Data Center

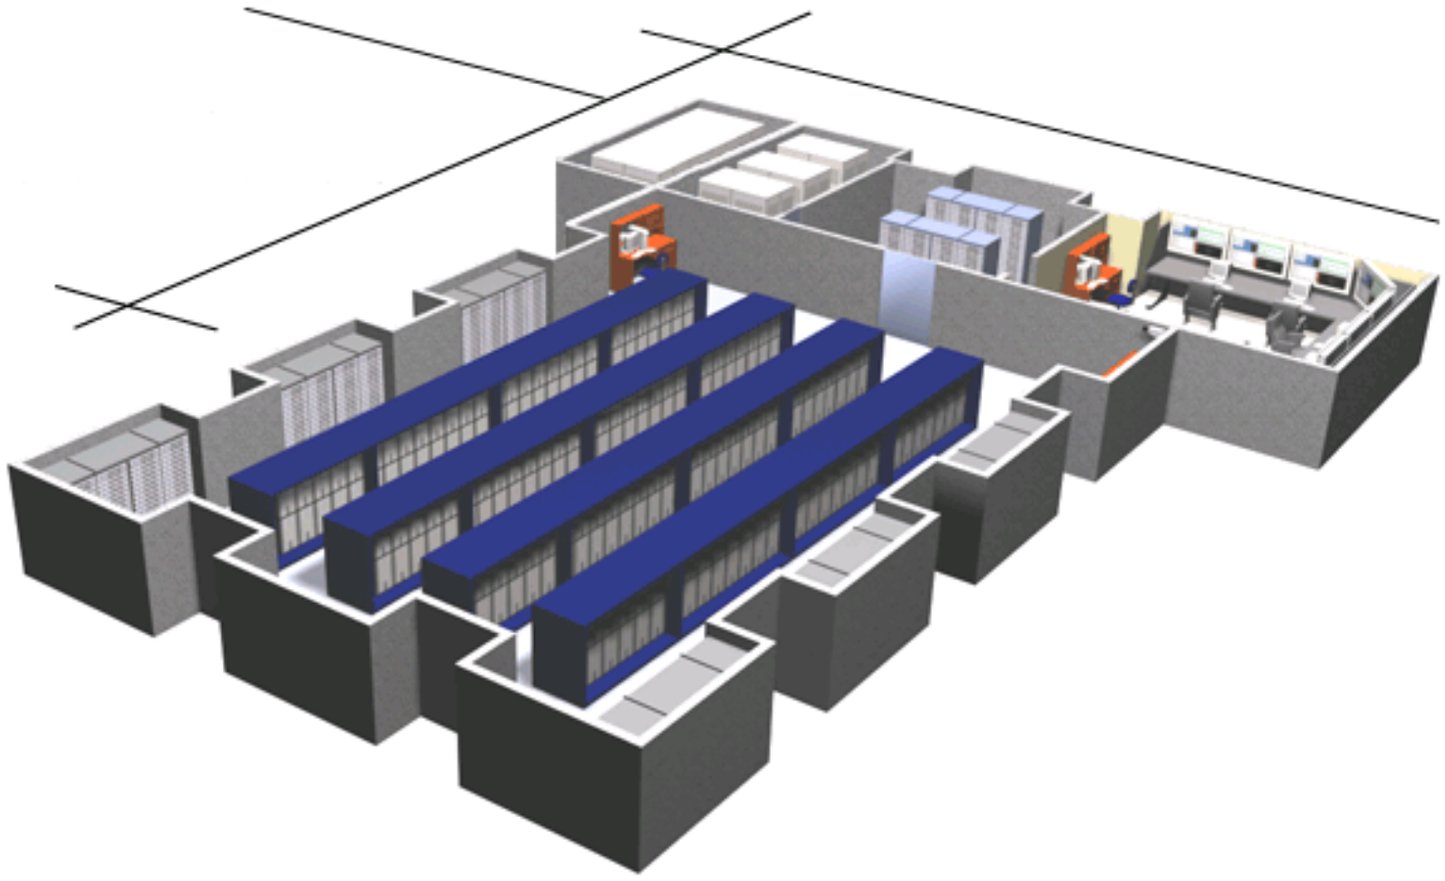

What is a Datacenter?

A data center is a facility used to house computer systems and associated components, such as telecommunications and storage systems.
It generally includes redundant or backup power supplies, redundant data communications connections, environmental controls (e.g., air conditioning, fire suppression) and security devices

Source: http://en.wikipedia.org/wiki/Datacenter

What do we really want?

- **We want a lot of computing power**
  - Corporate infrastructure: employees need to read email
  - Academic computing: students need to do assignments
  - Scientific computing: NASA needs to do galaxy simulations
  - Distributed computation: Google needs to index the web

We buy a cluster...

- **Why we don't just use...**
  - Individual PCs for everyone?
  - A computer cluster where people can do assignments?
- **But this is very inefficient and has issues...**
  - Rooms to store these computers
  - Requirements for shared storage
  - Staff to maintain machines (TCO)
  - Cabling
  - Remote access?
  - Poor amortization of resources

We use a single huge room...

- **Now TCO is reduced...**
- **But...**
  - What sorts of machines do you use? Not PCs: too much space.
  - PCs get hot. You need to keep them cool.
  - What about amortization of resources?
  - Remote Access strategy?
  - Shared storage?
  - Cabling?

We need availability

- **Datacenters have strict standards for reliability and availability**
  - Tier 1: 99.671% Availability: 28 hours of downtime/year
  - Tier 2: 99.741% Availability: 22 hours of downtime/year
  - Tier 3: 99.982% Availability: 1.5 hours of downtime/year
  - Tier 4: 99.995% Availability: 26 minutes of downtime/year

Then, a datacenter is...

- **A datacenter is a collection of computers and storage**
- **A datacenter has many purposes**
  - Provide single location for compute infrastructure for big organizations
  - Provide hosting facilities so companies don't have to manage their own IT
  - Perform large distributed computations (Google Warehouse computing, scientific simulations)
  - Provide scale-out on-demand IT services (e.g., accommodating temporary usage spikes)

Datacenter - Computer Analogy

- **There are many issues that influence the performance/reliability of a desktop**
  - Choice of CPU
  - Cache designs and memory hierarchies
  - OS
  - Storage design (e.g., RAID0-RAID5, etc.)
  - Power/Cooling
- **Similar issues are present in a datacenter.**
  - Storage subsystem
  - Networking infrastructure
  - Power distribution
  - Reliability guarantees (mission-critical?)

Datacenter Design Rules

- **Standard for Datacenter design: TIA-942**
- **Basic principles**
  - Use racks and rack-mount form factor for high density of compute resources
  - Heat harms components: control it
  - Keep cabling manageable
  - Obey FCC rules for EMI/RFI

Rack

- **What goes inside a rack?**
  - Servers
  - Storage
  - Cables
  - Power distribution units (PDUs)
  - Fans
  - Switches
  - KVM

Rack Front/Back

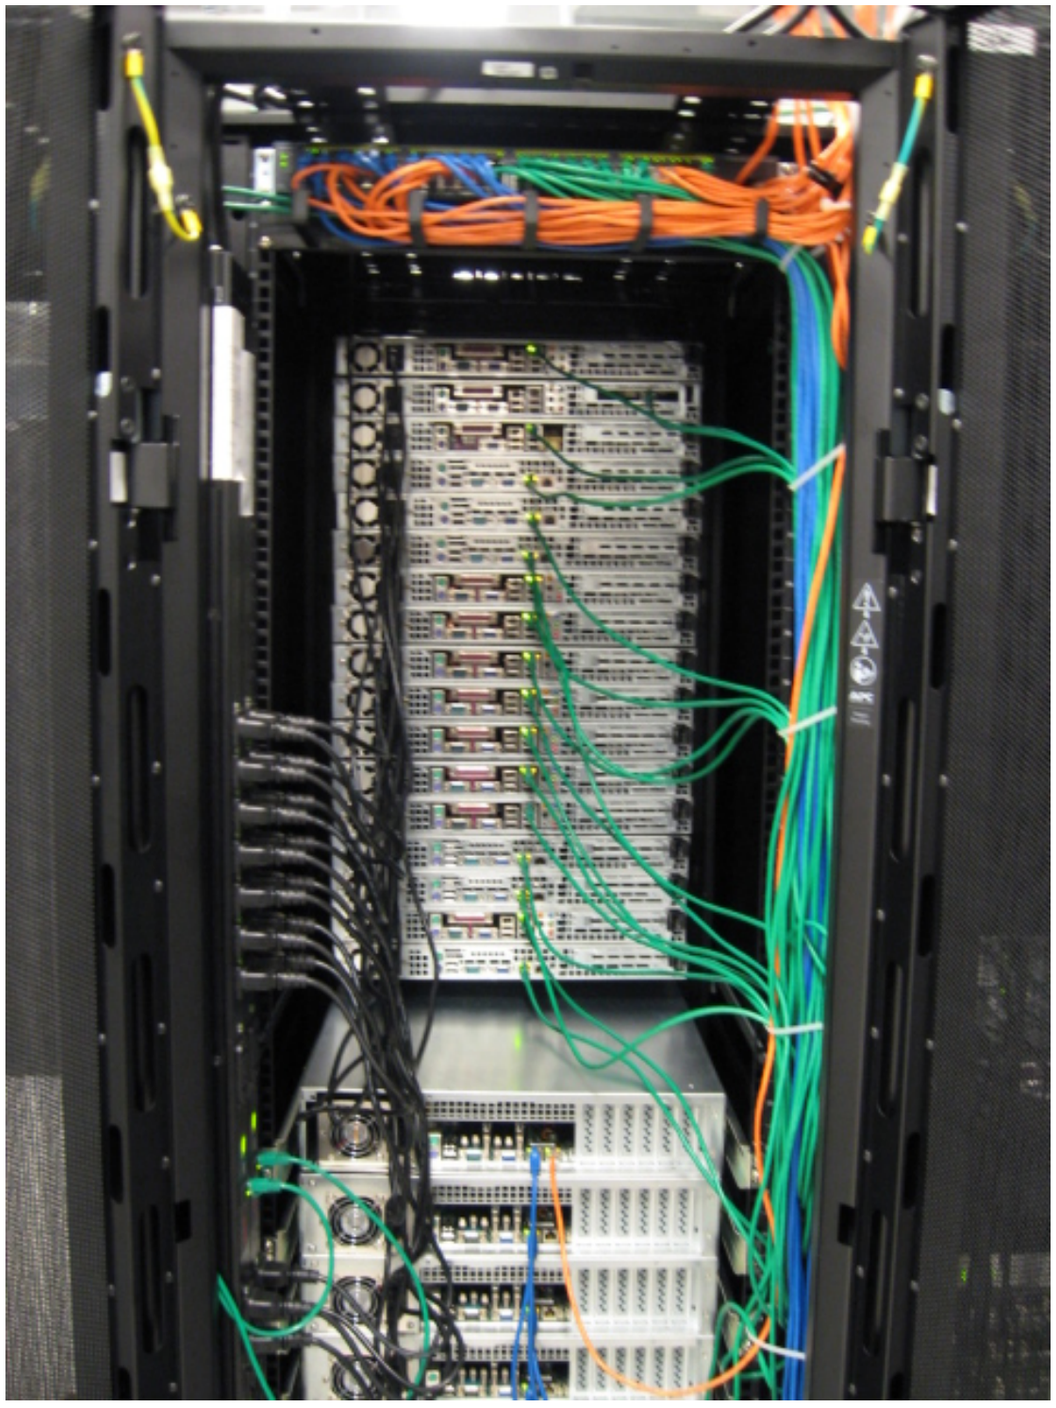

Racks

Cooling

- HVAC: Heating, Ventilation and Air Conditioning
- Inflow: Cool Air from HVAC.
- Outflow: Warmed Air. Eventually cooled by HVAC.
- Never let warmed air re-circulate to inflow or mix with cool air
- Same idea as in a standard PC chassis
- Hot Aisle/Cold Aisle

Cooling

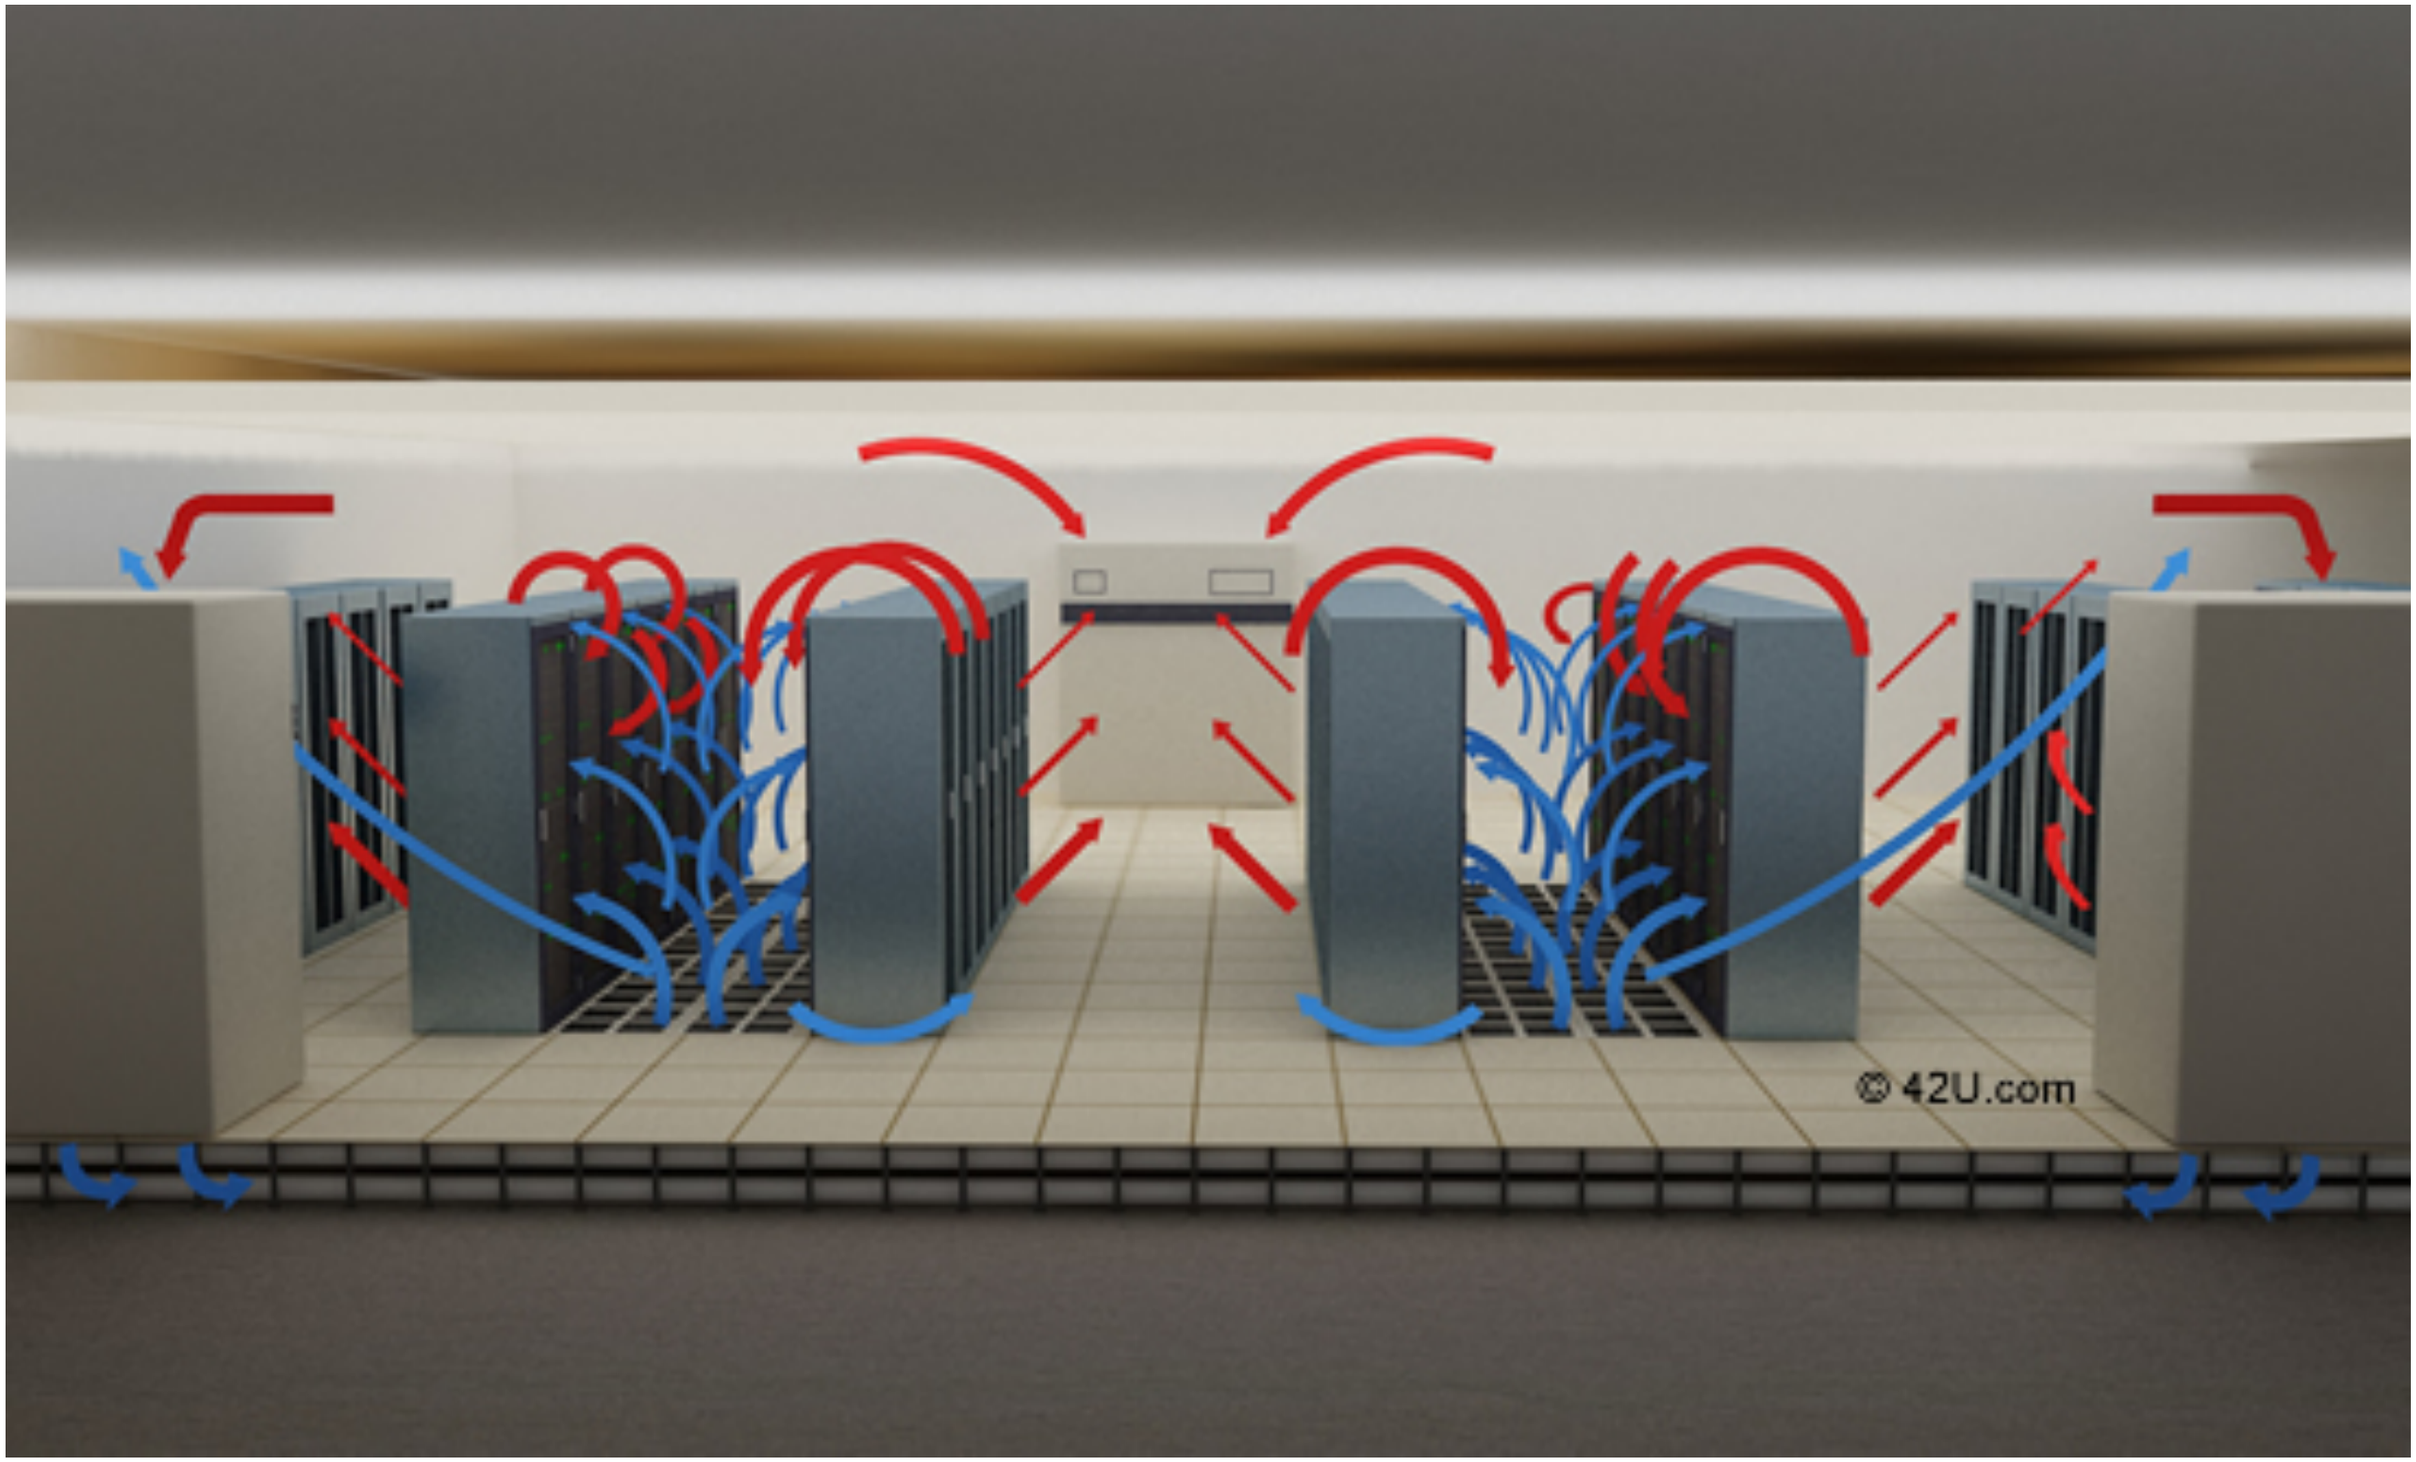

HVAC

Datacenter Server

- Similar to a simple desktop
- Different form factor (rack unit, blade)
- Different number of CPUs from desktop
- Different applications (less graphics)
- Headless, with KVM or console access
- Different peripherals (networking, shared storage)

Server vs Blade

Rack Server
2 units

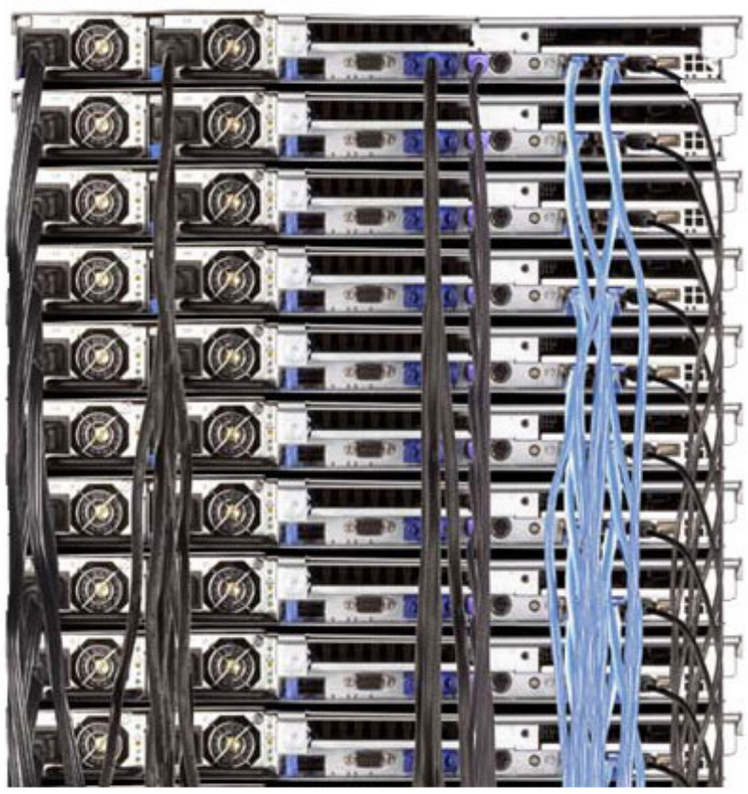

Blade Server
10 units

Datacenter Storage

- **Local storage**
  - e.g., SCSI
  - Only one host can access at a time
  - Block-oriented
- **Shared storage**
  - Multi-host access
  - Many protocols
  - SAN: block-oriented
  - NAS: file-oriented
  - iSCSI: block-oriented

Shared Storage

- **SAN (storage-attached network)**
  - Remote storage looks like a local SCSI block device
  - Use special fiber-optic cables for high-speed
  - Example: Fibre Channel
- **NAS: network-attached storage**
  - Server attached to disks
  - Remote hosts "mount" files on server: file-based access
  - Example: NFS, AFS
- **iSCSI**
  - Remote storage looks like a local SCSI block device
  - Uses TCP/IP as the transport mechanism.
  - Cheaper than SAN/FC (cabling, HBA cost)

Storage in a Datacenter

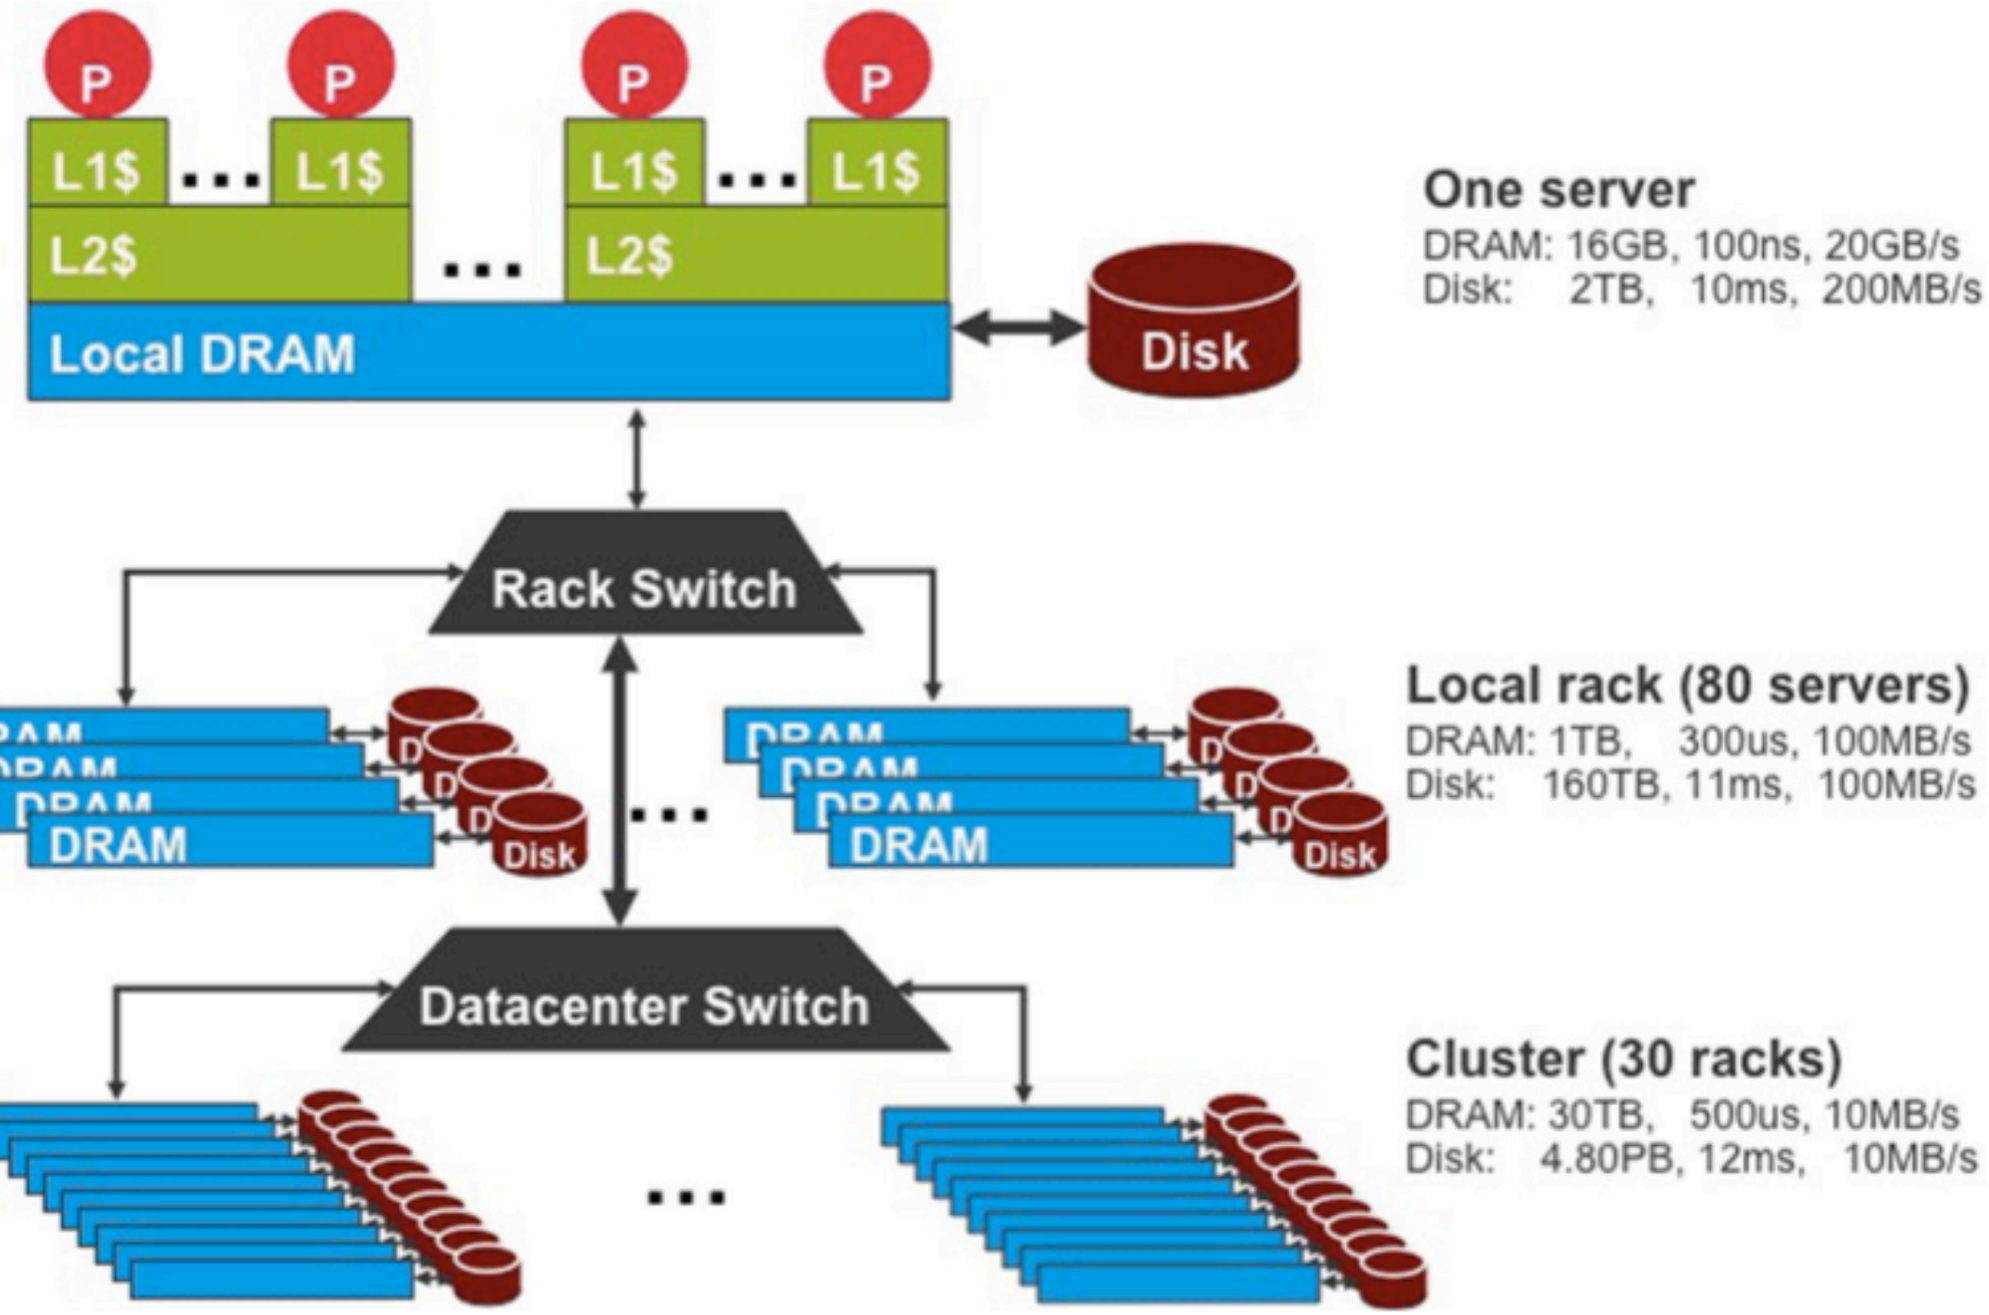

Switches and Cabling

- **Network Switches**
  - Connect servers to each other
  - Connects servers to NAS devices.
  - Connect switches to other switches (hierarchy)
  - Used for communication within a rack and between racks

- **Storage Switches**
  - Connects servers to their storage
    - Host to SAN devices
    - Host to iSCSI device.
  - Connects storage switches to other switches.
  - Usually use fiber-optic cabling.

- **Network**
  - Cabling
    - Fiber vs. Copper
  - Interfaces
    - Ethernet: 10/100 and Gigabit
    - Advanced: Myrinet, Infiniband, Quadrics

Power Distribution

- **Components/Issues**
  - PDU: Power distribution unit
  - Power inside a rack
  - Power across racks
  - Redundancy and surge protection (UPS)

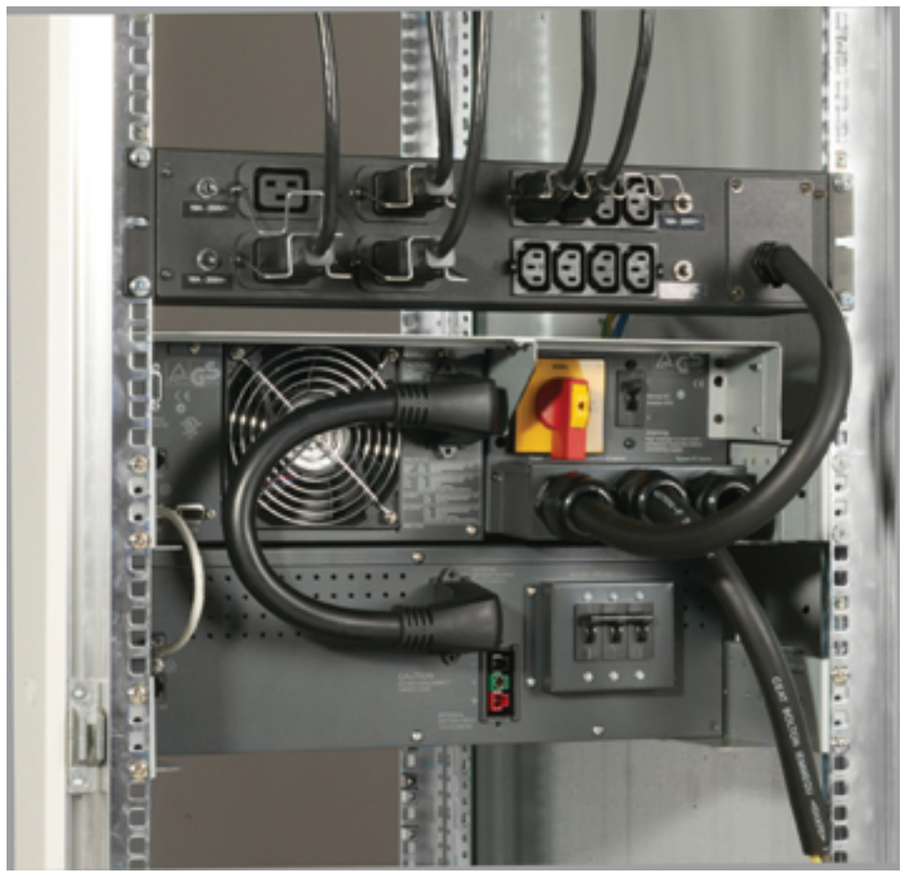

Energy Infrastructure

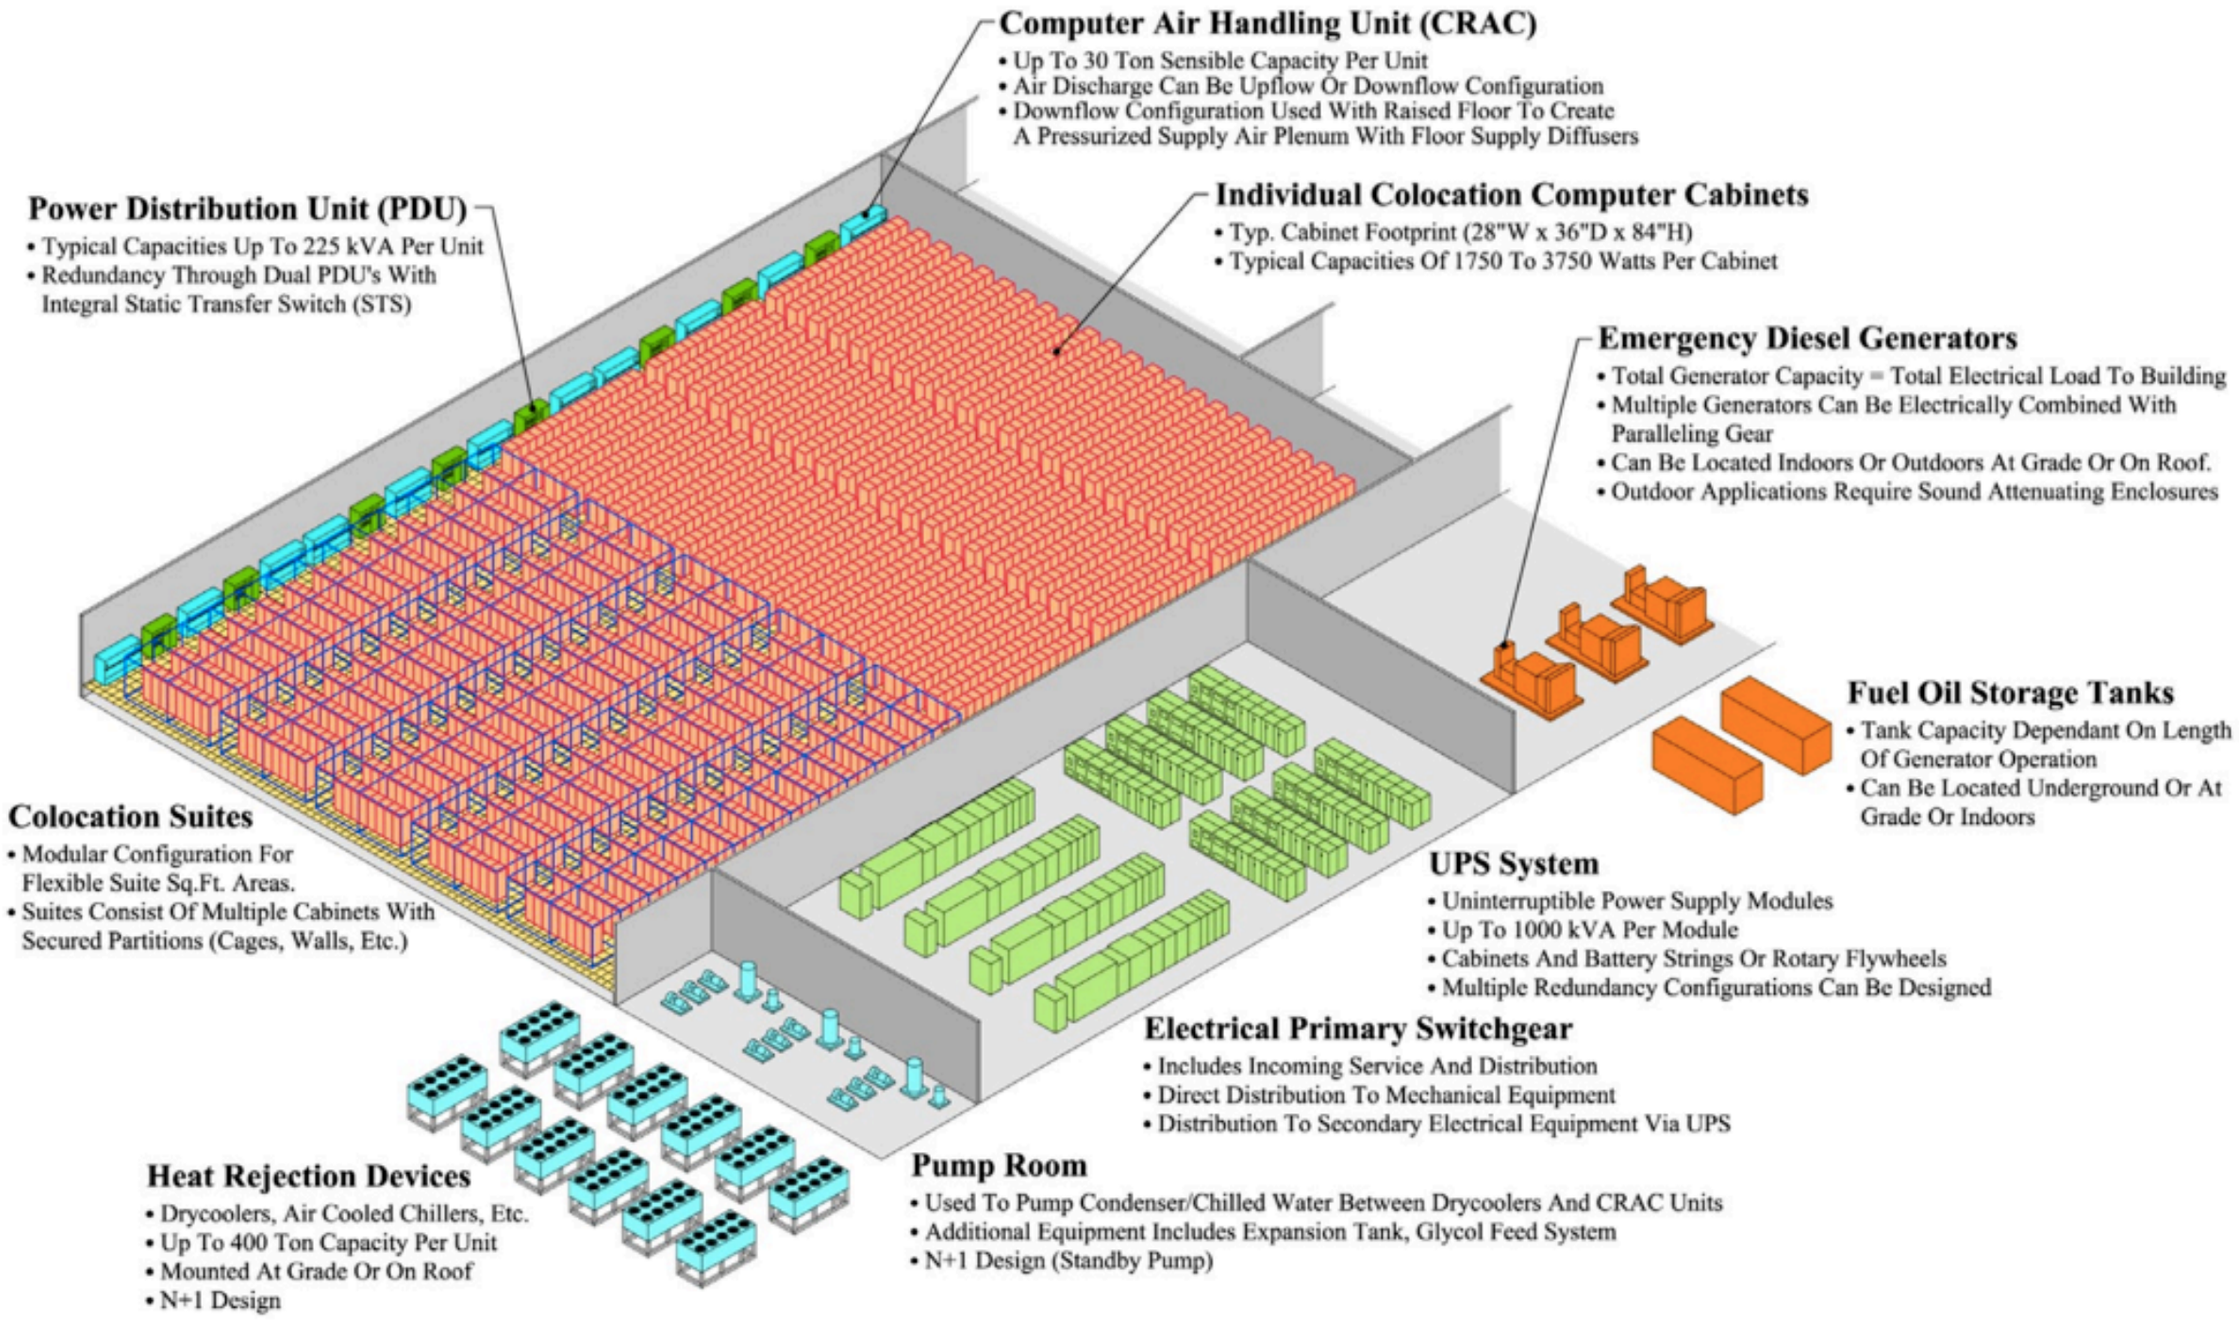

Energy Usage

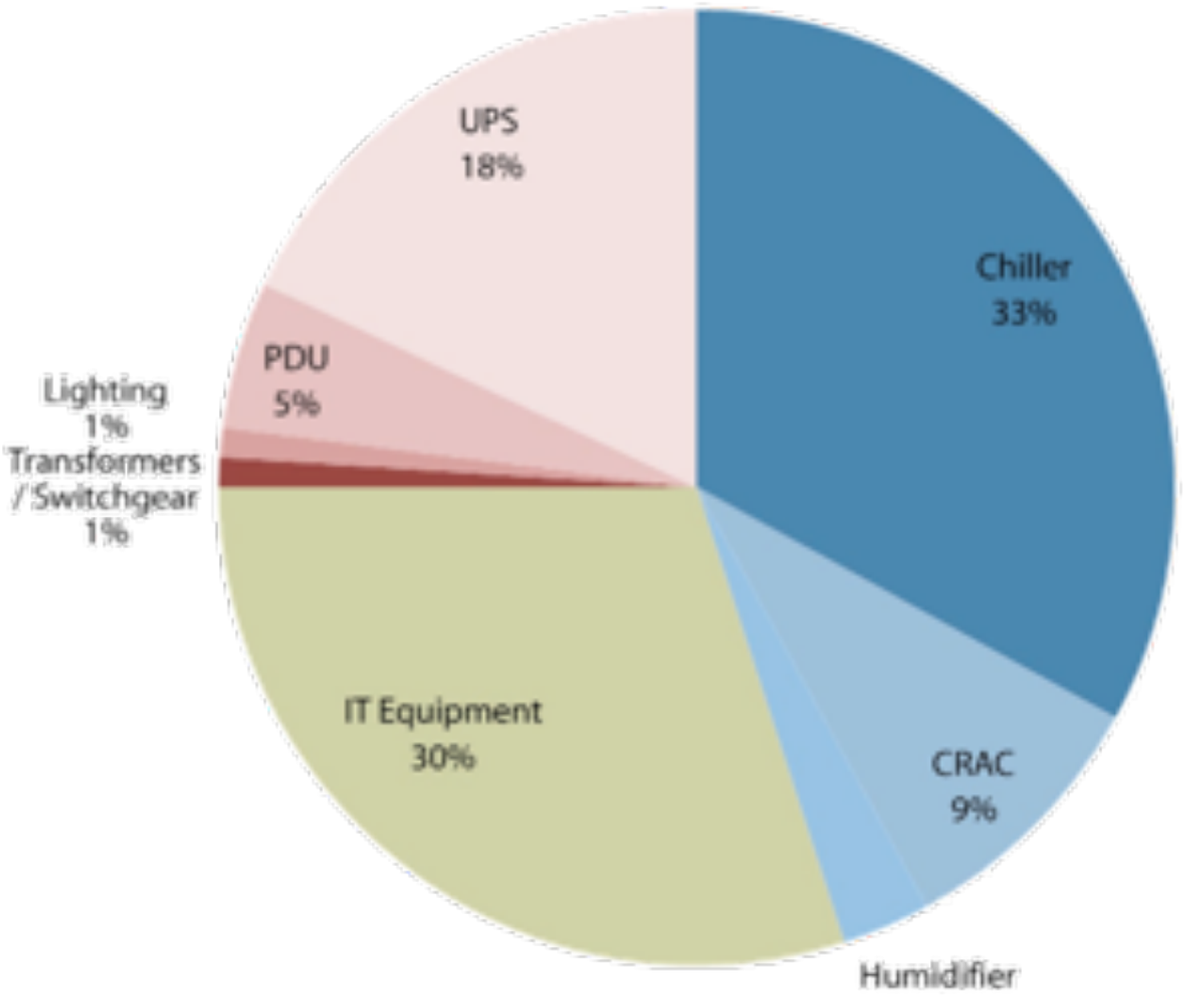

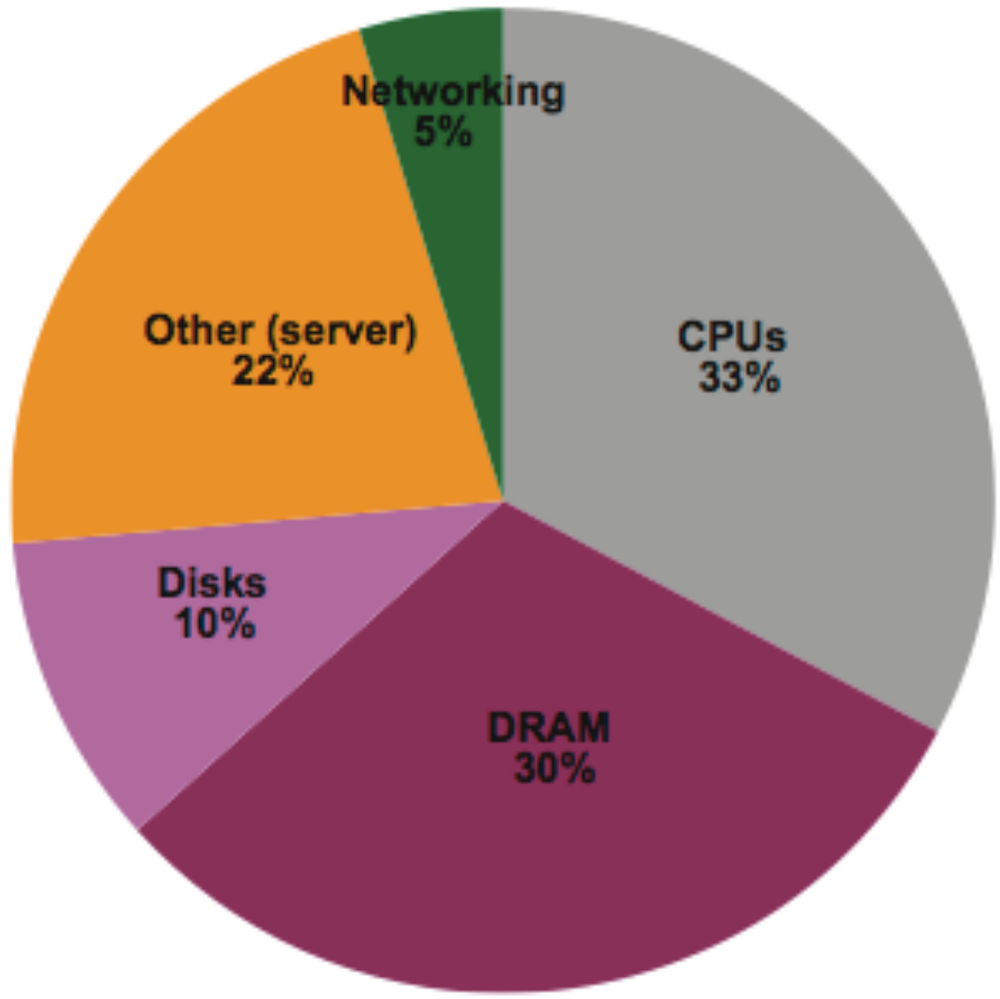

Energy Efficiency

DCPE

Data Center Performance Energy

$$\text{Efficiency} = \frac{\text{Computation}}{\text{Total Energy}} = \underbrace{\left(\frac{1}{\text{PUE}}\right)}_{\text{(a)}} \times \underbrace{\left(\frac{1}{\text{SPUE}}\right)}_{\text{(b)}} \times \underbrace{\left(\frac{\text{Computation}}{\text{Total Energy to Electronic Components}}\right)}_{\text{(c)}}$$

PUE = Power Usage Effectiveness

SPUE = Server Power Usage Effectiveness

Power Usage Effectiveness

- The ratio of total building power to IT power, i.e. the power consumed by the actual computing equipment (servers, network equipment, etc.)
  - Relates to the facility
- 85% of current datacenters were estimated in 2006 to have a PUE of greater than 3.0
  - the building's mechanical and electrical systems consume twice as much power as the actual computing load;
- only 5% have a PUE of 2.0

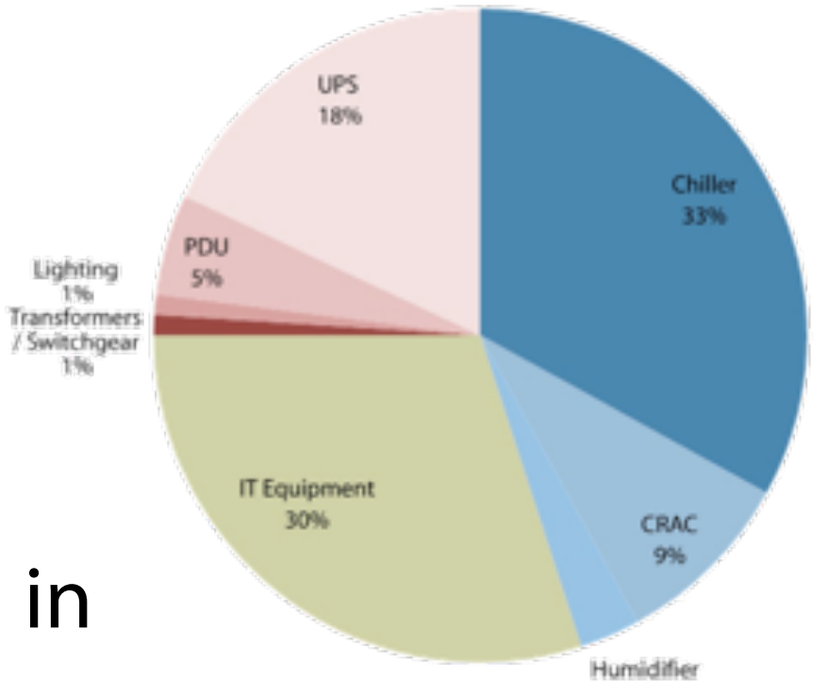

PUE of 24 Datacenters (2007)

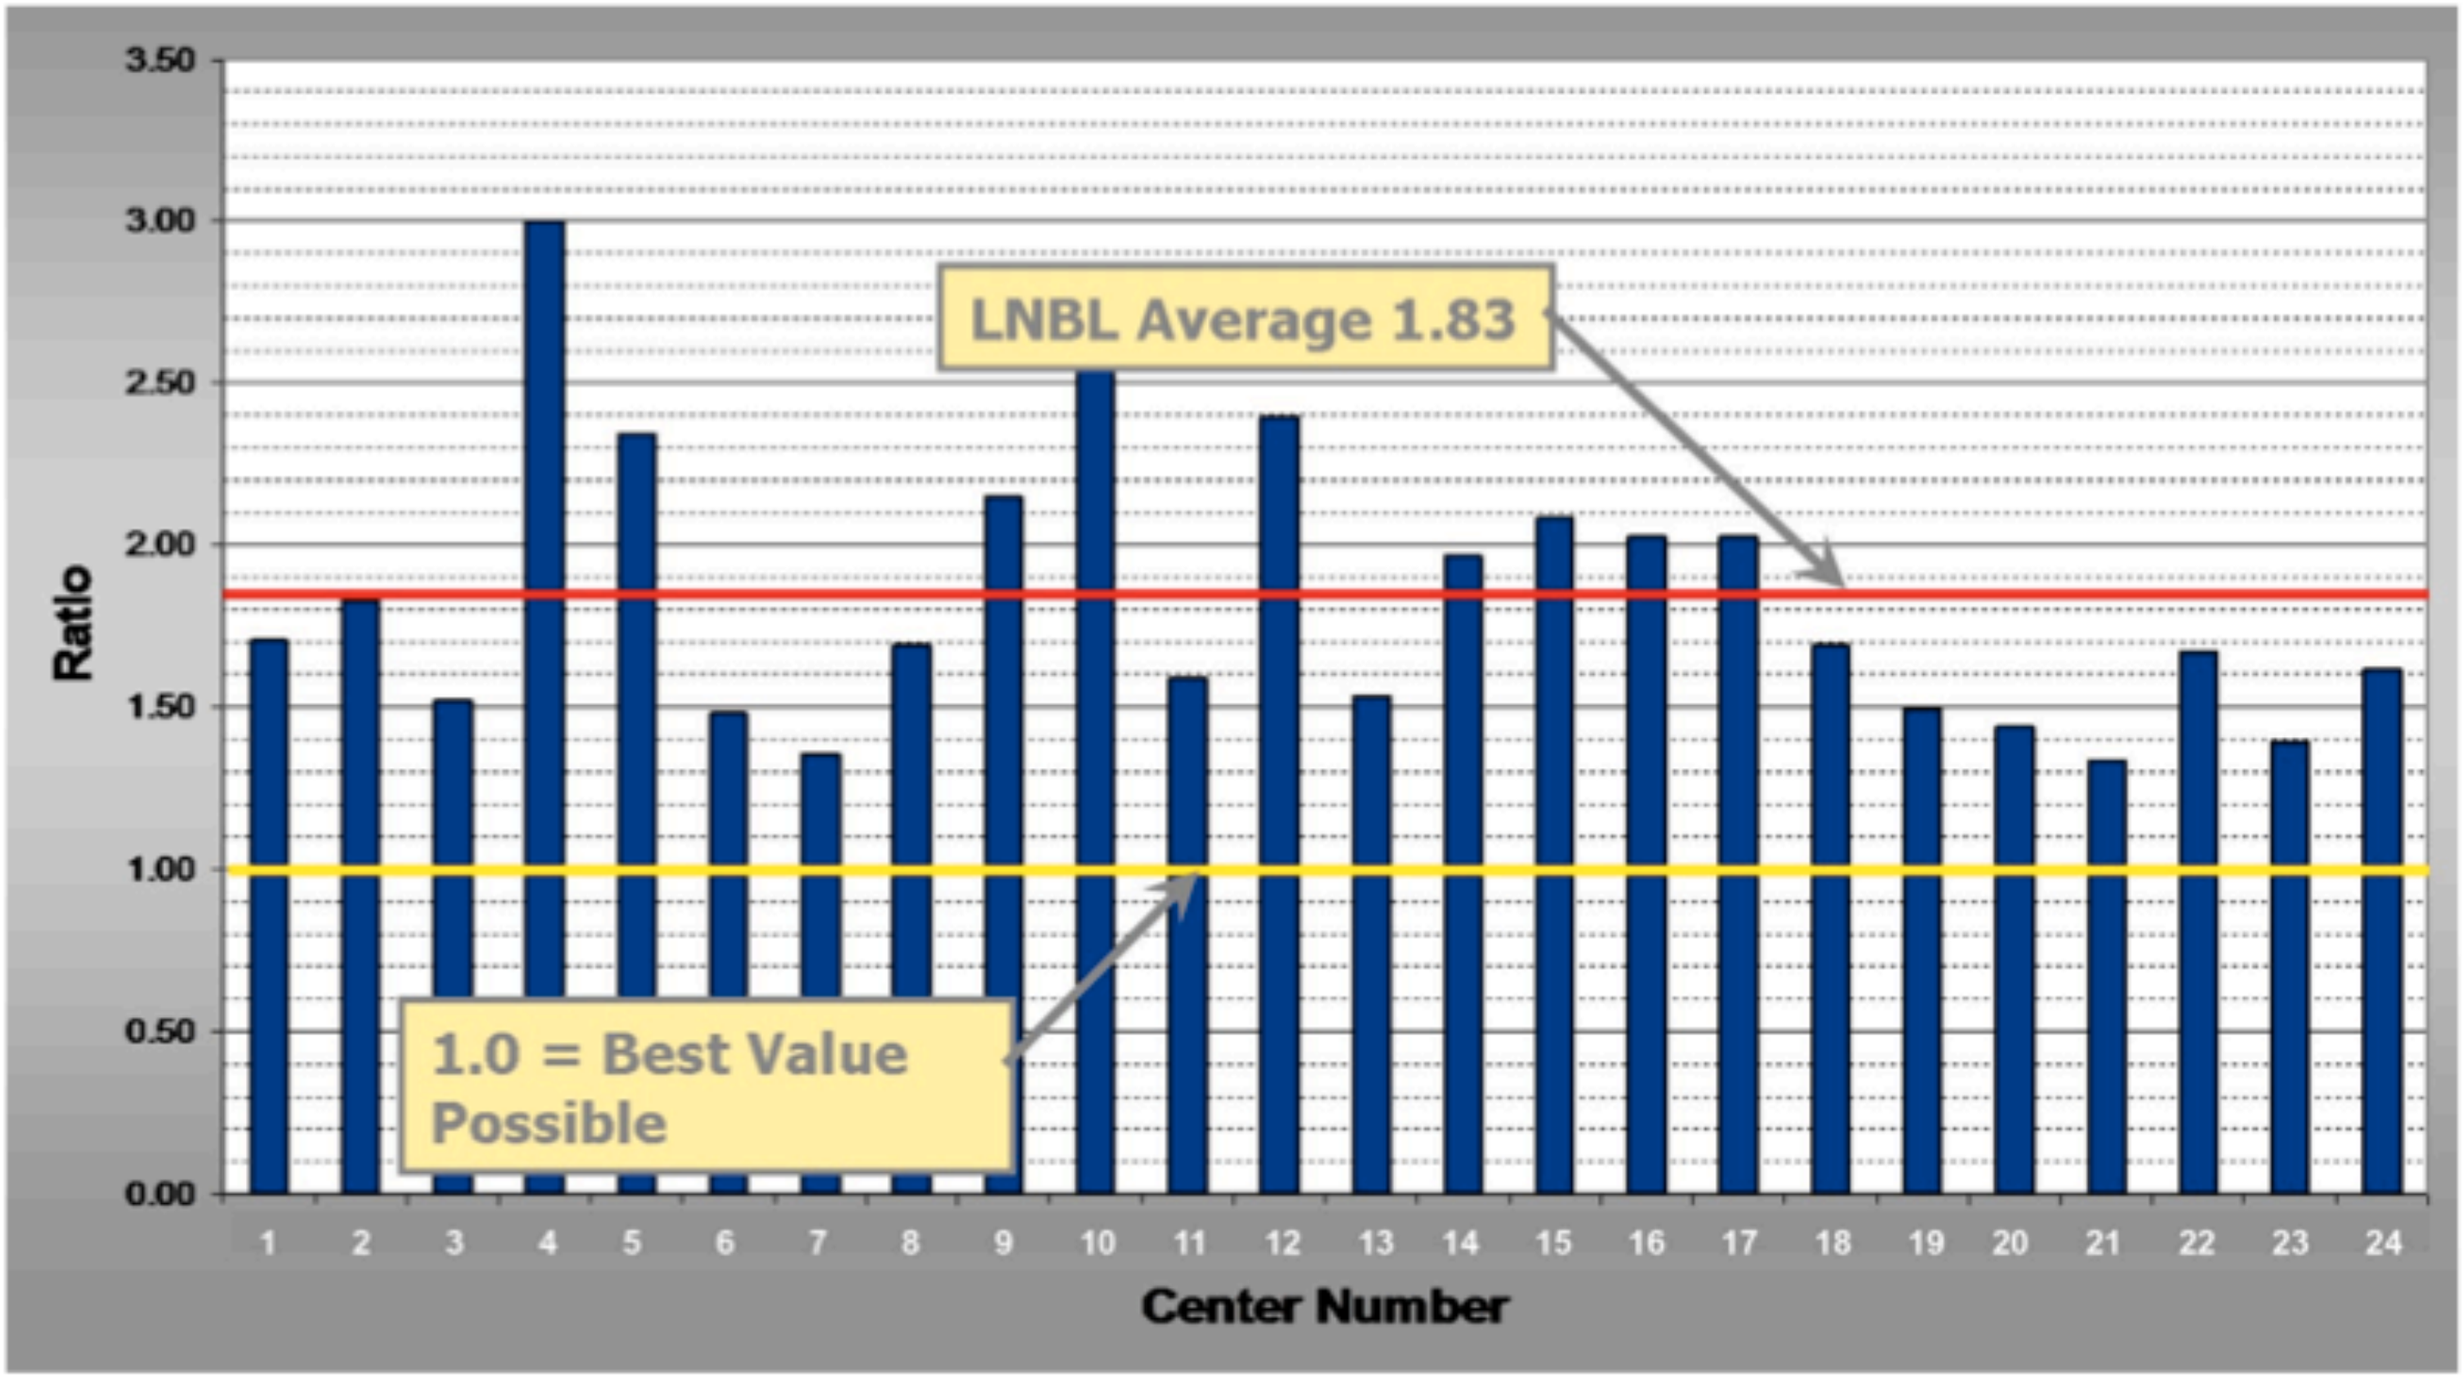

Improving PUE

- **Careful air flow handling**
  - The hot air exhausted by servers is not allowed to mix with cold air, and the path to the cooling coil is very short so that little energy is spent moving cold or hot air long distances.
- **Elevated cold aisle temperatures**
  - The cold aisle of the containers is kept at about 27°C rather than 18–20°C.
- **Use of free cooling**
  - Several cooling towers dissipate heat by evaporating water, greatly reducing the need to run chillers. In most moderate climates, cooling towers can eliminate the majority of chiller runtime. Google's datacenter in Belgium even eliminates chillers altogether, running on "free" cooling 100% of the time.
- **Per-server 12-V DC UPS**
  - Each server contains a mini-UPS, essentially a battery that floats on the DC side of the server's power supply and is 99.99% efficient. These per-server UPSs eliminate the need for a facility-wide UPS, increasing the efficiency of the overall power infrastructure from around 90% to near 99%.

Server Power Usage Effectiveness

- A server's energy conversion, ratio of total server input power to its useful power
- Useful power: power consumed by the electronic components directly involved in the computation: motherboard, disks, CPUs, DRAM, I/O cards, excluding all losses in power supplies, VRMs, and fans.
- SPUE ratios of 1.6–1.8 are common in today's servers; many power supplies are less than 80% efficient, and many motherboards use (voltage regulator modules) VRMs that are similarly inefficient, losing more than 30% of input power in electrical conversion losses.
- SPUE should be less than 1.2

Other Issues

- **Software Management**
- **Network Topology**
- **Status Monitoring**
- **Maintenance Costs**
- **Large Scale Management**
- **Power Consumption**

Reading Assignments

- Datacenter and Rack Standards
  - TIA-942 (Datacenter Standard)

    http://www.tiaonline.org/standards
  - EIA-310-d (Rack Standard):

    http://electronics.ihs.com/collections/abstracts/eia-310.htm

- Power Efficiency in Datacenters
  - Green Grid Metrics: Describing Datacenter Power Efficiency

    http://www.thegreengrid.org/~/media/WhitePapers/Green_Grid_Metrics_WP.pdf