University of CINCINNATI

Advanced Research Computing (ARC) Initiative

Six Month Status Report 1/1/2019 – 7/1/2019

User Profiles and Science/Scholarship supported

Faculty Advisory Committee
Meetings: April 30, June 13

Report/Recommendation due to Pat Limbach 7/1/2020

1. Value Proposition of UC HPC cluster – Why HPC at UC?
2. Operations/Funding model/Sustainability

Activities Completed

Operations

- All hardware received, installed and configured
- Network Connectivity established to the cluster
- AD Login installed and operational – AD groups defined

Research Facilitation and Support

- Open Source software installed
- HPC Cluster management, scheduling and connectivity established
- Intel Compilers purchased and installed
- ARC initiative website – splash page for marketing purposes and external audiences
- Benchmarking of hardware operations
- Scheduling policies established via Slurm
- Other Software Installed and configured per project needs
- Issues resolved – Omnipath, Slurm Pre-emption, Orphan processes, Mail forwarding, node failure
- Connection to existing license servers for MatLab and other tools – CEAS
- Training and workforce development opportunities for faculty, students and staff
- Testing Completed – 54 projects – facilitated scheduling, access, usage, optimization, software installation

Activities In-Process

- Globus subscription and implementation
- Purchase, install and implement Research Storage – scratch and project storage
- IDRAC – remote management tools installed/configured
- 'Open OnDemand' installation (web-front end from Ohio Supercomputer Center for Data Scientists, AI, ML)
- Training and workforce development opportunities for faculty, students and staff
- Require 'Intro to UC HPC ' and 'Advanced Linux' workshops before using HPC cluster
- Investigate Software License Server

Chair, Faculty Advisory Committee – Prashant Khare
Office of Research, Project Sponsor – Phil Taylor
HPC System Admin – Larry Schartman
Research Facilitator, Interim Director– Jane Combs
Linux System Admin – Adam Steele

University of CINCINNATI

Advanced Research Computing (ARC) Initiative

Six Month Status Report 1/1/2019 – 7/1/2019

- Implement XDMOD – user statistics module which connects to Coeus (or other grant mgmt. system)

Usage

- 31 projects requested access to the HPC Cluster
- CEAS Aerospace, CEAS Mech Eng, CEAS Civil Eng, CEAS EECS, A&S Chem, A&S Math, LCB OBAIS, Institute for Interdisciplinary Data Science (IIDS), CoM Env Health, CECH Criminal Justice

Average Wait Time for Jobs to Run

Issues/Risks

- New users to ARC (data science, humanities, etc.) who have not used HPC clusters before require significant support and assistance optimizing their codes to run in a multi-core environment.
- Reconfiguration of UCScienceNet is required to fully meet data transfer and security needs
- Omnipath is being phased out by Intel. Need to identify a mitigation plan

Chair, Faculty Advisory Committee – Prashant Khare
Office of Research, Project Sponsor – Phil Taylor
HPC System Admin – Larry Schartman
Research Facilitator, Interim Director– Jane Combs
Linux System Admin – Adam Steele

Updated 6/14/2019

University of CINCINNATI

Advanced Research Computing (ARC) Initiative

Six Month Status Report 1/1/2019 – 7/1/2019

Hardware

Currently, *ARC* is equipped with 50 teraFLOPS of peak CPU performance and 2 NVIDIA Tesla V100 GPU nodes (224 teraFLOPS deep learning peak performance) connected with high-performance 100 GB/s Omnipath (OPA) interconnect, a significant step forward in both bandwidth and latency. *

Pilot HPC Cluster Hardware – 36 total nodes

- **CPU Nodes -** Intel Xeon Gold 6148 2.4G, 20C/40T, 10.4GT/s, 27M Cache, Turbo, HT (150W) DDR4-2666 RAM- 192GB per node, DDR4-2666
- **GPU Node -** NVIDIA Tesla V100 32G Passive GPU with 2 nodes
- **ZFS Storage Node** – 96TB raw storage (initially configured to offer 43TB)
- **Omnipath HPC Networking infrastructure –** Maximum Ominpath bandwidth between nodes = 100Gbps
- **Miscellaneous** (including cables, racks, PDUs, hardware enclosures)

Additional HPC Cluster Hardware – requested in NSF MRI proposal submitted 1/20/2019

The proposed system, *Discovery Cluster* consists of the following:

- Sixty-five (65) dual-socket Intel Xeon Gold 6148 Skylake (37) and Cascade Lake (28) processors with 192 Gbytes of DDR4 memory per node. Each microprocessor chip has 40 cores at 2.4 GHz (3.7 GHz turbo), providing a total of 90 teraFLOPS for the acquired system. This will provide a peak performance of almost 140 teraFLOPS.
- Six (6) NVIDIA Tesla V100 GPU nodes with 32 Gbytes of CoWoS Stacked HBM2 memory per node. Each GPU node consists of 640 Tensor cores, providing a total 672 teraFLOPS deep learning peak performance. Once acquired, the peak performance for deep learning at *ARC* will become 900 teraFLOPS

Chair, Faculty Advisory Committee – Prashant Khare
Office of Research, Project Sponsor – Phil Taylor
HPC System Admin – Larry Schartman
Research Facilitator, Interim Director– Jane Combs
Linux System Admin – Adam Steele

Updated 6/14/2019