



AVX-512  
PERFORMANCE  
MPI PYTHON  
DATA ANALYTICS XEON  
XEON PHI  
MACHINE LEARNING  
VECTORIZATION  
THREADING

## INTEL® CODE MODERNIZATION WORKSHOPS 2016

AMSTERDAM - SEPTEMBER 20+21  
PRAGUE - OCTOBER 4+5  
MUNICH - OCTOBER 25+26

REGISTER NOW AT [WWW.INTELDEVCONFERENCE.COM](http://WWW.INTELDEVCONFERENCE.COM)

[www.hocomputer.de](http://www.hocomputer.de) - [info@hocomputer.de](mailto:info@hocomputer.de) - Tel: (+49) / 0221 / 76 20 86

\* 2016 h.o.-COMPUTER Software GmbH, Amsterdamer Str. 91, D-50735 Köln, HRB 22609. Nachdruck, Vervielfältigung oder Publikation in elektronischen Medien nur mit ausdrücklicher, schriftlicher Genehmigung. Druckfehler, Preisänderungen, Versionswechsel und Irrtümer vorbehalten. Alle Angaben ohne Gewähr, Verkauf nur an gewerbliche Abnehmer (B2B). © 2016 Intel Corporation Intel, the Intel logo, Pentium, Itanium, Intel Xeon and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Weitere Informationen zur Leistungssteigerung und Optimierung mit Intel Softwareprodukten erhalten Sie hier: <http://software.intel.com/en-us/articles/optimization-notice> \*



Liebe Leserinnen und Leser,

Europa wächst immer weiter zusammen und auch unser Kundenkreis ist schon lange nicht mehr auf den deutschsprachigen Raum begrenzt. Auch aus diesem Grund ist dieser Newsletter, vom Editorial abgesehen, ganz in englischer Sprache. Da es heute vor allem um die neue Version 2017 von Intel Parallel Studio XE geht, sollte dies kein Problem sein - es ist sowieso meist wenig sinnvoll etwa Compilerfeatures zu übersetzen.

Seit mehr als 25 Jahren sind wir Ihr verlässlicher Partner wenn es um HPC-Programmierung geht. Als einer der umsatzstärksten Intel Elite Reseller in ganz EMEA stehen wir Ihnen für alle Fragen rund um die Intel Compiler und Tools zur Verfügung.

Auch in diesem Herbst Jahr sind wir wieder auf zahlreichen Veranstaltungen vertreten oder organisieren diese selbst. Besuchen Sie uns einfach, z.B. in Amsterdam, Prag oder München.

Für heute alles Gute und viele Grüße aus Köln!

Dear readers,



Europe is growing and gets closer and closer together. Since several years we act and feel more like a European than a local German company and of course many of our customers speak other languages than German. So it is high time to provide this newsletter in English.

ho-COMPUTER is in business for more than 25 years now. We are focused on high performance compilers and tools, mostly in C++ and Fortran. Since a long time we are one of the largest so called Intel Software Elite reseller in EMEA, providing excellent support to our worldwide customers before and after sales.

I am looking forward to hear from you soon or see you face-to-face at one of our events in Amsterdam, Prague or Munich! All the best from Germany

Harald Odendahl, Managing Director

## Application Optimization Workshop with Hands-on Training - 26.10.

This specialized 1-day software optimization training will be run by tools experts from Intel Software and Services Group. The training targets Software developers using C/C++ & Fortran and related analysis tools (for example VTune™ Amplifier XE) and focuses on industry oriented hands-on labs using Intel Architecture based CPUs (x86).

Notebooks will be provided by Intel in order to analyze data collected. You will learn how to achieve significant performance improvements and create more reliable serial and parallel applications.

The event will take place in Garching close by to Munich at the LRZ (Leibniz Rechenzentrum) on Wednesday, 26.

10. from 9 am to 5pm. Lunch and coffee will be provided.

We would like to point out that the training will be taking place at the same location as the Code Modernization Workshop, running in parallel.

Please register at [hocomputer.de/et](http://hocomputer.de/et) quickly as seats are limited. This event is free of charge.



# Intel Parallel Studio XE 2017 - 3 different editions

## What's Included in Intel Parallel Studio XE

	Composer Edition <sup>1</sup>	Professional Edition <sup>1</sup>	Cluster Edition
Intel C++ Compiler	✓	✓	✓
Intel® Fortran Compiler	✓	✓	✓
Intel Distribution for Python <sup>2</sup>	✓	✓	✓
Intel Math Kernel Library	✓	✓	✓
Intel Data Analytics Acceleration Library	✓ (C++ Only)	✓	✓
Intel® Threading Building Blocks (C++ only)	✓ (C++ Only)	✓	✓
Intel® Integrated Performance Primitives	✓ (C++ Only)	✓	✓
OpenMP	✓	✓	✓
Intel VTune Amplifier XE <sup>3</sup>		✓	✓
Intel® Advisor		✓	✓
Intel® Inspector		✓	✓
Intel® MPI Library <sup>3</sup>			✓
Intel Trace Analyzer and Collector			✓
Intel® Cluster Checker			✓
Rogue Wave IMSL® Library <sup>4</sup>	Bundled and Add-On	Add-On	Add-On
Operating System (Development Environment)	<ul style="list-style-type: none"> <li>Windows (Visual Studio)</li> <li>Linux (GNU)</li> <li>OS X<sup>5</sup> (XCode)</li> </ul>	<ul style="list-style-type: none"> <li>Windows (Visual Studio)</li> <li>Linux (GNU)</li> </ul>	<ul style="list-style-type: none"> <li>Windows (Visual Studio)</li> <li>Linux (GNU)</li> </ul>



<sup>1</sup> Available with a single language (C++ or Fortran) or both languages.

<sup>2</sup> Available on Windows, Linux, and OS X.

<sup>3</sup> Available bundled in a suite or standalone.


<sup>4</sup> Available as an add-on to any Windows Fortran suite or bundled with a version of the Composer Edition.

<sup>5</sup> Available as single language suites on OS X.

# Intel Parallel Studio XE 2017 - Composer Edition

## Performance Awaits Your Application

As processors evolve, it's becoming more and more crucial to both vectorize (use AVX or SIMD instructions) and thread software to realize the full performance potential of the



processor. In some cases, code that is vectorized and threaded can be more than 175X faster than unthreaded/unvectorized code—and about 7X faster than code that is only threaded or vectorized. And that gap is growing with every new processor generation.

## What's New in Intel Parallel Studio XE 2017

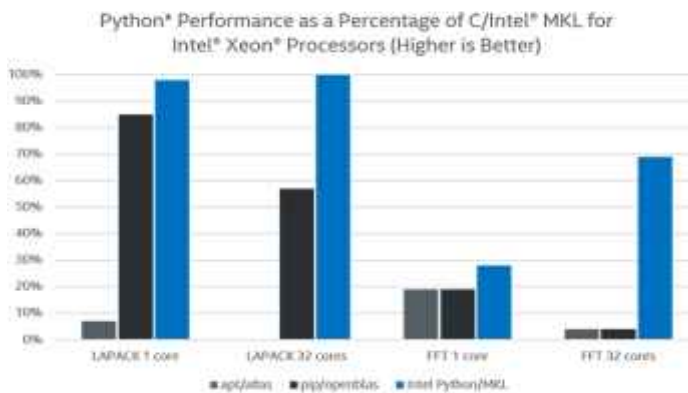
- Scale to next-generation platforms including latest **Intel Xeon Phi™** processor.
- Optimizations for **AVX-512**, high bandwidth memory, and explicit **vectorization** for compiler and analysis tools.
- Faster **deep learning** on Intel architecture. Optimize deep learning model training with **Intel Math Kernel Library neural network functions**. Intel Data Analytics Acceleration Library (**DAAL**) introduces new

deep learning functionality, which is faster than popular open source alternatives.

- Faster Python application performance. **Intel Distribution for Python** provides easy access to high-performance Python, powered by native Intel Performance Libraries, in an integrated distribution package. Identify the bottlenecks in Python and mixed native and **Python code using Intel VTune™ Amplifier XE**.
- Quick first look. **Snapshot features** of Intel VTune Amplifier XE and Intel Trace Analyzer and Collector help you quickly assess application performance.

## Intel Distribution for Python

The Intel Distribution for Python is a performance-oriented distribution that delivers robust solutions for faster performance from your Python applications, in an easy-to-access integrated package. It integrates multiple techniques to power high-performance Python such as linking widely used Python packages (e.g., NumPy, SciPy, scikit-learn) with Intel Performance Libraries including **Intel Math Kernel Library** (Intel MKL), enhanced threading with Intel Threading Building Blocks (Intel TBB), and easy access to the conda environment, numba, Cython and advanced data analytics with pyDAAL.





# Intel Parallel Studio XE 2017 - Professional and Cluster Edition

## Professional Edition

Includes everything in Composer Edition plus

- Advanced performance profiler (Intel VTune Amplifier XE) to tune application performance, scalability and memory access.




- Intel Advisor, a vectorization optimization and threading advisor tool for C, C++, C#, and Fortran applications, gives you a performance estimate before you invest significant effort in implementation - implement only the options that have a high return on investment. Intel Advisor is a breakthrough vectorization and threading design tool for software architects.

Function Call Sites and Loops	Vector Issues	Self Time	Total Time	Type	GFLOPS	Why No Vectorization	Vectorized Loops						
loop in 5242 at loop0011172	1 Possib...	0.0081	0.0081	Vectorized Do...	0.1911	0.10	1 vectored... AVX2	0.7%	6.3e+0	999	42-1	Dis	
loop in 5242 at loop0011172	1 Possib...	2.9704	2.9704	Scalar	0.0671	0.0833	vectorize...			999			Dis
loop in 5242 at loop0011172	1 Possib...	0.0904	0.0904	Vectorized Do...	4.3338	0.1250	AVX2			8	1	42	FMA
loop in 5242 at loop0011172	1 Possib...	0.0201	0.0201	Scalar	0		vectorize...			1			Dis
loop in 5210 at loop0011140	2 Possib...	2.3058	2.3058	Scalar	0.1521	0.0652	vectorize...						FMA
loop in 5128 at loop0014447	1 Assu...	1.0604	1.0604	Scalar	0.3701	0.1667	vector de...						
loop in 5343 at loop0012300	1 Assu...	1.0204	1.0204	Scalar			vector de...						
loop in 5333 at loop0012301	1 Possib...	0.8804	0.8804	Vectorized Do...	2.2748	0.1250	AVX2	0.6%	2.7e+0	8	42-1	FMA	

- Memory and thread debugger (Intel Inspector) to efficiently find memory errors corruption, data races and more in C, C++, and Fortran applications without special recompiles.

## Cluster Edition

Includes everything in the Professional Edition plus



- Accelerate applications performance on Intel architecture-based clusters with multiple fabric flexibility.
- Profile MPI applications to quickly finding bottlenecks, achieving high performance for parallel cluster applications with Intel MPI, a performance-optimized MPI library supporting the MPI 3.1 standard.
- Verify that cluster components continue working together throughout the cluster life cycle (Intel Cluster Checker).

## Intel® Performance Snapshot

The advertisement features a runner in the background and three main sections: 'Application Performance Snapshot', 'Storage Performance Snapshot', and 'MPI Performance Snapshot'. Each section asks a question and provides a list of metrics to be analyzed. For example, the Application Performance Snapshot asks 'Will your application have better vectorization & threading?' and lists metrics like CPU Utilization, Memory Bandwidth, and MPI Utilization. The Storage Performance Snapshot asks 'Will faster storage improve performance?' and lists metrics like Storage Utilization, CPU Utilization, Memory Utilization, and Network Utilization. The MPI Performance Snapshot asks 'Will optimizing settings & tuning improve cluster performance?' and lists metrics like MPI in Computer Time, MPI Bandwidth, CPU Rate & Memory Usage, and FLOPS.

# INTEL® CODE MODERNIZATION WORKSHOPS FALL 2016

## Performance Awaits Your Application

The workshop is intended to help software developers create, modernize and tune parallel code to run with highest performance on Intel's latest processors, including Intel Xeon and Intel Xeon Phi (aka Knights Landing, KNL).

Besides learning from experts during in depth technical sessions and real world case studies, you will also meet and network with your peers from the industry and with software development experts and leaders from Intel.

The new Intel Xeon Phi (aka Knights Landing) was officially launched end of June and we are now able to present and discuss in detail how to develop,

## DAY 1

## AGENDA

<b>08:30</b>	<b>10:00</b>	Registration and Breakfast, <b>Welcome – Day 1</b> <b>DELIVERING LEADING PERFORMANCE IN COMPUTING, ANALYTICS AND MACHINE LEARNING – ALL ABOUT INTEL® XEON™ &amp; XEON PHI™</b> Learn about the parallel architecture, technical advances and features of the latest and future Intel processors, especially Intel® Xeon™ and Intel® Xeon Phi™ (aka Knights Landing, KNL).
<b>10:15</b>	<b>10:45</b>	<b>TRENDS IN EFFICIENT PARALLEL COMPUTING AND PERFORMANCE</b> With all the advances in massively parallel and multi-core computing with CPUs and accelerators it is often overlooked whether the computational work is being done in an efficient manner. This presentation shows the well-known laws of parallel performance from the perspective of a system builder. It also covers through the use of real case studies, examples of how to program for energy efficient parallel application performance.
<b>10:45</b>	<b>11:15</b>	<b>VECTORIZATION – NEW APPROACHES FOR PARALLELISM AT CORE LEVEL (SIMD)</b> Vectorization is one of the critical elements to maximize parallel performance. In this session we will show how to get started with vectorization and avoid common pitfalls. We will review cases where automatic vectorization fails, providing tips and best known methods for effective vectorization. This session is also illustrated with a couple of real-world case examples by using Intel® Parallel Studio XE suite.
<b>11:45</b>	<b>12:30</b>	<b>CODING HIGH PERFORMANCE PYTHON® FOR DATA-INTENSIVE ALGORITHMS</b> This talk will introduce the recently released Intel® Distribution for Python which delivers high performance acceleration for scientific computing, data analytics, and machine learning. Learn how NumPy/SciPy can now leverage the full performance potential of parallel CPU architecture by linking performance libraries like Intel® MKL (Math Kernel Library), Intel® MPI (Message Passing Interface), Intel® TBB (Threading Building Blocks) and Intel® DAAL (Data Analytics Acceleration Library).
<b>12:30</b>	<b>13:15</b>	Lunch
<b>13:15</b>	<b>14:15</b>	<b>PROFILING AND TUNING PYTHON® CODE FOR MAXIMUM PERFORMANCE</b> Performance bottlenecks and underutilization of compute resources result in higher cost per performance unit and watt. In this session you will learn how to identify and eliminate hotspots and realize performance potential in Python code using Intel® VTune™ Amplifier XE.
<b>14:15</b>	<b>15:00</b>	<b>ENABLING MACHINE LEARNING ON INTEL ARCHITECTURE</b> In the first part of this session we will discuss the overall ecosystem of Machine Learning algorithms, challenges and opportunities brought by the Big Data era, and how they are addressed by Intel software solutions. In the second part of the presentation we will review capabilities of Intel libraries that address Big Data problems. The rest of the session devotes to deep overview of Intel® Data Analytics Acceleration Library, software solution that provides the building blocks for all stages of the data processing, from data acquisition till modelling and scoring. Architectural aspects, content of the capabilities, and performance of the library on Intel's highly parallel CPUs will be discussed in this session also.
<b>15:30</b>	<b>16:15</b>	<b>ENABLING DEEP LEARNING ON INTEL ARCHITECTURE</b> The first part of the session will discuss the capabilities available in Intel SW to enable high performance in Deep Learning applications. The rest of the session will review the features of two libraries, Intel® Math Kernel Library and Intel® Data Analytics Acceleration Library for support of Deep Learning, the performance of popular frameworks such as Caffe enabled with Intel SW, and explain similarities and differences of two solutions.
<b>16:15</b>	<b>17:00</b>	<b>Networking with drinks &amp; finger food</b>
<b>17:30</b>	<b>20:00</b>	



# AMSTERDAM - SEPTEMBER 20+21

## PRAGUE - OCTOBER 4+5

## MUNICH - OCTOBER 25+26

REGISTER NOW AT [WWW.INTELDEVCONFERENCE.COM](http://WWW.INTELDEVCONFERENCE.COM)

### AGENDA

### DAY2

08:30	09:30	Registration and Breakfast
09:30	10:15	<b>WHATS NEW IN IPS XE 2017?</b> Learn about some of the new features of Intel Parallel Studio XE 2017. We'll also take a look at some upcoming technologies, such as the roofline model support in Advisor
10:15	11:00	<b>CASE STUDY – Solving the N-Body Problem on INTEL® XEON PHI™ (KNL)</b> Porting and optimizing of an n-body algorithm to the newest generation of Intel® Xeon Phi™ processors. Tuning techniques such as scalar optimizations, vectorization with structures of arrays and memory optimizations will be explained and their effect demonstrated.
11:30	12:30	<b>TUTORIAL - PROFILING AND IMPROVING ADVANCED VECTORIZED CODE</b> Using the latest enhancements to Intel Vectorization Advisor, this session shows how to detect vectorisation issues and strategies to overcome them. Included in this session is a practical hands-on example using DL_MESO Lattice Boltzmann code from Daresbury Labs.
12:30	13:30	Lunch Break
13:30	14:15	<b>CASE STUDY – PARALLELIZING THE BLACK-SCHOLES MODEL</b> In this session, we will demonstrate how to go from serial code to optimized parallel code with a classic algorithm (the Black-Scholes option pricing model) using open standards, libraries and tools, to finally efficiently run on multi-core and many-core CPUs, including the latest Intel® Xeon Phi™ (KNL).
14:15	15:00	<b>TUTORIAL: APPLYING LIBRARIES AND TOOLS TO REAL WORLD EXAMPLE</b> In this tutorial we will demonstrate capabilities of Intel SW for development of the recommendation system, the typical example of the application that requires distributed processing of the data. The tutorial will show how to use Python API of Intel® Data Analytics Acceleration Library, its capabilities together with Intel® MPI communication technology to train the respective model and produce recommendation for specific users. The tutorial will be completed with the analysis of the performance and scalability of the Intel SW based application.
15:30	16:30	<b>PREPARING FOR THE LATEST GENERATION OF INTEL ARCHITECTURE BEFORE HAVING ACCESS TO HARDWARE</b> In this session we'll show how to prepare for the latest generation of Intel architecture before you have access to the real hardware. We will give practical examples of how to prepare for the Intel® Xeon Phi™ Knights Landing to make sure that your code is 'KNL ready'.
16:30	17:00	<b>CASE STUDY - IMPORTANCE OF CURRENT SOFTWARE TOOL CHAINS AND LIBRARIES</b> This session will show the impact on performance brought by up-to-date compilers and high performance libraries. Multiple examples of common software and custom code will be used to demonstrate how speed-ups can be achieved on various target platforms by using current software.

optimize and modernization code using the full spectrum of its features and capabilities, whether you are programming in C/C++, Python or Fortran.

Applications range from general High Performance Computing, Numerical Simulation to Machine Learning, Analytics, and Artificial Intelligence. We'll cover case studies from the area of Life Sciences, Financial Services, Manufacturing and Energy. The Agenda may be slightly different at each location.

# Pricelist Fall 2016

All prices are in Euro and subject to German VAT (if applicable) and exchange rate fluctuations.

Order online at [shop.hocomputer.de](http://shop.hocomputer.de) and save 5% on all prices listed below!

<b>Intel Parallel Studio XE</b>	<b>New, commercial Windows / Linux</b>	<b>UPGRADE, com. Windows / Linux</b>	<b>New, academic Windows/Linux</b>	<b>UPGRADE, acad. Windows / Linux</b>
<b>Prof. Ed Fort. &amp; C++</b> 1NU	2299 / 2299	1849 / 1849	1149 / 1149	949 / 949
2 User Floating	11499 / 11499	9199 / 9199	5479 / 5749	4599 / 4599
5 User Floating	22999 / 22999	18399 / 18399	11499 / 11499	9199 / 9199
<b>Prof Ed Fortran</b> 1NU	1899 / 1899	1549 / 1549	949 / 949	799 / 799
2 User Floating	9499 / 9499	7599 / 7599	4749 / 4749	3799 / 3799
5 User Floating	18999 / 18999	15199 / 15199	9499 / 9499	7599 / 7599
<b>Prof. Ed. C++</b> 1NU	1599 / 1599	1299 / 1299	799 / 799	649 / 649
2 User Floating	7999 / 7999	6399 / 6399	3999 / 3999	3199 / 3199
5 User Floating	15999 / 15999	12799 / 12799	7999 / 7999	6399 / 3699
<b>Cluster Edition</b> 1NU	2949 / 2949	2399 / 2399	1499 / 1499	1049 / 1049
2 User Floating	14749 / 14749	11799 / 11799	7399 / 7399	3599 / 3599
5 User Floating	29499 / 29499	23599 / 23599	14749 / 14749	7699 / 7699
<b>Comp.Ed Fort. &amp; C++</b> 1NU	1199 / 1449		499 / 599	
2 User Floating	4199 / 5099	- n / a -	1699 / 2049	- n / a -
5 User Floating	8999 / 10899		3599 / 4349	
<b>Comp. Ed. Fortran</b> 1NU	849 / 999		399 / 449	
2 User Floating	2999 / 3499	- n / a -	1349 / 1599	- n / a -
5 User Floating	6399 / 7499		2899 / 3399	
<b>Composer Ed. C++</b> 1NU	699 / 699		249 / 249	
2 User Floating	2449 / 2449	- n / a -	849 / 849	- n / a -
5 User Floating	5249 / 5249		1849 / 1849	

**1NU: single user** = 1 named user (system locked), may be used by one fixed human person.  
**2/5 User Floating** = 2/5 can use the software simultaneously (uses FlexLM).  
 Academic Licences are only available for Degree Granting Institutions. All licences are delivered by Email / Download. Our AGB are apply.