

Testbeds and Large-Scale Production Petaflops-class Hybrid Systems

Description:

Early access to hybrid systems at multiple sites is highly desired to jumpstart this technology and applications development for it. The HMC community needs several small early systems at multiple sites for testing purposes very soon, likely within a year. These would be used for prototyping of applications, libraries, and system and runtime software. These small testbeds should be a least 4 nodes or more for use in architectural testing but likely need to be in the 10 to 100 node size to be useful for applications development and prototype science testing. These machine need to provide open access to the entire community of HMC interested parties.

In addition to actual testbeds, very early access to hardware simulators would be useful before such hardware actually becomes available. Once actual testbeds become available, simulators are viewed as not quite as useful but can still be invaluable to work out difficult implementation details and track down correctness bugs or performance issues seen on real systems.

The Applications scientists want to see actual DOE plans for one or more big open production systems of significant scale in order to be motivated to work on hybrid apps. They want to see a system 10x larger than the largest of the current TOP500 systems (Jaguar at 2.3 PF and Roadrunner at 1.4 PF) in about 3 years. It was suggested that smaller interim hybrid systems of 100 Teraflops to 1 Petaflop in 2 years could provide testing at reasonable scale and for entry level hybrid production, but the scientists need to ultimately see and benefit from the larger system to enable Science at a competitive scale in 3 years.

Topic	Urgency	Duration	Responsiveness	Applicability	Timeline
Hardware Simulators	Critical	Near	Moderate	Narrow	Immediate
Multi-site testbeds (4 to 100 nodes each)	Critical	Near?	High	Narrow	Immediate
100 TF to 1 PF interim system (needed option?)		Near?	High	Narrow	Soon
15 to 25 PF Production machine			High	Science	Soon

Relations to other TCs

- Testing of architectural features and design space of hybrid devices
- Early Programming Models evaluations
- Application prototyping and algorithm development
- Testing to enable and validate Performance Prediction Models

Related Projects

- Proposed ORNL system
- Cerrillos (“small” 160 TF Roadrunner system) at LANL
- Lincoln & AC at UIUC
- NSF Keeneland project
- Telsa-10 and Tesla-20 “Fermi” clusters

Specific comments from each Technical Committee’s report-out in regards Testbed systems:

Apps:

- Hardware simulators are useful before hardware is available
- As soon as hardware is available, we need a few prototype nodes per site, preferably one per developer
- Small testbeds of 10-100 nodes within a year
- Leadership platform that is 10x more powerful than today’s fastest supercomputers within 2-3 years

Programming Models:

- Testbeds: a large variety of small systems to test cross-platform applicability
- Clusters: useful to evaluate programming models (e.g. PGAS), but only up to a point
- Stability of development and execution environments

Performance:

- No specific testbed comments in the report-out

Architectures:

- Open access to entire community
- Multiple sites
- Application and system software development
- Production systems of significant scale (100-1000TF)
- Hardware evaluation systems (~50TF) x 1-3
- Four-node systems x N for specific, small-scale development use