



## June 2019 Newsletter: ISC19 Special Edition

### Intel® Developer Connect

[Intel® Developer Connect @ Hotel Mövenpick](#) in Frankfurt, Germany June 17–19 is the MUST attend technical training event to learn what's next in HPC and AI. All sessions are open to the public, but seats are filling up fast! [View the full schedule & register NOW](#) to connect with Intel® architecture experts and industry leaders in technical sessions & hands-on tutorials, including:

## THE CONFLUENCE OF HPC AND AI – A CONVERGED JOURNEY WITH INTEL

### Build, Accelerate, Scale and Deploy AI Applications on Intel HPC

Monday, June 17 | 11:00 a.m.–2:00 p.m.

Driven by an exponential increase and availability in volume and diversity of data, Artificial Intelligence (AI) is transforming many businesses around the globe by enabling them to drive operational efficiencies and build new products and services. AI has also begun to reshape the landscape of scientific computing and enabling scientists to address large problems in ways that were not possible before. Intel collaborates with customers and partners worldwide to build, accelerate, scale and deploy their AI applications on Intel based HPC platforms. We share with you our strategy, roadmap and insights on several customer AI use cases we have enabled, the orders of magnitude performance acceleration we have delivered via popular open-source software framework optimizations, and the best-known methods to advance the convergence of AI and High Performance Computing on Intel® Xeon™ Scalable Processor based servers.

## TEXAS ADVANCED COMPUTING CENTER (TACC) TUTORIAL ON FRONTERA

Programming for Frontera, TACC's Cascade Lake Petascale System

Tuesday, June 18 | 12:00–5:00 p.m.

Visualizing black holes, Intel® Xeon™ Scalable Processors and the great state of Texas. What do they have in common? The Texas Advanced Computing Center (TACC) where Intel® Xeon™ Scalable Processors were part of the computing resources used in the recent breakthrough producing the first image of a black hole. The latest super computer at TACC, known as Frontera, utilizes the Intel's newest architecture, Cascade Lake (CLX) which scales up to 28 cores per CPU, or, in the case of Cascade Lake-AP, 48 cores per chip. The hardware is only part of the story to achieving breakthrough performance on a petascale system; knowing how to optimize and tune software is critical as well. Join team members from TACC and Intel to learn how to optimize and tune applications through applying vectorization/threading techniques, optimized caching, balancing MPI computation and communications and much more. Hands on labs will be performed remotely on Frontera. Please bring a laptop that is WIFI ready to participate in the labs.

## ISC'19 Intel Activities

**The Convergence of HPC and Artificial Intelligence:** The convergence of HPC and AI facilitates powerful ways to tackle previously impossible design, engineering, and scientific challenges. At ISC, Intel will highlight technologies that accelerate this convergence, including the only data center processors with built-in AI acceleration. Visit the Intel booth June 17–19 to explore how Intel® technology is enabling this convergence. Bookmark [Intel's ISC19 mobile agenda](#) to your phone and keep up to date on speakers, trainings, keynotes, and more...check back often for the latest and greatest from Intel.

## Highlights



### Optimization Techniques

#### [Intel® CPU Outperforms NVIDIA\\* GPU on ResNet-50 Deep Learning](#)

[Inference](#): demonstrated the effectiveness of Intel Xeon processors with optimized DL software, and achieved the throughput of ResNet-50 7878 images per second on Intel Xeon Platinum 9282 processors.

[Effectively Using Your Whole Cluster](#): provide a step-by-step methodology to improve the performance of [SPECFEM3D\\_GLOBE\\*](#) with Intel technology.

[New Breakthroughs to Overcome Memory and Storage Barriers](#): how new and affordable persistent memory enables developers to create flexible designs for managing workloads that meet optimal costs and balance performance tradeoffs.

[Unlocking Deep Learning Performance with nGraph](#): nGraph supports a range of deep learning frameworks and hardware back-ends and delivers dramatic performance improvements on a range of platforms.



### Case Studies

[Deploying AI Frameworks on Secure HPC Systems with Containers](#): first results on scaling out the training process achieve close to 94% scaling efficiency up to 128 nodes on Intel Xeon processors.

[Evaluation of Intel Memory Drive Technology \(IMDT\) Performance for Scientific Applications](#): some applications DDR4/Optane hybrid configuration outperforms DDR4 setup by up to 20%.

[Deeply-Pipelined FPGA Clusters Make DNN Training Scalable](#): as a result, storage demand is reduced to the point where only on-chip memory is used for the convolution layers. On average, 6.36x higher energy efficiency than comparable GPU servers.

[Wind farm interference and terrain interaction simulation by means of an adaptive actuator disc](#): wind farm simulation makes parallel computation mandatory. In order to carry out the flow simulations, the free and open source computational fluid dynamic (CFD) software OpenFOAM is used.

[Software-Defined Visualization with Intel Rendering Framework – No Special Hardware Needed](#): explores how the Intel Rendering Framework can deliver better performance at a higher degree of fidelity without having to invest in extra hardware.



### Scientific Breakthrough

[Fast, flexible, and functional: 4 real-world AI deployments at enterprise scale](#): learn how companies are taking advantage of hardware enhancements and hardware optimizations to deploy enterprise-scale AI on existing infrastructure.

[Cashing In on HPC and AI in the Financial Services Industry](#): drawing on the power of artificial intelligence applications and high-performance computing systems, financial services companies are taking their businesses to new heights.

[How AI is Helping Us Better Understand the Environment](#): AI-empowered systems provide improved understanding of complex, changing patterns in the environment—from ocean currents and rainfall patterns to animal migration and human population trends.

[Build Faster AI Solutions with the Intel-Optimized ONNX Runtime](#): developers can train a model with any popular framework (including PyTorch and TensorFlow), convert it to ONNX format, and inference efficiently across a wide range of hardware with ONNX Runtime.

[Pixar Animation Studios Releases RenderMan 22.5](#): Pixar released RenderMan version 22.5, featuring a completely rewritten RenderMan for Houdini plugin along with important new features and performance gains.

## Speaker & Publication Opportunities

There are several opportunities for you to share your learnings, best practices and techniques around the benefits you've received in leveraging Intel® architecture. We would like bring to your attention some key abstract submission deadlines for 2019 conferences and workshops. Feel free to submit abstracts to all that interest you.

Submission Deadline	Event
July 1, 2019	<a href="#">ALCF Data Science Program</a>
July 8, 2019	<a href="#">IXPUG Annual Conference 2019</a>
July 31, 2019	SC19 <a href="#">Research Posters</a>
July 31, 2019	SC19 <a href="#">Scientific Visualization &amp; Data Analytics Showcase</a>
July 31, 2019	SC19 <a href="#">Bird of A Feather(BoF)</a>
July 31, 2019	SC19 <a href="#">Early Career Program Applications</a>
July 31, 2019	SC19 <a href="#">Doctoral Showcase</a>
July 31, 2019	SC19 <a href="#">Exhibitor Forum</a>
July 31, 2019	SC19 <a href="#">HPC Impact Showcase</a>

## Global Event & Training Opportunities

We encourage you to participate in any of the upcoming global training and free webinar opportunities.

Date	Location	Event
June 4-5, 2019	Edinburgh, UK	<a href="#">2nd UK OpenMP Users Conference</a>
June 5, 2019	CERN, Switzerland	<a href="#">Speeding up Scientific Codes in HPC Architectures by Code Modernization: Lessons Learned</a>
June 10-13, 2019	TACC, Austin	<a href="#">Scientific Visualization</a>
June 10-14, 2019	Houston, Texas	<a href="#">Rice University Tapia Center Workshop in Computational Chemistry</a>
June 10-15, 2019	Long Beach, CA	<a href="#">ICML 2019</a>
June 15-21, 2019	Long Beach, CA	<a href="#">CVPR 2019</a>
June 16-20, 2019	Frankfurt, Germany	<a href="#">ISC 2019</a>
June 18-19, 2019	Beijing, PRC	<a href="#">O'REILLY + Intel Artificial Intelligence Conference</a>
June 25-28, 2019	TACC, Austin	<a href="#">Designing and Administering Large-scale Systems</a>
July 9, 2019	NERSC, Berkeley, CA	<a href="#">Big Data Summit 2019: AI and HPC Convergence for Science</a>
July 9, 2019	Munich, Germany	<a href="#">AI Dev Conference-Intel Summit Series'19</a>
July 11, 2019	Webinar	<a href="#">Scaling Distributed TensorFlow Training with Intel's nGraph Library on Xeon® Processor Based HPC Infrastructure</a>
July 11, 2019	Shenzhen, PRC	<a href="#">AI Dev Conference-Intel Summit Series'19</a>
July 9-12, 2019	TACC, Austin	<a href="#">Workflows and Reproducibility in Scientific Computing</a>

July 17-21, 2019	Cookeville, TN	<a href="#">Integrating Parallel and Distributed Computing in Introductory Programming Classes Workshops</a>
July 22-31, 2019	TACC, Austin	<a href="#">MoISSI 2019 Software Summer School</a>
July 23-26, 2019	TACC, Austin	<a href="#">Computational Science in the Cloud</a>
July 28-August 1, 2019	Chicago, IL	<a href="#">PEARC 19</a>
August 5-8, 2019	TACC, Austin	<a href="#">Machine Learning Foundations</a>
August 7, 2019	Bangalore, India	<a href="#">AI Dev Conference-Intel Summit Series'19</a>
August 8-9, 2019	Arlington, Virginia	<a href="#">MoISSI "Teach the Teachers" Instructor Training</a>
August 26-30, 2019	Göttingen, Germany	<a href="#">EURO-PAR 2019</a>
September 9 -12, 2019	San Jose, CA	<a href="#">O'REILLY +AI Artificial Intelligence Conference</a>
September 9-11, 2019	Auckland, New Zealand	<a href="#">OpenMPCon 2019</a>
September 11-13, 2019	Auckland, New Zealand	<a href="#">IWOMP 2019</a>
September 15 – 20, 2019	San Antonio, TX	<a href="#">SEG 19</a>
September 17-18, 2019	Mountain View, CA	<a href="#">AI Hardware Summit 2019</a>
September 17-19, 2019	TACC, Austin	<a href="#">HPC Leadership</a>
September 24-27, 2019	Geneva, Switzerland	<a href="#">IXPUG Annual Conference 2019</a>
October 4, 2019	Tokyo, Japan	<a href="#">AI Dev Conference-Intel Summit Series'19</a>
October 14-17, 2019	London, UK	<a href="#">O'REILLY + Intel Artificial Intelligence Conference</a>
October 27-November 2, 2019	Seoul, South Korea	<a href="#">International Conference on Computer Vision(ICCV)</a>
October 29, 2019	London, England	<a href="#">AI Dev Conference-Intel Summit Series'19</a>
November 17 -20, 2019	Brisbane, Australia	<a href="#">SIGGRAPH Asia 2019</a>
December 2-8 , 2019	Vancouver, CA	<a href="#">NeurIPS 2019</a>

## More News

Check out the latest Intel® news:

- [Intel Xe GPU Supports Ray Tracing Through Hardware](#)
- [Persistent Memory: The Keystone to Fast Data Access at Scale with a Lower TCO](#)
- [Intel Unveils Clear Linux OS Update at Open Source Summit](#)
- [Intel Unified Clusters Bring Together AI, Analytic, and Traditional HPC Workloads](#)

© 2019, Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.

To subscribe to the Intel PCC mailing list, please register [HERE](#). To [unsubscribe](#) from other Intel communications, please reply to those directly. Or contact us at this address: Intel Corporation, 2200 Mission College Blvd., M/S SC3-37, Attn: Unsubscribe/Privacy, Santa Clara, CA 95054. Intel Corporation has never engaged in the practice of sharing information about individual subscribers or sharing it with third parties. [Intel Privacy Policy](#)