

Aurora Site Update

Kevin Harms
Argonne National Laboratory



Aurora

Leadership Computing Facility
Exascale Supercomputer

Peak Performance
 ≥ 2 Exaflops DP

Intel GPU
**Intel® Data Center
GPU Max**

Intel Xeon Processor
**Intel® Xeon® CPU
Max**

Platform
HPE Cray-Ex

Compute Node

2 Xeon Intel® Xeon® CPU Max
processors
6 Intel® Data Center GPU Max
Node Unified Memory Architecture
8 fabric endpoints

GPU Architecture

Intel XeHPC architecture
High Bandwidth Memory Stacks

Node Performance

>130 TF

System Size

>9,000 nodes

Aggregate System Memory

>10 PB aggregate System Memory

System Interconnect

HPE Slingshot 11
Dragonfly topology with adaptive routing

Network Switch

25.6 Tb/s per switch (64 200 Gb/s ports)
Links with 25 GB/s per direction

High-Performance Storage



220 PB
 ≥ 25 TB/s DAOS bandwidth

Software Environment

- C/C++
- Fortran
- SYCL/DPC++
- OpenMP offload
- Kokkos
- RAJA
- Intel Performance Tools



Aurora Storage Architecture



The Aurora open-source storage strategy strongly favors cooperation:

- DAOS: object storage system for in-fabric high-performance platform storage (the first of its kind on a DOE leadership system!)
- Lustre: parallel file systems for facility-wide access and data sharing


Namespace integration will make it easier for users to manage data.

- 1024 DAOS server nodes, each with:
 - 16 x 512GB persistent memory
 - 16 x 15.3TB NVMe drives
 - 2 x HPE Slingshot NICs
 - Dual CPU with 512 GB RAM

Aurora DAOS Status

- Aurora compute nodes available for testing by ALCF, Intel and ECP users
- Intel team testing with Aurora compute nodes now
 - Looking at application scaling
 - Performance analysis
 - I/O testing (more later)
- Lots of 2.4 release testing
- Planning on DAOS partition with 20-40 DAOS nodes
 - Available and run in production like fashion
 - Setup to have DAOS available on job execution via parameters provided to scheduler
 - Evaluate DAOS as /home and /soft file systems
 - /home is our traditional home file system for building and basic work
 - /soft is a volume for locally built and installed software, supported by the wider ALCF staff
- DAOS Foundation

The DAOS Foundation




- **Intel**, in partnership with **Argonne National Lab**, **Enakta Labs**, **Google**, and **HPE** have launched the DAOS Foundation
 - Democratizing and accelerating DAOS for HPC and AI/ML applications
 - Together, members drive future DAOS development, pooling resources and knowledge
 - Independent and future-proof of any single company, product, or technology
 - Intel remains committed to DAOS

Sunspot

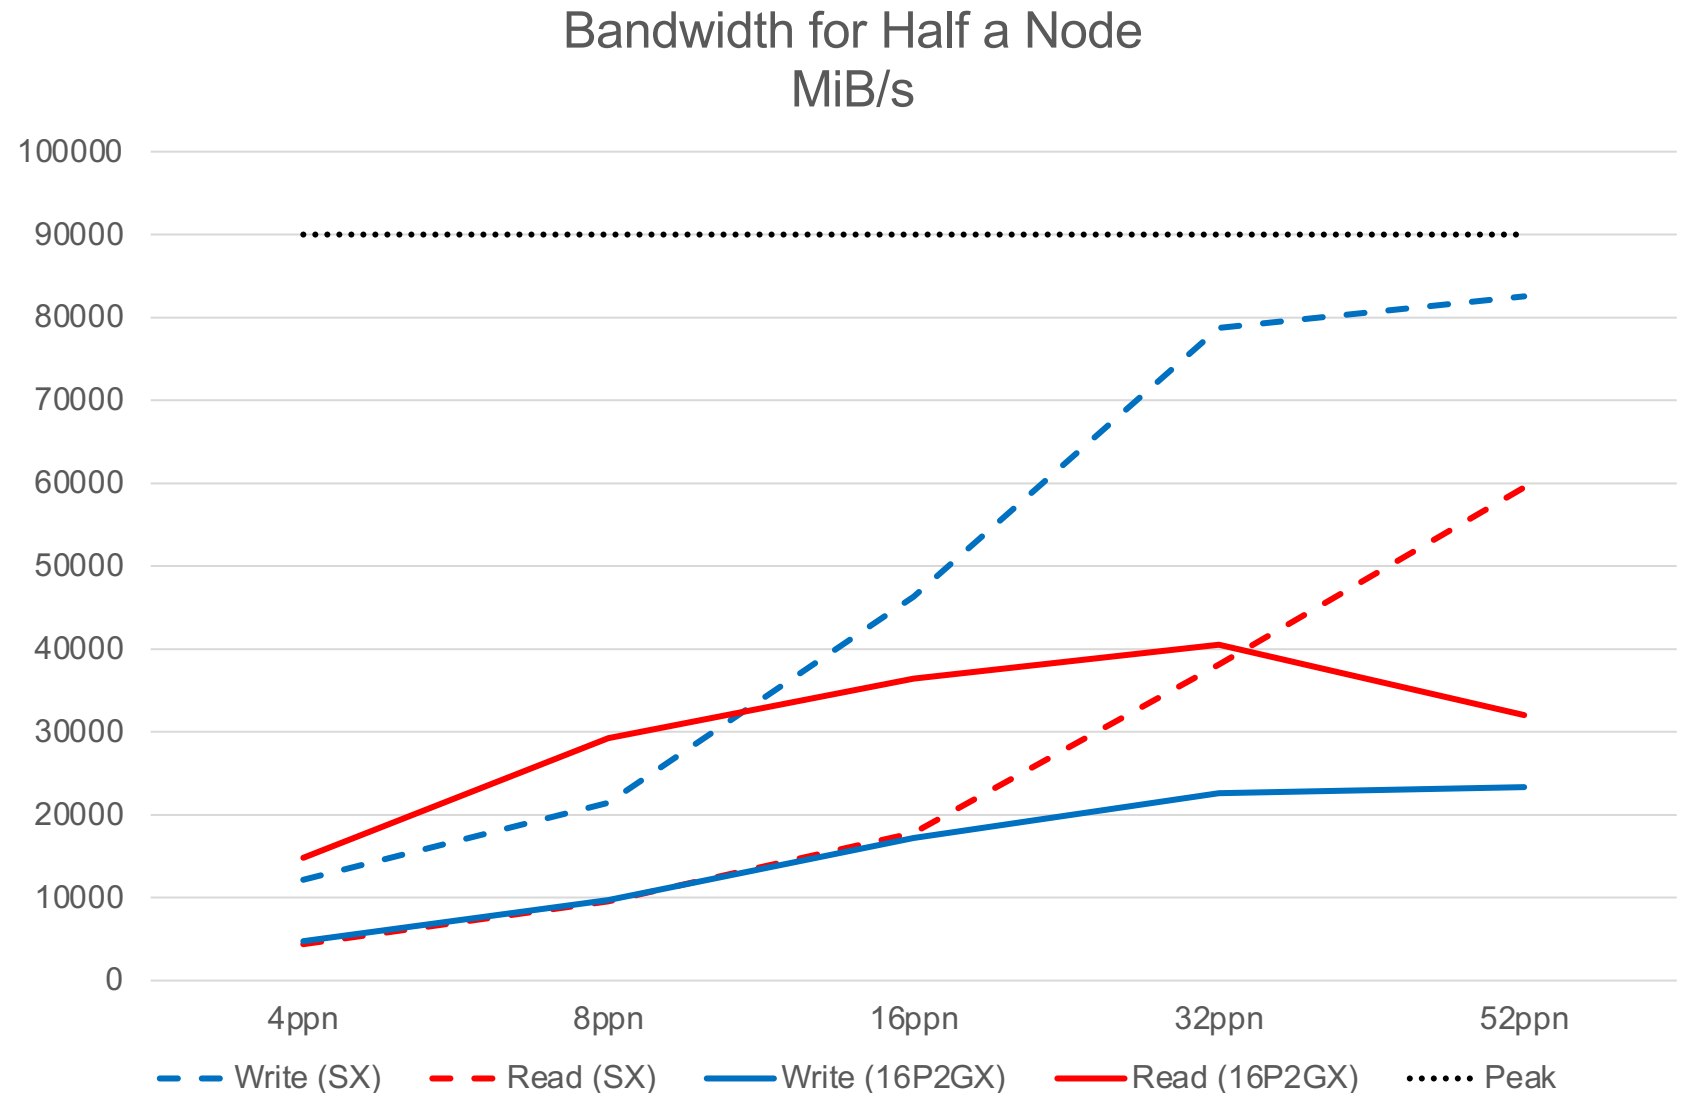
<https://www.alcf.anl.gov/support-center/aurora/getting-started-sunspot>

- ALCF's Test and Development system
 - Think of it as a baby Aurora
- Two compute racks / groups
 - 128 compute nodes
- DAOS deployment
 - 20 DAOS nodes
 - Identical server configuration to Aurora
 - Allows running EC16+2 – 18 nodes with 2 nodes for failover
- Production environment for DAOS at ALCF
 - Follow pool and container usage plan for Aurora
 - 1 pool per project
 - ACL limits pool to project members
 - Users create containers
 - Suggested default data protection of EC16+2 on containers
 - Running DAOS v2.2



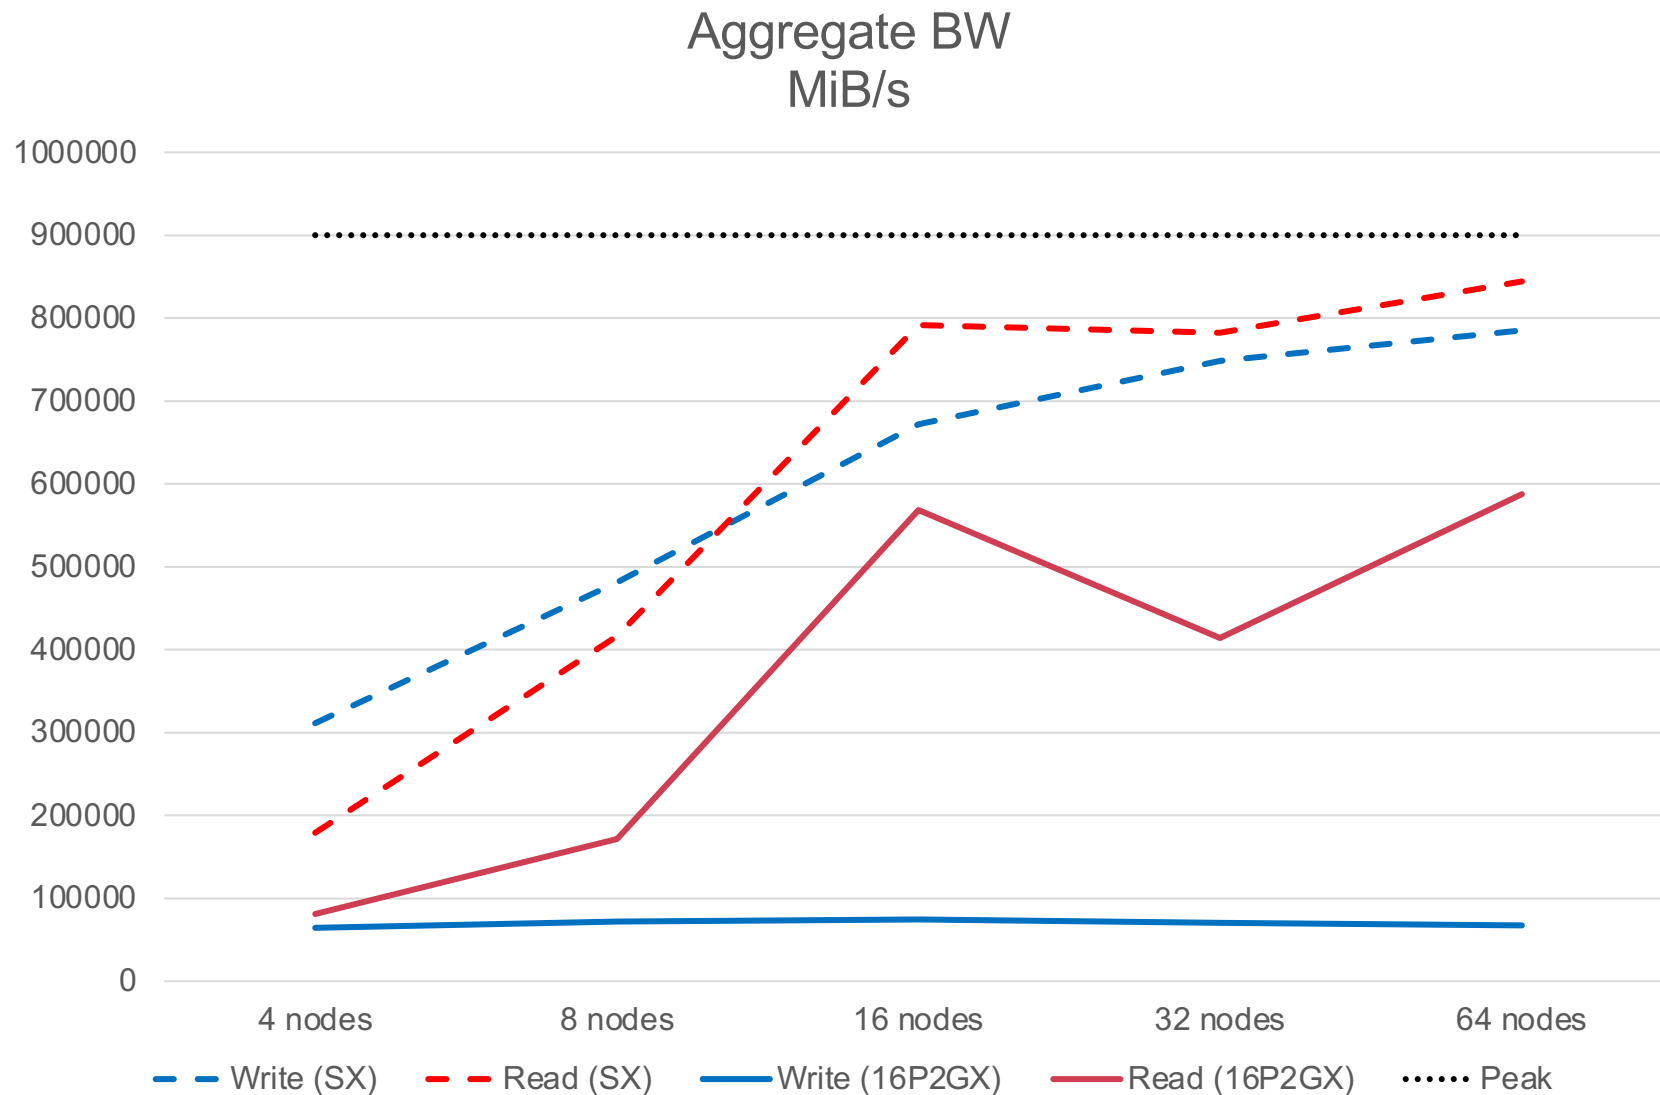
IOR Scaling per Process

- Each step uses 1 process per NIC across 4 NICs
— 1 MiB I/O request
- No redundancy
— dfs.oclass=SX
- Redundancy
— dir.oclass=EC_16P2GX



IOR per Node Scaling

- Running up to 64 nodes
 - Using 8 NICs per node
- No redundancy
 - 1 MiB I/O size
 - dfs.chunk_size=128k
 - dfs.oclass=SX
- Redundancy
 - 2 MiB I/O size
 - dfs.chunk_size=2M
 - dir.oclass=EC_16P2GX



Sunspot Results and Issues

- Sunspot is running DAOS v2.2
 - We would like to upgrade but blocked by other issues unrelated to DAOS or storage
- Sunspot performance within expectations
 - Issues with redundant performance (fixed in v2.4)
- Sunspot had several science users testing
 - Mostly using POSIX dfuse
 - A few with MPI-IO test cases
 - Some successes, some failures
- Rebuild would enter infinite rebuild loop when EC data was present
 - Known fixed in 2.4
- Servers would start reporting ENOMEM for internal operations
 - Need to evaluate on 2.4
 - DAOS test team doesn't see this
- Determined the need for a dedicated DAOS management node
 - Simplifies management scheme to use centralized management point that isn't central admin node
 - Don't perform management operations on the server node itself

IO-500 Results

- New IO-500 SC23 BoF submission
 - <https://io500.org>
 - First production submission for Aurora!
 - Tolerate 1 server failure
 - Using 16+1 and replication
 - Still only using a portion of the available Aurora resources
 - 300 clients
 - 642 servers
- **IO500: The High-Performance Storage Community**
 - Wednesday, November 15, 2023
 - 12:15 – 1:15pm MST
 - Room 607



<https://io500.org>

Acknowledgements

This research used resources of the Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory and is based on research supported by the U.S. DOE Office of Science-Advanced Scientific Computing Research Program, under Contract No. DE-AC02-06CH11357.