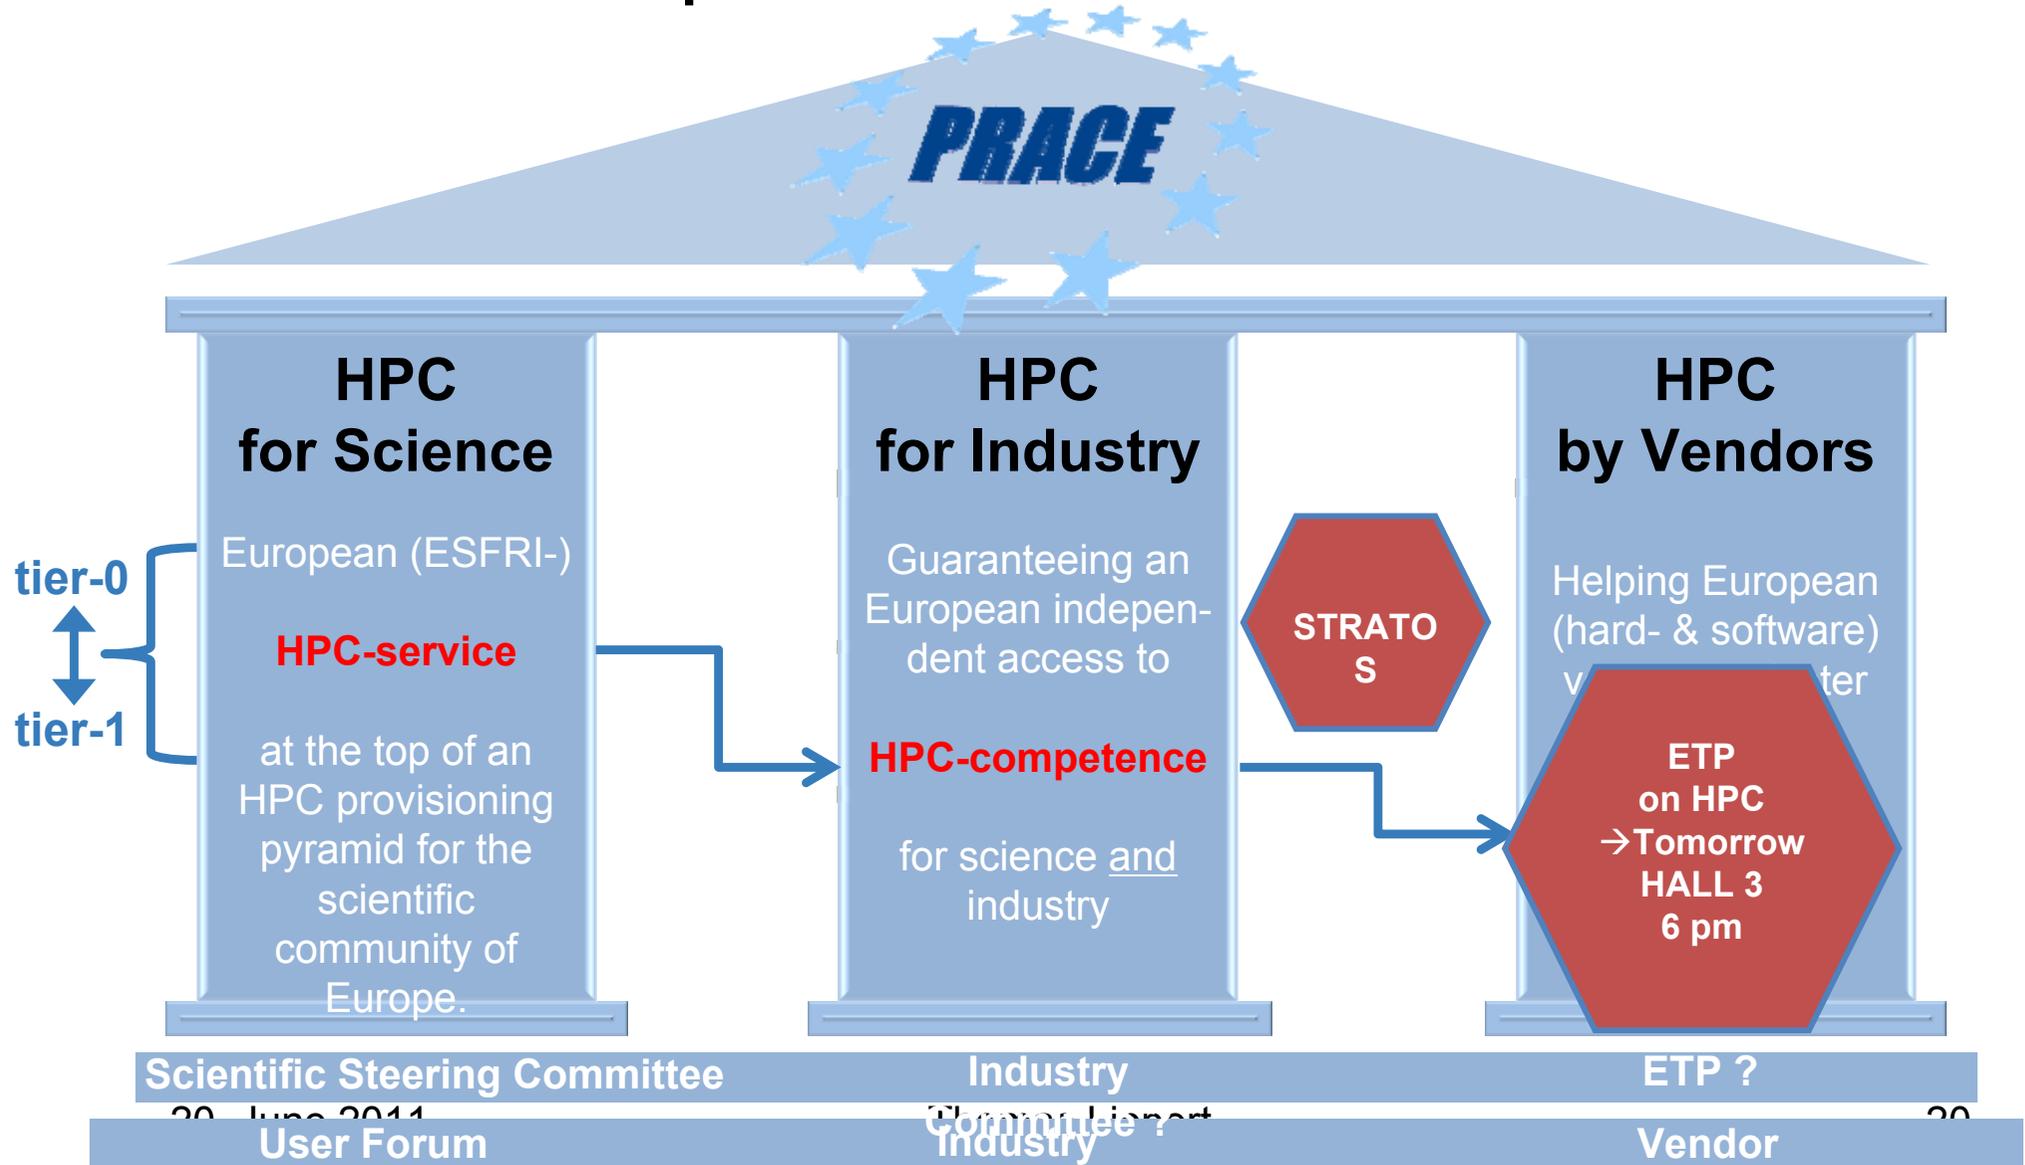


European HPC and PRACE



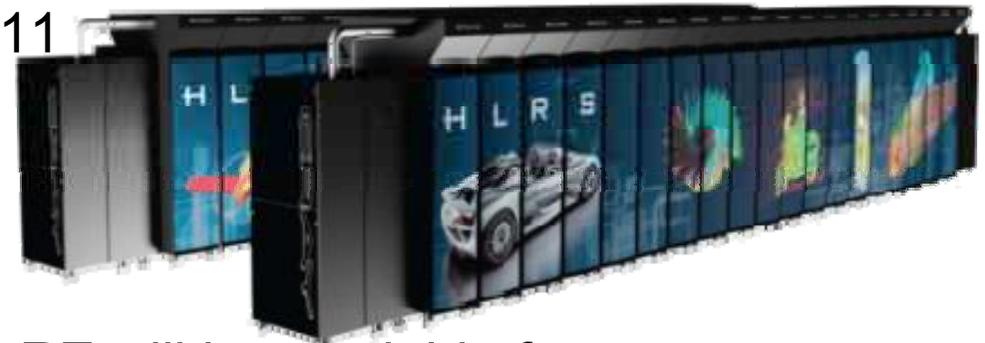
PRACE Tier-0 Systems

- 1st Tier-0 System provides cycles since August 1, 2010
 - Jugene: BlueGene/P in GCS@Juelich
 - 72 Racks, 1 PFlop/s Peak
 - 35% of capacity provided to PRACE
- 2nd Tier-0 System announced by GENCI on October 5, 2010
 - Curie: Bull Cluster with Intel CPUs operated by CEA
 - 1.6 PFlop/s peak in Oct. 2011 (1st step in 2010)
 - Largest fraction of capacity provided to PRACE
- Next Procurements (in alphabetical order)
 - BSC, CINECA, GCS@HLRS, GCS@LRZ
 - Procurement plan based on analysis of user requirements and market



Announced PRACE Tier-0 Systems

- Hermit
 - 1 Petaflop/s Cray (HERMIT) at HLRS will be available for European scientists starting in the end of 2011
 - An upgrade to 4-5 Petaflop/s is planned for 2013
- SuperMUC
 - 3 Petaflop/s IBM (Super-MUC) at LRZ will be available for European scientists starting in mid 2012
- Italy (Cineca) has announced its contribution for 2012
- Spain (BSC) will follow 2013



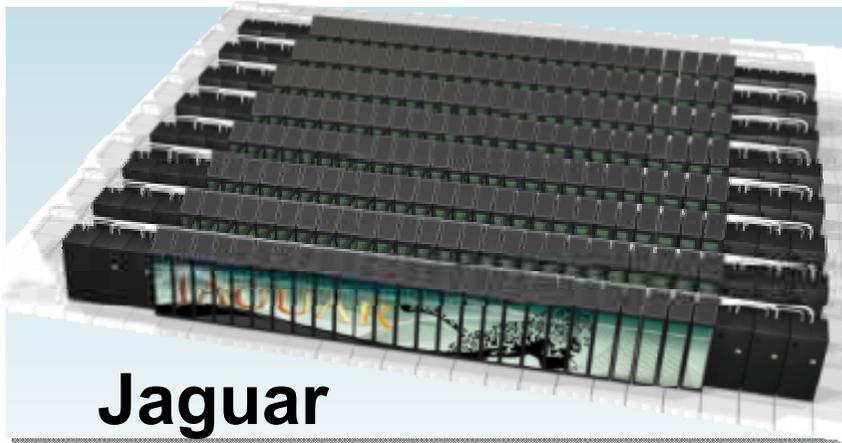
MSU "Lomonosov" supercomputer, 2011



- Peak: 1.37 Petaflops
- Linpack: 674 Petaflops
- Intel Xeon Processors
- Nvidia X2070 GPU
- Infiniband & 10 GigE

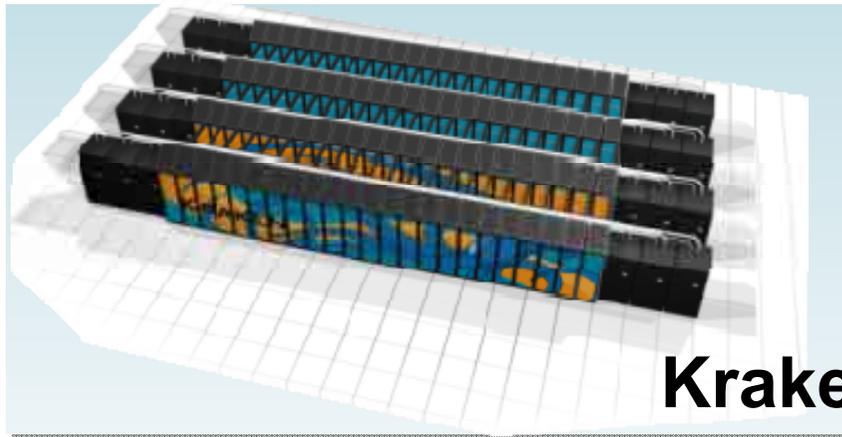
M.V.Lomonosov
1711 – 1765

Notable Petascale Systems at ORNL



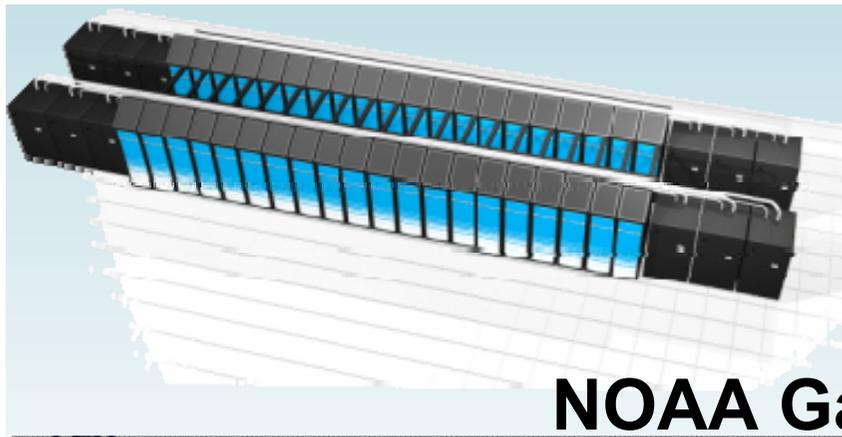
Jaguar

Peak performance	2.33 PF/s
Memory	300 TB
Disk bandwidth	> 240 GB/s
Square feet	5,000
Power	7 MW



Kraken

Peak performance	1.17 PF/s
Memory	151 TB
Disk bandwidth	> 50 GB/s
Square feet	2,300
Power	3.5 MW



NOAA Gaea

Peak Performance	1.1 PF/s
Memory	248 TB
Disk Bandwidth	104 GB/s
Square feet	1,600
Power	2.2 MW



#3



Dept. of Energy's most powerful computer



#11

National Science Foundation's most powerful computer



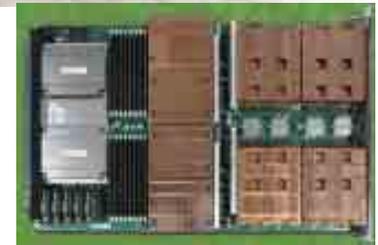
#37

National Oceanic and Atmospheric Administration's most powerful computer

Cray Systems



- Cray XK6
 - Cray's Gemini interconnect, AMD's multicore scalar processors, and NVIDIA's GPU processors
 - Scalable up to 50 Pflops of combined performance
 - 70 Tflops per cabinet
 - 16 core AMD Opteron 6200 (96 /cabinet) 1536 cores/cabinet
 - NVIDIA Tesla 2090 (96 per cabinet)
 - OS: Cray Linux Environment
- Cray XE6
 - Cray's Gemini interconnect, AMD's Opteron 6100 processors
 - 20.2 Tflops per cabinet
 - 12 core AMD Opteron 6100 (192/cabinet) 2034 cores/cabinet
 - 3D torus interconnect



IBM Systems



- BG/P – 1 Pflops system
 - 294,912-processor, 72-rack system
 - High bandwidth optical interconnect
 - Four 850MHz PowerPC 450 processors per BG/P Chip
 - Scaled upto an 884,736-processor, 216-rack cluster 3 Pflops
- BG/Q – Next Generation – 20 Pflops system
 - 2012 time frame
 - 4-way HT PowerPC A2 based chips with 16 cores
 - Sequoia system to be installed at LLNL by 2012
 - Mira system at ANL by 2012
- Blue Waters
 - 2011-12 Time Frame
 - Power 7 Processor
 - Bandwidth Monster, tuned for high efficiency
 - NCSA 10PFlops



SGI Systems

- SGI Altix UV

- 256 sockets (2,048 cores, 4096 threads) with architectural support to 262,144 cores (32,768 sockets) – Nehalem and Westmere EX
- Up to 16TB of global shared memory in a single system image
- 5th generation NUMALink interconnect with upto 120 Gb/s bandwidth and latency <1 microsecond.
- Paired node 2D torus
- Based on Intel Xeon 7500 series
- 18.5 Tflops per cabinet



Altix UV 1000

- SGI Altix ICE

- X86 based, very compact form factor, custom cluster
- Various topologies based on Infiniband

