Bring your code to RISC-V accelerators with SYCL

Charles Macfarlane

27th June, 2024

codeplay

Enabling AI & HPC To Be Open, Safe & Accessible To All

Established 2002 in **Edinburgh, Scotland**.

Grown successfully to around 100 employees.

In 2022, we became a **wholly owned subsidiary** of Intel.

Committed to expanding the **open ecosystem** for heterogeneous computing.

Through our involvement in oneAPI and SYCL governance, we help to **maintain and develop** open standards.

Developing at the forefront of **cutting-edge research**.

Currently involved in two research projects - **SYCLOPS** and **AERO**, both funded by the Horizon Europe Project.

What is SYCL and oneAPI?

and why is it necessary?

The Demand for Multiple Architectures

Data-centric computing continues to rise in demand, driven by the exponential growth in generative AI and accelerated computing

This work requires increasingly efficient performance, with workloads now spread across new and diverse hardware – GPUs, CPUs, FPGAs and specialist AI accelerators, including **RISC-V**

48% of developers target heterogeneous systems that use more than one kind of processor or core[1]

[1] Evans Data Global Development Survey Report 22.1, June 2022

The Challenges of Multiple Architectures

This rising demand for diverse hardware creates several challenges

- No common programming language or APIs
- Inconsistent tool support across platforms
- Each platform requires unique software investment

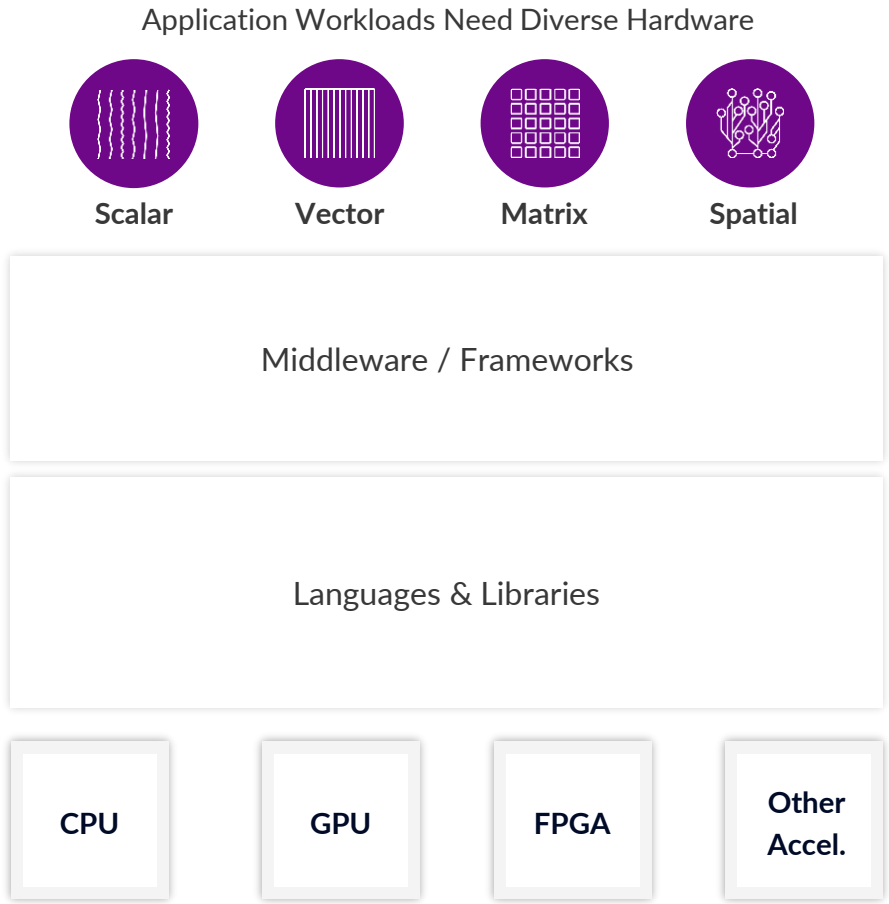

Introducing **oneAPI**

Unified programming to simplify development across diverse architectures

- Unified and simplified language and libraries for expressing parallelism
- Uncompromised native high-level language performance
- Based on ISO C++ and the Khronos-maintained, SYCL open standard
- Interoperable with existing HPC programming models

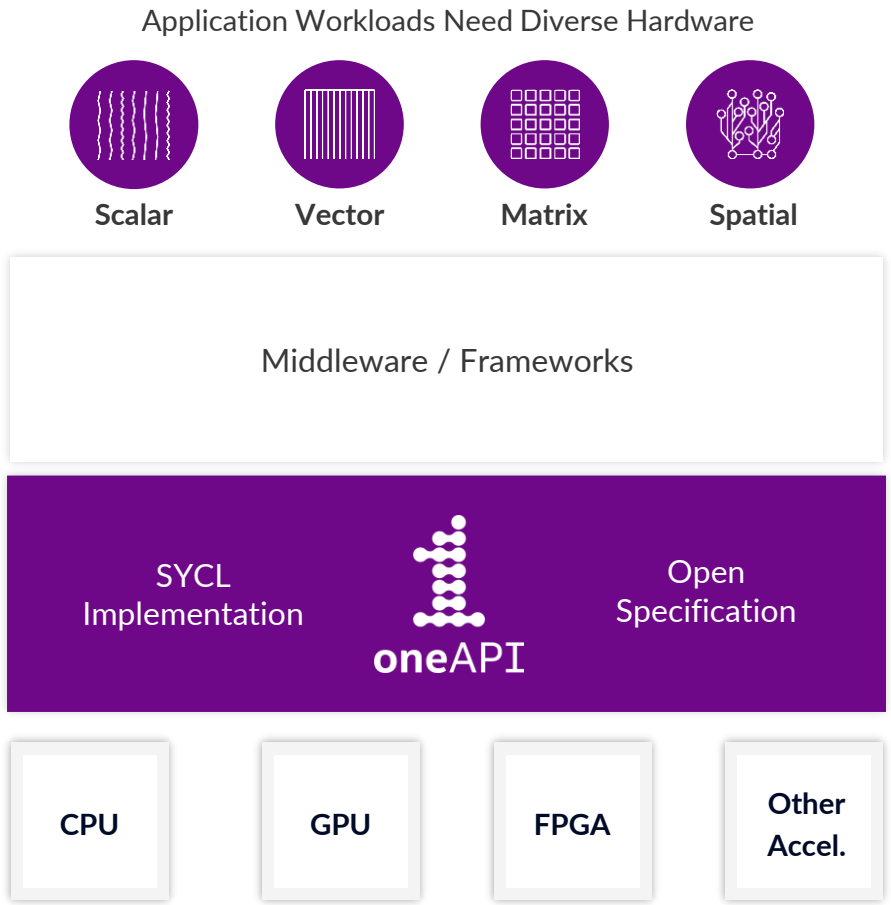

oneAPI for NVIDIA and AMD GPUs

Codeplay's plugins add support for **NVIDIA** and **AMD** GPUs to the Intel oneAPI Base Toolkit

Develop code using SYCL and run on AMD, Intel and NVIDIA GPUs

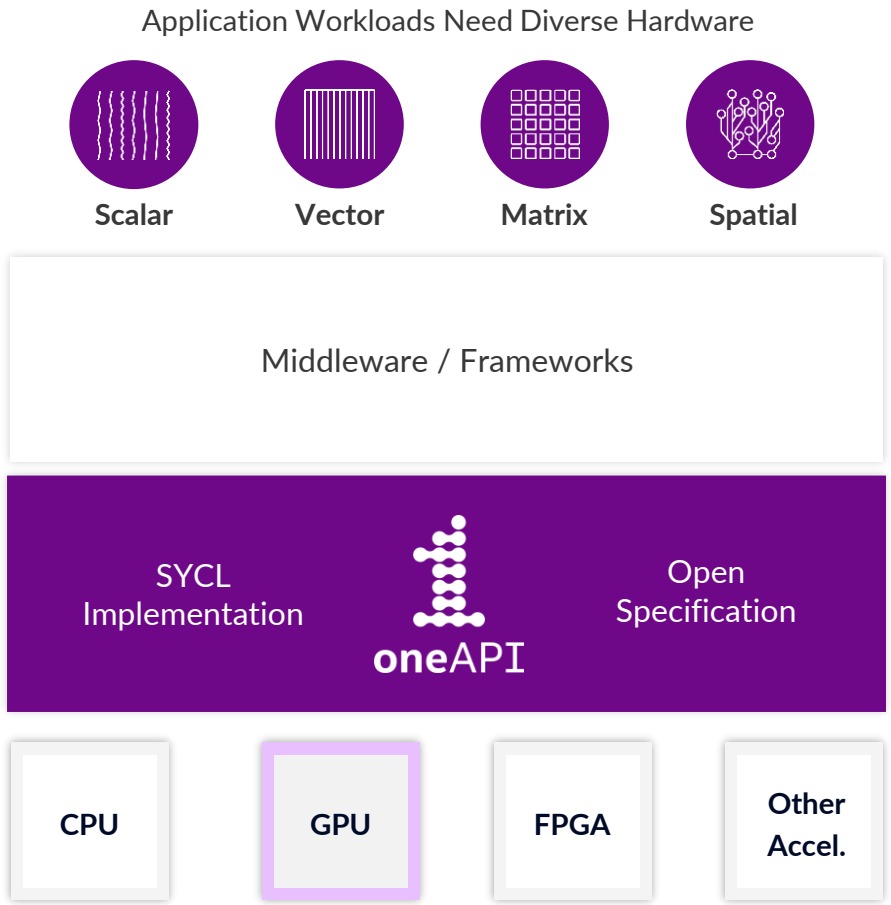

The oneAPI Construction Kit

The oneAPI Construction Kit brings **SYCL** and **oneAPI** to new and specialist accelerators - such as **RISC-V**

The oneAPI Construction Kit works by enabling the CPU to **offload compute-intensive kernels** to the **custom accelerator**

This is done with a **single, open standards programming language**, SYCL

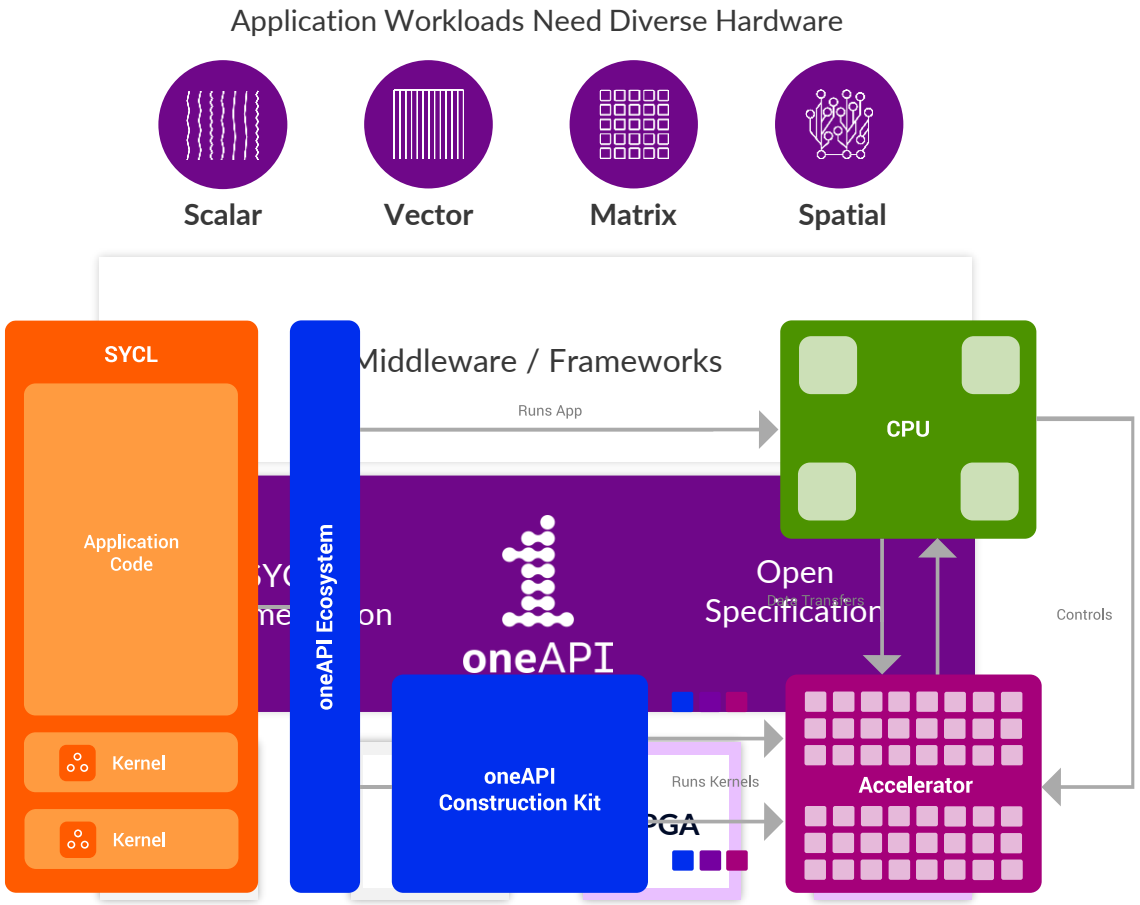

The oneAPI Construction Kit & RISC-V

bringing open standards to new hardware

Vision of the oneAPI Construction Kit

Open-Standards, Driven by Industry Collaboration

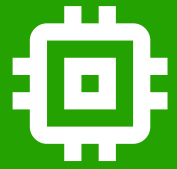

Enable SYCL and oneAPI for your RISC-V Accelerator

Take Advantage of a Huge Existing Ecosystem

Case Study
The Metis AIPU

- AI accelerator processor from Axelera
- Designed for edge computing
- Embracing open standards for programming

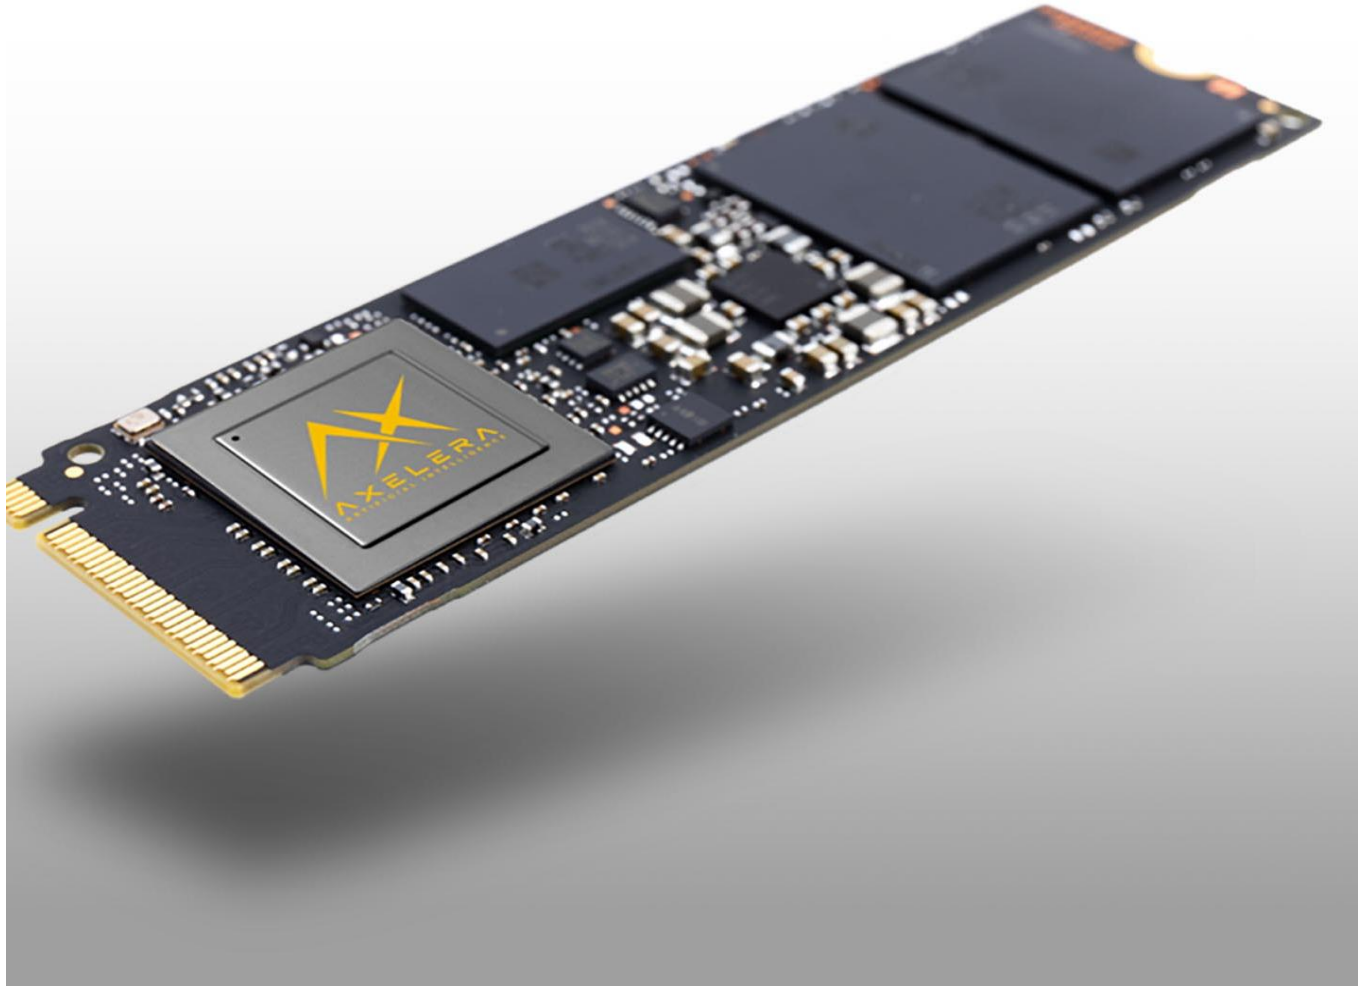

Open Standards Software

The oneAPI Construction kit provides the foundations to enable a software ecosystem

- Direct parallel programming in C++ with SYCL
- Optimized libraries for commonly used domains such as BLAS, Fourier Transform and neural networks
- Frameworks and Applications can be used without minimal changes

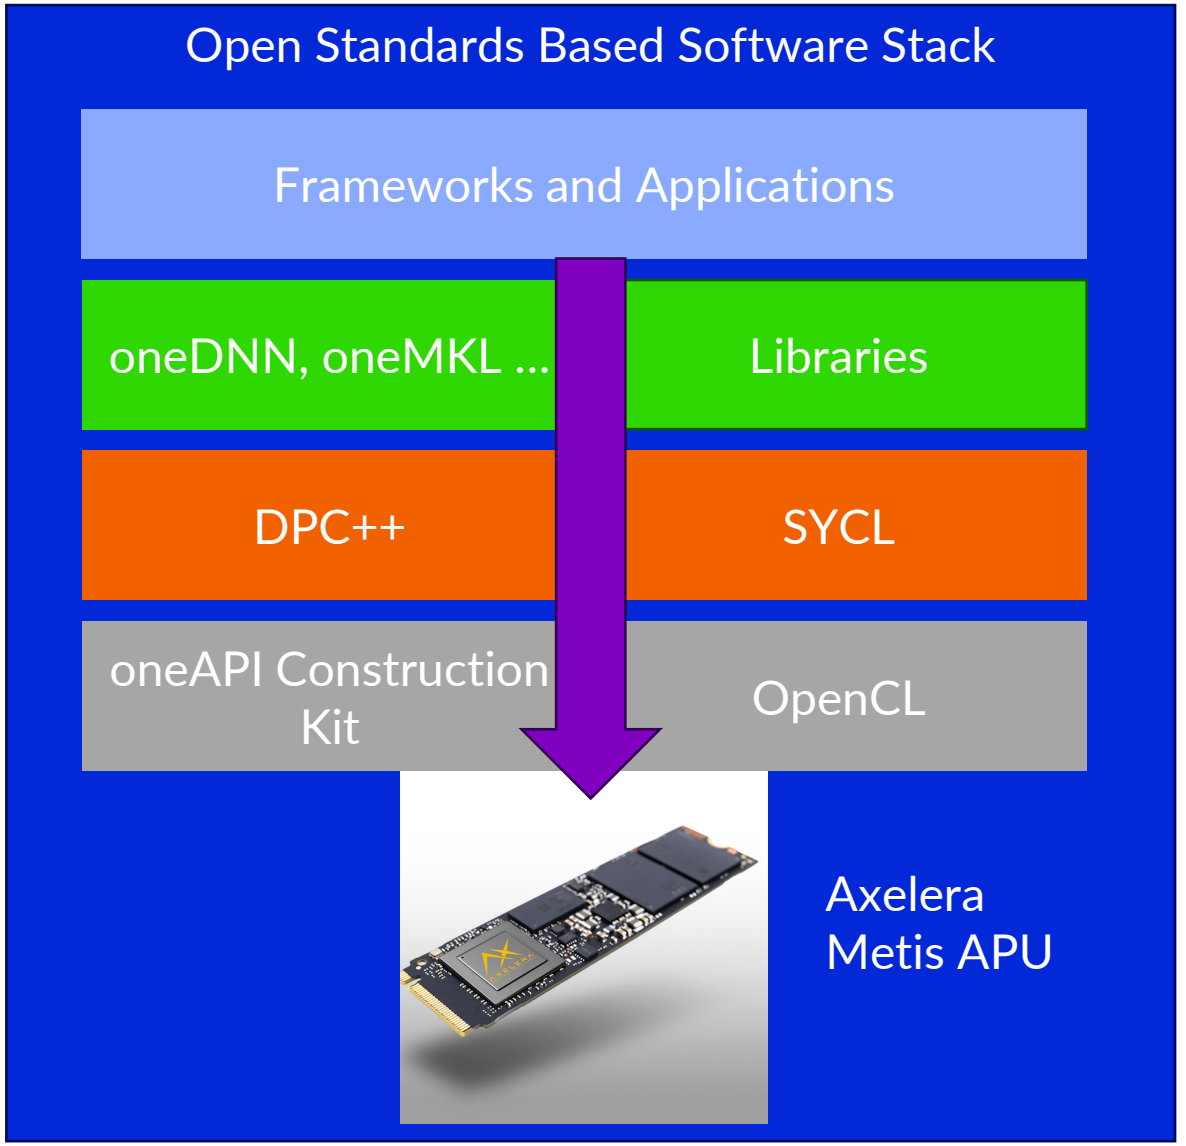

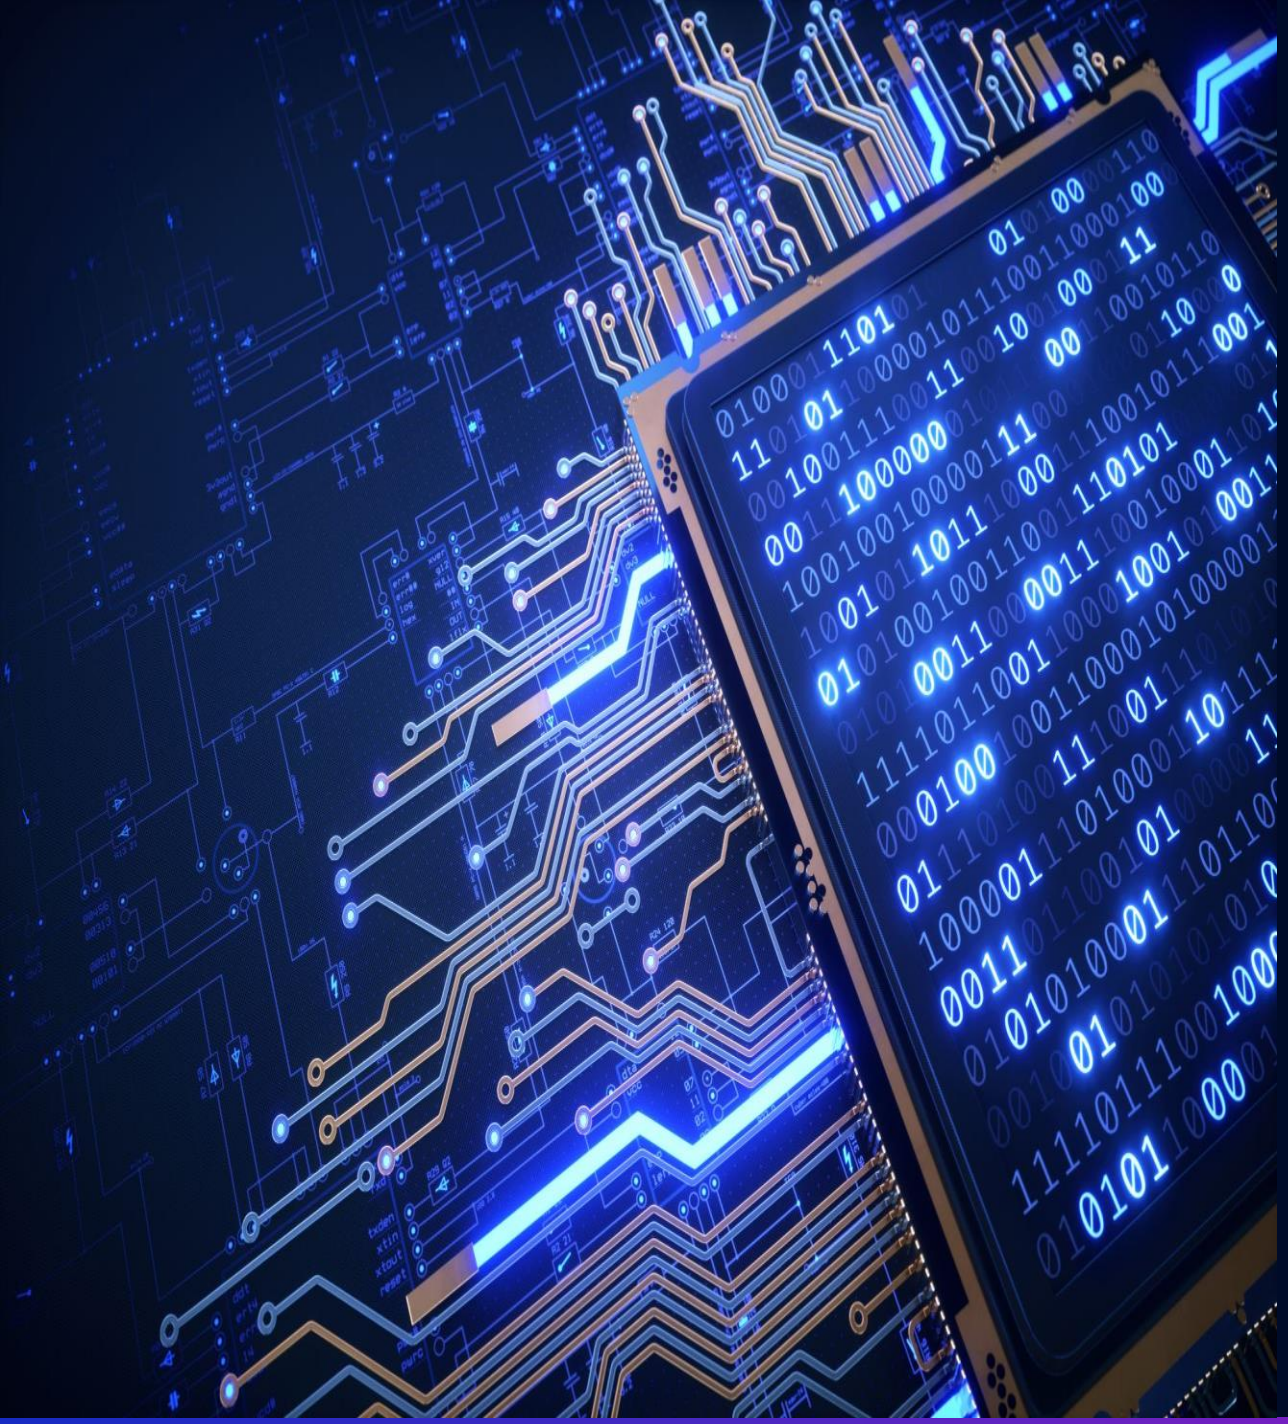

“

The integration with the oneAPI Construction Kit was straightforward, making it quick and easy to bring the whole open standards oneAPI ecosystem to Metis AIPU.

Manuel Mohr, Staff Software Engineer of Axelera

The Unified Acceleration Foundation

goals of the foundation, and how you can get involved

Unified Acceleration Foundation (UXL)

UXL
Unified Acceleration Foundation

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**

Use case focus: AI, HPC, Edge AI and Edge Compute

UXL
Unified Acceleration Foundation

UXL Foundation Members

Steering Members

arm
FUJITSU
Google Cloud
Imagination
intel
Qualcomm
SAMSUNG
vmware by Broadcom

Contributor Members

AKHETONICS
AXELERA ARTIFICIAL INTELLIGENCE
CloudsAI
codasip
CODE RECKONS Science to the CORE
EMBECOSM®
FIXSTARS Speed up your Business
GMAC Intelligence
InspireSemi™
MARKET POTENTIAL
UNTETHER AI
XIANGDIXIAN COMPUTING TECHNOLOGY 象帝先计算技术

oneAPI Specification and Projects

- **Initial contribution:** oneAPI Specification & Open Source

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

Khronos Group and UXL Foundation Liaison Agreement

UXL Foundation and Khronos Collaborate on the SYCL Open Standard for C++ Programming of AI, HPC and Safety-Critical Systems

June 10, 2024 by The Khronos Group and The UXL Foundation sycl

In a world where AI, HPC and Safety-Critical acceleration is shifting toward heterogeneous architectures that integrate

- Cross participation of members
- Assigned executive liaisons for Open Source and Safety Critical activities
- Exchange of requirements

Pioneering with RISC-V & SYCL

SYCLOPS

Funded by the European Union

UK Research and Innovation

Bringing together the SYCL and RISC-V Open Standards

The SYCLOPS Consortium

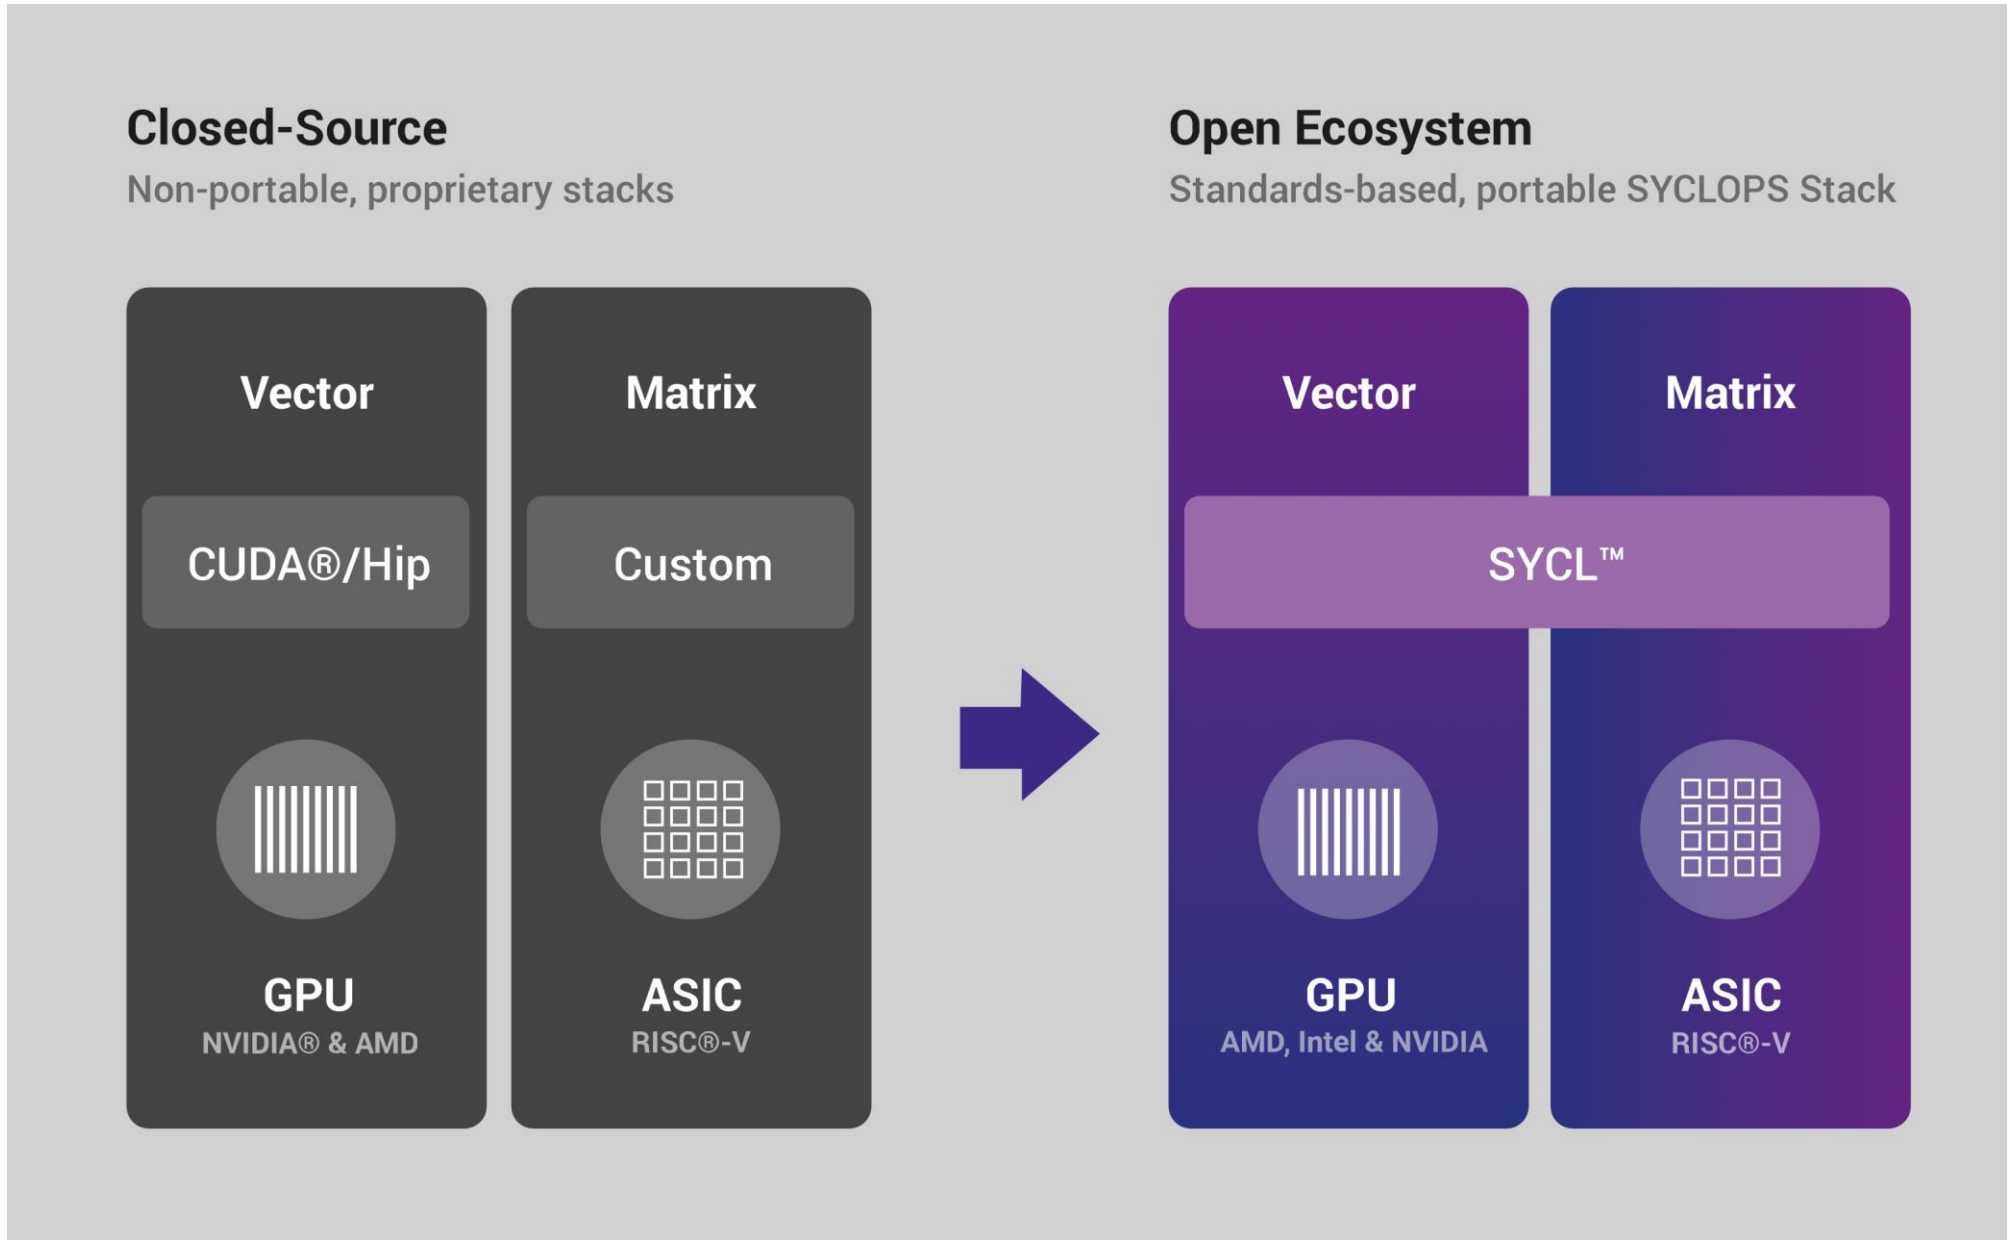
Closed-Source
Non-portable, proprietary stacks
Vector
Matrix
CUDA®/Hip
Custom
GPU
NVIDIA® & AMD
ASIC
RISC®-V
Open Ecosystem
Standards-based, portable SYCLOPS Stack
Vector
Matrix
SYCL™
GPU
AMD, Intel & NVIDIA
ASIC
RISC®-V

Funded by the European Union
UK Research and Innovation

AERO

The AERO Consortium

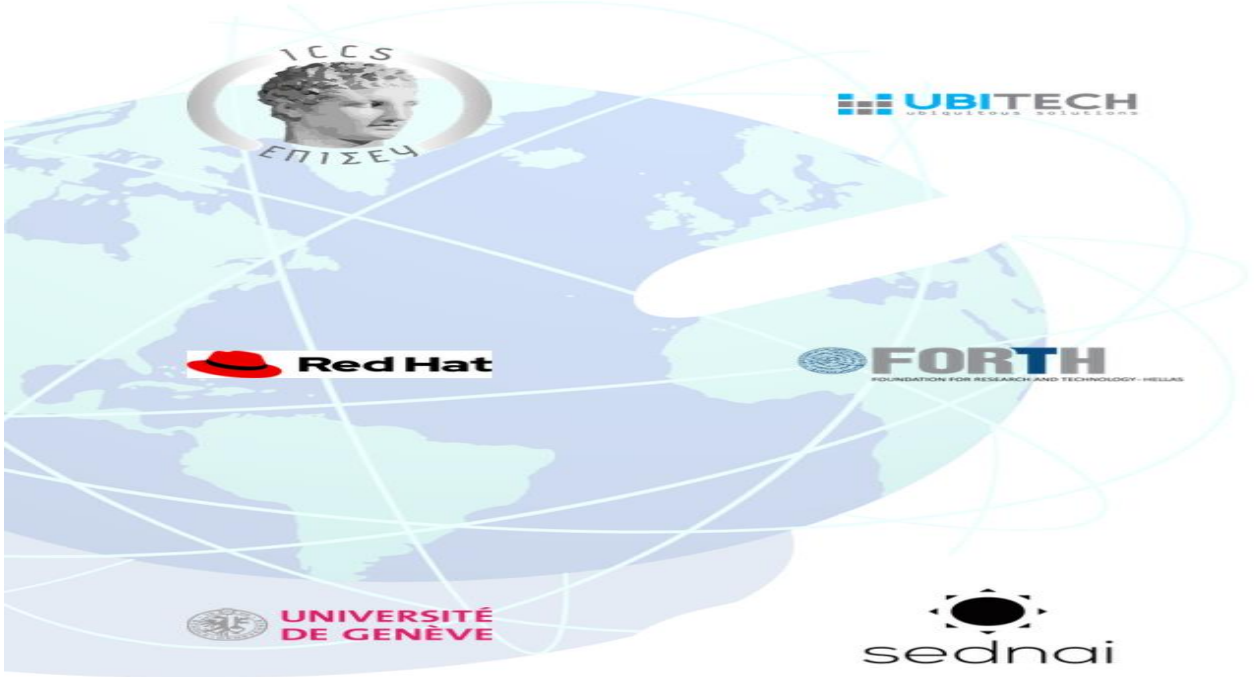

codeplay

Social Media

Don’t forget to follow us for the latest updates!

@codeplaysoft

@codeplaysoftware

codeplay-software

codeplay-software

Disclaimers

A wee bit of legal

Performance varies by use, configuration and other factors.

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.

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

UXL
Unified Acceleration Foundation

Unified Acceleration Foundation

Scan QR code or visit **oneapi.io**