



Unified Acceleration (UXL) Foundation

2024 Annual Report

uxlfoundation.org

Introduction

3

Mission & Leadership

4

2024 Milestones

5

Our Members

6

Quotes from our Members

7

Working Groups

8

Special Interest Groups (SIGs)

10

Technical Leadership

11

UXL Foundation Ecosystem

12

Events & Outreach

13

Outlook

14

Introduction



It's been around 20 years since the emergence of Linux and rapid growth of the data centre began, when increasing amounts of computing power were made available to organizations around the world in what we now call the cloud. The advent of cloud computing made it possible for companies to scale in ways that were not possible before. We are currently seeing a similar explosion in growth, but now cloud and edge computing are taking advantage of highly parallel processors to perform complex calculations. The top supercomputers are now dominated by architectures that predominantly use GPUs. Data centres and embedded/edge systems are integrating CPUs and GPUs from multiple vendors alongside specialised processors for domains such as AI. This change has reached all industry verticals from finance and healthcare through to autonomous vehicles and manufacturing.

What this means is that there is a growing demand for software

acceleration using architectures with accelerators from multiple vendors.

The UXL Foundation was formed to govern and evolve an open standards-based way to develop software for these processors regardless of vendor, and to deliver the building blocks needed to accelerate the vast majority of applications.

The inaugural year of the UXL Foundation has seen us establish a growing member base, adopt best practices for open source development, and extend projects to bring new target processors and features. Participation is increasing across the whole community through our Special Interest Group (SIG) meetings with a range of topics covered and feedback received from the community to drive the UXL Foundation in the right direction for everyone. The Working Groups have been busily tackling some of the key areas needed for the broad collaboration that is required.

This report will highlight the work and progress that has been made by the UXL Foundation during 2024.

Steering Committee

Chairperson



Andrew Wafaa
Arm



Ramesh Radhakrishnan
Broadcom



Rod Burns
Codeplay



Masahiro Doteguchi
Fujitsu



Dr Priyanka Sharma
Fujitsu



Evgeny Drapkin
GE Healthcare



Penporn Koanantakool
Google



Dave Murray
Imagination



John Melonakos
Intel



Manuj Sabharwal
Qualcomm



Hanwoong Jung
Samsung

Mission & Leadership



Building a multi-architecture and multi-vendor software ecosystem for all accelerators



Unifying the heterogeneous compute ecosystem around open standards



Building on and expanding open-source projects for accelerated computing

We have had many conversations with members of the community, even before the formation of the foundation, over the past two years and there is a common theme. Software developers are looking for a way to run their code with high levels of performance on the latest and future processors; and hardware developers want to enable all the programming frameworks needed for their latest processors. These common interests drive our mission, and our focus in 2024 has been to ensure that everything is in place to meet the requirements of our members and the community. This is done alongside the rapid development of new features and maintenance of the projects.

The first work area focus has been to ensure that the open source projects and the specification are set up to


seamlessly allow contributions from members across the community, regardless of where they work or who they are. The UXL Foundation projects adopted a scorecard of behaviours defined by leaders in the open source community, and these form best practices for the projects. We have used the OpenSSF and CHAOSS projects to guide us.

A second work area is to set up the mechanisms required to bring feedback and requirements to the UXL Foundation and its projects; and subsequently enable the community to act on this feedback. We have been doing this through work packages led by different members of the community.

2024 Milestones



30 Members & Growing
2 Affiliate Partners



- > 300 Attendees at SIG Meetings
- > 450 Views of UXL Webinar
- > 400 attendees for oneAPI DevSummit hosted by UXL Foundation



5 Active Work Packages

Our Members

The foundation was formed in September 2023 with 8 Steering Members and since then we have added 1 new Steering Member, 3 new General members and 21 Contributor members, as of December 2024. Our newest Steering Member is GE HealthCare and the General Members are Micron, CDAC and SK Hynix.



If you are interested in joining the UXL Foundation visit the website or email membership@uxlfoundation.org.

Our Members



						Steering

General					
---------	--	--	--	--	--

						Contributing

Affiliate					
-----------	--	--	--	--	--

Quotes from our Members



Intel is excited to support and expand UXL through our active involvement in the special interest groups, contribution to the open specifications, and participation in open-source projects. Our Intel® oneAPI products implement key elements of the UXL specification. Through some of that work over the past year, we've made it possible for the newest Intel technology to be used anywhere from personal laptops to some of the most powerful supercomputers in the world. This means a developer can create their software on the latest Intel processors including the world's leading AI supercomputer Aurora using more than 63,000 Intel Data Center GPU Max Series processors. Achieving this level of scale is only possible because we build on the open source foundation established by UXL that allows multiple vendors to collaborate and expand the capabilities of heterogeneous computing.

Paul Petersen (Intel)

Through the UXL Foundation, Samsung aims to build a robust software ecosystem for breakthrough memory technologies, such as PIM, and for this purpose, have performed SYCL vendor extensions and are readying Khronos extensions. We will continue our efforts into next year to effectively expand the basis for oneAPI projects so that PIM technology can be used not only in SYCL but also in oneDNN and oneMKL.

Seungwon Lee (Samsung)

Broadcom is partnering with industry chip and software makers to make it easier for its large installed base of enterprise customers to adopt AI. The efforts in UXL to build a multi-vendor software ecosystem for all accelerators and unify the heterogeneous computing sphere through open standards are key for wider adoption of AI in the enterprise.

Ramesh Radhakrishnan (Broadcom)

During the past year, Arm has led the open source working group and provided expertise for oneDNN and other libraries in response to growing aarch64 adoption across sectors. In the coming year, Arm will bring new resources and infrastructure to the UXL Foundation, underlining our commitment to open-source software and standards and to simplifying every aspect of heterogeneous compute for developers.

Andy Wafaa (Arm)

oneAPI is fundamental to the success of accelerated computing on heterogeneous hardware. Imagination is actively developing a software stack based on oneAPI for AI at the edge and promoting the wider adoption of the standard as the solution for cross-platform edge-accelerated computing. Through our position on the UXL Foundation Steering Committee we are engaging with strategic partners on proof-of-concept projects as well as developing benchmarks and solidifying developer migration flows with SYCLomatic.

Dave Murray (Imagination Technology)

Fujitsu, a founding Steering Committee member of the UXL Foundation, drives innovation in open-source projects such as oneDAL and oneDNN, for enhancing performance and energy efficiency, ensuring scalable AI solutions that aligns with our vision for sustainable high-performance computing, showcasing our commitment to democratizing AI and advancing global digital transformation.

Dr Priyanka Sharma (Fujitsu)

UXL Working Groups (1)



The UXL Foundation Working Groups are led by Andy Wafaa (Arm) and John Melonakos (Intel). We are deeply grateful for Robert Cohn's (Intel) contributions during the past year and wish him all the best in his new role.

During the first year the Working Group has initiated some key work packages and implemented a lightweight system for organising and tracking these packages. The key areas of focus have been on establishing the best practices for the open source projects, security best practices, public build infrastructure and migrating the projects into the UXL Foundation infrastructure. This section will outline the achievements and challenges from these groups.

The Working Group oversees 7 projects - oneCCL, oneDAL, oneDNN, oneDPL, oneMKL, oneTBB and oneAPI Construction Kit alongside the oneAPI specification. Three key areas of activity around the projects during this year are outlined in the sections below.


Open Source Best Practices

A major goal for the foundation is to increase contributions to the open source projects from a wide variety of individuals and organisations. To ensure the projects are in the best place for this to happen and consistent practices are in place, the Working Group established a work package for "open source best practices" ...

based on the guidance defined by the CHAOSS project and Red Hat. All projects were audited against a set of criteria, including the presence of open decision making, sufficient documentation and clearly defined leadership. This was a successful activity and all of the criteria has now been met. This puts the projects in the best position to accept contributions. We plan to monitor these best practices through 2025 to ensure that they are meeting the needs of the community.

Security Best Practices

Security was identified by our members as a crucial area for the foundation projects, and it was necessary to establish processes and best practices across them. Indeed many organizations can only use open source projects if they follow specific security practices. A work package was set up to ensure consistency and best practices were in place. The OpenSSF scorecard has been used for this activity and the scores for projects is now above the average for the industry across open source projects.



UXL Working Groups (2)



The group will continue to work in this area to close gaps and monitor the channels for security incidents. We have not yet set up a formal security team but this is a consideration for 2025. We would welcome the involvement of security experts from our membership and the broader community.

Build Infrastructure

Our goal as a foundation is ambitious, and we would like to test and build as many architectures as we are able to. Currently the public infrastructure for this is limited and projects are relying on limited resources

alongside infrastructure provided by internal networks.

A work package was established to set up the initial infrastructure for CI and testing and to document how this can be done. As an example the oneDNN project is already using infrastructure from Arm for CI, and next year we expect to add more. It's crucial that the community gets involved in making more resources available to the foundation. We invite members to join the Working Group to discuss how to expand the availability of resources for the projects. The foundation also intends to fund some additional infrastructure next year.

Track Our Progress & Contribute

You can keep track of our progress by visiting the GitHub backlog.

github.com/orgs/uxlfoundation/projects/3/views/1

The screenshot shows a GitHub backlog for the 'Security Work Package' project. The backlog is organized into columns for 'No Status', 'Todo', 'In Progress', and 'Done'. It lists several tasks related to oneDNN, such as 'oneDNN Test CodeQL for non-C/C++ code' and 'oneDNN Implement Fuzzing [double-check]', with their respective due dates and assignees.

Status	Task	Assignee	Due Date
No Status	oneDNN		
Todo	oneDNN Test CodeQL for non-C/C++ code	oneDNN	Q4 2024
In Progress	oneDNN Test Coverage on open source repos (C/C++ code)	oneDNN	Q2 2024
In Progress	oneDNN Implement Fuzzing [double-check]	oneDNN	Q4 2024
Done	oneDNN Adopt Vulnerability Policy (Security.md)	oneDNN	
Done	oneDNN Enable OpenSSF Scorecard GitHub Badge	oneDNN	Q3 2024

Special Interest Groups (SIGs)



The Special Interest Groups meet on at minimum a quarterly basis, with the notes and presentations available through the UXL Foundation GitHub repository. Here are some highlighted activities from the SIGs this year.

Artificial Intelligence SIG

The team at Fujitsu integrated openBLAS and openRNG backends to the oneDAL project. Adding new architecture to the oneDNN project was made easier through a generic backend for GPUs. oneDNN performance was improved on AArch64 through collaboration between Arm and the project maintainers. oneDAL and scikit-learn integration was achieved through a collaboration between Fujitsu and Intel. oneDNN integration with Llama.cpp and PyTorch brought large language model acceleration on GPUs.

Hardware SIG

The oneAPI Construction Kit was presented to the group, and later it was agreed it would be contributed to the UXL Foundation. Specification changes and new APIs were presented for Level Zero. Imagination Technologies presented their customer journey and how they use software alongside hardware design. The group is looking to investigate and discuss how new processors can integrate with the oneAPI specification and projects.

Language SIG

Integration of the UXL projects and specification with different languages was a topic, with a Python integration presented.

The group discussed new SYCL extensions for SYCL Graphs and memory technologies such as PIM.

In order to help the community understand how the layers of languages and instructions fit together, a presentation and discussion was held on SPIR-V extensions required for the DPC++ compiler.

Math SIG

There was a presentation from developers working on the GROMACS project and their use of SYCL alongside oneMKL. The ArrayFire projects were introduced with a potential for closer collaboration with the UXL Foundation. New Sparse and DFT APIs were presented and discussed for inclusion in the oneAPI specification. The group was asked for feedback on the re-structuring of the oneMKL project, and renaming to oneMath.

Technical Leadership



As a maintainer for the oneMKL project I work on various areas including Nvidia targets. In the past year I have worked with others to contribute two open source backends for the project covering SPARSE, BLAS and FFT math domains. I hope to work with more of the community to expand the number of open source backends that are available in the future.

Romain Biessy
Codeplay, oneMKL Code Owner

Collaborating closely with the Fujitsu team helped us to overcome the challenges of accepting the first major contributions to the project. As result the oneDAL project has an Arm target and this has helped Fujitsu achieve performance on ARM up to 40x compared to before. We will continue to work with the community to further extend oneDAL.

Nikolay Petrov
Intel, oneDAL Code Owner

During the year we have seen increasing activity on the open source project repositories from different organisations. Some projects have also formalised representation from multiple organisations as code maintainers and defined roles. It is possible to participate actively in all of our projects with different roles.

Now that all projects have adopted a transparent Request for Comments (RFC) process, open decision making is now available to the whole community, and regular information is shared on the current and future plans. A future goal is to publish public roadmaps.

The Fujitsu team has been working closely with the oneDAL project maintainers to add support for further open source components. With a goal to accelerate AI workloads on the Monaka supercomputer, the team adapted the oneDAL project and integrated this with openBLAS and openRNG.

The performance results were 40x compared to the previously used ML algorithms. You can read more of the details in their blog post.

The oneDNN project welcomed new maintainers to the project from Arm, Amazon and Codeplay alongside the core team from Intel. The oneDAL project welcomed maintainers from Fujitsu and Rivos. Alongside this, the project now includes a generic GPU back end that is open source and is intended to enable new targets for the oneDNN project.

The oneMKL project has 30 years of history, coming from the MKL product developed by Intel, and this year the project is evolving to meet the needs of the community. It is adapting naming and interfaces and integrating open source components for BLAS and FFT. Next year it is expected there will be further open source contributions to the project for other math components.

UXL Project Ecosystem



Artificial Intelligence

The rapid growth in use of AI continues and there are an increasing number of frameworks taking advantage of the UXL Foundation projects.

The PyTorch has a new backend that uses oneMKL and oneDNN alongside SYCL to accelerate AI workloads on Intel GPUs. But this open standards based implementation is designed to be used on any target that supports oneAPI and SYCL and we welcome our community to experiment with this using other GPUs and accelerators.

Llama.cpp, OpenXLA, TensorFlow and SciKit learn are also making usage of oneDNN, oneMKL and oneDAL to accelerate AI workloads on a range of hardware targets. As new projects in AI emerge we invite them to make use of the open standards based path to hardware provided by the oneAPI and SYCL specifications.

High Performance Computing

The HPC community has been adopting libraries based on the oneAPI specification for years now, and there is broad adoption within the scientific research community.

The GROMACS project is one of the most popular projects in the world and is achieving good levels of performance across architectures through the use of oneMKL and SYCL. NWChemX, Ginkgo, AmRex, TMM and others are being used to deploy researchers code across supercomputers using GPUs from AMD, Intel and Nvidia.

Alongside this there are commercial deployments of software making good use of the capabilities of the UXL Foundation projects, and this is crucially bringing cross-architecture capabilities to codebases.



UXL Foundation GitHub



UXL Foundation Get Started Guide



Open Source Working Group Repository



The UXL Foundation has participated in and hosted many events this year. Our first was a webinar hosted by the Linux Foundation. The oneAPI Dev Summit hosted by the UXL Foundation was our first official event and this brought together over 400 of our members and the broader community to share knowledge and discuss their latest work. This illustrated the importance of open source and industry collaboration with panel discussions on the past, present and future of open standards and open source.

We also held our first in person event at the SC24 conference where our Steering members shared their work and outlook for the foundation.

This was a highlight of the year as it was the first opportunity outside virtual calls and events.

The UXL Foundation also was represented at the Open Source Summits, the Linux Foundation Member Summit, as well as Embedded World, ISC, the PyTorch conference and RISC-V Summits.

During the year there were at least 16 presentations, exhibitions or dedicated events where we were represented thanks to our members.

Our marketing committee has worked together this year to expand our promotion, social media activities and publish key materials like templates and presentation decks.

Outlook



The first year of the UXL Foundation has involved a lot of work to ensure that our members and the broader community are able to effectively contribute to and evolve our open source projects and specification. This means open decision making, defined roles for projects, open communications and transparent build infrastructure amongst all the other required best practices to ensure success.

Our progress in this area has brought significant improvements that are already enabling the sort of contributions we would like to see. The teams at Fujitsu and Arm are making contributions, not only with code, but also leadership roles in the projects. There are strong expectations that more of our members follow suit.

The Autoware alliance has brought together a group of our members, working collectively on delivering a demonstration of the Open Autonomous Driving project using some of the UXL Foundation libraries to accelerate parts of this on a GPU. I look forward to seeing this demonstration in action and how we can move towards usage in real vehicles.

During 2024 we kicked off the first UXL Foundation work packages, driven by the Working Groups, and I expect to see the number of work packages, RFCs and technical activities increase in 2025. There are many opportunities for our members, and the broader community, to lead on. Some of the areas I envisage activities to thrive include RISC-V (CPUs and accelerators), a standards based backend for PyTorch using UXL libraries, open source math kernel contributions, and use of the generic backend for oneDNN.

I thank all our members for their efforts this year, and for being part of this bold movement to build a multi-architecture, multi-vendor software ecosystem for all accelerators using open standards and open source.

I look forward to working with you further in the coming year.



Rod Burns

Steering Committee Chair
rod@codeplay.com