

IPXPUG Asia January 2021

oneAPI Industry Initiative for Accelerated Computing

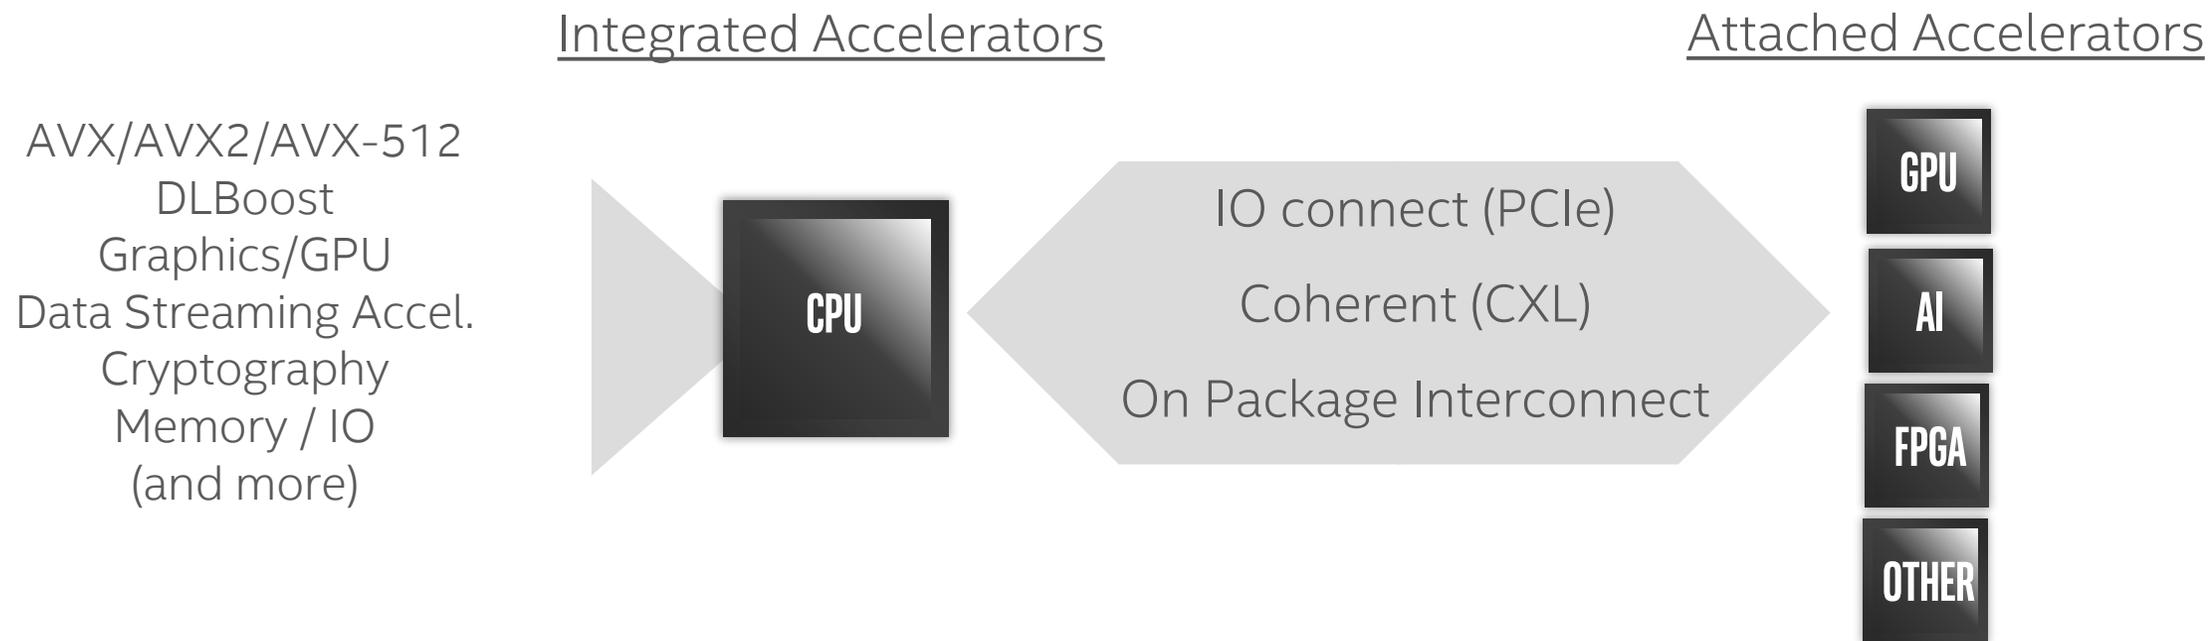
Joe Curley – oneAPI Products, Solutions, & Ecosystem

21 – January 2021

The Intel logo is located in the bottom left corner of the slide. It consists of the word "intel" in a lowercase, sans-serif font, with a registered trademark symbol (®) to its upper right. The logo is white and stands out against the dark blue background. To the left of the logo, there is a decorative graphic of several overlapping squares in various shades of blue, arranged in a stepped pattern.

intel®

Acceleration Technologies



We are experiencing a “boom” in accelerated computing technology

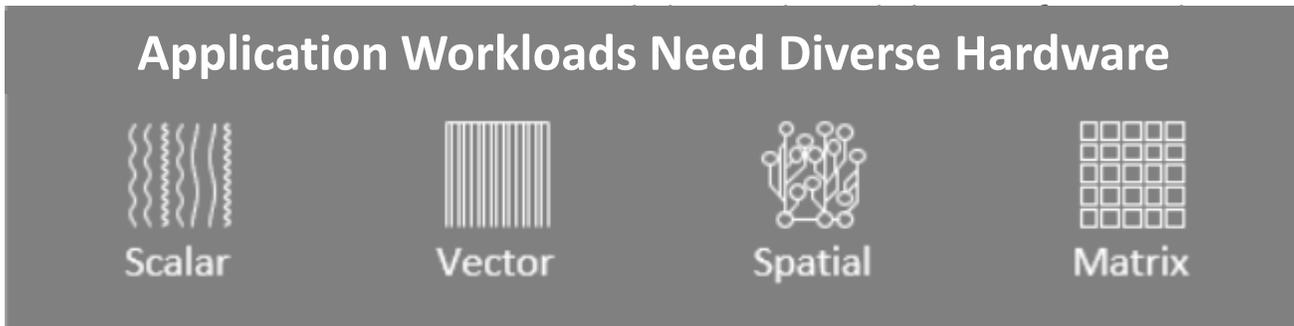
In the host. In the interconnect. In the node. In the Rack. In the DataCenter.

Programming Challenges

for Multiple Architectures

Growth in specialized workloads

Variety of data-centric hardware required



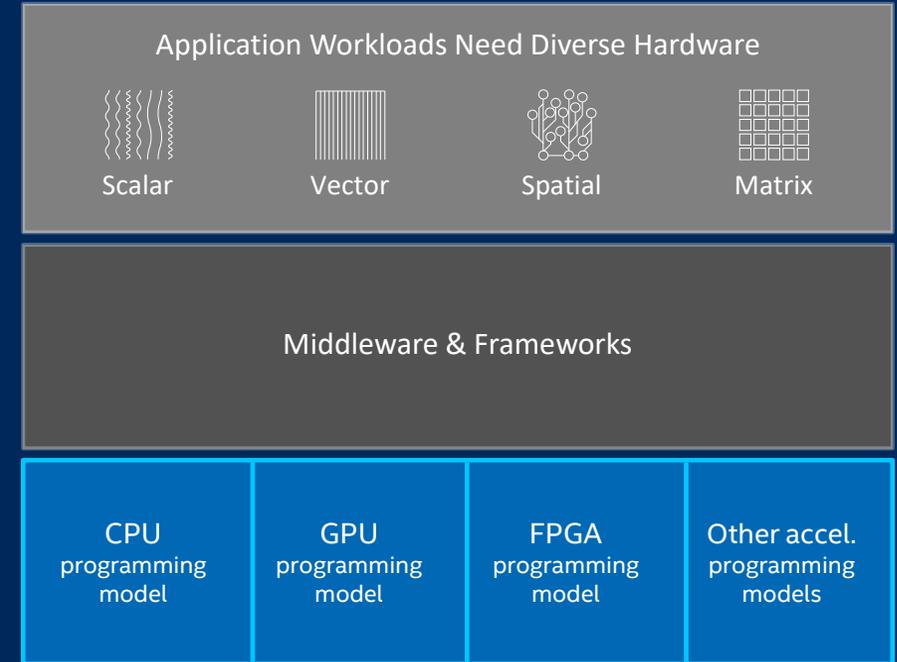
Programming Challenges

for Multiple Architectures

Growth in specialized workloads

Variety of data-centric hardware required

Separate programming models and toolchains for each architecture are required today



Programming Challenges

for Multiple Architectures

At least 4 accelerated micro architectures

Times

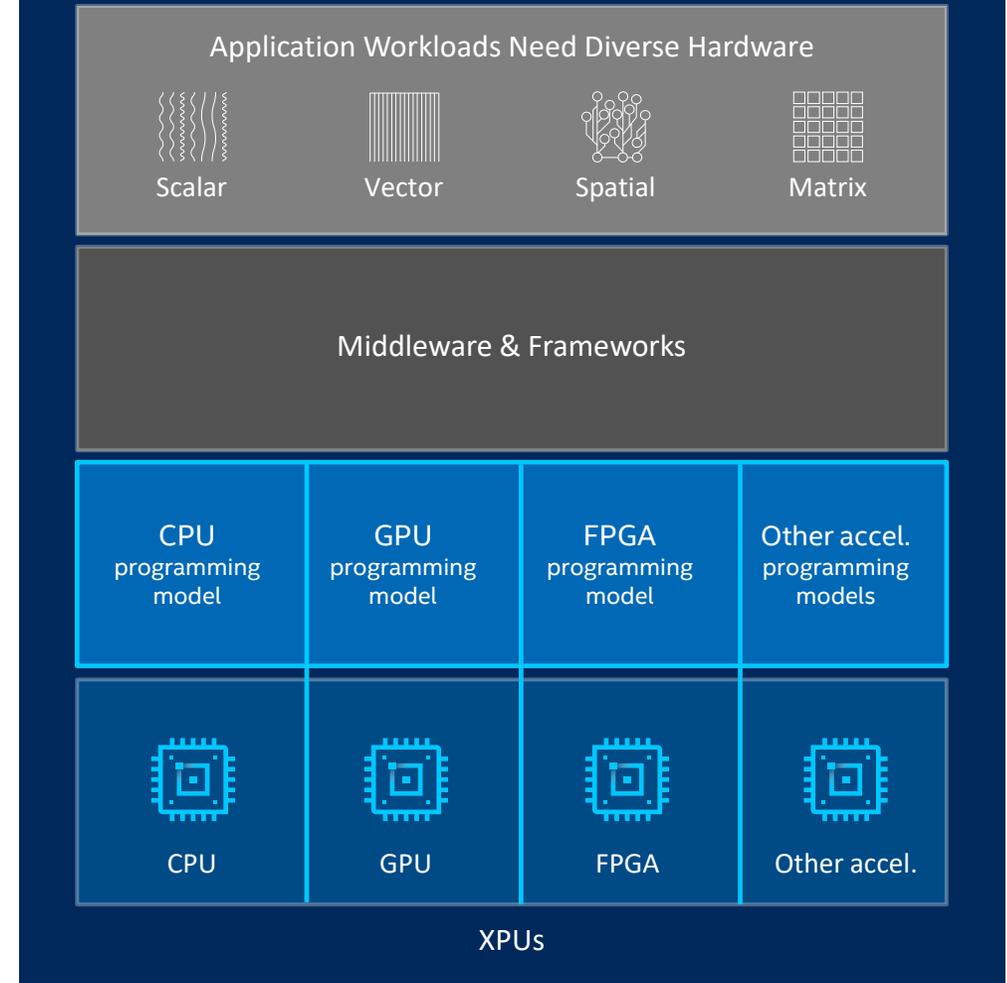
n vendors

Times

1 tool chain per implementation

equals...

...Challenges



oneAPI Industry Initiative



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

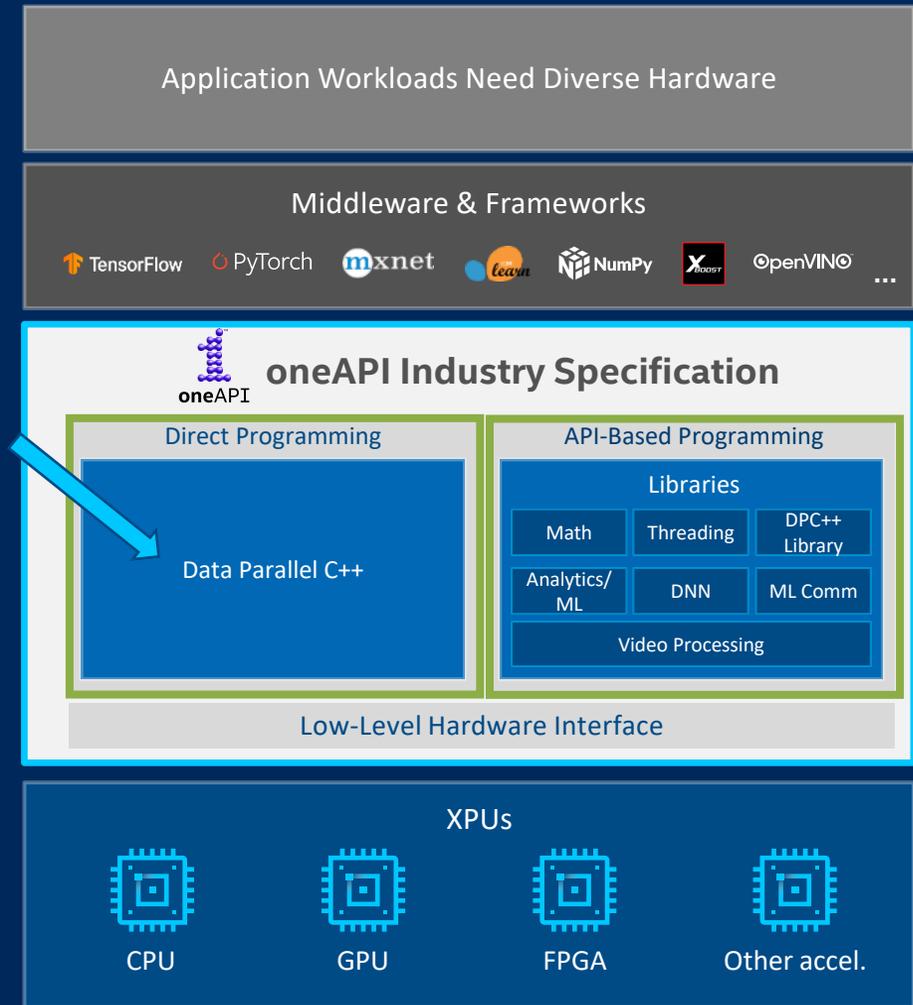
Visit oneapi.com for more details

Open to promote community and industry collaboration

Enables code reuse across architectures and vendors

Frees developers from challenges of proprietary programming models

A cross-architecture language based on C++ and SYCL standards



Data Parallel C++

Standards-based, Cross-architecture Language

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

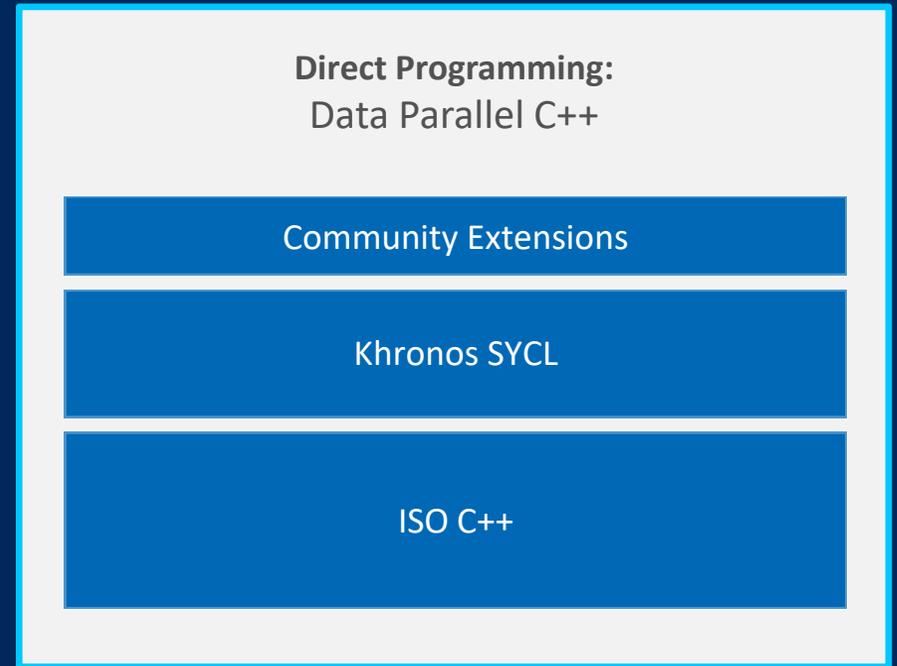
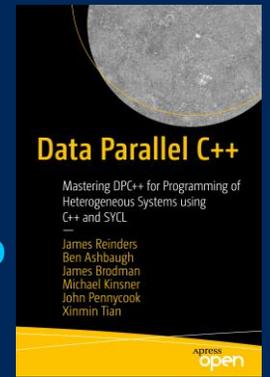
- Delivers C++ productivity benefits, using common, familiar C and C++ constructs
- Adds SYCL from the Khronos Group for data parallelism and heterogeneous programming
 - Modern C++ like language
 - Open Industry standard
 - Multi-Vendor support, growing ecosystem

Community Project Drives Language Enhancements

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

The open source and Intel DPC++/C++ compiler supports Intel CPUs, GPUs, and FPGAs. Codeplay announced a [DPC++ compiler that targets Nvidia GPUs](#).

Book (PDF) Download
tinyurl.com/DataParallelCpp



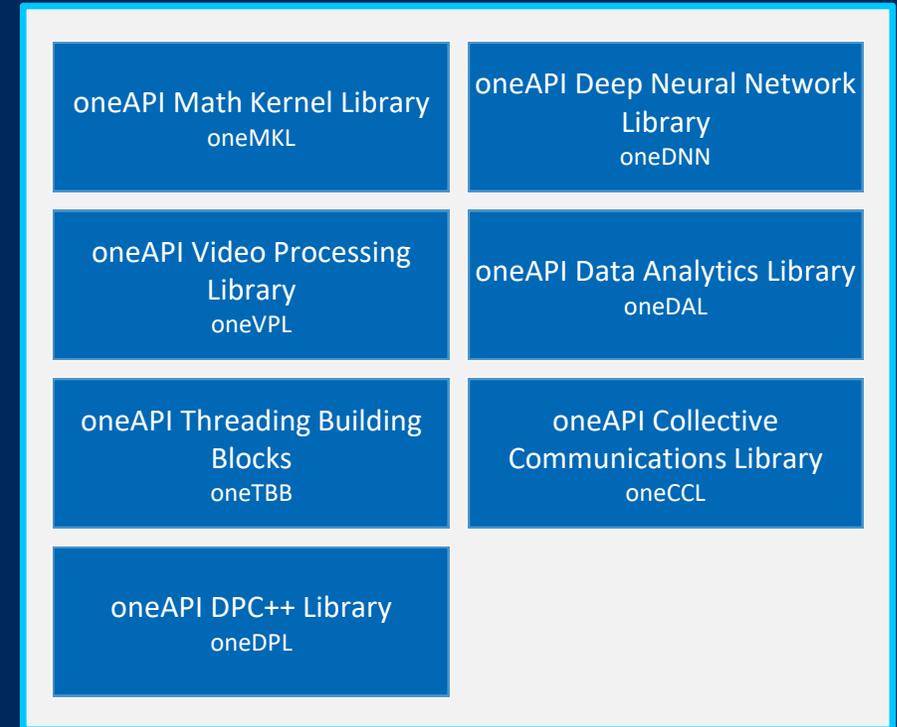
Powerful oneAPI Libraries

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

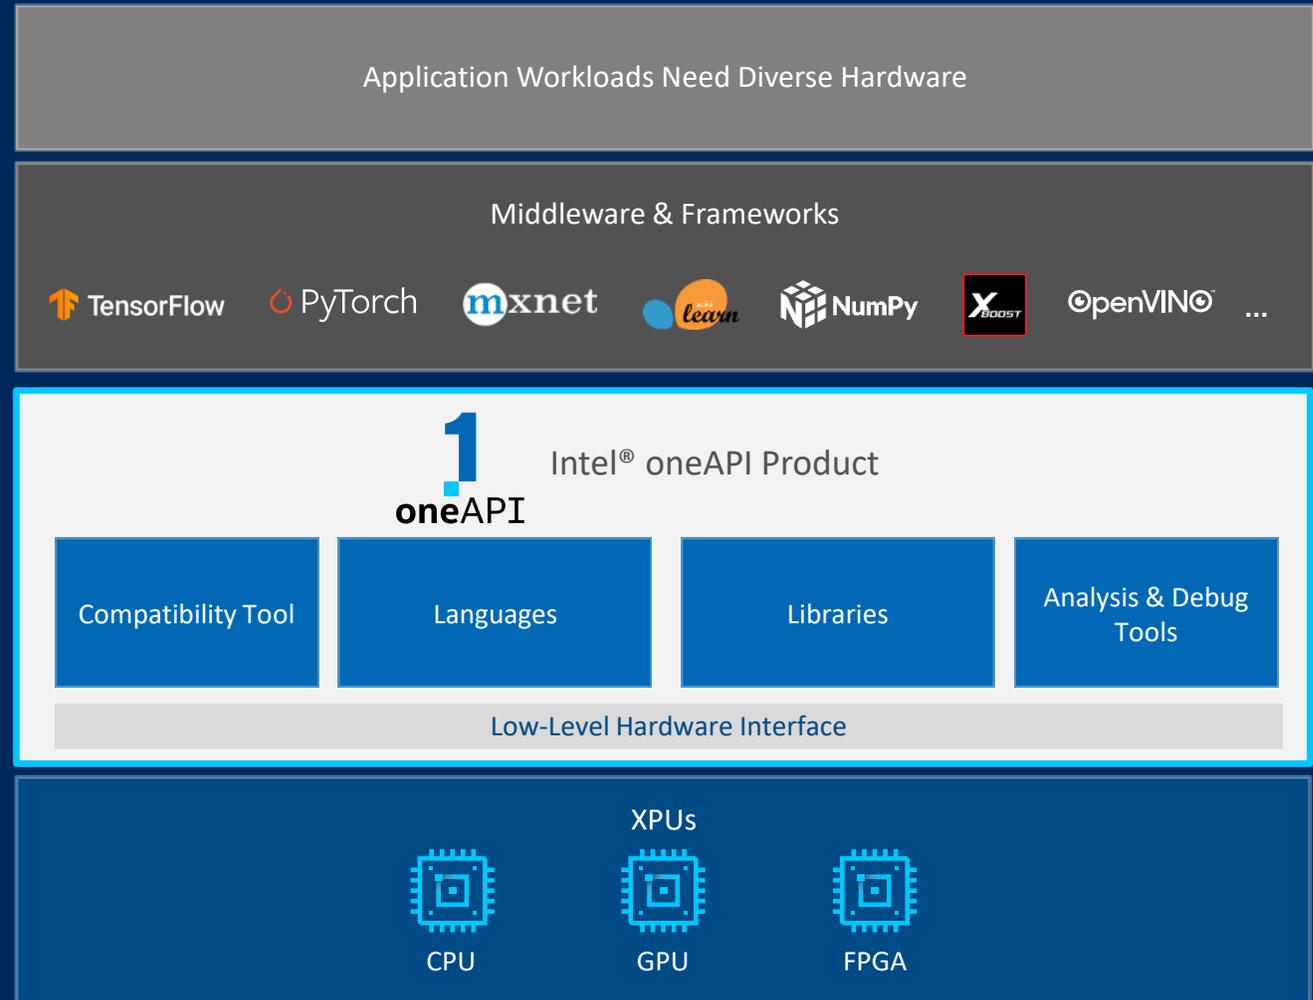


Intel® oneAPI Product

Built on Intel's Rich Heritage of CPU Tools Expanded to *XPU*s

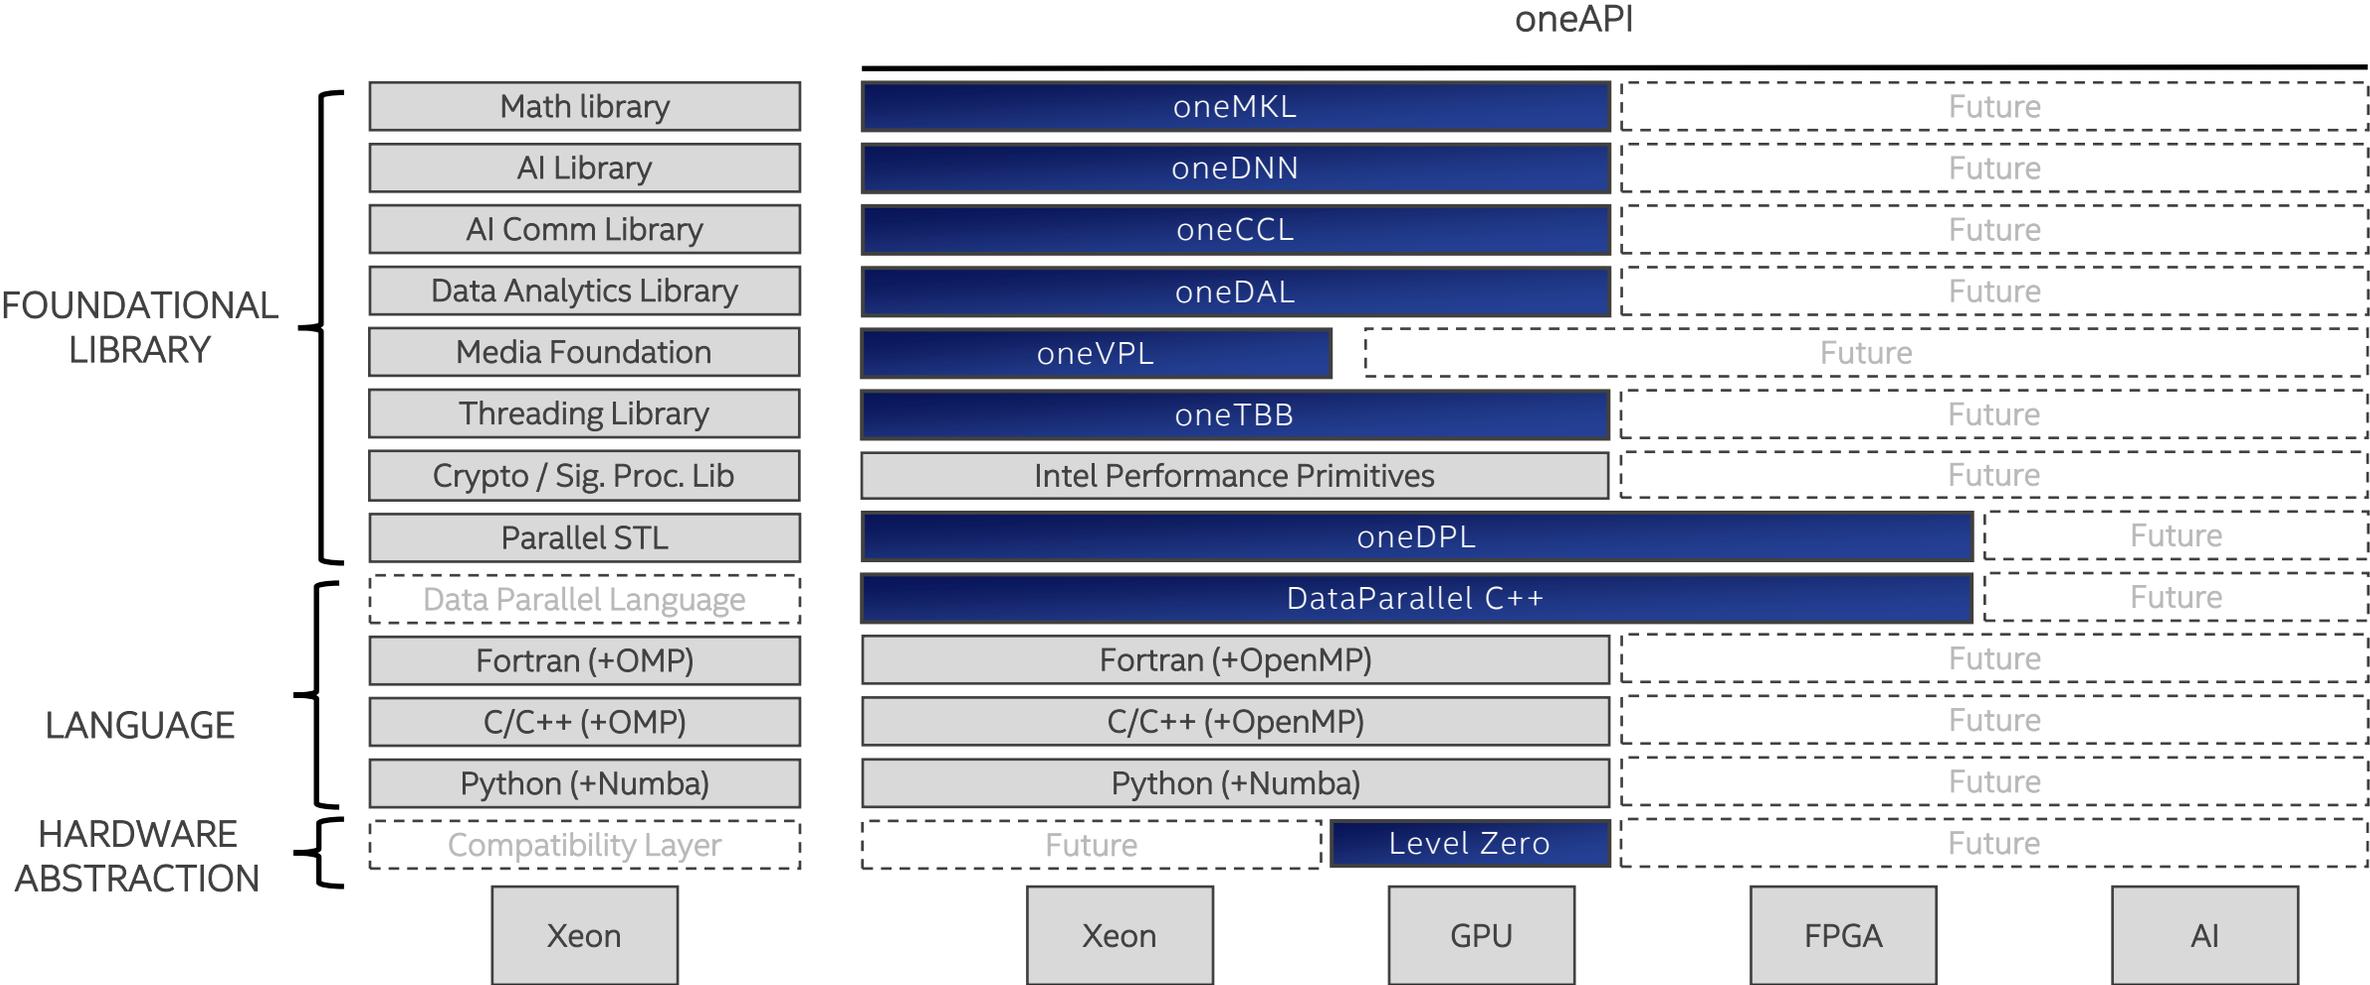
A complete set of advanced compilers, libraries, and porting, analysis and debugger tools

- Accelerates compute by exploiting cutting-edge hardware features
- Interoperable with existing programming models and code bases (C++, Fortran, Python, OpenMP, etc.), developers can be confident that existing applications work seamlessly with oneAPI
- Eases transitions to new systems and accelerators—using a single code base frees developers to invest more time on innovation



[Available Now](#)

Intel oneAPI 2021.1



First Customer Ship December 2020 - CPU, GPU, direct programming of FPGA

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 Rendering Toolkit

Create performant, high-fidelity visualization applications

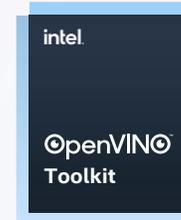
Toolkits powered by oneAPI

Data Scientists & AI Developers



Intel® AI Analytics Toolkit

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



Intel® Distribution of OpenVINO™ Toolkit

Deploy high performance inference & applications from edge to cloud

Latest version is 2021.1

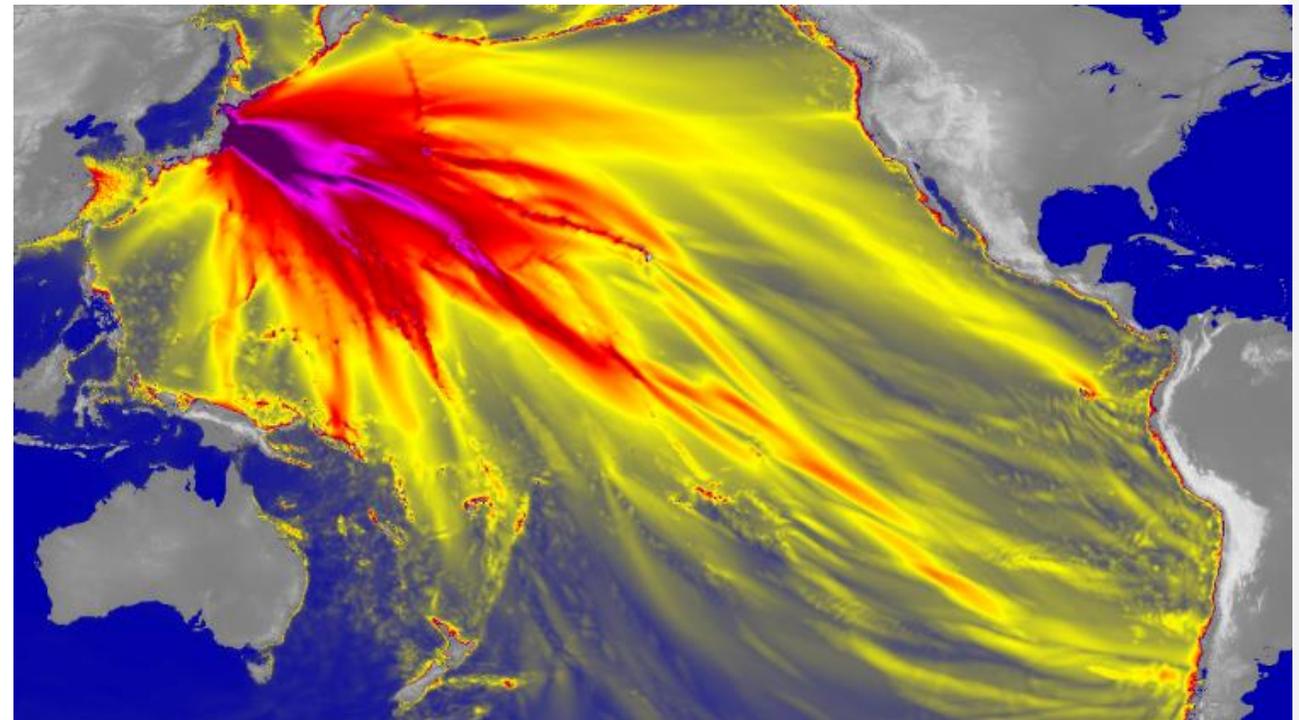
oneAPI Development Example



ZIB ported *EasyWave* application from CUDA to DPC++ delivering performance across multiple architectures

- Ported EasyWave written in CUDA to Data Parallel C++ efficiently using the Intel® DPC++ Compatibility Tool
- Achieved strong performance across Intel CPU, GPU and FPGA architectures, and within 5% of CUDA performance on Nvidia P100

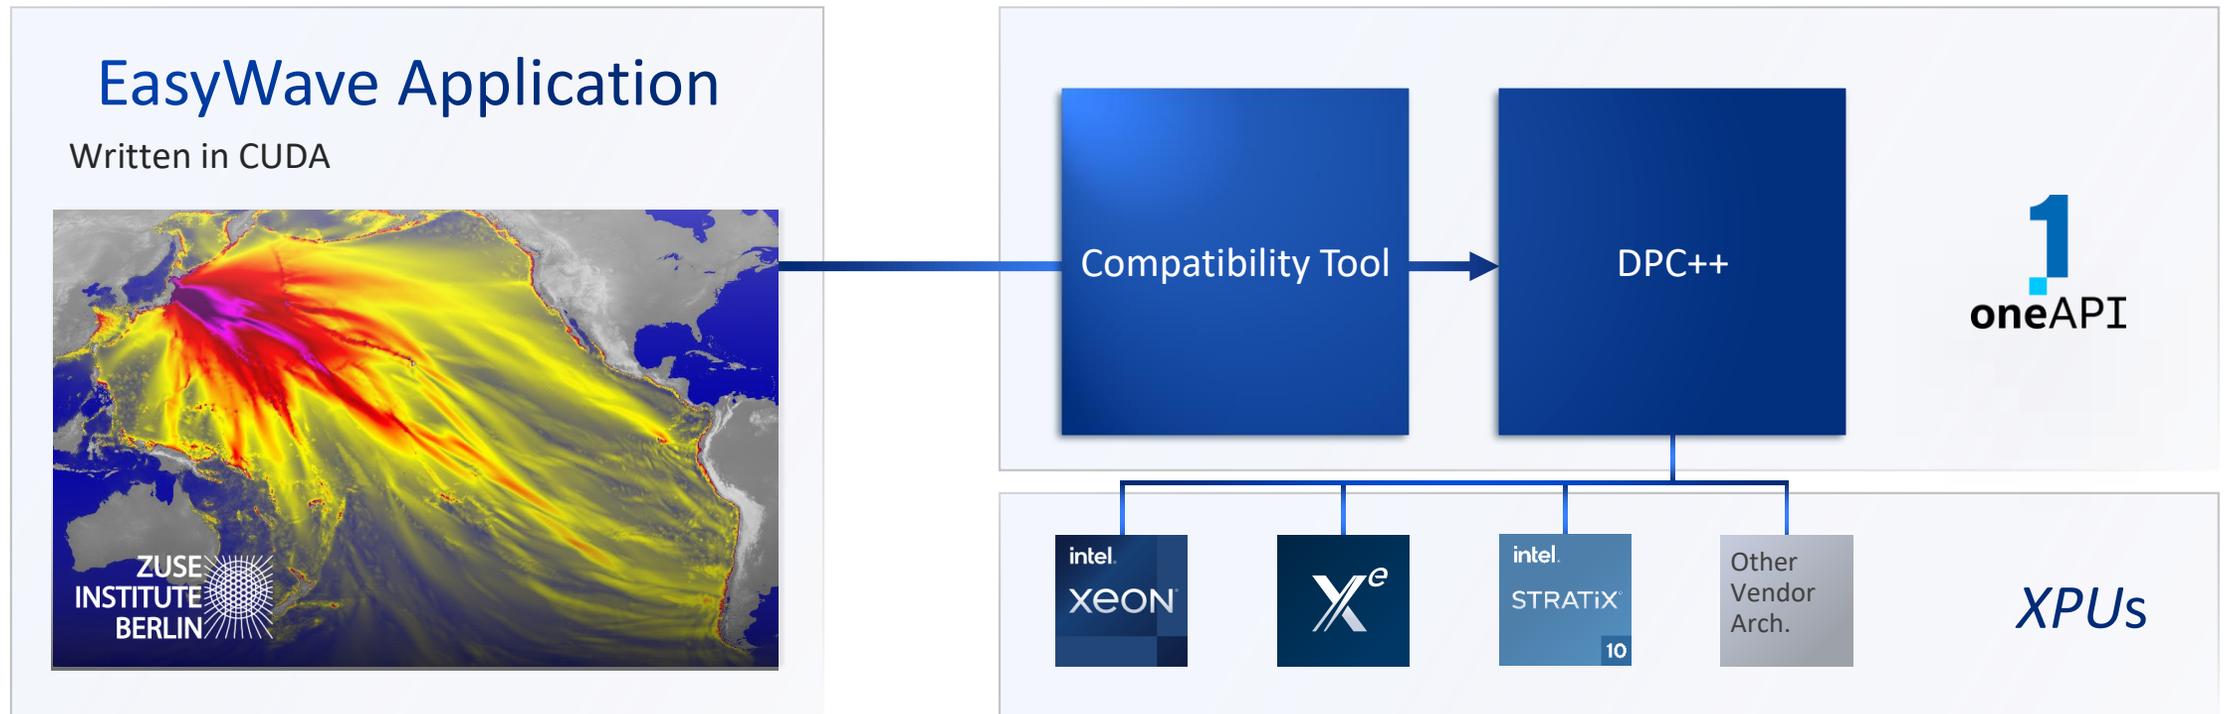
oneAPI used across XPUs



Visualization of EasyWave tsunami simulation application
Courtesy Zuse Institute Berlin (ZIB)

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

oneAPI Development Example - Workflow





GROMACS + oneAPI: Heterogenous Acceleration Molecular Dynamics from CUDA & OpenCL to SYCL



```

// Add corrections to Lagrange multipliers
float sqrtmu_sol = sqrtReducedMass*sol;
lagrangeScaled += sqrtmu_sol;

// Save updated coordinates for next iteration
// Dummy constraints are skipped
if (!isDummyThread)
{
    float3 tmp = rc*sqrtmu_sol;
    atomicAdd(&gm_xp[i], -tmp*invse);
    atomicAdd(&gm_xp[j], tmp*invse);
}

if (updateVelocities && !isDummyThread)
{
    float3 tmp = rc*invdt*lagrangeScaled;
    atomicAdd(&gm_v[i], -tmp*inverseM);
    atomicAdd(&gm_v[j], tmp*inverseM);
}

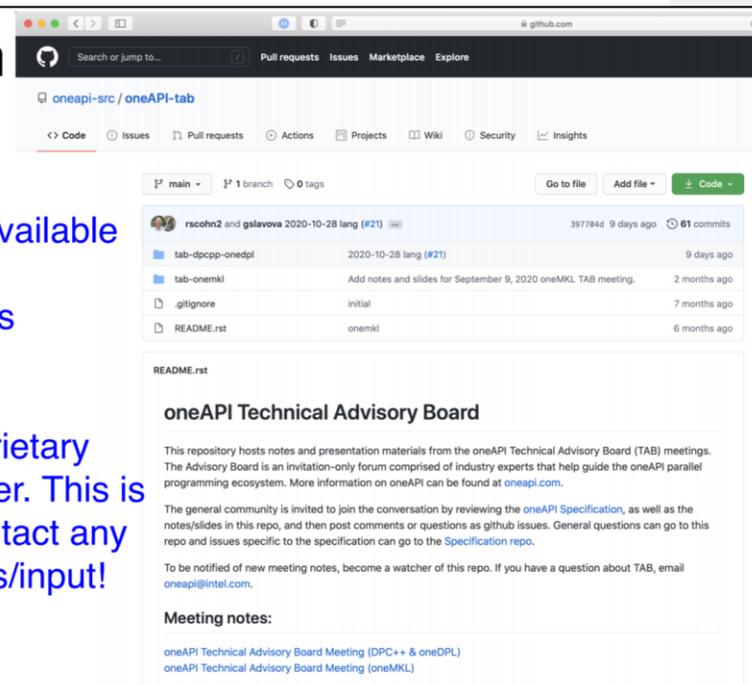
```

Erik Lindahl
Stockholm University
KTH Royal Institute of Technology
<https://twitter.com/eriklindahl>



“...give this (presentation) from the opposite of what Joe said yesterday, because I’m a developer... 😊
- Erik Lindahl, KTH Royal Institute of Technology

Open really means open



Open source implementation available

The oneAPI & SYCL extensions laundry is washed in public

We are not replacing one proprietary implementation/API with another. This is a strong community effort - contact any of us if you have opinions/ideas/input!

Video: <https://www.oneapi.com/videos/gromacs/>

Productivity. Openness. The realities of getting performance from XPU's

Intel® oneAPI Toolkits

oneAPI Toolkits – free download

Get Started Quickly

Code Samples, Quick-start Guides, Webinars, Training

software.intel.com/oneapi

Try in the DevCloud

Easy to register, access is free for 120 days with possibility of an extension

<https://intelsoftwaresites.secure.force.com/devcloud/oneapi>

Intel® Xeon® Scalable processors

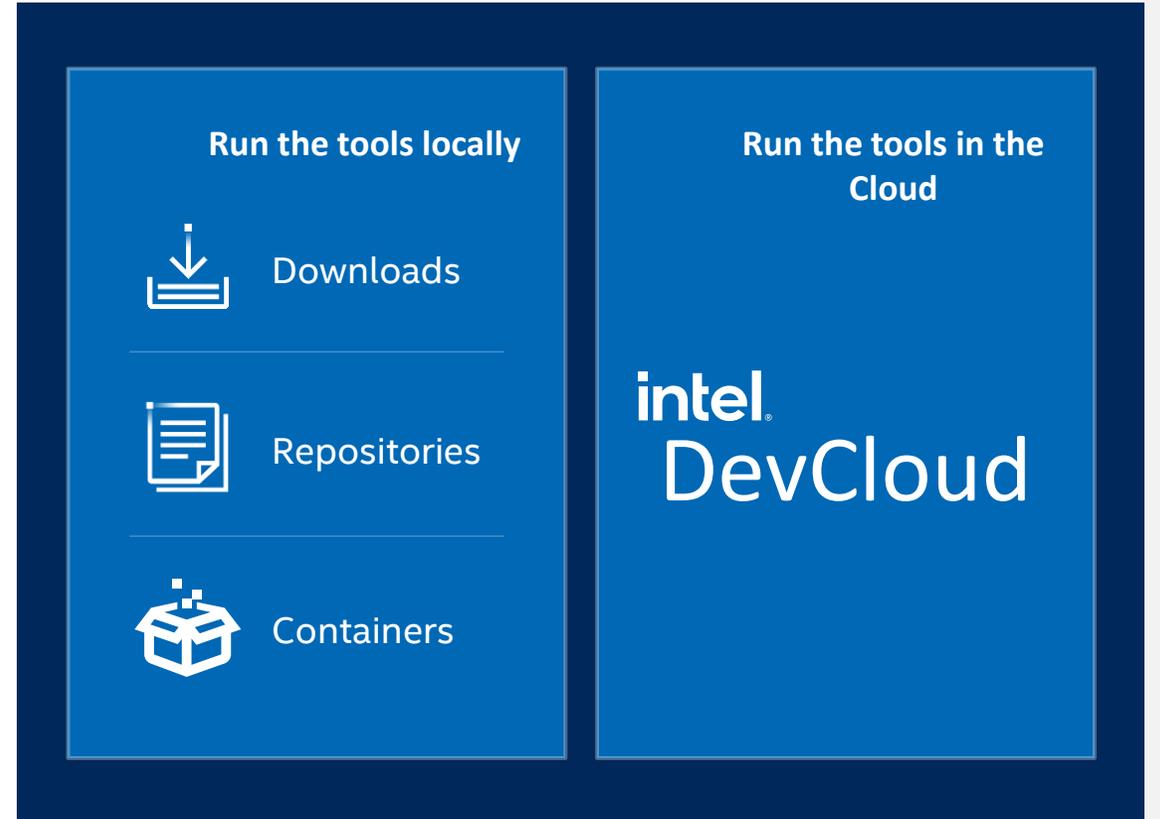
Intel® Xeon® Scalable processors with Intel Graphics Technology

Intel® Iris® Xe MAX GPU

Intel® Arria® 10 FPGAs

Intel® Stratix® 10 FPGAs

1
oneAPI



Call to action

Join us – we're at the beginning of a great journey

- Let's work together to establish open XPU computing
- Give feedback: www.oneapi.com

- Try out Intel's implementation:
 - Download (it's free!)
<https://software.intel.com/content/www/us/en/develop/tools/oneapi.html>
 - Or try in the Intel DevCloud (it's easy)
<https://intelsoftwaresites.secure.force.com/devcloud/oneapi>



oneAPI Ecosystem – add your contribution!



These organizations support the oneAPI initiative 'concept' for a single, unified programming model for cross-architecture development. It does not indicate any agreement to purchase or use of Intel's products.

*Other names and brands may be claimed as the property of others.
oneAPI

intel®

Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

All product plans and roadmaps are subject to change without notice.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

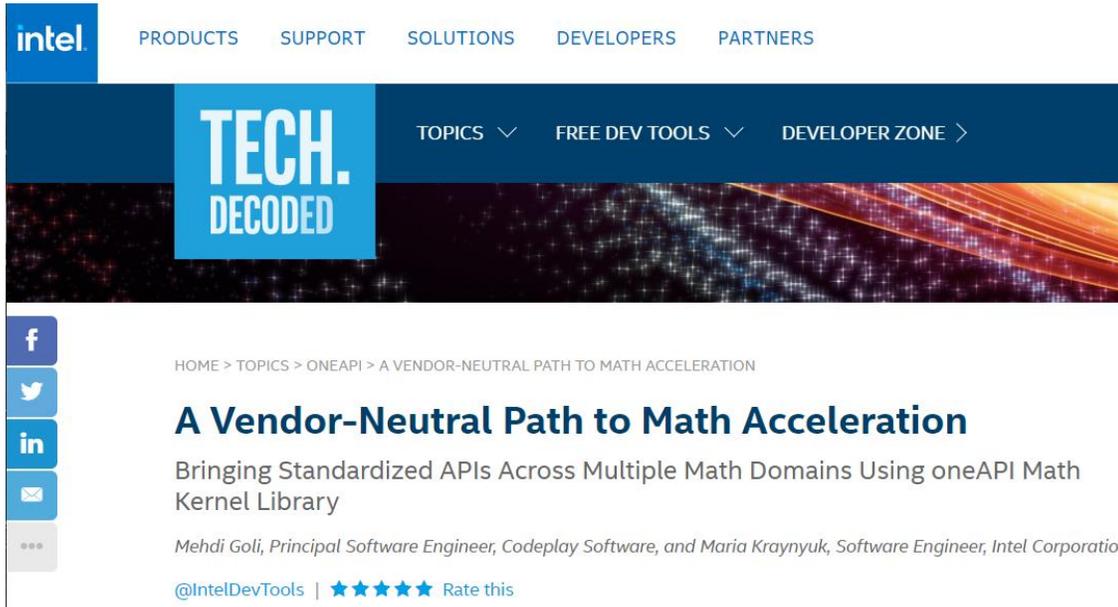
Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Useful Links:

- oneAPI specification:
 - <https://www.oneapi.com/documentation/>
- Intel's Jeff Hammond on SYCL
 - <https://www.youtube.com/watch?v=nwDxODr8mI8>
- Ecosystem developer talks at oneAPI developer summit 2020
 - <https://www.oneapi.com/events/devcon2020/>
- Dataparallel C++ book (free digital version)
 - <https://link.springer.com/content/pdf/10.1007%2F978-1-4842-5574-2.pdf>

Standardizing Library Interfaces



<https://techdecoded.intel.io/resources/a-vendor-neutral-path-to-math-acceleration/#gs.qqv92p>

