

the Exascale Report™



The voice of the emerging exascale community

<http://TheExascaleReport.com>
December 2012

2013: Time to stop talking about Exascale

A letter to the emerging exascale community

By: Bill Gropp, General Chair, SC13

Everyone reading this is a believer in the power of computing. We take for granted that the computing power of the highest performing systems needs to continue to grow at the same rate in order to meet the needs of society. Yet this is not obvious to others.

Worse, a focus on an arbitrary performance measure (ExaFLOPS) rather than the ability to solve the most pressing problems can (and does!) label us as focused on the technology rather than what can be accomplished with it. In turn, this makes it harder for us to make the case that Exascale deserves funding over, say projects in big data.

So, if we stop talking about Exascale, what should we be talking about? To me, its obvious — we need to better discuss and describe the challenges and opportunities that we face, from basic science to commercial applications, and the essential role that high-end computing plays in solving those problems. Those problems need to be identified first, then their computational challenges identified and quantified, to show that there is a need for significantly faster systems.

Too many of the efforts to date have assumed an exascale system and then asked "What could you do with it"? The danger here is that approaches that don't really need Exascale are proposed — such as simpler, sub-optimal algorithms or ensemble studies. These do more harm than good, as an outsider can easily debunk the claimed need for Exascale.

Don't get me wrong; I believe that we need much more powerful systems to solve the problems that we face, whether it is understanding the basic operation of life and the universe or developing the technology for a sustainable environment. But the need for Exascale must clearly come from the problems.

We can start by moving away from a focus on FLOPS (especially as measured by benchmarks that we know are misleading but that the greater computer science community thinks we still take seriously) and focus on solving the hardest, toughest, most challenging computational problems. This also provides the best guidance and rationale for the development of the new technologies needed to realize the much faster machines we all believe are essential. Yes, not having such a simple metric as ExaFLOPS makes it harder to quantify the goals, but we all know that an effective HPC system can't be described by a single number.

To make Exascale a reality, we need to stop talking just about Exascale.

William Gropp
General Chair, SC13

Director, Parallel
Computing Institute

Deputy Director for
Research
Institute for Advanced
Computing Applications
and Technologies

Paul and Cynthia Saylor
Professor of Computer
Science

University of Illinois
Urbana-Champaign



© 2012 The Exascale Report™