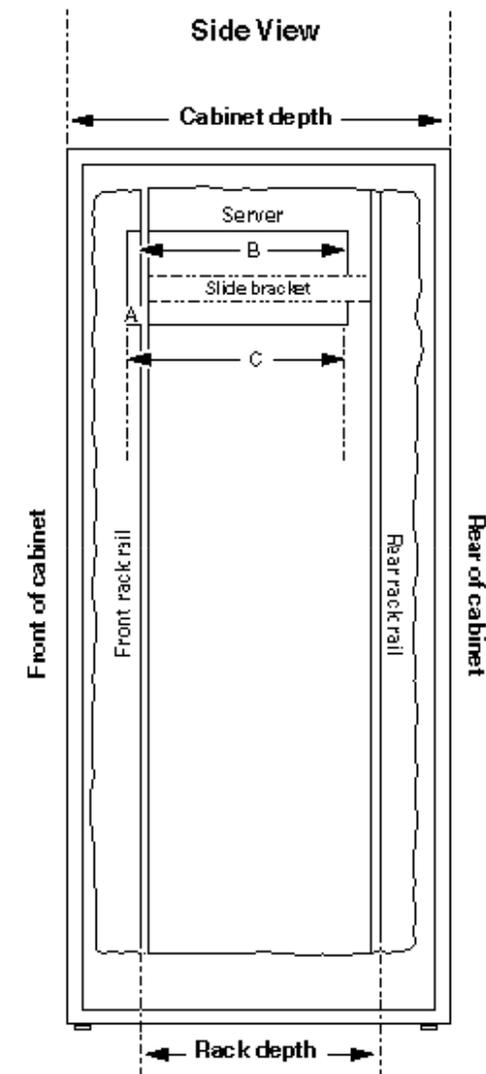
1 rack unit = 1.75 inches = 4.45 cm



Datacenter: rack and cabinet size

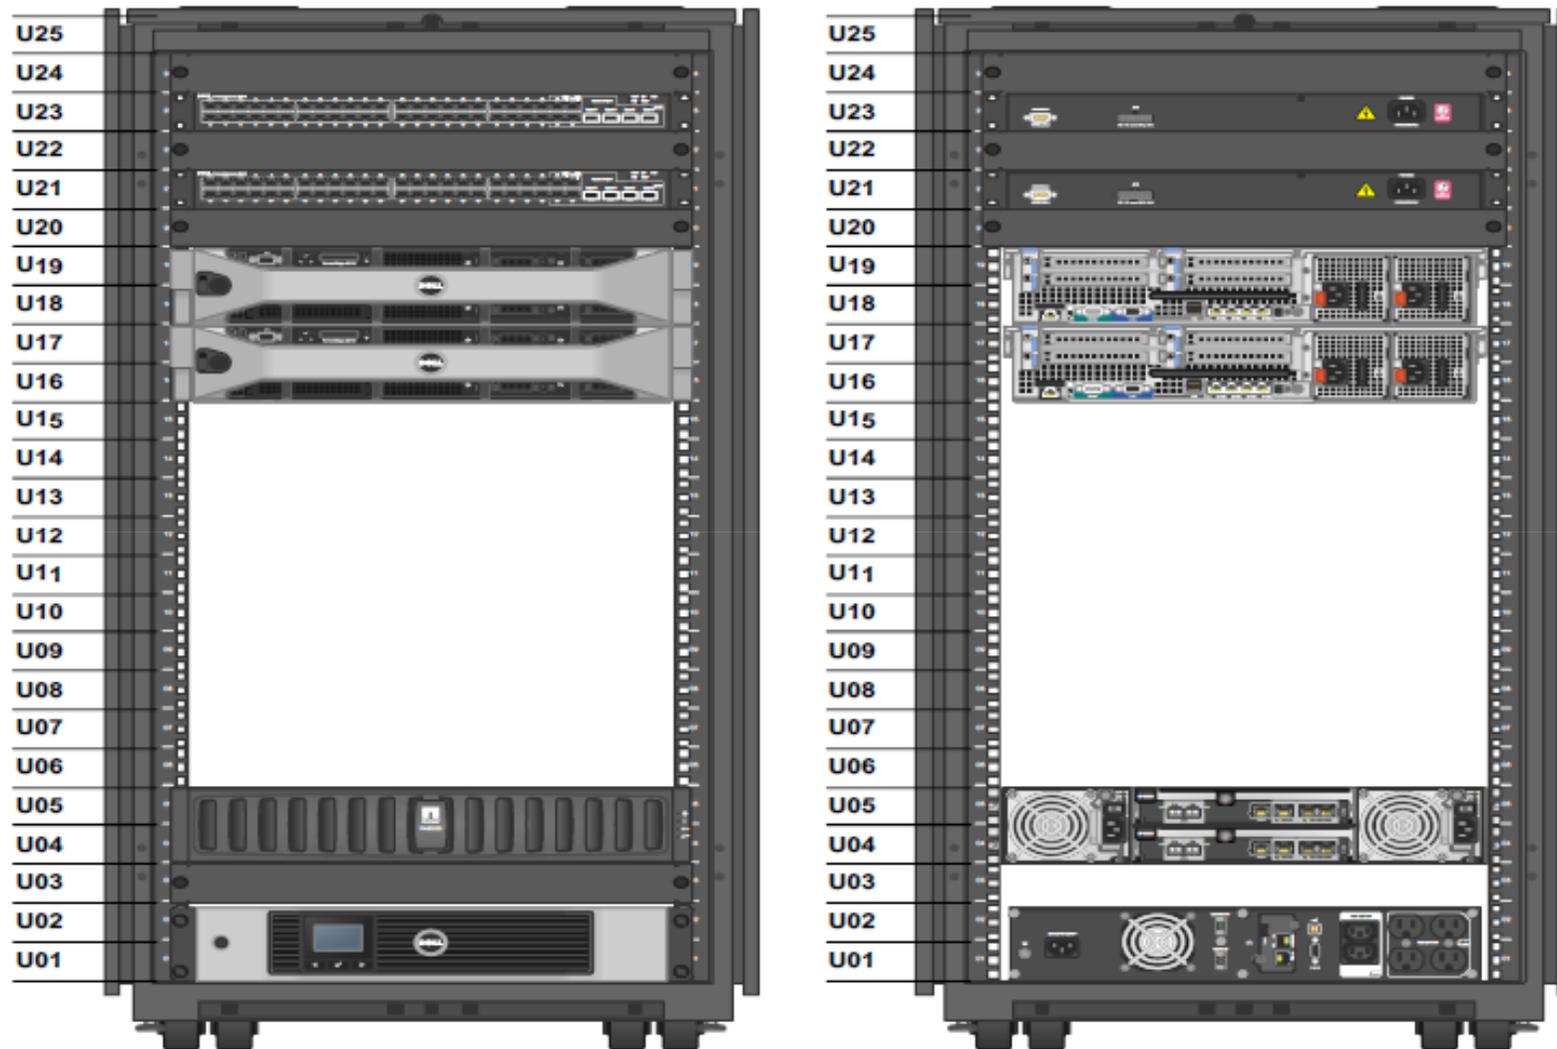


- ▶ A cabinet includes
 - ▶ the internal rack
 - ▶ sidewalls
 - ▶ power strips
 - ▶ front door
- ▶ Rack standard width
 - ▶ 19 inches (48.26 cm) → most common
 - ▶ 23 inches (58.42 cm)
 - ▶ 24 inches (60.96 cm)
 - ▶ 30 inches (76.2 cm)
- ▶ Cabinet depth
 - ▶ 34 inches (86.36 cm)
for servers length up to 28 inches (71.12 cm)
 - ▶ 39 inches (99.06 cm)
for servers length up to 33 inches (83.82 cm)



A = Depth of the server in front of the front rack rail
B = Depth of server from the forward-most part of the front rack rail to the rear-most part of the server
C = Total depth of server

25U rack: front-rear view



- ▶ With respect to traditional tower servers, rack-mountable servers bring the following advantages:
 - ▶ **Manageability:** racks allow easier management of servers and other IT equipments in particular wrt cabling
 - ▶ **Consolidation:** racks allow reduction of volumes, leading to a higher density of computational power → but make servers cooling more challenging
 - ▶ **Security:** organizing servers in racks lead also to higher physical security; server cabinets equipped with lockable doors allow a higher protection against unauthorized access to physical equipments, data theft and sabotages
 - ▶ **Expandability:** racks make easier to expand systems by adding new equipments
 - ▶ **Modularity:** the adoption of cabling standards, the fact the racks size are standardized, etc, make it easier to integrate into the datacenter equipments from different vendors as well extend in modular way the amount of resources in the datacenter

Rack mountable servers: sample configurations



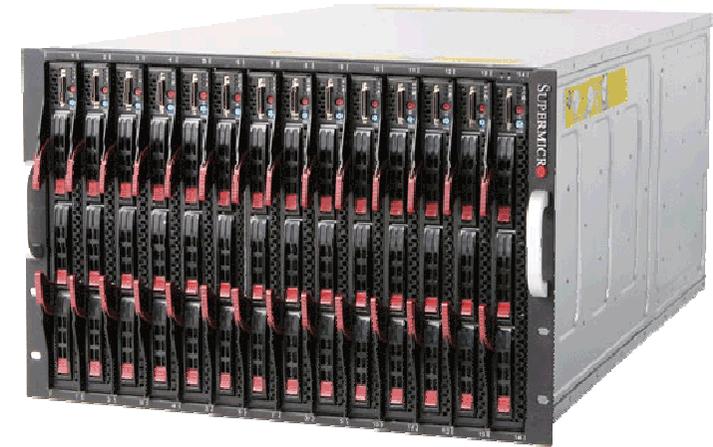
	IBM System x3650 M4 7915	HP ProLiant DL380 Gen9 E5-2650v3	DELL PowerEdge R630
CPU	Intel Xeon E5-2680 2,70GHz (64-bit)	Intel Xeon E5-2650 v3 2,30 GHz (64-bit)	Intel Xeon E5-2640 v3 2,60 GHz (64-bit)
Number of cores	8	10	8
Number of CPUs	max 2	max 2	max 2
RAM slots	24 DDR3 SDRAM	24 DDR4 SDRAM	24 DDR4 SDRAM
RAM max	24*4=96 GB UDIMM or 24*32=768 GB LRDIMM	24*16=384 GB RDIMM or 24*32=768 GB LRDIMM	24*16=384 GB RDIMM or 24*32=768 GB LRDIMM
Storage controller	Serial Attached SCSI (SAS)	Serial Attached SCSI (SAS)	Serial Attached SCSI (SAS)
Storage	6 x 3,5" SATA slots or 8/16 x 2,5" SAS slots	8 x 2,5" (SFF) slots	8 x 2,5" (SFF) slots
Network Interface Cards	4x1Gb Ethernet Intel I350AM4	4x1Gb HP Embedded + 2x10Gb-T Flexible LOM	4x1Gb Ethernet Broadcom 5720
Remote control system	Integrated Management IMM2	HP iLO	Intelligent Platform Management Interface (IPMI)
Size	2U	2U	1U



Blade servers



- ▶ A blade server is a stripped-down server computer with a modular design optimized to minimize the use of physical space and energy
- ▶ Blade servers have many components removed to save space and minimize power, while still having all the functional components to be considered a computer
- ▶ Unlike a rack-mountable server, a blade server needs a *blade enclosure*, which can hold multiple blade servers, providing services such as power, cooling, networking, various interconnects and management
- ▶ Blades and enclosure together form a *blade system*
- ▶ Blade solutions may achieve a higher density (no. of computers / rack) with respect to rack mountable servers



A sample blade server system



▶ Chassis

- ▶ Bays to host blade compute nodes and network devices (I/O modules)
- ▶ Redundant power supplies for all the blades



Lenovo Flex System Enterprise Chassis

▶ A blade compute node



Lenovo Flex System x240 M5

Processor(s):	2 Intel Xeon processors E5-2600 v3 series, up to 36 cores per node
Memory:	24 DDR4 LP, 1.5 TB maximum with 64 GB LRDIMM
Internal Storage:	2 x hot-swap 2.5-inch (SAS/SATA/SSD/PCIe)
Network Connectivity:	2 x 10 GbE (default configuration) 2-port, 4-port and 6-port x 10 GbE adapters (optional) 2-port and 4-port x 8/16 Gb Fibre Channel adapters (optional) 2 port x QDR/FDR InfiniBand adapters (optional)

- ▶ A **white box server** is a data center computer that is not manufactured by a well-known brand name vendor
- ▶ The term “white box” indicates that the equipment is generic or brand-agnostic
- ▶ White box servers can be found in large data centers run by giant Internet companies like Google or Facebook
- ▶ White box servers are usually purchased in bulk quantities from suppliers called **original design manufacturers (ODMs)**
- ▶ An ODM typically builds their servers with commercial off-the-shelf (COTS) components that can be assembled in slightly different ways or upgraded to provide the customer with a degree of system customization
- ▶ Examples of ODMs producing white box servers: Quanta and Hon Hai's Foxconn
- ▶ The white box approach has also been used for networking devices
 - ▶ E.g. Pica8 white box switches
- ▶ An increasing number of white box products are built using designs promoted by the **Open Compute Project**

- ▶ The Open Compute Project is an organization that shares designs of data center products among companies, including:
 - ▶ Facebook, Apple, Microsoft, Rackspace
 - ▶ Cisco, Juniper Networks
 - ▶ Goldman Sachs, Fidelity, Bank of America
- ▶ Components of the Open Compute Project include:
 - ▶ **Open racks** have the same outside width (600 mm) and depth as standard 19-inch racks, but are designed to mount wider chassis with a 537 mm width ($\cong 21''$)
 - ▶ This allows more equipment to fit in the same volume and improves air flow
 - ▶ Compute chassis sizes are defined in multiples of an OpenU, which is 48 mm, slightly larger than the typical Rack Unit
 - ▶ **Server compute nodes** (based on Intel or AMD processors)
 - ▶ **Open network switch**
 - ▶ **Open Vault storage building blocks** offering high disk densities
 - ▶ 277 VAC power distribution

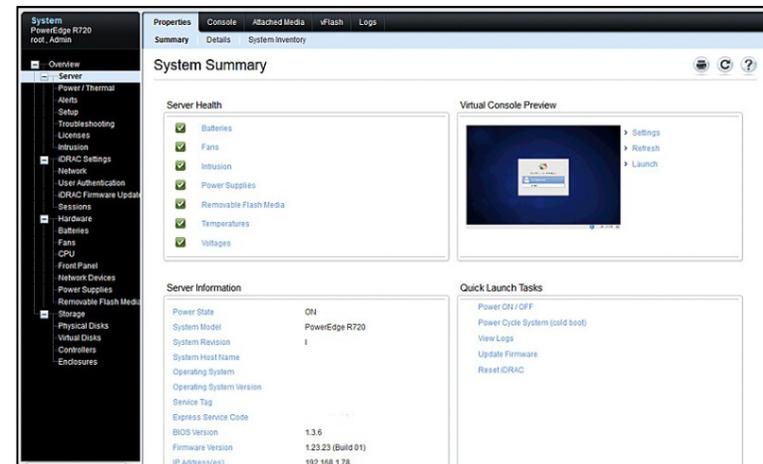
Rack: switch KVM



- ▶ To access server consoles, a rack may include a KVM console (keyboard, video, mouse) connected to servers' I/O ports (VGA video, USB or PS2 for keyboard and mouse) through a KVM switch
- ▶ A KVM switch allows to share the same console among all the servers in a rack



KVM console in a rack



Web-based remote management interface

- ▶ In modern datacenters, servers remotely managed by system administrators from the DC's NOC through special remote management interfaces
 - ▶ iLO, IPMI, etc.
- ▶ Through remote management interfaces it is possible a complete control over servers (including power status and console) from a (web-based) management GUI

Datacenter: floor and cable trays



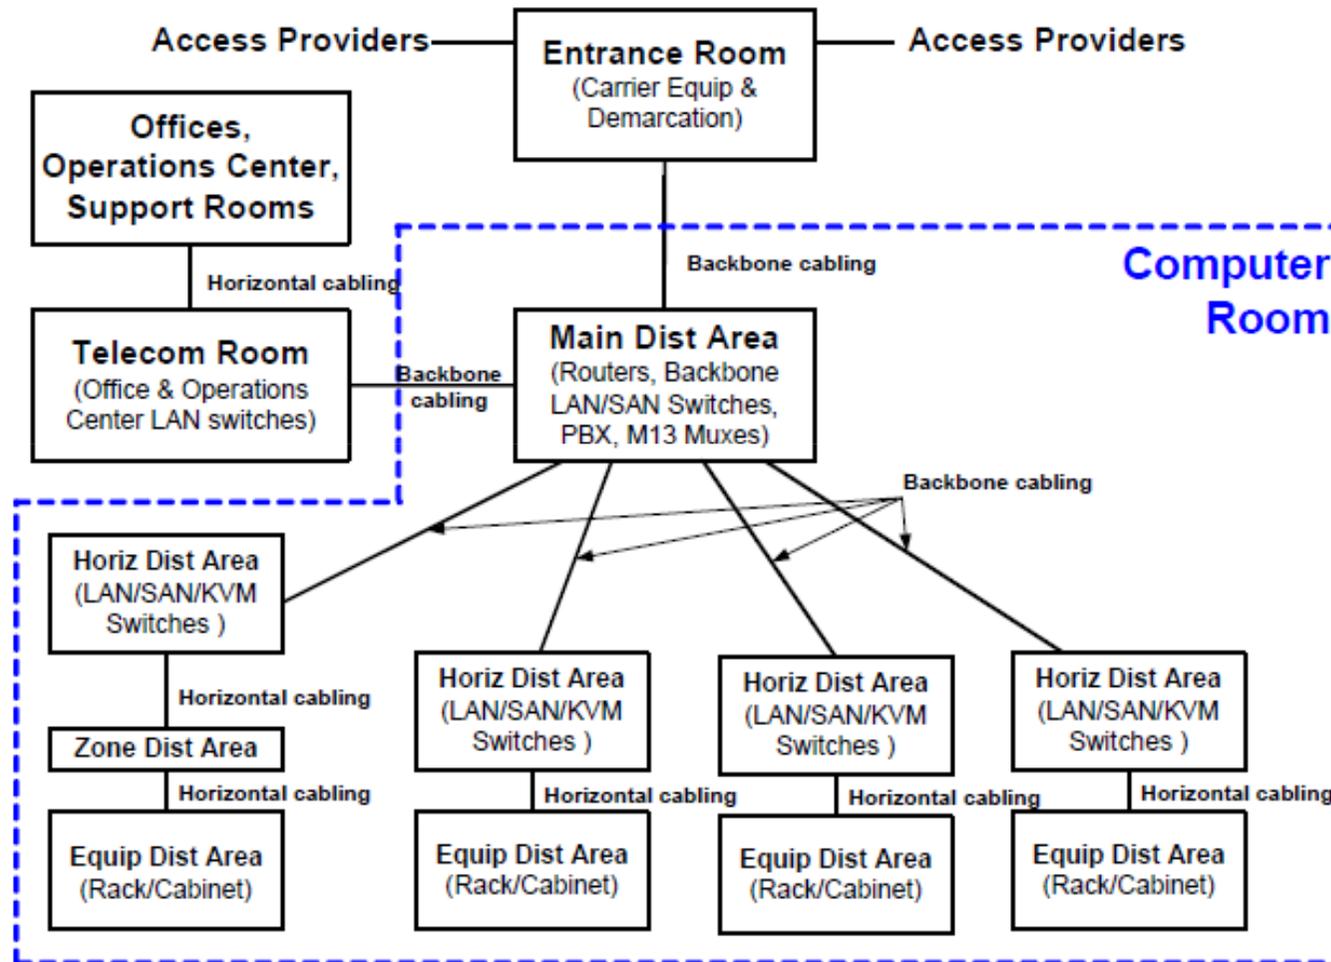
Overhead
Cable Tray

Floor Tile

Perforated
Floor Tile

- ▶ ***Telecommunications Infrastructure Standard for Data Centers***
- ▶ **Approved by ANSI and EIA in april 2005**
- ▶ **It defines the requirements and rules to follow when a datacenter is designed & built**
- ▶ **Provides guidelines for DC organization in so-called *work areas***
 - ▶ **Entrance Room (ER)**
 - ▶ **Operations Center (OC), Offices and Support Rooms**
 - ▶ **Telecommunications Room (TR)**
 - ▶ **Computer Room (CR)**
 - ▶ **Main Distribution Area (MDA) → Main Cross-connect (MC)**
 - ▶ **Horizontal Distribution Area (HDA) → Horizontal Cross-connect (HC)**
 - ▶ **Zone Distribution Area (ZDA) → Zone Outlet (ZO) or Consolidation Point (CP)**
 - ▶ **Equipment Distribution Area (EDA) → Equipment cabinets and racks**

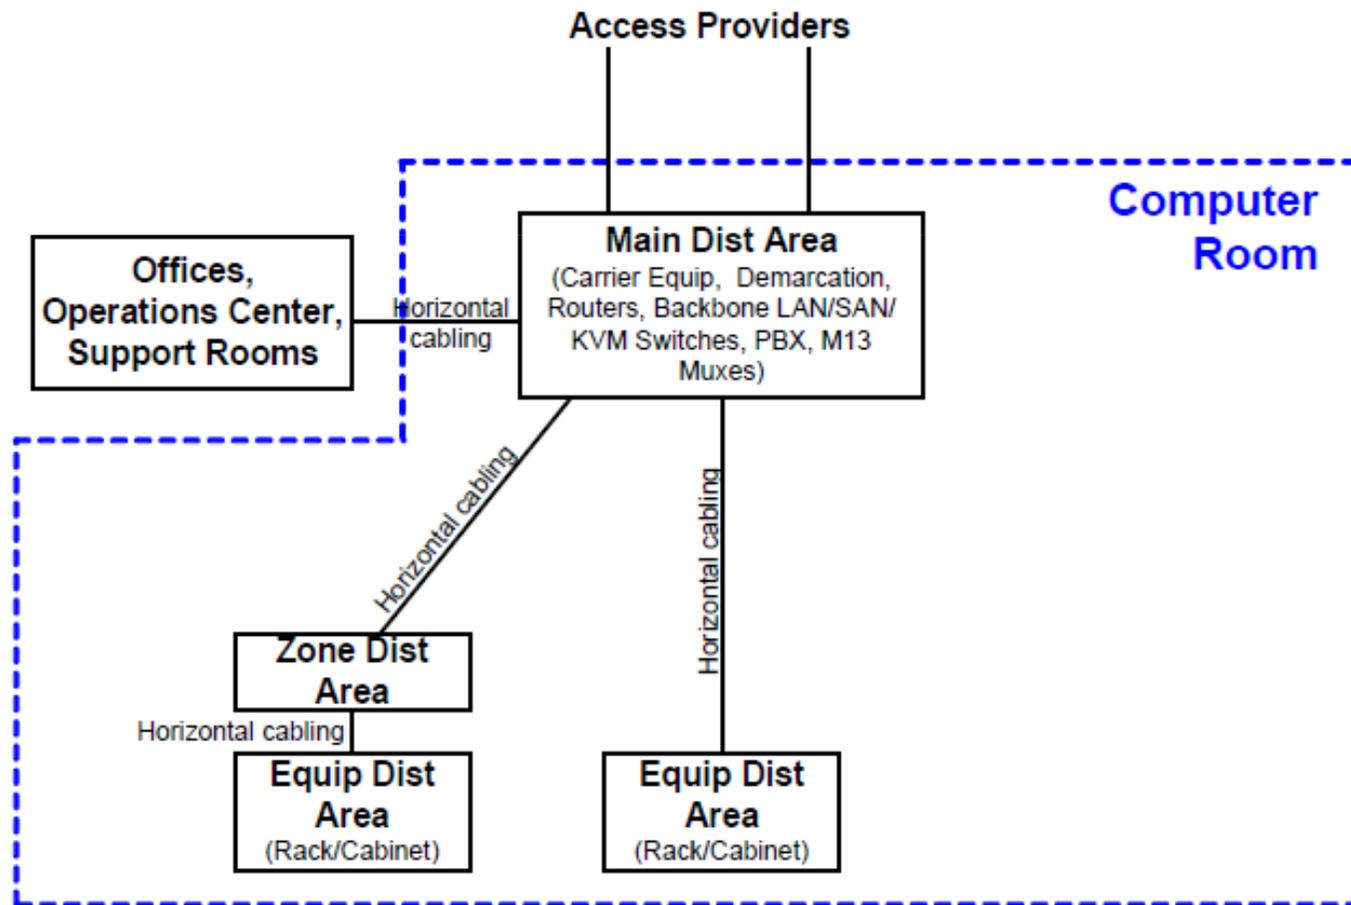
TIA-942: basic datacenter organization



TIA-942: reduced datacenter organization



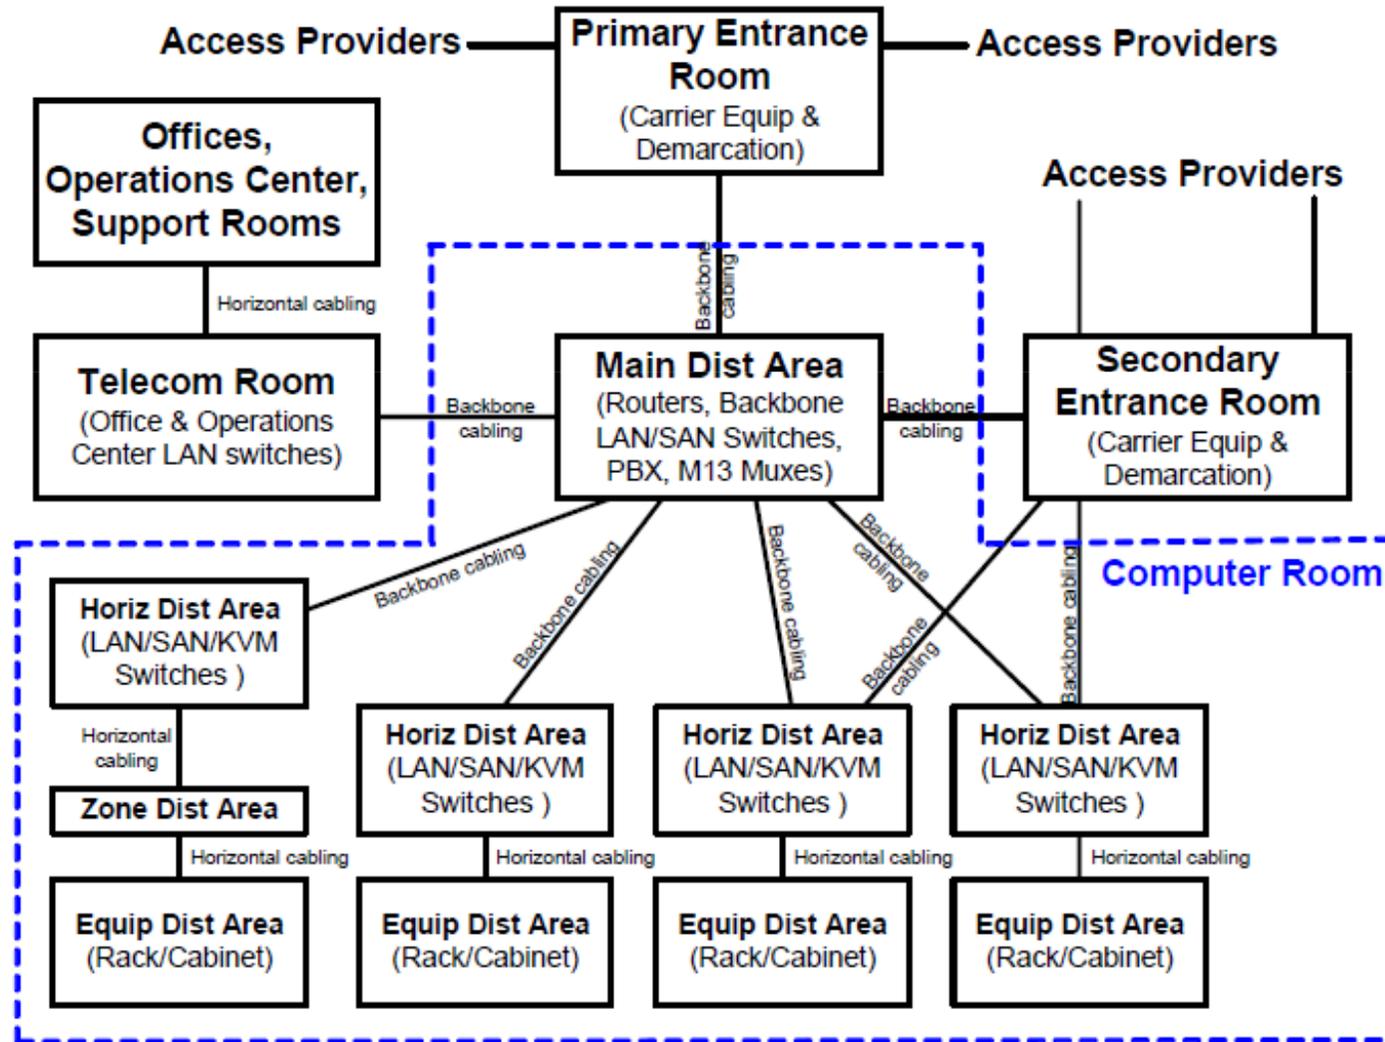
- ▶ In small scale datacenters:
 - ▶ Main Cross-connect and Horizontal Cross-connect can be consolidated in a single Main Distribution Area, possibly as small as a single cabinet or rack
 - ▶ Telecommunications Room for cabling to the Support Areas and the Entrance Room may also be consolidated into the Main Distribution Area



TIA-942: large datacenter organization



- ▶ Multiple Entrance Rooms
- ▶ All other parts are replicated for redundancy

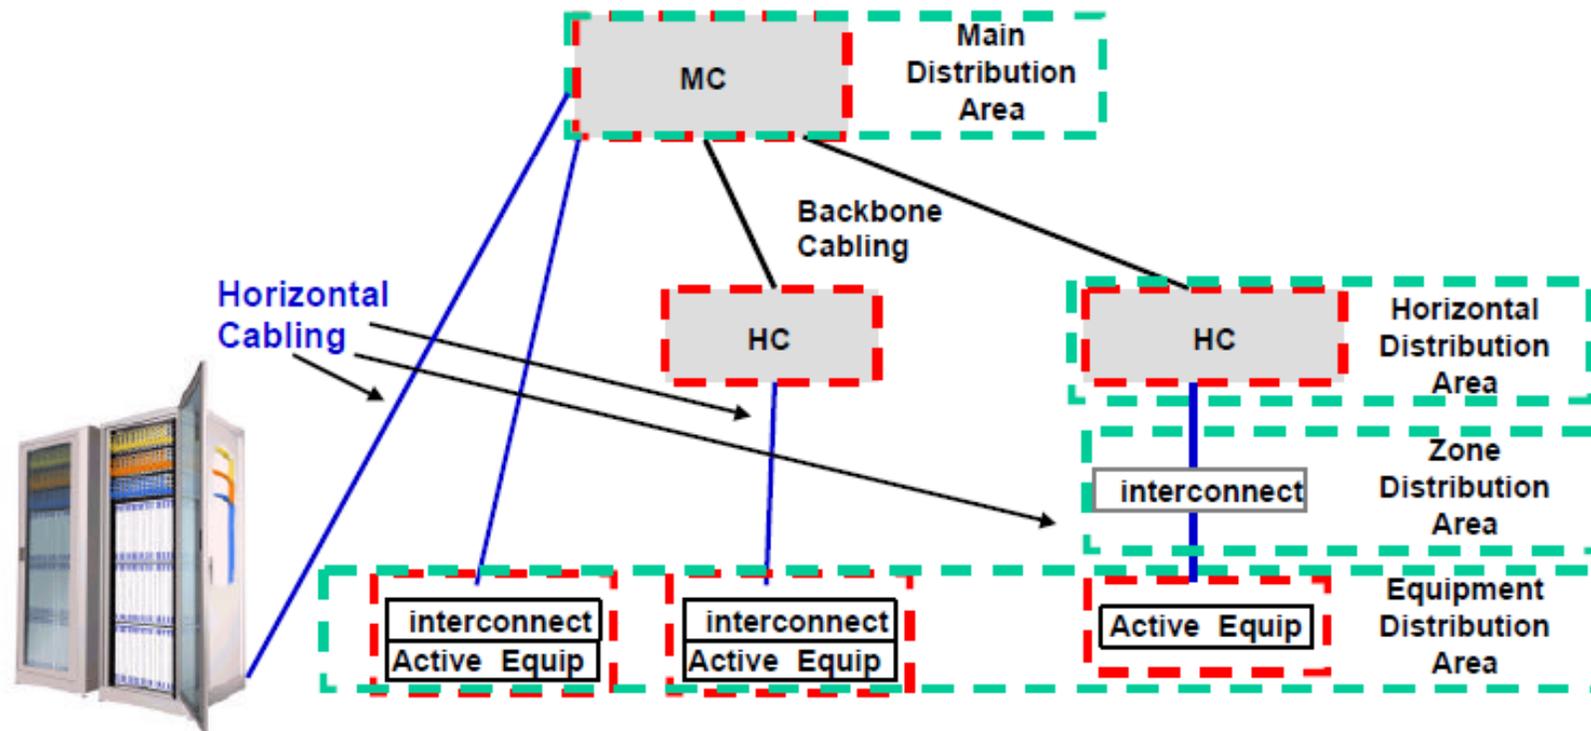


- ▶ UTP cable - 100 Ω (ANSI/TIA/EIA-568-B.2)
 - ▶ UTP Categoria 6 (ANSI/TIA/EIA-568-B.2-1)
- ▶ Multimodal optical fiber (ANSI/TIA/EIA-568-B.3) - 62.5/125 μm or 50/125 μm
 - ▶ Multimodal optical fiber 50/125 μm optimized for 850 nm laser (ANSI/TIA-568-B.3-1)
- ▶ Single mode optical fiber (SMF) (ANSI/TIA/EIA-568-B.3)
- ▶ Coax cable 75-ohm
 - ▶ Telcordia Technologies GR-139-CORE

TIA-942: Horizontal cabling



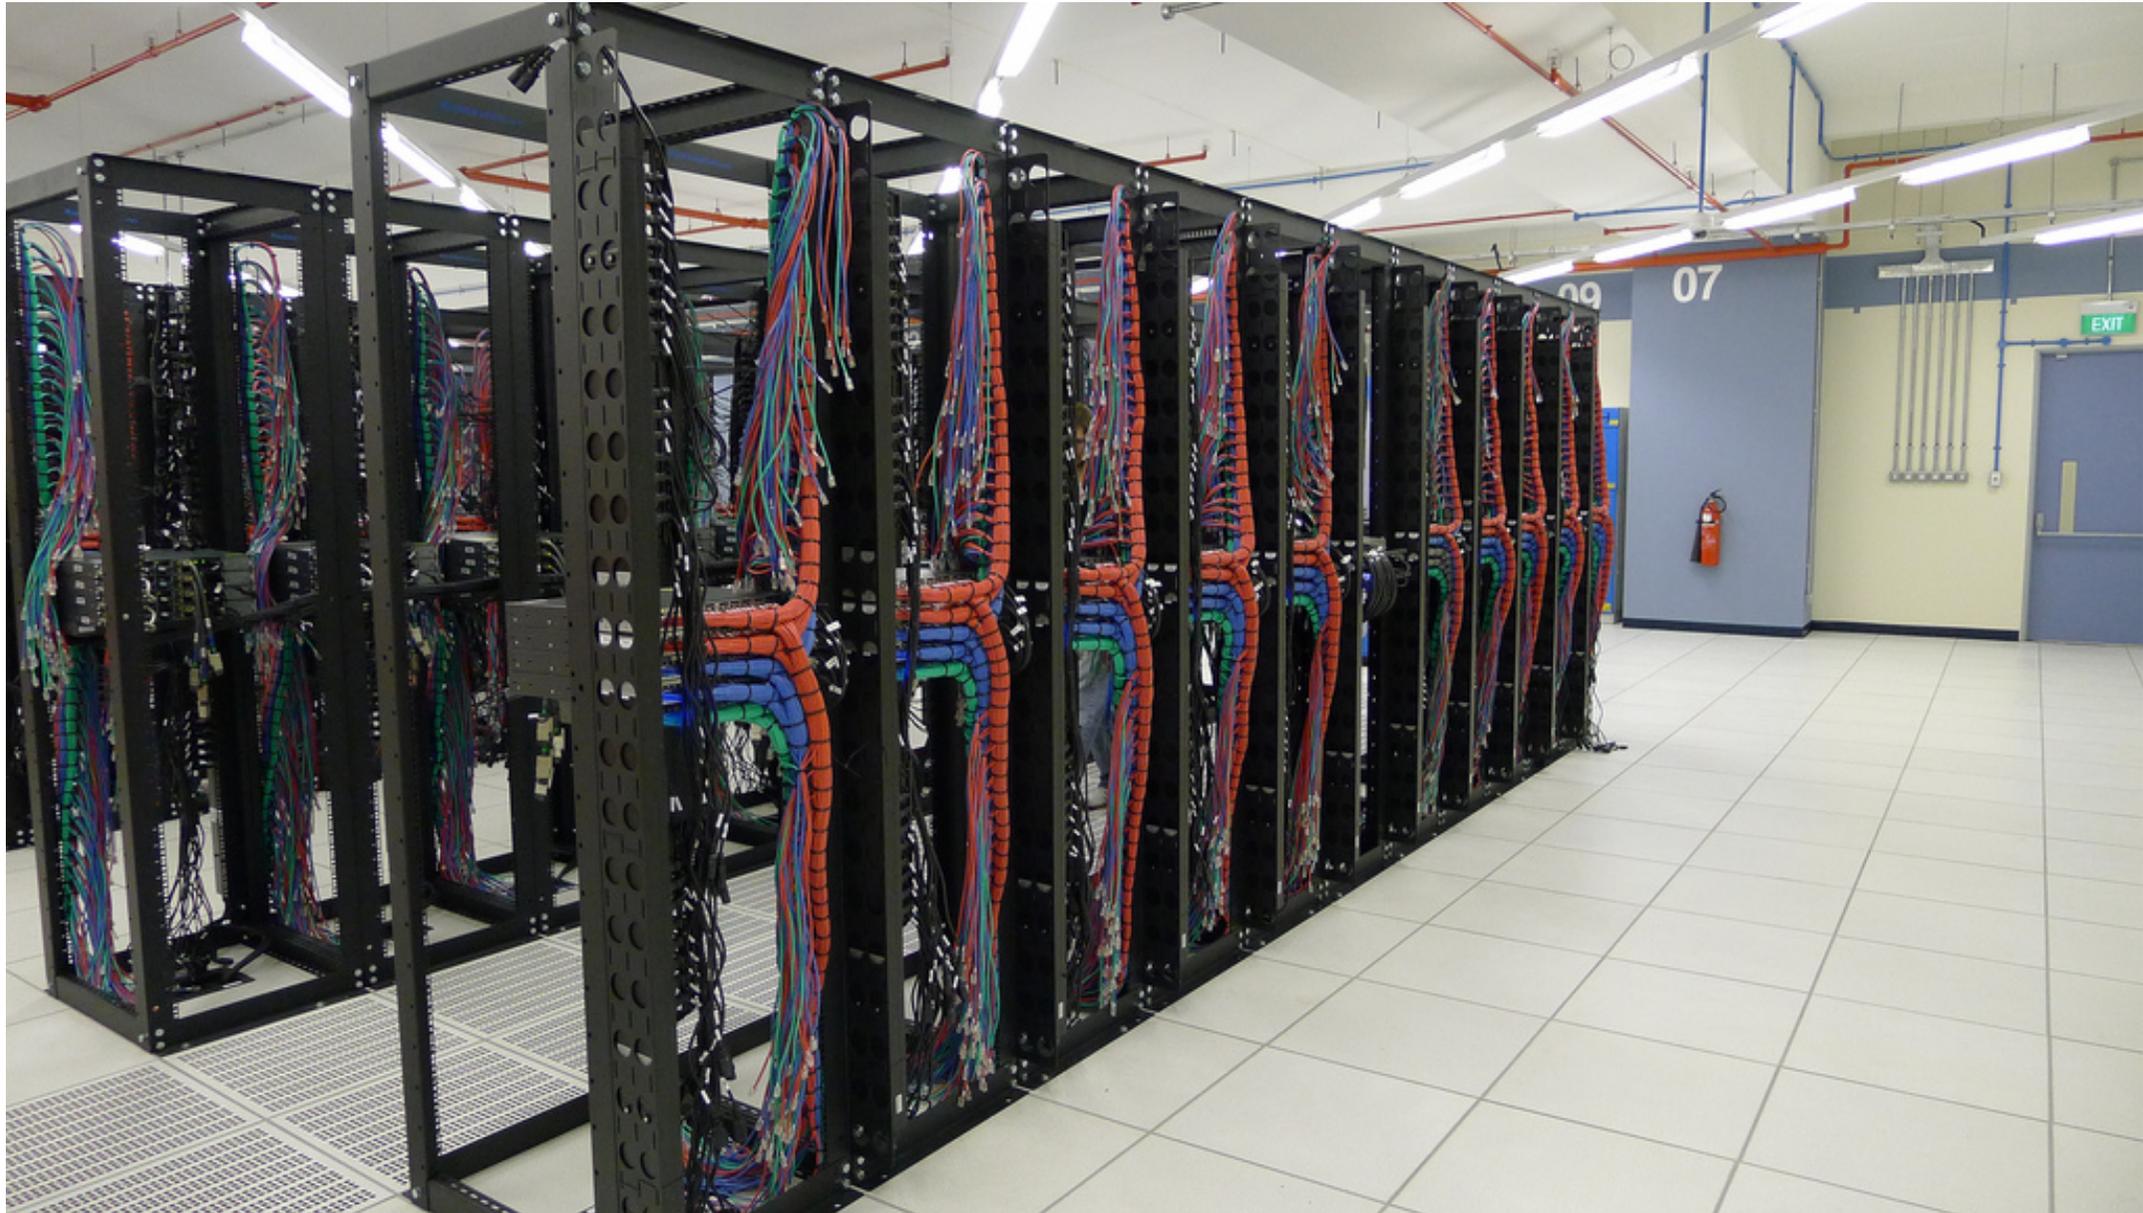
- ▶ Horizontal cabling is the cabling from the Horizontal Crossconnect (HC) in the Main Distribution Area (MDA) or Horizontal Distribution Area (HDA) to the outlet in the Equipment Distribution Area (EDA) or Zone Distribution Area (ZDA)
- ▶ Backbone cabling includes cabling from MDA to ER and HDA





- ▶ The maximum horizontal distance is 90 m independent of media type
- ▶ The maximum channel distance including equipment cords is 100 m
- ▶ The maximum cabling distance in a data center not containing a horizontal distribution area is:
 - ▶ 300 m for an optical fiber channel including equipment cords
 - ▶ 100 m for copper cabling including equipment cords

Datacenter: rack cabling



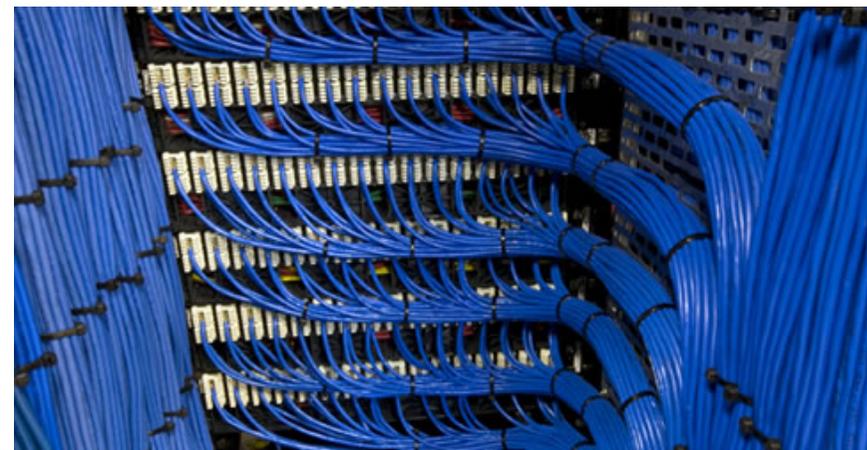
Cabling: patch panels



- ▶ A *patch panel* is a panel, that houses cable connections. One shorter patch cable will plug into the front side, whereas the back holds the connection of a (typically much longer) cable that is punched permanently
- ▶ Patch panels are produced for both copper cables (UTP) and fiber optic cables
- ▶ In a DC, links between the various parts of a DC are usually terminated onto patch panels (*structured cabling*)
- ▶ Patch panels's front ports and switch ports are connected by *patch cords*, i.e. cables with connectors at both sides
- ▶ In so called Top-of-Rack configurations, servers are usually directly connected to ToR switches with no patch panels in the middle



Patch panel for UTP cables: front view

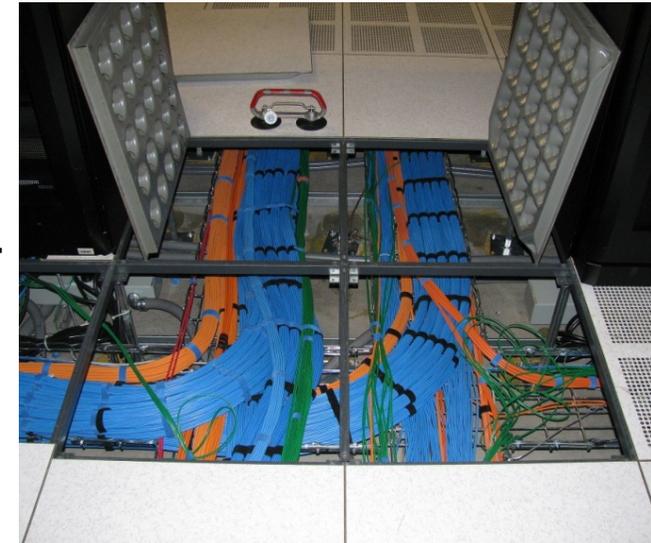


Patch panel for UTP cables: rear view

Datacenter: raised floor



- ▶ A raised floor consists of a gridded metal framework or substructure of adjustable-height supports (called "pedestals") that provide support for removable (*liftable*) floor panels (*tiles*), which are usually 2×2 feet or 60×60 cm in size
- ▶ A void space (*underfloor plenum*) is created below the floor for establishing pathways for piping, wires, cable, as well as to uniformly distribute conditioned air
- ▶ Many types of commercially available floors exist that offer a wide range of structural strength and loading capabilities
- ▶ Stringerless raised floors
 - ▶ Consists of an array of pedestals that provide the necessary height for routing cables and also serve to support each corner of the floor panels
 - ▶ No mechanical attachments between the pedestal heads
 - ▶ Pros: Maximum accessibility to the space under the floor
 - ▶ Cons: significantly weaker than stringered raised floors
- ▶ Stringered raised floors
 - ▶ Consists of a vertical array of steel pedestal assemblies uniformly spaced on two-foot centers and mechanically fastened to the concrete floor (*slab*)
- ▶ Structural platforms
 - ▶ Consists of members constructed of steel angles or channels that are welded or bolted together to form an integrated platform for supporting equipment
 - ▶ Structural platforms may or may not contain panels or stringers
 - ▶ Pros: permits equipment to be fastened directly to the platform

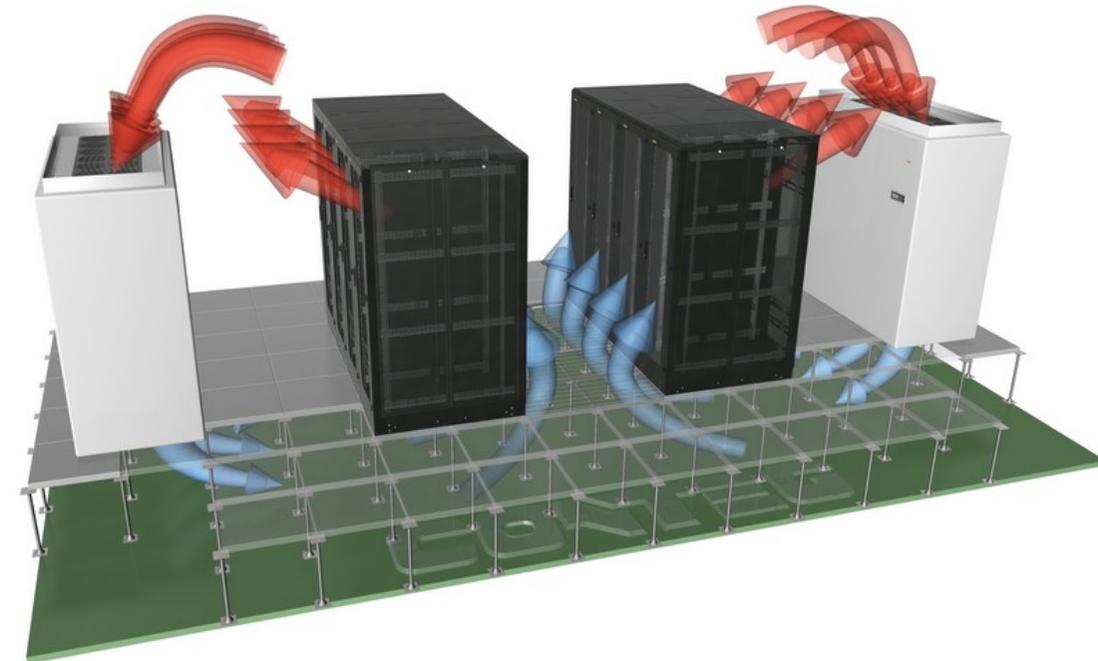


"Under Floor Cable Runs Tee" by Robert.Harker

Datacenter: racks layout



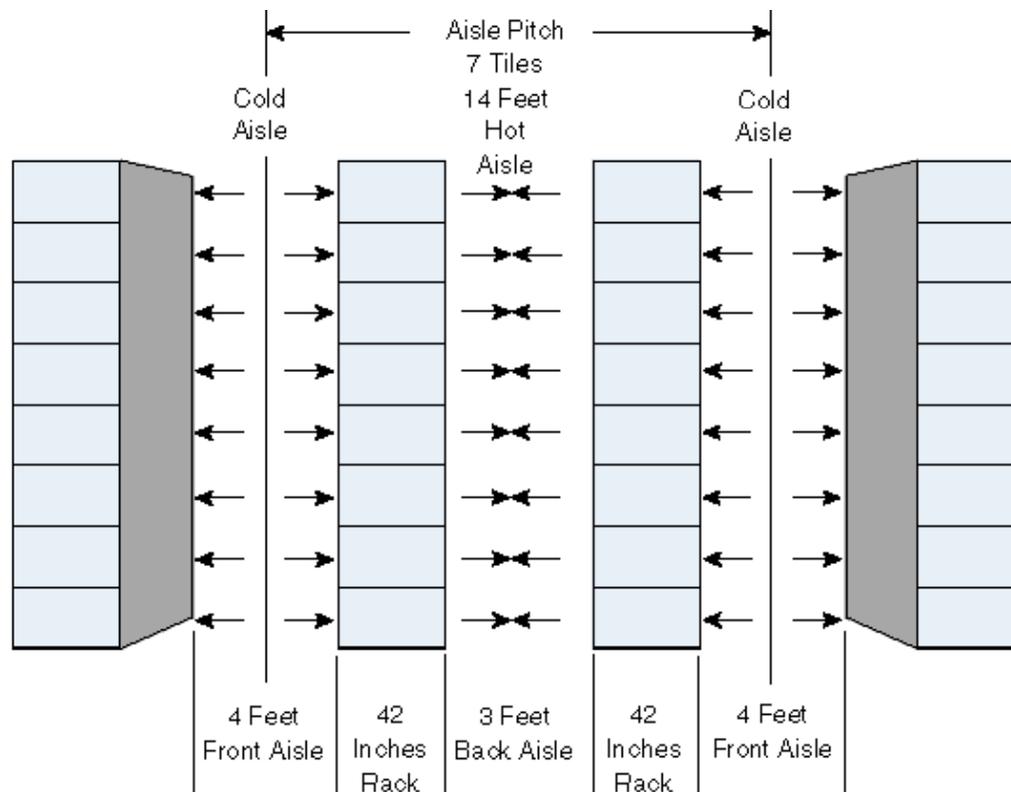
- Racks form rows separated by aisles
- Typical *hot-and-cold aisle layout*:
 - Servers in adjacent rows face-to-face
 - Air conditioning systems, positioned around the perimeter of the room or at the end of hot-aisles, push *cold air* under the raised floor
 - Cold air flows from the raised floor through perforated tiles
- Cold aisles are 2 floor tiles large
 - Standard floor tile size: 60x60 cm
 - Typical cold aisle width: 120 cm
 - Typical hot aisle width: 90 cm



Datacenter: aisle pitch



- ▶ *Aisle pitch* is the distance between center of a cold aisle and the center of the next cold aisle
- ▶ Data centers often use a seven-tile aisle pitch
 - ▶ Typical cold aisle width: 4 feet (\approx 120 cm)
 - ▶ Typical hot aisle width: 3 feet (\approx 90 cm)
 - ▶ Space left for rack depth: 42 inches = 3.5 feet (\approx 105 cm)
 - ▶ Typical aisle pitch: 7 tiles \rightarrow 14 feet (\approx 420 cm)

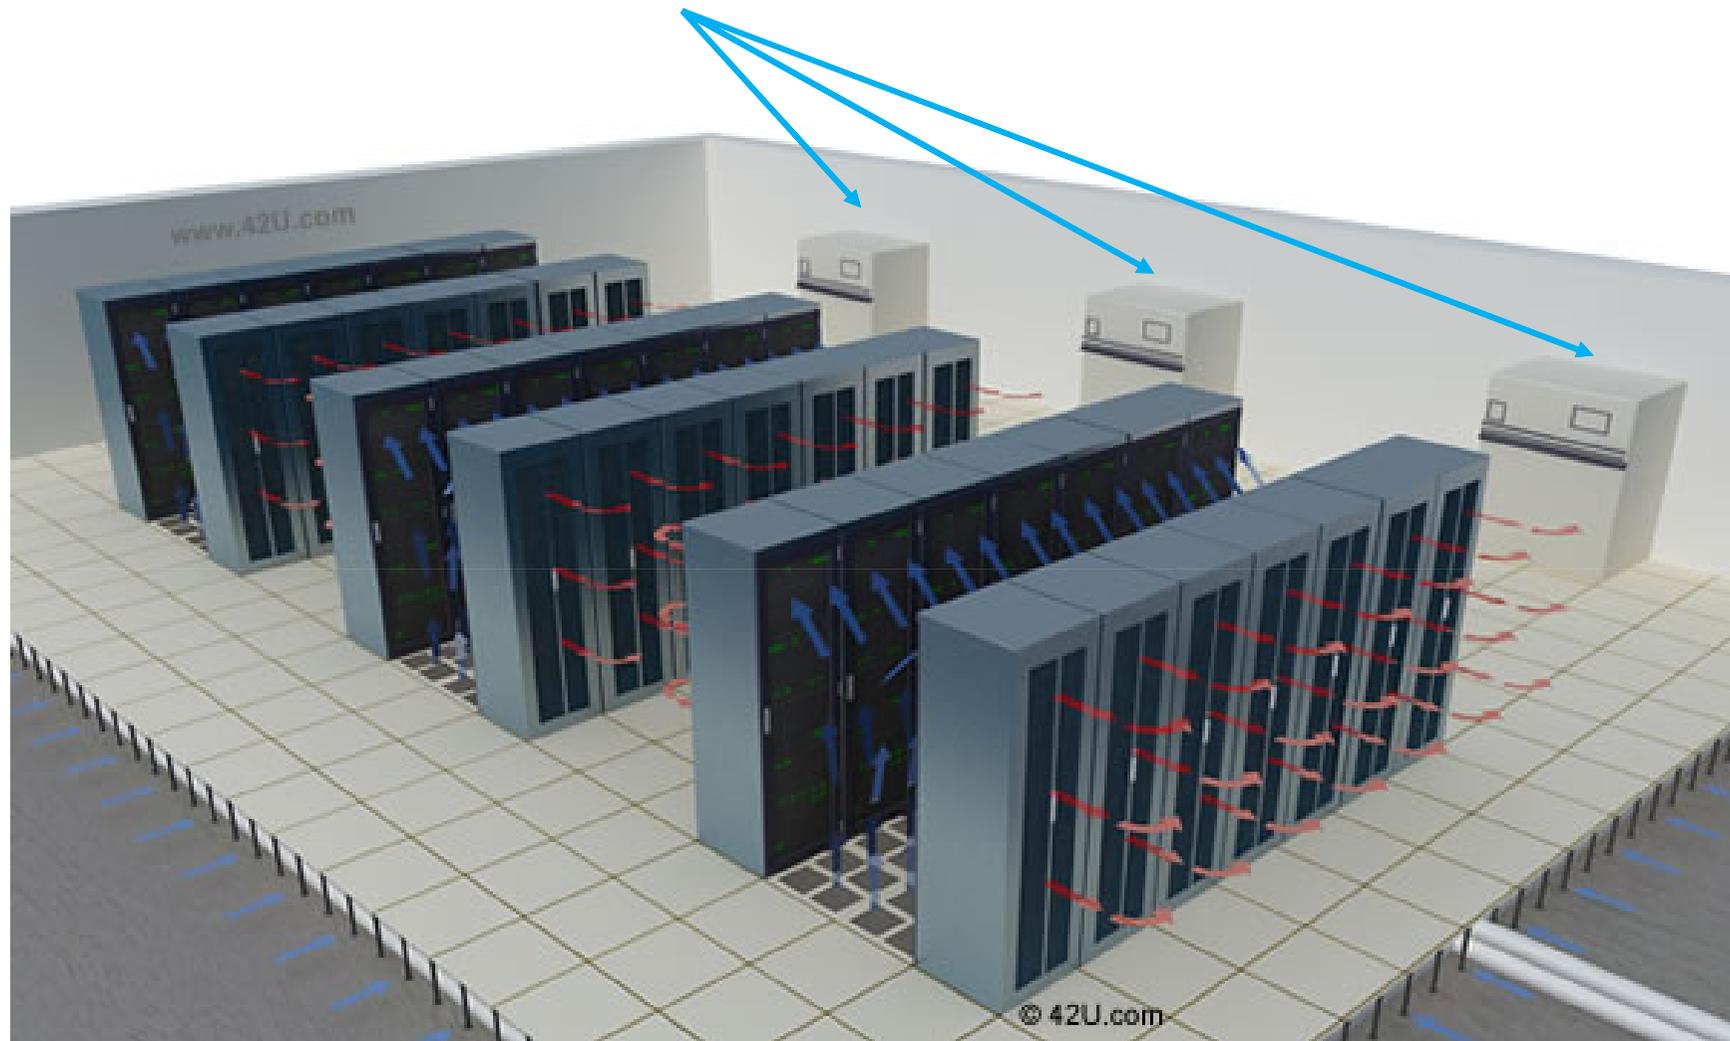


Seven-tile aisle pitch

Datacenter: air conditioning (1)



Heating, ventilation, and air conditioning (HVAC) units push cold air under the raised floor



Datacenter: air conditioning (2)

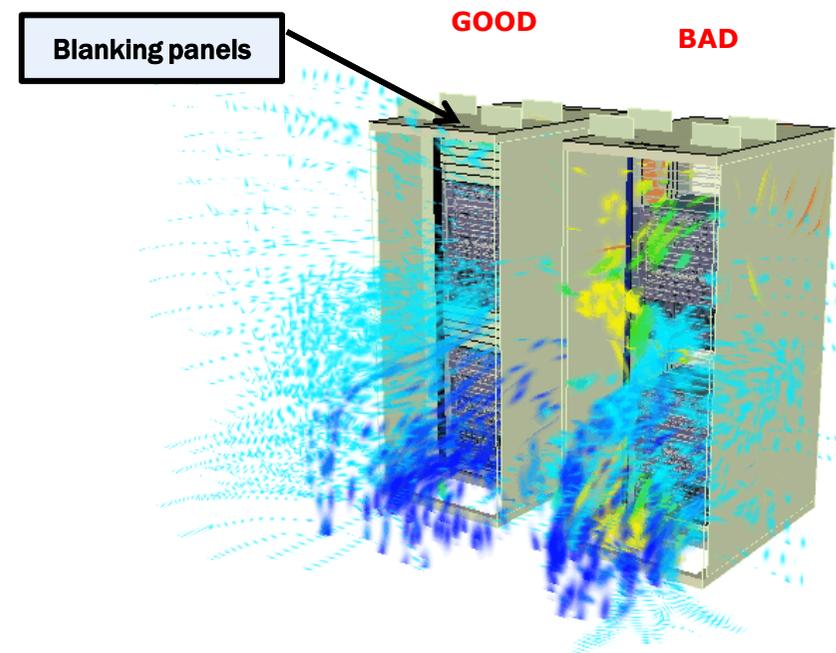
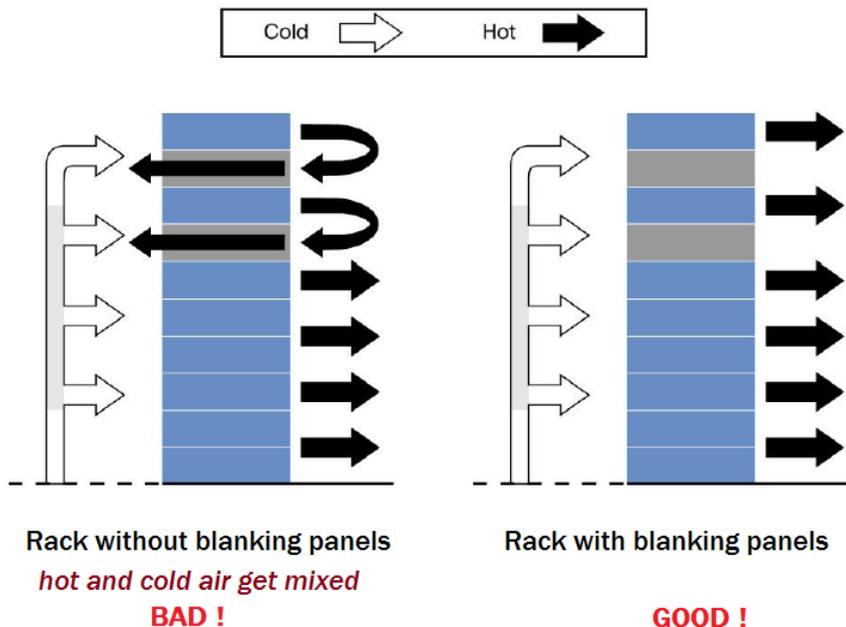


- ▶ Purpose of air conditioning system: keep in the equipment room the environmental optimal conditions for IT devices
 - ▶ Guidelines provided by the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) Technical Committee 9.9
 - ▶ http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf
 - ▶ Acceptable temperature range: 18 – 27 ° C (where should it be measured ?)
- ▶ Achieving this goal requires removing the heat produced by the IT devices and transferring it to the external environment
- ▶ The amount of energy consumed by the air conditioning system depends on:
 - ▶ the amount of heat produced by IT equipments
 - ▶ Temperature difference between DC's equipment room and the external environment
- ▶ Until recently, most of DCs used only air cooling systems
 - ▶ Suitable up to 1.500 / 2.000 W*m²
- ▶ Recent trend for DCs: increase device density in the equipment room
 - ▶ This make air conditioning more challenging
- ▶ In some high-density DCs water cooling is also employed
 - ▶ More challenging to design and build
 - ▶ In this case, power supply cables need to run on overhead trays and not in the underfloor for security reasons (eg. *flooding of the equipment room in case of water leaks from cooling pipes*)

Datacenter: air conditioning (3)



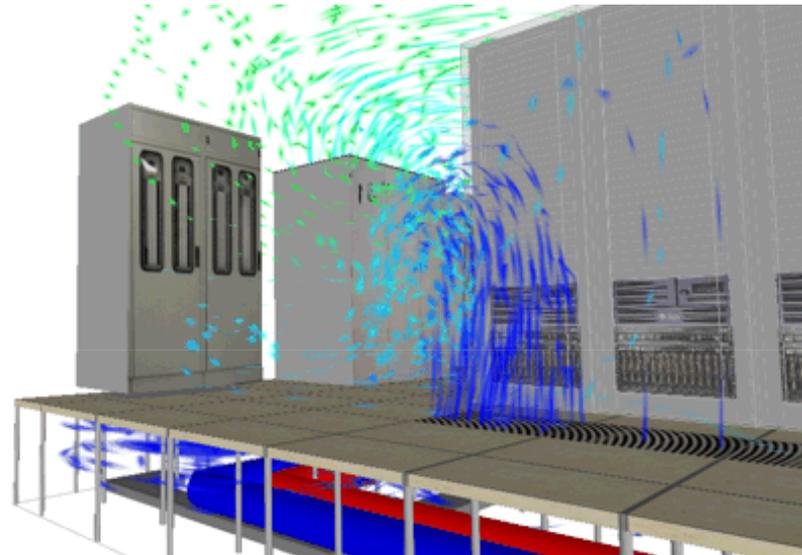
- ▶ The purpose of the hot-and-cold-aisle layout is to minimize the mix of hot air and cold air flows in order to increase the efficiency of the air conditioning system
- ▶ When such a mix occurs, the air conditioning system has to pump cold air at an even lower temperature to obtain the same air conditioning effect in the equipment room
- ▶ Best practices:
 - ✓ Fill racks with heaviest and hottest equipments at the bottom (lowest part of racks)
 - ▶ for stability reasons and to make them closer to perforated tiles in the floor (from which cold air flows)
 - ✓ If a rack is not fully filled with equipments, fill void positions with *blanking panels*
 - ▶ blanking panels prevent hot air from flowing from the front of racks and getting mixed with cold air



Datacenter air conditioning: containment (1)



- ▶ Air conditioning inefficiency:
 - ▶ Cold air flowing from the raised floor through perforated tiles does not flow into the racks but directly reaches the CRAC unit

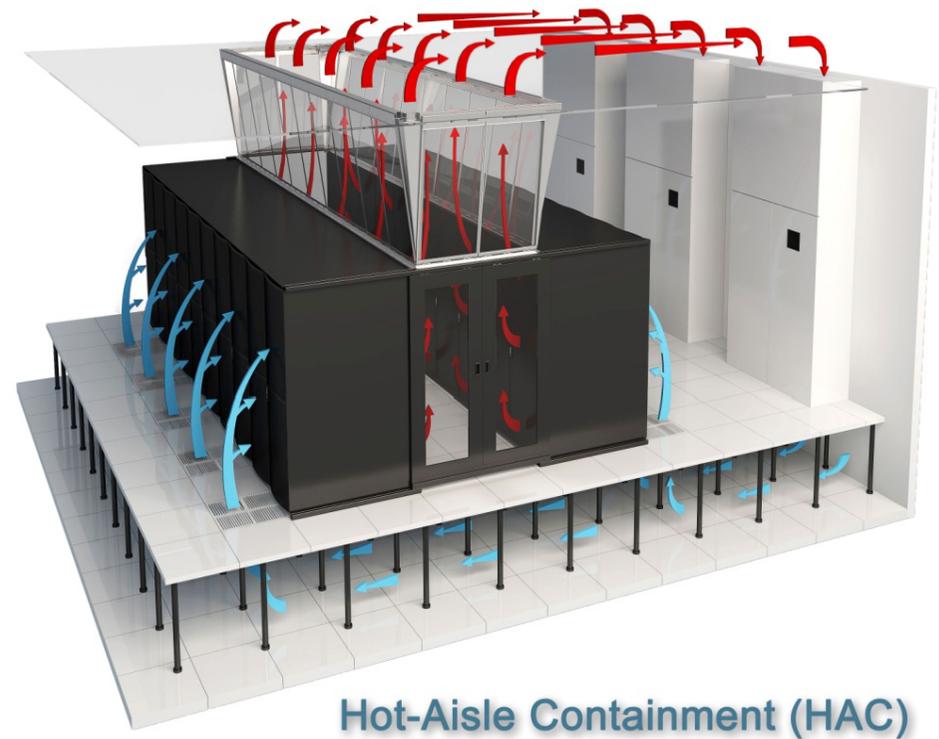
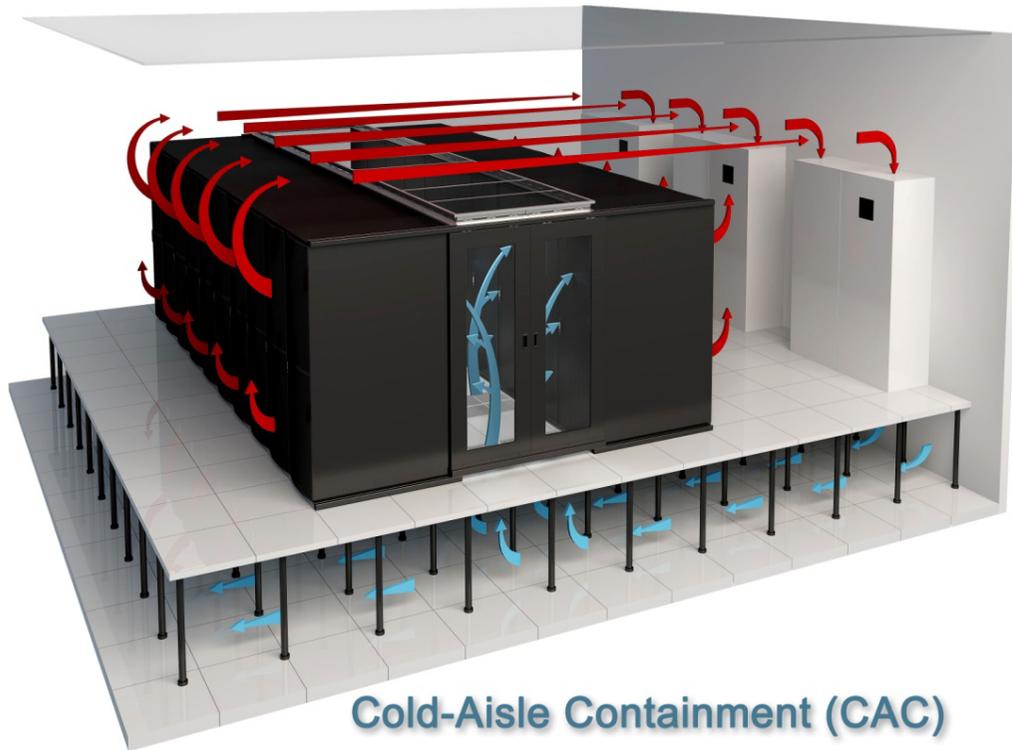


- ▶ Solution: prevent cold and hot air mix through physical barriers (“containment”)

Datacenter air conditioning: containment (2)



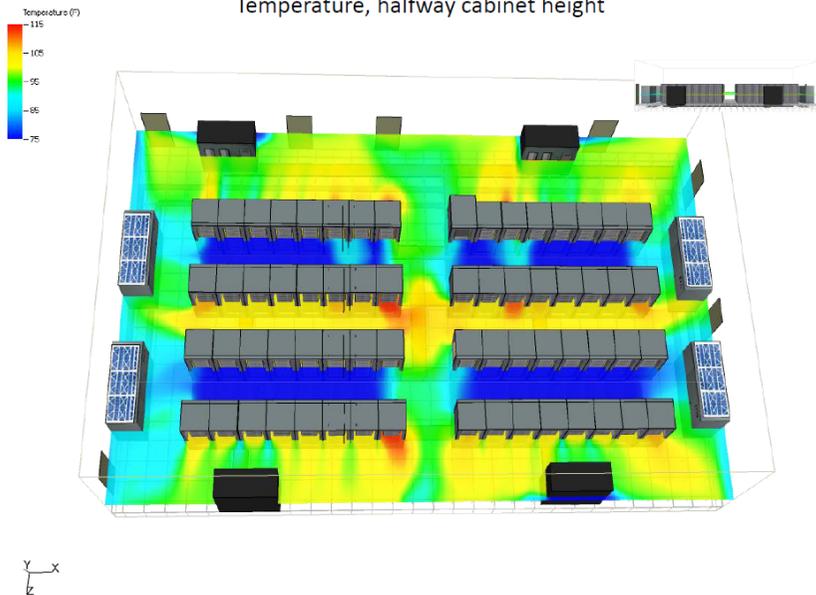
- ▶ Cooling containment solutions prevent hot and cold air flows to mix by means of physical barriers



► Comparison of temperatures without and with cold-aisle containment

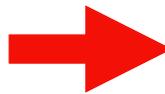
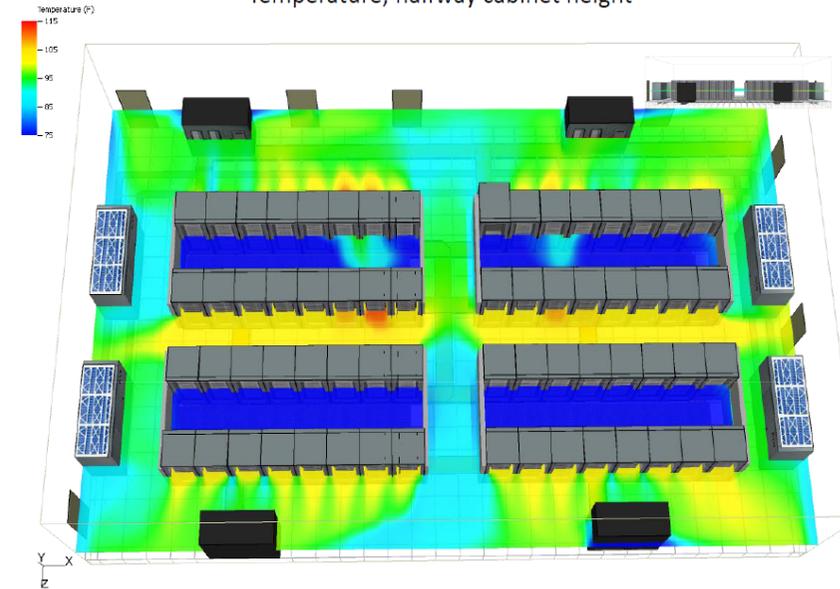
Standard (Hot Aisle/Cold Aisle)

Temperature, halfway cabinet height



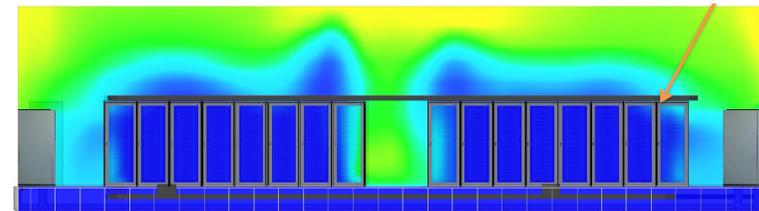
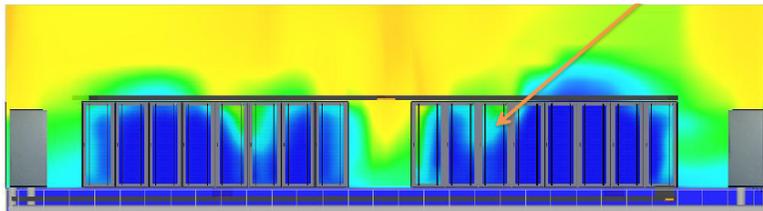
Standard Cold Aisle Containment

Temperature, halfway cabinet height



GREAT LAKES
CASE & CABINET

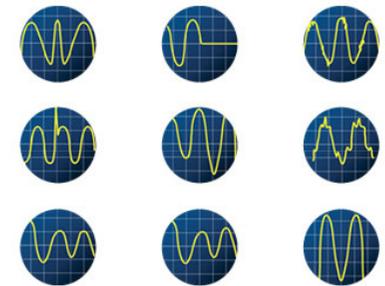
GREAT LAKES
CASE & CABINET



Datacenters: UPS system



- ▶ An UPS is a system that supplies electric energy, in the form of alternate current, to IT equipments with no interruptions in time, in case of a black-out
 - ▶ Near-instantaneous protection from power interruptions
- ▶ Different types of UPS systems:
 - ▶ Static UPSes (equipped with batteries)
 - ▶ Dynamic UPSes (*flywheel*)
- ▶ An UPS also conditions incoming power provided by the electric grid so to protect IT equipments from “power disturbances”
 - ▶ IEEE Standard 1159-1995, "IEEE Recommended Practice for Monitoring Electrical Power Quality" describes many power quality problems
 - ▶ Power disturbances classified into 7 categories based on wave shape: Transients, Interruptions, Sag / Undervoltage, Swell / Overvoltage, Waveform distortion, Voltage fluctuations, Frequency variations
- ▶ An UPS intervention is required for the few minutes that are needed to start an *emergency generator* (e.g. a Diesel engine generator)

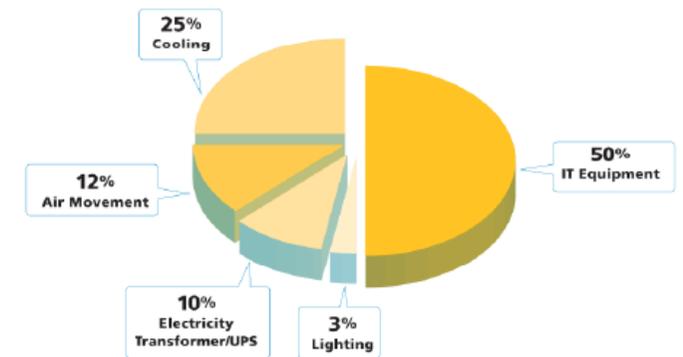
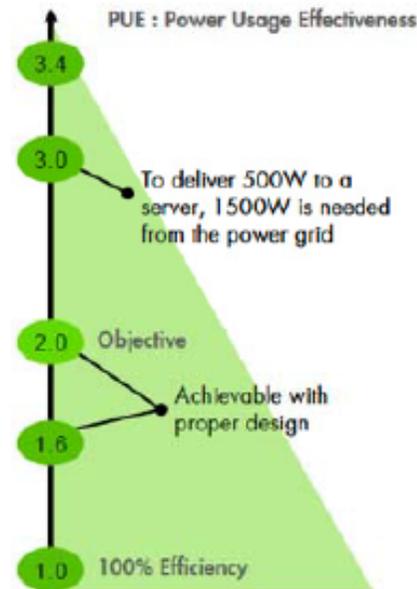
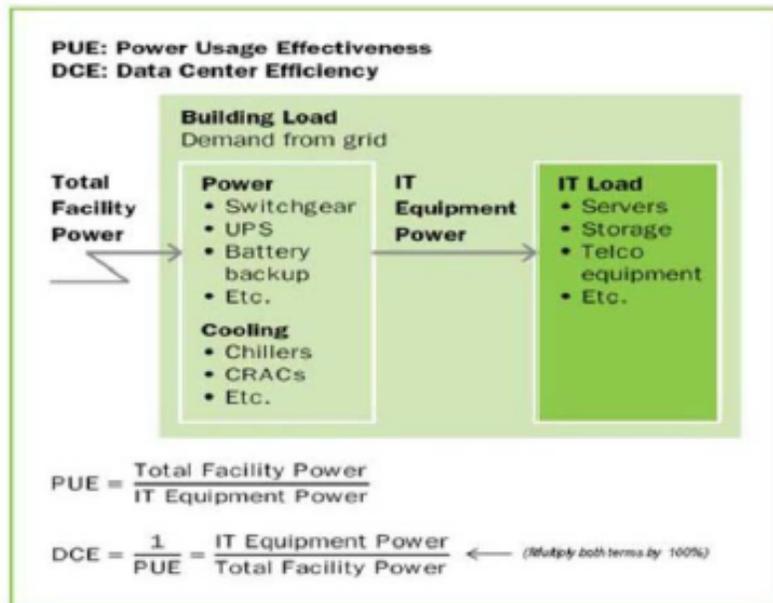




- ▶ Datacenters are believed to consume about 3 per cent of the global electricity supply and account for about 2 per cent of total greenhouse gas emissions, giving the sector around the same carbon footprint as the airline industry
- ▶ In 2014, datacenters consumed about 70 terawatt-hours of electricity in US only, representing 2% of the country's total energy consumption
 - ▶ That's equivalent to the amount consumed by about 6.4 million average American homes that year
- ▶ In 2015, datacenters used 416.2 terawatt-hours of electricity worldwide
 - ▶ UK's total national consumption was about 300 terawatt hours
- ▶ In 2015, less than 20% of the electricity globally used by cloud computing service providers currently came from renewable sources
(Click-Clean Greenpeace report)

DC: Power Usage Effectiveness (PUE)

- ▶ A datacenter's PUE index is a measure of how efficiently it uses energy
- ▶ PUE is the ratio of total power required by the whole DC to power needed by IT equipments only
 - ▶ Excess energy mainly consumed by air conditioning and emergency power systems
 - ▶ Ideally, PUE should be 1, actually it may be 3 or more for DCs whose design did not follow energy efficiency criteria
 - ▶ Nowadays, DCs are designed to achieve a PUE between 1.8 and 1.1 (at best)
 - ▶ In 2015, Facebook's Prineville data center reached a PUE of 1.078
- ▶ Inverse of PUE is the DC energy Efficiency (DCE or DCIE)

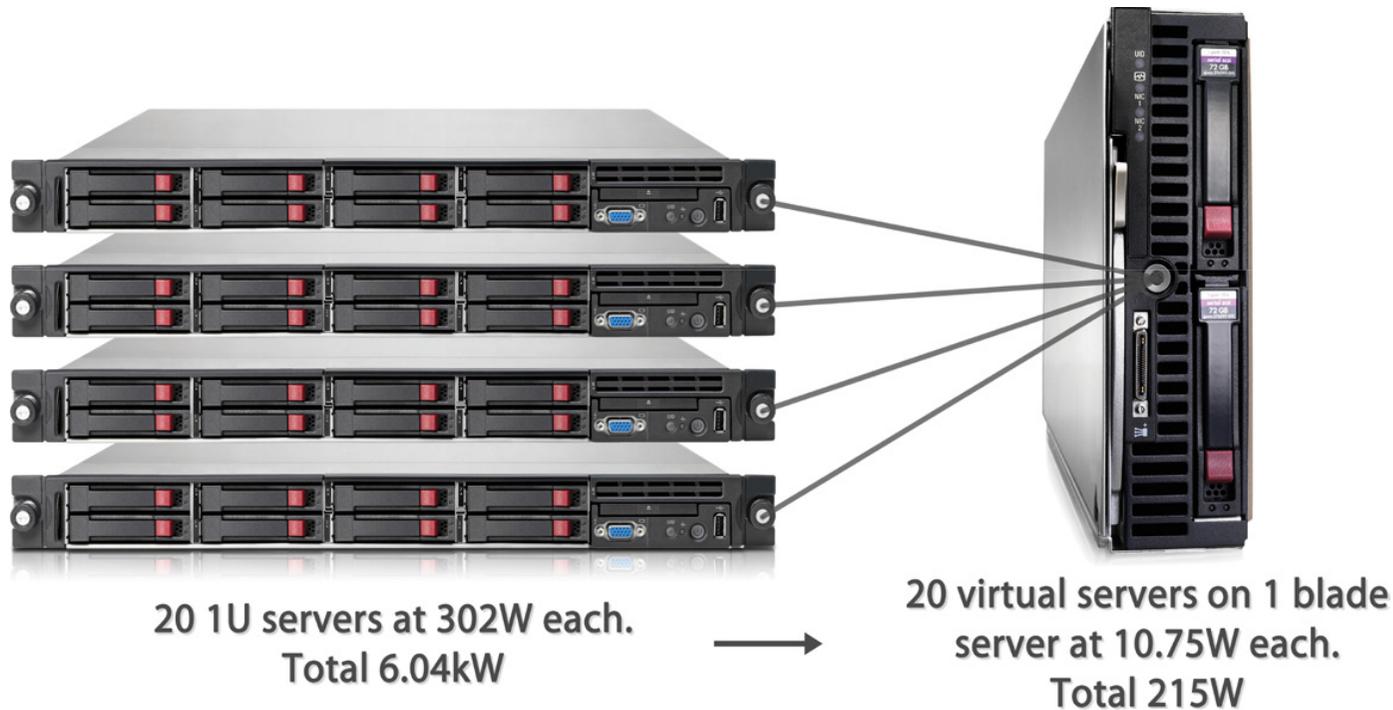


Source: EYP Mission Critical Facilities Inc., New York

Reducing energy costs: role of virtualization



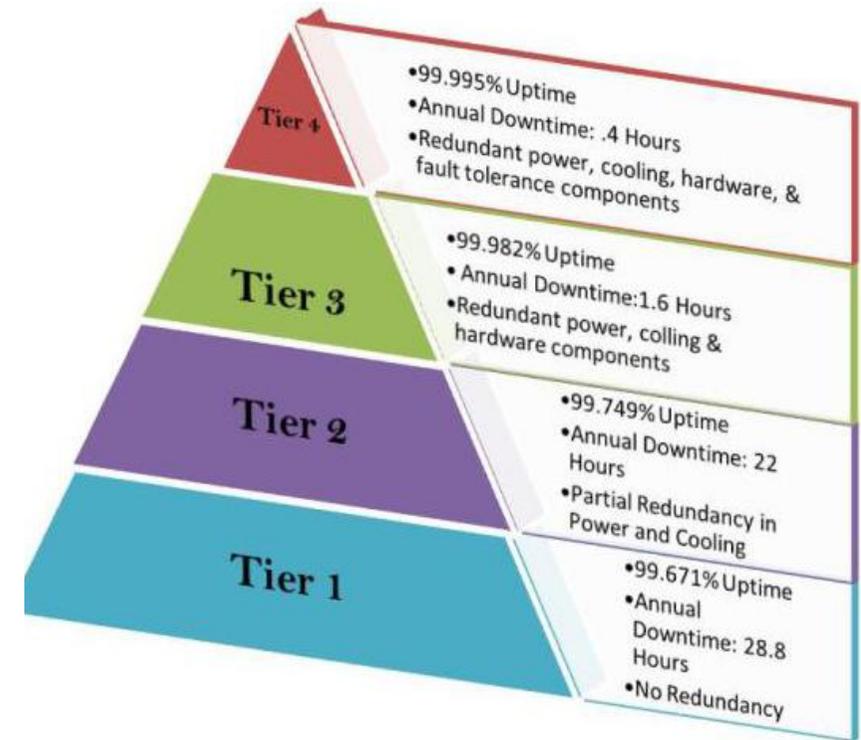
- ▶ Modern datacenters heavily rely on virtualization to consolidate servers load and reduce the energy consumed by servers



Datacenter tiers: Tier 1, 2, 3 and 4



- ▶ Defined by the Uptime Institute
 - ▶ Uptime Institute LLC is a consortium of companies based in Santa Fe, New Mexico, offering professional-services related to the design and operation of datacenters since 1993
- ▶ Tier 1
 - ▶ Basic level of data center reliability
 - ▶ Errors and failures in the systems and equipment lead to a malfunction of the entire data center
 - ▶ Also, the data center is interrupted for maintenance and repair works
 - ▶ No redundant power supplies and no uninterruptible power supply (UPS)
 - ▶ Uptime: 99.671% → Downtime in 1 year: 28.8 hours

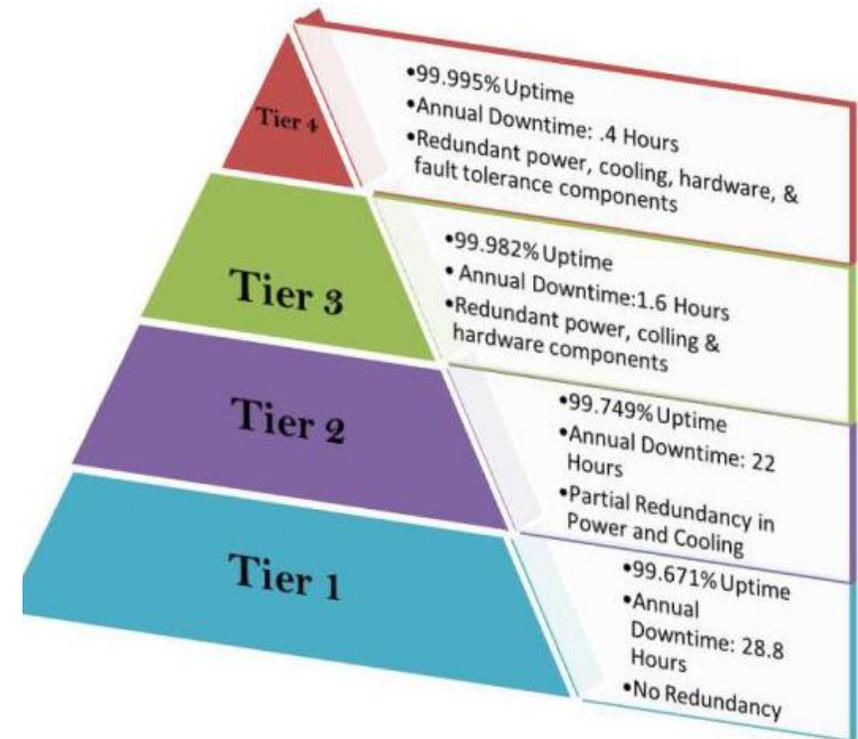


Datacenter tiers: Tier 1, 2, 3 and 4



▶ Tier 2

- ▶ Small level of redundancy and little excess capacity
- ▶ Susceptible to outages due to planned and unplanned failures of the equipment in the data center
- ▶ Technical repair work requires a stop in data center
- ▶ Redundant power supplies
- ▶ Cabling with raised floor
- ▶ Uptime: 99.749% → Downtime in 1 year: 22 hours

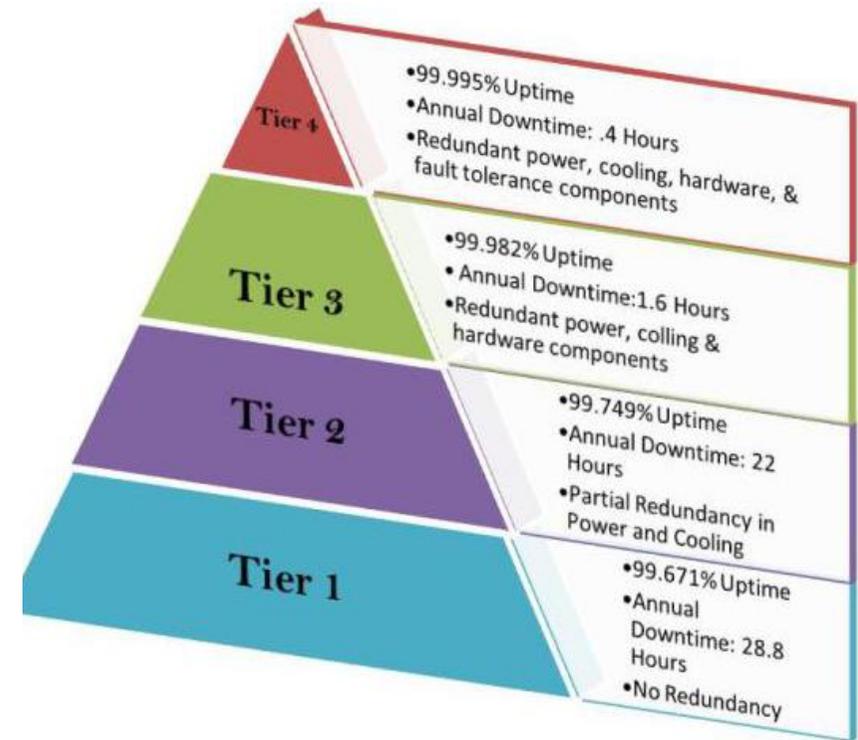


Datacenter tiers: Tier 1, 2, 3 and 4



▶ Tier 3

- ▶ Allows for repair and maintenance work without stopping the data center
- ▶ All IT equipment must be dual-powered
- ▶ Redundant power supplies, cooling and uninterruptible power supply (UPS)



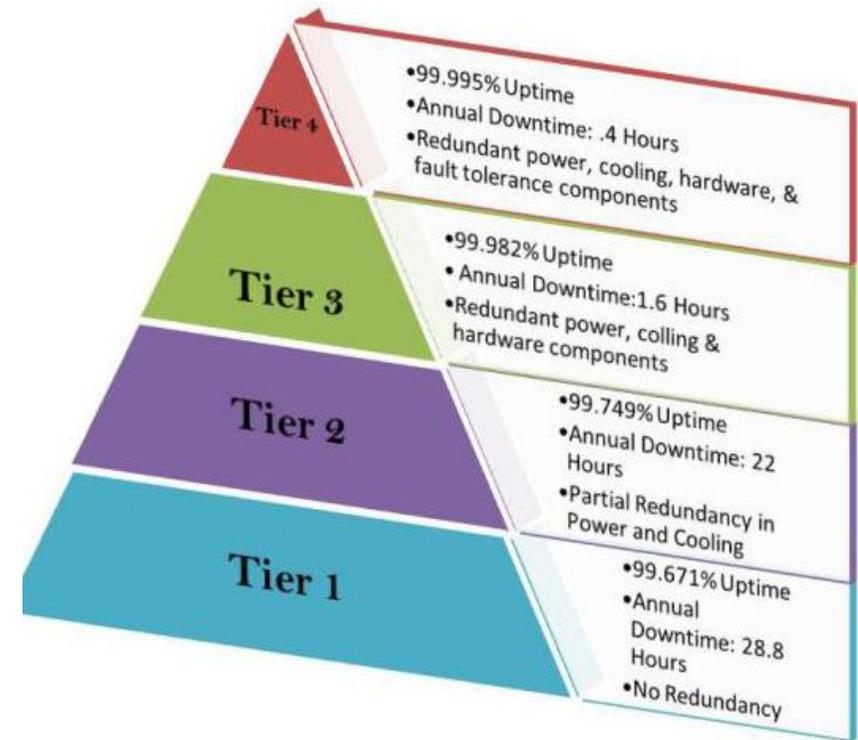
- ▶ Uptime: 99.982% → Downtime in 1 year: 1.6 hours

Datacenter tiers: Tier 1, 2, 3 and 4



▶ Tier 4

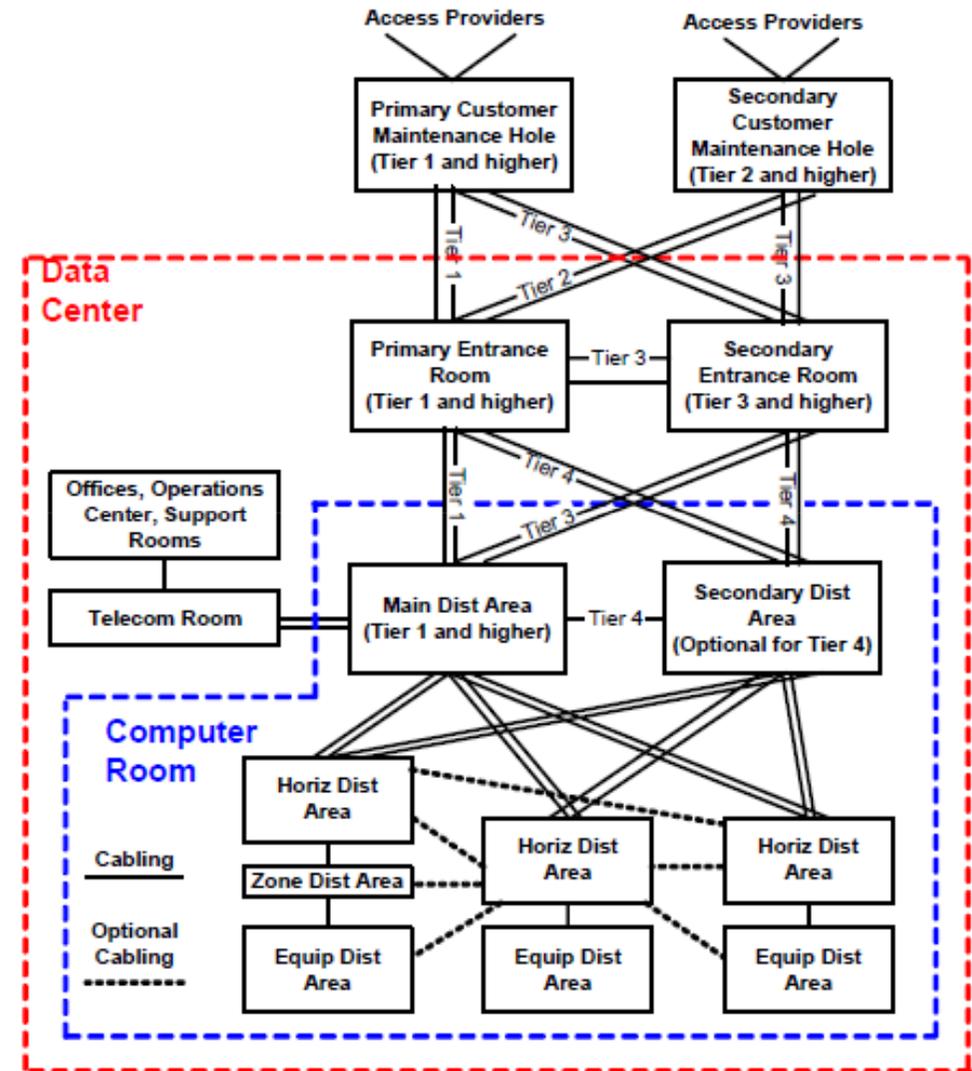
- ▶ Fault-tolerant data center with redundancy of all systems
- ▶ Allows scheduled and unscheduled work without stopping the data center
- ▶ Redundant power supplies and redundant uninterruptible power supply (UPS)
- ▶ All cooling equipments independently dual-powered, including chillers and heating, ventilating and air-conditioning (HVAC) systems
- ▶ Uptime: 99.995% → Downtime in 1 year: 0.4 hours



Datacenter cabling: redundancy



- ▶ Redundancy is required to obtain a higher tolerance to faults
- ▶ Tier 2, 3 and 4 datacenters have to be designed by duplicating facilities and connections



Modular datacenters: Container-based DCs



- ▶ Modular Data Center (MDC)
- ▶ Modular composable unit: *container*
 - ▶ A containerized *datacenter only needs*:
 - ▶ Power supply
 - ▶ Network links to the rest of the world
 - ▶ Available in small-size (4 racks, power need < 1MW), medium-size (10 racks, 1-4 MW), large-size (20 racks, > 4MW)
 - ▶ Each container is equipped with an air conditioning system on its own
 - ▶ Very high energy efficiency (PUE $\approx 1,1$)
 - ▶ Large scale datacenters have been built by assembling containerized DCs



- ▶ Also suitable for emergency use
 - ▶ Mobile Data Center

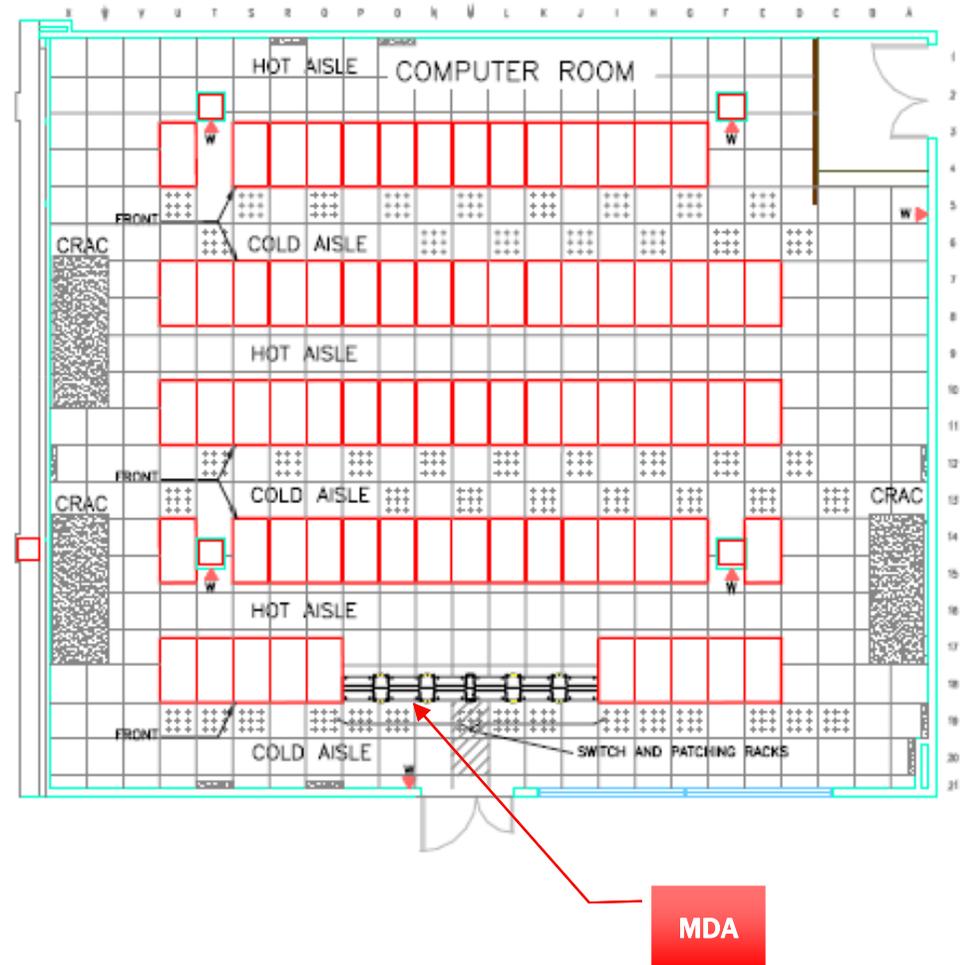


- ▶ The overall datacenter architecture (servers, storage, network, air conditioning, etc.) is influenced by the application
- ▶ Depending on their use, we may classify DCs as follows:
 - ▶ Commercial DCs and Cloud Computing DCs
 - ▶ Es: Google, Facebook, Amazon, Microsoft Azure,
 - ▶ Datacenters hosting scientific applications (HPC, Grid Computing)
 - ▶ OLCF: Oak Ridge Leadership Computing Facility
 - ▶ ALCF: Argonne Leadership Computing Facility
 - ▶ XSEDE: Extreme Science & Engineering Discovery Environment
 - ▶ NWSA: National Center for Atmospheric Research Wyoming Super Computing Center
- ▶ For parallel scientific applications, latency is crucial for performance. Hence, in DCs hosting this kind of applications, besides the traditional Ethernet general-purpose communication infrastructure, another low-latency network infrastructure is built → Eg. based on Infiniband
- ▶ In HPC datacenters, server cooling is usually more challenging due to a higher power density → water-based cooling may be used for servers
 - ▶ Commercial datacenters power density: 1-6 kW/rack
 - ▶ HPC datacenters power density: up to 30 kW/rack

Small scale datacenter: sample organization



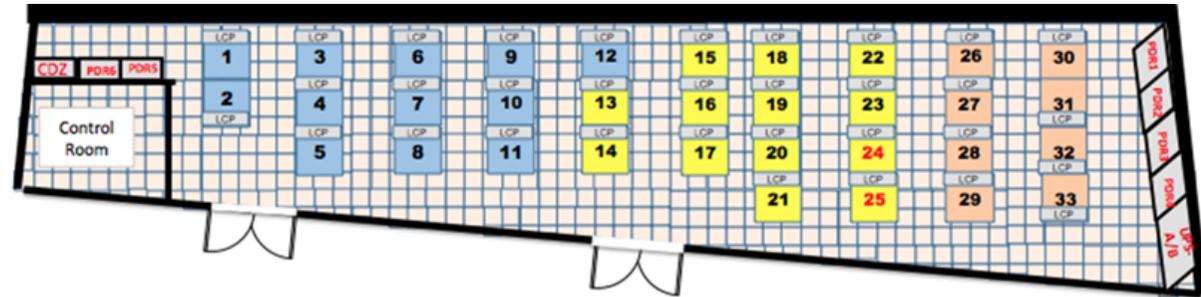
- ▶ Surface extension of about 180 m²
- ▶ 73 server racks to form the EDA
- ▶ 6 racks form the MDA
- ▶ Arrangement of cabinets obtained by alternating hot and cold aisles



University of Napoli Federico II datacenter



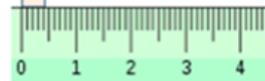
- ▶ Located in Monte S. Angelo campus
- ▶ Funded by projects SCOPE, ReCAS and ATLAS
- ▶ 33 rack
- ▶ 304 server blade DELL M600
- ▶ 19 server DELL Power Edge 1950
- ▶ 10 server DELL Power Edge 2950 with FC connection to the SAN system
- ▶ SAN storage system
 - ▶ EMC² (100 TB)
 - ▶ Dell Equallogic PS5000XV (24 TB)



Numero di rack: 33



40x40 cm



Scala 1:100

