The Abdus Salam
International Centre
for Theoretical Physics

United Nations
Educational, Scientific and
Cultural Organization

IAEA
International Atomic Energy Agency

2494-19

Workshop on High Performance Computing (HPC) Architecture and Applications in the ICTP

14 - 25 October 2013

Overview of Renewable energy in the data centres: status & outlook

C. Onime
ICTP, Trieste

Overview of Renewable energy in the data centres: status & outlook

Clement Onime

onime@ictp.it

Overview

- Introduction
- Techniques & solutions
- Backup energy
- Alternative energy

INTRODUCTION

Introduction

- Energy is central in data centers
  - Power consumption well over that of normal office
- Global CO2 emission
  - IT data centers account for about 3% GHG *but expected to double by 2020**
- Global need to conserve energy
  - Running out of fossil fuels, need to conserve energy
- HPC = High(er) Performance Computing = High(er) Power Computing

Statistics about data center energy costs

- Energy bills are over 10% of running costs in a data center
- Power consumption (U.S.A)
  - 2006: 61billion kWh
  - 2011: 100 billion kWh
- Costs (U.S.A)
  - 2006: 4.5 billion USD
  - 2011: 7.4 billion USD

courtesy of thegreengrid.org

Data centres

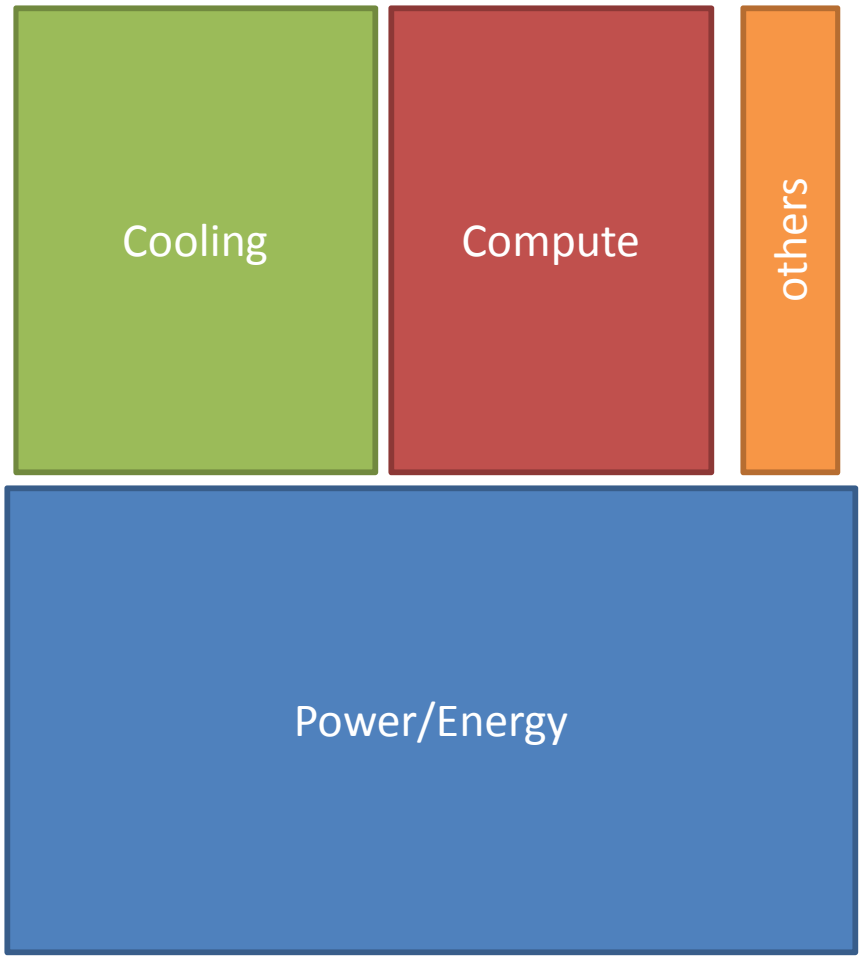

- 4 functional activities
  - Power
  - Cooling
  - Compute
    - And storage
  - Others
    - Communications & networking devices

Data center - Devices

- Cooling
  - Room based cooling, in-row or in-rack units
- Redundant power
  - Power GRID (backup & alternative energy sources),
  - UPS (and battery banks)
- Compute
  - Servers and storage
- Redundant communication
  - Switches, modems etc.

Metrics

- Power Usage Efficiency

  PUE = Total data center power/ IT equipment power

  (1 <= PUE <= ∞), lower is better

  – Also partial pPUE
- Data Center infrastructure Efficiency

  – DCiE = 1 / PUE
- Data center power density

  – DCD = Power of all equipment / area

TECHNIQUES & SOLUTIONS

Cooling systems

- Types
  - Computer Room Air Conditioning (CRAC)
    - Traditional room based cooling
  - Enclosed racks
    - Cold air vents on the floor
  - In-Row/in-rack
    - Hot-air containment
  - Embedded
    - Within server or processor boards
- Costs
  - Initial
  - running

Data center cooling technologies

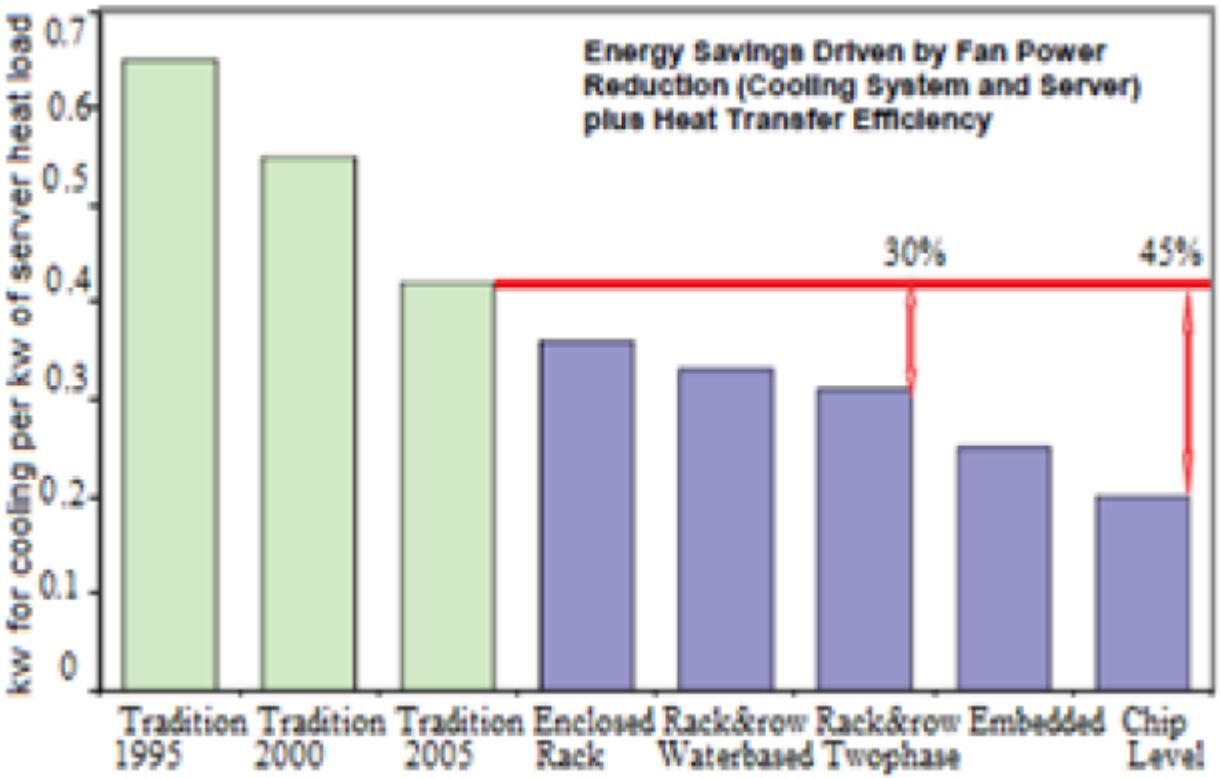

Source: Partha Pratim Ray / Indian Journal of Computer Science and Engineering Vol. 1 No. 4 333-339

Alternative cooling solutions

- Basic principles of cooling servers
  - Good Air flow
    - extractors & filters
- Avoid direct sunlight
  - Shield windows with reflectors/heat shields
- Avoid hot or heated parts of buildings
  - Top-most floor
  - Ground floor

Virtualization

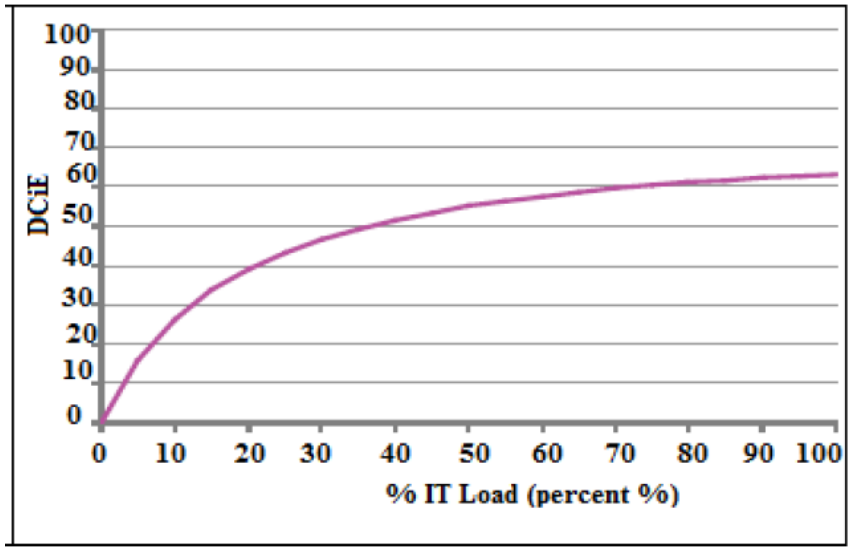

- Increases computing efficiency
- More computing per watt
- Server consolidation
- Over-sizing is necessary to obtain energy gains

Common Virtualization tools

	Qemu	Xen	Virtualbox	Vmware ESXi
Type	Platform	hypervisor	platform	Hypervisor
Virtual network switch	no	yes	yes	Yes
Host O.S.	Linux & Windows	none	Linux, windows, OS/X, solaris	None
Cost	Free	Free	Free	Free
License	GPL	GPL	GPL/ Proprietary component	Proprietary
Management interface	Virt-manager	Virt-manager	Built-in/ proprietary	Windows client or perl cmdline utilities

Power management

- CPU (transistor gating)
- Turn off portions of CPU when not in use.
  - Coarse or fine
- Clock gating
  - Reduce activity per clock cycle
- Frequency gating
  - Reduce number of clock cycles
- Power gating
  - Turn off cores not in use.
- CPU freq, GPU power state, HDD standby, network speed management, etc..

Power management and consumption

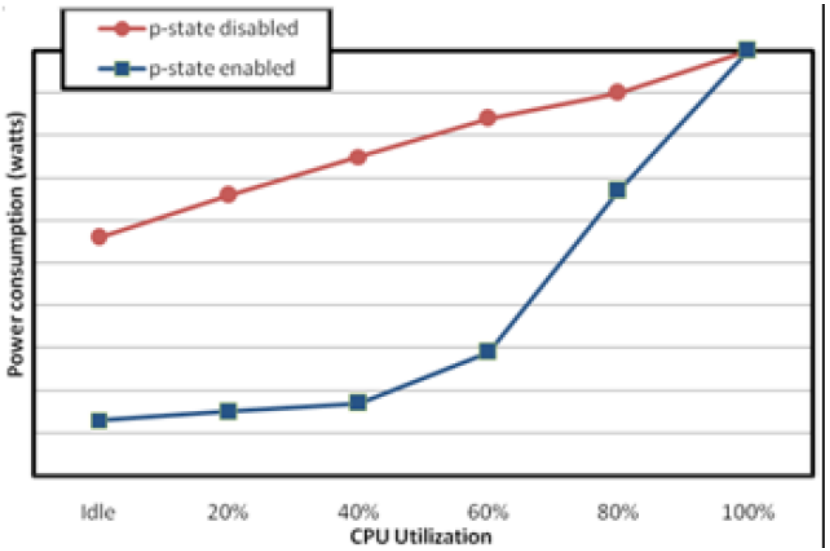

Source: Partha Pratim Ray
Indian Journal of Computer Science and Engineering
Vol. 1 No. 4 333-339

Source: The GREEN-NET Framework

Using power-gating in HPC

- Enable power management in BIOS
- Install daemon on each worker node to manage clock speed.
- Change performance profile using prolog & epilog scripts of resource manager (pbs_mom)
  - Set High Performance at start of job
  - Set Dynamic or low speed at end of job
  - Minimal latency to jobs as nodes are always on.
- Also possible to force processors to always low performance at all times. Reduces generated heat.
  - Lower high speed network, run infiniband @ ddr instead of qdr

Energy Aware Reservation Infrastructure/whole node gating

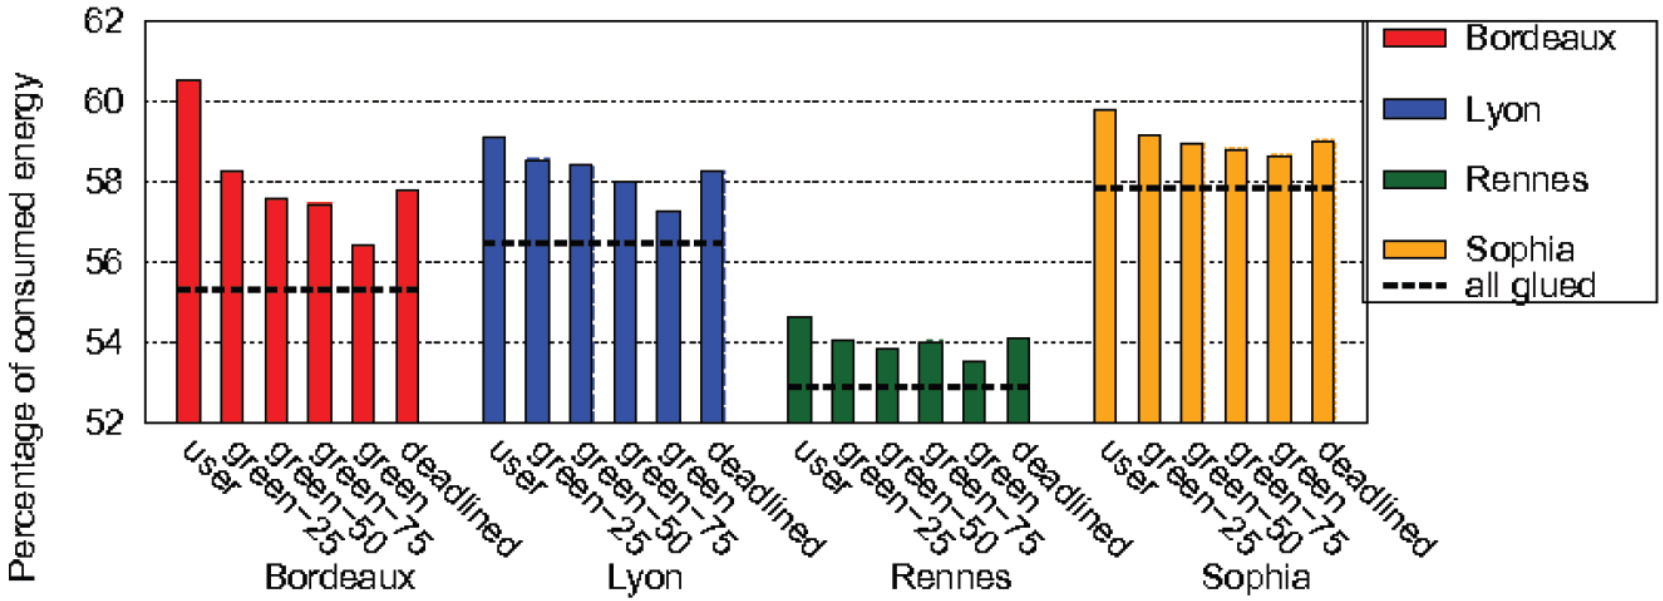

All glue = theoretical limit of energy savings

EARI uses a predictive algorithm to turn on needed nodes

Source: The GREEN-NET Framework

Whole node gating

- IPMI remote management to switch on/off cluster nodes.
- Simple crontab job on master node.
  - Fixed time node gating
    - Fixed period or fixed duration.
  - On-demand gating
    - Integration with resource manager & scheduler

UPS based whole node gating

- Variable energy based on UPS remaining time
- Nodes management
  - Using ssh for power off & eth-wake to power on
- On power fail
  - Switch off un-used nodes
  - All compute nodes off at UPS < 8% left
  - Masternode halt (no power off) cycle at < 4%
    - Hardware based watchdog reset set for 15 minutes
- On power return
  - Wait until battery level > 20% before power on nodes.

backup energy sources

- Battery based systems
  - UPS + GRID
    - Load should be about 50% for best result
    - Only server/compute protection no cooling
- Mixed energy solutions
  - UPS + Diesel generators + GRID
  - Generators also handles cooling
  - Running costs for diesel
    - 100kw @ ¾ load ≈ 20 - 24 liters/hour

ALTERNATIVE ENERGY

Wind turbine @ 50 kw

- Design lifetime = 20 years
- Start speed = 3m/s
- Rated speed = 10 m/s
- Stop speed = 50 m/s
- Height = 12 – 30m

Solar @ 20 kw

- Design lifetime = 25 years
- 80 units (1.6m x 1.02m)
  - 35mm – 50mm
- Batteries
  - 60pcs for 20kw
  - lifetime = 5 years

MicroHydro

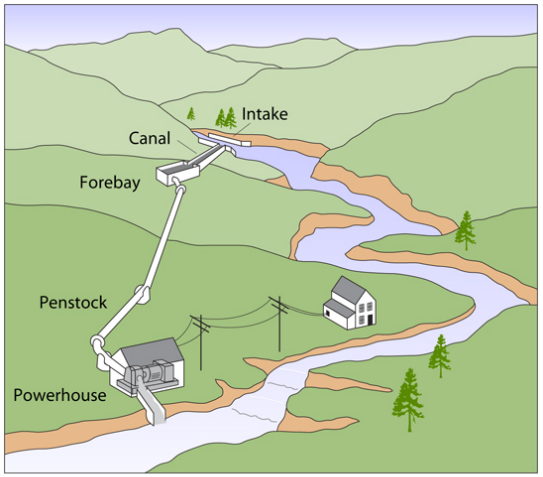

In this microhydropower system, water is diverted into the penstock. Some generators can be placed directly into the stream.

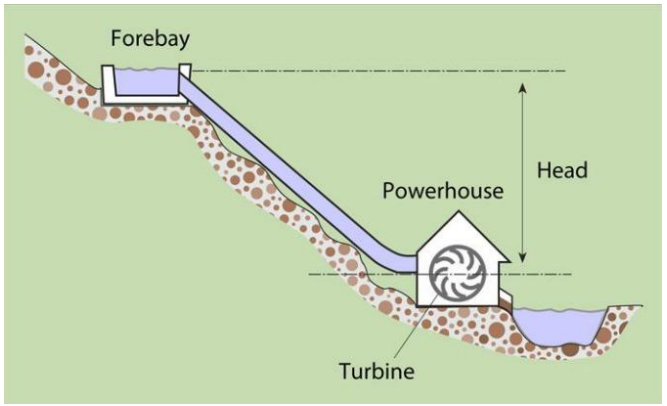

- Head drop
  - Minimum of 61cm
- Flow rate
  - Faster is better
- Quite seasonal
- Estimated output
  - 0.082 * H(m) * F(l/m)
  - 0.1 * H (ft) * F(g/m)

That’s all folks!!
Thank you