

CSCS

Swiss National Supercomputing Centre



HPC Systems at CSCS

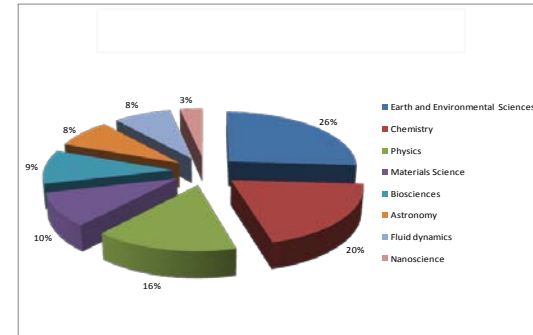
HPC Advisory Council Switzerland, 15th March 2010

Neil Stringfellow, CSCS



Client Portfolio

- Academic Researchers

- Swiss Universities
- ETH Domain



- MeteoSwis



- CHIPP



- Swiss High-energy physics community




- *Paul Scherrer Institute jointly invests in largest CSCS machine*

Client Portfolio - Requirements

- Academic Researchers
 - Swiss Universities
 - ETH Domain





- MeteoSwiss

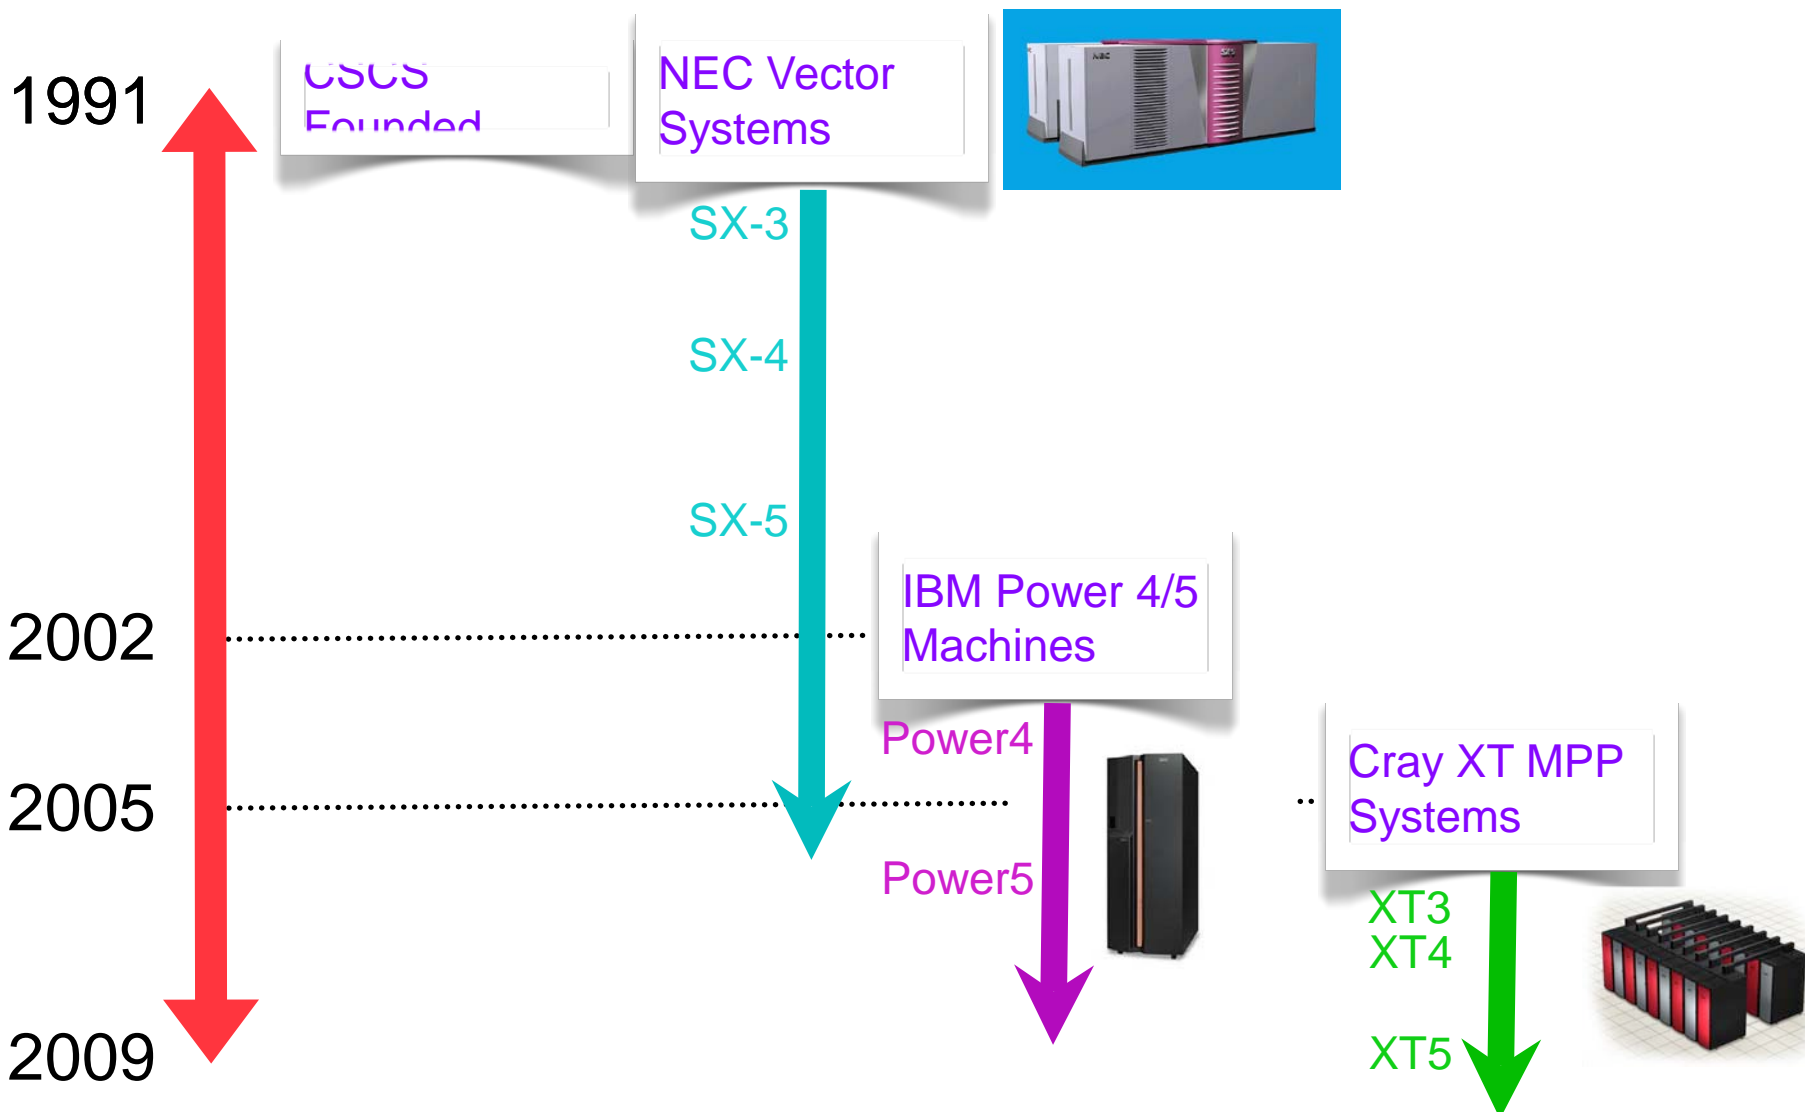


- CHIPP

- Swiss High-energy physics community



History of HPC Machines at CSCS



CHIPP Community

- Swiss high-energy physics community
- Experiments analyse data from Large Hadron Collider at CERN

Need just
under 1
Petabyte of
Disk

Refresh disk
monthly – need
~3 Gigabits of
internet
bandwidth

Operate in a
grid – have
clearly defined
machine
specifications

Sun Cluster - Phoenix

- 960 processors
 - 60 nodes, each with 4 x AMD Opteron quad-core 2.3 GHz
- 960 Terabytes of memory
 - 1 Gigabyte per core
- 560 Terabytes of disk
- Gigabit Ethernet
- Mainly for integer based, rather than floating point calculations



Upgrade currently taking place

- Processors become 2.53 GHz Intel quad-core
- Disk becomes attached as a common parallel file system
- Additional QDR Infiniband interconnect

MeteoSwiss Requirements


- Need high computational power
 - Must run a full forecast suite in 30 minutes
- Require reliability
 - A forecast is produced every three hours
- Need high-availability of external access
 - Observational data is being constantly updated
 - Products (forecasts) need to be disseminated quickly

CSCS Uses failover machines to ensure the reliability

Failover machine may be slower than full machine, but forecast will be delivered, with a small delay.

The Cray XT4s – Buin/Dole

- Originally one system, later split into two machines for redundancy
 - Original machine was first Cray XT to be used for high-resolution operational weather forecasting
- Two Cray XT4 systems have a total of 1728 processors for HPC
 - 432 AMD Opteron quad-core compute nodes running at 2.3 GHz



- Buin has 1032 processors
- There are 30 Terabytes of scratch space
 - Internal Lustre parallel file system
- Buin is used to run the high-resolution suite of MeteoSwiss eight times per day

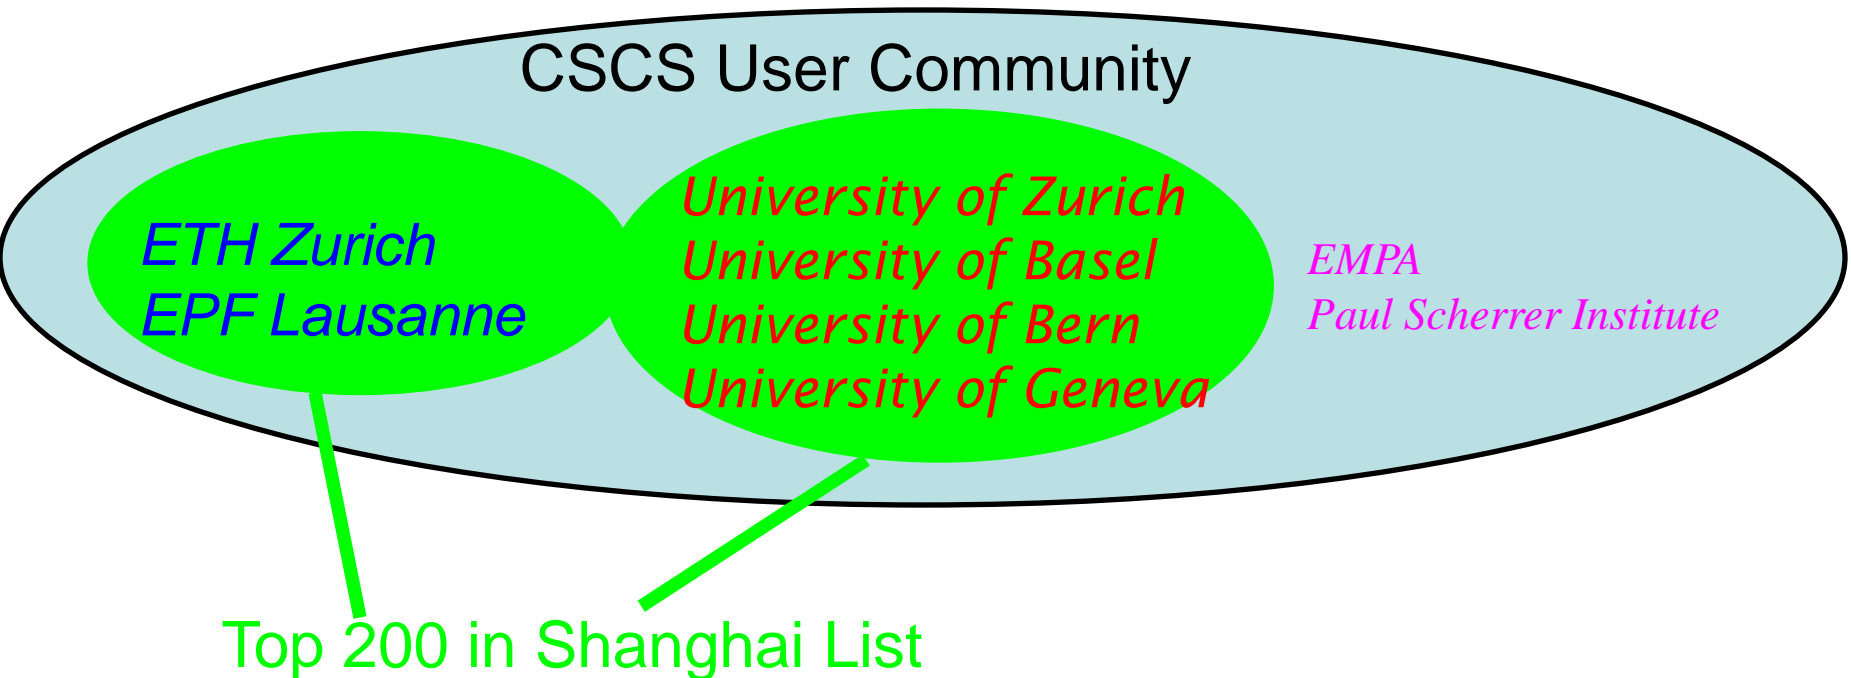


- Dole has 688 processors
- There are 14.5 Terabytes of scratch space
 - Internal Lustre parallel file system
- Dole is used for MeteoSwiss application development and acts as a failover for Buin



Academic Researchers

- CSCS core business is to provide HPC facilities and expertise to the Swiss research community
- Requirements for raw compute power, data storage and data analysis

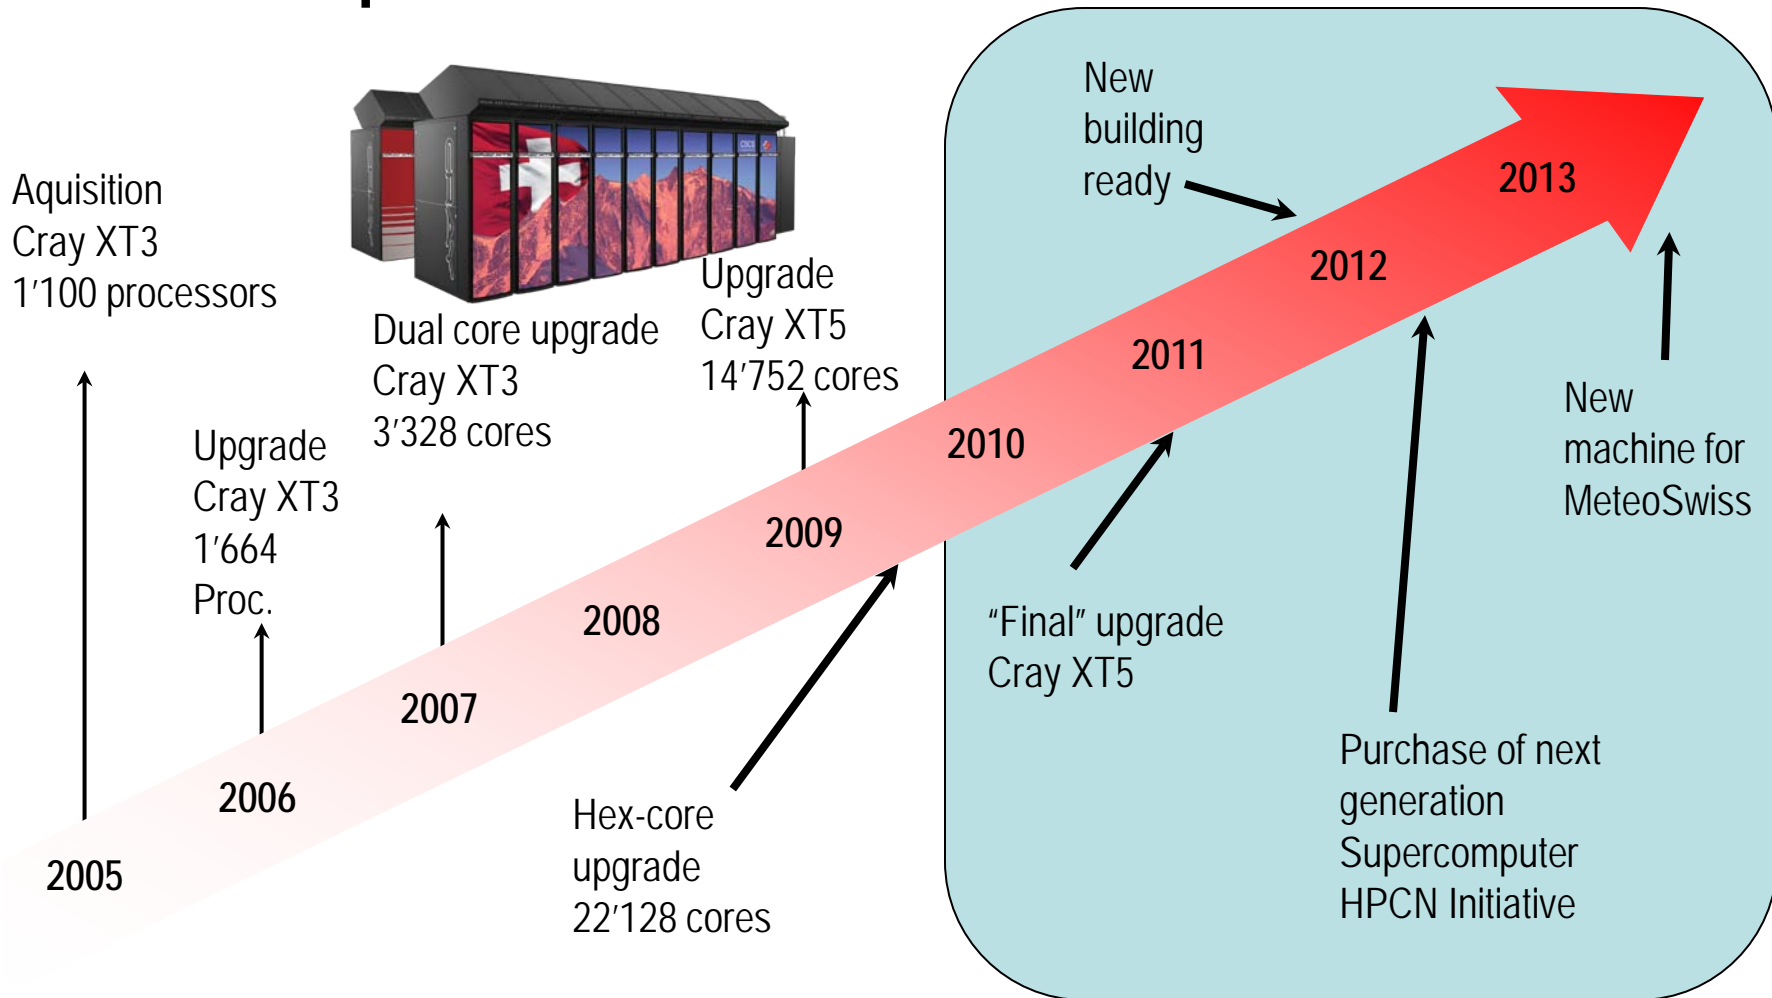


Cray XT5 – Monte Rosa

- 22,168 processors
 - 1844 twelve-way nodes
 - 2 AMD 2.4 GHz “Istanbul” Opterons per node
- Peak performance 212 Tflop/s
 - Linpack 167 Tflop/s
- 29 Terabytes of memory
 - 16 Gigabytes per node
 - 1.33 Gbytes per processor core
- 287 Terabytes of scratch file system
 - ~ capable of 14 GB/s sustained write bandwidth
- 21st on Top500 list in Nov. 2009
 - 4th most powerful system in Europe
- In 2H2009, ran at 85% utilisation
 - ~ 20% of jobs required > 50% of machine



Growth plan 2010-2013



www.hp2c.ch