

# the Exascale Report™



*The voice of the emerging exascale community*

<http://TheExascaleReport.com>  
May 2013

## Intel: The Path to Exascale is Inevitable

*By Mike Bernhardt*

Preparing to design, build, deploy and use exascale systems requires a holistic approach and a delicate balance – an equilibrium – of three building blocks, each of them so critical to the development of an exascale ecosystem that they all might be considered keystones. However, by themselves, none of these will get us to exascale. It is the balance of the three, working in harmony that will enable the next great era of high performance computing,

The three exascale building blocks are:

- Technology
- Human Resources
- Funding

With anticipated timelines moving farther out, and no committed national funding program for exascale in place, there is skepticism that we can pull these building blocks together and actually field an exascale system by the end of the decade. Yet, amidst the chaotic arguments around exascale progress, we've found a calming, confident voice to bring a balance and insight to this discussion.

According to Diane Bryant, Senior Vice President and General Manager, Datacenter and Connected Systems Group at Intel, *"We've done it before and we can do it again. We just have to repeat history – and I think we are pretty well situated to do that."*

Not everyone agrees. Some people say we need revolutionary changes beyond anything we've done in the past. And some industry opinion shapers are starting to say there is no way we can reach an exaflops milestone by the end of the decade.

Bryant makes an effective argument that this is semantics. To her way of thinking, what seems to be 'revolutionary' is 'evolutionary' to others. Perhaps what we refer to as evolution in HPC is actually made up of breakthroughs and revolutionary milestones our industry has come to take for granted.

Inarguably, exascale systems will require significant breakthroughs – revolutionary changes – in power management and memory architectures that go far beyond what we can hope to accomplish through an evolutionary process of simply scaling up. And this does not even scratch the surface of the complex system software and application development challenges that are required.

Regardless, raising the bar on HPC capabilities has always been driven by new breakthroughs, and it has always been expensive.

According to Bryant, "As an industry, both public and private, we are going to invest to get us to exascale. Everyone is going to invest. It's inevitable."

In this feature Exascale Report interview, we talk at length with Intel's Diane Bryant to better understand her positive and enthusiastic perspective on the inevitable realization of an exascale milestone for HPC.

© 2013, The Exascale Report™

# the Exascale Report™



## *The voice of the emerging exascale community*

**The Exascale Report:** *Recent workgroups have concluded that U.S. HPC investments are stagnant or declining while other countries are increasing their investments in this area. Several publications have written about the lack of U.S. Government exascale funding, and the likelihood that exascale systems will appear in other countries first – most likely China.*

*From Intel's perspective, does this even matter? What impact would this have on the nation's scientists and researchers and what would it mean to U.S. competitiveness?*

**BRYANT:** The U.S. government's investment in the move to exascale is very important. The U.S. government has been a driving force year after year after year in advancing the state of computing. That's important for the industry. It's important for the United States, for public interest, and for economic development. And, at large, we will collectively, holistically continue the move to exascale.

Government can't do this without industry, and industry can't do this without the support of the government. Intel invested \$10 billion in R&D last year, and we continue to make a significant investments year over year in R&D. We take the High Performance Computing and the supercomputing space extremely seriously. If for no other reason, technology starts at the top and like water, falls down. Everyone benefits from that trickle down – from the highest supercomputing

investment down to commercial HPC down to the desktop. So we will continue to invest.

From an international perspective, you certainly do see other governments – other geographies – recognizing the criticality of an investment by the government in supercomputing. It's actually very encouraging. It reinforces what the U.S. government has been doing year after year, decade after decade.

Certainly China's investment in supercomputing has been impressive, and India now has their five-year plan and a path to petaflops level computing. Despite the recent dispute between the U.S. and Russia regarding supercomputing technology, Russia is moving forward in HPC R&D as well. So this phenomenon is worldwide. The international activity just confirms how important it truly is that the U.S. government continues to invest. We think it's inevitable that the path to exascale is going to happen.



Diane Bryant  
Intel Senior Vice President

Does it really matter if the U.S. is not the first nation to field an exascale-class system?

I think it matters if the U.S. Government is not investing to win. Who actually "wins" is perhaps less important, but the investment in computing and the investment in supercomputing is critical in core. As to who has the number one system or who has the fastest system at any one time – I'm not sure that matters as much as knowing that the United States is continuing to advance its footprint and position in supercomputing. As I said, there

# the Exascale Report™



## *The voice of the emerging exascale community*

are so many aspects of supercomputing that make it critical to any given nation – from national security to public interest – healthcare, weather forecasting, etc. There are so many elements that go into it. And then just look at what this can do for economic development. We need to know that the U.S. government is continuing to invest in advancing the state of computing.

In terms of supercomputing leadership, the number one position at any point in time, as you look back on the Top 500 supercomputer list over the past four or five updates, the lead position has toggled between two or three countries. Is that what matters? No, I think what matters is the continued advancement.

The level of innovation and advancement that is needed to continue this advancement is a significant research effort from everyone involved. If you look at what we've (Intel) been doing, we just launched the XEON PHI co-processor last year to significantly improve the performance in parallel workloads. We're investing heavily in fabrics – in the interconnect and integrating them into the processors. These are big investments that were driven by research. So you really need a very healthy research arm fueling your engineering efforts in order to continue driving the investment. That research piece is critical.

**TER:** *Most people agree, the transition from terascale to petascale was accomplished by an evolutionary process or approach that relied primarily on performance improvements driven by increased transistor density and of course, higher clock rates.*

*Many in HPC, but not all, believe the transition to exascale requires a "revolutionary" approach to deal with the overwhelming technical constraints. It seems Intel's exascale strategy is owned by the product groups, and therefore, is, or will be, based on whatever technology Intel can bring to market as commodity products – which is a more evolutionary strategy.*

### **Exascale Report Factoid**

*Ten years ago China had no systems on the TOP500 list of the world's fastest supercomputers.*

*Today, China has 68 systems on this list*

*Help us understand how Intel views the exascale roadmap.*

**BRYANT:** That is an interesting question. Intel has a world class research arm – Intel Labs – and they look five years out and beyond. And then we have our development arm which is current day to five years. From the development or product side, we work very

collaboratively with the labs on bringing technology that they are investing in back into the development flow. So we have an incredibly structured product life cycle at Intel that starts with the research arm.

There has actually been a lot of innovation from the labs that has been funneled into our products. In fact, their success is in part measured on how many of their technologies make it into the product roadmap and so we have a very close collaboration. For example, let's take something as simple as virtualization technologies. This is something we now all take for granted in standard servers – everyone has virtualization technology – and that technology was developed in Intel's Labs and then moved into our product lifecycle and then into our products. That is a very robust cycle – so of course the product group is responsible for the final product launch, but we reach back into the labs to pull the technology.

# the Exascale Report™



## *The voice of the emerging exascale community*

As far as evolutionary vs. revolutionary, if you go back to when Intel entered into the server business in 1997, the top Intel Architecture-based server was 1.3 teraflops. Now fast forward to last year. The top Intel supercomputer was 2 petaflops. So from 1.3 teraflops to 2 petaflops, we increased the performance over that time by 1,500 times while increasing the power only 4 times and reducing the cost of compute 100 times.

It's easy to look back and say that was all very evolutionary, but to my way of thinking that was revolutionary – a 1,500X improvement in computing with only a 4X increase in power. So now we have to go from petaflops to exaflops. We “only” have to get a 50X improvement in performance at a 2X increase in power.

So, it's not as aggressive as what we've achieved historically. Sitting here today and looking back at what we accomplished, we all tend to refer to it as evolutionary. And I do believe the way we did it historically is the way we'll continue to do it. It is innovation at every single layer in the stack, and last year when we launched our Sandy Bridge product we doubled the number of floating point operations in a single clock cycle.

You tell me. Is that revolutionary or evolutionary? It felt pretty revolutionary for us.

But that's just one piece of the puzzle.

As you innovate at the transistor level, thanks to Moore's Law, at the CPU architecture, at the system architecture, at the software, hardware, solution level – all those innovations stack up to exascale computing and it takes all of us innovating at every one of those levels to get us there.

I think it's comforting to know. We've done it before and we can make the leap to exascale at only two times the power. We just have to repeat history and I think we are pretty well situated to do that.



**TER:** *I recently saw a [YouTube video](#) of Intel Fellow Eric Dishman \* talking about the future world of customized or personalized medicine. Getting there requires affordable gene mapping and that certainly sounds like an application area that will require extreme levels of computation. So in the case of the healthcare segment, how do Intel's HPC and exascale strategies align with your goals for changing healthcare as we know it today?*

**BRYANT:** Our investment in supercomputing and the high performance computing space is extremely well aligned with our investment in the healthcare space. The health industry has been slow to adopt technology, but thanks to the whole healthcare reform act and Obamacare, it is going to have to happen. You can't meet those performance goals if you haven't digitized your environment, so there is a big push now to modernize the health industry and move to a technology-based system. But this all comes with complexities.

# the Exascale Report™



## *The voice of the emerging exascale community*

There is complexity in terms of the massive amounts of data. Whether it's claims data or clinical data or consumer generated data or genome data, we are facing unprecedented, massive amounts of data. And the cost is still too high to be practical. If we're going to allow or enable personalized treatments for every individual, we need a far more cost effective capability. These are challenges that require solutions and we hope to address these through our HPC strategy.

Prior to this job, I was Intel's CIO. I ran our IT organization and for four years I got to see just how complex it is to deploy some of these solutions. To achieve a transformation of the healthcare industry, it has to get cheaper and it has to get easier. And we have to be able to manage the massive amounts of data.

Those attributes align incredibly well with our focus around high performance computing. We continue to invest at the high end, always searching for new innovations that deliver greater and greater performance at lower and lower cost of operations. So we have a very strong strategy. We're looking first at genomics and specifically the analytics around genomics. It's a great high performance computing problem. The goal is to get to where, per patient, you can sequence and analyze your genome for less than \$1,000 a person. We think this is the target to go after – where it becomes generally available. That process today is probably over \$300,000. So we have a long way to go. That's the cost element.

And then that solution stack is very complex. So we have to work with the ecosystem to drive down the complexity as well. We've done a lot in this area already. We actually talked publicly about it around a breakthrough we had working with Dell. So between us and Dell and The Translational Genomics Research Institute ([TGen](#)), we took

TGen's software stack, their analytics stack, and we parallelized it to run better on a Xeon Phi solution, and we took the process of sequencing the full genome, which is 5 terabytes of data and processed a full genome. We took it from seven days down to four hours. So we dramatically simplified it and reduced the cost but there's still a lot of work to be done. That level of industry collaboration completely leverages the work that's going on in the HPC space and the innovation that's going on there: drive down the cost, reduce the complexity and manage the massive amounts of data required to do analytics in the genome space.

It's a very well-aligned and an exciting area for us. The New England Journal of Medicine published a study that said the probability of you as a patient receiving a science-based treatment is no better than a coin toss. So there is a long way to go in terms of applying scientific analysis and modeling to healthcare. It's a big data problem and it's a high performance computing problem.

It's an exciting place – to take those innovations and apply them to real-world problems.

So, the sequencing of the genome is step one. Next is having the HPC environment and the analytics to compare that sequenced genome against clinical databases of other genomes to identify the mutants and then identify what therefore would be the correct treatment. Once that process is simplified and available, it will unleash such better treatment for the general population. That will be a huge breakthrough.

**TER:** *If we try to understand a holistic view of the exascale challenge, we could distill it down to three critical building blocks that need to progress in harmony. First, the technology of course, next the human resources or assets needed to develop, deploy, operate, and create applications, and third is funding. Any of these could be brought to bear*

# the Exascale Report™



## *The voice of the emerging exascale community*

*without the others, but unless the three work together, an exascale milestone does not seem likely.*

*How does Intel approach these three factors as a holistic challenge?*

**BRYANT:** Well, from a human resource perspective we have shifted our investment in our engineering and research teams. We made recent acquisitions of the Qlogic fabric team and the CRAY fabric team, recognizing that historically, as a processor company, we're not going to be able to drive all the innovations required to get to exascale computing without looking at the system at a higher level. It requires a balanced system of compute, storage and I/O. And that means hiring and building talent around systems architectures.

We have great talent that's come in from the IBM Blue Gene team, with Alan Gara \* joining our team along with some other folks, so we can take a systems view of the challenge. We're also making investments on the storage and memory side.

We certainly have moved from being CPU-centric to being systems-centric and building the core competency required to do that.

Even one level up, as we take advantage of Moore's Law and we integrate more functionality on to a single die, the complexity also moves up one layer to the integration between software and hardware.

We are also working internally as well as with the ecosystem partners on how to create a tighter architecture between hardware and software and have that core competency.



*"We're changing the makeup of our talent pool as we look at what it takes to move to exascale computing."*

So we're changing the makeup of our talent pool as we look at what it takes to move to exascale computing. We are absolutely addressing the human resource side.

On the technology side, you know we've got a very clear beat rate of technology investment and that beat rate goes on. The \$10 billion of R&D that Intel spent last year and we will continue to spend is maintaining that beat rate of technology funding, it has been a collaborative funding effort between the investments we make, the investments the ecosystem partners make and the investments the U.S. government makes.

It is that collaborative funding model that has allowed us to go from teraflops to petaflops and it is that same model that will allow us to go from petaflops to exaflops.

# the Exascale Report™



## *The voice of the emerging exascale community*

**TER:** *This is more of a statement for you to comment on than a question - Intel has numerous investments offshore and has invested in exascale research labs in Europe. With the exception of maybe one or two nations, it seems like wherever the first exascale system gets launched, Intel will likely be a part of it.*

**BRYANT:** Certainly we would like to believe that. We believe we are leaders in technology and I think we've demonstrated that over the past. We plan to continue investing so this shouldn't change.

You are absolutely right that we have exascale labs in places around the world in order to leverage talent. In most cases, the exascale labs are a collaborative effort with local universities. That university talent is incredibly valuable to us. Not just what they can directly contribute, but then they become Intel's next generation of employees and help drive our investment cycle internally as well as within the university base. Our investment in these exascale research labs is a collaborative effort within the universities as well as with other governments.

**TER:** *You've seen a lot of changes in technology and process during your career. So what do you see as The Next Big Thing?*

**BRYANT:** So let me start off by saying I am an engineer. I've worked at Intel now for 28 years. As I moved from being the person who drew the transistors on the silicon to a general management role, what excites me the most is the application of technology to real-world problems, and today it's the application of exascale to changing the world.

When we talk about what we can do to the healthcare industry, thanks to the investments in supercomputing, that's what I think will be the big game changer. The application of the technology

that we are all collectively investing in, and applying that technology to real-world change.

That's what I think the big breakthrough of the future will be. And we'll all stand back 10 years from now and say, 'Remember when... Remember when it was no better than a coin toss as to whether your treatment was applicable to your disease or not. Remember when.'

###

**Diane Bryant** is a Senior Vice President and General Manager of the Datacenter and Connected Systems Group at Intel Corporation  
[http://newsroom.intel.com/community/intel\\_newsroom/bios?n=Diane%20M.%20Bryant&f=searchAll-4](http://newsroom.intel.com/community/intel_newsroom/bios?n=Diane%20M.%20Bryant&f=searchAll-4)

**Alan Gara** \* see  
<http://insidehpc.com/2011/04/25/inside-track-blue-gene-architect-alan-gara-leaving-ibm-for-intel/>

**Eric Dishman** \* see  
[http://www.ted.com/speakers/eric\\_dishman.html](http://www.ted.com/speakers/eric_dishman.html)

© 2013, The Exascale Report™