the Exascale Report™
A Global Perspective
The voice of the emerging exascale community

http://TheExascaleReport.com
December 2012

The Truth(s) About Exascale

By: Mike Bernhardt
December 2012

The Truth(s) About Exascale

As I reflect on 2012, and most recently, on several very interesting and colorful conversations from SC12, I am feeling encouraged that we have indeed made key progress in our journey toward a new and exciting world – one where science is driven forward by extreme levels of computation.

I believe 2013 will be a pivotal year for extreme scale research; it will solidly move us closer to realizing much needed scientific discovery made possible by unprecedented levels of computation, for which the next milestone is, of course, the exascale level.

Learning is a journey and progress can be measured in many different ways – and in this case, I believe realization is progress.

A growing number of people in the community are lining up to accept and articulate what I refer to as several important 'truths' about exascale.

While we have not yet attempted to actually build physical exascale-class systems, we have indeed made tremendous progress in conceiving and evaluating ideas and approaches that will move us in the right direction for continued strategic research in areas such as power management, architectural design, memory and storage. Much of this research will lead to improvements in general purpose HPC solutions - delivering the additional advantage of advancing our capabilities in petascale computing.

Exascale, like the previous quests for teraflops and petaflops, is a journey not to be taken solely for the sake of developing new computing technology. We must not lose sight of the true purpose of our quest for exascale-level computation – the underlying need to move technology forward in order to make possible new scientific advances that will have a profound impact on all aspects of life on this planet.

Path Along the Stream by Fortunato Ornaghi

Progress can be measured by what we learn – and also what we eliminate. **Progress can be measured by reflecting on where we've been and by gaining a better view of the path ahead.**

FACT: The research undertaken to get us to exascale will result in more efficient implementation of petaflops.

And, lest we forget, petascale provides the stepping stones on the path to exascale.

the Exascale Report™
A Global Perspective
The voice of the emerging exascale community

The truth is (#1) – the journey, and what we learn going down that path, is far more importnant than the achievement of building any one or two systems.

Determination will show us the way

HPC market dynamics have changed from the early 1990s when government funding was the primary impetus pushing the advanced computing market forward.

In contrast to those days, the diminishing commitment and funding for high performance computing research in the U.S over the past several years has taken a serious toll on morale among the nation's leading scientists and researchers. And this has happened while several other countries have announced significant funding commitments targeted at HPC and exascale-level research.

This is something my colleagues and I witnessed at SC12 in Salt Lake City – a polar difference in the morale of the U.S. attendees compared to the international attendees.

In general, the non-U.S. attendees were far more upbeat, enthusiastic, and positive.

And while the U.S. Department of Energy and the Department of Defense HPC Modernization Program were a limited presence, foreign government-sponsored booths and technical demonstrations were in strong attendance.

Speaking in general terms, morale throughout the U.S. government HPC research community is very low. These are the researchers and scientists who historically have been part of the "push" equation to help drive the community and U.S. competitiveness forward. Now, many of them feel abandoned as they face uncertainty around programs and even their jobs.

Despite the low morale and the uncertainty of their own future employment, the scientists and researchers I interviewed from the national labs and research organizations made it very clear that HPC progress has not stopped. While it may be moving forward at a painfully slow pace, it is indeed moving forward. Many of these scientists have had to put their fears and frustrations aside and work with what they have – changing the very dynamics of the advanced computing market while they wait for the U.S. leadership to catch up.

The 33rd president of the United States, Harry S. Truman once said, *"America was not built on fear. America was built on courage, on imagination and an unbeatable determination to do the job at hand."*

From Push to Pull

When there is no force 'pushing' us forward, we need to use our combined strength to 'pull' ourselves forward.

When policy makers and funding agencies eventually determine that extreme scale computing is the foundation of progress, and reaching exascale-level computation is a 'necessity' – only then will we see a change in attitude and budget allocations.

In the meantime, it's up to the scientists and researchers and community leaders to cobble together whatever collaboration they can. Many of you have heard the very old and wise saying, "Necessity is the mother of invention."

The U.S. has enjoyed the distinction of being the world leader in advanced computing and technology for many years. Even though the bragging rights for world's fastest computer have changed hands a few times, leaders always point

the Exascale Report™
A Global Perspective
The voice of the emerging exascale community

out that it's only temporary – and are quick to also point out there is little value in "stunt machines" – i.e., machines designed to achieve high performance on specific benchmarks but of limited utility in real-world applications.

The conversation completely misses the real significance of why more computational power is needed.

"Over the past several years, it has been the unswerving dedication and determination of many dedicated scientists and researchers, driven by the realization that we must never slow down or give up on our quest for quantum leaps in computational progress, that has created a ground swell of demand. The demand is for new science – not just for bigger and faster machines."

"Today, with limited funding for HPC research programs, the sheer determination and dedication of hundreds of scientists and researchers are really what is keeping the U.S. in the game."

The truth is (#2) – If we can't get there by pushing the technology forward, we need to work together and pull it forward. Determination will lead the way.

***Barge Haulers on the Volga*,** 1870–73 (State Russian Museum, St. Petersburg) by Ilya Yefimovich Repin

Exascale: Wrong Name – Wrong Metric

This is something I've been hearing more frequently. When you ask most people what they think Exascale is all about, they immediately describe a massive computer system with its own power generating plant and a small army of technicians.

Perhaps the conversation in the community has been part of the problem in that too much discussion has focused on the race to build a system – and not the necessity of scientific discovery.

If we drop the exascale label from these discussions, and start talking about Extreme Scale Computing, we conjure up much different images.

Intel and IBM figured this out quite some time ago. Intel has their Extreme Scale Computing research as part of the massive Intel Labs, and IBM built their Deep Computing Institute under the umbrella of IBM Research. In both of these cases, exascale is a research milestone.

Readers – please take note. This view needs to be spread to larger audiences, which include policy makers, funding agencies, and the general public. Setting an artificial date for when we think we will hit exascale-level computation should be seen only as putting a stake in the ground to keep us moving forward and to remind everyone that there should be a sense of urgency – the same sense of urgency we should apply to finding a cure for AIDS or cancer. It's a race only in that we need to get there as fast as we possibly can – but not for the purpose of claiming bragging rights.

According to Bill Harrod, the Research Division Director of the U.S. Department of Energy's Office of Advanced Scientific Computing Research, *"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. One of my goals is marketable technology. That's where and how we can have an impact on science – and on the economy."

Following his keynote presentation at SC12, I was able to spend some time talking with Bill to get his views on the topic of a 'race to exascale' – a phrase used widely in the public media and recently in several pieces of Congressional communication.

According to Harrod, *"I think an emphasis on a 'race to exascale' is counterproductive. The focus should be on the ability to deliver the science, not*

Bill Harrod, DOE

the ability to "ace" the Linpack benchmark, which is the current standard benchmark for assessing computational speed, and the current criterion for ranking the TOP500. Linpack is a nice benchmark – I have no problems with it – but it is only one component of the complex set of criteria for assessing the performance and utility of a computing system. When you call the exascale effort a race, you encourage the development of stunt machines"

"We need to get away from this focus on Linpack as the sole criterion for assessing the "goodness" of a computing system. Although it is a good gauge of raw peak speed, and the sole criterion for ranking the TOP500 international supercomputers, it was never meant to be a gauge of a computer's utility and worth to real-world applications and problems. And we need to stop thinking of exascale as a race. Development of exascale technology is far beyond 'business as usual', or, rather, 'research as usual'. SC12 has only served to emphasize that fundamental exascale science is greatly benefitted by international collaboration. As a global community, we need cooperate and share basic science results. Given the current funding situation in the U.S., we in the U.S. need to educate policy makers, funding agencies, and the general public to elucidate the importance of this most critical of enabling technologies for future progress."

"Because of its prominence in the world of supercomputing, the TOP500 ranking will always keep alive the notion of the race to win the top spot. We have to recognize that other countries could very well beat us in terms of bringing an exaflops-level stunt machine to market first. We accept that. As I've said, we must not view this as a race – this is an important step for science and we need to do it right."

Editor's Note: With several countries believing they can beat the U.S. in terms of achieving exaflops, I seriously doubt we will sway the marketing forces from their discussion of a 'race', and therefore, we face the challenge of repositioning perception of what an exascale 'race' is really all about. If we must talk about a race, think of it more like the international 'race for the cure' – a race with an urgent and vitally important purpose. If we are going to describe a race, let's make sure it's the right race.

the Exascale Report™
A Global Perspective
The voice of the emerging exascale community

The Tortoise and the Hare by Jan Wildens

This oil-on-canvas titled The Tortoise and the Hare, from Aesop's "Fables" gives us a serene and pastoral scene consisting of a path and a brown hare dashing in from the bottom left corner. The tortoise is absent – in fact one might not even know the subject matter were it not for the title – and therein lies the genius. At the twilight of the race, the tortoise is already far and gone, and it is too late for the frightened hare.

The Tortoise and the Hare has often been interpreted as 'slow and steady wins the race.'

= = == = = = =

Harrod continued, *"We are taking a systematic, methodical approach towards the realization of exascale-level computation so that the resulting system will be one that has generality and can best serve science. As I've said, I am focused not so much on "a system" but more on the enabling exascale technology, which I believe will scale down – and have a far reaching impact all across technical computing and scientific discovery – not just at the extreme high end."*

An important message for all of us to take away: Exascale is a milestone of extreme scale computing – it is not an end point – and it is definitely not about any one system.

As expressed by Michael Resch, Director at HLRS and a Professor at the University of Stuttgart:

"The TOP 500, while it has some useful purposes within the community, can draw us in the wrong direction. Outside of the HPC community,

Mike Bernhardt, The Exascale Report, interviewing Michael Resch of HLRS

politicians may look at it and question why their investments in important systems and technology have not placed the systems they have funded at the top of this list. Internally, it can be useful to see the changes in parallelism and microcomputer architectures, but not to gauge scientific progress – and clearly not to determine what country or organization is truly leading in terms of practical, important HPC systems or the application of powerful systems to solving useful problems."

the Exascale Report™
A Global Perspective
The voice of the emerging exascale community

If we focus on the technology spawned from all the exascale research, we may very well win the science race, and that's the only race that counts.

The truth is (#3) – *Winning a race to build an exascale stunt machine is not a worthwhile goal, and stunt machine benchmarks will serve no purpose in helping us explore the frontiers of science.*

Our thanks to Bill Harrod for taking the time to talk with us following his keynote at SC12.

While this publication has no qualms with pointing fingers and telling it like it is, we see Harrod as a responsible, dedicated and pragmatic leader. It is our opinion that under his direction, and despite the inability of Congress and others within the administration to offer and defend an adequate, long-term budget, progress in extreme scale computing will continue to march forward – driven by sheer determination and passion. But even a solid leader like Harrod can only hold things together for so long. HPC funding is crucial.

Without the necessary funding, the U.S. efforts will start to erode as scientists and researchers will very likely start migrating from government positions to private industry as they seek stability.

We agree with Harrod. The urgency of what we can accomplish with extreme levels of computation has been lost in the discussions of a 'race' to build the world's fastest supercomputer.

Moving computational science forward is not a short-term or one-time project. It is in fact long-term, or more accurately, ongoing and never ending. Economic progress will be driven by innovation, and innovation is fueled by both determination – and technology.

Taking a more solid and steady approach to funding HPC and extreme scale computing programs will ultimately result in the creation of more widely applicable technology and will have a far greater impact on advancing science – and that's a race in which we all want to participate.

We close with this final comment from SC12 Keynote Speaker, Dr. Michio Kaku. *"If the U.S. is serious about turning around its economy, they need to invest heavily in technology. Technology drives everything."*

Indeed.

<<<<< >>>>>

We'll explore more of the "Truths About Exascale" in the next issue.

To all of our readers, we wish you a very **Happy New Year**. May 2013 carry you beyond your wildest algorithmic dreams.

1251 Goff Rd. ~ Forest Grove, OR 97116 ~ info@TheExascaleReport.com