

CERN Openlab

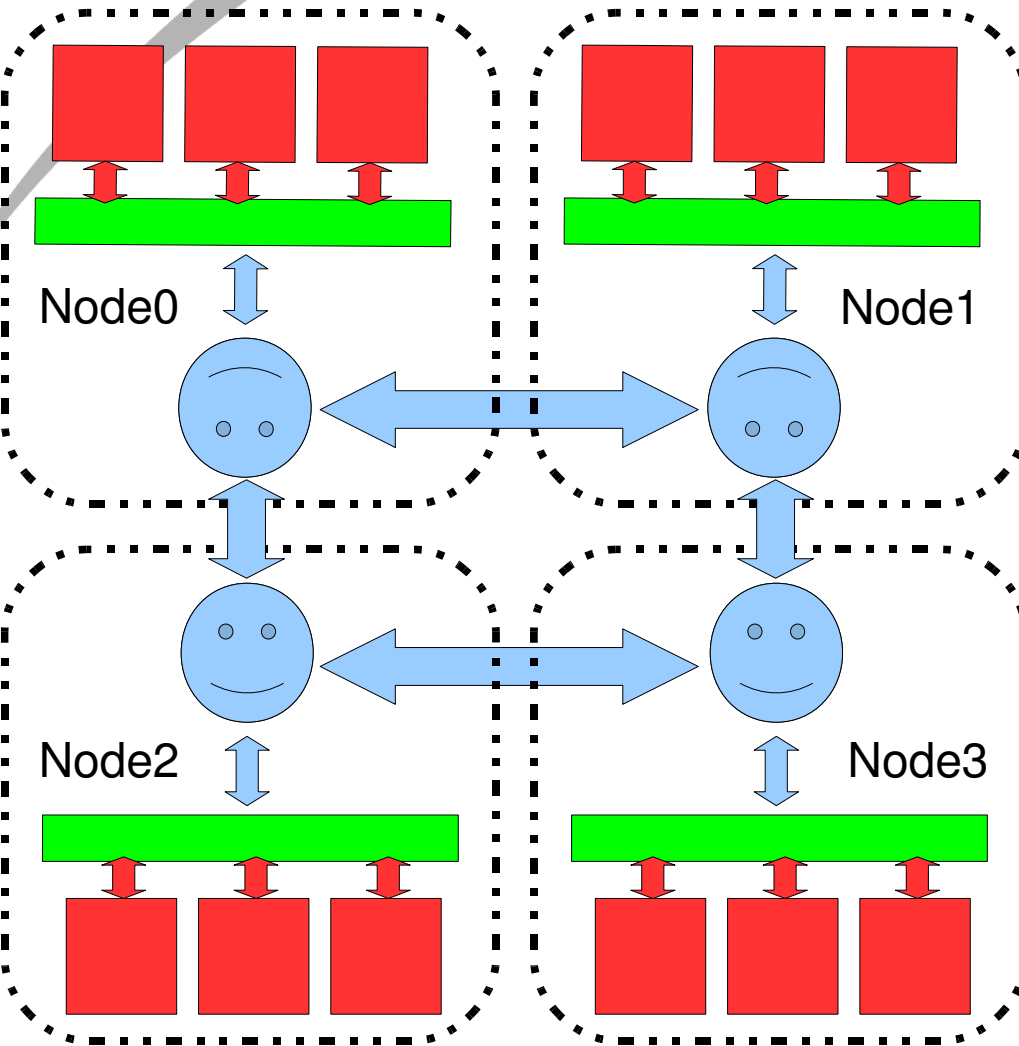
Platform CC
Minor review



Julien Leduc
Openlab Fellow




NUMA architecture divides in multiple nodes with access to local and remote memory, at a cost for remote memory.



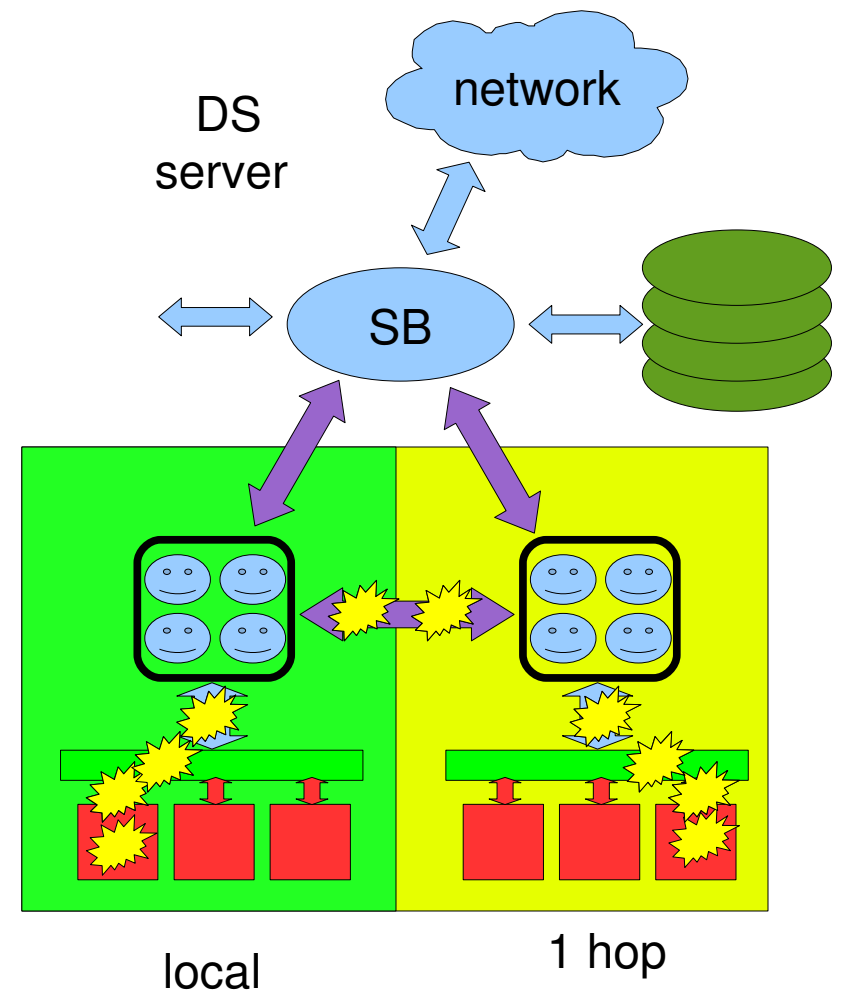
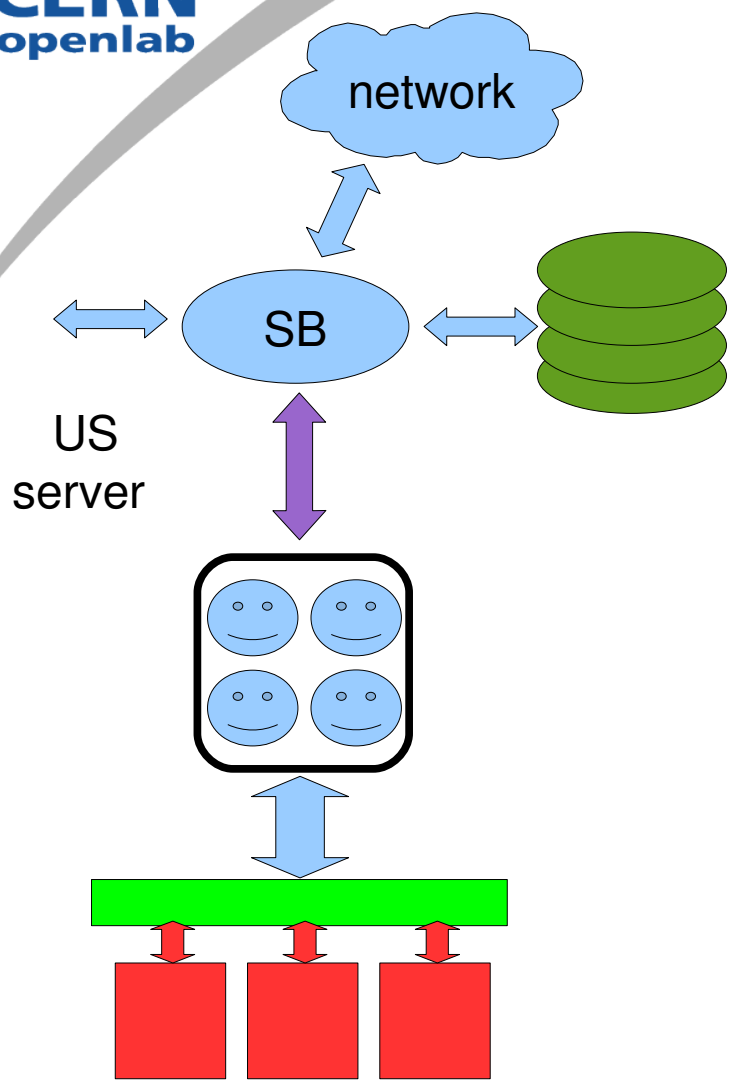
- Increasing performance now means more and more cores
- Both CPUs and memory don't see a boost in frequency anymore
- NUMA nicely solves these issues, offering a scalable design for future many-cores systems

Nehalem microarchitecture

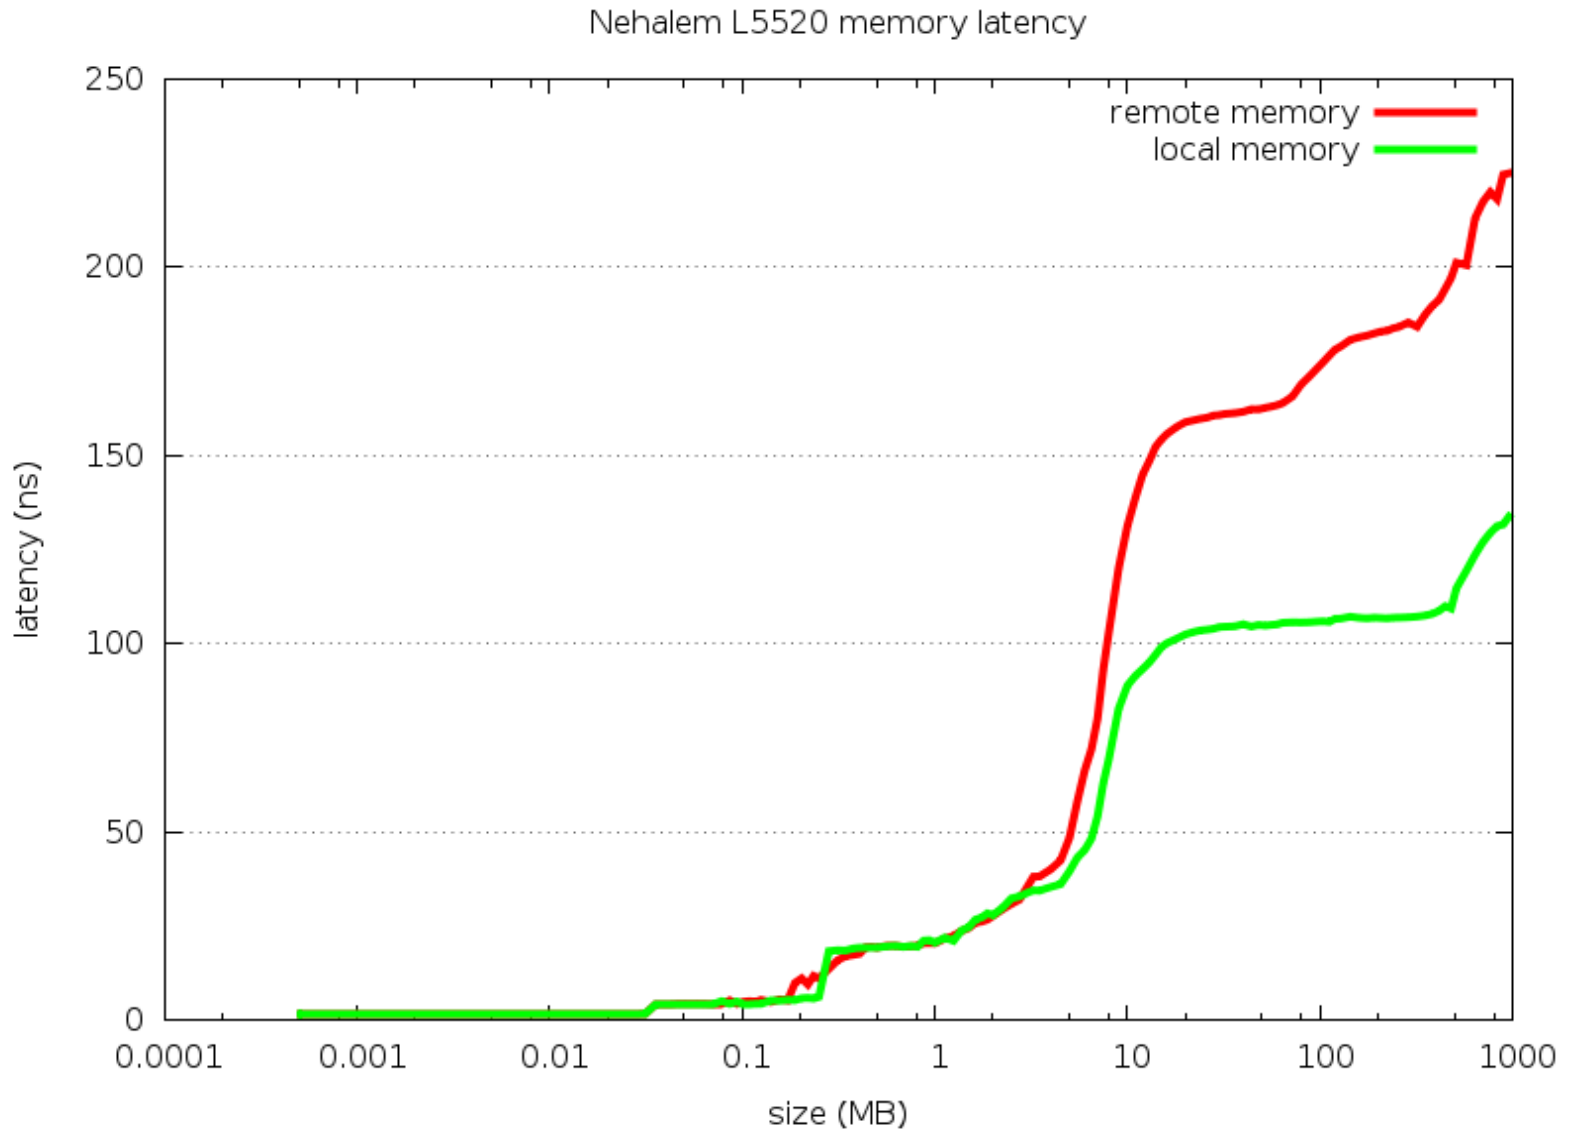
- Nehalem microarchitecture is equipped with an Integrated Memory Controller, and some QuickPath Interconnect links (1 for workstations, 2 for DS servers, 4 for MS servers)
- Microarchitecture allows scalable design for servers



Nehalem microarchitecture Single and Dual Sockets designs

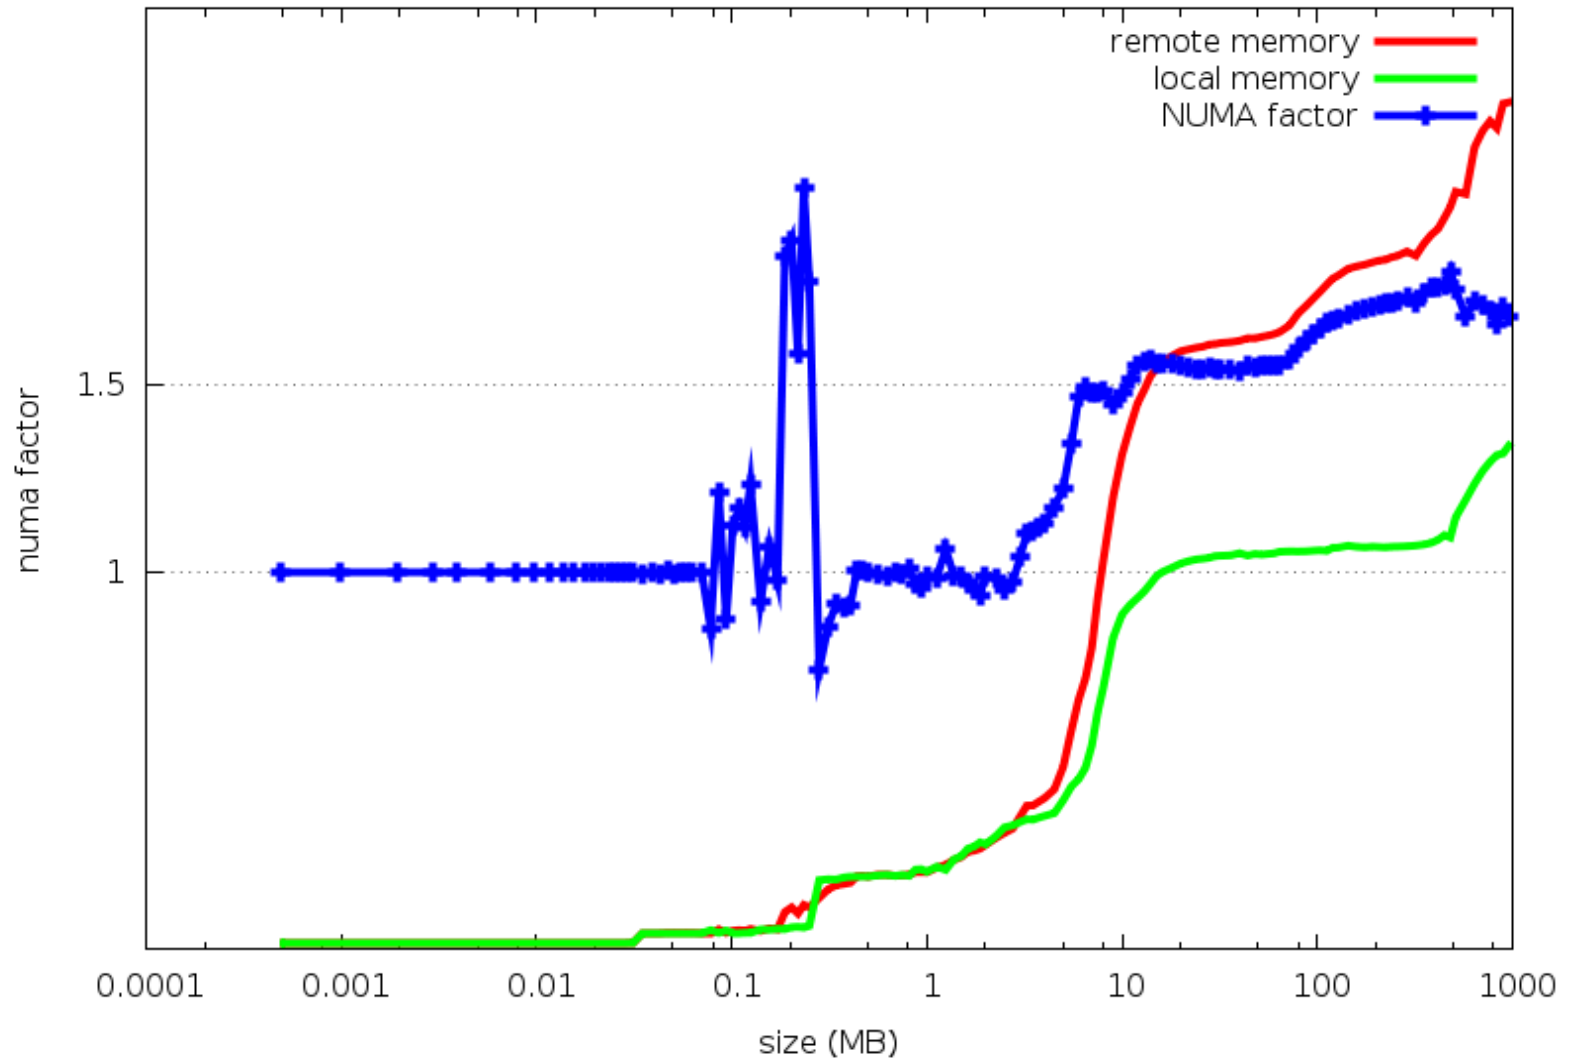


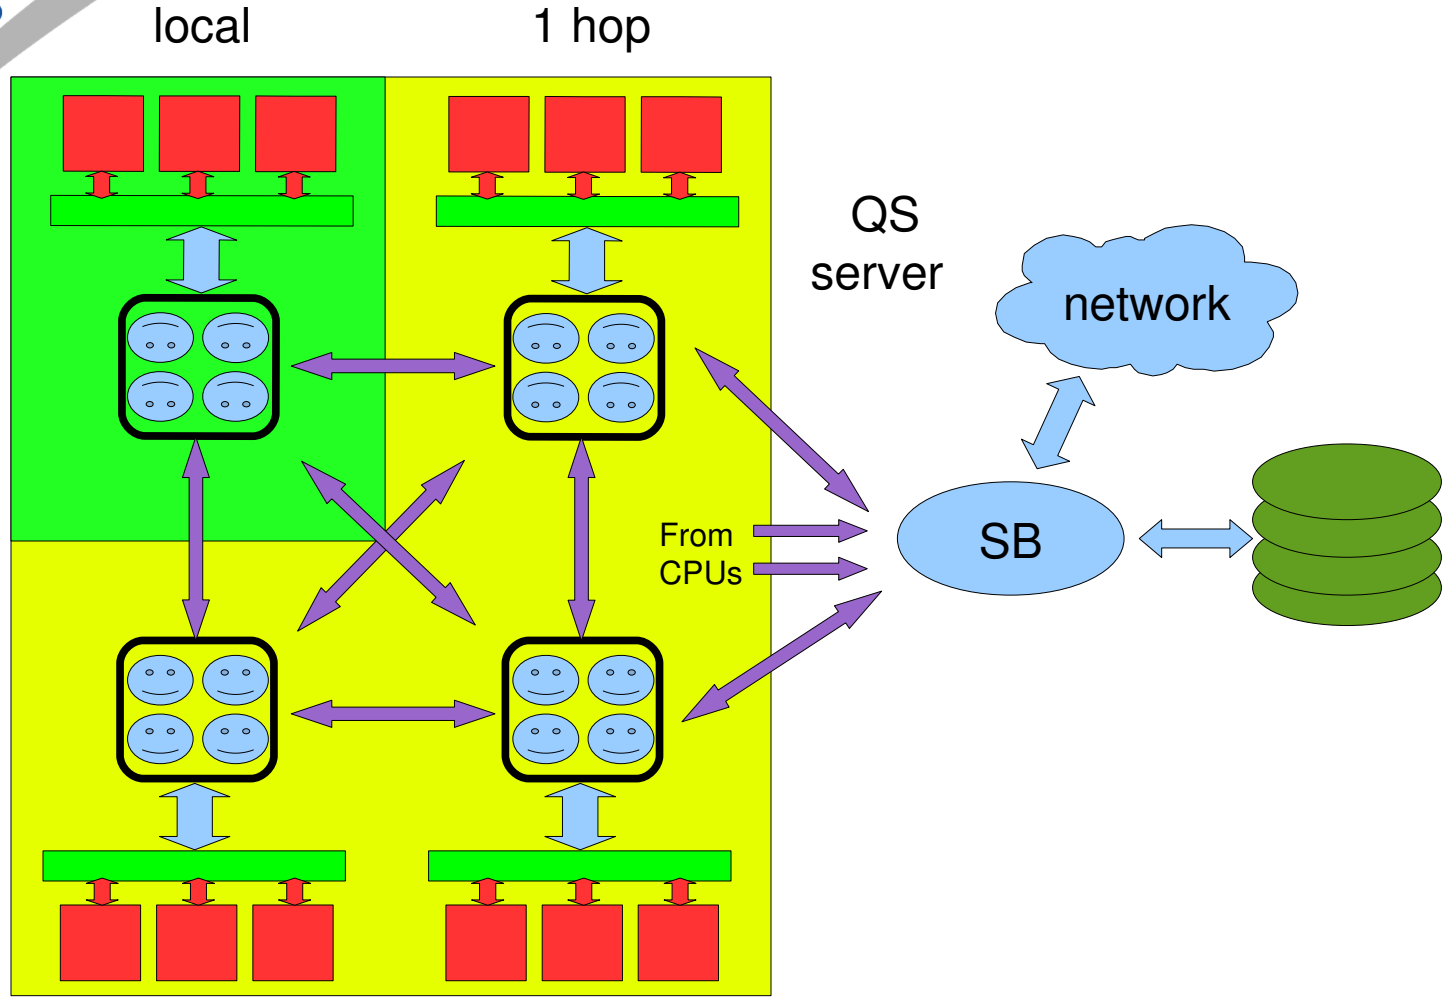
Memory latency measurements



Memory latency measurements

Nehalem L5520 memory latency





- Computational Fluid Dynamics benchmarks using StarCCM software on Nehalem EX
- Physical problem was a helium leak inside an underground tunnel
 - Compare performances between
 - Corei7 (SP NHM system)
 - NHM-EP (DP system)
 - NHM-EX (QP system)
 - Preliminary results have just been produced and need some investigations



Questions?