intel + Argonne NATIONAL LABORATORY + Hewlett Packard Enterprise + U.S. DEPARTMENT OF ENERGY

Aurora Supercomputer

Argonne National Laboratory teams with Intel to advance science and engineering to speed discovery with world-changing technology

Exascale = a billion billion (a quintillion) operations per second

Artificial Intelligence

Analytics

HPC Simulation

1 second
The time it takes Aurora to solve a math problem that would take 40 years if all the people on Earth each did **one calculation every 10 seconds.**

600 tons
The weight of Aurora, which equals that of an **Airbus 380.**

300 miles
The length of optical cable used in Aurora could reach **from Los Angeles to San Jose, California.**

10,000 square feet
The amount of floor space for Aurora, which **equals to 4 tennis courts.**

8 minutes
The time it takes Aurora to store enough characters to write **a stack of books that could reach the moon.**

34,000 gallons per minute
The rate of water moving through the **cooling loop.**

Aurora Early Science Program

Simulation
Simulation allows researchers to create virtual representations of complex physical systems or processes that are too small or large, costly, or dangerous to study in a laboratory.

Data
The use of advanced data science techniques and tools to gain insights into massive datasets produced by experimental, simulation, or observational methods.

Learning
A form of artificial intelligence, machine learning refers to a set of algorithms that uses training data to identify relationships between inputs and outputs, and then generates a model that can be used to make predictions on new data.

Exascale Research

Cancer Research

Neuroscience

Aircraft Design

Fusion Energy

Catalyst Research

High Energy Physics

Intel Inside

Building the Foundation for Exascale Computing

Aurora Node Architecture

2 Future Intel® Xeon™ Scalable Processors
"Sapphire Rapids"

6 Xᵉ Architecture Based GPUs
"Ponte Vecchio"

oneAPI
Unified programming model

Unparalleled I/O Scalability across Nodes
8 fabric endpoints per node, DAOS

Leading Performance
HPC, data analytics, AI

All-to-All Connectivity within Node
Low latency, high bandwidth

Unified Memory Architecture
Across CPUs and GPUs

Packaging
Foveros and EMIB

Unmatched Exascale-Class Storage Performance
Exascale systems require a completely rearchitected storage infrastructure. Aurora will benefit from the fastest High Performance Computing (HPC) storage on the planet – based on Intel® Optane™ persistent memory and the open source Distributed Asynchronous Object Storage (DAOS) framework, which together have enabled systems to achieve #1 ranking on the IO500 list.

Additional Details
• Aurora will have more than 230 petabytes of storage with 25TB/s access rates
• Interconnect: HPE Slingshot
• Topology: Dragonfly
• Network switch: 64-port switch, 25GB/s per direction