Datacenter Efficiency Trends

Cary Roberts

Tellme, a Microsoft Subsidiary

Power Slashing Tools

P3 Kill-A-WATT

- Measures min, max, and instantaneous voltage and current draw
- Calculates power (watts), kilowatt hours used, and power cost
- Under $25

Power Slashing Tools

EE Devices Watts Up?

- Measures min, max, and instantaneous voltage and current draw
- Calculates power (watts), kilowatt hours used, and power cost
- Serial port for exporting measured data
- $130

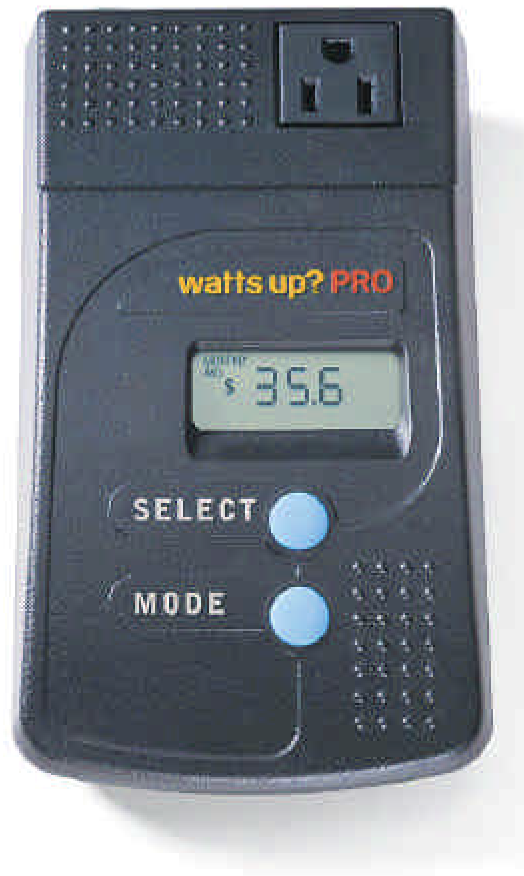

Datacenter Inefficiencies

UPS and Power Systems

- Conversion losses
- Partial loads significantly decrease efficiency of older UPS designs
- Power distribution and transformer losses
- Redundancy strategies strand power

Datacenter Inefficiencies

Cooling Systems

- Older systems, particularly DX (dry cool), are very inefficient
- Chillers without VFDs have large power draw step function
- Pump losses
- CRAC unit fan motor and impeller inefficiencies
- Hot aisle / cold aisle air mixing issues
- Recirculation and bypass

Datacenter Inefficiencies

Servers

- Power supplies
- DC-DC conversion for CPU and RAM
- CPU gate leakage
- Fan control and impeller design
- Inadequate power capping capability

Servers - Power Supplies

- Typical power supply is 65-75% efficient
- Manufacturers finally delivering 80%+ efficient power supplies
- Some manufactures delivering 90-93% efficient designs, but single supply rail with point of load DC-DC conversion – does not account for DC-DC conversion losses
- 80 PLUS Program server test protocol still in draft

Servers - Power Supplies

Example: Dual quad core Intel w/8GB RAM and 4 SATA disks

- 255 watt DC load when heavily utilized
- 300 watt AC load with 85% efficient supply
- 392 watt AC load with 65% efficient supply

30% more servers per rack!

Servers – DC-DC Conversion

- System board manufacturers designing DC-DC conversion based on cost, not efficiency
- Extremely wide range of available CPUs makes optimizing for a particular power draw difficult
- DC-DC converter is inefficient when at low or no load (CPU idle)
- Work is being done on ultra high frequency switching converters fabbed on 90-180nm CMOS processes

Servers – Fan Technology and Speed Control

- Traditional axial fan designs, especially 40mm units for 1U servers, achieve poor efficiency
- New designs such as ducted fans, PAX Scientific, and backward curved centrifugal impellers can reduce the percentage of server energy consumed by fans
- Intelligent temperature monitoring with feedback from CPUs, disks, power supply, and inlet/outlet temperatures can more accurately control fan speed and delta T

Servers – Fan Technology and Speed Control

Servers – Fan Technology and Speed Control

Example: Shared fan design – four large fans provide air flow for 48 servers

- Removed two 18 watt 80mm fans per server
- Shared fan assemblies with 110 watt backward curved impeller reduce per server fan load to under 10 watts
- Fan failure is recoverable by increasing speed of remaining fans
- Lower cost and complexity of feedback and control system

Servers – Power Capping

- Peaks in server power draw require over-provisioning of power for normal loads
- Power reserved for peaks is “stranded” at the circuit, rack, and PDU level
- Capping can be accomplished by shedding non-critical application load
- Capping can be accomplished at the CPU - IBM Power6, AMD Barcelona
- Non-utilized RAM can be “put to sleep”

Cooling Systems – Chiller Plant

- Variable frequency drives dramatically reduce power consumption for partial loads, up to 30%
- Magnetic bearings decrease friction and reduce chiller motor power draw
- Cold water and ice storage can be used for load shifting or for redundancy

Cooling Systems – Air Handling

- Air mixing significantly contributes to elevated rack inlet temperatures
- Controlling air mixing and eliminating bypass can reduce CRAC fan energy up to 75% and allow the chilled water plant temperature to be raised 15F or more
- Majority of CRAC units using outdated fan impeller technology and no VFDs
- Localized heat exchangers can reduce fan energy consumption

Cooling Systems – Air Side Economizers

- Air side economization is the intake of cooler outside air and exhaust of heated air
- Benefits
  - Allows for “free” cooling 65% or more of the year in some climates
- Disadvantages
  - Does not integrate well with traditional underfloor cooling in datacenters
  - Humidity control issues
  - Requires more aggressive filter maintenance

Cooling Systems – Localized Cooling Solutions

Supplemental

- Rejects only partial heat load
- Used for spot cooling
- Usually affixed to rear of rack
- Most commonly water cooled
- Downsides:
  - Water on datacenter floor
  - Flexible hoses on door units failure prone
  - Higher pump losses than traditional CRAC

Cooling Systems – Localized Cooling Solutions

Primary, water cooled

- Rejects full heat load
- Used for spot cooling
- Can be affixed to rear of rack
- Can be mounted between racks
- Custom rack base unit available
- Downsides:
  - Water on datacenter floor
  - Flexible hoses on door units failure prone
  - Higher pump losses than traditional CRAC

Cooling Systems – Localized Cooling Solutions

Primary, refrigerant based

- Rejects full heat load
- Used for spot cooling
- Can be mounted between racks
- Can be mounted above racks
- Downsides:
  - Limited vendor choice
  - Requires floor space for refrigerant to water heat exchanger and pump unit

Cooling Systems - Liebert XD Solution

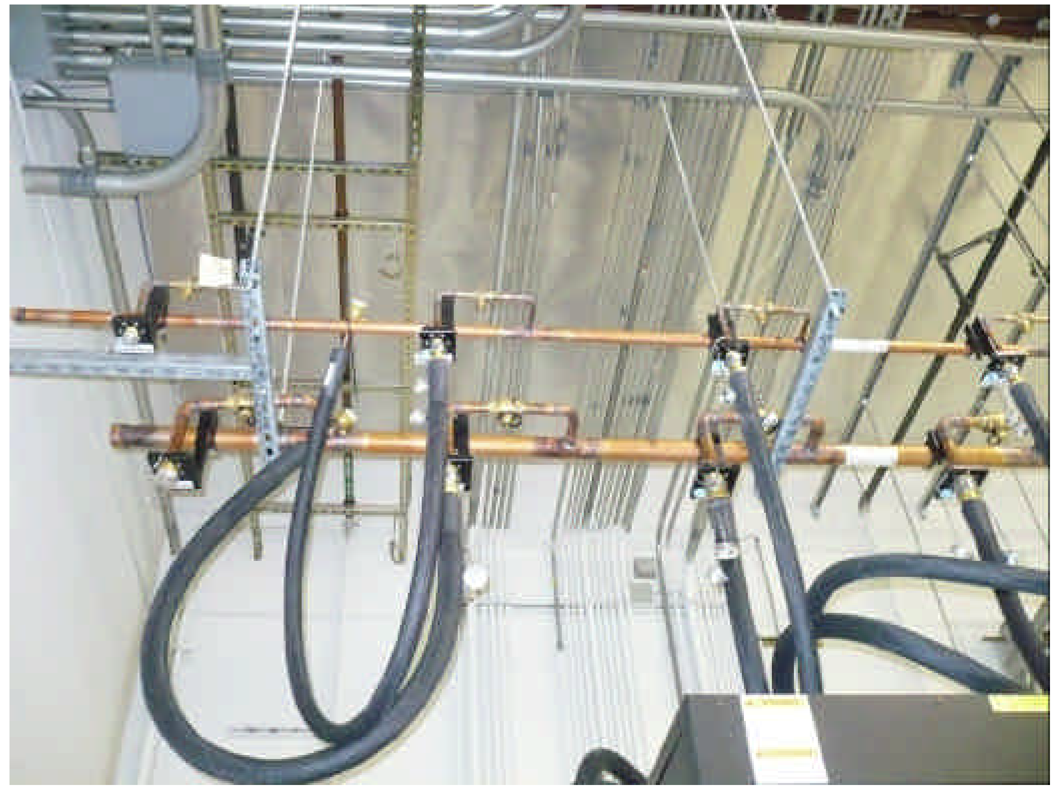

Cooling Systems - Liebert XD Solution

<u>Benefits</u>

- R-134a refrigerant based system (no water in/near racks)
- Up to 20kw of heat rejection per rack
- 160kw (45 tons) of heat rejection per XDC pump/refrigerant heat exchanger
- Reduces energy consumed by CRAC fans
- Modular – add capacity incrementally
- ~18 month ROI for new installations over traditional CRAC units assuming $.09/kwh

Cooling Systems - Liebert XD Solution

Downsides

- Higher initial hardware and install cost
- No monitoring of fan failure
- No temperature monitoring on cabinet heat exchangers
- Requires < 45F (7C) water temperatures
- No provisions for humidity control - XDC modulates refrigerant flow to maintain coil temperature above dew point
- More difficult to design in redundancy

UPS – Conversion Losses

- Dual conversion most popular
- Efficiency varies with load
- Newer designs 90-93% efficient at partial loads
- Rotary UPSes becoming prevalent for new construction
- Rotary UPSes typically 97% efficient

UPS – Power Distribution

- Distribution losses 3-5% of power consumed
- Power transformers consume valuable raised floor space and generate additional heat load
- Copper prices skyrocketing
- New power distribution strategies include 600VAC, 540VDC, and in row rotary UPS and/or DC-DC converters
- DC distribution simplifies redundancy and eliminates complex/costly transfer switches

Links

http://hightech.lbl.gov/

http://www.80plus.org/

http://www.eere.energy.gov/datacenters

http://www.cfroundtable.org/

http://energystar.gov/index.cfm?c=prod_development.server_efficiency

Google - Power Provisioning for a Warehouse-sized Computer