

VISIT NVIDIA AT SC14



Subscribe now and receive weekly updates from your #1 source for HPC news!

SUBSCRIBE



(<http://www.hpcwire.com>)
 Since 1986 - Covering the Fastest Computers in the World and the People Who Run Them

Top News from Leading HPC Solution Providers (<http://www.hpcwire.com/>)

-  (<http://tci.taborcommunications.com/sponsor-altair-2>)
-  (<http://tci.taborcommunications.com/sponsor-asetek>)
-  (<http://tci.taborcommunications.com/sponsor-aspen>)
-  (<http://tci.taborcommunications.com/sponsor-atipa>)
-  (<http://tci.taborcommunications.com/sponsor-boston>)
-  (<http://tci.taborcommunications.com/sponsor-bull>)
-  (<http://tci.taborcommunications.com/sponsor-chelsio>)
-  (<http://tci.taborcommunications.com/sponsor-convey>)
-  (<http://tci.taborcommunications.com/sponsor-cray>)
-  (<http://tci.taborcommunications.com/sponsor-cyclecomputing>)
-  (<http://tci.taborcommunications.com/sponsor-ddn>)
-  (<http://tci.taborcommunications.com/sponsor-eurotech>)
-  (<http://tci.taborcommunications.com/sponsor-extremenetworks>)
-  (<http://tci.taborcommunications.com/sponsor-fujitsu-2>)
-  (<http://tci.taborcommunications.com/sponsor-hp-2>)
-  (<http://tci.taborcommunications.com/sponsor-ibm>)
-  (<http://tci.taborcommunications.com/sponsor-intel>)
-  (<http://tci.taborcommunications.com/sponsor-mellanox>)
-  (<http://tci.taborcommunications.com/sponsor-nec>)
-  (<http://tci.taborcommunications.com/netapp>)
-  (<http://tci.taborcommunications.com/sponsor-numascale>)
-  (<http://tci.taborcommunications.com/sponsor-nvidia>)
-  (<http://tci.taborcommunications.com//21812/2014-04-25/513mh>)
-  (<http://tci.taborcommunications.com/sponsor-pgi>)
-  (<http://tci.taborcommunications.com//21812/2014-11-03/69dc5>)
-  (<http://tci.taborcommunications.com/re-store>)
-  (<http://tci.taborcommunications.com/sponsor-sas>)
-  (<http://tci.taborcommunications.com/sponsor-scalemp>)
-  (<http://tci.taborcommunications.com/sponsor-Seagate-2>)
-  (<http://tci.taborcommunications.com/sponsor-sgi>)
-  (<http://tci.taborcommunications.com/sponsor-supermicro>)
-  (<http://tci.taborcommunications.com/sponsor-univa>)


 Select Language ▼

Subscribe (<http://www.hpcwire.com/subscribe/>) (<http://www.facebook.com/pages/HPCwire/115532028467957?ref=ts>)

(<http://twitter.com/hpcwire>)

Translation disclaimer (<http://hpcwire.wp-engine.com/about-hpcwire/terms-of-use/#translation>)

- [Home](http://www.hpcwire.com/) (<http://www.hpcwire.com/>)
- [News](#)
- [Topics](#)
- [Sectors](#)
- [Exascale](http://www.hpcwire.com/topic/exascale-2/) (<http://www.hpcwire.com/topic/exascale-2/>)
- [Resources](#)
- [Specials](#)
- [Market Watch](http://markets.hpcwire.com/) (<http://markets.hpcwire.com/>)
- [Events](#)
- [Job Bank](http://www.hpcwire.com/job-bank/) (<http://www.hpcwire.com/job-bank/>)
- [About](#)

Visit additional Tabor Communication Publications November 14, 2014



(<http://www.enterprisetech.com>)

CORAL Signals New Dawn for Exascale Ambitions

Nicole Hemsoth

Just when it started to look as though the architectural course had been set for the next wave of large-scale supercomputers, today offered quite a shakeup to the standard.

Off The Wire	Most Read	Podcasts
November 17, 2014		
HPCwire Reveals Winners of the 2014 Readers' and Editors' Choice Awards at SC14 (http://www.hpcwire.com/off-the-wire/hpcwire-		



Top News from Leading HPC Solution Providers
(<http://www.hpcwire.com>)



(<http://tci.taborcommunications.com/sponsor/nvidia>)



(<http://tci.taborcommunications.com/sponsor/hp-2>)



(<http://tci.taborcommunications.com/sponsor/sgi>)



(<http://tci.taborcommunications.com/sponsor/intel>)



(<http://tci.taborcommunications.com/sponsor/fujitsu-2>)



(<http://tci.taborcommunications.com/sponsor/04-25/5/3mh>)



(<http://tci.taborcommunications.com/sponsor/boston>)



(<http://tci.taborcommunications.com/sponsor/eurotech>)



(<http://tci.taborcommunications.com/sponsor/atipa>)



(<http://tci.taborcommunications.com/sponsor/Seagate-2>)



(<http://tci.taborcommunications.com/sponsor/cyclecomputing>)



(<http://tci.taborcommunications.com/sponsor/chelsio>)



And it's not just the amount spent to turn a novel architecture into a pre-exascale reality, although to be fair, it's rare indeed to see a lump \$325 million deal from the Department of Energy to fund new systems with an extra \$100 million added to support extreme scale technologies under the FastForward initiative.

Aside from the sheer investment figures, the fascinating part of what's happening is architectural—and therefore, important in terms of what this means for how centers think about energy consumption, prioritization of extreme scale scientific and security challenges, and of perhaps to some degree, the slightly less dominant position of the U.S. in terms of its national supercomputing capability.

While many expected these first two of the new pre-exascale systems to come out of the CORAL collaboration between Oak Ridge, Lawrence Livermore, and Argonne national laboratories to follow the trends set by Titan and other accelerated Intel-powered x86 machines, those expectations were upended by IBM in today's announcement about a new class of systems sporting GPUs via a close collaboration among OpenPower members IBM, NVIDIA and Mellanox.

Before we delve into an early overview of the systems, it's worth noting that the very status of IBM's role in the future of supercomputing had been called into question over the last year, making this a rather surprising announcement in its own right. From selling off their core HPC-oriented server business to Lenovo to quietly bringing the Blue Gene era to a close, it seemed that their interests were shifting toward a more general Power-based approach for all datacenters—not just HPC with its unique subsets of system choices.

To be fair though, this is still what they're doing. The massive procurement is for systems that are not exactly distinct HPC offerings per se, but rather more advanced and forward-looking variants on the overall OpenPower push to upend Intel's dominance. However, with the addition of key technologies from Mellanox and NVIDIA, specifically the latter's NVLink interconnect, the new generation, which we heard for the first time today is called "Power9" IBM has found a way to maintain an edge at the high end while refining the Power approach to the wider datacenter market as these technologies mature and are put to the test at scale...and massive scale, at that.

The result of all of this are two systems that will be installed in the 2017 time frame. One, called Summit, will be housed at Oak Ridge National Laboratory and will be dedicated to large-scale scientific endeavors ranging from climate modeling to other open science initiatives. The other, called Sierra, is set to be installed at Lawrence Livermore with emphasis on security and weapons stockpile management.

Both are GPU-accelerated systems that have fewer nodes for all the performance they're able to pack in due to the collaboration between NVIDIA and its Volta architecture, which for those who follow these generations, is two away from where we are now with Pascal expected in 2016. The key here is the NVLink interconnect, which is set to push new limits in terms of making these "data centric" supercomputers IBM is espousing as the next step beyond supercomputers which have traditionally been valued according only to their floating point capabilities.

We will be exploring the technology in a companion piece that will immediately follow this one and offer a deeper sense of the projected architecture from chip to interconnect. However, to kick off this series, we wanted to provide a touchstone for these first inkings at what exascale-class systems might look like in the U.S. in the years to come.

One thing is for sure, these are packing a lot of punch in a far lessened amount of space. The Summit system at Oak Ridge is expected to push the 150 to 300 peak petaflop barrier, but according to Jeff Nichols, one of the most remarkable aspects of the system is how they were able to work partners IBM, NVIDIA, and Mellanox to create an architecture that can be boiled down to a much smaller number of nodes for far higher performance and a much larger shared memory footprint.

At this stage, Summit will be 5x or more the performance of Titan at 1/5 the size, weighing in at just around 3400 nodes.

"This shared memory capability and lower node count is important to our developers going forward," he said. "I can say as a computational chemist myself that developers love having fewer nodes to manage and more shared memory per node to work with."

The "data-centric" approach that IBM has been wrapping around for this announcement in particular is another key feature of the Summit system said Nichols. In addition to having the 5x to 10x performance boost using accelerators, which are already in play at Oak Ridge National Lab on the Titan machine, the capabilities for managing vast amounts of complex simulation data is critical. "We can ingest more data, more varieties of data, and explore modeling and simulation data in new ways that we couldn't do even with Titan," he explained. "As we move toward exascale, and this is certainly an early step towards that, we do feel that we have a good path forward in terms of how we'll develop and deploy future systems along this architectural path" with both computational and data centric needs in mind.

As NVIDIA's Sumit Gupta told us today that each of these nodes is so powerful that four of them alone today would make the Top 500. "You probably need a couple of racks of servers to get into the Top 500 but GPU performance will advance so much that we'll get that with just four nodes. The central reason why the largest supercomputers are using accelerators is that CPU alone is too much power. A 150 petaflop system today would be half the power of Vegas—and that isn't going to improve much."

Gupta added that NVLink, which will explore in depth in a follow-up technical piece, is central because the CORAL collaborators wanted a fast processor but required a data movement paradigm that would allow data to be handled quickly without extra hops. The

reveals-winners-2014-readers-editors-choice-
(<http://www.hpcwire.com>)

THE HPC STORAGE LEADER **SC14**
New Orleans | HPC
LA | mclatters

**ACCELERATE RESULTS
SIMPLY AND COST EFFECTIVELY**

Hear first-hand from our User Group why **DDN** is the preferred storage in HPC.

REGISTER NOW ▶

Along These Lines



DOE Funds Exascale Interconnect R&D
(<http://www.hpcwire.com/2013/12/06/doe-funds-exascale-interconnect-rd/>)



Why Iterative Innovation is the Only Path to Exascale
(<http://www.hpcwire.com/2014/04/14/iterative-innovation-path-exascale/>)



This Week in HPC News
(<http://www.hpcwire.com/2014/02/27/week-hpc-news-6/>)



Asia's Place in the Exascale Pack
(<http://www.hpcwire.com/2013/12/06/asia-place-exascale-pack/>)

HPC Tweets

IBM Big Data & HPC @ibmhpc 8m
Attend "Hybrid and Public Cloud Optimized for Technical and High Performance Computing" today at #SC14, 7:15AM-8:15AM CDT#ibmhpc #hpc

Kenneth Hoste @kehoste 14m
Is there any streaming of any of the #SC14 sessions?

Life Sciences Hub @lshubwales 1h
1 week today until the @HPCWales "Making Big Data work for you" event, 09:15 - 15:30, Register here: eventbrite.co.uk/e/making-big-d-...
Retweeted by HPC Wales
Show Summary

HPC Wales @HPCWales 44m
Thought provoking keynote from SPECIFIC @Tatsteeltd starting @LCRIMarine Annual Conference - buildings as power stations #sustainability



(http://tci.taborcommunications.com/sponsor/pgi)



(http://tci.taborcommunications.com/sponsor/ddn)



(http://tci.taborcommunications.com/sponsor/store)



(http://tci.taborcommunications.com/sponsor/mellanox)



(http://tci.taborcommunications.com/sponsor/ibm)



(http://tci.taborcommunications.com/sponsor/altair-2)



(http://tci.taborcommunications.com/sponsor/bull)



(http://tci.taborcommunications.com/sponsor/aspenn)



(http://tci.taborcommunications.com/sponsor/extremenetworks)



(http://tci.taborcommunications.com/sponsor/cray)



(http://tci.taborcommunications.com/sponsor/supermicro)



(http://tci.taborcommunications.com/sponsor/netapp)



(http://tci.taborcommunications.com/sponsor/univa)



(http://tci.taborcommunications.com/sponsor/nec)

traditional CPU and GPU connected traditionally over PCIe has been great for classical high performance computing, he noted, but with high throughput computing users at that scale need the processors to be able to move data efficiently from point to point.

These features are key for the weapons stockpile program that is central to national security where the Sierra system will offer a massive increase in performance and efficiency at Lawrence Livermore. This machine is expected to offer in excess of 100 peak petaflop performance.

As LLNL's Mike McCoy said today, "Simulation is critical to our stockpile program—it's critical for us to make sure we never have to return to nuclear testing. But our 3D weapons codes easily run over a million lines not to mention the databases they employ. At the end of the day, key national security decisions are made based on these calculations but the question is always how do we know these systems are going to do the work we need?"

In answering his own question he explained the way value of the partnership of OpenPower members. "This is not an off the shelf approach—the partnerships are strong and we share the risk in development and deliver platforms that can rapidly come into production and serve our needs. This effort is achieved through a systems integration approach and there will be tight integration between the vendors and code development teams which is called codesign—this has been interestingly enough applied into the past and led to advances like the Blue Gene L that led to advances and performed. This partnership represents a huge opportunity to deliver these and future first gen exascale systems.

We've displaced an Intel-based system at ORNL and we haven't been there for a number of years. It's a nice achievement for us," said IBM's Dave Turek in a conversation today. But the real value in this news is how it could represent the first seismic shift away from the FLOPS-centric approach to large-scale systems to one that takes the problems of data to heart at the core. "We are aided here not because of anything other than what we're seeing in terms of the evolution of the marketplace through direct measurement how necessary it is to simultaneously deal with analytics in concert with modeling and simulation. If you look at an example like seismic processing and you go back ten years, the bulk of the time would have been dedicated to the algorithm and making it faster but what's transformed the conversation is the radical influx of data. Now when you inspect the infrastructure that's being deployed in examples like this, there's a tremendous amount of mundane data sorting and managing that's taking up the compute.

Just as efforts like this have bolstered IBM's supercomputing products overtime, this new collaboration represents a shift for the company. IBM has in fact established an entirely new HPC roadmap—all around the concept of data centric computing. With these systems, the balance of performance, data movement, memory, and overall footprint are balanced with the needs of the new generations of highly scalable codes under development now with assistance from NVIDIA and IBM.

Follow up with us during your SC travels over the weekend and on Monday for more detail about the architectural features we've been able to tease out of a few conversations with IBM, NVIDIA, Mellanox and others.

Share this:

- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=twitter&nb=1) 23
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=facebook&nb=1) 38
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=google-plus-1&nb=1)
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=linkedin&nb=1) 16
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=pocket&nb=1)
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=reddit&nb=1)
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=pinterest&nb=1)
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=tumblr&nb=1)
- (http://www.hpcwire.com/2014/11/14/coral-signals-new-dawn-exascale-ambitions/?share=stumbleupon&nb=1)

Only registered users may comment. Register using the form below.

Check off newsletters you would like to receive *

- HPCwire
- EnterpriseTech
- Datanami
- Technology Conferences & Events
- Advanced Computing Job Bank
- Technology Product Showcase

Feature Articles

NASA Pushes Long-Range Climate Model Limits with SGI
(http://www.hpcwire.com/2014/11/13/nasa-pushes-long-range-climate-model-limits-sgi/)

The Discover system at NASA's Center for Climate Simulation was designed with scalability and flexibility in mind, starting with its original Read more... (http://www.hpcwire.com/2014/11/13/nasa-pushes-long-range-climate-model-limits-sgi/)

Monday Twitter Roundup
(http://www.hpcwire.com/2014/11/13/monday-twitter-roundup/)

In celebration of SC14, we've decided to put together a daily list of some the top tweets from the event. For those unable to attend, we hope Read more... (http://www.hpcwire.com/2014/11/13/monday-twitter-roundup/)

Why the Top 500 Standstill Won't Last Forever
(http://www.hpcwire.com/2014/11/13/why-the-top-500-standstill-wont-last-forever/)

Traditionally, one of the most exciting opening elements of the annual SC event is the announcement of the list of the Top 500 supercomputers on Read more... (http://www.hpcwire.com/2014/11/13/why-the-top-500-standstill-wont-last-forever/)

Read more HPCwire Features... (http://www.hpcwire.com/ca)

Short Takes

Weekly Twitter Roundup
(http://www.hpcwire.com/2014/11/13/weekly-twitter-roundup-21/)

Here at HPCwire, we want to help keep the HPC community as up-to-date as possible on some of the most captivating news items that were Read more... (http://www.hpcwire.com/2014/11/13/weekly-twitter-roundup-21/)

Big Data Compels HPC Adoption in Life Sciences
(http://www.hpcwire.com/2014/11/13/big-data-compels-hpc-adoption-life-sciences/)

Expect a lot of talk at SC14 this year to revolve around big data. Ari E. Berman, Ph.D., Director of Government Services and Principal Read more... (http://www.hpcwire.com/2014/11/13/big-data-compels-hpc-adoption-life-sciences/)

UK Project Tackles Bike Helmet Safety
(http://www.hpcwire.com/2014/11/13/uk-project-tackles-bike-helmet-safety/)

There are certain HPC projects that stand out for their ability to help humankind in practical ways. One recent example of such a project Read more... (http://www.hpcwire.com/2014/11/13/uk-project-tackles-bike-helmet-safety/)



BEST PRACTICES

USER PRESENTATIONS

PEER NETWORKING

TECHNOLOGY BRIEFINGS

REGISTER TODAY ▶



HPC Job Bank

- System Support Manager- Applications and Infrastructure - BHP Billiton Petroleum (http://jobs.hpcwire.com/jobdetails.c?id=1970)
- GUI Developer - Texas Advanced Computing Center (http://jobs.hpcwire.com/jobdetails.c?id=1960)
- Visit the HPCwire Job Bank (http://www.hpcwire.com/job-bank/)



Featured Events

- SC14 (http://www.hpcwire.com/event/s)
- SC14 hpc matters. sct4.supercomputing.org
- November 16 -
- November 21