

SOS10 Capability panel

Dr. Marie-Christine Sawley, CEO CSCS

sawley@cscs.ch

Maui, 7-9 March 2006



CSCS present infrastructure



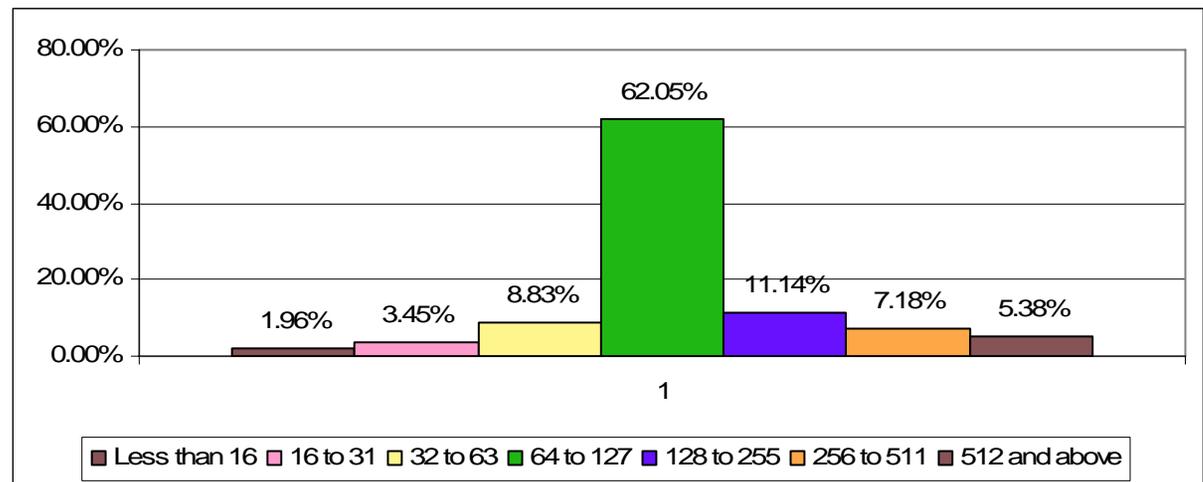
New system: "Horizon"



1'100 AMD Opteron @2.6 GHz
5.7 Tflops Peak performance
The system went into production on
January 18th, 2006

A few indicators

- The system is being used up to 85%
- The majority of jobs use between 64 and 127 processors

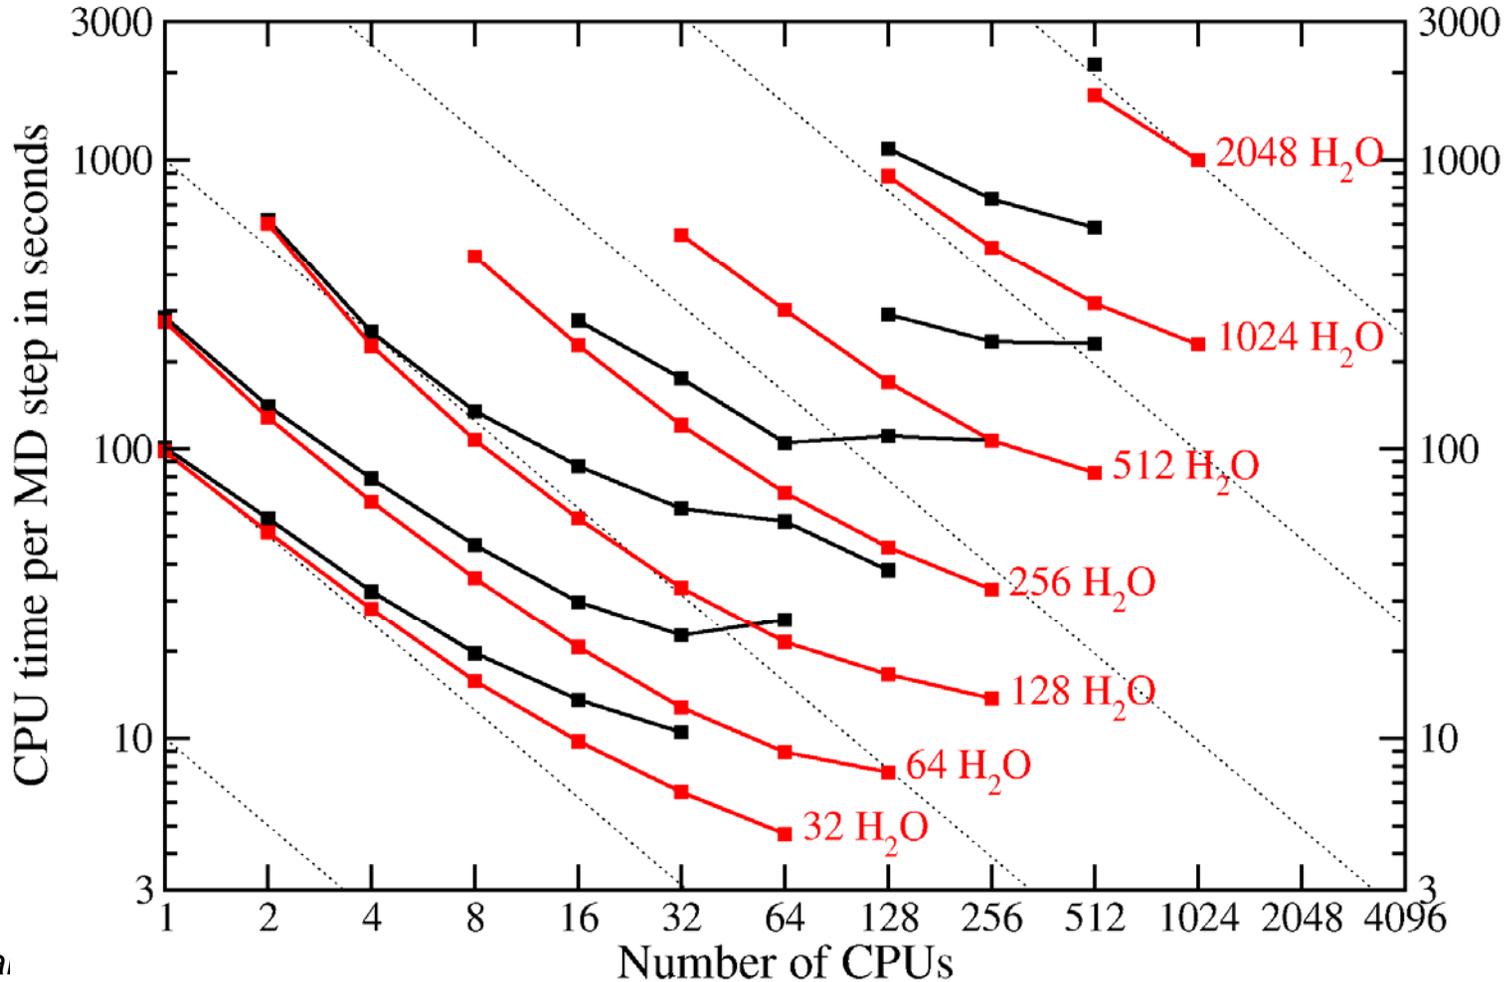


Early users.....and good scaling

Courtesy of Dr. Matthias Krack, group of Prof. Parrinello

CP2K water benchmark

Cray XT3 without Lustre (CSCS) Cray XT3 with Lustre (CSCS)



Offering new possibilities

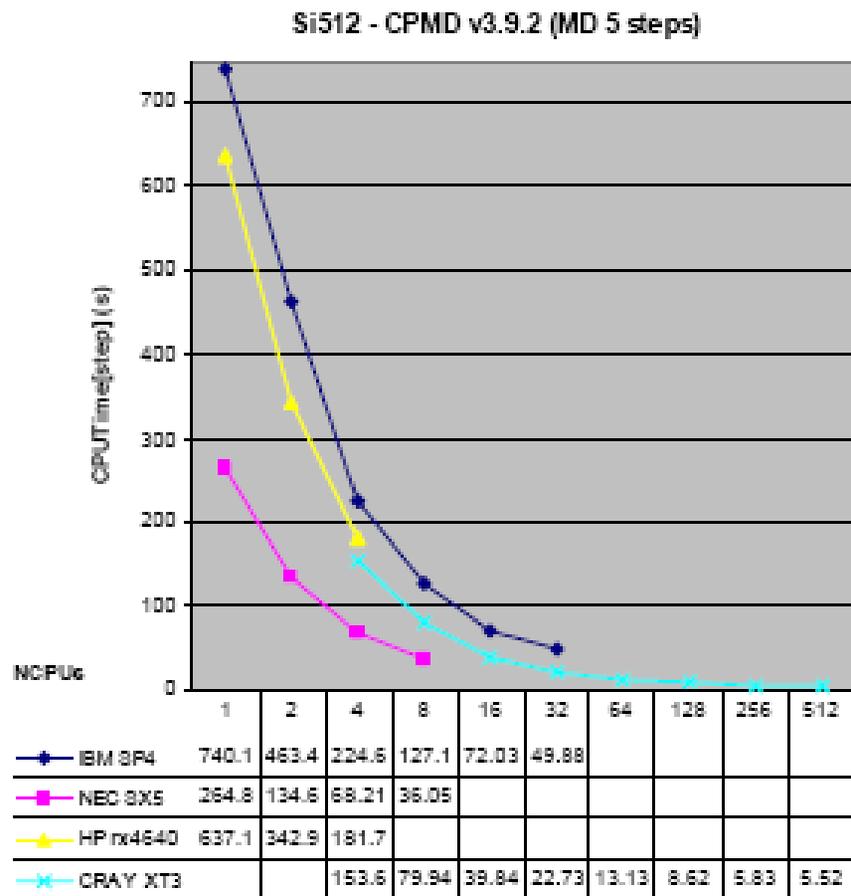
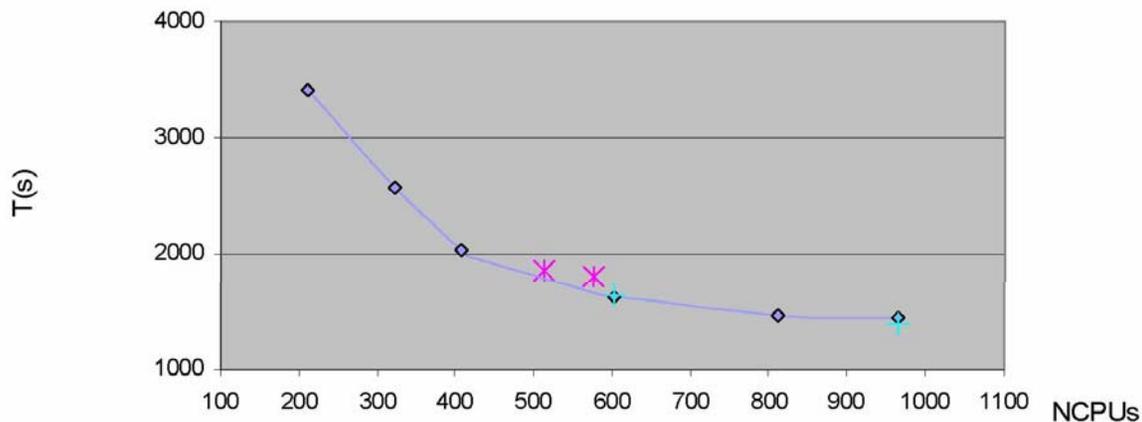


Figure 16

Research in Meteorology : early results for higher resolution

aLMo Bench: Timings



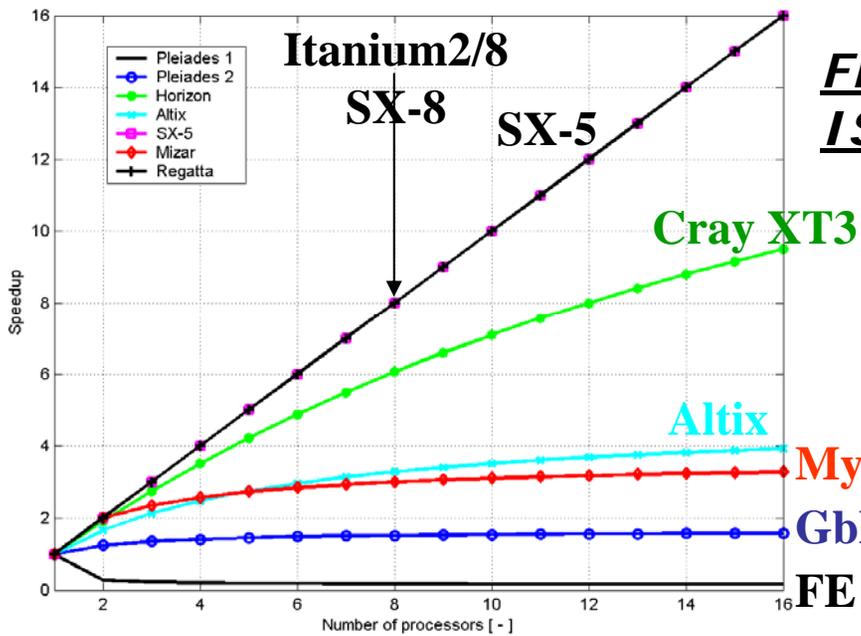
	210	322	406	512	576	602	812	966
—◇— TIME(s) ded.	3405	2568	2044			1631	1473	1438
* TIME(s) shared				1849	1801			
+ TIME(s) ded. + inline						1639		1387

Capability vs capacity

- Capability system offers **high scalability** and favours jobs using a major portion of the resources
- Capacity is the ability to handle a **large** number of jobs efficiently, mostly middle size
- In a HPC centre, both systems exhibit **unique features**
 - large memory; very high bandwidth; vector processors, high performance communication network; etc...

The definition we prefer

- *Capability computing* refers to the use of the most powerful supercomputers to solve the largest and most demanding problem with novel or extreme applications
- The main figure of merit in capability computing is
 - the **size** of the problem
 - or **time** to solution
 - if not **feasibility** of solution
- Usually only for a few, **very large** projects
- Catch is **turnaround**: the balance is hard to find, but excessive turnaround can kill a capability system attractiveness



From R. Gruber: eco computing and ISS presentation,

Myrinet Opteron cluster
GbE Xeon cluster
FE P4 cluster

Get result in	Capability	Capacity	Infrastructure investment costs
4 months	SX-8 (1) Itanium2 (4), XT3 (8)	any other node	150 \$
1 month	SX-8 (1)	Itanium2 (4), XT3 (8)	Around 400 \$.
1 week	SX-8 (4 procs)	?	Around 2500 \$
1 day	unfeasible	unfeasible	

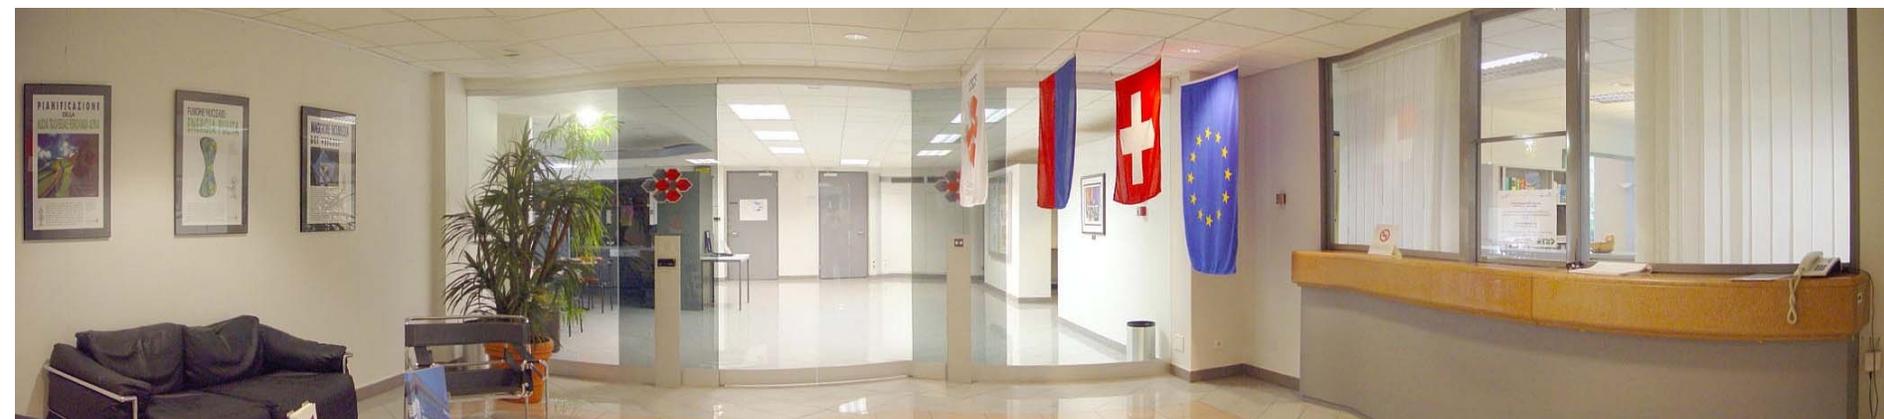
Investment strategy

- Looking at the **whole** portfolio of applications, to differentiate the tiers of the pyramid
- Evaluation based on a mix of **total cost of ownership**, **performance** on selected applications and synthetic benchmarks, and if needed, **time to solution** (typical: NWS application)
- For having the best return on investment, applications scaling to about one eighth to one quarter of the system are **capability jobs** if they use efficiently special features such as
 - **communication network, memory size, memory bandwidth or vectorization**
- The others are **capacity** jobs; in order to host capacity jobs, you need a system with high throughput, both capacity or capability
- **Some users need to be educated to understand where the premium stands, make the right choice, or accept to recode their application**

Users need to choose in which category they work

–...“ You do not hire a an America’s Cup boat –with the support team- to sail between Geneva and Lausanne at the cost of a standard boat; besides, it may be uncomfortable....”





UNICORE
FORUM

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zürich

eGee
Enabling Grids
for E-science

Working for enabling scientific discovery in world class scientific centres in Switzerland

March 7-9/sos10/mcs