# Introducing the oneAPI Community Forum

Rod Burns, VP Ecosystem
Codeplay Software

AI DevSummit, December 2022

oneAPI

# Programming Challenges for Parallel Computing

- No common programming language or APIs for parallel workloads, lacks maturity or support

- Significant investment for software developers migrating to new hardware

- Heterogeneous architectures are multi-vendor

oneAPI

Application Workloads Need Diverse Hardware

Scalar    Vector    Matrix    Spatial

Middleware / Frameworks

Languages & Libraries

CPU    GPU    FPGA    Other Accel.

# oneAPI Mission

- Industry defined, <u>open standard-based</u> APIs for accelerated devices

- Bring <u>open source</u> implementations for wide adoption

- Gather industry leaders and contributors to support <u>the mission</u>

- Enable <u>diverse</u> processor designs

## oneAPI Community Forum

Define a standards-based open specification

Foster open-source implementations of the specification

## oneAPI Projects
## Open Source, Open Standard

Libraries

Programming Model

Hardware Interface

| CPU | GPU | FPGA | Other Accel. |

# What is the oneAPI Community Forum?

**1**

A cross industry group of hardware and software experts

**2**

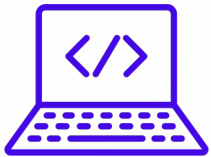Defines standard interfaces for accelerator computing

**3**

ÆĪ ǦĜ Ǧ·
ÎI Ě ËǦĜ Ĝ Ĩ·
ĨĚ ËG Ĥ Ĝ ÊĢ
İ Ħ ÍĠ Ĝ F·F ÍĦĪ I Î

**4**
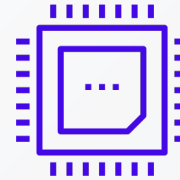
Drives the future of open-standard accelerator computing

oneAPI

# Benefits

oneAPI

### For Software Developers

- Develop with open standards for accelerator computing

- Single code base for multiple processors targets

- Standards and industry defined libraries

- Enable an existing ecosystem of software and educational resources

- Future proof your software

### For Hardware Developers

- Fast track into market with an existing ecosystem

- Share development cost with open-source implementations

- Enable an existing ecosystem of software and educational resources

- Leverage an existing tested and optimized toolchain

## Free and based on open standards

# Contribute to the oneAPI Forum

- oneAPI Community Forum Steering Committee

- Join and lead SIGs and Working Groups

- Submit proposals for features and changes
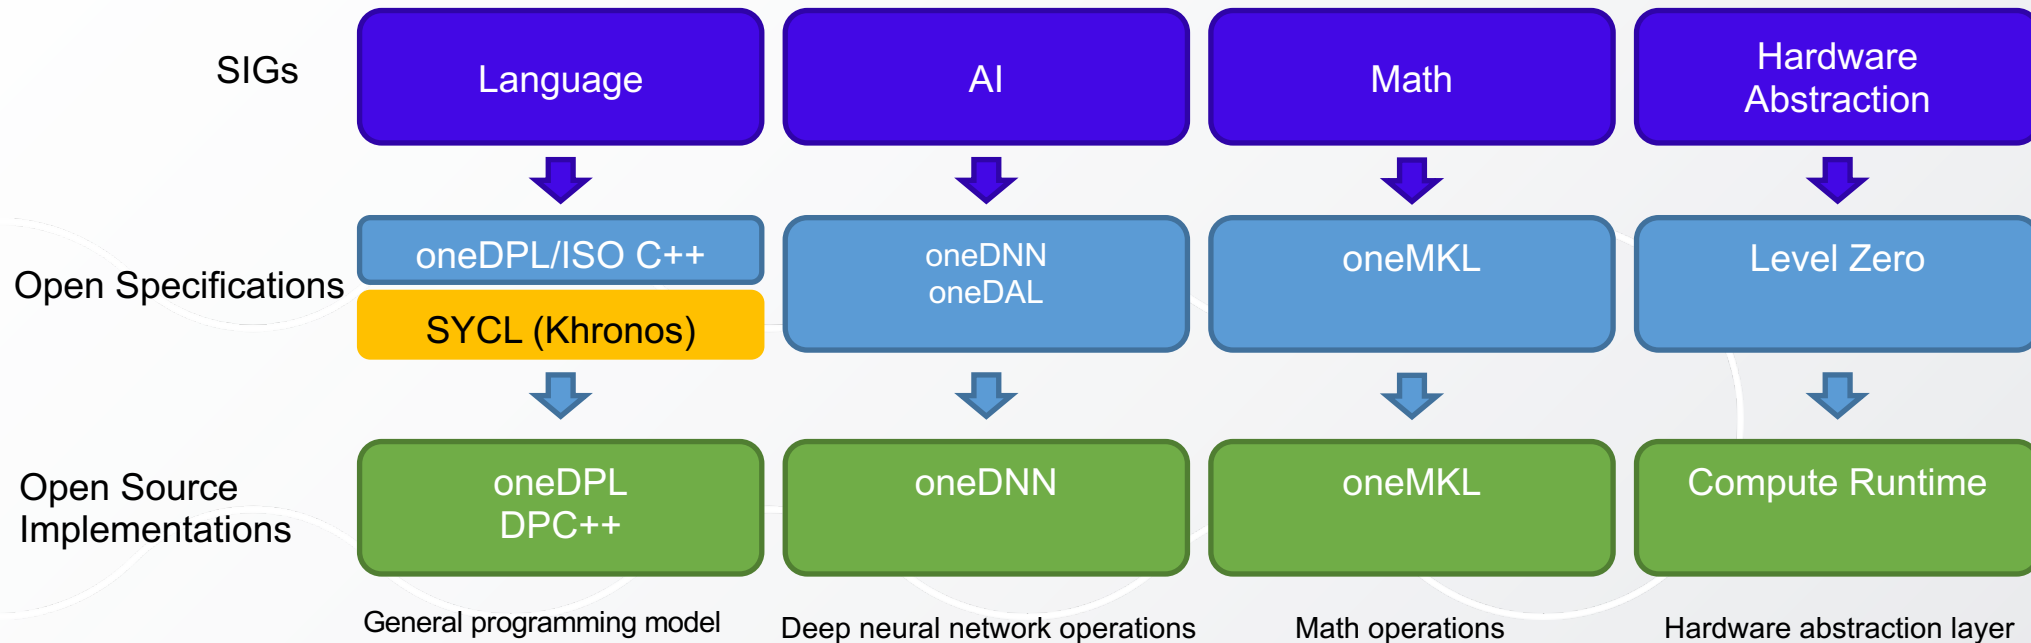
- Vote on proposals

Drive the future of programming for heterogeneous architectures

https://www.oneapi.io/

# Special Interest Groups

Special Interest Groups influence the specifications and implementations

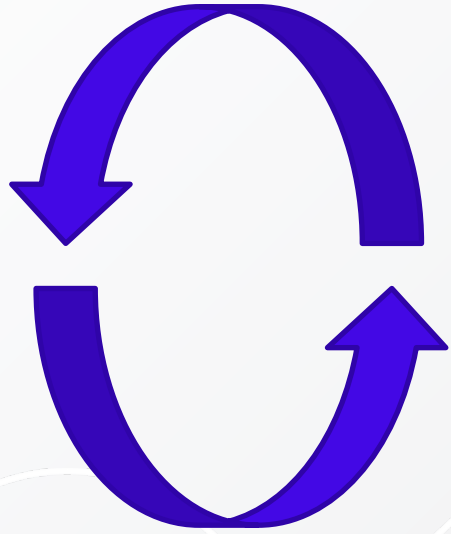| SIGs | Language | AI | Math | Hardware Abstraction |
|---|---|---|---|---|
| Open Specifications | oneDPL/ISO C++  <br> SYCL (Khronos) | oneDNN  <br> oneDAL | oneMKL | Level Zero |
| Open Source Implementations | oneDPL  <br> DPC++ | oneDNN | oneMKL | Compute Runtime |
|  | General programming model | Deep neural network operations | Math operations | Hardware abstraction layer |

# Join the oneAPI Community Forum

Influence the spec

## Browse the specification

**DPC++**
**oneAPI** Data Parallel C++ (SYCL)

**oneDPL**
**oneAPI** Data Parallel C++ Library

**oneDNN**
**oneAPI** Deep Neural Network Library

**oneCCL**
**oneAPI** Collective Communications Library

**Level Zero**
**oneAPI** Level Zero

**oneDAL**
**oneAPI** Data Analytics Library

**oneTBB**
**oneAPI** Threading Building Blocks

**oneVPL**
**oneAPI** Video Processing Library

**oneMKL**
**oneAPI** Math Kernel Library

**Ray Tracing**
**oneAPI** Ray Tracing

**https://www.oneapi.io/spec/**

## Join Special Interest Groups

AI

Math

Language

Hardware Abstraction

**oneapi.io/community**

## Explore the open source projects

Contribute to open source

**oneAPI-SRC**
oneAPI open source projects

🧑‍🤝‍🧑 240 followers   📍 Hillsboro OR   🔗 https://oneapi.com   ✉ webadmin@linux.intel.com

🏠 Overview   📒 Repositories 38   🗂 Projects 2   📦 Packages   👥 Teams 34   👤 People 329

Pinned

**https://github.com/oneapi-src**

oneAPI

# oneAPI Ecosystem

Organizations are designing new processors based on the oneAPI spec to foster the HPC ecosystem and deliver numerous capabilities. Advancements using oneAPI are improving the work carried out by developers and researchers working in heterogeneous compute environments.

Additionally oneAPI helps universities and labs deliver open source codes for a variety of applications and use cases.


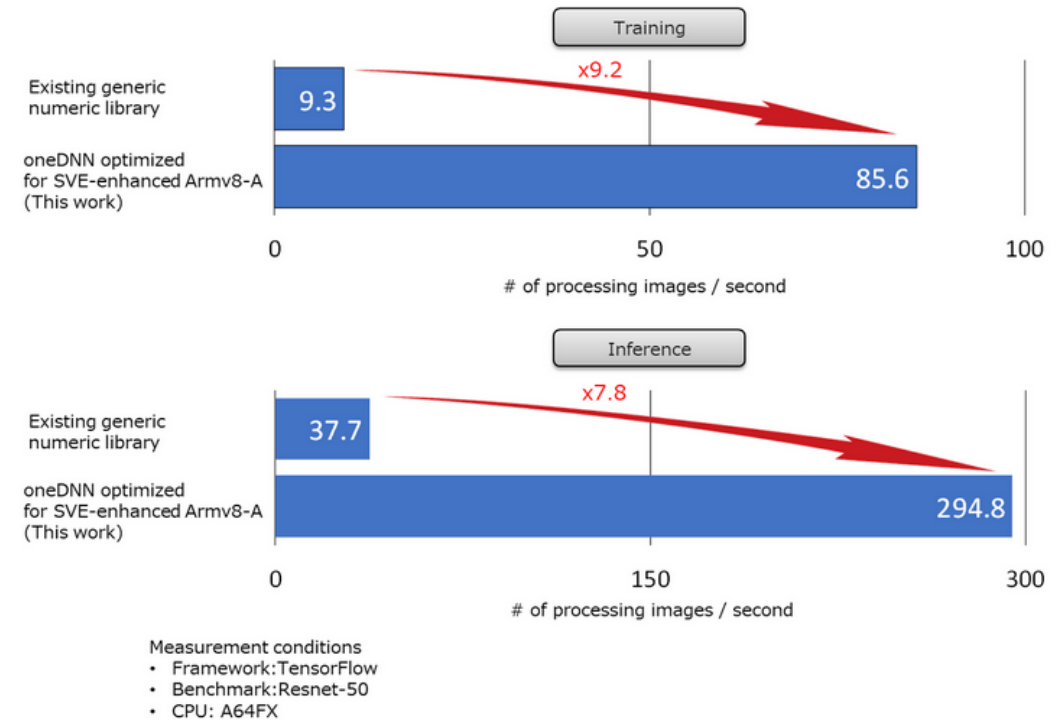
CoEs

# Case Studies

# Fujitsu Laboratories

- Adapted oneDNN library for Fugaku Arm CPU

- Achieved significant performance improvements

- Contributed to oneDNN open source project



A Deep Dive into a Deep Learning Library for the A64FX Fugaku CPU - The Development Story in the Developer's Own Words

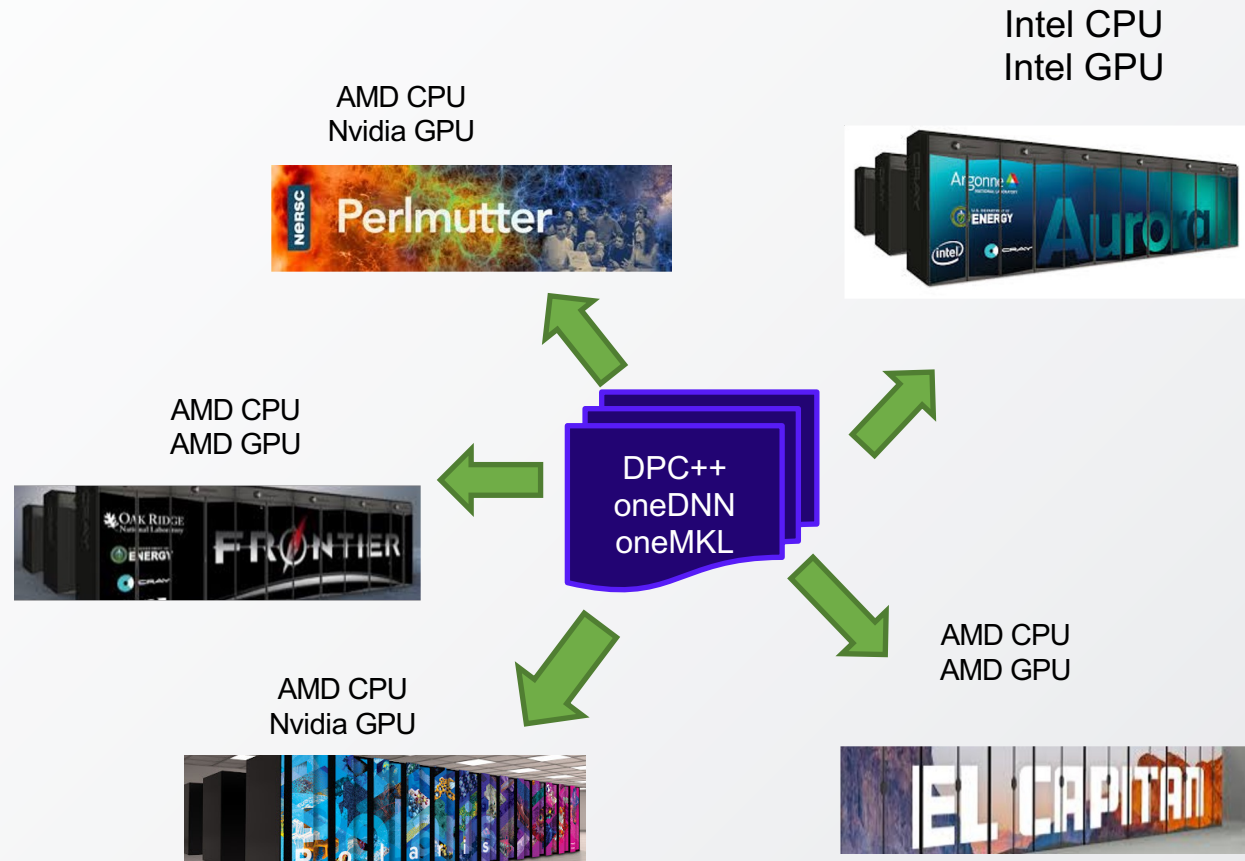https://blog.fltech.dev/entry/2020/11/19/fugaku-onednn-deep-dive-en

# Argonne, Lawrence Berkeley and Oak Ridge

- US National Laboratories are ensuring researchers can target new supercomputers using a common programming model

- Partnerships enable the SYCL implementation DPC++ and oneDNN on Intel, Nvidia and AMD GPUs

AMD CPU
Nvidia GPU


Perlmutter

Intel CPU
Intel GPU


Aurora

AMD CPU
AMD GPU


FRONTIER

DPC++
oneDNN
oneMKL

AMD CPU
Nvidia GPU

AMD CPU
AMD GPU


EL CAPITAN

# Get Involved!

Email [oneapi@codeplay.com](mailto:oneapi@codeplay.com) to get in touch and find out how to join.