Intel Collaborates with CERN to Support Upgraded LHC Experiments

By Sean Thielen

November 4, 2016

Editor's note: In this contributed feature, Sean Thielen details the technical collaboration between CERN and Intel in preparation for the center's next major upgrade, scheduled to commence in late 2018. The joint team has deployed a two-socket Xeon-FPGA (Stratix V) proof-of-concept machine that makes use of a hybrid package design and will soon be testing a newer implementation that will leverage the Arria 10 FPGA and a faster interconnect.

Much of the media attention given to the particle accelerator experiments that happen at the European Organization for Nuclear Research, known as CERN, is focused on the Large Hadron Collider (LHC). It's no surprise, given the LHC is the world's largest, most complex machine, unravelling some of the toughest scientific problems by accelerating particles (protons or heavy ions) and making them collide in a gigantic 27-kilometer ring. But the work that happens immediately after particles collide in the LHC is not only critical to science, it's also quite interesting and important from a computing and data processing perspective. After all, the creation of particles or results in the LHC is only significant if scientists can quickly isolate them from millions of inconsequential signals for further study. That means ongoing advancements in trigger and data acquisition systems are essential to fully reaping the rich potential of the LHC. And, as you can imagine, the networking and computing challenges are extreme in nearly every dimension.

Historically, CERN's trigger and data acquisition systems have relied heavily on custom technology. For the next upgrade cycle scheduled to start around the end of 2018, however, CERN engineers and scientists were aware that the scalability limitations and costs of their custom solutions could start to limit progress. This led to a new collaborative project with Intel, through CERN openlab (http://openlab.cern/), focused on exploring the feasibility of complementing one of the LHC's detectors with off-the-shelf data acquisition, data movement, and data filtering technology from Intel, including an FPGA platform. If the proof of concept is successful, it could have a significant impact on the future design and efficiency of trigger systems, including the remaining LHC detectors and other scientific instruments.

The Large Hadron Collider tunnel is located 100 meters underground on the Franco-Swiss border, near Geneva. Source: CERN

Preparing for enormous data growth

Simply phrased, the LHC fires high-energy particle beams at each other in a 27-kilometer ring. The detectors on the LHC include tracking devices that plot the trajectory of particles following collisions, as well as calorimeters that measure their energy, which helps to narrow down their identity. Niko Neufeld, a deputy project leader at CERN who works on the LHCb experiment, likens the computing challenges (http://openlab.cern/news/high-throughput-computing-lhcb-experiement) for the LHC experiments to solving millions of small puzzles involving up to a billion proton-proton collisions every second to retain the most interesting ones for deeper analysis.

Given the task at hand, the near-detector "online" computing challenges at CERN have always been extreme. And when CERN upgrades the LHC and detectors from late 2018 to early 2021 as part of its regular upgrade cycle, the data rates running through the various systems will jump significantly. Neufeld provides perspective on the jump: "Network scaling needs have really grown. For example, the largest networks we currently run at CERN have total bandwidth of around 800 gigabits per second. Following the upgrade work, our networks will need to carry between 40 and 50 terabits of data per second. If you compare that to a Google data center, it may not sound impressive, but for a scientific instrument it's a huge step in terms of bandwidth." Neufeld said that the computing challenges have also grown quite complex. "We cannot simply scale up computing by a factor of 100 because we have whatever Moore's law gives us… We may be able to grow our computing farms by a factor of 1 or 2, but not much more. The rest has to come from more clever processing models," he added.

According to Neufeld, the detector teams face a host of challenges in preparing for the data-rate jump. Neufeld said that the customized trigger (hardware and software) on the front-end of the existing detector system had a long and expensive refresh cycle. "The engineering resources for ASICs [application-specific integrated circuits] and the FPGAs [Field Programmable Gate Arrays] in high energy physics are limited compared to industry, and the tight integration with the detectors makes upgrades outside of our major maintenance periods impractical," Neufeld explained. "We thought that moving to a more software-centered approach using off-the-shelf technology could greatly reduce these limitations and expand the developer base that is available in our community. Physicists are usually knowledgeable in some programming language, however, HDL [hardware description language] is a different challenge with a very long and steep learning curve."

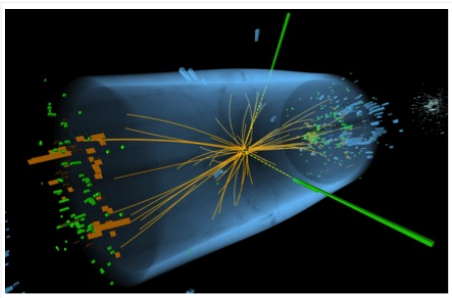

The Compact Muon Solenoid (CMS) is one of two large general-purpose detectors on the

Olof Bärring, a deputy group leader of computing facilities at CERN, added that cost and energy considerations were also an important part of the equation. The lab needs to continue addressing greater computing, data moving and storage challenges with a more or less flat budget and within the datacenter's existing energy envelope.

LHC. The image above captures a candidate proton-proton event as a part of the CMS search for the Higgs boson. Source: CERN

Exploring options for three critical challenges

In 2014, the former CTO of CERN openlab, Sverre Jarp, invited Intel's Karl Solchenbach, director Exascale Labs Europe, and Steve Pawlowski, former vice president of advanced computing solutions at Intel, to discuss CERN's technical challenges with near-detector online computing to see if any Intel technology might be useful in addressing them. During the meeting, Neufeld presented three main challenges, and the participants worked together to map technologies to the challenges.

Challenge one: real-time or near-real-time data-processing with very short (order of 10 microseconds) latencies

The Intel team considered an Intel Xeon processor and Altera Stratix V FPGA integration along with another spatial architecture, and ended up recommending the Intel Xeon processor/FPGA configuration. Neufeld says that the reasoning was that the Intel Xeon processor/FPGA configuration should allow LHC experiments to potentially replace part, and in some cases all, of the custom electronics used in the first step of online data filtering. "That meant we would be able to use off-the-shelf hardware programmed using 'high-level' general purpose languages instead of HDLs, which was an important step for us," he explained.

Challenge two: very high throughput local area networks

For data transport, Intel recommended investigating the potential of Intel Omni-Path Architecture (Intel OPA) as an alternative to deep-buffer Ethernet switching. Neufeld said the main driver for considering Intel OPA over traditional approaches was cost. "It is certainly technically possible to build the network in a more traditional way, but it has become prohibitively expensive, given our budget," he explained.

Challenge three: the need of massive data-processing for data-reduction

For the filtering of detector data in software, the new Intel Xeon Phi processor was an attractive potential solution, given that the filtering process itself is quite parallel and individual collisions in the LHC are statistically independent.

Not your average POC

Once the CERN and Intel teams agreed that the identified off-the-shelf solutions for each challenge had potential, the hard work of proving the viability of each solution in the next-generation data acquisition environment needed to begin in earnest. In 2015, CERN and Intel decided to expand their existing collaboration through CERN openlab and signed an agreement for a joint three-year project called the High Throughput Computing Collaboration (HTCC). As of writing, the HTCC project has reached the halfway point and the core team, which includes seven CERN scientists and one Intel engineer on site, has made progress in each key challenge area.

The main CERN datacenter. The 110,000 processor cores and 10,000 servers hosted in its three rooms run 24/7. Source: CERN

Networking milestones

In the networking area, Neufeld said a big part of the challenge is simply getting access to a system big enough to test software. "To properly prepare our software, we need access to complete supercomputers. Fortunately, Intel can provide access to clusters that are up to the job," he explained. Neufeld said that the team recently had an important success while running its software on an Intel OPA cluster with more than 500 nodes. "We have already been able to achieve a full-duplex, high-throughput transit of 70 terabits per second of data flying through the cluster, so that's already half of what we need by the time of the upgrades."

FPGA-related milestones

When CERN updates the LHC and the experiments from late 2018 to early 2021, its detectors will support trigger-free readouts. The LHC generates up to around 1 billion collisions per second in the experiments, and the goal is to read them all. Moving forward, a flexible software-based trigger system running in a large (up to 4000 nodes) computing farm will select the interesting collisions. In the meantime, CERN is investigating which technology options are the best fit for accelerating its algorithms. For its initial FPGA proof of concept, the HTCC team deployed a two-socket Intel Xeon processor/FPGA machine, which included the following hardware connected by the Intel QuickPath Interconnect:

- Intel Xeon CPU E5-2680 v2
- Altera Stratix V GX A7 FPGA with 234,720 adaptive logic modules (ALMs)

The team is particularly interested in the potential compute and power efficiency gains that are possible with using OpenCL in a combined CPU and FPGA system. With respect to the FPGA platform testing, Neufeld said that the team dealt with host of technical challenges because it wanted to do a meaningful comparison of OpenCL and to get a sense of the costs using a high-level framework. "Just for illustrative purposes, it took us two weeks to set up a new kernel using OpenCL compared to three months to complete an equivalent Verilog implementation, and we had a very skilled engineer on that job," explained Neufeld. "In the end, I think it was a good investment because we needed to prove to the electronic engineers that the new technology actually provided a less painful way to get the results we need."

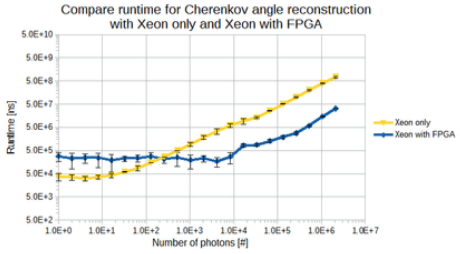

(https://6lli539m39y3hpkelqsm3c2fg-

Following a series of test cases, such as sorting and calculating the Mandlebrot fractal, to understand the potential of the Xeon/FPGA system, the HTCC team developed an FPGA fine-tuned for the rigors of RICH (ring-imaging Cherenkov) reconstruction. Only then did it begin doing LHC-specific workload analysis.

In coming months, the HTCC team will also be testing a newer system that is built using a combined Intel Xeon CPU and FPGA in a single package. It will include the new high-performance Arria 10 FPGA from Altera as well as a faster interconnect of the CPU and FPGA.

Intel Xeon Phi processor milestones

wpengine.netdna-ssl.com/wp-content/uploads/2016/11/Cherenkov-angle-reconstruction-Intel-CERN-429x.png)
Cherenkov angle reconstruction is used for particle identification in the detectors. Based on initial tests the Xeon/FPGA machine shows promise for processing greater numbers of photons after the LHC upgrade. Source: CERN

Neufeld said that testing of the Intel Xeon Phi processor platform will be a major focus of the HTCC team for the next year or so. He noted that like everyone else, the team needs to figure out how to adapt well to the new architecture and different level of parallelism. To achieve this the HTCC team has been working with Intel engineers on benchmarking and understanding the different algorithms' implementations using Intel analysis tools, such as Intel VTune Amplifier XE and Intel Advisor XE, as well as different performance models, such as the roofline model.

"In addition to the inclusion of bootable sockets, in-package memory and high main-memory bandwidth, what is particularly interesting with Intel Xeon Phi processors is the integrated fabric and its potential to quickly distribute workloads to where they fit best. We will test both data movement aspects on the Intel Xeon Phi processor as well as the distribution of the algorithms between the Intel Xeon and Intel Xeon Phi processor using the fabric as an interconnect," Neufeld added.

Pushing the boundaries of precision

When asked what the progress made on the data acquisition systems for the LHC might mean for wider applications, Neufeld said the value is all about greater precision for complex experiments. "To some extent, it's just statistics. Either you really increase the amount of data and the precision by a significant factor, or you stop doing it," Neufeld explained. "This work should lead to an important jump in precision. For example, the LHCb experiment's online collection and analysis system currently selects just 1 million of the 40 million bunches of protons that cross in the accelerator every second, with the others being discarded based on less-precise hardware-calculated signatures. After the LHCb upgrade, the number of collisions is set to grow yet further, and we will look at all of them in the software, in order to take the best physics out of there. And since there are actually only a couple of milliseconds to do that for each collision, it's really quite a leap forward."t direct discussions with the Intel development team will continue to be invaluable to getting things right on projects as the team races to meet its deadline for the start of the upgrade work.

Off The Wire Industry Headlines

November 9, 2016

- Bright Computing to Highlight Cluster Management Solutions SC16 (https://www.hpcwire.com/off-the-wire/bright-computing-highlight-cluster-management-solutions-sc16/)
- Cray Joins iEnergy, the Oil & Gas Industry's Community for Exploration and Production (https://www.hpcwire.com/off-the-wire/cray-joins-ienergy-oil-gas-industrys-community-exploration-production/)
- DDN Storage Solutions Deployed by Aalto University (https://www.hpcwire.com/off-the-wire/ddn-storage-solutions-deployed-aalto-university/)
- Engility to Showcase HPC Innovation at SC16 (https://www.hpcwire.com/off-the-wire/engility-showcase-hpc-innovation-sc16/)
- Atos Tackles Storage Bottlenecks for HPC With Bull Director for HPSS (https://www.hpcwire.com/off-the-wire/atos-tackles-storage-bottlenecks-hpc-bull-director-hpss/)
- Katharine Frase, Cognitive Computing Expert, to Address the Future of AI at SC16 Keynote (https://www.hpcwire.com/off-the-wire/katharine-frase-cognitive-computing-expert-address-future-ai-sc16-keynote/)
- ExaFLOW Project Secures Funding (https://www.hpcwire.com/off-the-wire/exaflow-project-secures-funding/)
- Texas Multicore Technologies Joins the OpenPOWER Foundation (https://www.hpcwire.com/off-the-wire/texas-multicore-technologies-joins-openpower-foundation/)
- Nor-Tech to Roll Out Clusters Integrated With Intel's Orchestrator Platform (https://www.hpcwire.com/off-the-wire/tech-roll-clusters-integrated-intels-orchestrator-platform/)
- SC16 Sessions to Highlight HSA Momentum (https://www.hpcwire.com/off-the-wire/sc16-sessions-highlight-hsa-momentum/)
- Blue Waters Research Integrates Data to Predict E. Coli Behavior (https://www.hpcwire.com/off-the-wire/blue-waters-research-integrates-data-predict-e-coli-behavior/)

November 8, 2016

- Cray Reports Third Quarter 2016 Financial Results (https://www.hpcwire.com/off-the-wire/cray-reports-third-quarter-2016-financial-results/)
- BIOS IT Partners With vScaler to Deliver Optimized HPC Offerings for the Cloud (https://www.hpcwire.com/off-the-wire/bios-partners-vscaler-deliver-optimized-hpc-offerings-cloud/)
- San Diego State to Make First Appearance at Student Cluster Competition (https://www.hpcwire.com/off-the-wire/san-diego-state-make-first-appearance-student-cluster-competition/)
- DDN Unveils New Security Solution for Lustre (https://www.hpcwire.com/off-the-wire/ddn-unveils-new-security-solution-lustre/)
- Cycle Computing's CycleCloud V6 Now Available (https://www.hpcwire.com/off-the-wire/cycle-computings-cyclecloud-v6-now-available/)
- Caringo to Exhibit Cloud & Object Storage Platform at SC16 (https://www.hpcwire.com/off-the-wire/caring-exhibit-cloud-object-storage-platform-sc16/)
- SDSC to Host High-Speed, Large Data Transfer Experiment at SC16 (https://www.hpcwire.com/off-the-wire/sdsc-host-high-speed-large-data-transfer-experiment-sc16/)
- DDN Announces New Products and Updates (https://www.hpcwire.com/off-the-wire/ddn-announces-new-products-updates/)
- Blue Waters Simulates Largest Membrane Channel Made of DNA Origami (https://www.hpcwire.com/off-the-wire/blue-waters-simulates-largest-membrane-channel-made-dna

(http://text=cern-suppor lhc-experir

(http://u=http cern-suppor lhc-experir

(http://mini=tr cern-suppor lhc-experir

(http://url=htt cern-suppor lhc-experir

(http://url=htt cern-suppor lhc-experir

Subscribe to HPCwire's Weekly Update!

Be the most informed person in the room! Stay ahead of the tech trends with industy updates delivered to you every week!

(https://www.hpcwire.com/subscribe/)

THE LATEST EDITOR'S PICKS MOST POPULAR

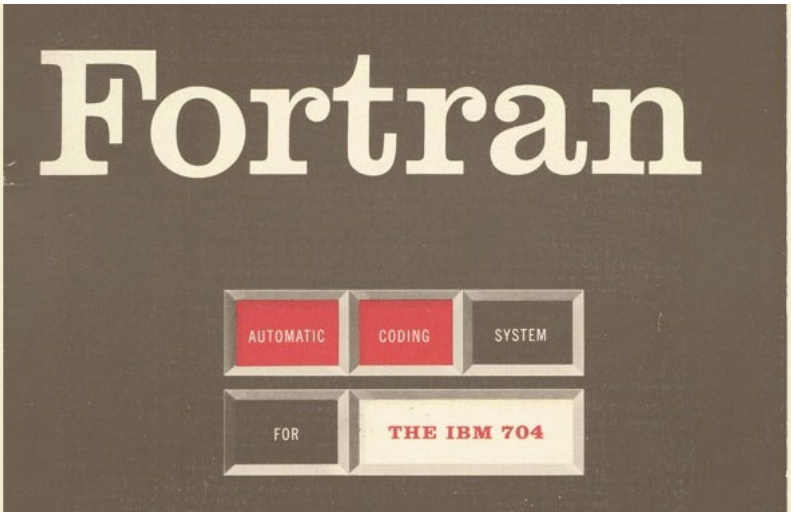

SC16 Preview: Modernizing, Modularizi Fortran Codes

(https://www.hpcwire.com/2016/11/09/sc16-preview-modernizing-fortran-codes/)

Nover

The programming language Fortran has been in existence and a large percentage of software that runs on HPC syste worldwide is written, at least in part, in Fortran. Since it has existence for so long, and so many scientists and engineer early in their careers and have utilized it extensively, much is not up to modern standards of software development. Re (https://www.hpcwire.com/2016/11/09/sc16-preview-moder fortran-codes/)

By Mark Potts

(http://twitter.com/intent/tweet?status=SC16+Preview%3A+Modernizing%2C+Modularizing+Fortran+Codes+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F09%2Fsc16 modernizing-fortran-codes%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F09%2Fsc16-preview-modernizing-fortran-codes%2F&title=SC16+Preview%3A+Modernizing%2C+Modularizing+Fortran+Codes&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F09%2Fsc16-preview-modernizing-fortran-codes%2F&title=SC16+Preview%3A+Modernizing%2C+Modularizing+Fortran+Codes/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F09%2Fsc16-preview-modernizing-fortran-codes%2F)

A BoF to Boost Your Supercomputing Outreach Skills

(https://www.hpcwire.com/2016/11/09/bof-boost-supercomputing-outreach-skills/)

In this Q&A, University of Edinburgh researcher Nick Brown discusses an upcoming SC BoF, titled "HPC Outreach: Promoting Supercomputing to the Next Generation organized by Brown and two of his colleagues, focuses not only on outreach to different groups of people, but also some of the questions around diversity and ways to activities that are inclusive to all. Read more… (https://www.hpcwire.com/2016/11/09/bof-boost-supercomputing-outreach-skills/)

By Tiffany Trader

Tackling the Co-design 3.0 Puzzle – Nev Thinking Needed

(https://www.hpcwire.com/2016/11/09/tackling-co-des
puzzle/)

Co-design has long been a vibrant discussion point in the H
community. The need to coordinate development across ha
software, and system architecture in the face of constraints
declining Moore's Law is a given. Read more…
(https://www.hpcwire.com/2016/11/09/tackling-co-design-pu

By John Russell

Heading into SC16 CENATE Flexes its Growing Muscle

(https://www.hpcwire.com/2016/11/08/heading-sc16-cenate-flexes-growing-muscle/)

Nove

In September, the Center for Advanced Technology Evaluation (CENATE) at Pacific Northwest National Laboratory (PNNL) took possession of NVIDIA's DGX-1 GPU-b
(Pascal 100) supercomputer. Read more… (https://www.hpcwire.com/2016/11/08/heading-sc16-cenate-flexes-growing-muscle/)

By John Russell

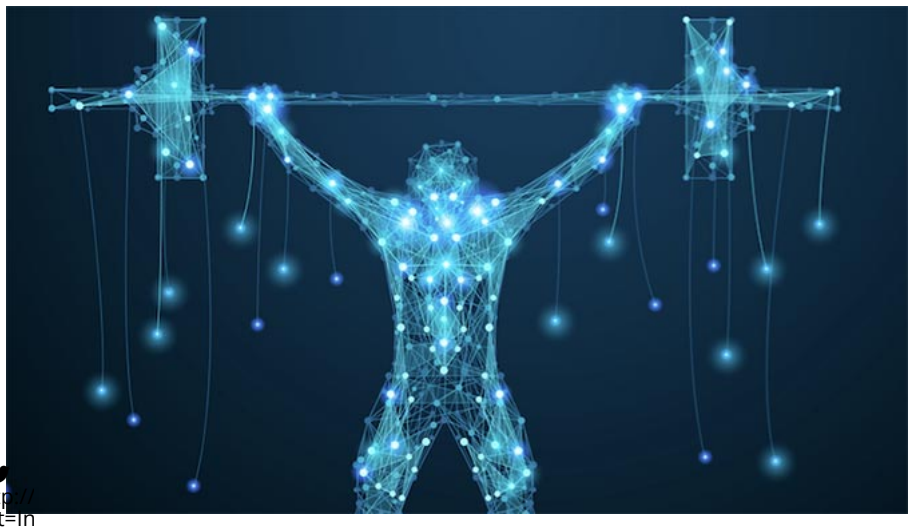

(http://text=In cern-suppor lhc-experir (http://twitter.com/intent/tweet?status=Heading+into+SC16+CENATE+Flexes+its+Growing+Muscle+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fheading-sc16-flexes-growing-muscle%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fheading-sc16-cenate-flexes-growing-muscle%2F&title=Heading+into+SC16+CENATE+Flexes+its+Growing+Muscle&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fheading-sc16-cenate-flexes-growi (http://u=http cern-suppor lhc-experir muscle%2F&title=Heading+into+SC16+CENATE+Flexes+its+Growing+Muscle/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fheading-sc16-cenate-flexes-growing-muscle%2F)

SC16: Taking Diversity and Inclusivity Seriously

(https://www.hpcwire.com/2016/11/08/sc16-taking-diversity-inclusivity-seriously/)

Diversity has always been important to the SC Conference and this year they have taken their actions the creation of the Diverse HPC Workforce Committee. In this contributed Q&A, Kim McMahon asks S John West and Trish Damkroger, Chair of the Diverse HPC Workforce committee about their respectiv what motivates them to volunteer for this effort. Read more… (https://www.hpcwire.com/2016/11/08/s diversity-inclusivity-seriously/)

By Kim McMahon

(http:// mini=tr cern-suppor lhc-experir (http:// url=htt cern-suppor lhc-experir (http://twitter.com/intent/tweet?status=SC16%3A+Taking+Diversity+and+Inclusivity+Seriously+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fsc16-taking-diversinclusivity-seriously%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2 (http:// url=htt cern-suppor lhc-experir taking-diversity-inclusivity-seriously%2F&title=SC16%3A+Taking+Diversity+and+Inclusivity+Seriously&source=https%3A%2F%2Fwww.hpcwire.com (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fsc16-taking-diversity-inclusivity-seriously%2F&title=SC16%3A+Taking+Diversity+and+Inclusivity+Seriously/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fsc16-taking-diversity-inclusivity-seriously%2F)

DDN in Product Line Refresh, Multi-leve Security for Lustre

(https://www.hpcwire.com/2016/11/08/ddn-product-lir

Ahead of next week's SC16 conference, high end storage DataDirect Networks today announced a raft of product line including a new burst buffer appliance and updates to its bl storage appliances, along with what DDN says is the fastes security solution for Lustre aimed at HPC-level performanc with increased security requirements, such as financial ser genomics and government organizations. Read more… (https://www.hpcwire.com/2016/11/08/ddn-product-line-refr

By Doug Black

(http://twitter.com/intent/tweet?status=DDN+in+Product+Line+Refresh%2C+Multi-level+Security+for+Lustre+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fddn-product-line-refresh%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fddn-product-line-refresh%2F&title=DDN+in+Product+Line+Refresh%2C+Multi-level+Security+for+Lustre&source=https%3A%2F%2Fwww.hpcwire.com/)

(http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fddn-product-line-refresh%2F&title=DDN+in+Product+Line+Refresh%2C+Multi-level+Security+for+Lustre/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F08%2Fddn-product-line-refresh%2F)

STREAM Benchmark Author McCalpin Traces System Balance Tren

(https://www.hpcwire.com/2016/11/07/mccalpin-traces-hpc-system-balance-trends/)

Nover

When Dr. John D. McCalpin introduced the STREAM benchmark in 1991, it had already become beco that peak arithmetic rate was not an adequate measure of HPC system performance for many applica then, CPU performance has continued to outpace memory performance measures, leading to the pro memory speed gap, known as the memory wall. In an invited talk at SC16, McCalpin will trace the hist changing “balances” between computation, memory latency, and memory bandwidth and will explore the next-generation of HPC systems. Read more… (https://www.hpcwire.com/2016/11/07/mccalpin-tra system-balance-trends/)

By Tiffany Trader

(http://twitter.com/intent/tweet?status=STREAM+Benchmark+Author+McCalpin+Traces+System+Balance+Trends+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fmccalpin-traces-hpc-system-balance-trends%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fmccalpin-traces-hpc-system-balance-trends%2F&title=STREAM+Benchmark+Author+McCalpin+Traces+System+Balance+Trends&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fmccalpin-traces-hpc-system-balance-trends%2F&title=STREAM+Benchmark+Author+McCalpin+Traces+System+Balance+Trends/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fmccalpin-traces-hpc-system-balance-trends%2F)

OpenACC Expands Community, Reveals Roadmap Details

(https://www.hpcwire.com/2016/11/07/openacc-reveals-roadmap-details/)

In advance of the SC16 expo in Salt Lake City next week, the OpenACC standards group today welco member NSSC-Wuxi and highlighted a number of important developments for the directives-based pr standard. Ahead of the announcement, HPCwire spoke with Michael Wolfe, technical director of Open Duncan Poole, OpenACC president and director of platform alliances for accelerated computing at Nv more… (https://www.hpcwire.com/2016/11/07/openacc-reveals-roadmap-details/)

By Tiffany Trader

(http://twitter.com/intent/tweet?status=OpenACC+Expands+Community%2C+Reveals+Roadmap+Details+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fopenacc-reveals-roadmap-details%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fopenacc-reveals-roadmap-details%2F&title=OpenACC+Expands+Community%2C+Reveals+Roadmap+Details&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fopenacc-reveals-roadmap-details%2F&title=OpenACC+Expands+Community%2C+Reveals+Roadmap+Details/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F07%2Fopenacc-reveals-roadmap-details%2F)

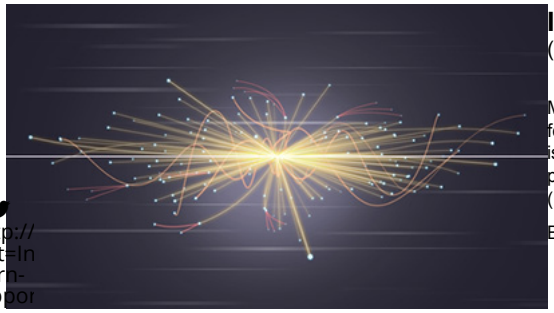

Intel Collaborates with CERN to Support Upgraded LHC Experiment

(https://www.hpcwire.com/2016/11/04/intel-cern-support-lhc-experiments/)

N

Much of the media attention given to the particle accelerator experiments that happen at the Europea for Nuclear Research, known as CERN, is focused on the Large Hadron Collider (LHC). It's no surpris is the world's largest, most complex machine, unravelling some of the toughest scientific problems by particles (protons or heavy ions) and making them collide in a gigantic 27-kilometer ring. Read more.. (https://www.hpcwire.com/2016/11/04/intel-cern-support-lhc-experiments/)

By Sean Thielen

(http://text=In cern- suppor lhc- experir (http://twitter.com/intent/tweet?status=Intel+Collaborates+with+CERN+to+Support+Upgraded+LHC+Experiments+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2support-lhc-experiments%2F) (http://u=http cern- suppor lhc- experir (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F0cern-support-lhc-experiments%2F&title=Intel+Collaborates+with+CERN+to+Support+Upgraded+LHC+Experiments&source=https%3A%2F%2Fwwv (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2Fintel-cern-support-lhc-experiments%2F&title=Intel+Collaborates+with+CERN+to+Support+Upgraded+LHC+Experiments/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2Fintel-cern-support-lhc-experiments%2F)

(http:// mini=tı cern- suppor lhc- experir

IDC's Conway Sets Stage for SC16 Precision Medicine Panel

(https://www.hpcwire.com/2016/11/04/idc-conway-sc16-precision-medicine-panel/)

Kicking off SC this year is what promises to be a fascinating panel – HPC Impacts on Precision Medici Future–The Next Frontier in Healthcare. In this pre-SC16 Q&A, Steve Conway, research VP in IDC's H Computing group and moderator of the panel, sets the stage. HPC, of course, has been transforming medicine for nearly two decades. Read more… (https://www.hpcwire.com/2016/11/04/idc-conway-sc1 medicine-panel/)

By John Russell

(http:// url=htt cern- suppor lhc- experir (http://twitter.com/intent/tweet?status=IDC%26%238217%3Bs+Conway+Sets+Stage+for+SC16+Precision+Medicine+Panel+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11conway-sc16-precision-medicine-panel%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2Fidc-conway-sc16-precision-medicine-panel%2F&title=IDC%26%238217%3Bs+Conway+Sets+Stage+for+SC16+Precision+Medicine+Panel&source=https%3A%2F%2Fwww.hpcwire.con (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2Fidc-conway-sc16-precision-medici panel%2F&title=IDC%26%238217%3Bs+Conway+Sets+Stage+for+SC16+Precision+Medicine+Panel/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F04%2Fidc-conway-sc16-precision-medicine-panel%2F)

Is Entrepreneurship Flagging in America? Not in HPC

(https://www.hpcwire.com/2016/11/03/entrepreneurship-flagging-america-not-hpc/)

N

In this guest commentary, veteran HPC professional Stephen Perrenod makes the case that entrepre heart of HPC's success and critical to HPC's future as a driving force of the digital economy. Read mo (https://www.hpcwire.com/2016/11/03/entrepreneurship-flagging-america-not-hpc/)

By Stephen Perrenod, OrionX

(http://twitter.com/intent/tweet?status=Is+Entrepreneurship+Flagging+in+America%3F+Not+in+HPC+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fentrepreneu america-not-hpc%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fentrepreneurship-flagging-america-not-hpc%2F&title=Is+Entrepreneurship+Flagging+in+America%3F+Not+in+HPC&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fentrepreneurship-flagging-america hpc%2F&title=Is+Entrepreneurship+Flagging+in+America%3F+Not+in+HPC/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fentrepreneurship-flagging-america-not-hpc%2F)

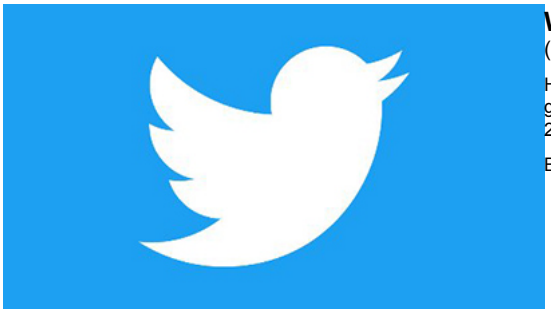

Weekly Twitter Roundup (Nov. 3, 2016)

(https://www.hpcwire.com/2016/11/03/weekly-twitter-roundup-nov-3-2016/)

Here at *HPCwire*, we aim to keep the HPC community apprised of the most relevant and interesting n get tweeted throughout the week. Read more… (https://www.hpcwire.com/2016/11/03/weekly-twitter-r 2016/)

By Thomas Ayres

(http://twitter.com/intent/tweet?status=Weekly+Twitter+Roundup+%28Nov.+3%2C+2016%29+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fweekly-twitter-roun 2016%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fweekly-twitte 3-2016%2F&title=Weekly+Twitter+Roundup+%28Nov.+3%2C+2016%29&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fweekly-twitter-roundup-nov-3- 2016%2F&title=Weekly+Twitter+Roundup+%28Nov.+3%2C+2016%29/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fweekly-twitter-roundup-nov-3-2016%2F)

Sandia to Referee IARPA Visual Pattern Recognition Project

(https://www.hpcwire.com/2016/11/03/sandia-referee-iarpa-project/)

Much has been rightly made of the human brain's ability to discern patterns from the seemingly visual the world around us. Read more… (https://www.hpcwire.com/2016/11/03/sandia-referee-iarpa-project

By John Russell

(http://twitter.com/intent/tweet?status=Sandia+to+Referee+IARPA+Visual+Pattern+Recognition+Project+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fsandia-r project%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fsandia-refe project%2F&title=Sandia+to+Referee+IARPA+Visual+Pattern+Recognition+Project&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fsandia-referee-iarpa-project%2F&title=Sandia+to+Referee+IARPA+Visual+Pattern+Recognition+Project/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F03%2Fsandia-referee-iarpa-project%2F)

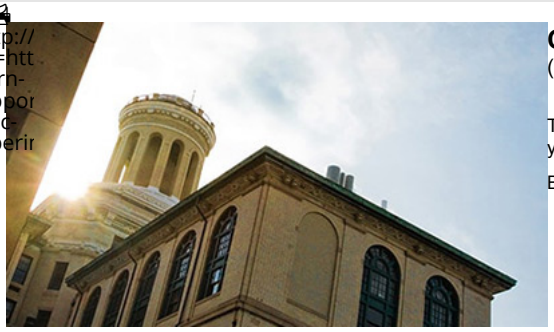

Carnegie Mellon Receives $10M for Center to Explore AI Ethics

(https://www.hpcwire.com/2016/11/02/carnegie-mellon-center-explore-ai-ethics/)

N

The prospect of creating artificial intelligence has created a stir inside and outside science communitie years. Read more… (https://www.hpcwire.com/2016/11/02/carnegie-mellon-center-explore-ai-ethics/)

By John Russell

(http://twitter.com/intent/tweet?status=Carnegie+Mellon+Receives+%2410M+for+Center+to+Explore+AI+Ethics+https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F02%2F mellon-center-explore-ai-ethics%2F) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F02%2Fcarnegie-mellon-center-explore-ai-ethics%2F&title=Carnegie+Mellon+Receives+%2410M+for+Center+to+Explore+AI+Ethics&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F02%2Fcarnegie-mellon-center-explore-ai-ethics%2F&title=Carnegie+Mellon+Receives+%2410M+for+Center+to+Explore+AI+Ethics/) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2016%2F11%2F02%2Fcarnegie-mellon-center-explore-ai-ethics%2F)

Dell Knights Landing Machine Sets New STAC Records

(https://www.hpcwire.com/2016/11/02/dell-knights-landing-machine-sets-new-stac-records/)

The Securities Technology Analysis Center, commonly known as STAC, has released a new report characterizing the performance of the Knight Landing-based Dell P C6320p server on the STAC-A2 benchmarking suite, widely used by the financial services industry to test and evaluate computing platforms. The Dell machine has se both the baseline Greeks benchmark and the large Greeks benchmark. Read more… (https://www.hpcwire.com/2016/11/02/dell-knights-landing-machine-sets-new-st

By Tiffany Trader

Obama Announces New Working Group on US Semiconductor Indu

(https://www.hpcwire.com/2016/11/02/obama-announces-new-working-group-us-semicond

On Monday President Obama announced the formation of a working group to examine methods for g strengthening the U.S. semiconductor industry. It's not clear the U.S. chip industry is in trouble but cer abroad are gaining steam. Notably China is focusing on developing its own chips – not least in respon restrictions on select U.S. chips to China. Also, Fujitsu has announced its post K machine will be ARM more… (https://www.hpcwire.com/2016/11/02/obama-announces-new-working-group-us-semiconduct

By John Russell

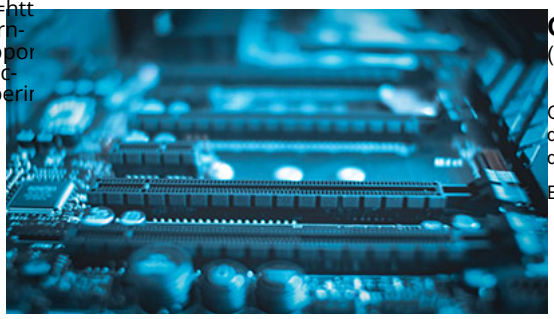

Graphcore Promises AI-Speedup with 'Intelligent Processing Unit'

(https://www.hpcwire.com/2016/11/01/graphcore-100x-speedup-pcie-plugin/)

N

Graphcore emerged from stealth mode today with news of a $30 million Series A round to help financ development of its machine learning (ML) and deep learning acceleration solutions, including a PCIe directly into a server's bus. Read more… (https://www.hpcwire.com/2016/11/01/graphcore-100x-spee

By Alex Woodie

HPC Career Notes (Nov. 2016)

(https://www.hpcwire.com/2016/11/01/hpc-career-notes-nov-2016/)

In this monthly feature, we'll keep you up-to-date on the latest career developments for individuals in the high performance computing community. Read more… (https://www.hpcwire.com/2016/11/01/hpc-career-notes-nov-2016/)

By Thomas Ayres