



CERN openlab Explores New CPU/FPGA Processing Solutions

(http://
text=C
By Linda B
fpga-
proces
tested-
at-
cern%
April 14, 2017

Editor's note: In this contributed feature, Linda Barney describes the ongoing technical collaboration between CERN and Intel to develop a co-packaged Xeon/FPGA processor.

(http://
u=http
fpga
proces
tested-
at-
cern%
At CERN, the European Organization for Nuclear Research (<http://home.cern/about>), physicists and engineers are probing the fundamental structure of the universe. The Large Hadron Collider (LHC), which began working in 2008, is the world's largest and most powerful particle accelerator (<http://home.cern/about/how-accelerator-works>); it is housed in an underground tunnel at CERN. Niko Neufeld is a deputy project leader at CERN who works on the Large Hadron Collider beauty (<http://lhcb-public.web.cern.ch/lhcb-public/>) (LHCb) experiment, which explores what happened after the Big Bang that allowed matter to survive and build the Universe we inhabit today.

(http://
mini=ti
fpga-
proces
tested-
at-
cern%
CERN experiments produce an enormous amount of data with forty million proton collisions every second, which leads to primary data rates of 40 terabytes per second," says Neufeld when speaking on a recent FPGA vs. CPU panel (<http://www-ssl.intel.com/content/www/us/en/high-performance-computing/architecting-hpc-with-fpgas-video.html>). "This is an enormous amount of data and there are a number of technical challenges in our work. We use a number of processing solutions including central processing units (CPUs), field-programmable gate arrays (FPGAs), and graphic processing units (GPUs), but each of these solutions have some limitations. We are collaborating with Intel in experimenting with a co-packaged Intel Xeon processor plus FPGA Quick Path Interconnect (QPI) processor in our LHCb research to try to determine which technology provides the best results."

(http://
url=htt
fpga-
proces
tested-
at-
cern%
CERN collaborates with leading ICT companies and other research institutes through a unique public-private partnership known as 'CERN openlab'. Its goal is to accelerate the development of cutting-edge solutions for the worldwide LHC community and wider scientific research. Through a CERN openlab project known as the 'High-Throughput Computing Collaboration,' researchers are investigating the use of various Intel technologies in data filtering and data acquisition systems.

(http://
url=htt
fpga-
proces
tested-
at-
cern%
Figure 1. CERN researchers shown in the Large Hadron Collider tunnel in front of the LHCb detector. Courtesy of CERN (courtesy CERN).



(<https://6lli539m39y3hpkelqsm3c2fg-wpengine.netdna-ssl.com/wp-content/uploads/2017/04/Fig1-CERN-researchers-LHC-collider-tunnel.jpg>)

Figure 1. CERN researchers shown in the Large Hadron Collider tunnel in front of the LHCb detector. Courtesy of CERN (courtesy CERN).

Introducing the co-packaged Intel CPU / FPGA Processor

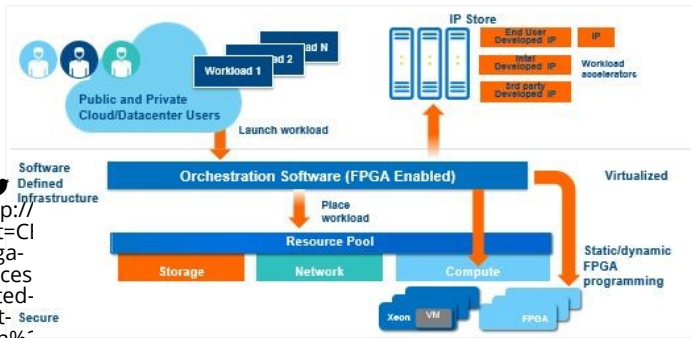
Today the CPU and FPGA are used as discrete chips in a solution – with an Intel Xeon processor and an FPGA which is typically attached via a PCIe interconnect to the CPU. And the development environment is also discrete using independent development tools from Intel and tools such as OpenCL and C++. Intel is working toward a common workflow and development flow to better integrate FPGAs.

"FPGA typically uses a higher level machine abstraction language (such as Verilog and VHDL) which have a painful low-level hardware programming model for most people. As a next step, Intel has a solution that co-packages the CPU and FPGA in the same Multichip Chip Product (MCP) package to deliver higher performance and lower latency than a discrete solution," states Bill Jenkins, Intel Senior AI Marketing Manager. The Intel MCP is supported by a cross-platform development framework like OpenCL that can be used to develop applications for both the CPU and FPGA. The Intel solution includes a fully unified intellectual property (IP) and development suite, including languages, libraries and development environments. The roadmap to a unified development flow leverages common tools and libraries to support both FPGA and Intel Xeon processor + FPGA systems along with an expansive ecosystem network of Intel and vendors working on independent development tools for demanding workloads such as HPC, imaging identification, security and big data.

Abstracting away FPGA Coding

Intel is building an abstraction layer (as part of the product containing the Intel CPU and FPGA in the same MCP package), called the Orchestration Software layer. This layer and the higher level IP and software models help make development less complex so that developers don't need to code specifically to the FPGA. The FPGA-enabled Orchestration software layer abstracts away the API to communicate with the FPGA as shown in the following example.

(http://text=Ci fpga-proces tested-at-cern%



(https://6lli539m39y3hpkelqsm3c2fg-wpengine.netdna-ssl.com/wp-content/uploads/2017/04/Fig2-CERN-example-of-Intel-implementation.jpg)

Figure 2. Example of Intel implementation of user IP implemented into FPGA via an abstraction Orchestration software layer

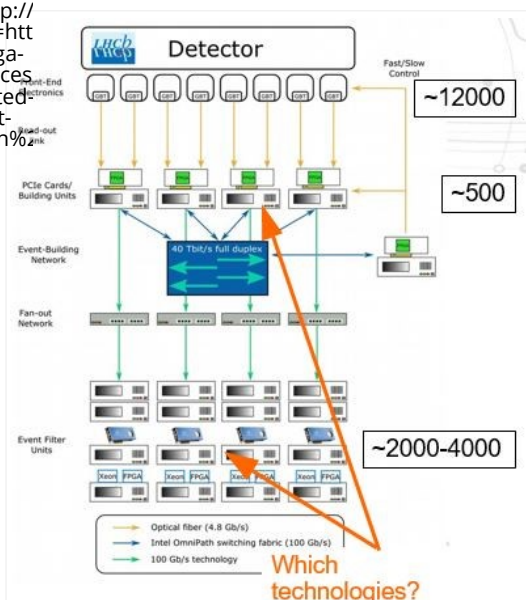
There is a cloud-based library of functions and end-user IP that have been pre-compiled and built that is loaded into the FPGA at runtime. The user first launches a workload from the host and it goes into the Orchestration software which pushes a function into the FPGA. This produces a bitstream that is pre-compiled on the FPGA to bring the data in—it is almost like a fixed architecture I/O interface.

In the example scenario, users simply download the image from the abstraction Orchestration software layer to the FPGA and it is ready to run without compilation. "With the abstraction Orchestration software layer," Jenkins explained, "Intel is abstracting away all the difficulties of FPGA programming using machine language tools while enabling all the higher level Intel frameworks including the Intel Trusted Analytics Platform (TAP) and Intel Scalable System Framework (SSI) and tying the FPGA into the frameworks. Intel is developing this approach for a variety of markets including visual understanding, analytics, enterprise, Network Function Virtualization (NFV), VPN, genomics, HPC and storage."

Large Hadron Collider High-Energy Physics Research at CERN

Neufeld indicates that the experiments at CERN — through what they refer to as 'online computing' — require a first-level data-filtering to reduce the data to an amount that can be stored and processed on more traditional processing units such as Intel Xeon processors. Figure 3 shows a schematic view of the future LHCb readout system. At the top level, there is a detector and optical fiber links, which transfer data out of the detector. CERN uses FPGAs to acquire data from the detector. There are also large switching fabrics, as well as clusters of processing elements including CPUs, FPGAs, and GPUs to reduce the amount of data. One of the questions the CERN team is testing is "Which technologies should we use and which provide the best performance and lowest energy usage results?"

(http://url=htt fpga-proces tested-at-cern%



(https://6lli539m39y3hpkelqsm3c2fg-wpengine.netdna-ssl.com/wp-content/uploads/2017/04/Fig3-CERN-schematic-diagram.jpg)

Figure 3. Schematic diagram showing future LHCb first-level data-filtering system. Courtesy of CERN.

CERN Tests Complex Cherenkov Angle Reconstruction Calculation

CERN has extensive experience using FPGAs in their research work. "We typically use FPGAs in our research to run algorithms looking for simple integer signatures, or for other less complicated calculations. When we heard about the Intel Xeon / FPGA combined processor, we chose a test using a complex algorithm to do a Cherenkov angle reconstruction of light emission in a particle detector, which is not typically performed on an FPGA. This involves tracing a light particle — photon — through a complex arrangement of optical reflection and deflection systems. Our test case used a rich PID algorithm to calculate the Cherenkov angle for each track and detection point. This is a complex mathematical calculation that involves hyperbolic functions, roots, square roots, etc., as shown in Figure 4. It is one of the most costly calculations done in online reconstruction," states Neufeld.

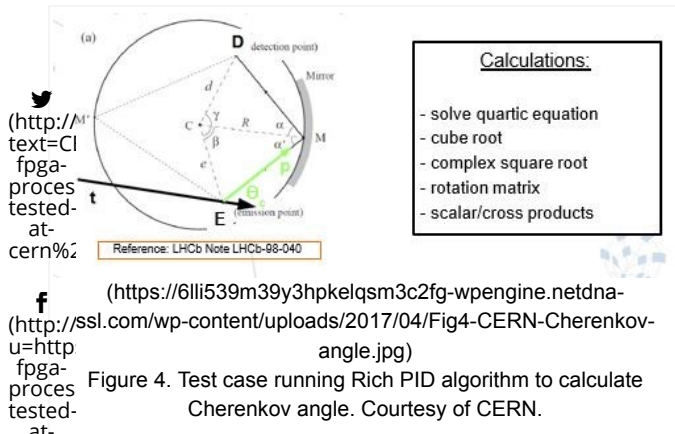


Figure 4. Test case running Rich PID algorithm to calculate Cherenkov angle. Courtesy of CERN.

Coding the Cherenkov Angle Reconstruction in Verilog versus OpenCL

The CERN team first implemented the Cherenkov angle reconstruction by coding it in the Verilog HDL. The team wrote a 748 clock-cycle long pipeline in Verilog, along with additional blocks developed for the test including: cubic root, complex square root, rotational matrix, and cross/scalar product. It was a lengthy task doing this coding in Verilog with 3,400 lines of code. With all test benches, the implementation took 2.5 months.

Next, the team recoded the Cherenkov angle code using the OpenCL (https://en.wikipedia.org/wiki/OpenCL) and the BSP (board support package) designed to work across a variety of hardware platforms. Because OpenCL is an abstraction language, it required only 250 lines of code and took two weeks of coding. Not only was coding in OpenCL much faster but the performance results were similar. Figure 5 shows the results of the Verilog versus OpenCL implementation.

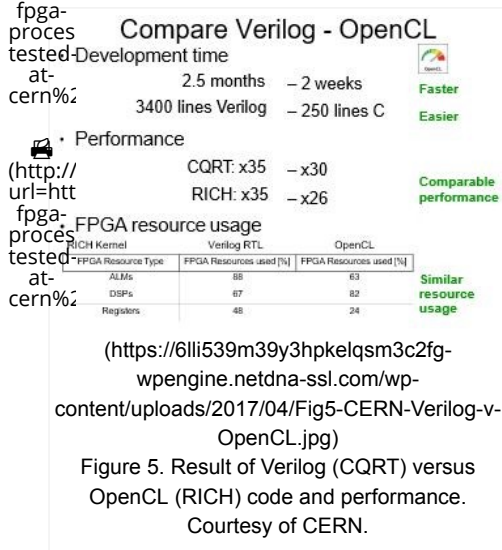


Figure 5. Result of Verilog (CQRT) versus OpenCL (RICH) code and performance. Courtesy of CERN.

CERN Compares Co-packaged Intel Xeon – FPGA Processor against Nallatech PCIe Stratix V FPGA Board

To test performance of the Verilog code, the CERN team used a commercially available Stratix V GXA7 FPGA board / Nallatech 385 board for testing. They achieved an acceleration of a factor up to six with the Stratix – Nallatech PCIe board. However, they found a bottleneck in data transfer—they could not keep the pipeline busy because the PCIe card was limited to an eight-lane interface. Next, the CERN team did tests with the Cherenkov angle code comparing a Nallatech FPGA Board with the co-packaged Intel Xeon/FPGA QPI processor.

CERN Systems used in Test

FPGA system used by CERN	Intel/FPGA system used by CERN
Nallatech 385 board	Two socket system:
FPGA: Altera Stratix V GX A7	<ul style="list-style-type: none"> First: Intel® Xeon® E5-2680 v2 Second: Intel/Altera Stratix V GX A7 FPGA <ul style="list-style-type: none"> 234'720 ALMs, 940'000 Registers, 256 DSPs
<ul style="list-style-type: none"> 234'720 ALMs, 940'000 Registers 256 DSPs 	
Programming model : OpenCL	Programming model : Verilog & OpenCL
Host Interface: 8-lane PCIe Gen3	Host Interface:
<ul style="list-style-type: none"> Up to 7.5GB/s 	<ul style="list-style-type: none"> High-bandwidth and low latency
Memory: 8GB DDR3 SDRAM	Memory: Cache-coherent access to main memory
Network Enabled with (2) SFP+ 10GbE ports	

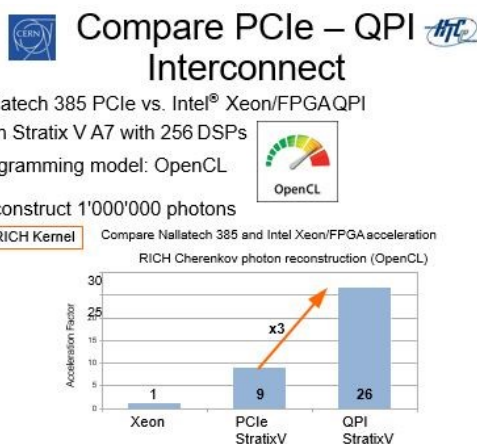
(http://text=fpga-process-tested-at-cern%)

f (http://u=http-fpga-process-tested-at-cern%)

in (http://mini=fpga-process-tested-at-cern%)

g (http://url=http-fpga-process-tested-at-cern%)

(https://6lli539m39y3hpkelqsm3c2fg-wpengine.netdna-ssl.com/wp-content/uploads/2017/04/Fig6-CERN-test-results-comparison.jpg)
 Figure 6. Test results from the CERN team comparing Intel Xeon CPU, PCIe Stratix V FPGA and Intel Xeon processor/FPGA QPI. Courtesy of CERN.



CERN Plans to do Future Testing using co-packaged Intel Xeon/ Intel Arria10 FPGA Processor

“Our CERN team found the results of using the co-packaged Intel Xeon processor/Stratix V QPI processors to be very encouraging. In addition, we find the programming model with OpenCL attractive and it will be mandatory for the High-Energy Physics (HEP) field. Intel will be launching a co-packaged Intel Xeon processor / Intel Arria 10 FPGA processor in the future. We want to do other experiments with the co-packaged Intel Xeon processor/ Arria 10 FPGA. We expect that the high-bandwidth interconnect and modern Arria 10 FPGA card will provide high performance and performance per Joule for HEP algorithms,” states Neufeld.

















Linda Barney is the founder and owner of Barney and Associates, a technical/marketing writing, training and web design firm in Beaverton, OR.

Share this:

[\(//www.reddit.com/submit?url=https://www.hpcwire.com/2017/04/14/xeon-fpga-processor-tested-at-cern/\)](https://www.reddit.com/submit?url=https://www.hpcwire.com/2017/04/14/xeon-fpga-processor-tested-at-cern/)

Leading Solution Providers

Altair (http://tci.taborcommunications.com/sponsor-altair-2)	Applied Micro (http://tci.taborcommunications.com/APM)
Asitek (http://tci.taborcommunications.com/sponsor-asetek)	Aspen Systems (http://tci.taborcommunications.com/sponsor-aspen)
ASRock (http://tci.taborcommunications.com/sponsor-asrock)	Atipa (http://tci.taborcommunications.com/sponsor-atipa)
Bull (http://tci.taborcommunications.com/sponsor-bull)	Cray (http://tci.taborcommunications.com/sponsor-cray)
DDN Storage (http://tci.taborcommunications.com/sponsor-ddn)	Dell EMC (http://tci.taborcommunications.com/sponsor-dell)
Gdell (http://tci.taborcommunications.com//21812/2016-07-28/5fzz61)	Hewlett Packard Enterprise (http://tci.taborcommunications.com/sponsor-hp-3)



(http://text=Cl fpga-proces tested-at-cern%

f Off The Wire

Industry Headlines



(http://u=htt

April 18, 2017

- Intel Says Goodbye to Intel Developer Forum (<https://www.hpcwire.com/off-the-wire/intel-says-goodbye-intel-developer-forum/>)
- PSSC Labs Announces New Release of CBeST Cluster Management Software Stack (<https://www.hpcwire.com/off-the-wire/pssc-labs-announces-new-release-cbest-cluster-management-software-stack/>)
- LTO Program Breaks Year-Over-Year Records for Tape Shipment (<https://www.hpcwire.com/off-the-wire/lto-program-breaks-year-year-records-tape-shipment/>)
- PRACE 14th Call for Proposals Awards Nearly 2000M Core Hours (<https://www.hpcwire.com/off-the-wire/prace-14th-call-proposals-awards-nearly-2000m-core-hours/>)
- Mellanox 25Gb/s Ethernet Adapters Chosen By Major ODMs (<https://www.hpcwire.com/off-the-wire/mellanox-25gbs-ethernet-adapters-chosen-major-odms/>)

(http://April 17, 2017

- OpenPOWER Foundation Announces Developer Congress focused on AI (<https://www.hpcwire.com/off-the-wire/openpower-foundation-announces-developer-congress-focused-ai/>)
- Baidu Advances AI in the Cloud with Latest NVIDIA Pascal GPUs (<https://www.hpcwire.com/off-the-wire/baidu-advances-ai-cloud-latest-nvidia-pascal-gpus/>)
- DOE's INCITE Program Seeks Advanced Computational Research Proposals for 2018 (<https://www.hpcwire.com/off-the-wire/incite-program-seeks-advanced-computational-research-proposals-2018/>)
- DDN Names Jessica Popp General Manager of IME Business Unit (<https://www.hpcwire.com/off-the-wire/ddn-names-jessica-popp-general-manager-ime-business-unit/>)

April 14, 2017

- Supermicro to Share Third Quarter Fiscal 2017 Financial Results (<https://www.hpcwire.com/off-the-wire/supermicro-share-third-quarter-fiscal-2017-financial-results/>)

April 13, 2017

- Larry Smarr Talks Machine Intelligence at Jackson State (<https://www.hpcwire.com/off-the-wire/larry-smarr-talks-machine-intelligence-jackson-state/>)
- Intel Promotes Three Executives (<https://www.hpcwire.com/off-the-wire/intel-promotes-three-executives/>)
- Engility to Pursue NASA Advanced Computing Services Opportunity (<https://www.hpcwire.com/off-the-wire/engility-pursue-nasa-advanced-computing-services-opportunity/>)

April 12, 2017

- Tutorials Schedule Announced for PEARC17 (<https://www.hpcwire.com/off-the-wire/tutorials-schedule-announced-pearc17/>)

(http://url=htt

fpga-

proces

tested-

at-

cern%

HPC Job Bank

Storage Engineer/Senior Storage Engineer - NCSA (<http://careers.hpcwire.com/jobdetails.cfm?jid=2424>)

View this Career Listing (<http://careers.hpcwire.com/jobdetails.cfm?jid=2424>)

HPC Data Storage Administrator - Scientist 2/3 - Los Alamos National Laboratory (<http://careers.hpcwire.com/jobdetails.cfm?jid=2457>)

View this Career Listing (<http://careers.hpcwire.com/jobdetails.cfm?jid=2457>)

More Career Resources [▶▶ \(http://careers.hpcwire.com\)](http://careers.hpcwire.com)

Subscribe to HPCwire's Weekly Update!

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

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

- THE LATEST
- EDITOR'S PICKS
- MOST POPULAR



Facebook Open Sources Caffe2; Nvidia, Intel Rush to Optimize

(<https://www.hpcwire.com/2017/04/18/facebook-open-sources-caffe2-nvidia-intel-jumpstart-optimizations/>)


From its F8 developer conference in San Jose, Calif., today, Facebook announced Caffe2, a new open-source, cross-platform framework for deep learning. Caffe2 is to Caffe, the deep learning framework developed by Berkeley AI Research and community contributors. Read more... (https://www.hpcwire.com/2017/04/18/facebook-caffe2-vidia-intel-jumpstart-optimizations/)

By Tiffany Trader

<http://twitter.com/intent/tweet?status=Facebook+Open+Sources+Caffe2%3B+Nvidia%2C+Intel+Rush+to+Optimize+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Fopen-sources-caffe2-vidia-intel-jumpstart-optimizations%2F> <http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Ffacebook-open-sources-caffe2-vidia-intel-jumpstart-optimizations%2F&title=Facebook+Open+Sources+Caffe2%3B+Nvidia%2C+Intel+Rush+to+Optimize&source=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Ffacebook-open-sources-caffe2-vidia-intel-jumpstart-optimizations%2F&title=Facebook+Open+Sources+Caffe2%3B+Nvidia%2C+Intel+Rush+to+Optimize/> <https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Ffacebook-open-sources-caffe2-vidia-intel-jumpstart-optimizations%2F>

(http://
text=C
fpga-
proces
tested-
at-
cern%

f
(http://
u=http
fpga-
proces
tested-
at-
cern%



in Knights Landing Processor with Omni-Path Makes Cloud Debut

(http://www.hpcwire.com/2017/04/18/knights-landing-processor-omni-path-makes-cloud-debut/)

Cloud specialist Rescale is partnering with Intel and HPC resource provider R Systems to offer first-ever cloud access to Xeon Phi "Knights Landing" processors. Infrastructure is based on the 68-core Intel Knights Landing processor with integrated Omni-Path fabric (the 7250F Xeon Phi). Read more... (https://www.hpcwire.com/2017/04/18/knights-landing-processor-omni-path-makes-cloud-debut/)

By Tiffany Trader

<http://twitter.com/intent/tweet?status=Knights+Landing+Processor+with+Omni-Path+Makes+Cloud+Debut+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Fknights-landing-processor-omni-path-makes-cloud-debut%2F> <http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Fknights-landing-processor-omni-path-makes-cloud-debut%2F&title=Knights+Landing+Processor+with+Omni-Path+Makes+Cloud+Debut&source=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Fknights-landing-processor-omni-path-makes-cloud-debut%2F&title=Knights+Landing+Processor+with+Omni-Path+Makes+Cloud+Debut/> <https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F18%2Fknights-landing-processor-omni-path-makes-cloud-debut%2F>

(http://
url=htt
fpga-
proces
tested-
at-
cern%




IBM Coils Anaconda Around Power Processor

(https://www.hpcwire.com/2017/04/17/ibm-coils-anaconda-around-power-processor/)

IBM, which recently extended support for the Anaconda data science platform to its open source mainframe, takes another step this week by offering the data platform Cognitive Systems platform in collaboration with Anaconda developer Continuum Analytics. Read more... (https://www.hpcwire.com/2017/04/17/ibm-coils-anaconda-around-power-processor/)

By George Leopold

<http://twitter.com/intent/tweet?status=IBM+Coils+Anaconda+Around+Power+Processor+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F17%2Fibm-coils-anaconda-around-power-processor%2F> <http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F17%2Fibm-coils-anaconda-around-power-processor%2F&title=IBM+Coils+Anaconda+Around+Power+Processor&source=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F17%2Fibm-coils-anaconda-around-power-processor%2F&title=IBM+Coils+Anaconda+Around+Power+Processor/> <https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F17%2Fibm-coils-anaconda-around-power-processor%2F>



CERN openlab Explores New CPU/FPGA Processing Solutions


(<https://www.hpcwire.com/2017/04/14/xeon-fpga-processor-tested-at-cern/>)

Through a CERN openlab project known as the 'High-Throughput Computing Collaboration,' researchers are investigating the use of various Intel technologies in data acquisition systems. Read more... (<https://www.hpcwire.com/2017/04/14/xeon-fpga-processor-tested-at-cern/>)

By Linda Barney

Share on social media: [Twitter](http://twitter.com/intent/tweet?status=CERN+openlab+Explores+New+CPU%2FFPGA+Processing+Solutions+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F14%2Fxeon-fpga-processor-tested-at-cern%2F), [LinkedIn](http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F14%2Fxeon-fpga-processor-tested-at-cern%2F&title=CERN+openlab+Explores+New+CPU%2FFPGA+Processing+Solutions&source=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F14%2Fxeon-fpga-processor-tested-at-cern%2F), [Facebook](http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F14%2Fxeon-fpga-processor-tested-at-cern%2F&title=CERN+openlab+Explores+New+CPU%2FFPGA+Processing+Solutions/), [Google+](https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F14%2Fxeon-fpga-processor-tested-at-cern%2F)

HPE Extreme Performance Solutions



HPC-Driven Weather Simulations Improving Forecasting Capabilities

(https://www.hpcwire.com/solution_content/hpe/weather-climate/hpc-driven-weather-simulations-improving-forecasting-capabilities/)

In September of 1938, a massive hurricane (<http://sos.ri.gov/virtualarchives/items/show/291>) traversed the Atlantic Ocean and made landfall in New England. Due to incorrect forecasting, the storm struck farther north and with greater intensity than had been predicted, leaving residents and authorities with virtually no warning or time to prepare. Read more... (https://www.hpcwire.com/solution_content/hpe/weather-climate/hpc-driven-weather-simulations-improving-forecasting-capabilities/)


Visit the



(https://www.hpcwire.com/solution_channel/hpe/)

Previous:

- AI and Deep Learning Come to Wall Street (https://www.hpcwire.com/solution_content/hpe/financial-services/ai-deep-learning-come-wall-street/)
- Enhancing the Speed and Security of Energy Production with Remote Visualization (https://www.hpcwire.com/solution_content/hpe/energy-oil-gas/enhancing-speed-energy-production-remote-visualization/)
- HPE Servers Demonstrate Critical Speed for HFT Traders on STAC-N1 Test (https://www.hpcwire.com/solution_content/hpe/financial-services/hpe-servers-demonstrate-critical-speed-hft-traders-stac-n1-test/)



DOE Supercomputer Achieves Record 45-Qubit Quantum Simulation

(<https://www.hpcwire.com/2017/04/13/doe-supercomputer-45-qubit-quantum-computer/>)

In order to simulate larger and larger quantum systems and usher in an age of "quantum supremacy," Read more... (https://www.hpcwire.com/2017/04/13/doe-super-qubit-quantum-computer/)

By Tiffany Trader

🐦 (http://twitter.com/intent/tweet?status=DOE+Supercomputer+Achieves+Record+45-

Qubit+Quantum+Simulation+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F13%2Fdoe-supercomputer-45-qubit-quantum-computer%2F (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F13%2Fdoe-supercomputer-45-qubit-quantum-computer%2F&title=DOE+Supercomputer+Achieves+Record+45-Qubit+Quantum+Simulation&source=https%3A%2F%2Fwww.hpcwire.com/) **f**

(http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F13%2Fdoe-supercomputer-45-qubit-quantum-computer%2F&title=DOE+Supercomputer+Achieves+Record+45-Qubit+Quantum+Simulation/) **G+** (https://plus.google.com/share?

url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F13%2Fdoe-supercomputer-45-qubit-quantum-computer%2F

(http://

url=htt

text=Cl

fpga-

proces

tested-

at-

cern%

f

(http://

u=http

fpga-

proces

tested-

at-

cern%



in Hyperscalers Emerging From 'Hype Phase'

(http://
url=https://www.hpcwire.com/2017/04/12/hyperscalers-emerging-hype-phase/)

mini-

fpga-

proces

tested-

at-

cern%

The number of companies fitting the description of "hyperscale" now accounts for 68 percent of the cloud infrastructure services market, a researcher found. Read more (https://www.hpcwire.com/2017/04/12/hyperscalers-emerging-hype-phase/)

By George Leopold

🐦 (http://twitter.com/intent/tweet?

(http://

url=https://www.hpcwire.com/2017/04/12/hyperscalers-emerging-hype-phase%2F

emerging-hype-phase%2F) **in** (http://www.linkedin.com/shareArticle?

mini-

fpga-

proces

tested-

at-

cern%

status=Hyperscalers+Emerging+From+%E2%80%98Hype+Phase%E2%80%99+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fhyperscalers-emerging-hype-phase%2F&title=Hyperscalers+Emerging+From+%E2%80%98Hype+Phase%E2%80%99&source=https%3A%2F%2Fwww.hpcwire.com/) **f**

(http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fhyperscalers-emerging-hype-phase%2F&title=Hyperscalers+Emerging+From+%E2%80%98Hype+Phase%E2%80%99/)

G+ (https://plus.google.com/share?

url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fhyperscalers-emerging-hype-phase%2F)

(http://

url=htt

fpga-

proces

tested-

at-

cern%



Penguin Takes a Run at the Big Cloud Providers

(https://www.hpcwire.com/2017/04/12/penguin-takes-run-big-cloud-providers/)

HPC specialist Penguin Computing recently re-ran benchmarks from a study of its larger brethren and says the results show its 'public cloud' – Penguin on Demand (Penguin among the leaders in cost and performance. Read more... (https://www.hpcwire.com/2017/04/12/penguin-takes-run-big-cloud-providers/)

By John Russell

🐦 (http://twitter.com/intent/tweet?

status=Penguin+Takes+a+Run+at+the+Big+Cloud+Providers+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fpenguin-takes-run-big-cloud-providers%2F) **in** (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fpenguin-takes-run-big-cloud-providers%2F&title=Penguin+Takes+a+Run+at+the+Big+Cloud+Providers&source=https%3A%2F%2Fwww.hpcwire.com/) **f**

(http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fpenguin-takes-run-big-cloud-providers%2F&title=Penguin+Takes+a+Run+at+the+Big+Cloud+Providers/)

G+ (https://plus.google.com/share?

url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F12%2Fpenguin-takes-run-big-cloud-providers%2F)



DDN Claims Faster Object Storage 'Cross-Region' S3 Connectivity than AWS

(<https://www.hpcwire.com/2017/04/11/ddn-claims-faster-object-storage-cross-region-s3-connectivity-aws/>)

Advanced-scale storage specialist DataDirect Networks (DDN), whose solutions populate the upper-rung of the Top500 supercomputing list, today announced new "cr" features to its WOS object storage platform that support multi-site AWS connectivity as well as new data protection and disaster recovery capabilities. Read more... (<https://www.hpcwire.com/2017/04/11/ddn-claims-faster-object-storage-cross-region-s3-connectivity-aws/>)

By Doug Black

(<http://twitter.com/intent/tweet?status=DDN+Claims+Faster+Object+Storage+%E2%80%98Cross-Region%E2%80%99+S3+Connectivity+than+AWS+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fddn-claims-faster-object-storage-cross-region-s3-connectivity-aws%2F>) **in** (<http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fddn-claims-faster-object-storage-cross-region-s3-connectivity-aws%2F&title=DDN+Claims+Faster+Object+Storage+%E2%80%98Cross-Region%E2%80%99+S3+Connectivity+than+AWS&source=https%3A%2F%2Fwww.hpcwire.com/>) **f** (<http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fddn-claims-faster-object-storage-cross-region-s3-connectivity-aws%2F&title=DDN+Claims+Faster+Object+Storage+%E2%80%98Cross-Region%E2%80%99+S3+Connectivity+than+AWS/>) **G+** (<https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fddn-claims-faster-object-storage-cross-region-s3-connectivity-aws%2F>)

Leading Solution Providers

Altair (<http://tci.taborcommunications.com/sponsor-altair-2>) **applied micro** (<http://tci.taborcommunications.com/APM>)

ASETTEK (<http://tci.taborcommunications.com/sponsor-asetek>) **Aspen Systems** (<http://tci.taborcommunications.com/sponsor-aspen>)

ASRock (<http://tci.taborcommunications.com/sponsor-asrock>) **atipa** (<http://tci.taborcommunications.com/sponsor-atipa>)

Bull (<http://tci.taborcommunications.com/sponsor-bull>) **CRAY** (<http://tci.taborcommunications.com/sponsor-cray>)

DDN STORAGE (<http://tci.taborcommunications.com/sponsor-ddn>) **DELL EMC** (<http://tci.taborcommunications.com/sponsor-dell>)

GDEL (<http://tci.taborcommunications.com//21812/2016-07-28/5fzz61>) **Hewlett Packard Enterprise** (<http://tci.taborcommunications.com/sponsor-hp-3>)

IBM (<http://tci.taborcommunications.com/sponsor-ibm>) **inspur** (<http://tci.taborcommunications.com/sponsor-inspur>)

intel (<http://tci.taborcommunications.com/sponsor-intel>) **Lenovo** (<http://tci.taborcommunications.com/sponsor-lenovo>)

Mellanox (<http://tci.taborcommunications.com/sponsor-mellanox>) **Microsoft** (<http://tci.taborcommunications.com/sponsor-microsoft>)

NEC (<http://tci.taborcommunications.com/sponsor-nec>) **PENGUIN COMPUTING** (<http://tci.taborcommunications.com//21812/2014-04-25/5l3mh>)

PGI (<http://tci.taborcommunications.com/sponsor-pgi>) **PSCLABS** (<http://tci.taborcommunications.com/sponsor-PSCLabs>)

PURE STORAGE (<http://tci.taborcommunications.com/sponsor-purestorage>) **re-store-2** (<http://tci.taborcommunications.com/re-store-2>)

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

SUPERMICRO (<http://tci.taborcommunications.com/sponsor-supermicro>) **VERNE GLOBAL** (<http://tci.taborcommunications.com/verneglobal>)



(http://
text=Cl
fpga-
proces
tested-
at-
cern%

f
(http://
u=http
fpga-
proces
tested-
at-
cern%

Dutch Researchers Build Little Green Machine II Out of IBM Minsky Servers

(https://www.hpcwire.com/2017/04/11/dutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers/)

(http://
mini-
fpga-
proces
tested-
at-
cern%

(http://
url=htt
fpga-
proces
tested-
at-
cern%

(http://
url=htt
fpga-
proces
tested-
at-
cern%

The newest Dutch HPC cluster is only four "pizza boxes" high, small enough to be carried around on a bicycle, but Little Green Machine II, as the new system is called, is used for research in oceanography, computer science, artificial intelligence, financial modeling and astronomy. Read more... (https://www.hpcwire.com/2017/04/11/dutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers/)

By Tiffany Trader

(http://twitter.com/intent/tweet?

status=Dutch+Researchers+Build+Little+Green+Machine+II+Out+of+IBM+Minsky+Servers+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fdutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers%2F)

in (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fdutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers%2F&title=Dutch+Researchers+Build+Little+Green+Machine+II+Out+of+IBM+Minsky+Servers&source=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fdutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers/)

(http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fdutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers%2F&title=Dutch+Researchers+Build+Little+Green+Machine+II+Out+of+IBM+Minsky+Servers/)

G+ (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fdutch-uni-builds-little-green-machine-ii-out-of-ibm-minsky-servers%2F)

A ♠

冷扑大师 VS 中国龙之队

人工智能和顶尖牌手巅峰表演赛

2017.4.6-4.10 海南

联合主办: 创新工场 SINOVIATION VENTURES RSC 海南生态软件园

承办: 传奇扑克 LEGEND POKER 协办: 赢之星 首席顾问: 杜悦

支持单位: 亞美娛樂 信息娱乐第一 DXRACER 王境 品淘商业评论

(http://text=Cl fpga-proces tested-at-cern%

f (http://u=http fpga-proces tested-at-cern%

in (http://min=ti fpga-proces tested-at-cern%

s (http://url=htt fpga-proces tested-at-cern%

(http://url=htt fpga-proces tested-at-cern%

CMU's Latest Poker Program Wins in China

It's now become old hat – AI poker algorithm beats humans. Carnegie Mellon University is at the forefront of creating this new class of crafty AI gamblers. Read more. (https://www.hpcwire.com/2017/04/11/cmus-latest-poker-program-wins-china/)

By John Russell
<http://twitter.com/intent/tweet?status=CMU%E2%80%99s+Latest+Poker+Program+Wins+in+China+https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-china%2F> [in](http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-china%2F&title=CMU%E2%80%99s+Latest+Poker+Program+Wins+in+China&source=https%3A%2F%2Fwww.hpcwire.com/) (http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-china%2F&title=CMU%E2%80%99s+Latest+Poker+Program+Wins+in+China&source=https%3A%2F%2Fwww.hpcwire.com/) [f](http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-poker-program-wins-china%2F&title=CMU%E2%80%99s+Latest+Poker+Program+Wins+in+China/) (http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-poker-program-wins-china%2F&title=CMU%E2%80%99s+Latest+Poker+Program+Wins+in+China/) [G+](https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-poker-program-wins-china%2F) (https://plus.google.com/share?url=https%3A%2F%2Fwww.hpcwire.com%2F2017%2F04%2F11%2Fcmus-latest-poker-program-wins-china%2F)