

# the Exascale Report™



*The voice of the emerging exascale community*

<http://TheExascaleReport.com>  
March 2013

## MORE Truths About Exascale

*By: Mike Bernhardt and Doug Black  
March 2013*

*“Cuts to basic research threaten our scientific leadership and economic competitiveness at a time when investments in science and technology are more important than ever to our Nation’s future prosperity. At a time when other nations are intensifying their investments in R&D--recognizing the centrality of such investments to their prosperity, security, and international competitiveness; America, once the world’s undisputed leader in R&D, is hobbling its research efforts.”*

Director of the Office of Science,  
William F. Brinkman  
U.S. Department of Energy  
Subcommittee on Energy & Water Appropriations  
U.S. House of Representatives  
March 5, 2013  
Speaking on the impact of sequestration

This article expands on an earlier editorial titled, “The Truth(s) About Exascale.”

<http://theexascalereport.com/content/2013/truths-about-exascale-0>

With feedback from the emerging exascale community, we are pleased to present the next three ‘truths about exascale.’

*Unattributed quotes appearing in The Exascale Report have been provided by sources wishing to remain anonymous.*

**The truth is (#4) – The U.S. is unsure about Exascale.**

In our previous article on this topic, Bill Harrod, Research Division Director of the U.S. Department of Energy’s Office of Advanced Scientific Computing Research, stated, *“I am most excited about the capabilities that will be enabled by exascale technology. The future building of exascale systems will be a natural consequence of these capabilities.”*

Bill emphasized a key point that we seem to keep missing. It’s all about the science race – not the technology race. However, in our humble opinion, the two races are conjoined – never to be separated.

*“For the U.S. to be unsure about exascale / extreme scale indicates a lack of the basic understanding that pushing the extreme end will reap tremendous benefits all along the way – or a disregard for the importance of advancing scientific discovery with some sense of urgency.”*

Is the U.S. unsure about exascale? This discussion has been brewing since SC12 when the U.S. Department of Energy announced it had awarded a grant to the Council on Competitiveness to study the effects of extreme computing on U.S. competitiveness.

<http://www.compete.org/news/entry/2429/department-of-energy-awards-grant-to-council-on-competitiveness-to-study-the-effects-of-extreme-computing-on-u.s.-competitiveness/>

The question begs to be asked. Do we really need a three-year, \$914,000 study to identify the potential impact on U.S. competitiveness of extreme scale, or exascale computing? Obviously someone at the Department of Energy thinks we do. However, quite a few people think

# the Exascale Report™



## *The voice of the emerging exascale community*

it's a waste of time and money. What is there to question – or research here? The Council's own white paper from four years ago (March, 2009) states, "For U.S. leading manufacturers, to out-compete is to out-compute." Sure seems like extreme scale computing would fit in the "out-compute" category quite nicely.

The following quote refers to the DOE award to the Council on Competitiveness for studying the potential impact of extreme scale computing on U.S. competitiveness. This is from one of our many sources in the nation's capital wishing to remain anonymous, and seems to reflect the opinion of a large number of HPC stakeholders we recently talked with.

*"This is a ridiculous and sad use of much needed funds. Pushing forward toward any levels of extreme scale computing will reap tremendous, positive benefits for science and industry. It is the forward-looking, fresh slate research required to investigate extreme scale computing that will lead to breakthroughs in available HPC systems and technology. Industrial competitiveness, not just in the U.S., but on a global scale, is directly linked to technology. The drive toward Extreme Scale computing will force us to push the boundaries in many areas of technology from compute to memory to storage to data movement. This research grant is equivalent to the U.S. government's mythical or real \$600 hammer. It's like asking should we have cleaner air? Should we have cleaner water? Should we try to cure cancer? To say it's a poor use of government research money is a gross understatement."*

When I asked one political representative (who does not like to be called a politician) why he doesn't feel there is a compelling need to push for more aggressive exascale / extreme scale funding, he responded with this question:

*"With limited research dollars available, would funding be better spent on Cloud Computing or*

*Big Data where there are more commercial implications?"*

Well, not meaning to over simplify this question, but Cloud Computing and Big Data research will be driven forward quite aggressively by private industry. These both represent shorter term revenue possibilities and complement the product strategies of many companies. However, that being said, all computing paradigms will benefit from the research conducted under the umbrella of extreme scale / exascale research.

The entire exascale discussion is unfortunately morphing into 'where else could we put our money' discussions and having a negative impact on the HPC community – slowing down our progress – and impacting competitiveness across the board.

We are seeing a tension – a conflict – among the various groups who equate HPC and exascale with 'big iron' and those who argue that the exascale journey will produce numerous technology advances and drive both advanced computing and scientific discovery forward. The tension is most apparent among those who feel Big Data research (and funding) should take precedence over HPC.

To address this coin toss mentality, William Gropp, Director of the Parallel Computing Institute at the University of Illinois Urbana-Champaign gives us three points to consider:

*"First, HPC has become a key part of science and engineering and access to HPC systems, at all scales, is essential to continued progress*

*Second, other aspects of computing, such as "Big Data", are emerging as new transformational areas that will require their own, sometimes large scale, infrastructure*

# the Exascale Report™



## *The voice of the emerging exascale community*

*And Third, These are not incompatible. Big Data requires significant compute capabilities. While some operations on data are well-suited to simple commodity clusters, others may not be and require a more tightly coupled system.”*

Gropp concludes with, *“It shouldn’t be either-or.”*

The U.S. is certainly facing its share of economic challenges and those are only compounded with the recent burden of sequestration. Budget line items are being cut throughout government organizations, and critical research programs are in peril of being arbitrarily frozen or eliminated.

We asked Gropp to give us more of his perspective on this:

*“In a fixed budget, adding to big data means taking from somewhere else. I think we should be questioning this assumption of flat budgets (at the level of investment in computing). As scientists learn to do more with different kinds of computing, it makes sense to increase the investment in that infrastructure - not starve one as it becomes mature and widely used to support the next wave. In the current budget climate this will be hard to achieve, but trying to survive by freezing or reducing every separate line in the budget, without investing in ideas that may create new knowledge and capabilities, is the road to ruin.”*

Another of our HPC community luminaries, Dan Reed, also believes this doesn’t have to be a coin toss.

We offer this direct quote from Dan Reed’s recent blog that appeared on the Communications of the ACM site, <http://cacm.acm.org/blogs/blog-cacm/161741-exascale-software-just-a-few-orders-of-magnitude/fulltext>

*“The exascale hardware and software challenges are real. Do we pursue incremental extensions of current practices or step back and explore more radical and fundamental options? Each has*

*different advantages and disadvantages, which suggests we should probably pursue both, recognizing the costs. To be sustainable, an exascale research and development program must lead to cost effective and usable systems that are an integral part of the mainstream of semiconductor and software industries.”*

There is an underlying message here beyond the need to move on from this coin toss mentality. We can’t get to exascale – or the scientific discovery it will enable in any reasonable time frame without adequate government investment.

**The truth is (#5) – Private industry alone can’t get us there.**

If the U.S. government doesn’t step up to provide adequate funding for extreme scale / exascale research and development, can private industry carry the ball?

*“The level of research necessary to reach exascale-level computation by 2020 is far too cost prohibitive for any profit-oriented company to pursue. Private industry needs to focus on profitability and that means bringing products to market that have a chance of impacting the bottom line immediately.”*

Anonymous source,  
.gov organization

Intel represents what appears to be a case study to reinforce this discussion.

Intel has been a champion of exascale for several years and still stands by its position that the company will help field the first exascale-class systems by the end of the decade. The company has stated numerous times in the past that their position is to lead the world into the era of exascale – but they have also said they can’t do it

# the Exascale Report™



## *The voice of the emerging exascale community*

on their own – despite having some of the deepest research pockets on the planet.

Recently, we learned that Intel's R&D engine, Intel Labs, has dismantled its Extreme Scale Computing research program.

So what does that mean? Is Intel following the lead of the U.S. government and giving up on exascale?

*“Not at all,”* according to Wilf Pinfeld, Intel's Director of Extreme Scale and Government Research Programs.

*“Work on Exascale remains a priority in both Intel Labs and the product business. Transition of technology from Labs to product group is a good sign that the Labs are influencing product direction and that there is good communication and tech transfer happening between these entities. From an Intel labs' perspective, we are working on exciting projects and are committed to transferring relevant work to future Intel products.”*

Let's look at this a little more closely.

Like most commercial companies, Intel feeds its massive technology research engine with revenue from the product groups. If the government doesn't provide funding for extreme scale research that allows private industry to throw significant resources at it, then companies such as Intel will try to get us there with product roadmaps that drive shorter-term revenue goals – which may or may not get us to that next big milestone.

The lack of government funding forces the HPC community down an evolutionary path – one on which we rely on product sales to fuel the R&D engines.

Most HPC community leaders agree - breakthroughs in current HPC technology – the types of breakthroughs needed to move scientific discovery forward – will only be made possible when we look beyond current product offerings and supplement R&D budgets to encourage revolutionary thinking.

**The truth is (#6)** – The discussion of the computation barrier to exascale has been pushed to the back shelf.

So, while we get all wrapped up in these discussions about funding, and the typical technical barriers such as power requirements, we seem to have placed the discussion of the exascale computation barrier on a back shelf.

To this point, we hear from one of the most widely recognized voices of the HPC community, Professor Thomas Sterling at Indiana University.

*“Exascale computational science will require a sustained concurrency of approximately a billion operation issues (or completions) per cycle (or per nanosecond) on any single application.”*

He goes on to point out that more will be required to provide the necessary overhead functionality and to hide system data movement latencies.

Sterling believes this is a foundation for going forward. If we don't hit the computation goal, all the rest is a moot point.

Think about it. Where are we investing all of our energy and discussion today? Politics, budgets, and repeated studies that have been going on for decades. Sad – but true. Isn't computer technology supposed to be about computation?