Omni-Path - Bridges New User Communities

Richard Underwood

HPC System Administrator

11/12/2018

- *Available at no cost for open research and courses and by arrangement to industry*
- *Easier access for CMU and Pitt faculty through the Pittsburgh Research Computing Initiative*
- 28,628 Intel Xeon CPU cores
- 216 NVIDIA GPUs: 64 K80, 64 P100, 88 V100
- 17 PB storage (10 PB persistent, 7.3 PB local)
- 277 TB memory (RAM), up to 12 TB per node
- 44M core-hours, 442k GPU-hours, and 343k TB-hours allocated quarterly
- Serving 1,573 projects and ~7500 users

- ***Bridges-DL:*** NVIDIA DGX-2 Enterprise AI system + 9 HPE 8-Volta Apollo 6500 Gen10 servers: total of 88 V100 GPUs

Bridges **converges HPC, AI, and Big Data to empower new research communities, bring desktop convenience to advanced computing, expand remote access, and help researchers to work more intuitively.**

- Funded by NSF award #OAC-1445606 ($20.9M), Bridges emphasizes usability, flexibility, and interactivity
- Available at no charge for open research and coursework and by arrangement to industry
- Popular programming languages and applications: Python, Jupyter, R, MATLAB, Java, Spark, Hadoop, …
- 856 compute nodes containing Intel Xeon CPUs and 128GB (800), 3TB (42), and 12TB (4) of RAM each
- 216 NVIDIA Tesla GPUs: 64 K80, 64 P100, **(new) 88 V100 configured to balance capability & capacity**
- Dedicated nodes for persistent databases, gateways, and distributed services
- The world's first deployment of the Intel Omni-Path Architecture fabric

NSF

Hewlett Packard Enterprise

intel NVIDIA

Carnegie Mellon University PSC University of Pittsburgh

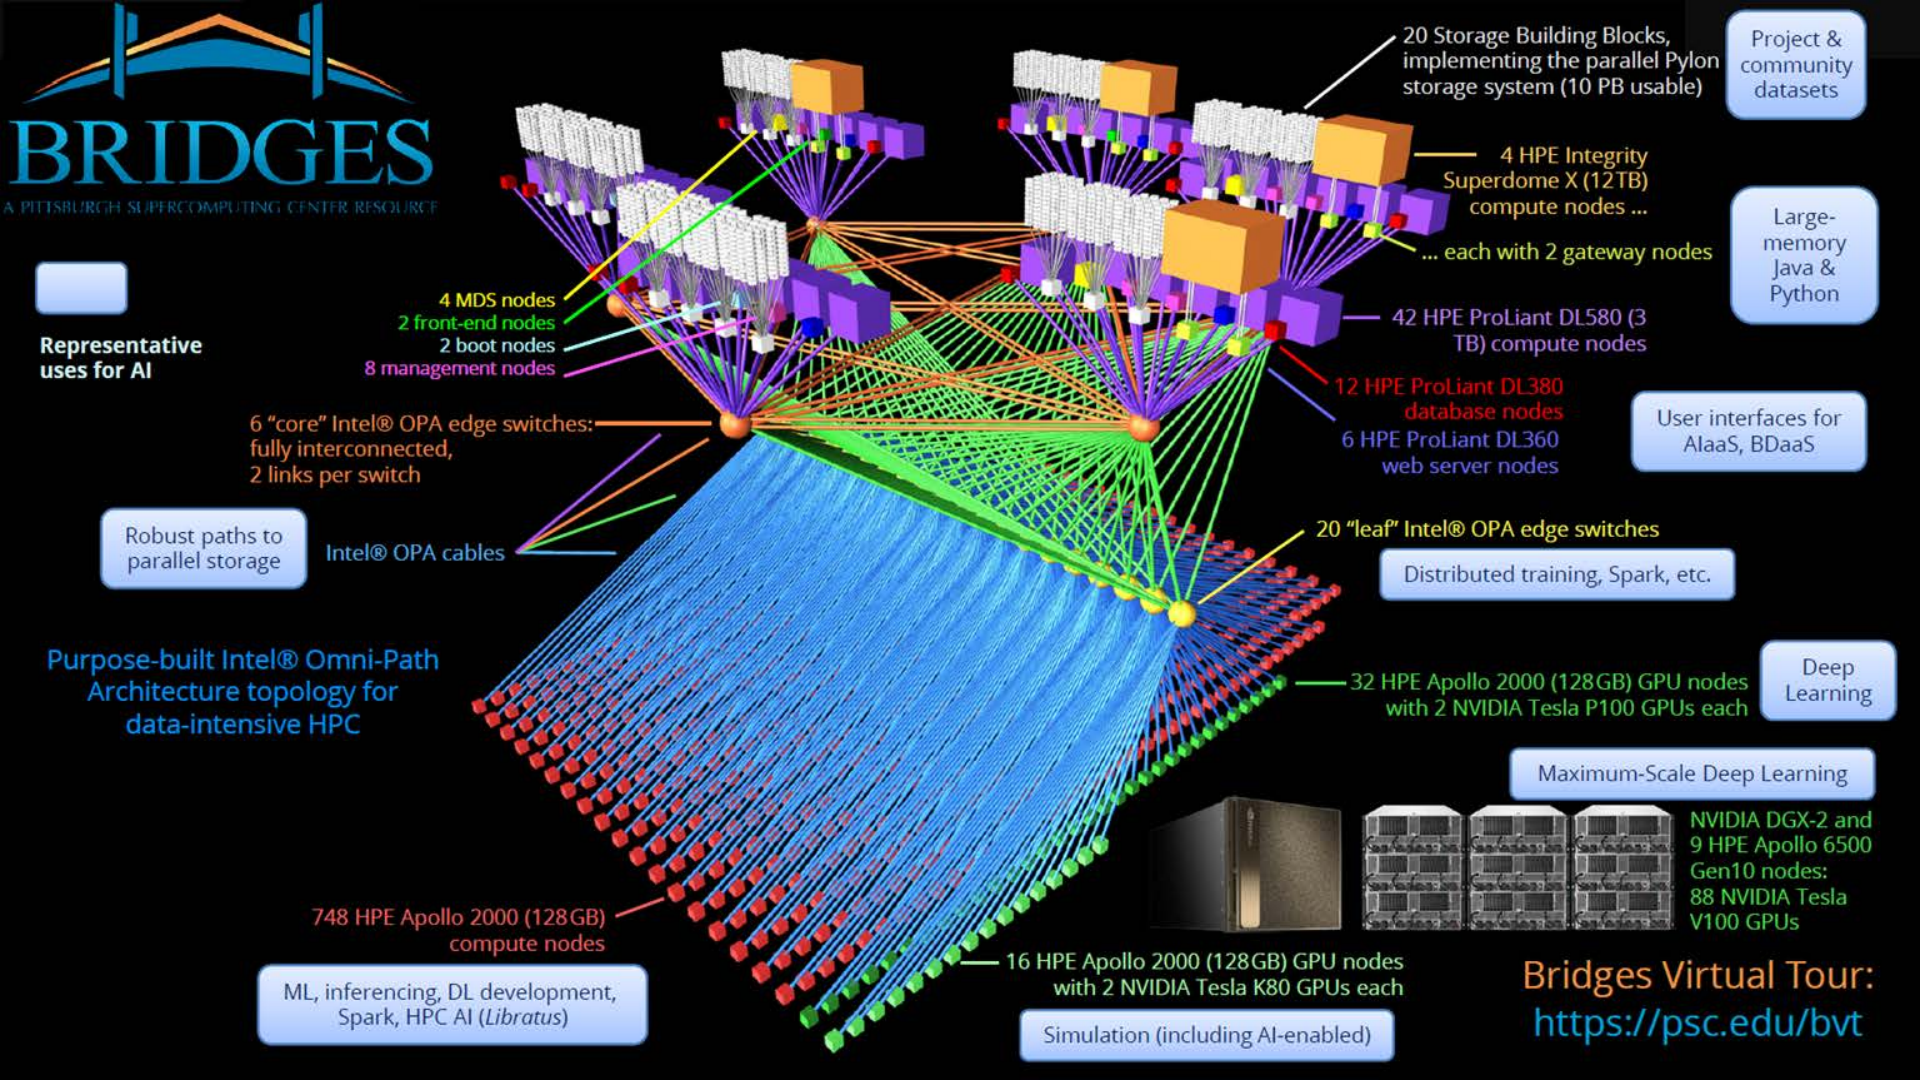
BRIDGES
A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE
Representative uses for AI
20 Storage Building Blocks, implementing the parallel Pylon storage system (10 PB usable)
Project & community datasets
4 HPE Integrity Superdome X (12TB) compute nodes ...
... each with 2 gateway nodes
Large-memory Java & Python
4 MDS nodes
2 front-end nodes
2 boot nodes
8 management nodes
42 HPE ProLiant DL580 (3 TB) compute nodes
12 HPE ProLiant DL380 database nodes
6 HPE ProLiant DL360 web server nodes
User interfaces for AIaaS, BDaaS
6 "core" Intel® OPA edge switches: fully interconnected, 2 links per switch
Robust paths to parallel storage
Intel® OPA cables
20 "leaf" Intel® OPA edge switches
Distributed training, Spark, etc.
Purpose-built Intel® Omni-Path Architecture topology for data-intensive HPC
32 HPE Apollo 2000 (128GB) GPU nodes with 2 NVIDIA Tesla P100 GPUs each
Deep Learning
Maximum-Scale Deep Learning
NVIDIA DGX-2 and 9 HPE Apollo 6500 Gen10 nodes: 88 NVIDIA Tesla V100 GPUs
748 HPE Apollo 2000 (128GB) compute nodes
ML, inferencing, DL development, Spark, HPC AI (Libratus)
16 HPE Apollo 2000 (128GB) GPU nodes with 2 NVIDIA Tesla K80 GPUs each
Simulation (including AI-enabled)
Bridges Virtual Tour:
https://psc.edu/bvt

PITTSBURGH SUPERCOMPUTING CENTER

Omni-Path on Bridges

- Works out of the box pretty easily
- Easy to expand the fabric
- Bridges fabric is loosely connected