

Published on *CERN openlab* (<http://test-static-05.web.cern.ch>)

[Home](#) > [Announcing the CERN openlab / Intel - Winter 2012 Computer Architecture and Performance Tuning Workshop](#)

Announcing the CERN openlab / Intel - Winter 2012 Computer Architecture and Performance Tuning Workshop [1]

Wednesday, 15 February, 2012



The aim of the lectures and exercises contained in this workshop is to give the attendees a practical introduction to performance optimization and monitoring on Linux, based on a good understanding of modern computer architectures. While the focus will be on C and C++, programmers of other languages will also benefit from the performance monitoring and hardware related classes. Non-expert users should feel more than welcome, as the course will be an excellent opportunity to improve their knowledge related to computer architecture and performance optimization. As a novelty, an interactive mini-lab on floating point computing will be introduced.

Performance optimization tools will be used during the course, enabling the participants to discover how the structure of the code influences its performance. An in-depth compiler study is foreseen, where [Jeff Arnold](#) [3] from the Intel compiler team will explain the pitfalls of floating point computation, and the key highlights and differences of GCC and the Intel compiler. The participants will be given the task of correlating performance figures with certain programming decisions. In addition, the participants will understand the limits of performance tuning and the ways to establish at which point inside those limits their workload is placed. The exercises will

be supported by demonstrating real world problems in production environments.

? Event highlights:

- Day 1, AM: Lectures
 - Lecture 1: Scalability in software and hardware
 - Lecture 2: Systematic benchmarking
 - Lecture 3: Compiler overview
 - Lecture 4: Understanding performance tuning

- Day 1, PM: Hands-on labs

- Performance tuning exercises
- C++ optimization
- NUMA memory systems
- Working with GNU and Intel compilers
- Floating point computation mini-lab
- Performance optimization exercises
- Compiler exercises

- Day 2, AM: Lectures

< >Parallelism through vectorization

Day 2, PM: Hands-on labs

? Pre-requisites:

- Good understanding of the C or C++ programming language
- Knowledge of Linux
- Basic understanding of modern computer architecture and compilers

The workshop is organized by CERN openlab and Intel for users affiliated with CERN, **free of charge at this time**. Registrations are based on a first-come first-served basis ? 40 seats are available for the lectures and the hands-on exercises ? please indicate which are of interest to you during registration. Please note that for the hands-on labs **you need to bring a portable computer** with you to access the exercise machines. Go to [EDH](#) ^[4] in order to register. All registrations will be confirmed by e-mail ? a registration in EDH is required and **does not yet guarantee a place in the workshop**.

- [Visit Us](#)
- [RSS Feeds](#)

DISCLAIMER: This Web page contains pointers to material related to the management of CERN openlab in the Information Technology Department at the European Organization for

Nuclear Research (CERN). Their use and distribution are regulated by the CERN copyright notice.



Source URL: <http://test-static-05.web.cern.ch/news/announcing-cern-openlab-intel-winter-2012-computer-architecture-and-performance-tuning-workshop>

Links

- [1] <http://test-static-05.web.cern.ch/news/announcing-cern-openlab-intel-winter-2012-computer-architecture-and-performance-tuning-workshop>
- [2] http://test-static-05.web.cern.ch/sites/test-static-05.web.cern.ch/files/styles/large/public/news/images/FP_workshop_poster_small.jpg?itok=Vx-P_w-x
- [3] https://openlab-mu-internal.web.cern.ch/openlab-mu-internal/00_News/News_pages/2008/08-11_Jeff/08-09_Jeff.htm
- [4] <https://edh.cern.ch/Document/Personnel/TRN/new?course=083TUN01>