Preparing for Science in the Exascale Era

As the future home to Aurora, the ALCF has been ramping up efforts to ready researchers and developers for science at exascale.

Science on day one. That's the goal for each new supercomputer introduced at the ALCF, including its upcoming exascale system, Aurora. But how do you prepare for a machine before it is even built?

With Aurora's arrival drawing closer, the ALCF has been ramping up its efforts to lay the groundwork for science in the exascale era.

The facility's preparatory activities include continued collaborations with Aurora vendors Intel and HPE on the testing and development of various components to ensure they meet the requirements of the ALCF user community; partnering with DOE's Exascale Computing Project (ECP) to develop a robust exascale computing ecosystem; facilitating the Aurora Early Science Program (ESP) to prepare key applications and software for the system; deploying early hardware and software tools to provide an environment for exascale code development; and hosting training events to educate the research community on various exascale tools and technologies.

The Aurora system's exaflop of performance combined with an ability to handle both traditional HPC and emerging AI workloads will enable researchers to address some of the world's most challenging scientific problems. Early science projects range from mapping the brain to simulating the universe at extreme scales to accelerating the discovery of new drugs and functional materials.

Aurora's innovative design will be based on Intel's Xeon Scalable processors, high-performance Intel X^{e} GPU compute accelerators, and Optane DC persistent memory. The system will rely on HPE's Shasta exascale-class architecture and Slingshot interconnect technology, which can provide concurrent support for advanced simulation and modeling, AI, and analytics workflows. Aurora will leverage historical advances in software investments along with increased application portability via Intel's oneAPI. The supercomputer will also be equipped with a new I/O platform called Distributed Asynchronous Object Storage (DAOS) to meet the needs of exascale workloads.

Exascale Collaborations

The process of planning and preparing for new leadership-class supercomputers takes years of collaboration and coordination. It requires partnerships with vendors and the broader HPC community to develop hardware, software, and storage technologies, as well as making facility enhancements to ensure the infrastructure is in place to power and cool the massive systems.

With Aurora, the ALCF has also benefitted from a close collaboration with the ECP, which has opened the door to working with a broader software portfolio and more extensive partnerships with other DOE labs than have been afforded in the past. Launched in 2016, the ECP is a multi-lab initiative aimed at accelerating the delivery of a capable exascale computing ecosystem that

The Argonne-Intel Center of Excellence (COE) has held multiple workshops to provide details and instruction on various aspects of the Aurora hardware and software environment.

encompasses applications, system software, hardware technologies, architectures, and workforce development. Argonne—one of the six ECP core labs—has a strong presence on the ECP leadership team and has several researchers engaged in ECP projects and working groups in the areas of application development, software development, and hardware technology.

ALCF staff members are working closely with more than 20 ECP application and software projects to prepare a diverse set of scientific applications to run on Aurora. Efforts have focused on porting various codes, mini-apps, frameworks, and libraries to Intel Gen9, DG1, and XeHP GPU hardware to evaluate and optimize performance. ALCF researchers are also developing algorithms, benchmarking tools, and use cases to help ensure applications can leverage the full potential of Aurora.

Another area of collaboration involves deploying and integrating ECP software, vendor software, and facility-based software environments to ensure exascale software stacks meet application requirements while allowing for optimal facility operations. Activities include developing a common Continuous Integration strategy to drive automation of recurring building and testing across DOE's exascale site environments; using the Spack package manager as a tool for build automation and final deployment of software; managing the ECP allocation program that awards time at DOE computing facilities for the testing and development of new features and functionality at scale; exploring and potentially enhancing the installation, upgrade, and monitoring capabilities of HPE's Shasta software stack; and enabling container support and tools on Aurora and other exascale systems.

Working in concert with the ECP, Argonne researchers are also contributing to the advancement of programming models (OpenMP, SYCL, Kokkos, Raja), language standards (C++), and compilers (Clang/LLVM) that are critical to developing efficient and portable exascale applications.

Furthermore, the ALCF continues to work closely with Intel and HPE on the testing and development of various components to ensure they can be leveraged effectively by the scientific computing community. By analyzing the performance of key benchmarks and applications on early hardware, ALCF researchers are developing a broad understanding of the system's architecture and capabilities. This effort helps to identify best practices for optimizing codes, and, ultimately, a roadmap for future users to adapt and tune their software for the new system. In addition, the team is providing feedback to Intel about the alpha/beta compilers and software stack for Aurora, including the gold release of its oneAPI toolkits.

Early Science

The Aurora Early Science Program (ESP), an ALCF initiative dedicated to preparing key applications for the scale and architecture of the exascale machine, is another critical piece of the puzzle. Through open calls for proposals, the ESP awarded pre-production computing time and resources to a diverse set of projects that are employing emerging data science and machine learning approaches alongside traditional modeling and simulation-based research. The research teams also field-test compilers and other software, helping to pave the way for other production applications to run on the system.

ESP research teams are investigating a wide range of computational research areas that will help prepare Aurora for primetime. This includes mapping and optimizing complex workflows; exploring new machine learning methodologies; stress testing I/O hardware and other exascale technologies; and enabling connections

AURORA/EXASCALE TRAINING EVENTS

Aurora Hackathons

06/22–06/23
High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

08/17–08/18
Exascale Computational Catalysis

09/17–09/18
Metascalable Layered Materials Genome

09/25 AND 09/28
Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

12/07
Simulating and Learning in the ATLAS Detector at the Exascale

12/08–12/10
Aurora COE Dungeon Session 1: Intensive Development on Pre-Aurora GPUs (Four Simulation Applications)

Aurora Webinars

03/25
DAOS: Next-Generation Data Management for Exascale

06/02
Preparing Applications for Aurora Using the Intel DPC++ Compatibility Tool

06/24
OpenMP Offload Capabilities in the oneAPI HPC Toolkit

09/30
Overview of the New Intel oneAPI Math Kernel Library

Aurora Workshops

02/25–02/27
Aurora COE Workshop 2: Aurora Software Development Kit and Hardware Overview

10/21–10/22
Aurora COE Workshop 3: Intel Hardware Update and Pre-Aurora GPU Development

Additional Events

07/15–07/17
ALCF-ECP CMake Workshop

09/01–09/02
2020 Performance, Portability, and Productivity at HPC Forum

10/13–10/16
Intel eXtreme Performance Users Group (IXPUG) Annual Meeting

Scheduled to arrive at Argonne in 2022, Aurora will provide researchers with a powerful tool for advancing scientific discovery and innovation.

to large-scale experimental data sources, such as CERN's Large Hadron Collider, for analysis and guidance.

With access to the early Aurora software development kit (a frequently updated version of the publicly available oneAPI toolkit) and Intel Gen9, DG1, and XeHP GPUs through Argonne's Joint Laboratory for System Evaluation, researchers participating in the ESP are able to test code performance and functionality using the programming models that will be supported on Aurora.

User Training

From workshops to webinars, the ALCF has hosted hundreds of researchers for a multitude of training opportunities designed to help them prepare for Aurora.

The facility launched its new Aurora Early Adopter webinar series late last year to introduce attendees to programming models, exascale technologies, and other tools available for testing and development work. Open to the public, the quarterly webinars have covered topics like Aurora's DAOS I/O platform, oneAPI's OpenMP offload capabilities, and the new oneAPI Math Kernel Library (oneMKL).

The Argonne-Intel Center of Excellence (COE) hosted two Aurora workshops this year aimed at ESP and ECP research teams. In February, around 100 attendees visited the ALCF for a three-day workshop focused on applications and software development for Aurora. The event included substantial hands-on time for attendees to work with ALCF and Intel experts on developing, testing, and profiling their codes, as well as presentations on the Aurora software development kit, the system's memory model, and available machine learning tools and frameworks. Attendees also had the opportunity to share programming progress, best practices, and lessons learned.

In October, the COE held a virtual Aurora workshop for ESP and ECP teams, providing another opportunity for researchers to receive updates on the latest Aurora hardware and software developments, connect with ALCF and Intel staff for assistance in developing and testing their codes, and share their experiences working with various exascale tools and programming models.

The COE also continued to host its series of intensive, hands-on sessions called "hackathons." These multi-day collaborative events, held virtually this year due to the COVID-19 pandemic, pair individual ESP teams with Argonne and Intel researchers to help advance their application development efforts using the Aurora software development kit.

As part of a broader effort to groom a new generation of supercomputer users, ALCF staff members organize and manage the annual Argonne Training Program on Extreme-Scale Computing (ATPESC). Supported by the ECP, this intensive, two-week program teaches attendees the key skills, approaches, and tools needed to design, implement, and execute computational science and engineering applications on current supercomputers and the exascale systems of the future. ATPESC has hosted more than 500 participants since it began in 2013.