Datacenter Design Issues

Topics

- Power
- Cooling
- Floor loading
- Wiring
- Density

Power

- Higher voltage = better efficiency
- Fewer conversions = better efficiency
- Power factor correction
- 3 phase power (polyphase)

Power factor

- The ratio of real power to apparent power.
  - Real power – capacity for performing work in a unit of time
  - Apparent power – current multiplied by voltage – can appear to be higher than real power due to inefficiencies, distortions and loading effects
  - A typical modern PC has a powerfactor of 90% or higher. Higher is better.
- Also measured as VA (UPS requirements) vs Watts (heat dissipation and power utilization)

High voltage power distribution

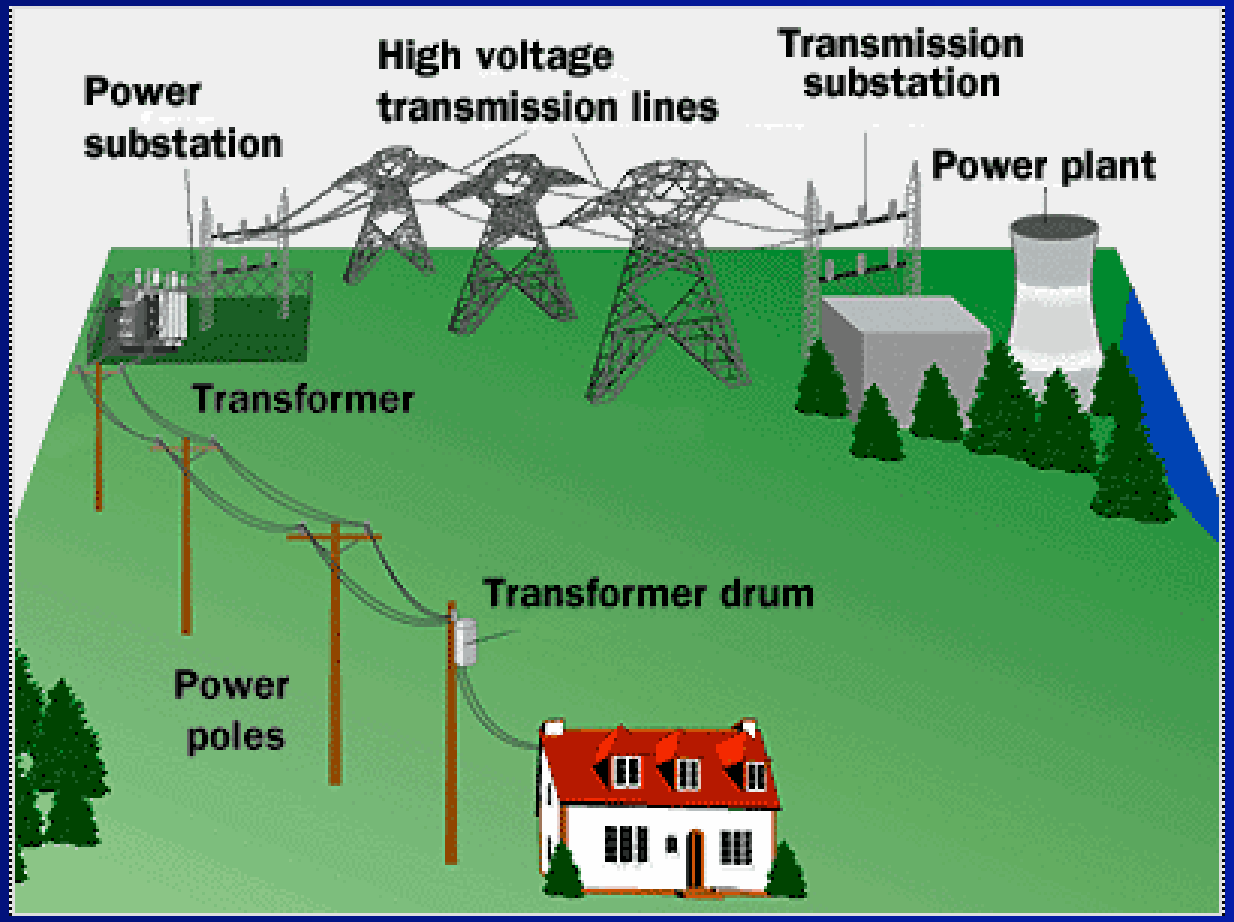

Stepping down

- 1MW =~ 1000 homes
- 1 utility generator =~ 4MW
- 1 medium-large datacenter =~ 10MW
- Transmission -> 23KV 3ph -> 12.5KV 3ph-> 480V 3ph -> UPS -> 480V 1-3 ph -> 208V 3ph -> 120V single phase -> 12V/5V ATX
- Each conversion uses 1-2%
  - The PS in your computer can be as much as 25% efficiency drop. (old ones suck worse)
    - Redudant PSUs are even worse
  - All loss is dissipated as heat which you have to remove, using more electricity

Switched mode power supply (SMPS)

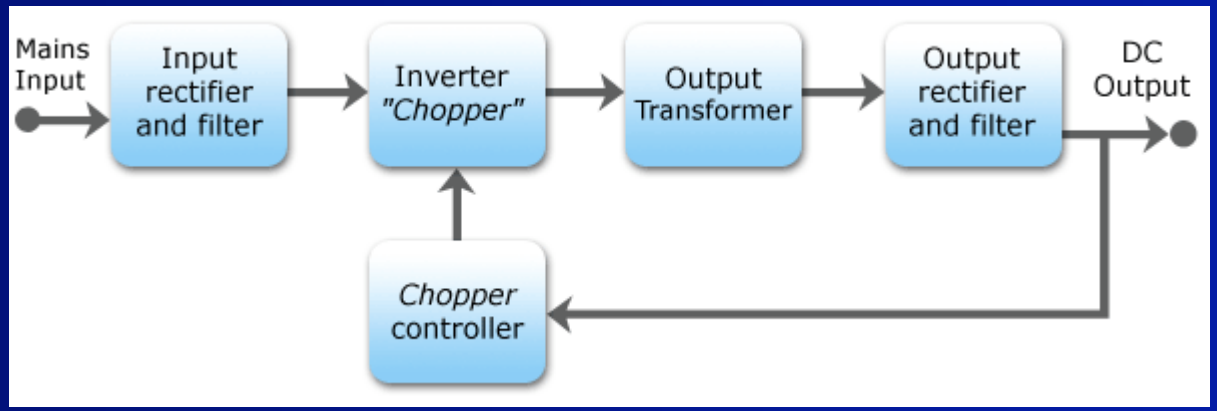

Input sampling happens at waveform peaks, shearing off the top
Of the sine wave and distorting the waveform with harmonics

sampling

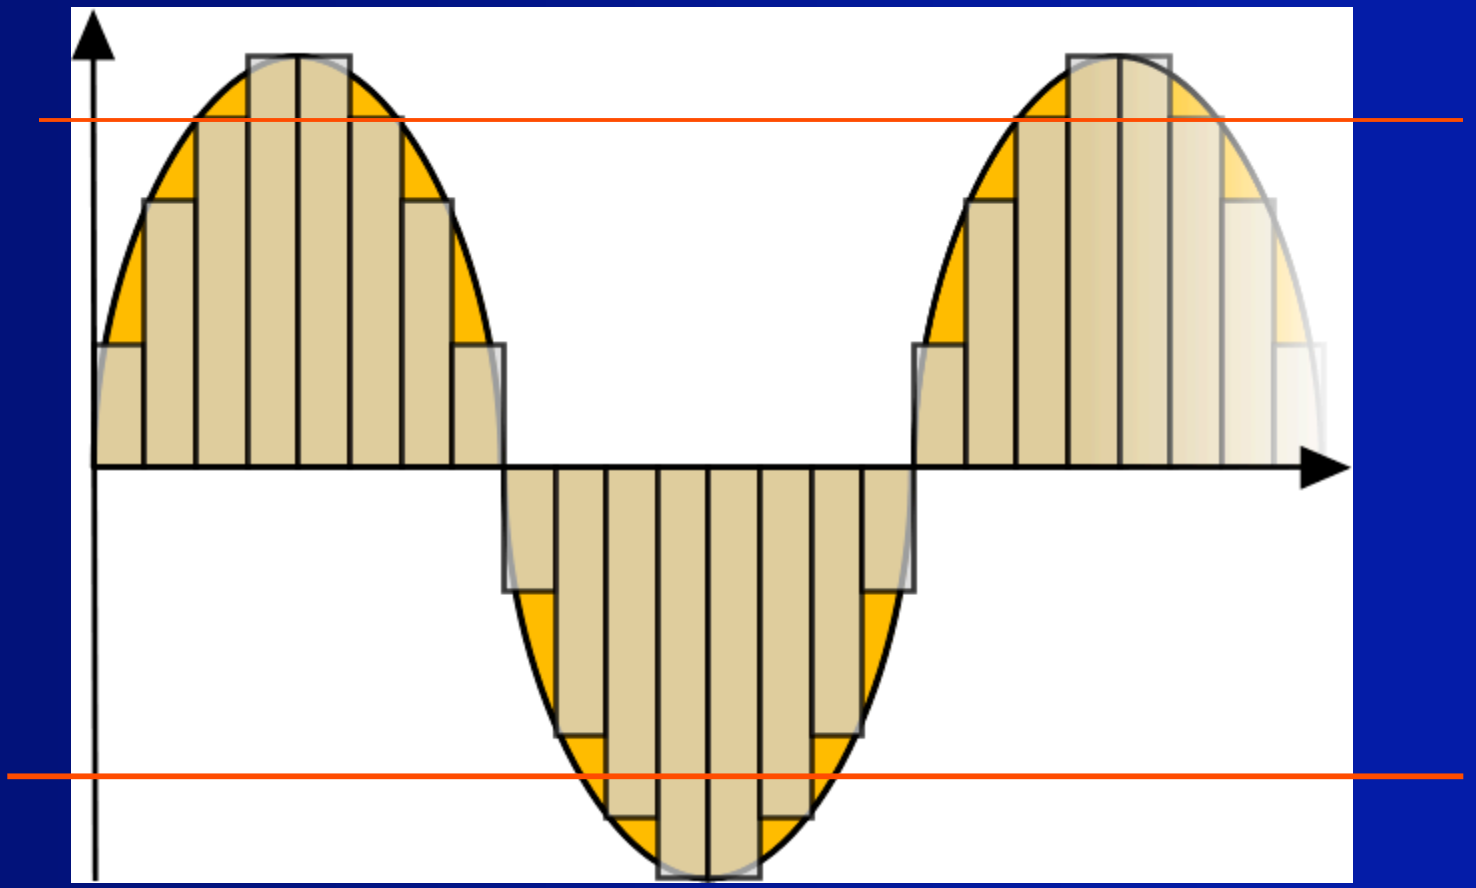

PSU outputs

- 12V -> external drive
- 12V -> DC/DC voltage converter
  - 1.5/2.5V -> memory controller
  - 0.8-1.8V -> microprocessor
- 5V -> internal drive
- 3.3V -> SDRAM
- 3.3V -> I/O (e.g. PCI)
- 3.3V -> graphics controller

PSU efficiency as function of load

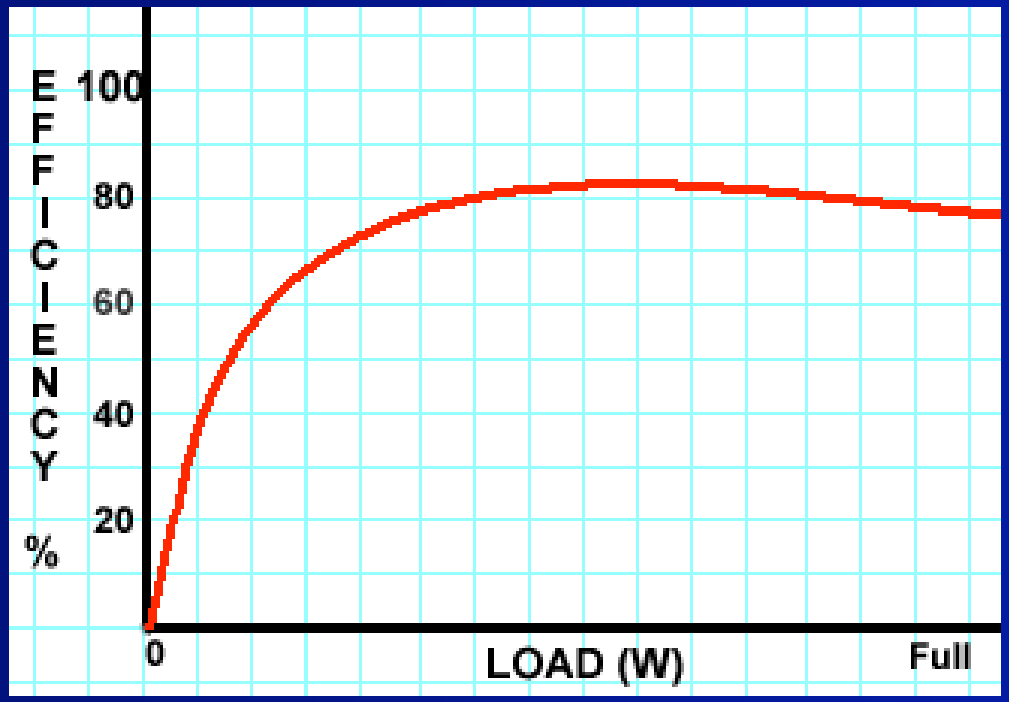

"Computer power supplies are generally about 70–75% efficient; to produce 75W of DC output they require 100W of AC input and dissipate the remaining 25W in heat. " - wikipedia

3 phase motor

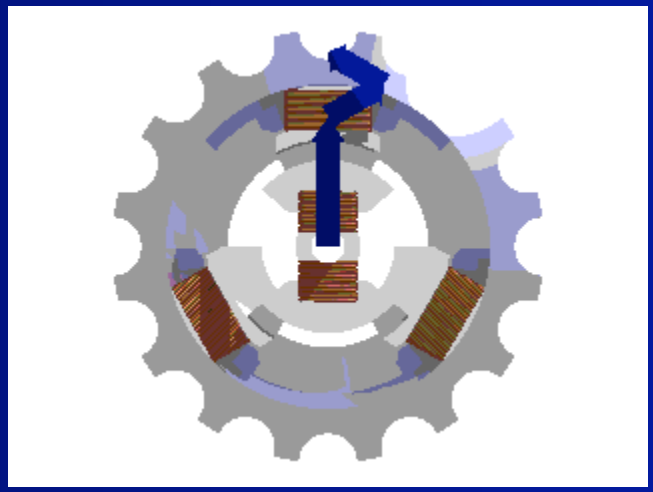

A 3 wire, three-phase system can provide 173% more power than the two conductors of a single-phase system.

3 phase power

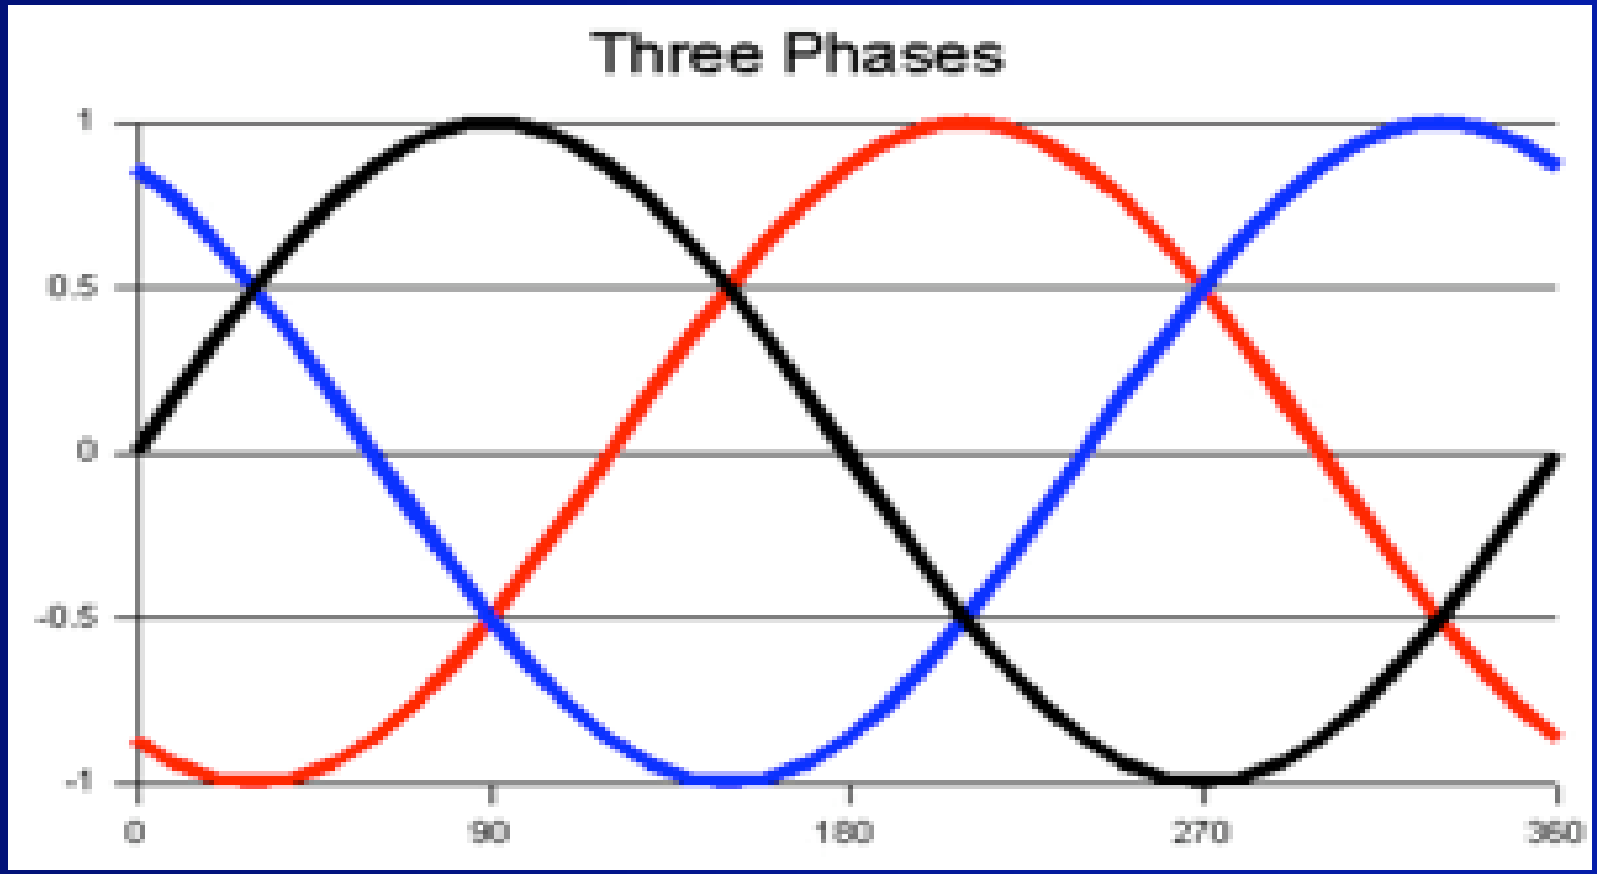

Simplified 4:2 3-phase transformer

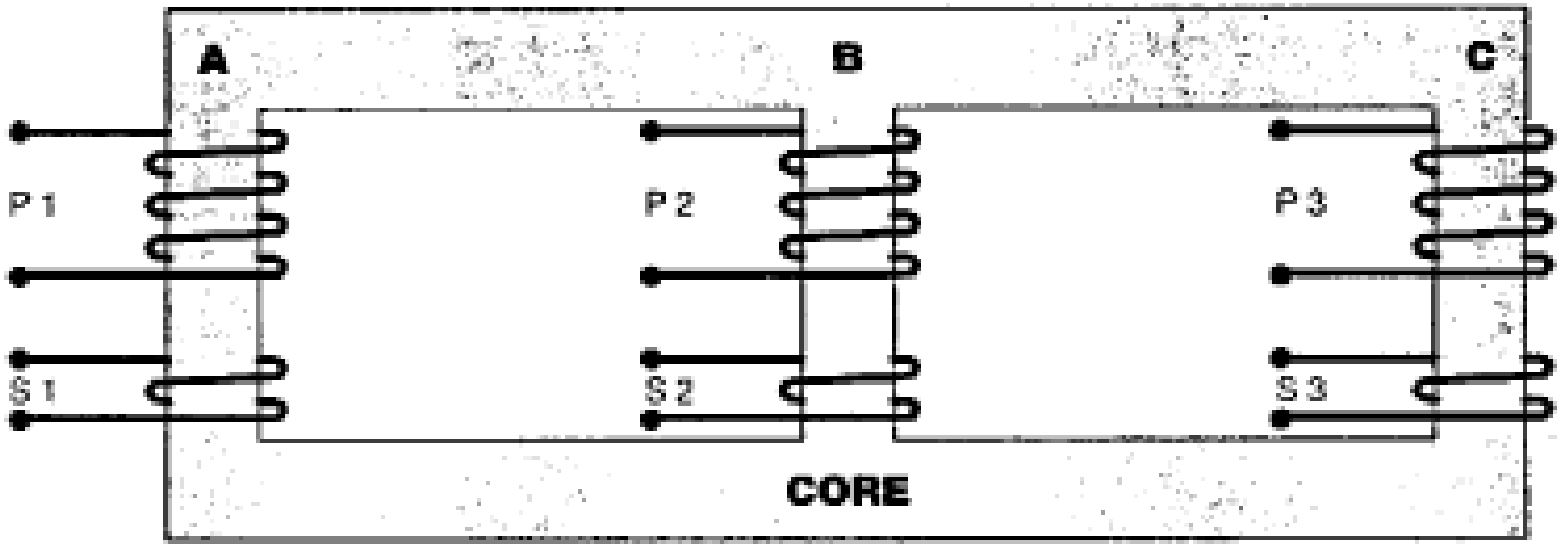

WYE vs Delta

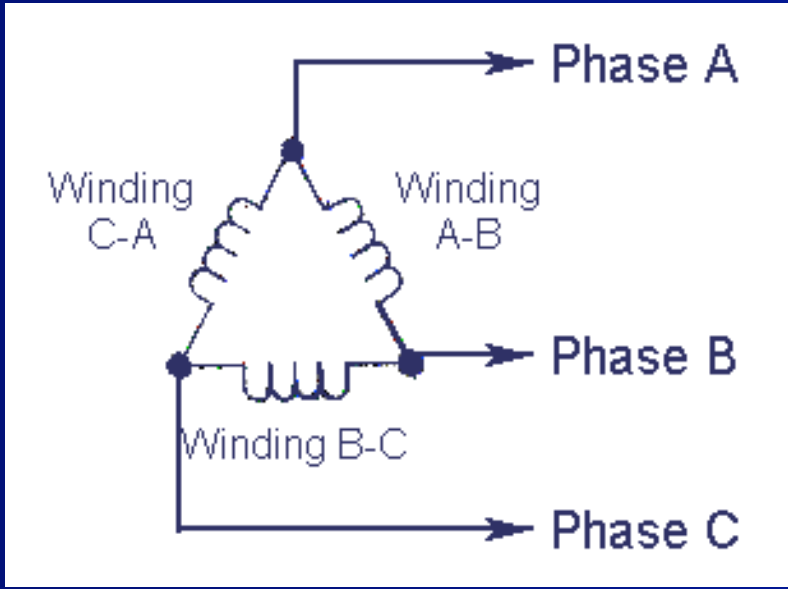

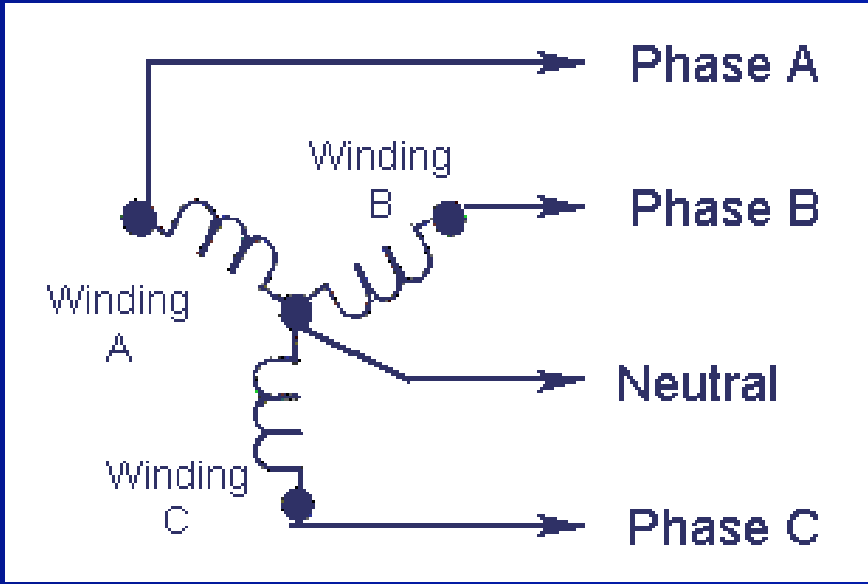

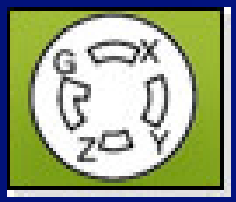

L16-30 (480)

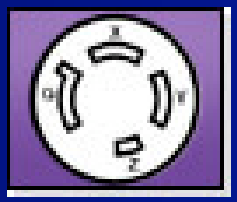

L16-20 (480)

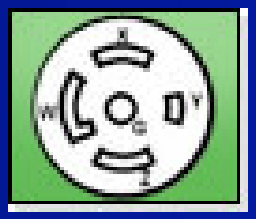

L21-30 (208/120)

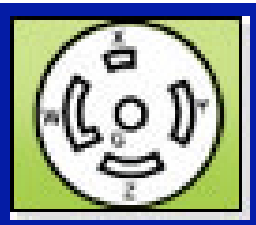

L22-30 (480/277)

408V to 208V via 120V

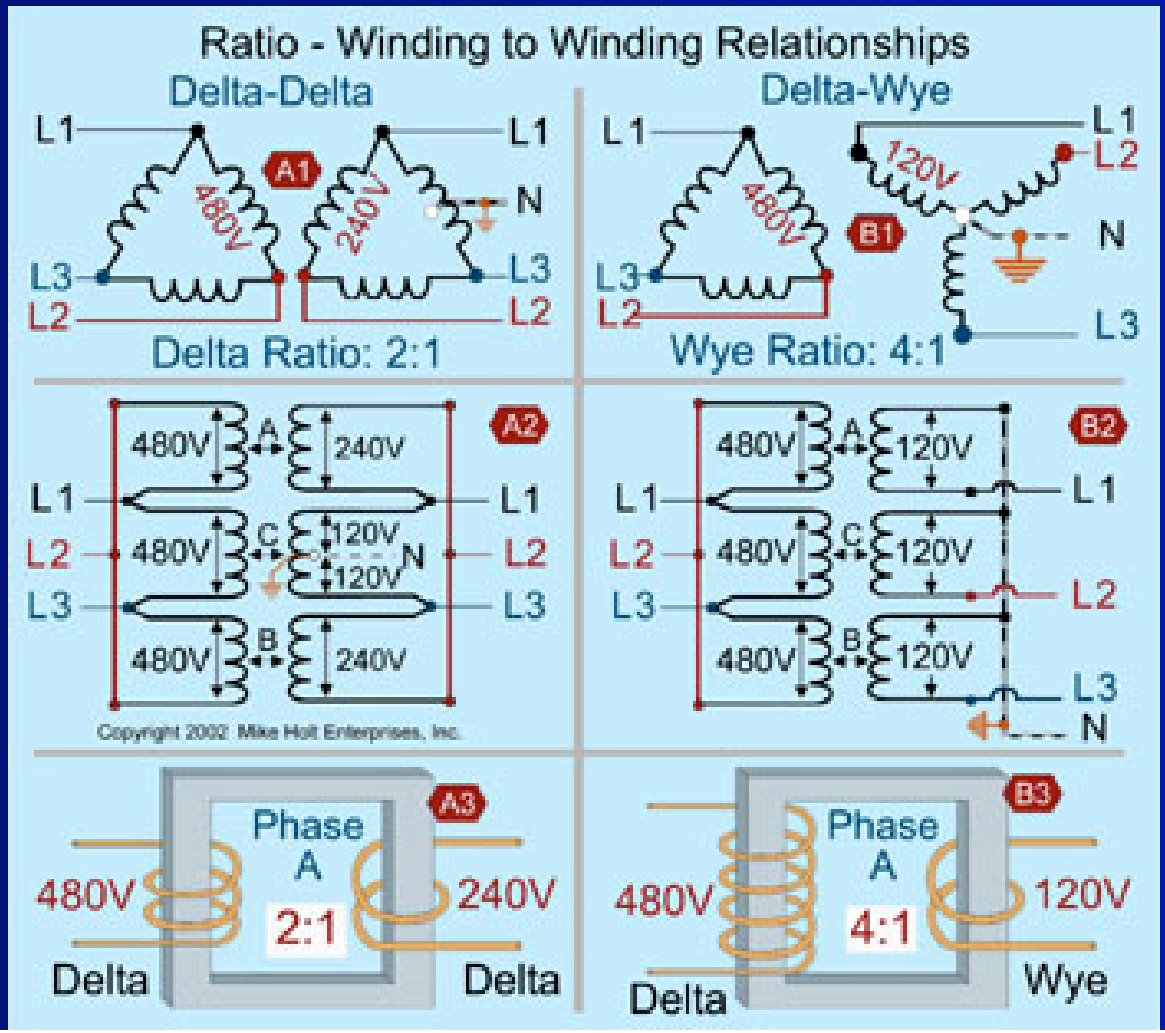

WYE

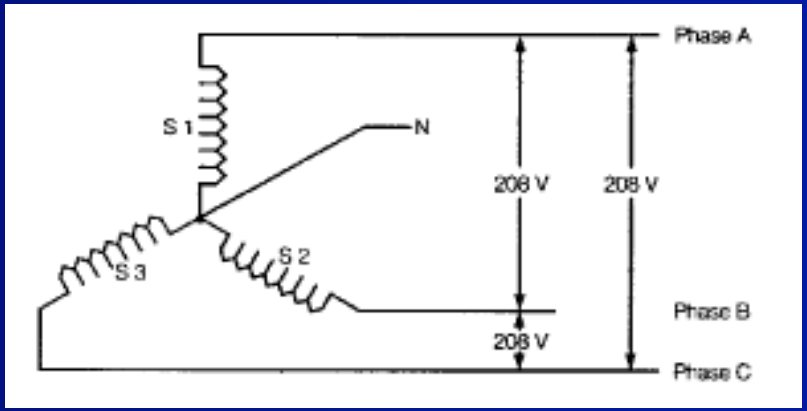
Phase A
S 1
N
208 V
208 V
S 3
S 2
Phase B
208 V
Phase C

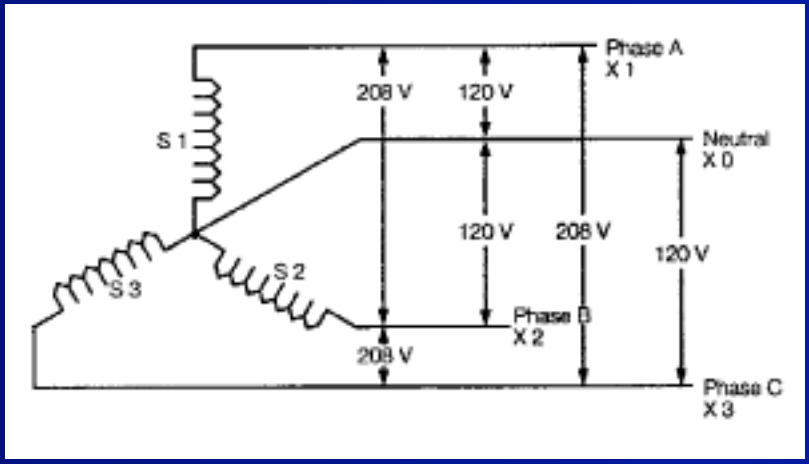
Phase A
X 1
208 V
120 V
S 1
Neutral
X 0
120 V
208 V
120 V
S 3
S 2
Phase B
X 2
208 V
Phase C
X 3

Busway

PDU

- The two most popular vendors of smart PDU gear are ServerTech and APC
  - monitoring, individual outlet control, traps, etc.

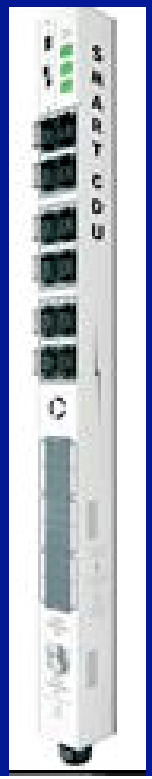

AC vs DC

- DC will use the same thickness of wire approximately 10% more efficiently than AC.
  - (higher voltages than you have in your datacenter)
  - UL listing?

Cooling

- Raising your datacenter 1 degree F can save you about 1 percent in operating costs.
- Companies like Intel are pioneering centers with 6 foot raised floors and cabinets with up to 40KW of heat load. (that's 40 1U machines that consume about 1KW each of electricity)
  - The average Dell PowerEdge 1950 with dual 120W Intel quadcore processors uses about 450W at maximum load (depending on memory config, cards, etc)
  - A Supermicro Twin with 2 machines each with two dualcore Intel 120W processors and 32G total ram uses about 900W at full load
- On average, a 10% reduction in power consumption or conversion loss results in 20% in savings (also 10% in cooling reduction)
- Water has 1000-2500 times the heat carrying capacity of air

Cooling Math

- 1 ton = 12000 BTU/hr
- 1 KWH = 3413 BTU/hr = 1.341 horsepower
- 1 Therm = 100,000 BTU
- 1 MMBTU = 1,000,000 BTU
- 1 watt = 1 Joule/sec
- 1 ton = 3.515 KW

Delta-T

- System temperature change
- Cold water and hot air = efficient cooling
  - E.g. air temperature near a large body of water
- Hotter hot aisle and colder water = higher deltaT and more efficient energy usage.
- Mixing hot and cold air results in extra work, low delta-T, and extra energy usage.

Racks

Liebert XDK
(17-25/25-35KW)
Aka Knuerr CoolTherm

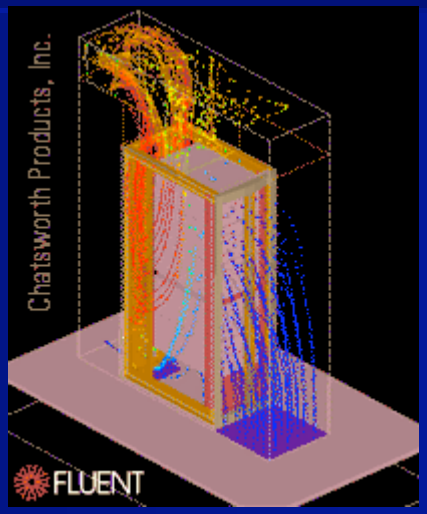

Chatsworth passive
(4-8KW)

Sanmina-sci ecobay
(25KW)

APC rear door (to 25KW non-redundant)

+ comfortable environment
+/- medium density
- Matched blower pressurization
+/- moderate electrical usage
+ high delta-T

More cabinet stuff

IBM rear door
(coolBlue)

Rittal LCP (30)
(for 42U™ racks)

HP (freaking huge)
Freon loop!

+ density
+ noise reduction
- cost
- water out near computers

Supplemental cooling

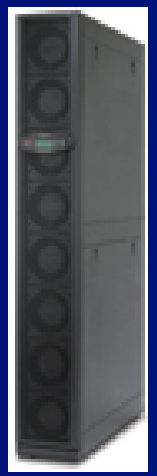

RC – to 30KW

RP 0-70KW

Inline SC – up to 7KW

APC inline cooling accessories

Hot aisle containment

Sandia cold air containment

Liebert - refrigerant/water heat exchangers
XDV - to 10KW
XDO – to 20KW

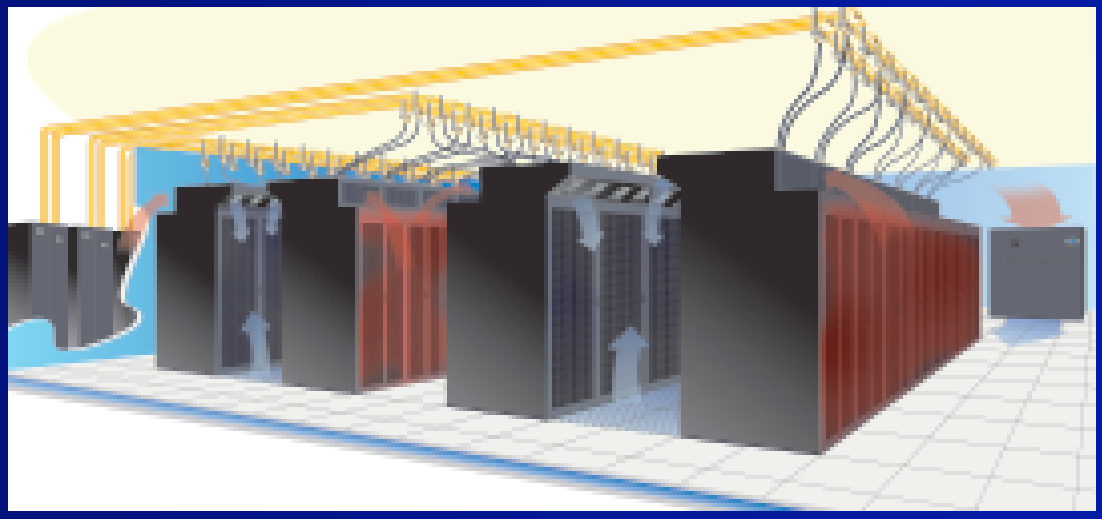

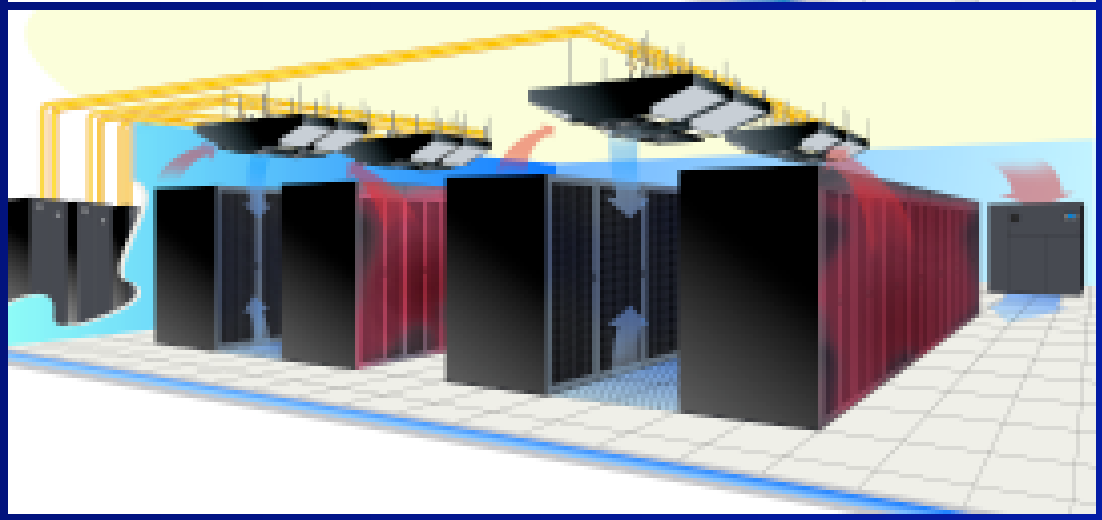

Flooring

- Rolling vs Static loading
- Cement/epoxy
  - + high weight load
    - Usually – consult your building engineer, especially for multi-floor
  - – Bad for chilled water
    - Possible high density implications
- Raised floor
  - – High load =~ $25/sqft
  - + place for piping of chilled water
  - +/- cabling?
  - - Mixing power and water may lead to unhappiness
  - + forced air distribution
    - high floor = high expense but builtin plenum)
  - - What about your elevator? (if necessary)
  - - Pushing heavy stuff up ramps sucks – lifts?
  - - whiskers

wiring

- Under raised floor? Overhead?
  - Tracing
  - Distance limitations
  - Rats nests
- BICSI standards
  - (Building Industry Consulting Service International)
- TIA standards

Cable management

- Copper wire ducts
- Ladders
- Fiber ducting
  - Bend radius
- Vertical vs horizontal

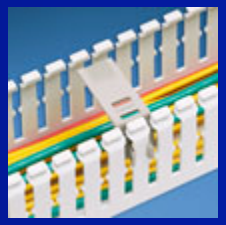

Density

- How many KW per rack? (conventional = 8KW)
- Hot aisle/cold aisle? Forced air return? Air mush? Inline row cooling? Centralized cooling?
  - Fluid dynamics
- Interconnect limitations? (Infiniband, Myrinet, etc)
- PDU limitations
- Power plant limitations
- Chilled water plant limitations

More density

- Console serving?
  - (do you need one?)
  - Cables coming out the wazoo
- Blade servers
  - Power per compute unit?
  + Cabling advantages
- Compute units per $$
  - T2000? Sicortex?
- Cpu clock speeds
  - AMD vs Harpertown (54xx) vs Clovertown (53xx)

Resources

- Uptime Institute (uptimeinstitute.org)
- Datacenter dynamics presentation (LBL/PGE) - **http://preview.tinyurl.com/2vlopz**
- **BICSI** datacenter standards & design documents ($$)
  - www.bicsi.org
  - ITSIM – Information transport systems installation manual
    - (nee) Telecommunications cabling installation manual (TCIM)
    - Standards, codes, tables, architecture, planning, termination, test
    - 888 pages
- TIA 942
  - Telecommunications infrastructure standard for datacenters
  - Fire, room layout, environment, change control, safety, security, IT infrastructure
- NFPA –70 National Electric Code
- http://hardware.slashdot.org/article.pl?sid=06/09/26/2039213
  - Google calls for simplified power supplies