



NCAR/UCAR
Intel Developer Tools Training
February 9th and 11th

February 9th, 9 – 11:30 am CST

oneAPI Overview: The oneAPI specification extends existing developer programming models to enable a diverse set of hardware through language, a set of library APIs, and a low-level hardware interface to support cross-architecture programming. To promote compatibility and enable developer productivity and innovation, the oneAPI specification builds upon industry standards and provides an open, cross-platform developer stack. Introduction to Intel oneAPI DevCloud

Intel Fortran and C++ Update:

A Tradition of Trusted Application Performance

The Intel® Fortran Compiler is built on a long history of generating optimized code that supports industry standards while taking advantage of built-in technology for Intel® Xeon® Scalable processors and Intel® Core™ processors. Staying aligned with Intel's evolving and diverse architectures, the compiler now supports GPUs.

Standards: The Path Forward

There are two versions of this compiler.

Intel Fortran Compiler: provides CPU and GPU offload support

Intel Fortran Compiler Classic: provides continuity with existing CPU-focused workflows

Both versions integrate seamlessly with popular third-party compilers, development environments, and operating systems.

Features

- Improves development productivity by targeting CPUs and GPUs through single-source code while permitting custom tuning
- Supports broad Fortran language standards
- Incorporates industry standards support for OpenMP* 4.5, and initial OpenMP 5.0 and 5.1 for GPU Offload
- Uses well-proven LLVM compiler technology and Intel's history of compiler leadership
- Takes advantage of multicore, Single Instruction Multiple Data (SIMD) vectorization, and multiprocessor systems with OpenMP, automatic parallelism, and coarrays
- Optimizes code with an automatic processor dispatch feature

Intel VTune Profiler: Intel® VTune™ Profiler optimizes application performance, system performance, and system configuration for HPC, cloud, IoT, media, storage, and more.

- CPU, GPU, and FPGA: Tune the entire application's performance—not just the accelerated portion.
- Multilingual: Profile Data Parallel C++ (DPC++), C, C++, C#, Fortran*, OpenCL™, Python*, Google Go* programming language, Java*, Assembly, or any combination.
- System or Application: Get coarse-grained system data for an extended period or detailed results mapped to source code.
- Power: Optimize performance while avoiding power- and thermal-related throttling.



NCAR/UCAR
Intel Developer Tools Training
February 9th and 11th

Intel MPI Library:

Intel® MPI Library is a multifabric message-passing library that implements the open-source MPICH specification. Use the library to create, maintain, and test advanced, complex applications that perform better on high-performance computing (HPC) clusters based on Intel® processors.

- Develop applications that can run on multiple cluster interconnects that you choose at run time.
- Quickly deliver maximum end-user performance without having to change the software or operating environment.
- Achieve the best latency, bandwidth, and scalability through automatic tuning for the latest Intel® platforms.
- Reduce the time to market by linking to one library and deploying on the latest optimized fabrics.

oneAPI MKL - Intel Math Kernel Library (MKL): accelerate math procession routines, increase application performance, and reduce development time. MKL includes: Linear algebra, fast flourier Transforms (FFT) Vector statistics and data fitting, vector math and miscellaneous solvers

February 11th 9 – 11:30 am CST

Intel oneAPI AI Toolkit/Intel Distribution for Python:

Deliver high-performance, deep-learning training on Intel® XPU's and integrate fast inference into your AI development workflow with Intel®-optimized, deep-learning frameworks for TensorFlow* and PyTorch*, pretrained models, and low-precision tools.

Achieve drop-in acceleration for data preprocessing and machine-learning workflows with compute-intensive Python* packages, Modin*, scikit-learn*, and XGBoost, optimized for Intel.

Gain direct access to analytics and AI optimizations from Intel to ensure that your software works together seamlessly.