

Case Study



High Performance Computing (HPC)
Intel® Xeon® Processors, codenamed Sapphire Rapids
with high bandwidth memory
Intel® Data Center GPUs, codenamed Ponte Vecchio
Intel® Optane™ Persistent Memory
Intel® oneAPI Programming Model

Argonne National Laboratory's Aurora 2+ ExaFLOPS System Will Enable New Science and Engineering

Offering performance projected to exceed two exaFLOPS of double precision compute performance, Aurora will empower previously unattainable research and engineering endeavors



Executive Summary

When delivered, Argonne National Laboratory's Aurora will be the nation's first Exascale HPC system built on Intel® Architecture.¹ With subcontractor Hewlett Packard Enterprise (HPE) and Intel, plus the support of the [U.S Department of Energy \(DOE\)](#), Aurora's performance is expected to exceed two exaFLOPS of double precision compute performance. With its extreme scale and performance levels, Aurora will offer the scientific community the compute power needed for the most advanced research in fields like biochemistry, engineering, astrophysics, energy, healthcare, and more.

Challenge

As a leading research agency in the United States, Argonne National Laboratory is at the forefront of the nation's efforts to deliver future Exascale computing capabilities. The [Argonne Leadership Computing Facility \(ALCF\)](#), the future home to Aurora, is helping to advance scientific computing through a convergence of HPC, high performance data analytics, and AI.

ALCF computing resources are available to researchers from universities, industry, and government agencies. Through substantial awards of supercomputing time and user support services, the ALCF enables large-scale computing projects aimed at solving some of the world's largest and most complex problems in science and



The Argonne Leadership Computing Facility (ALCF) will help drive simulation, data, and learning research to a new level when the facility unveils one of the nation's first Exascale machines built on Intel Architecture, Aurora.

engineering. Along with the desire to ensure competitiveness, the DOE and ALCF wanted to enable researchers to tackle challenges such as AI-guided analysis of massive data sets or full-scale simulations.

Solution

As prime, Intel built upon internal HPC system expertise and a tight partnership with HPC experts from Argonne and HPE as an integrator. Together, they will deliver the Exascale system, Aurora, expected to exceed two exaFLOPS of double precision compute performance.

The combined team has spent several years designing the system and optimizing it with specialized software and hardware innovations to achieve the performance necessary for advanced research projects. Other requirements for Aurora's design included components with long-term reliability and energy efficiency.

Upon its arrival, Aurora will feature several new Intel technologies. Each tightly integrated node will feature two Intel Xeon Processors, codenamed Sapphire Rapids with high bandwidth memory, plus six Intel Data Center GPUs, codenamed Ponte Vecchio. Each node will also offer scaling efficiency with eight fabric endpoints, unified memory architecture, and high bandwidth, low-latency connectivity. The system will support ten petabytes of memory for the demands of Exascale computing.

Aurora users will also benefit from Intel® Distributed Asynchronous Object Storage (DAOS) technology, which alleviates bottlenecks involved with data-intensive workloads. DAOS, supported on Intel Optane Persistent Memory, enables a software-defined object store built for large-scale, distributed Non-Volatile Memory (NVM).


The system will build upon HPE Cray Shasta supercomputer architecture, which incorporates next-generation HPE system software to enable modularity, extensibility, flexibility in processing choice, and seamless scalability. It will also include the HPE Slingshot interconnect as the network backbone, which offers a host of significant new features such as adaptive routing, congestion control, and Ethernet compatibility.

Cray ClusterStor E1000 parallel storage platform will support researchers' increasingly converged workloads by providing a total of 200 petabytes (PB) of new storage. The new solution

Spotlight on Hewlett Packard Enterprise

HPE combines computation and creativity, so visionaries can keep asking questions that challenge the limits of possibility. Drawing on more than 45 years of experience, HPE develops the world's most advanced supercomputers, pushing the boundaries of performance, efficiency, scalability, and sustainability. With developments like the HPE Cray Programming Environment for the HPE Cray EX supercomputing architecture and the HPE Slingshot interconnect, HPE continues to innovate new solutions for the convergence of data and discovery. HPE offers a comprehensive portfolio of supercomputers, high-performance storage, data analytics, and artificial intelligence solutions.

[Learn More.](#)



The Aurora supercomputer will be the United States' first Exascale system to integrate Intel's forthcoming HPC and AI hardware and software innovations including:

- Intel Xeon Processors, codenamed Sapphire Rapids with high bandwidth memory
- Intel Data Center GPUs, codenamed Ponte Vecchio
- > 230 Petabytes of storage based on Distributed Asynchronous Object Storage (DAOS) technology, bandwidth >25TB/S, and
- oneAPI unified programming model designed to simplify development across diverse CPU, GPU, FPGA, and AI architectures.


encompasses a 150 PB center-wide storage system, named Grand, and a 50 PB community file system, named Eagle, for data sharing. Once Aurora is operational, Grand, which is capable of one terabyte per second (TB/s) bandwidth, will be optimized to support the converged simulation science and new data-intensive workloads.

The Argonne team will depend on the oneAPI programming model designed to simplify development on heterogeneous architectures. oneAPI will deliver a single, unified programming model across diverse CPUs, GPUs, FPGAs, and AI accelerators.

Results

The team is currently working on ecosystem development for the new architecture. The ALCF formed the [Aurora Early Science Program \(ESP\)](#) to ensure the research community and critical scientific applications are ready for the scale and architecture of the Exascale machine at the time of deployment.

The ESP awarded pre-production time and resources to diverse projects that span HPC, high performance data analytics, and AI. Most of the chosen projects represent research so sophisticated that they have outgrown the capability of conventional HPC systems. Therefore, Aurora will help lead the charge into a new era of science where compute-intensive scientific endeavors not possible today become a reality.



Neurons rendered from the analysis of electron microscopy data. The inset shows a slice of data with colored regions indicating identified cells. Tracing these regions through multiple slices extracts the sub-volumes corresponding to anatomical structures of interest. (Image courtesy of Nicola Ferrier, Narayanan (Bobby) Kasthuri, and Rafael Vescovi, Argonne National Laboratory)

Next-Generation Science Requires Extreme HPC Systems

The projects first slated for time on Aurora represent some of the most difficult, and compute-intensive, endeavors. A few of the many projects accepted into the Aurora Early Science program include:

Developing Safe, Clean Fusion Reactors

Fusion, the way the sun produces energy, offers enormous potential as a renewable energy source. One type of fusion reactor uses magnetic fields to contain the fuel—a hot plasma including deuterium, an isotope of hydrogen derived from seawater. Dr. William Tang, Principal Research Physicist at the Princeton Plasma Physics Lab, plans to use Aurora to train an AI model to predict unwanted disruptions of reactor operation. Aurora will ingest massive amounts of data from present-day reactors to train the AI model. The model may then be deployed at an experiment, to trigger control mechanisms that prevent upcoming disruptions. Thanks to Exascale computing, the emergence of AI, and deep learning, Tang will provide new insights that will advance efforts to achieve fusion energy.

Neuroscience Research

Dr. Nicola Ferrier, a Senior Computer Scientist at Argonne, is partnering with researchers from the University of Chicago, Harvard University, Princeton University, and Google. The collaborative effort seeks to use Aurora to understand the bigger-picture brain structure and how each neuron connects with others to form the brain's cognitive pathways. The team hopes their arduous endeavor will reveal information to benefit humanity, like potential cures for neural diseases.

Seeking More Effective Treatments for Cancer

Dr. Amanda Randles, Alfred Winborne Mordecai and Victoria Stover Mordecai Assistant Professor in the Department of Biomedical Engineering at Duke University, and her colleagues developed the "HARVEY" system. HARVEY predicts the flow of blood cells moving through the highly complex human circulatory system. With her time on Aurora, Dr. Randles seeks to repurpose HARVEY to understand metastasis in cancer better. By predicting where metastasized cells might travel in the body, HARVEY can help doctors anticipate, early on, where secondary tumors may form.


Understanding the "Dark" Universe

The combination of stars, planets, gas, clouds, and everything else that is visible in the cosmos comprises a mere five percent of the universe. The other 95 percent consists of dark matter and dark energy. The universe is not only growing—its rate of expansion is accelerating. Dr. Katrin Heitmann, Physicist and Computational Scientist at Argonne National Laboratory, has big goals for her time with Aurora. Her research seeks to obtain a deeper understanding of the Dark Universe, which we know so little about today.

Designing More Fuel-Efficient Aircraft

Dr. Kenneth Jansen, Professor of Aerospace Engineering at the University of Colorado Boulder, pursues designs for safer, more performant, and more fuel-efficient airplanes by analyzing turbulence around an airframe. The variability of turbulence makes it difficult to simulate an entire aircraft's interaction with it. Each second, different parts of a plane experience different impacts from the airflow. Therefore, Dr. Jansen and his team need to evaluate data in real-time as the simulation progresses. HPC systems today lack the capability for the task, simulating airflow around a plane one-nineteenth its actual size, traveling at a quarter of its real-world speed.

Aurora will help Dr. Jansen and his team learn more about the



This simulation of a massive structure, a so-called cluster of galaxies, was run on Argonne's Theta system, as part of an earlier ESP. The mass of the object is 5.6×10^{14} Msun. The color shows the temperature, and white areas show the baryon density field. (Image courtesy of JD Emberson and the HACC team)

Spotlight on Argonne National Laboratory

Based in Illinois, Argonne National Laboratory is a multidisciplinary research center focused on tackling the most important questions facing humanity. With support from the US Department of Energy (DoE), Argonne collaborates with many organizations including corporate and academic institutions, and other labs around the nation, to enable scientific breakthroughs crossing disciplines like physics, chemistry cosmology, and biology. [Learn More.](#)

fundamental physics involved with full flight scale and real flight conditions. From there, they can identify where design improvements can make an important difference for in-flight characteristics.

Supporting CERN's Large Hadron Collider Project (LHC)

Dr. Walter Hopkins, a Physicist at Argonne, is a member of the ATLAS experiment, which is an international collaboration that studies the fundamental particles and forces that make up our universe. The ATLAS experiment images the results of proton collisions in CERN's Large Hadron Collider (LHC).

These images were used in the historic discovery of the Higgs Boson in 2012, which completed the Standard Model of Particle Physics. Over the next ten years the upgraded LHC and ATLAS experiment will collect 10x more data that will help answer questions that remain, for instance "What is dark matter?" or "How is gravity related to the electromagnetic, strong, or weak forces?" While the amount of data will increase by 10x, the amount of simulation needed for physics studies will increase by 100x, quickly outpacing current resources. This project is porting some of the more computationally intensive simulations to accelerators to address this increase. Additionally, the project is leveraging deep learning to expand the analytical reach of current particle identification algorithms. With this project, Aurora will become an important resource for discovery in the next phase of searching for new physics.

A Bright Future for Research

Exascale computing will empower researchers with a profound and transformative tool. Aurora's performance levels, scale, and the ability to process enormous data sets offer incredible potential. The system will help unlock mysteries, which have perplexed scientists and engineers for decades. Aurora will also enable unprecedented levels of innovation and discovery in engineering.

Additional Information

Learn more about [Intel HPC solutions](#)

Learn about [HPE Cray Supercomputing Technology](#)

Watch videos explaining the plans of the [Early Science Program Researchers.](#)

"HPE is honored to partner with Intel to build and deliver the first US exascale supercomputer to Argonne. It is an exciting testament to HPE Cray EX's flexible design and unique system and software capabilities, along with our HPE Slingshot interconnect, which will be the foundation for Argonne's extreme-scale science endeavors and data-centric workloads. The HPE Cray EX Supercomputer is designed for this transformative exascale era and the convergence of artificial intelligence, analytics, modeling, and simulation—all at the same time on the same system—at incredible scale."

— Peter Ungaro, senior vice president and general manager, HPC and AI, at HPE

Technical Components

[Intel Xeon Processors, codenamed Sapphire Rapids with high bandwidth memory](#)

[Intel Optane Persistent Memory](#)

[Intel Data Center GPUs, codenamed Ponte Vecchio](#)

[Intel oneAPI unified programming model](#)

[Intel Distributed Asynchronous Object Storage \(DAOS\) technology](#)

[HPE Cray EX Supercomputer](#)

[HPE Cray software stack](#)
optimized for Intel architectures

[HPE Slingshot interconnect](#)



¹ See the press announcement in the [Intel Newsroom](#).

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available security updates. See backup for configuration details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.