



How should petascale systems change what we are doing?

*Ray Bair, Director
Argonne Leadership Computing Facility
Argonne National Laboratory
and University of Chicago*

May 16, 2007



*Argonne National Laboratory is managed by
The University of Chicago for the U.S. Department of Energy*



Petascale Era - a Time of Many Changes and Challenges

- ASICs with burgeoning transistor budgets
 - Many cores, large caches, high perf memory interfaces, added functional units
 - What mix is best for petascale science applications?
- Power consumption is soaring
 - 18, 24, 40, 50, 100 KW per rack
 - What tradeoffs in equipment vs. power/heat densities optimize ownership costs in science facilities? Or mesh with facility schedules?
- Limited consensus on how to best feed power to computers today
 - HPC market moving to 480V 3-phase AC, but servers use 208VAC or 48VDC.
 - Some power supplies drop after short sags, others adhere to more robust standards
 - Can we drive to standards that minimize our TCO?

Petascale Era - a Time of Many Changes and Challenges

- Manufacturer plants are too small to house petascale systems
 - Cannot fully test system until delivered
 - Cannot test I/O without customer file system
 - How to enable vendors to complete development and testing at scale, in our facilities?
- It is hard to agree on acceptance tests
 - Manufacturer has their legacy tests and engineering tests
 - Customers have their own legacies and testing ideas
 - How can we get good coverage without duplicating work?
 - Who will develop the tests that exercise the new system scales?



Petascale Era - a Time of Many Changes and Challenges

- Reliability, Availability and Serviceability system are parochial
 - Compute engine has one RAS, file servers have another system, SANs have yet another approach, network switches another, building systems another
 - Can we converge to standards-based APIs and data management?
 - Can we build holistic RAS analysis and fault detection/prediction capabilities?
- Scalable systems software, libraries and tools require a lot of work to support sustained petascale and peak exascale systems
 - There are too few developers and too little time to have proprietary efforts for each brand or product line
 - How can the community organize itself around open source to meet the needs?

Many petascale challenges - just a few examples given

- How to best use the ASIC real estate?
- How to best engineer next gen. powering and cooling?
- How can vendors develop and test at full scale?
- What are the reliable petascale site planning principles and tools?
- Can acceptance testing be a science, with well understood goals and community tools?
- Can RAS standards enable holistic systems?
- What leverage can we get on the software for all these petascale systems?

Every one of those areas would benefit from collaboration

- Vendor-customer partnerships in petascale system design and development
- Joint TCO analyses of packaging, cooling and deployment
- Manufacturer-Community development of better testing protocols
- Open Source software for petascale systems, top to bottom
- ...

