Intel® oneAPI Toolkits Training at LRZ

Intel Tools Co-Design and Customer Engineering
Data Center SW, Xeon and AI
Intel Corporation

April 2026

Agenda

01	Opening
02	Intel oneAPI Toolkits: Unified Performance Across CPUs and GPUs
03	Why Intel Software Tools Still Matter for CPU and the latest features for Intel Xeon 6
04	Training Overview & Call to Actions (the latest version of oneAPI, performance bug, priority support)

Welcome to Intel oneAPI Tools Training

- Thank you to LRZ for hosting and to all attendees
- Team Introduction
- Two-day workshop: In-depth and hands-on practical training on Intel oneAPI tools

Intel - LRZ Partnership

Platform Adoption for HPC & AI | oneAPI Software Integration

Intel Platform Adoption

- SuperMUC-NG Phase 2 expands AI capabilities through Intel & Lenovo partnership
- Powered by Intel Xeon Scalable (Sapphire Rapids) + Intel Data Center GPU Max Series
- DAOS storage with 3rd Gen Intel Xeon Scalable processors
- Built on Intel's XPU portfolio and unified oneAPI software stack

Intel® oneAPI toolkits

- Intel oneAPI Base & HPC Toolkits integrated across LRZ systems via Intel-toolkit modules
- Intel LLVM compilers are defaults on SuperMUC-NG Phase 2
- GPU workloads require oneAPI models: OpenMP offload, SYCL, or OpenCL
- Joint LRZ–Intel workshops on oneAPI for HPC & AI targeting Intel platform

Intel® oneAPI Toolkit

Flexible, Comprehensive, Developer Software Stack – Powered by Intel, Optimized for Intel CPUs and GPUs

Intel Software Tools

Build and Run

- Intel Compilers (C/C++, SYCL, Fortran, OpenMP)
- Triton & SYCL*TLA
- oneDPL (DPC++ Library)
- oneMKL (Math Kernel Library)
- oneDAL (Data Analytics Library)
- oneTBB (Threading Library)
- IPP & Crypto Library (Performance Primitives)
- oneDNN (Deep Neural Network Library)

Analyze & Debug

- Intel® VTune™ Profiler
- Intel® Advisor
- Intel® Distribution for GDB*

Scale

- Intel® MPI Library
- Intel® SHMEM Library
- oneCCL (Collective Communications Library)

These tools integrate with existing orchestration and scheduling environments through standard runtimes, MPI, and communication libraries fitting naturally into production Xeon-targeted workflows.

HPC and AI Applications Optimized on Intel Platforms

Why Intel® tools matter on Intel® hardware

*If Intel® tools aren't part of your build, you may be leaving Intel® Xeon® performance on the table — our SPEChpc results show **up to 9.4%** from software tuning alone.*

THE TOOLS UPLIFT

BASE

22.4

generic build

▶

PEAK

24.5

Intel® oneAPI build

Same Intel® Xeon® silicon. Same workload.
What changed: the toolchain.

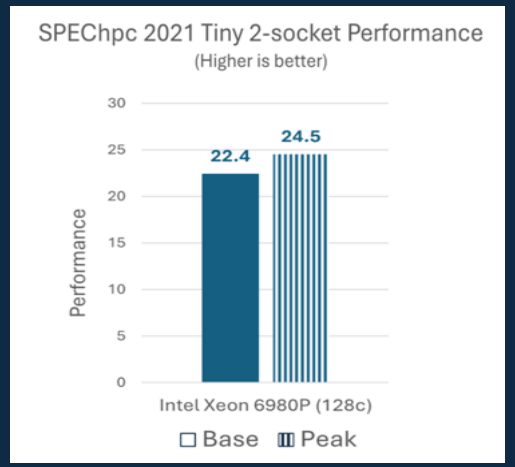

WHAT INTEL® TOOLS UNLOCK

Intel® oneAPI C/C++ (ICX) & Fortran (IFX)
AVX-512 vectorization · Intel® APX register doubling

Intel® oneMKL
AVX-512-tuned BLAS/LAPACK — drop-in faster than OpenBLAS

Intel® oneDNN
Intel® AMX kernels for FP16 · BF16 · INT8 AI workloads

Intel® MPI Library
Low-latency multi-node scaling on Xeon clusters

Intel® VTune™ Profiler
Hardware-counter analysis built for Intel® microarchitecture

Reference: Intel® Xeon® CPUs and High-Performance Computing Technologies Meet the Demands of HPC Workloads

SPEChpc™ 2021 Tiny, 2-socket peak/base. Intel® Xeon® 6980P 128-core with Intel® oneAPI Compiler 2025.2 + Intel® MPI Library 2021.16, Ubuntu 22.04.5 LTS. SPEC® and SPEChpc® are trademarks of the Standard Performance Evaluation Corporation. Full submission details on SPEC.org. Results may vary.

INTEL® XEON® 6980P — PERFORMANCE BENCHMARK VS. AMD EPYC 9755 (128-CORE)

BENCHMARK RESULTS

1.15×
faster on WRF
CONUS 2.5km · 6-hour forecast

1.33×
faster on MPAS-A
720 timesteps · 3-day forecast

Predict Faster. Predict Smarter. Real-Time Weather Simulation Powered by Intel® Xeon® Processors

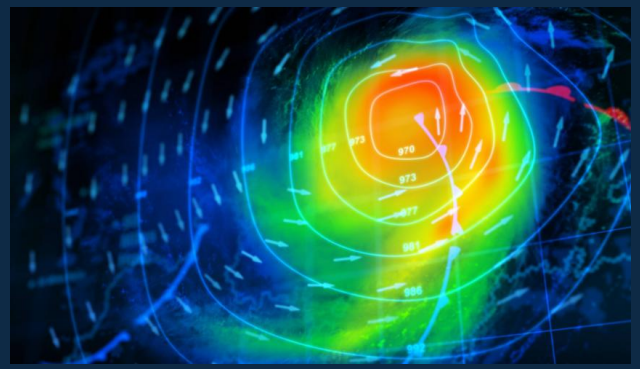

ENABLING INTEL® SOFTWARE

Intel® Fortran Compiler
v2025.0

Auto-vectorization and loop optimization tuned for WRF's compute-intensive atmospheric physics kernels

Intel® MPI Library
v2021.14 / v2021.15

High-efficiency message passing for multi-node MPAS-A runs — minimizes inter-core latency across 256 total cores

Intel® oneAPI HPC Toolkit
2025 release

Accelerate your workloads with oneAPI tools on Intel® platforms

Source: Intel internal testing, Dec 2024 – Mar 2025. Intel Xeon 6980P (128-core, 500W TDP) vs AMD EPYC 9755 (128-core, 500W TDP). Results may vary. Intel technologies may require enabled hardware, software, or service activation.

Intel Tools Advantage for Intel Platforms

Intel compilers, libraries, and analyzers uniquely designed to unlock full performance on Intel hardware.

- **Universal Coverage**: Supports all Intel platforms (Xeon, Core Ultra, Arc, and future GPUs)
- **Performance Advantage:**
  - Auto-optimizing code to use AMX, AVX-512, and AVX512-VNNI instructions
  - Harnessing on-die accelerators (AMX, QAT, DSA, IAA) with tuned libraries
  - Using power and performance analyzers (VTune, SoCWatch) to expose hardware bottlenecks
  - Enabling seamless acceleration, security, and vector optimizations
  - Delivering deeper HW/SW co-optimization
- **Future-Proof Development**: Enables code reuse across CPUs and GPUs, minimizing rewrite effort for new hardware generations
- **Developer-Friendly**: Supports familiar languages (C++, Fortran, Python) and standards to minimize development time, cost and maintenance.
- **Business Value**: Simplifies development, reduces complexity, and accelerates time-to-market—helping drive hardware adoption.

Intel® oneAPI Toolkit for Xeon 6

General Compute

Compilers

Intel® oneAPI DPC++/C++ Compiler

- Cache preloading optimizations for improved data access & latency
- Intel® Advanced Matrix Extensions-FP16 (AMX-FP16) instruction support

Intel® Fortran Compiler

- Backend performance optimizations for Xeon 6 + modern Fortran 2023 feature support

Intel® oneAPI Math Kernel Library (oneMKL)

2.5×
HPCG performance
(with MRDIMM)

- Performance improvements across BLAS, FFT and other routines for Intel Xeon 6
- Optimized for high-core-count Xeon 6 architectures

Intel® MPI Library

- Optimized for up to 128-core Xeon 6 with improved scale-out & scale-up performance
- MPI 5.0 Preview + ABI compatibility with legacy MPI standards via -mpi-abi (Linux)

Accelerate AI Workloads & Applications

AI Frameworks & oneDNN

- PyTorch, TensorFlow & leading framework optimizations upstreamed regularly, enabling speedups on Xeon 6
- oneDNN leverages AMX with FP16 support to accelerate deep learning models

up to 3×
Faster Llama 2 inference
vs. previous generation[2]

up to 1.5×
Better real-time AI inference
vs. AMD EPYC 9755[3]

Intel IPP and Cryptography Primitives Library

- Optimized for Intel Xeon 6, new APIs for floating-point complex addition & multiplication
- Warp Affine functionality with 16-bit integer and 32-bit floating-point data support

Intel® VTune™ Profiler

- New analysis views: hotspots, microarchitecture, memory access, I/O & platform, enabling fast resolution of performance bottlenecks on Xeon 6

Unlock best performance with latest Intel oneAPI 2025.3.1

Intel® oneAPI 2026.0 release

New Platform Support: Clearwater Forest & Wildcat Lake in full production. Early enabling for Crescent Island, Nova Lake, Diamond Rapids

PyTorch 2.12 Upstream Ready: Full oneAPI stack latest oneDNN, oneMKL, oneCCL & Compiler aligned as upstream dependencies

LLMs with oneDNN: Support for weight compression (int4/fp4/mxfp4/fp8), fused SDPA+GQA attention, and RMSNorm fusion

VTune Profiler: Support for native PyTorch model profiling, find performance bottlenecks at the model layer for PyTorch models running on Intel NPUs and GPUs

Distributed Training via oneCCL: aligned with PyTorch upstream for multi-node, multi-GPU scale-out via new IPC support for efficient cross-process GPU communication

SYCL Performance and Productivity: With SYCLBIN fast linking for lower-latency AOT workloads, reduced SYCL Runtime overhead for faster kernel submission, more expressive SYCL graph APIs for flexible parallel programming, and streamlined free function kernel development

Intel® oneAPI Toolkits Changes

Introducing

Two new product configurations - Targeted – April/May 2026

Intel oneAPI Toolkit NDA

Intel oneAPI Toolkit – Internal

Expected Outcome: More individual component early readiness for early platform engagements

Unifying

Combining Base and HPC toolkits into one

Intel® oneAPI Toolkit.

Expected Outcome: Simplifying bundle options and remaining commitments to ISVs and end users for public released platforms.

intel

Software Tools Roadmap 2026+ – Enabling the Future of Client and DC platforms

Platform Enablement & Tuning

- Readiness for next-gen CPU and GPU
- Enabling and performance tuning
- Libraries optimizations
- Runtimes & backend tuning

Programming Models & SW Stack Unification

- **Standard programming languages**
  C++23/26, Fortran 23, SYCL Next, Python
- **Performance libraries integrated with languages**
  Acceleration delivered through libraries without code rewrites
- **Unified compiler and library stack**
  Co-optimized to expose Intel platform capabilities
- **Open standards and upstream alignment**
  OpenMP, MPI, SYCL, MLIR, Triton

Analysis & Debugging Experience

- Expanded AI and Python profiling
- Unified analyzer workflows
- Prepare analyzers and debugger for next-gen Xeon & data center GPU

Next-gen CPU & GPU

Programming Models & Standards

Analyzers & Debugger

Intel Tools Strategy Update in 2026

Accelerate Intel platform wins with Intel Tools

1. **Ready early for Intel platform milestones**
   Align roadmaps to platform milestones to ensure early readiness and competitive advantages
2. **Simpler packaging + predictable releases (faster adoption)**
   Simplify toolkits and cadence to reduce friction and speed customer adoption.

3. **Enablement motion that scales**
   Accelerate targeted customer/ISV adoption through early enablement, awareness, and co-marketing.

Intel tools' primary goal is driving Intel hardware success.

intel

Key points we would like you to retain

1. **Intel® software tools: turning Intel® silicon features into time to market and time to result**
   Compilers, libraries, and analyzers unlock performance, efficiency, and determinism on Intel platforms.
2. **oneAPI tools are Intel platform-first and ready early**
   Aligned to Intel platform milestones, with early standalone tool packages when needed.
3. **Enablement + Support is part of the value**
   Intel Tools Customer Engineering provides enablement, training, and escalation for customers.

Key takeaway:
If performance, power, or scale matters to your workload, bring Intel® tools in from day one.