UXL Foundation: Drive an Open Standard Accelerator Software Ecosystem

UXL Foundation Steering Committee Members

UXL
Unified Acceleration Foundation

Agenda

UXL Foundation introduction	Rod Burns, Steering Committee Chair
UXL Foundation Projects **oneAPI Specification** **Working Groups and SIGs**	Andy Wafaa, Arm, Open Source Working Group Chair Robert Cohn, Intel, Specification Working Group Chair Penporn Koanantakool, Google Cloud, AI SIG Chair
What's happening in the UXL Foundation?	Our members
How to contribute	Rod Burns, Steering Committee Chair

Solving the Challenge of Diverse Hardware Acceleration

CPU

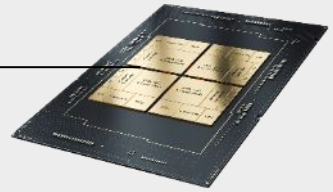

GPU

FPGA

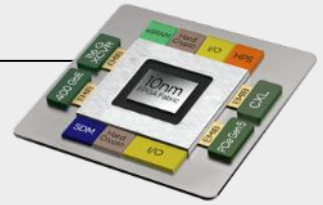

Other Accelerators

70% of developers target heterogeneous systems
that use more than one kind of processor or core[1]

Developer Challenges: Multiple Architectures, Vendors, and Programming Models

UXL
Unified Acceleration Foundation

Open, Standards-based, Multiarchitecture Programming

Video: What is oneAPI? Overview & Benefits

1. Evans Data Global Development Survey Report 23.1 2023

Commitment to Open, Scalable Acceleration

UXL
Unified Acceleration Foundation

Linux Foundation governed open industry foundation driving a vendor-neutral software ecosystem for multi-architecture accelerated computing

- oneAPI Specification
- Libraries deliver cross-vendor building blocks for developers

oneAPI

Open, Standards-based, Multi-architecture Programming

Performance | Productivity | Freedom from Vendor Lock-In

CPU

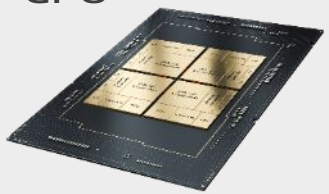

GPU

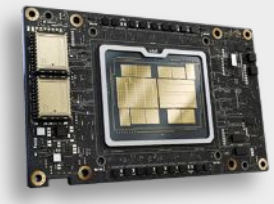

FPGA

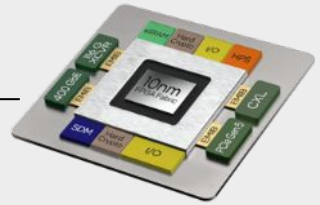

Other Accelerators

Open Standards Programming for All Accelerators

Mission

- Build a **multi-architecture multi-vendor software ecosystem** for all accelerators
- **Unify** the heterogeneous compute ecosystem **around open standards**
- Build on and expand **open source projects for accelerated computing**

Unified Acceleration Foundation

Steering Members

arm

FUJITSU

Google Cloud

Imagination

intel.

Qualcomm

SAMSUNG

vmware by Broadcom

Contributor Members

AKHETONICS

FIXSTARS Speed up your Business

AXELERA ARTIFICIAL INTELLIGENCE

CloudsAI

codasip

CODE RECKONS Science to the CORE

EMBECOSM®

GMAC Intelligence

MARKET POTENTIAL

XIANGDIXIAN COMPUTING TECHNOLOGY 象帝先计算技术

UXL Unified Acceleration Foundation

UXL Foundation Governance

Using Best Practices

Joint Development Foundation governance

SIGs: AI, Hardware, Language, Math, Safety Critical

Working Groups: Specification, Open Source

Anyone can participate in SIGs and Open Source Projects. Contributors to the Specification must sign the Membership Agreement or Non-Member Feedback Agreement.

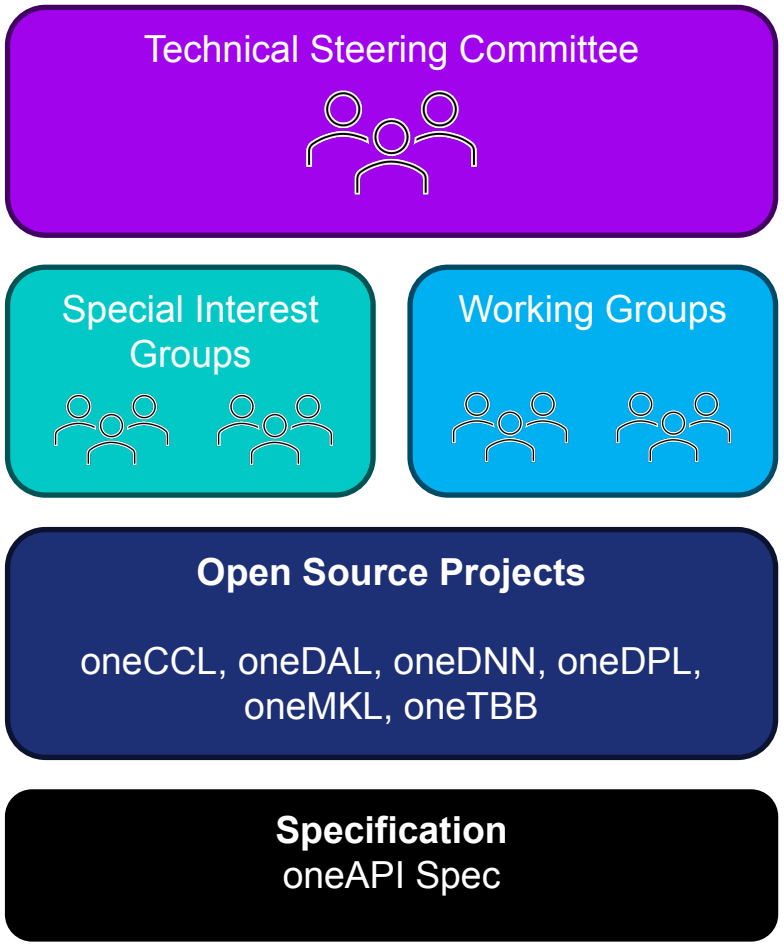

UXL
Unified Acceleration Foundation

Unified Acceleration Foundation

oneAPI Specification and Projects

KHRONOS GROUP

SYCL™

UXL Unified Acceleration Foundation

oneAPI™ Specification

- SYCL powers the UXL Foundation libraries

Heterogeneous, cross-vendor programming model

oneDPL
Data
Parallel C++ Library

oneDNN
Deep Neural
Network Library

oneCCL
Collective
Communications Library

oneDAL
Data
Analytics Library

oneTBB
Threading
Building Blocks

oneMKL
Math
Kernel Library

Building on Strong Foundations

	2019-2020	2021-2022	2023-2024
SPECIFICATION	▪ oneAPI provisional specification ▪ Technical Advisory Boards established	▪ oneAPI Specification delivered ▪ Created additional TABs ▪ Held oneAPI DevSummits	▪ Members join the UXL Foundation ▪ Specification migrated to UXL Foundation ▪ SIGs and Working Group established to coordinate feedback and work
OPEN SOURCE	▪ Open source implementations of oneAPI specification made available ▪ Initial targets for Intel processors	▪ Nvidia and AMD targets added to oneMKL and oneDNN ▪ Arm target added to onDNN ▪ Broad adoption of libraries	▪ Open Source projects migrated to UXL Foundation ▪ Arm target added to oneDAL

Projects Adopting UXL Foundation

Using libraries for cross-vendor portability

GROMACS

Project uses oneMKL math library to target multiple architectures

TensorFlow and PyTorch

Projects use oneDNN library for accelerated graph optimization

Ginkgo

Project uses oneMKL and oneDPL for optimized math and ISO C++ routines

US National Laboratories

Projects using oneMKL to deploy big science applications across supercomputers and for exascale

And many more…

Use the UXL Foundation Libraries Today

Download the project binaries via the oneAPI Base Toolkit

Or build the projects from open source

Target AMD, Arm, Intel and Nvidia processors

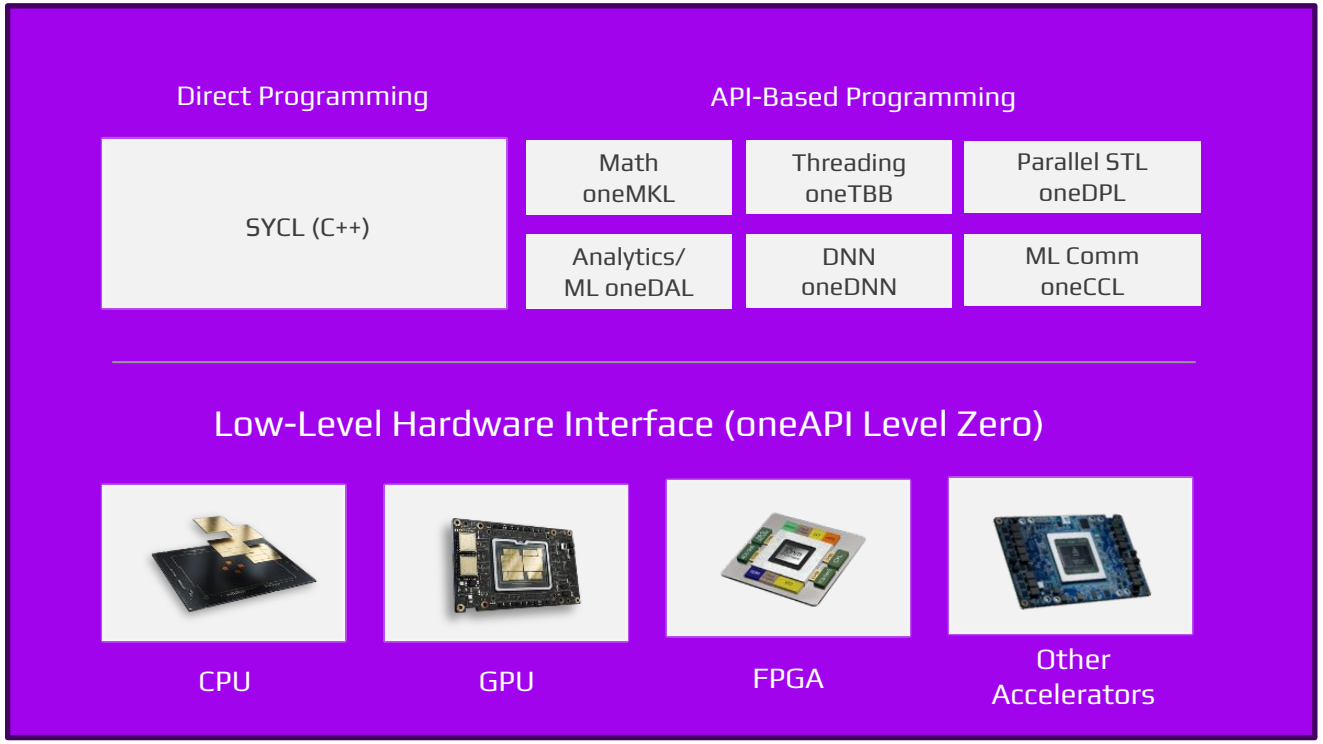

Building alliances

UXL Foundation

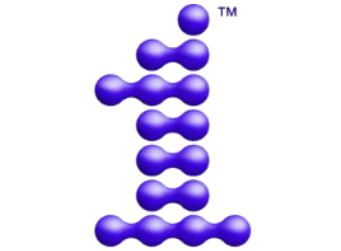

Project Dependencies

KHRONOS GROUP

SYCL™

Future Collaborations

THE AUTOWARE FOUNDATION

RISC-V®

PyTorch

LLVM Foundation

Defines the specification used by UXL Foundation projects

Integration of UXL projects and demonstrations

UXL
Unified Acceleration Foundation

UXL Foundation Working Groups

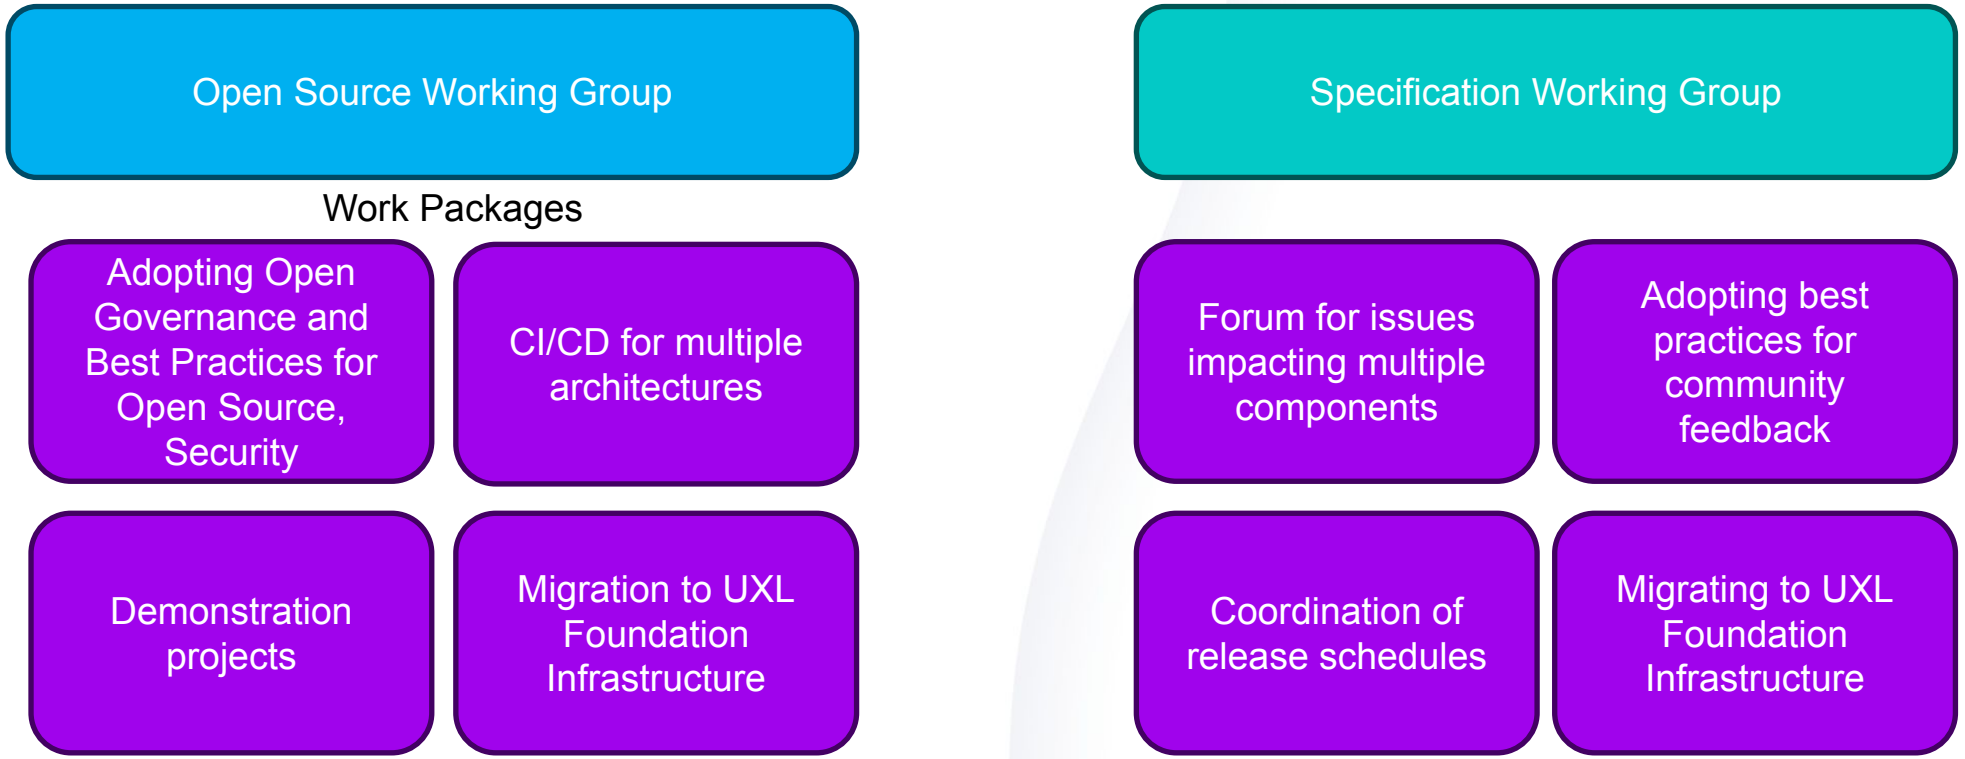

Become a member to join the Working Groups and help shape the projects

UXL
Unified Acceleration Foundation

UXL Foundation SIGs

Language

Enabling languages for accelerated computing

Hardware

Bringing more target processors to the UXL projects

Math

Optimized operators for commonly used functions

AI

Accelerated libraries and frameworks for machine learning and AI algorithms

Safety Critical

Enable or accelerate integration into safety critical systems, targeting markets such as automotive and avionics.

Fujitsu

Masahiro Doteguchi & Dr. Priyanka Sharma

Fujitsu's presence in OSS community since 2005, via open-source development in mission-critical systems and in the Supercomputer Fugaku & we are further committed to continuing with this legacy through FUJITSU-MONAKA (2nm Arm CPU)

Optimized and ported the oneDNN DL process library software for the Arm SVE instruction so that it can run at high speed on the Fugaku supercomputer.

Recent success towards porting of oneDAL on Arm to accelerate ML workloads on Arm. This is also one of the first OSS contributions to UXL foundation

Google Cloud

Google Cloud

Penporn Koanantakool

TensorFlow, JAX, and OpenXLA use oneAPI Deep Neural Network (oneDNN) library to accelerate performance on x86 and aarch64 CPUs, and Intel GPUs.

Google Cloud's Cloud HPC Toolkit includes oneAPI Math Kernel Library (oneMKL).

Imagination

Dave Murray

Imagination

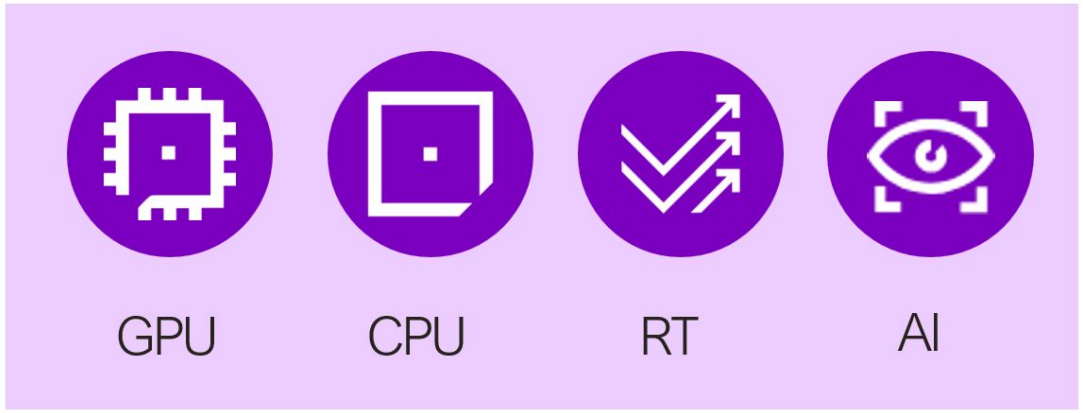

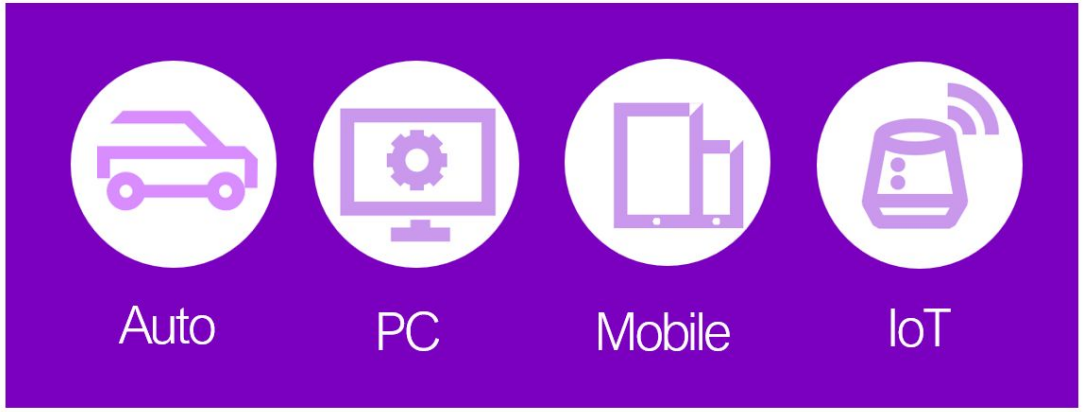

oneMKL

oneDNN

SYCLomatic

UXL
Unified Acceleration Foundation

+

THE AUTOWARE FOUNDATION

UXL
Unified Acceleration Foundation

Intel

Robert Cohn

- Unlocking the full potential for accelerated computing requires cross platform software eco-system based on open source & open standards
- Intel contributed projects to UXL for: math, AI, data parallel & distributed computing…
- With UXL and LF, we are fully embracing open governance to build strong developer communities for accelerated computing

Qualcomm

Qualcomm

Dr. Vinesh Sukumar

UXL
Unified Acceleration Foundation

Samsung

SAMSUNG

Hanwoong Jung

- Parallel programming model for memory-centric computing
  - SYCL extensions for PIM/PNM
- Deep learning compiler/runtime
  - Use oneDNN and unified runtime in oneAPI
- With UXL, we expect the open collaboration will bolster programming models and SW stacks to effectively accelerate AI/HPC applications on NPU, PIM/PNM, and RISC-V CPU/accelerator by Samsung

VMware

Ramesh Radhakrishnan

- Open Source and Standardization is a core tenant of VMware culture
- VMware offers customers choice and flexibility through our growing Private AI Ecosystem
- Streamlining interoperability in both hardware and software is key for adoption of new architectures

VMware Private AI Open Ecosystem


Falcon


Llama 2


Mistral


MPT


StarCoder


WizardML

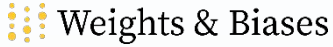

VMware Cloud Foundation™

Contribute

How to Contribute?

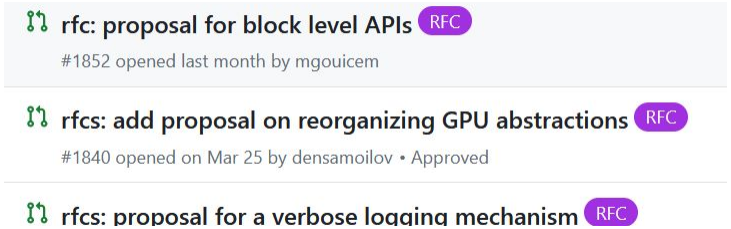

Bring your feedback

- Submit issues
- Join the SIGs
- Join the Working Groups
- Contribute to project RFCs

Contribute to projects

- New features
- Expand hardware support
- Documentation

Contribute resources

- Build infrastructure
- Testing

UXL
Unified Acceleration Foundation

Membership

Steering Member	General Member	Contributor Member
$20,000/year*	$5,000/year*	$0/year

Steering Member

- Seat on Steering Committee
- Influence direction
- Voting Rights

General Member

- Voting rights in Working Groups
- Influence project work packages
- Co-marketing

Contributor Member

- Join Working Groups
- Participate in work packages
-

* Linux Foundation Membership required

Join Us!

Check out the UXL Foundation & oneAPI specification

UXL
Unified Acceleration Foundation

- UXLFoundation.org
- oneAPI.io

Join Our Mailing Lists

Join Slack

Q&A