

Energy Efficient HPC systems



Architect of an Open World™

CAS 2K13 – Sept. 2013

Jean-Pierre Panziera
Chief Technology Director

personal note



Prévisions Bulletins Avalanches

Rechercher Ville, Code Postal... Choisir un département Tous les départements Choisir une station Choisir une station

France → Alpes du Nord

▲ Aujourd'hui
▲ Jeudi
▲ Vendredi
▼ Samedi
▶ Matin
▶ Après-midi

Prévisions du Samedi 7 septembre 2013
Alpes du Nord

Légende des pictogrammes

METEO FRANCE
Toujours un temps d'avance

NOUVEAU !

Complete solutions for Extreme Computing



bullx supercomputer suite

Production ready
supercomputers



Data Centers for all
organizations



Applications optimized
for hyper-parallel
supercomputers

extreme factory
stay lean: compute smart

HPC for every user
with public/private HPC cloud

Bull: from Supercomputers to Cloud Computing

Expertise & services

- HPC Systems Architecture
- Applications & Performance
- Energy Efficiency
- Data Management
- HPC Cloud

extreme factory
stay lean: compute smart

center for
excellence in parallel
programming


Software

- Open, scalable, reliable SW
- Development Environment
- Linux, OpenMPI, Lustre, Slurm
- Administration & monitoring

bullx **supercomputer suite**


Servers

- Full range development from ASICs to boards, blades, racks
- Support for accelerators



Infrastructure

- Data Center design
- Mobile Data Center
- Water-Cooling




Leading HPC technology with Bull



TERA100 – 2010

1st European PetaFlop-scale System


Rank #6



CURIE – 2011

1st PRACE PetaFlop-scale System

Rank #9




“C1” – 2013

1st Intel Xeon E5-2600 v2 System

Direct Liquid Cooling Technology


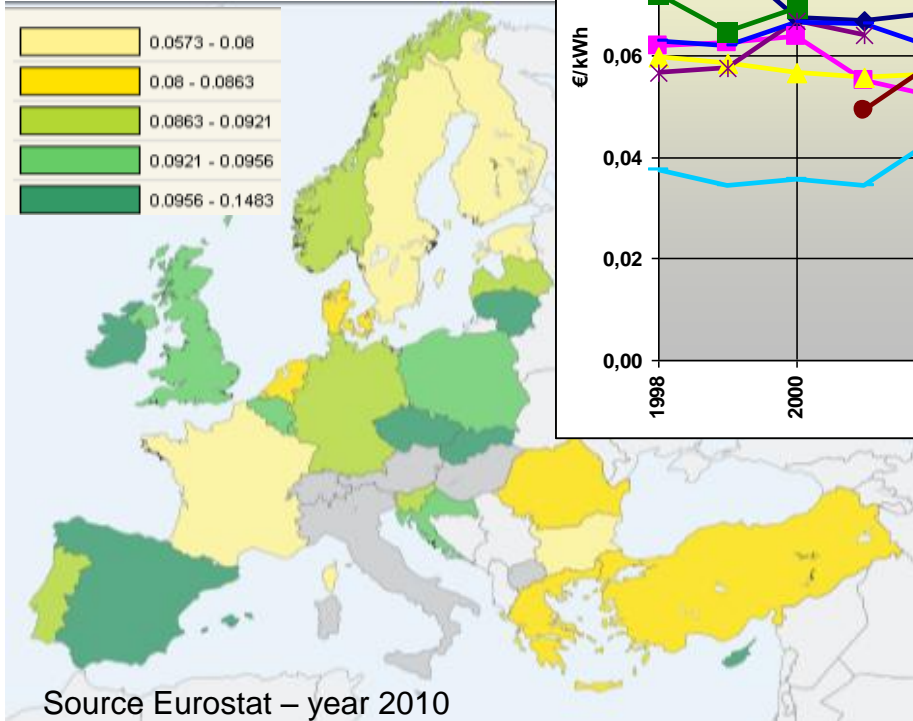


Energy (Electricity): a significant part of HPC budget




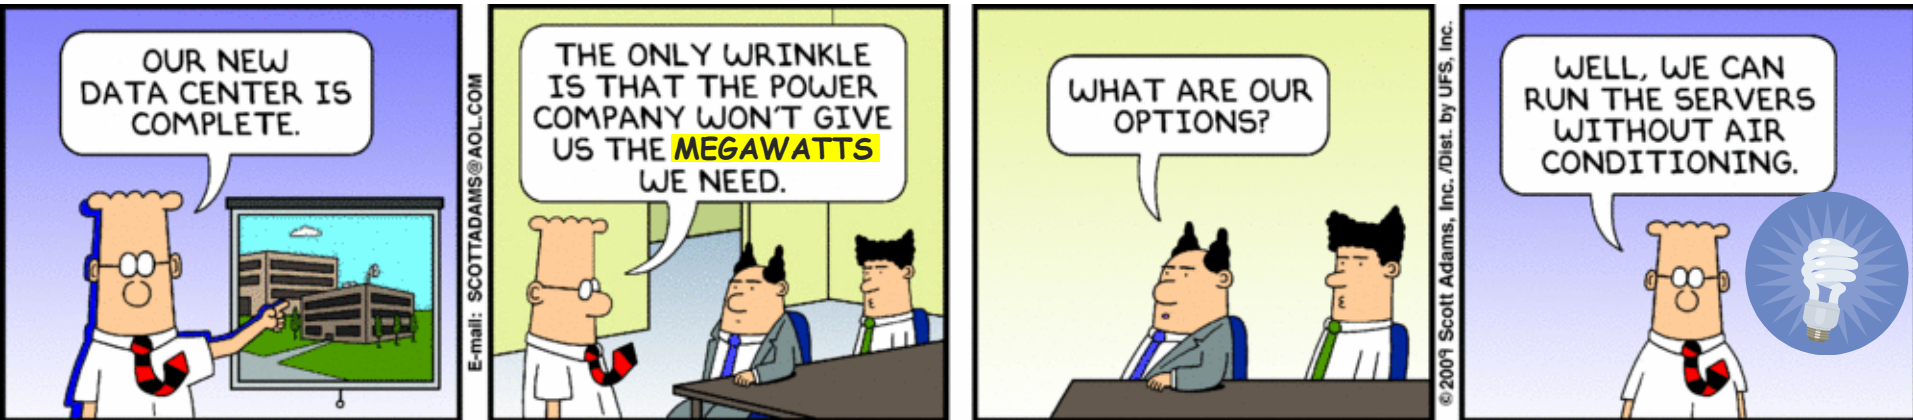
Industrial Electricity Prices in Europe

Electricity prices highly variable across Europe
Avg 0.11€/kWh




Electricity prices
Rising steadily
CAGR 12%

Power to the datacenter




bullx B700 series – Direct Liquid Cooling (DLC)



Direct Liquid Cooling DLC rack:


- Dual-pump unit (80 kW cooling capacity)
- Power Supply Unit + UltraCapacitor opt.
- Rack management (incl. Gigabit Ether.) +
- 5 chassis, each including:
 - 18 dual-processor nodes
 - Embedded 1st level InfiniBand switch
 - Extra Embedded Gigabit Ethernet switch
- Silent
- Extra-easy maintenance
- Optimized PUE (< 1.1)

Cooling & Power Usage Effectiveness (PUE)





Direct Liquid Cooling Infrastructure



- With hot water cooled servers, water chillers are not required anymore



Where do all these Watts go ?



Node consumption varies with workload



Wrt Linpack (max -> 100%)
Memory streaming 75%
Irregular memory access 55%
Idle 25%

Using turbo is never energy efficient

Power Management

Accounting

- Users billed separately for CPU, IO, ... and Energy
- Keep compute center electricity bill within budget


Control power

- Avoid running over capacity
- Allow for priority jobs
- Adjust power consumption with electricity cost

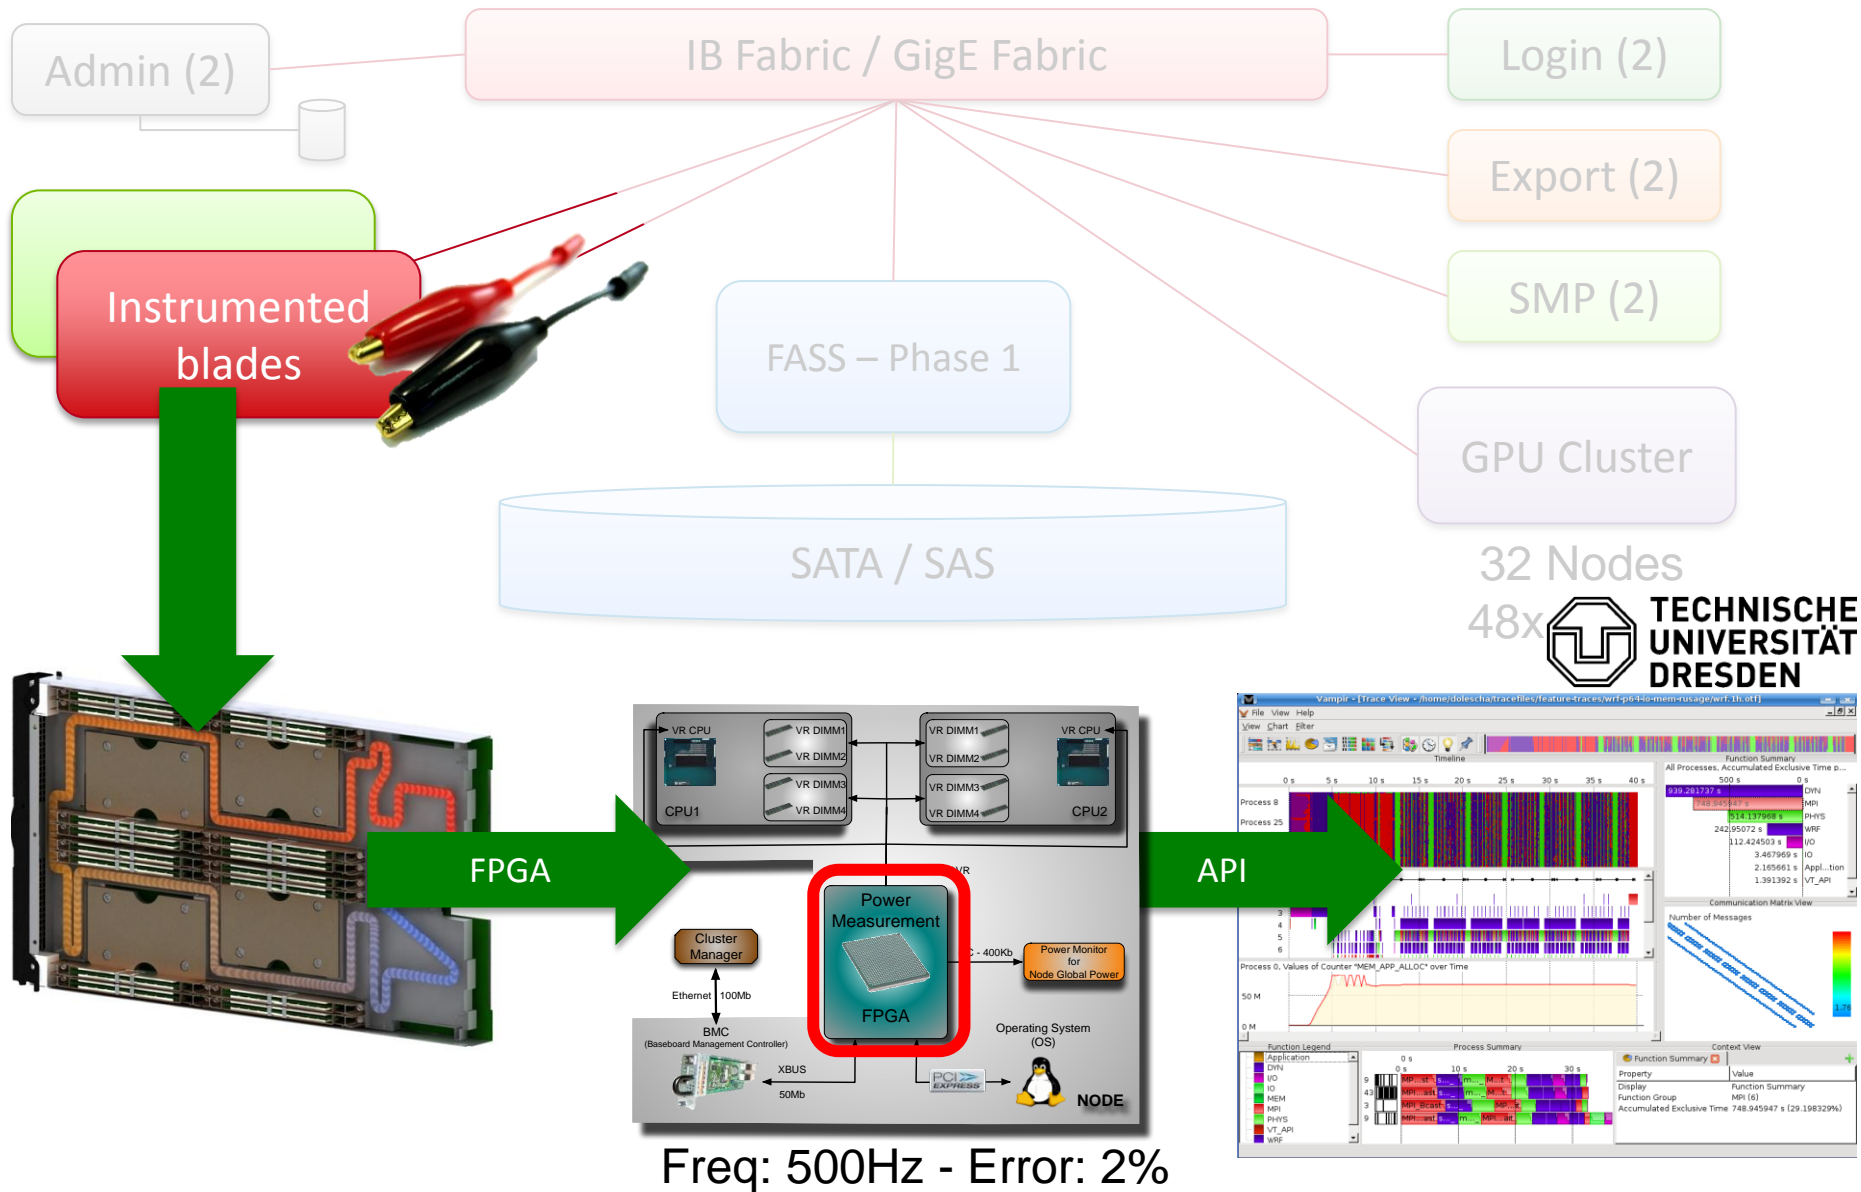
Energy consumption / cost optimization

- Fine & precise power monitoring
- Power data analysis
- Control all system resources power

... enter software



Bull - TU Dresden → high frequency monitoring



Energy efficient HPC systems ...

Green systems

- Interest driven by energy cost and green attitude
- Green systems start with Green components
 - CPUs – Power only where/when needed, throttle frequency
 - Memory, DDR4 saves 20-30%
 - PSUs, optimize AC/DC, DC/DC conversion
 - Interconnect...
- Direct Liquid Cooling save on CAPEX (chillers) & OPEX (electricity)
- Non-intrusive high definition power monitoring
- Power: another parameter in system usage optimization
- Energy aware batch scheduler



Architect of an Open World™
