



The 3rd Annual Japan Conference

April, 2023


The HPC-AI Advisory Council

- **Worldwide HPC and AI community organization, established in 2008**
- **More than 430 member companies / universities / research centers**
- **Bridges the gap between HPC and AI usage and its potential**
- **Provides best practices, education, technology demonstrations, development center**
- **Explores future technologies and future developments**

HPC Advisory Council Objectives

- HPC Technology
- Network of Expertise
- HPC Outreach
- High-Performance Center
- Education
- Best Practices

HPC-AI ADVISORY COUNCIL CENTERS



- The Council operates a cluster center with 11 clusters available for operation
- Providing free of charge access to variety of compute, network and storage technologies
- For more information: http://hpcadvisorycouncil.com/cluster_center.php



- Daytona_X AMD 8-node cluster
- Dual Socket AMD EPYC 7742 64-Core Processor @ 2.25GHz
- Mellanox ConnectX-6 HDR 200Gb/s InfiniBand/Ethernet
- Mellanox HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand
- Memory: 256GB DDR4 2666MHz RDIMMs per node
- Lustre Storage, NFS



- Dell C6400 32-node cluster
- Dual Socket Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz
- Mellanox ConnectX-6 HDR100 100Gb/s InfiniBand/VPI adapters
- Mellanox HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand
- Memory: 192GB DDR4 2666MHz RDIMMs per node
- Lustre Storage, NFS



- Supermicro SYS-6029U-TR4 / Foxconn Groot 1A42USF00-600-G 32-node cluster
- Dual Socket Intel(R) Xeon(R) Gold 6138 CPU @ 2.00GHz
- Mellanox ConnectX-6 HDR/HDR100 200/100Gb/s InfiniBand/VPI adapters with Socket Direct
- Mellanox HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand
- Memory: 192GB DDR4 2666MHz RDIMMs per node
- 1TB 7.2K RPM SSD 2.5" hard drive per node



- IBM S822LC POWER8 8-node cluster
- Dual Socket IBM POWER8 10-core CPUs @ 2.86 GHz
- Mellanox ConnectX-4 EDR 100Gb/s InfiniBand adapters
- Mellanox Switch-IB SB7700 36-Port 100Gb/s EDR InfiniBand switch
- Memory: 256GB DDR3 PC3-14900 RDIMMs per node
- 1TB 7.2K RPM 6.0 Gb/s SATA 2.5" hard drive per node
- GPU: NVIDIA Kepler K80 GPUs

Includes a wide range of architectures from over a decade of benchmarking


- 19 Best practices added in H2 2022:

Latest updates:

- MPAS (AMD Milan, HDR InfiniBand)
- DL-POLY (AMD Milan, HDR InfiniBand)
- STAR-CCM+ (Intel Cascadelake, HDR InfiniBand)
- MILC (AMD Milan, HDR InfiniBand)
- CESM (Intel Cascadelake, HDR InfiniBand)
- AMG (AMD Milan, HDR InfiniBand)
- AMG (Intel Icelake, HDR InfiniBand)
- GROMACS (AMD Milan, HDR InfiniBand)
- NAMD (AMD Milan, HDR InfiniBand)
- GROMACS (Intel Cascadelake/Icelake, HDR InfiniBand)
- Quantum ESPRESSO (AMD Milan, HDR InfiniBand)
- LS-DYNA (Intel Cascadelake, HDR InfiniBand)
- HPCG (Intel Icelake, HDR InfiniBand)
- HPCG (AMD Milan, HDR InfiniBand)
- NAMD (Intel Cascadelake/Icelake, HDR InfiniBand)
- CP2K (AMD Milan, HDR InfiniBand)
- CP2K (Intel Cascadelake/Icelake, HDR InfiniBand)
- Quantum ESPRESSO (Intel Cascadelake/Icelake, HDR InfiniBand)
- ANSYS Fluent (Intel Cascadelake, HDR InfiniBand)

Abaqus	ABySS	AcuSolve	Amber	AMG	AMR
B_EFF	BiFrost	BQCD	BSMBench	CAM-SE	CASTEP
CCSM	CESM	ChaNGa	CFX	COSMO	CP2K
CPMD	Dacapo	Desmond	DL-POLY	Eclipse	FLOW-3D
Fluent	GADGET	Graph500	GRID	GROMACS	Himeno
HIT3D	HOOMD	HPCC	HPCG	HYCOM	ICON
Lattice	LAMMPS	LS-DYNA	MetaComp	miniFE	MILC
MSC	MR-Bayes	MM5	MPQC	NAMD	Nekbone
NEMO	NEMO5	NWChem	Octopus	OpenAtom	OpenFOAM
OpenMX	OptiStruct	PARATEC	PFA	PFLOTRAN	Quantum
RADIOSS	RFD	SNAP	SPECFEM3D	STAR	VASP
VPS	WRF				

The New Universe of Scientific Computing



- Multi-Tenant RDMA
- Infrastructure Processing
- In-Network Computing
- Performance Isolation
- Zero Trust Security

Overlapping



In-Network Computing
Asynchronous Progress

Load Imbalanced



In-Network Computing
Synchronization

Jitter



In-Network Computing
Infrastructure Processing

Multi-Job
Performance



In-Network Computing
Performance Isolation

Championing Student Learning and Development

**Annual ISC-Student
Cluster Competition
and
DPU Hackathons**



**Annual APAC HPC-AI
Competition**



**Annual APAC Advanced
RDMA Programming
Workshop and
Competition**



Student / Organized by HPC-AI Advisory Council
RDMA Programming Competition

**Stanford Summer HPC
Graduate Student Series
and Stanford Summer
AI Workshop**

Stanford

Worldwide Conferences - Fostering Community and Learning

- Featuring leading experts in invited talks that address a range of interests
- Latest trends, cutting-edge technologies and breakthrough science and exploration
- All new best practices in applications, tools and techniques
- Open to the community

Winter

Stanford Conference



Spring

Swiss Conference
Japan Conference



Summer

Australia Conference



Autumn

UK Conference
China Conference



- **For more information**

- www.hpcadvisorycouncil.com
- info@hpcadvisorycouncil.com

- **Social**

- Twitter: @hpccouncil
- LinkedIn: <https://www.linkedin.com/groups/1732037/>
- YouTube: [youtube.com/user/hpcadvisorycouncil](https://www.youtube.com/user/hpcadvisorycouncil)



Thank You



All trademarks are property of their respective owners. All information is provided “As-Is” without any kind of warranty. The HPC Advisory Council makes no representation to the accuracy and completeness of the information contained herein. HPC Advisory Council undertakes no duty and assumes no obligation to update or correct any information presented herein