

Intel[®] DevCloud for oneAPI

ECP OpenMP Virtual Hackathon 2021



intel[®]

oneAPI Industry Initiative

Break the Chains of Proprietary Lock-in

Open to promote community and industry collaboration

Enables code reuse across architectures and vendors

Develop & Deploy Software with Peace of Mind

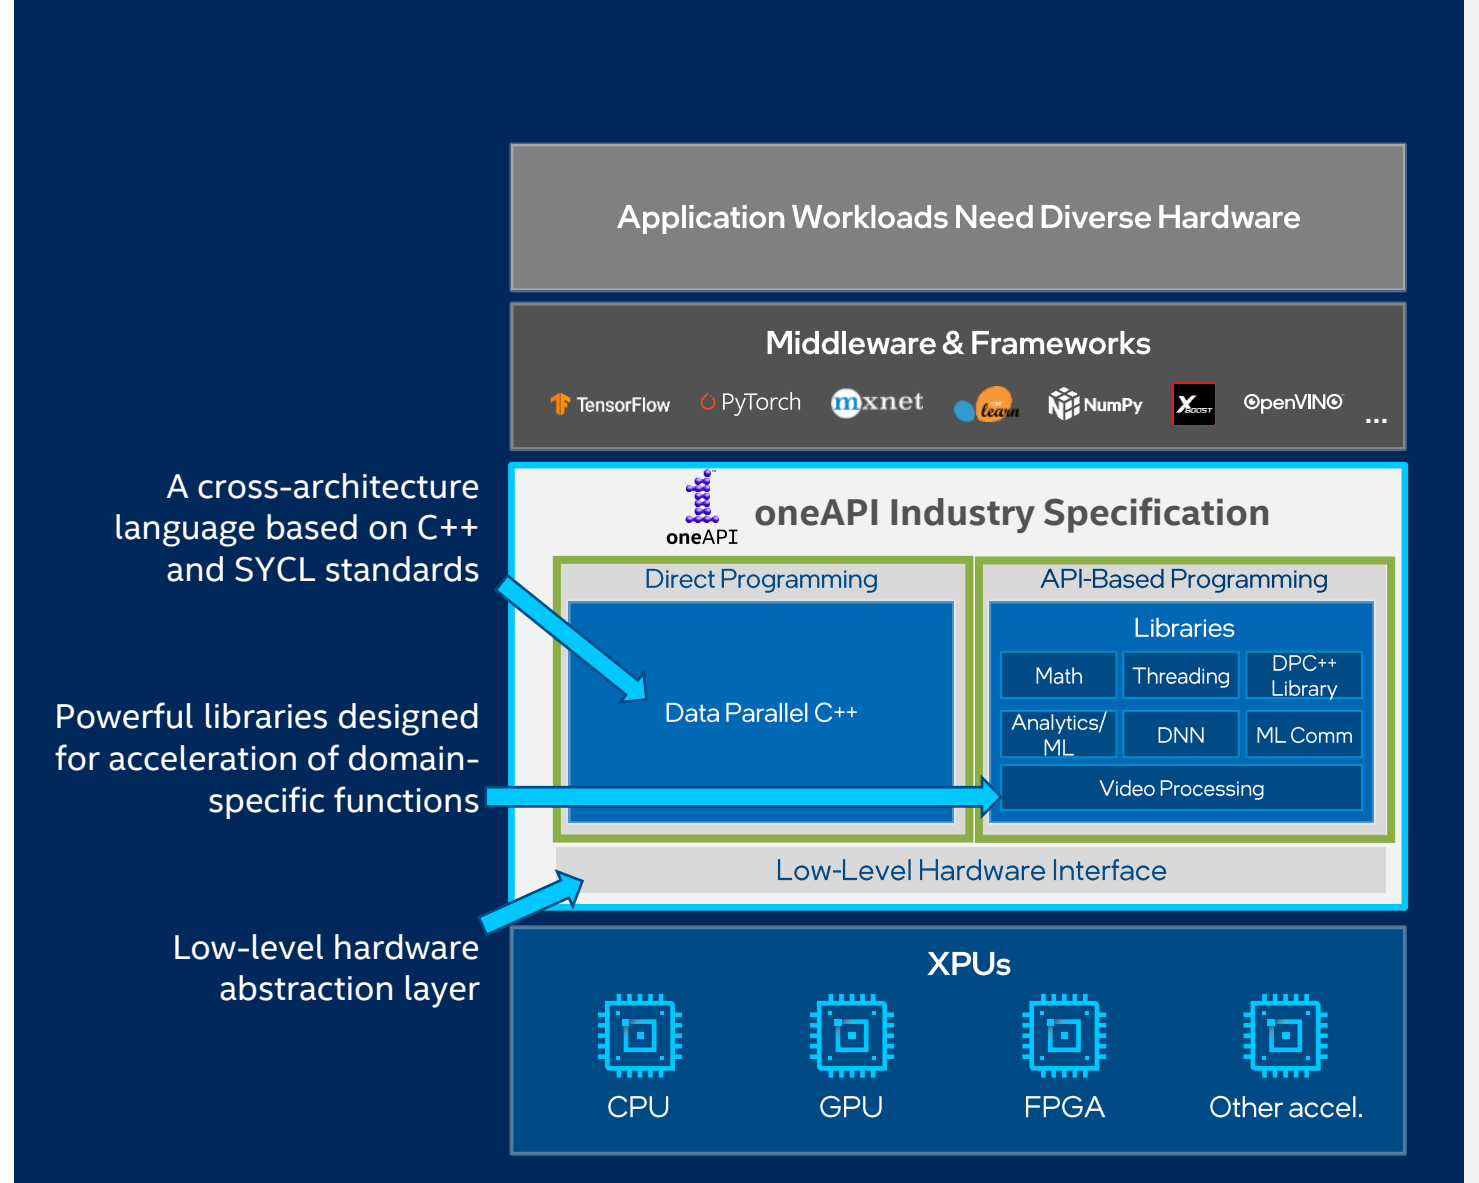
- Open industry standards provide a safe, clear path to the future
- Compatible with existing languages and programming models including C++, Python, SYCL, OpenMP, Fortran, and MPI



The productive, smart path to freedom for accelerated computing from the economic and technical burdens of proprietary programming models

Visit oneapi.com for more details

ECP OpenMP Virtual Hackathon 2021



Intel® oneAPI Toolkits


A complete set of proven developer tools expanded from CPU to XPU



Intel® oneAPI Base Toolkit


Native Code Developers

A core set of high-performance tools for building C++, Data Parallel C++ applications & oneAPI library-based applications




Add-on Domain-specific Toolkits

Specialized Workloads



Intel® oneAPI Tools for HPC

Deliver fast Fortran, OpenMP & MPI applications that scale




Intel® oneAPI Tools for IoT

Build efficient, reliable solutions that run at network's edge



Intel® oneAPI AI Analytics Toolkit

Accelerate machine learning & data science pipelines with optimized DL frameworks & high-performing Python libraries



Intel® oneAPI Rendering Toolkit

Create performant, high-fidelity visualization applications

Toolkit powered by oneAPI

Data Scientists & AI Developers



Intel® Distribution of OpenVINO™ Toolkit

Deploy high performance inference & applications from edge to cloud

Powerful oneAPI Libraries

Realize all the Hardware Value

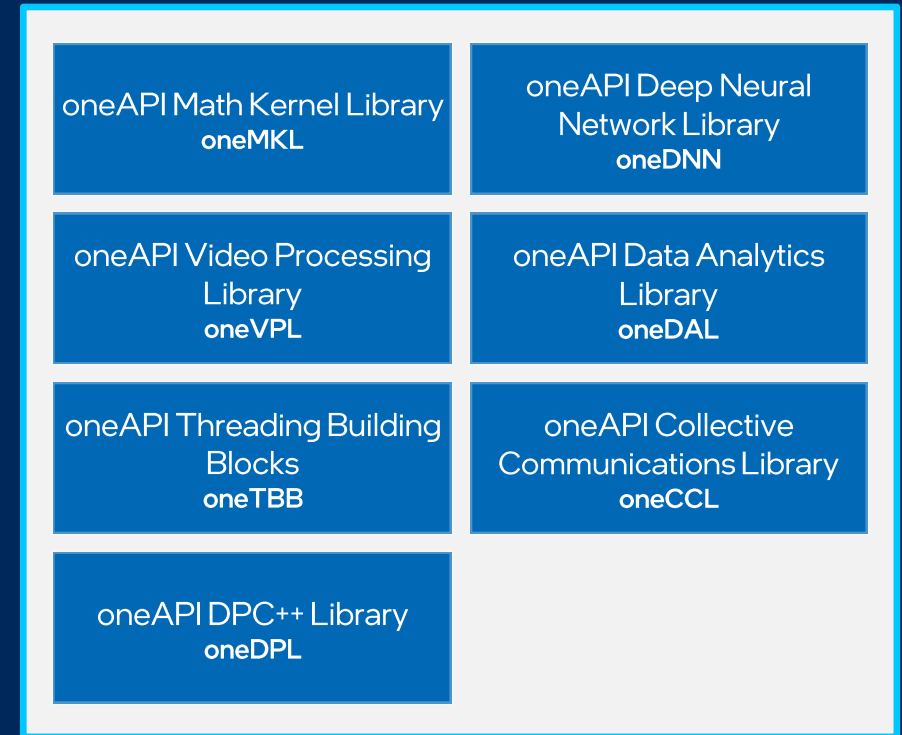
Designed for acceleration of key domain-specific functions

Freedom of Choice

Pre-optimized for each target platform for maximum performance

oneMKL

- Optimized Library for Scientific Compute
- Optimized applications for current and future Intel® CPUs, GPUs, and other accelerators
- Supports DPC++ and OpenMP Offload
- Functions include BLAS, LAPACK, sparse solvers, FFT, RNG, summary statistics, data fitting, vector math
- Seamless upgrade for users of Intel MKL



Intel® oneAPI Tools for HPC

Intel® oneAPI HPC Toolkit

Deliver Fast Applications that Scale

What is it?

A toolkit that adds to the Intel® oneAPI Base Toolkit for building high-performance, scalable parallel code on C++, Fortran, OpenMP & MPI from enterprise to cloud, and HPC to AI applications.

Who needs this product?

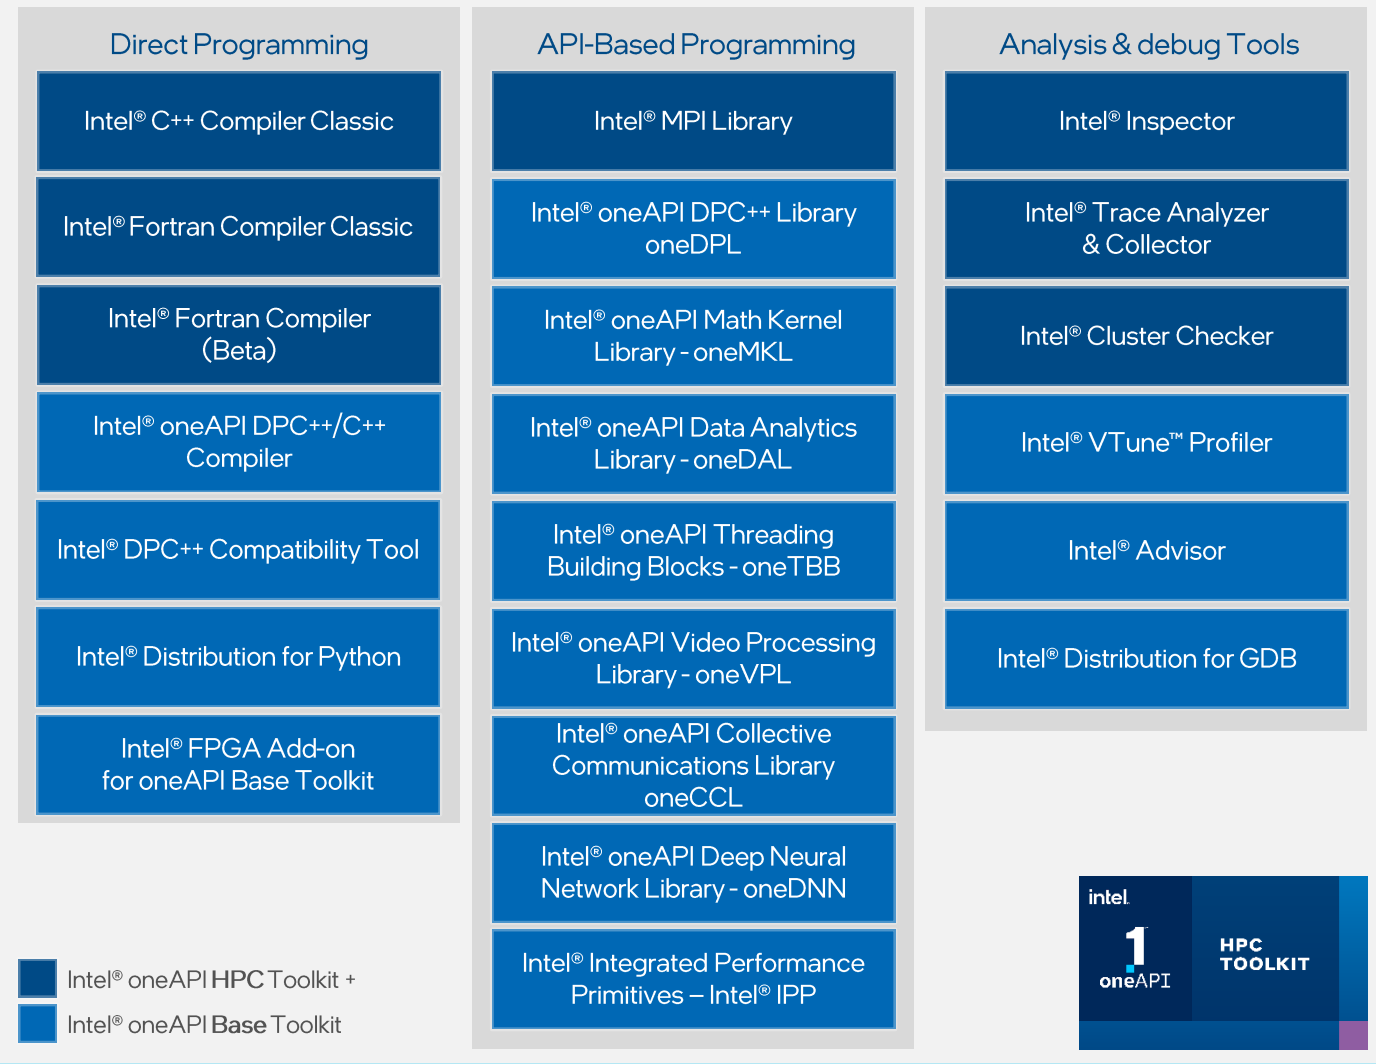
- OEMs/ISVs
- C++, Fortran, OpenMP, MPI Developers

Why is this important?

- Accelerate performance on Intel® Xeon® & Core™ Processors and Intel® Accelerators
- Deliver fast, scalable, reliable parallel code with less effort built on industry standards

Learn More: intel.com/oneAPI-HPCKit

Intel® oneAPI Base & HPC Toolkits



Analysis & Debug Tools

Get More from Diverse Hardware



Design

Intel® Advisor

- Efficiently offload code to GPUs
- Optimize your CPU/GPU code for memory and compute
- Enable **more** vector parallelism and improve efficiency
- Add effective threading to unthreaded applications



Debug

Intel® Distribution for GDB

- Multiple accelerator support with CPU, GPU and FPGA
- Enables deep, system-wide debug of Data Parallel C++ (DPC++), C, C++, and Fortran code



Tune

Intel® VTune™ Profiler

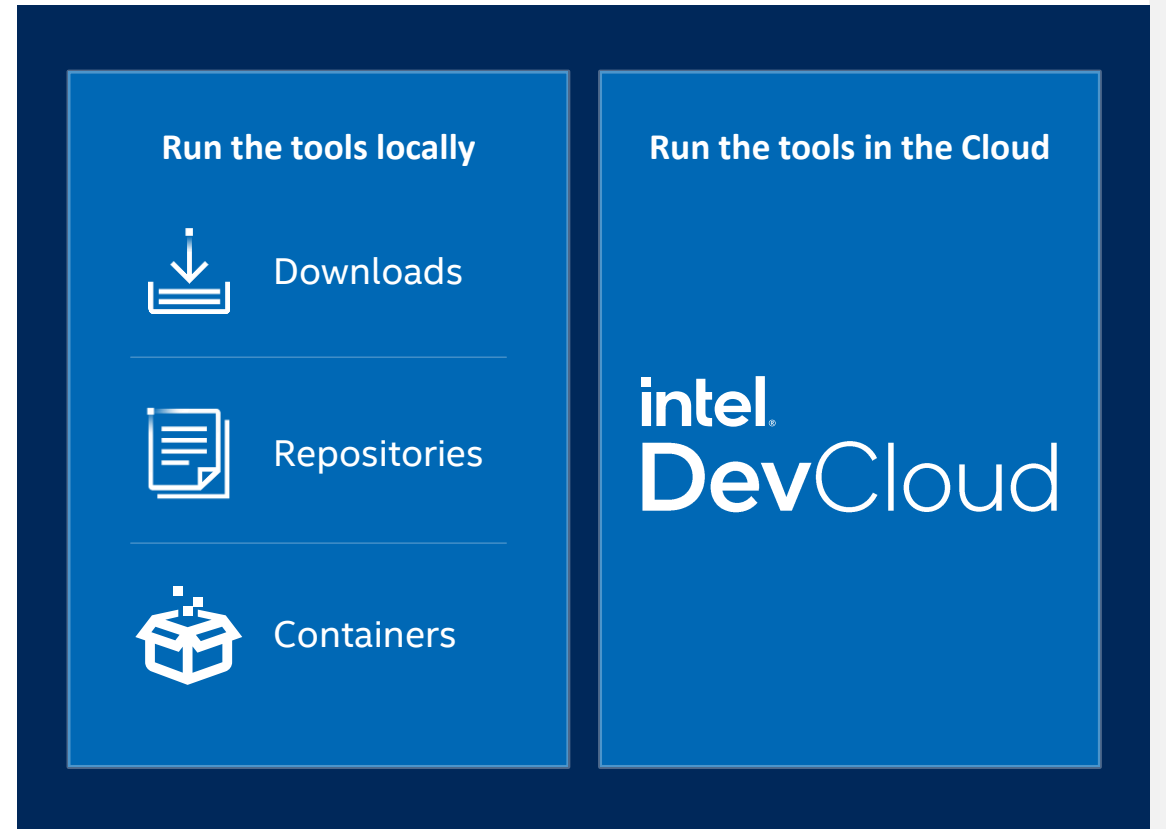
- Tune for GPU, CPU, and FPGA
- Optimize offload performance
- Supports DPC++, C, C++, Fortran, Python, Go, Java or a mix of languages

Intel[®] oneAPI Toolkits Free Availability

Get Started Quickly

Code Samples, Quick-start Guides, Webinars, Training

software.intel.com/oneapi



Run the tools locally

- Downloads
- Repositories
- Containers

Run the tools in the Cloud

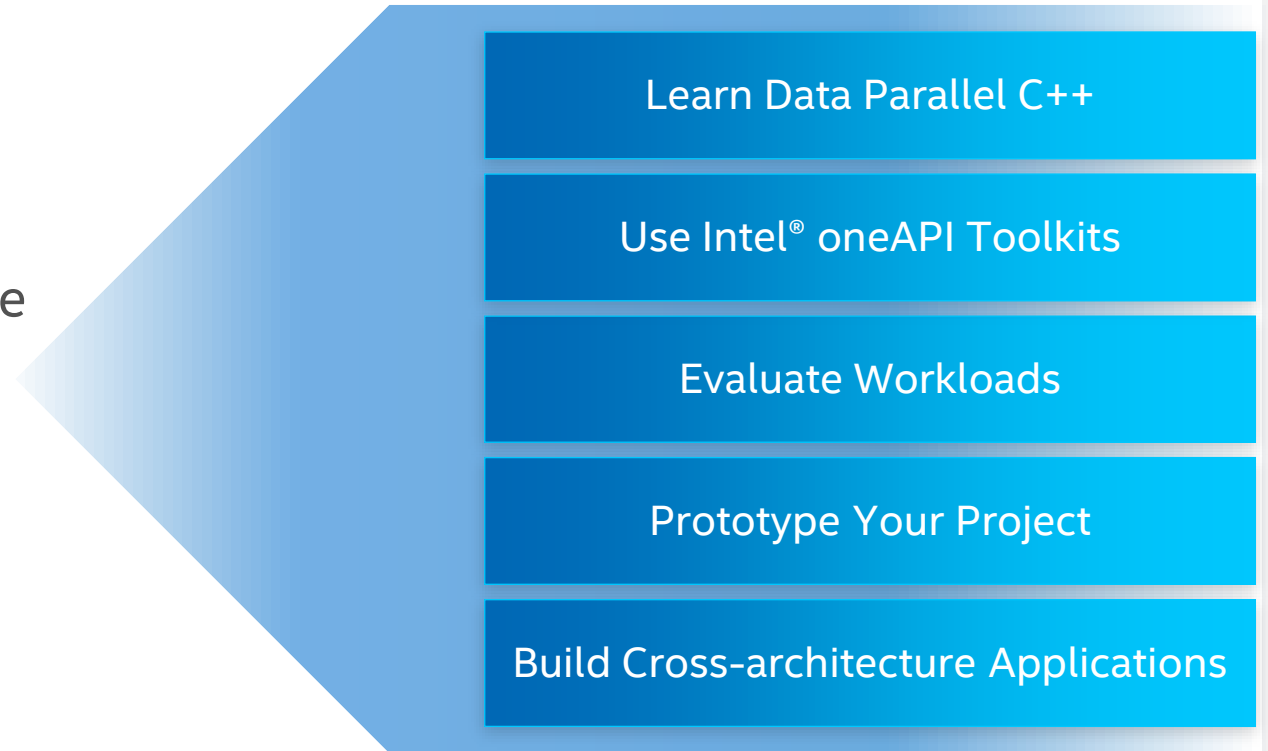
intel
DevCloud

Intel® DevCloud for oneAPI

Free Access, A Fast Way to Start Coding

A development sandbox to develop, test and run workloads across a range of Intel® CPUs, GPUs, and FPGAs using Intel's oneAPI software

For customers focused on data-centric workloads on a variety of Intel® architecture



No Downloads | No Hardware Acquisition | No Installation | No Set-up & Configuration

Get Up & Running in Seconds!

software.intel.com/devcloud/oneapi

Getting Started

Intel® DevCloud for oneAPI

Users sign up to get an account, & can get started using the DevCloud environment immediately

<https://software.intel.com/en-us/devcloud/oneapi>

Learn how to use Intel® oneAPI Toolkits, compilers, performance libraries & tools with latest Intel® processors and accelerators.

What you get

- 120 days of free access (can be extended)
- 200 GB of file storage
- 192 GB RAM

Included Toolkits

- Intel® oneAPI Base Toolkit
- Intel® oneAPI HPC Toolkit
- Intel® AI Analytics Toolkit
- Intel® oneAPI Rendering Toolkit
- Intel® oneAPI DL Framework Developer Toolkit
- Intel® Distribution of OpenVINO™ Toolkit
- + more

Hardware

CPU: Intel® Xeon® Scalable processors

GPU: Intel® Iris® Xe MAX GPU
Intel® Xeon® processors with
Intel® Graphics Technology/GPU

FPGA: Intel® Stratix® 10 FPGAs
Intel® Arria® 10 FPGAs

Featured Tools & Libraries

- Intel® OneAPI DPC++ Compiler & Library
- Intel® C++ & Intel® Fortran Compilers
- Intel® OneAPI Math Kernel Library
- Intel® OneAPI Data Analytics Library
- Intel® OneAPI DNN Library
- Intel® Distribution for Python
- Intel® VTune™ Profiler & Intel® Advisor
- Intel® FPGA Add-on for oneAPI Base Toolkit
- + many more...

intel Developer Zone Software

Search our content library... Support Sign in English

INTEL® DEVCLOUD Workloads

Share

A DEVELOPMENT SANDBOX FOR DATA CENTER TO EDGE WORKLOADS

Develop, test, and run your workloads on a cluster of the latest Intel® hardware and software. With integrated Intel® optimized frameworks, tools, and libraries, you'll have everything you need for your projects.

Try Out a Diverse Collection of Intel® Hardware

Expand your skills and experiment with this state-of-the-art cluster that offers capabilities such as natural language processing and time-series analysis, as well as edge acceleration hardware.

Develop with Intel Software Tools

Jump-start your projects without having to download, configure, or install the latest compilers, performance libraries, and tools from Intel.

oneAPI Compilers Profilers/Analyzers Libraries Debuggers

QUARTUS PRIME

OpenVINO

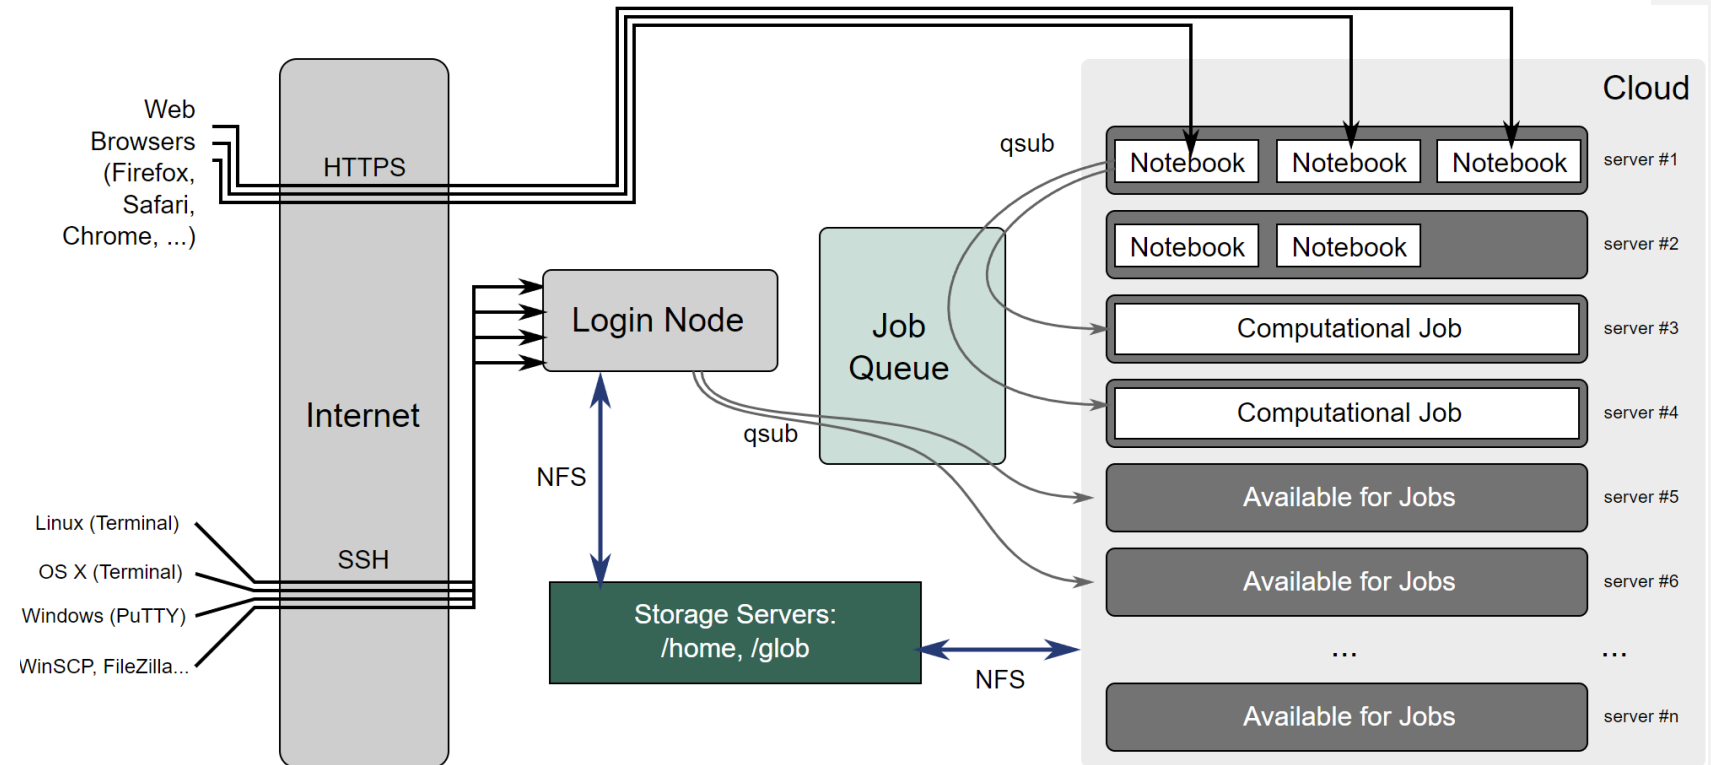
Connect and Use the DevCloud

■ Connection Options

- SSH
 - VSCode
- Jupyter

■ Job Submissions

- TORQUE



intel®

Backup



Intel® oneAPI Base Toolkit

Accelerate Data-centric Workloads

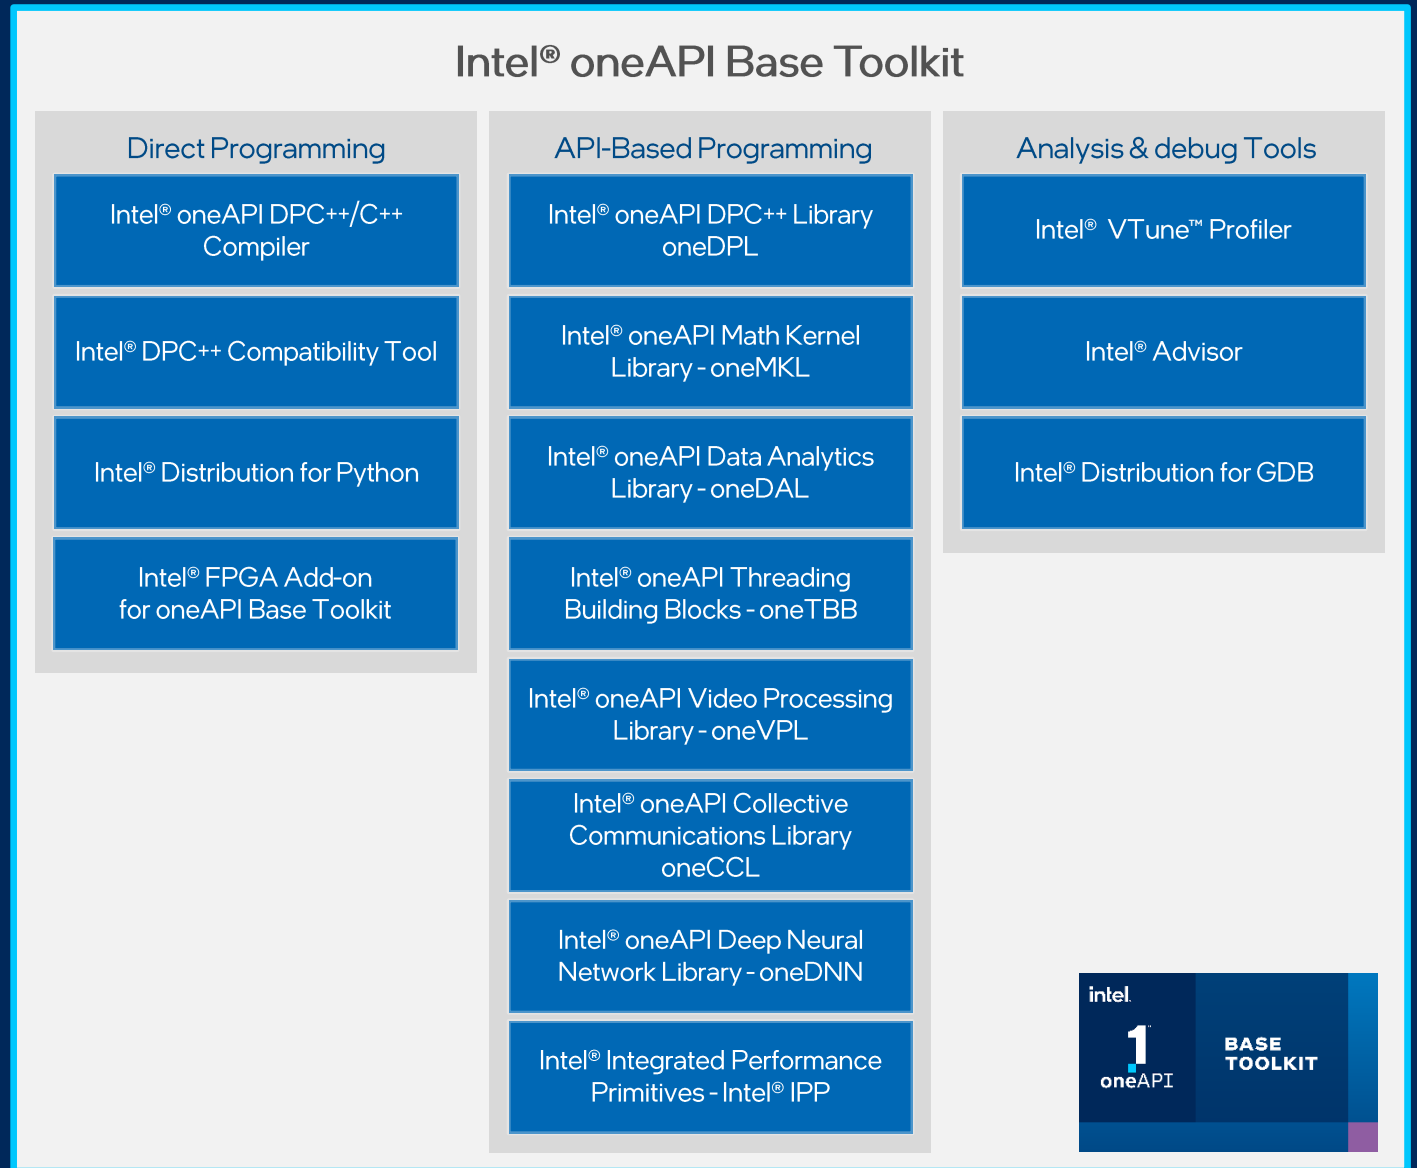
A core set of core tools and libraries for developing high-performance applications on Intel® CPUs, GPUs, and FPGAs.

Who Uses It?

- A broad range of developers across industries
- Add-on toolkit users since this is the base for all toolkits

Top Features/Benefits

- Data Parallel C++ compiler, library and analysis tools
- DPC++ Compatibility tool helps migrate existing code written in CUDA
- Python distribution includes accelerated scikit-learn, NumPy, SciPy libraries
- Optimized performance libraries for threading, math, data analytics, deep learning, and video/image/signal processing



Intel® oneAPI AI Analytics Toolkit

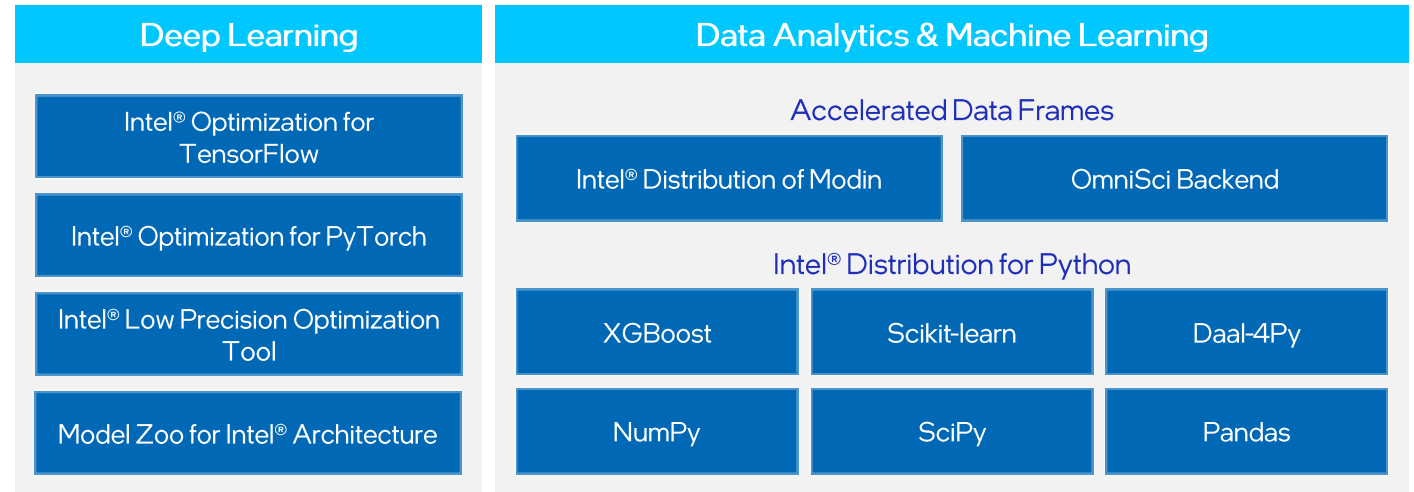
Accelerate end-to-end AI and data analytics pipelines with libraries optimized for Intel® architectures

Who Uses It?

Data scientists, AI researchers, ML and DL developers, AI application developers

Top Features/Benefits

- Deep learning performance for training and inference with Intel optimized DL frameworks and tools
- Drop-in acceleration for data analytics and machine learning workflows with compute-intensive Python packages




Samples and End2End Workloads



Supported Hardware Architectures¹

Hardware support varies by individual tool. Architecture support will be expanded over time. Other names and brands may be claimed as the property of others.

Get the Toolkit [HERE](#) or via these locations



Data Parallel C++

Standards-based, Cross-architecture Language

DPC++ = ISO C++ and Khronos SYCL and community extensions

Freedom of Choice: Future-Ready Programming Model

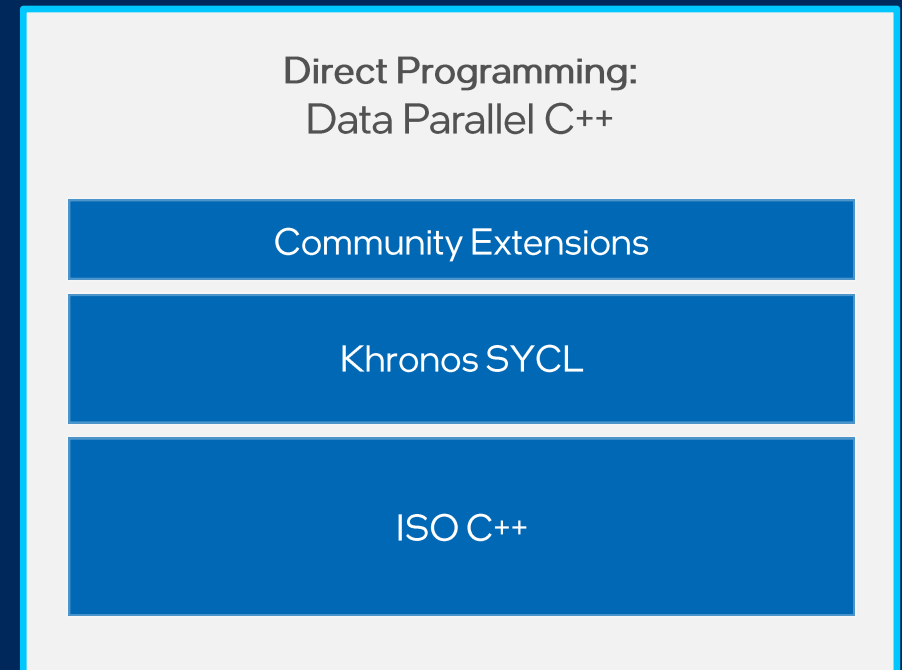
- Allows code reuse across hardware targets
- Permits custom tuning for a specific accelerator
- Open, cross-industry alternative to proprietary language

DPC++ = ISO C++ and Khronos SYCL and community extensions

- Delivers C++ productivity benefits, using common, familiar C and C++ constructs
- Adds SYCL from the Khronos Group for data parallelism and heterogeneous programming

Community Project Drives Language Enhancements

- Provides extensions to simplify data parallel programming
- Continues evolution through open and cooperative development



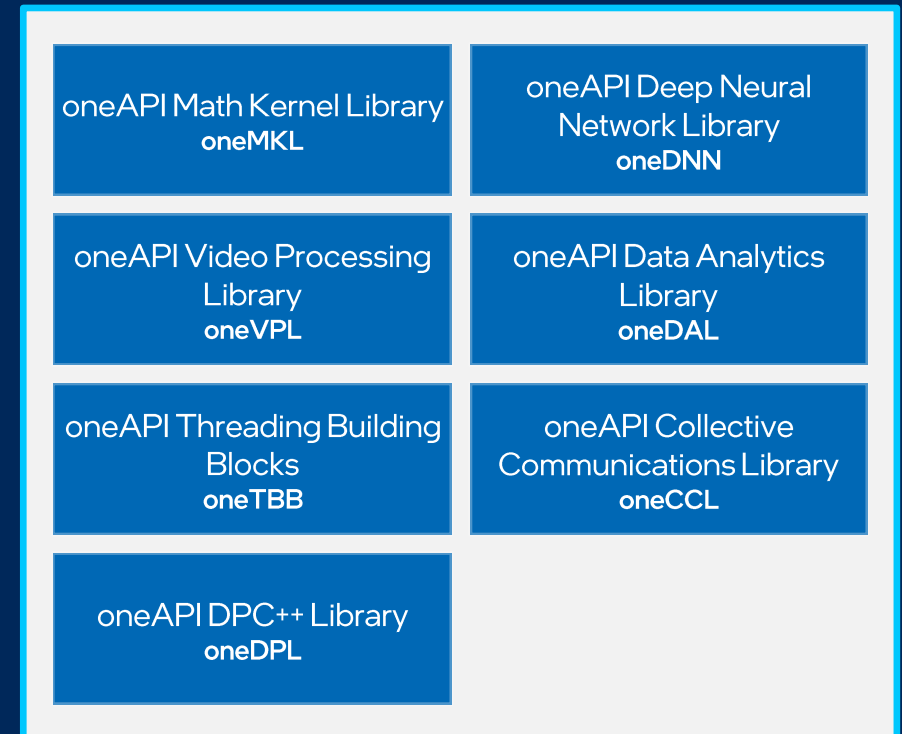
Powerful oneAPI Libraries

Realize all the Hardware Value

Designed for acceleration of key domain-specific functions

Freedom of Choice

Pre-optimized for each target platform for maximum performance



Intel® DPC++ Compatibility Tool

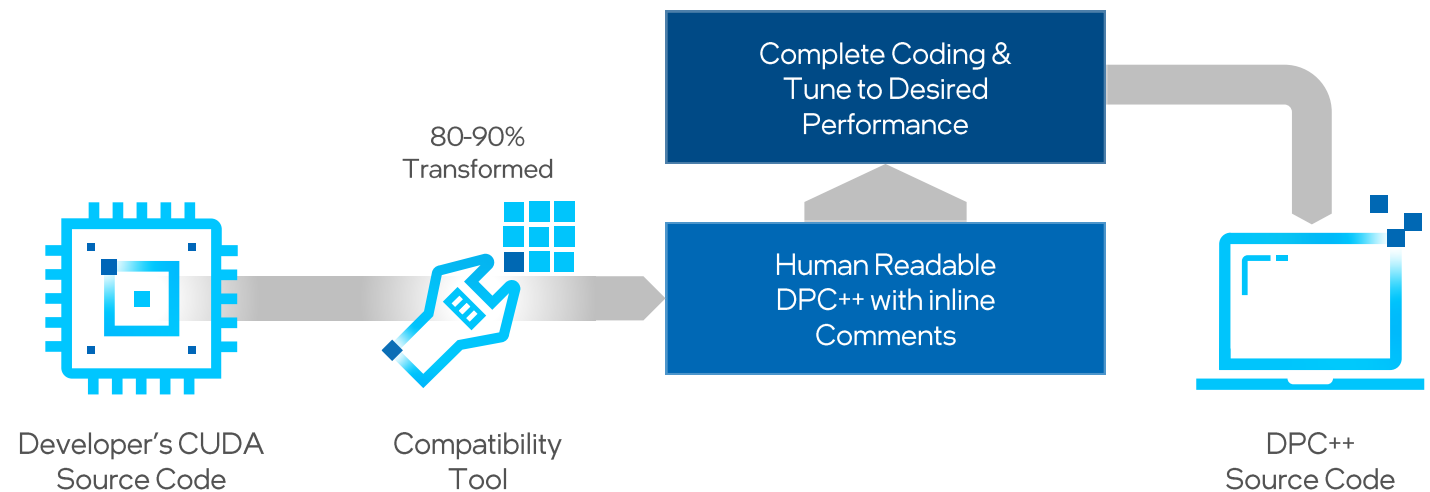
Minimizes Code Migration Time

Assists developers migrating code written in CUDA to DPC++ once, generating **human readable** code wherever possible

~80-90% of code typically migrates automatically

Inline comments are provided to help developers finish porting the application

Intel DPC ++ Compatibility Tool Usage Flow



Intel[®] oneAPI DPC++/C++ Compiler

Parallel Programming Productivity & Performance

Compiler to deliver uncompromised parallel programming productivity and performance across CPUs and accelerators

- Allows code reuse across hardware targets, while permitting custom tuning for a specific accelerator
- Open, cross-industry alternative to single architecture proprietary language

DPC++ is based on ISO C++ and Khronos SYCL

- Delivers C++ productivity benefits, using common and familiar C and C++ constructs
- Incorporates SYCL from The Khronos Group to support data parallelism and heterogeneous programming

Builds upon Intel's decades of experience in architecture and high-performance compilers

oneAPI DPC++/C++ Compiler and Runtime

DPC++ Source Code

Clang/LLVM

DPC++ Runtime



CPU



GPU



FPGA

Intel[®] oneAPI DPC++ Library

Accelerate DPC++ Kernels on Intel CPUs, GPUs & FPGAs

Optimized C++ Standard Algorithms

Contains 75 parallelized C++17 algorithms and utilities for efficient application development and deployment on a variety of hardware.

Based on parallel libraries that C++ developers are already familiar with

Incorporates popular libraries Parallel STL and Boost. Compute for easier developer adoption.

Integrated with Intel[®] DPC++ Compatibility Tool

Complements all oneAPI DPC++ components to simplify migration of developers' CUDA* code to DPC++ code.

Intel[®] oneAPI Deep Neural Network Library

Deliver High Performance Deep Learning


Helps developers create high performance deep learning frameworks

Abstracts out instruction set & other complexities of performance optimizations

Same API for both Intel CPUs and GPUs, use the best technology for the job

Supports Linux, Windows

Open sourced for community contributions



Intel® oneAPI Video Processing Library

Boost Media Performance

Boost media and video application performance with hardware-accelerated codecs and programmable graphics on Intel CPUs and Intel GPUs

Simple API that works the same on CPU and GPU

Using the API, developers have full control over codec visual quality and performance



Intel® oneAPI Collective Communications Library

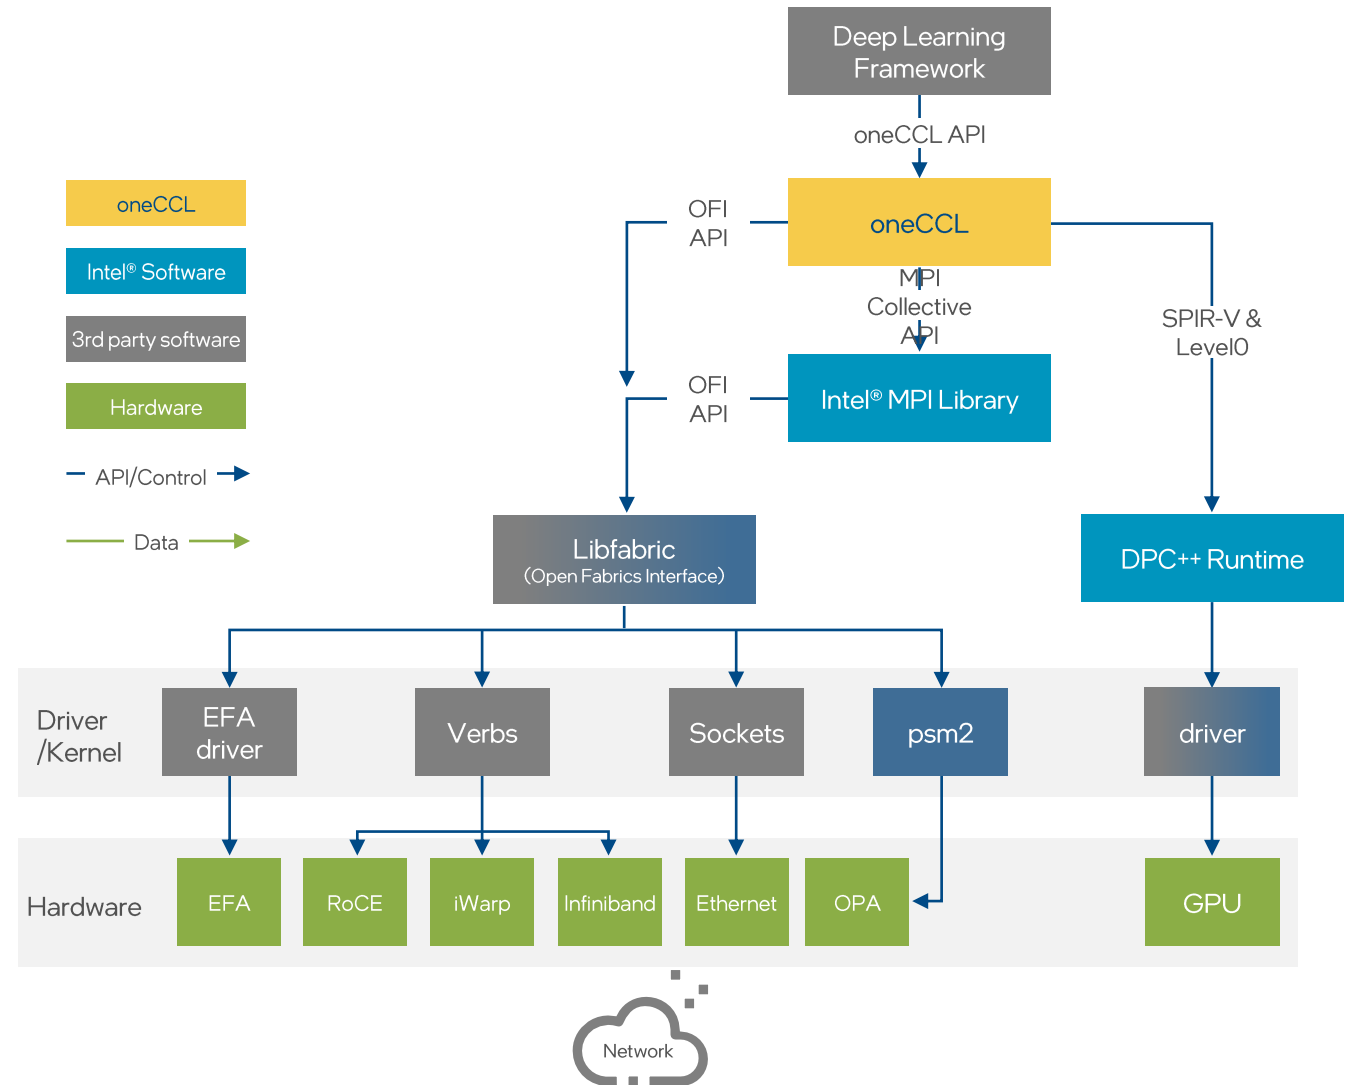
Optimize Communication Patterns

Provides optimized communication patterns for high performance on Intel CPUs & GPUs to distribute model training across multiple nodes

Transparently supports many interconnects, such as Intel® Omni-Path Architecture, InfiniBand, & Ethernet

Built on top of lower-level communication middleware-MPI & libfabrics

Enables efficient implementations of collectives used for deep learning training-all-gather, all-reduce, & reduce-scatter



Intel[®] VTune[™] Profiler

DPC++ Profiling—Tune for CPU, GPU & FPGA

Analyze Data Parallel C++ (DPC++)

See the lines of DPC++ that consume the most time

Tune for Intel CPUs, GPUs & FPGAs

Optimize for any supported hardware accelerator

Optimize Offload

Tune OpenMP offload performance

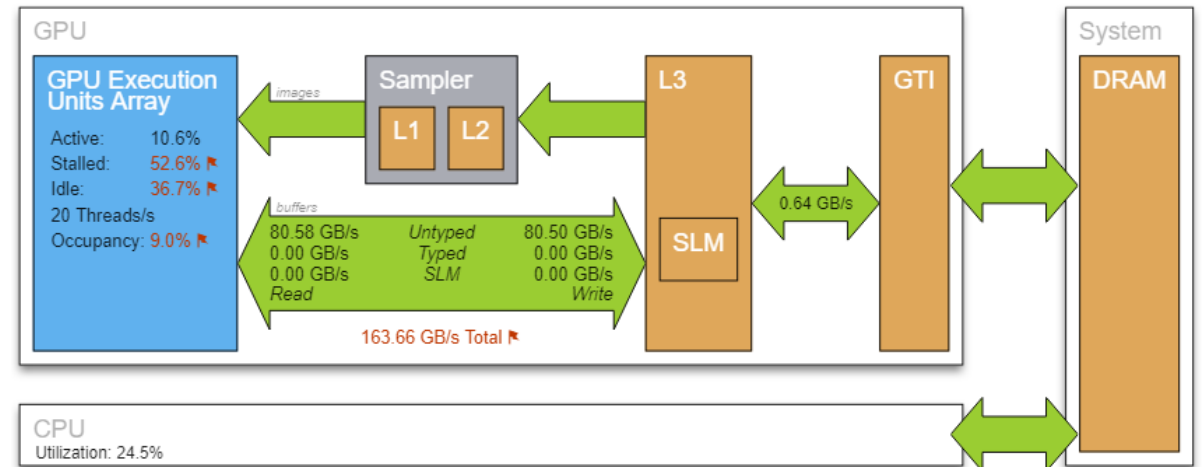
Wide Range of Performance Profiles

CPU, GPU, FPGA, threading, memory, cache, storage...

Supports Popular Languages

DPC++, C, C++, Fortran, Python, Go, Java, or a mix

Source	Assembly	GPU Instructions Executed by Instruction T...
158	<code>dx = ptr[j].pos[0] - ptr[i].pos[0];</code>	75,002,500
159	<code>dy = ptr[j].pos[1] - ptr[i].pos[1];</code>	12,500,000
160	<code>dz = ptr[j].pos[2] - ptr[i].pos[2];</code>	12,500,000
161		
162	<code>distanceSqr = dx*dx + dy*dy + dz*dz;</code>	87,500,000
163	<code>distanceInv = 1.0 / sqrt(distanceSqr);</code>	12,500,000
164		
165	<code>ptr[i].acc[0] += dx * G * ptr[j].ma</code>	162,503,750
166	<code>ptr[i].acc[1] += dy * G * ptr[j].ma</code>	150,000,000
167	<code>ptr[i].acc[2] += dz * G * ptr[j].ma</code>	150,000,000



Intel[®] Advisor

Design Assistant—Design for Modern Hardware

Offload Advisor

Estimate performance of offloading to an accelerator

Roofline Analysis

Optimize CPU/GPU code for memory and compute

Vectorization Advisor

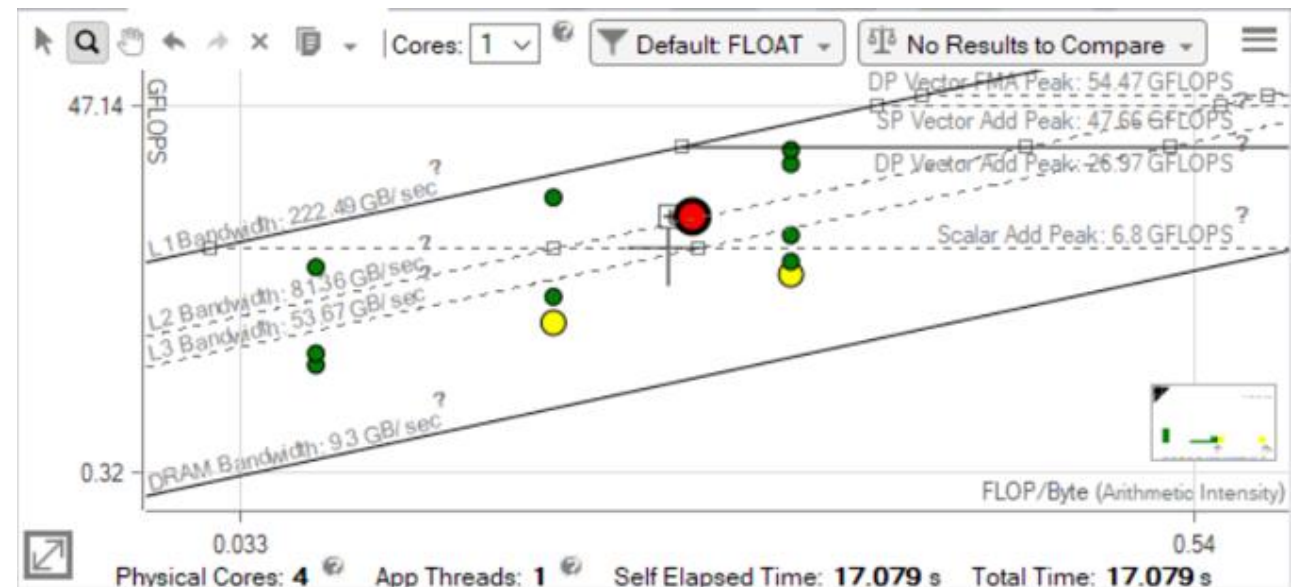

Add and optimize vectorization

Threading Advisor

Add effective threading to unthreaded applications

Flow Graph Analyzer

Create and analyze efficient flow graphs



Intel[®] Distribution of GDB

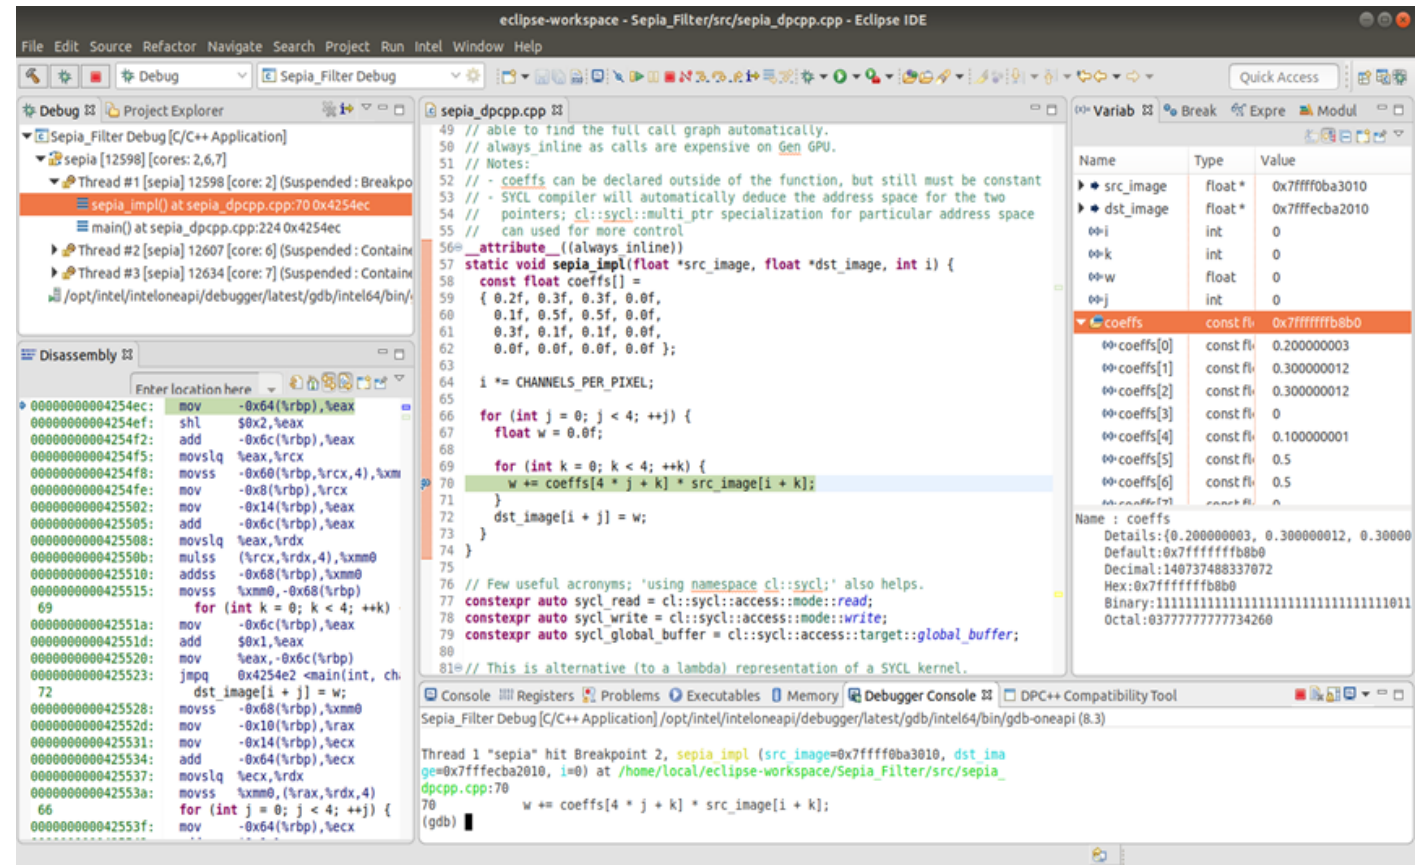
Data Parallel C++ Debug—Heterogeneous Application Debug

High-level language debug support

Multiple accelerator support:
Intel CPU, GPU, FPGA emulation

Auto-detect accelerator architecture
during application runtime

Non-proprietary open-source solution
based on GDB



oneAPI for FPGA

DPC++ Coding for Spatial Architecture

For Experienced FPGA Developers

Ease of Use

Experienced FPGA users can take advantage of a streamlined programming model using Data Parallel C++

Real Time Processing

Process data faster with deterministic low latency and high throughput

Runtime Analysis Support

Collect profiling data at runtime to analyze CPU and FPGA interaction with Intel® VTune™ Profiler

Device Specific Optimizations

One-day class provides experienced FPGA developers training to begin optimizing oneAPI code for FPGA

