FIO Fabric Infrastructure and Operations

CERN IT Department

Computer Centre Evolution: Infrastructure Issues

CERN openlab Board of Sponsors Meeting

April 26th 2007

Tony Cass

Many interesting discussions with Intel personnel gratefully acknowledged

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Agenda

CERN IT Department

- Power Growth
- Cooling
- Power Delivery
- Summary

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

FIO

Power Growth

CERN IT Department

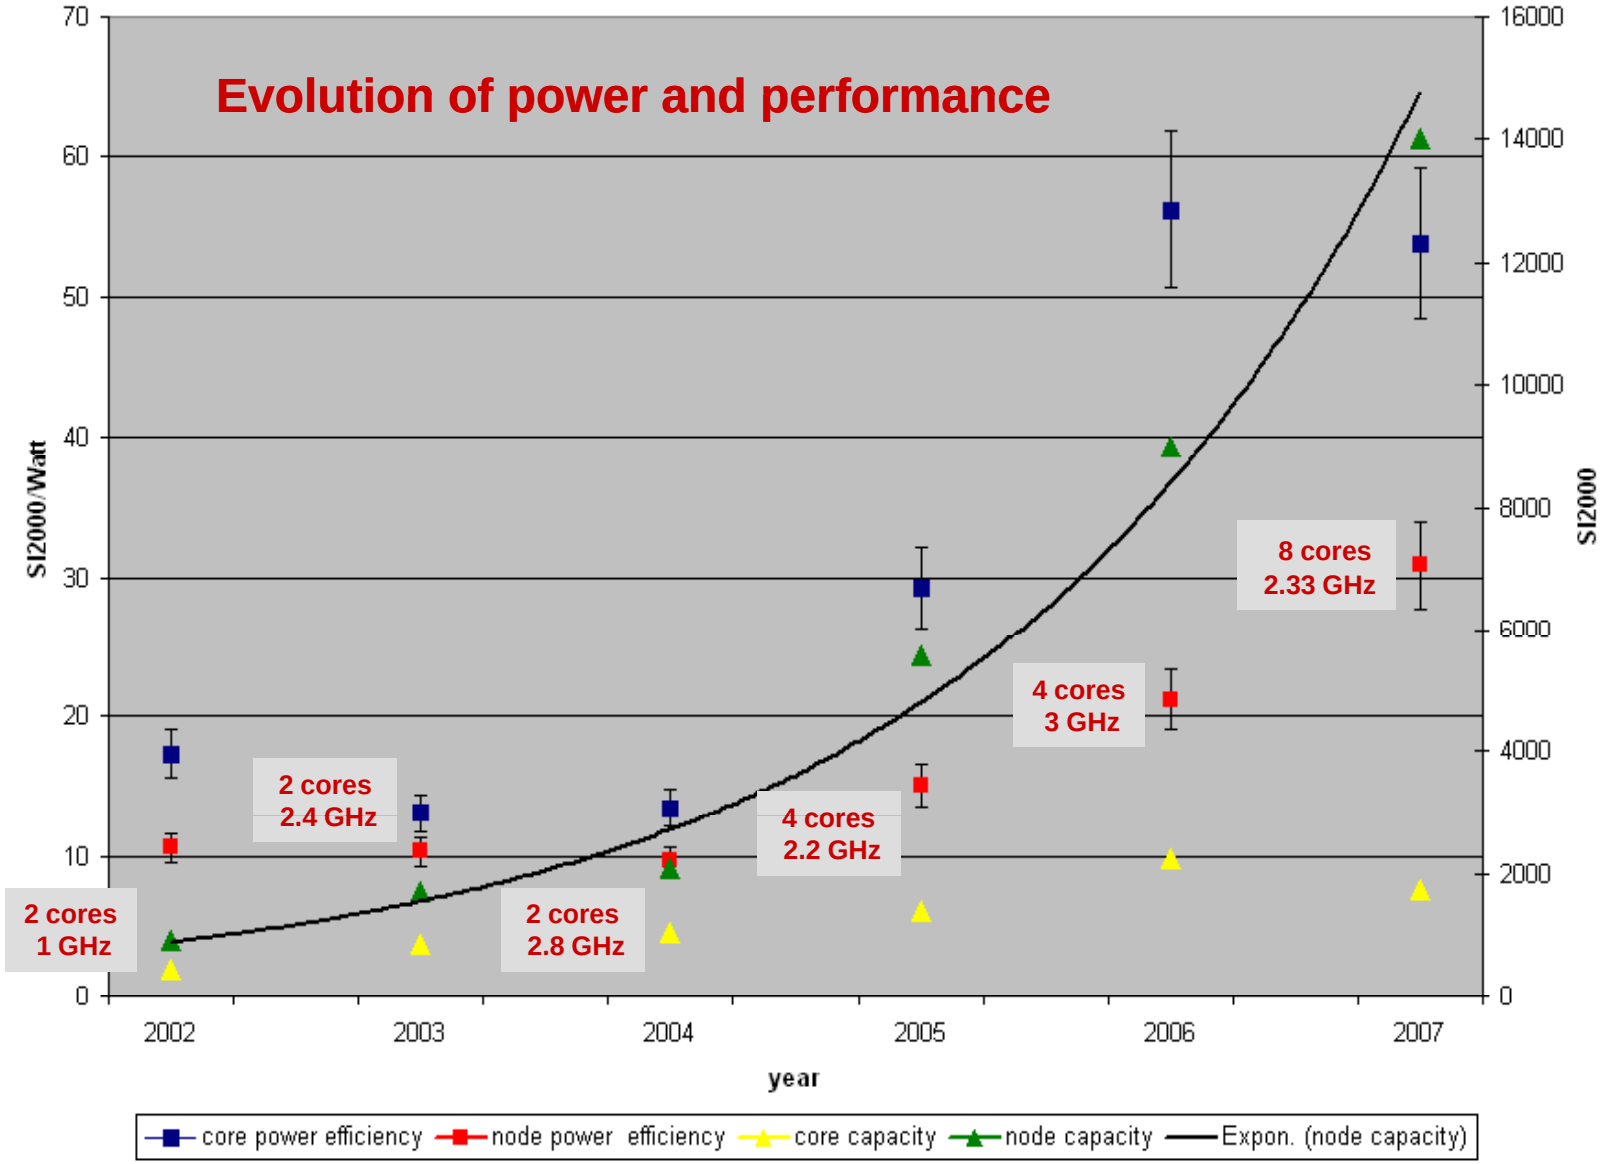

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Cooling Options

- Power in = Heat out
- How to remove?
  - Water
    - \+ efficient (over 3,000x better than air, volume for volume)
    - \- in-room distributions (if they exist) generally not adapted
    - \- leak risk relatively high for per-rack solutions
    - \- limited redundancy for closed rack solutions (failure or maintenance)
  - Air
    - \- inefficient
    - \+ flexible
    - \+ standard
  - Alternatives?
    - exist (e.g. CO_2), but are rather exotic and not considered further here.

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Air Cooling Issues

CERN IT Department

- Capability…
  - … needs to increase from ~3kW/m^2 to ~20kW/m^2
  - rigorous separation of hot and cold aisles is key
- Efficiency
  - Poorly designed cooling systems can double overall power demand for a computer centre.
  - Efficient centres can reduce the overhead to ~30%
    - minimise cold air leakage
    - maximise temperature difference between outlet and inlet air.
    - high outlet air temperature improves overall system efficiency
      - ambient external air can be used as part of the cooling cycle for much of the year.
- Redundancy and backup are essential
  - Temperatures can rise dramatically if the cooling fails!

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN

Power delivery

- High power density electrical installations are not a problem.
- Efficient power delivery to the processor is the issue.
  - At electricity costs of ~300k€/MW/year (400k$), small efficiency increases lead to large savings
    - a 1% increase in efficiency saves 30k€/yr for a 10MW supply.
  - If the efficiency reduces the heat dissipated, the savings can more than double
    - reduced heat load => reduced power demand for the cooling system.

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Where can efficiency increase?

CERN IT Department

- Motherboard/processor?
  - ? I'm not an expert, but reducing the range of voltages used would likely improve efficiency.
- System power supply
  - Easier to achieve efficiency in a larger power supply
    - a key advantage for blades over single systems
    - further improvements with an in-rack backplane?
      - but industry standard needed, not proprietary options
  - Harmonics & Harmonic reduction impact
    - passive filters reduce 3rd harmonics, but switch the nature of the load from inductive to capacitative which has a negative impact on the upstream distribution.
    - active filters avoid this impact, but are more costly.
- Upstream
  - Power line losses
    - a factor to consider when choosing a location
  - Machine room distribution: DC vs AC

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN

DC Distribution

- DC distribution reduces losses due to multiple power conversions
  - AC/DC then DC/AC in the UPS
  - Step down transformer before load
  - AC/DC in power supply
- Interesting option but
  - standardisation needed
  - supply architecture differences (and safety norms?) may reduce the efficiency gains in Europe as compared to the US.

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Distribution Architectures

CERN IT Department

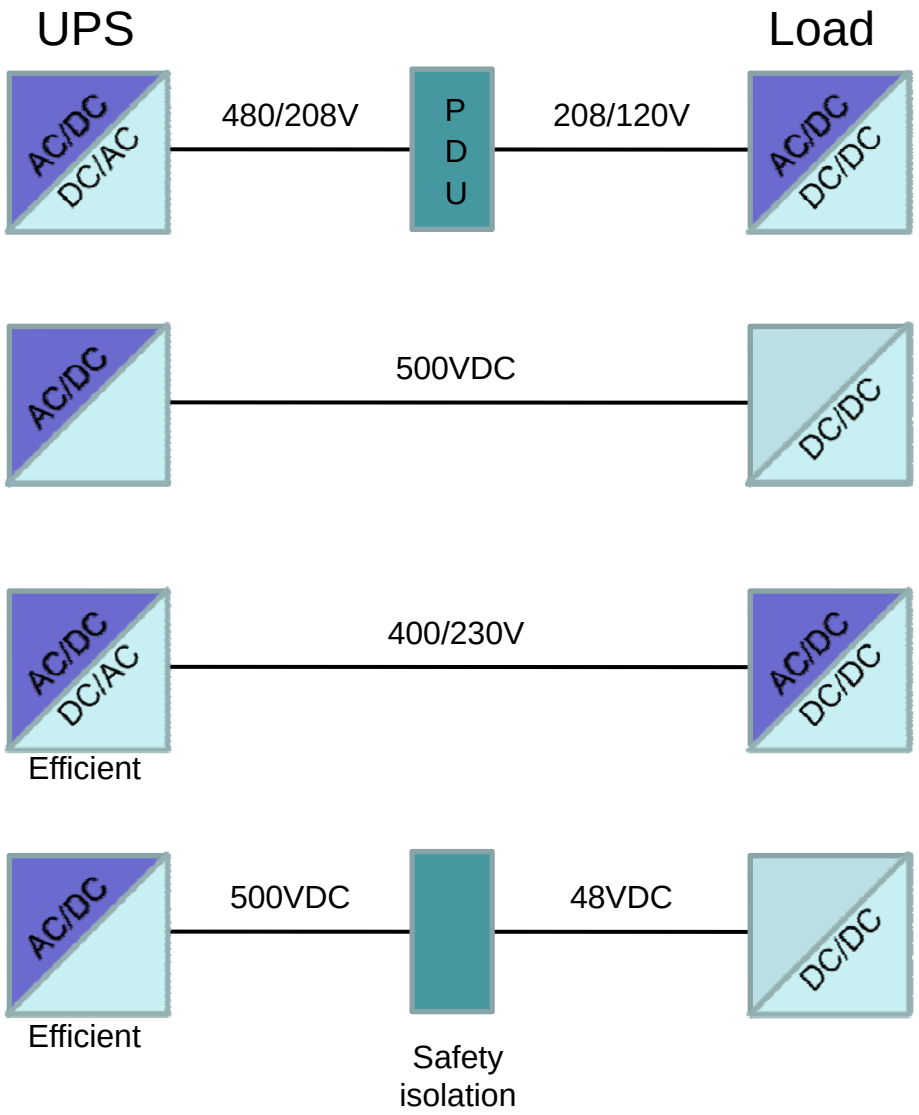

US AC Distribution

DC Distribution Option

EU AC Distribution

EU DC Distribution?

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Summary

- Computer Centre infrastructure must evolve if we are to meet likely demand for growth in compute capacity.
- Air cooling is possible for densities up to ~20kW/rack
  - but the option to use at CERN depends on the configuration of the building; if constructing new, favour long narrow rooms over square.
  - efficient solutions are essential to reduce overall power demand
- Power delivery efficiency is also increasingly important
  - upstream, c.f. interest in DC solutions
  - in the box

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Agenda and Summary

CERN IT Department

- Power Growth
  - looks set to continue; overall and per rack
- Cooling
  - the biggest challenge!
- Power Delivery
  - efficiency, efficiency, efficiency
- Summary

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it