

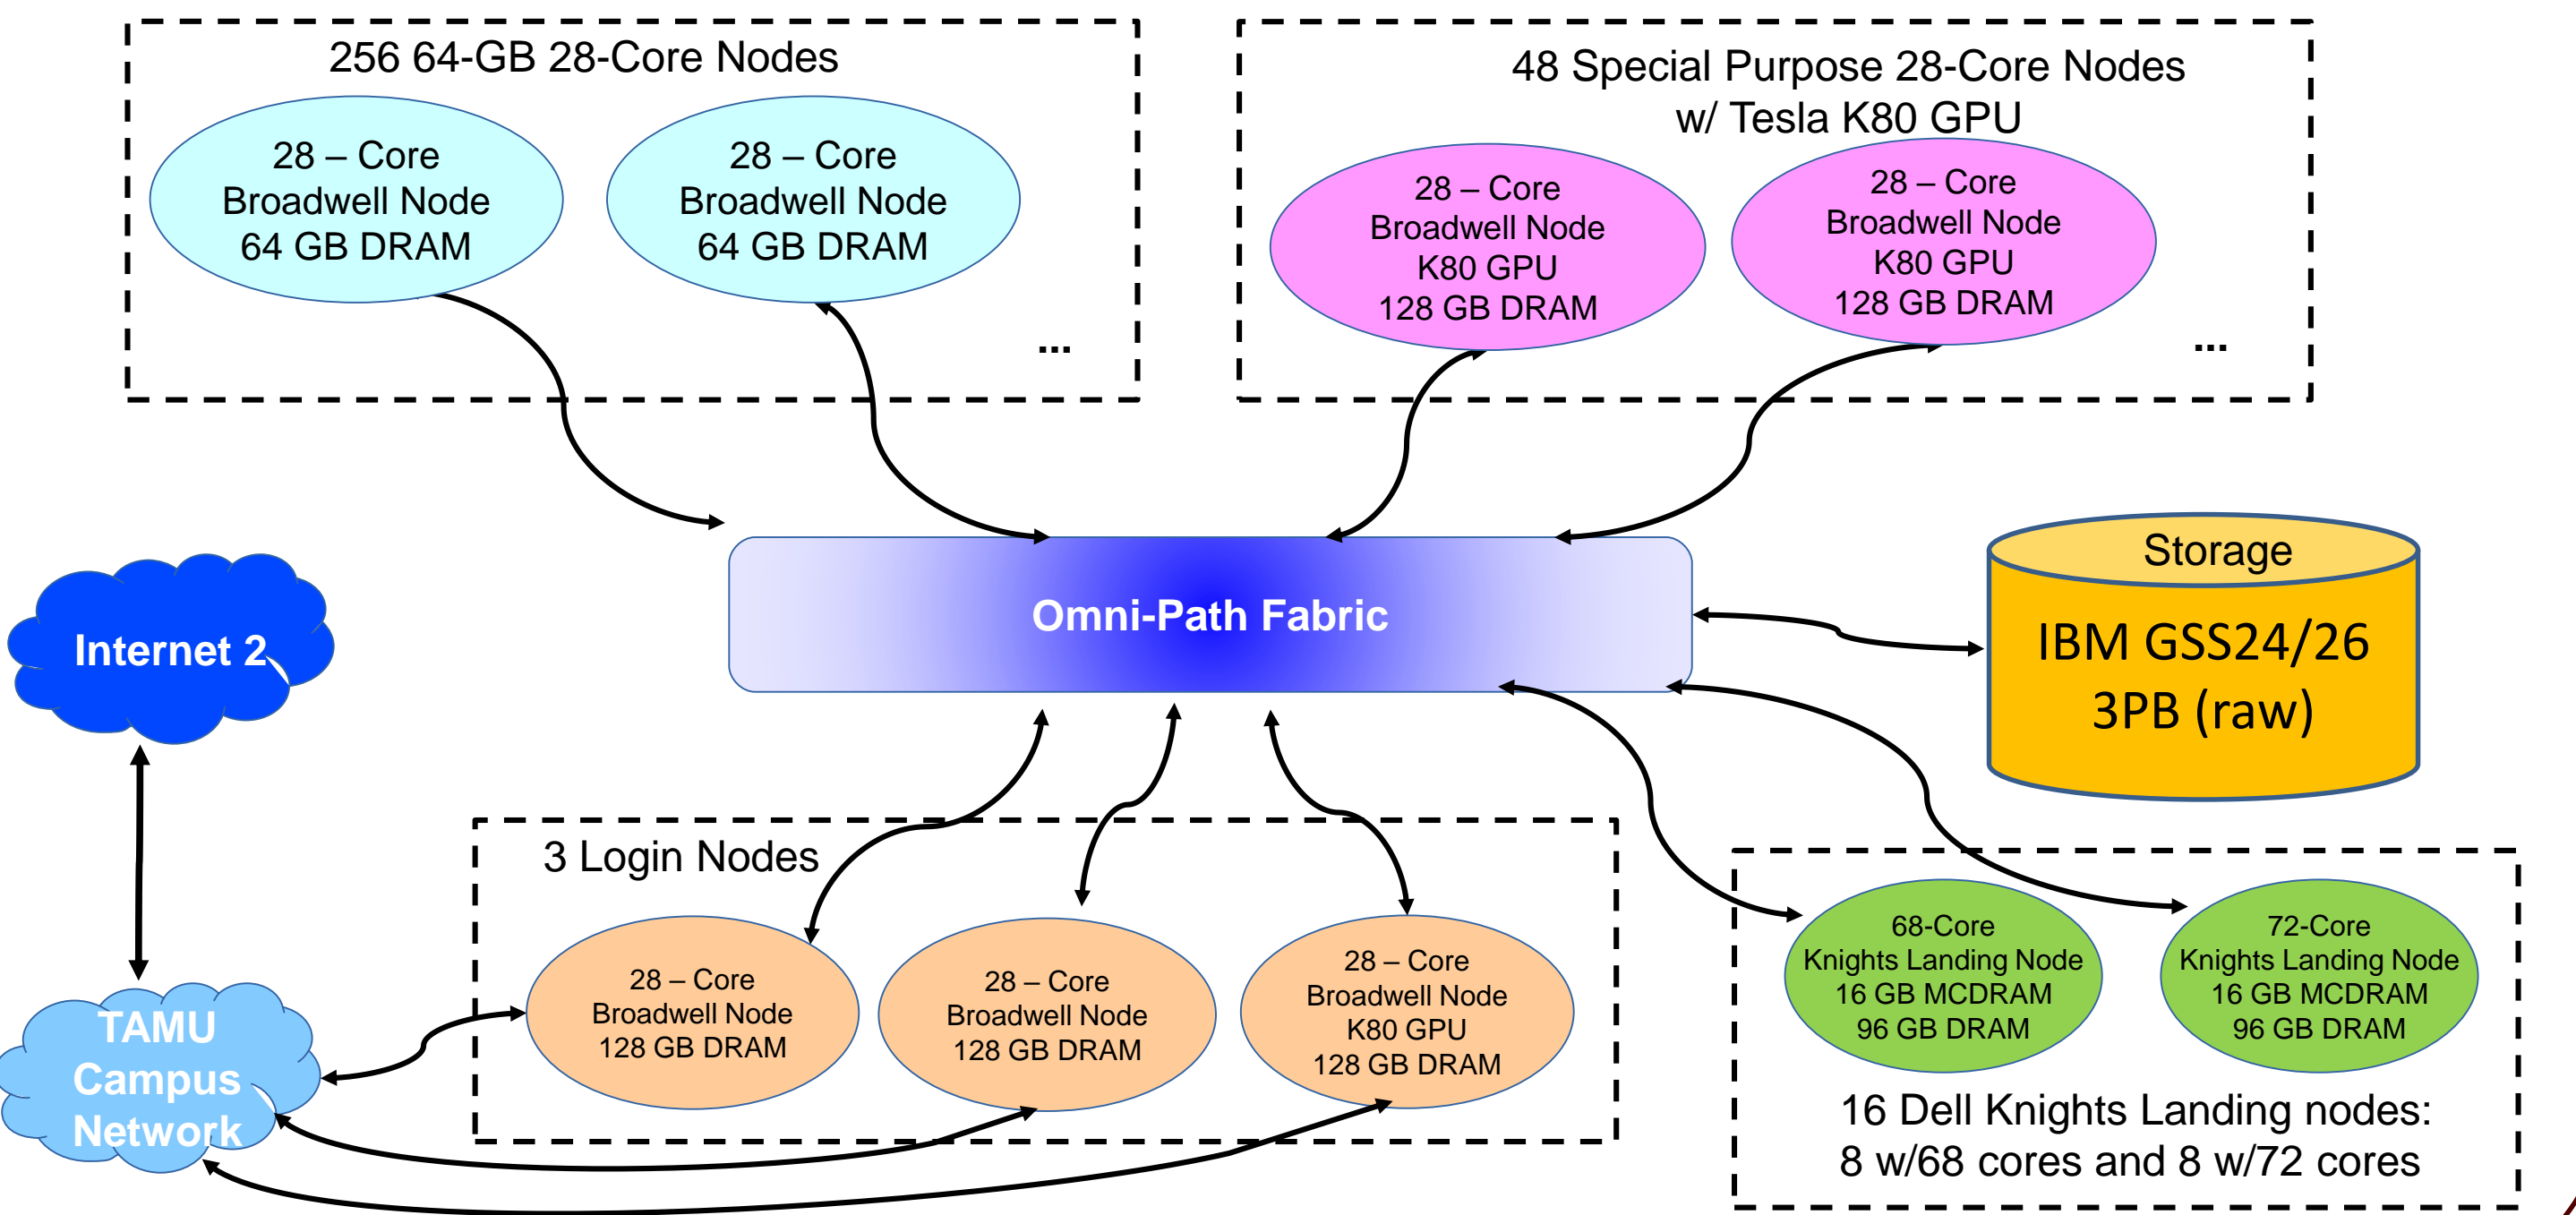


Terra: The 324-node Lenovo x86 Cluster




- Terra is a 8,624-Xeon-core and 1,120-Phi-core Lenovo/Dell commodity cluster with
 - 304 Lenovo compute nodes based on Intel's 64-bit 14-core *Broadwell* processors and 64 GB memory.
 - 16 Dell compute nodes with Intel Knights Landing, 16 GB of MCDRAM and 96 GB of memory (added in Fall 2017)
 - 8 compute nodes with 68-core 1.4 GHz processor
 - 8 compute nodes with 72-core 1.5 GHz processor
- Among these nodes, 48 Lenovo compute nodes have one *Tesla K80* GPU each and 128 GB memory.
- In addition, there are 3 login nodes (one with GPU), each with 128 GB of memory.
- The interconnecting fabric is based on the *Intel Omni-Path* in a two-level fat-tree topology.
- 3 PB (raw) IBM Spectrum Scale (GPFS) file system via GSS24/26 appliance.
- 1 node with 10 Gbps link to Internet 2 (Science DMZ)
- Slurm as scheduler

Terra Schematic:

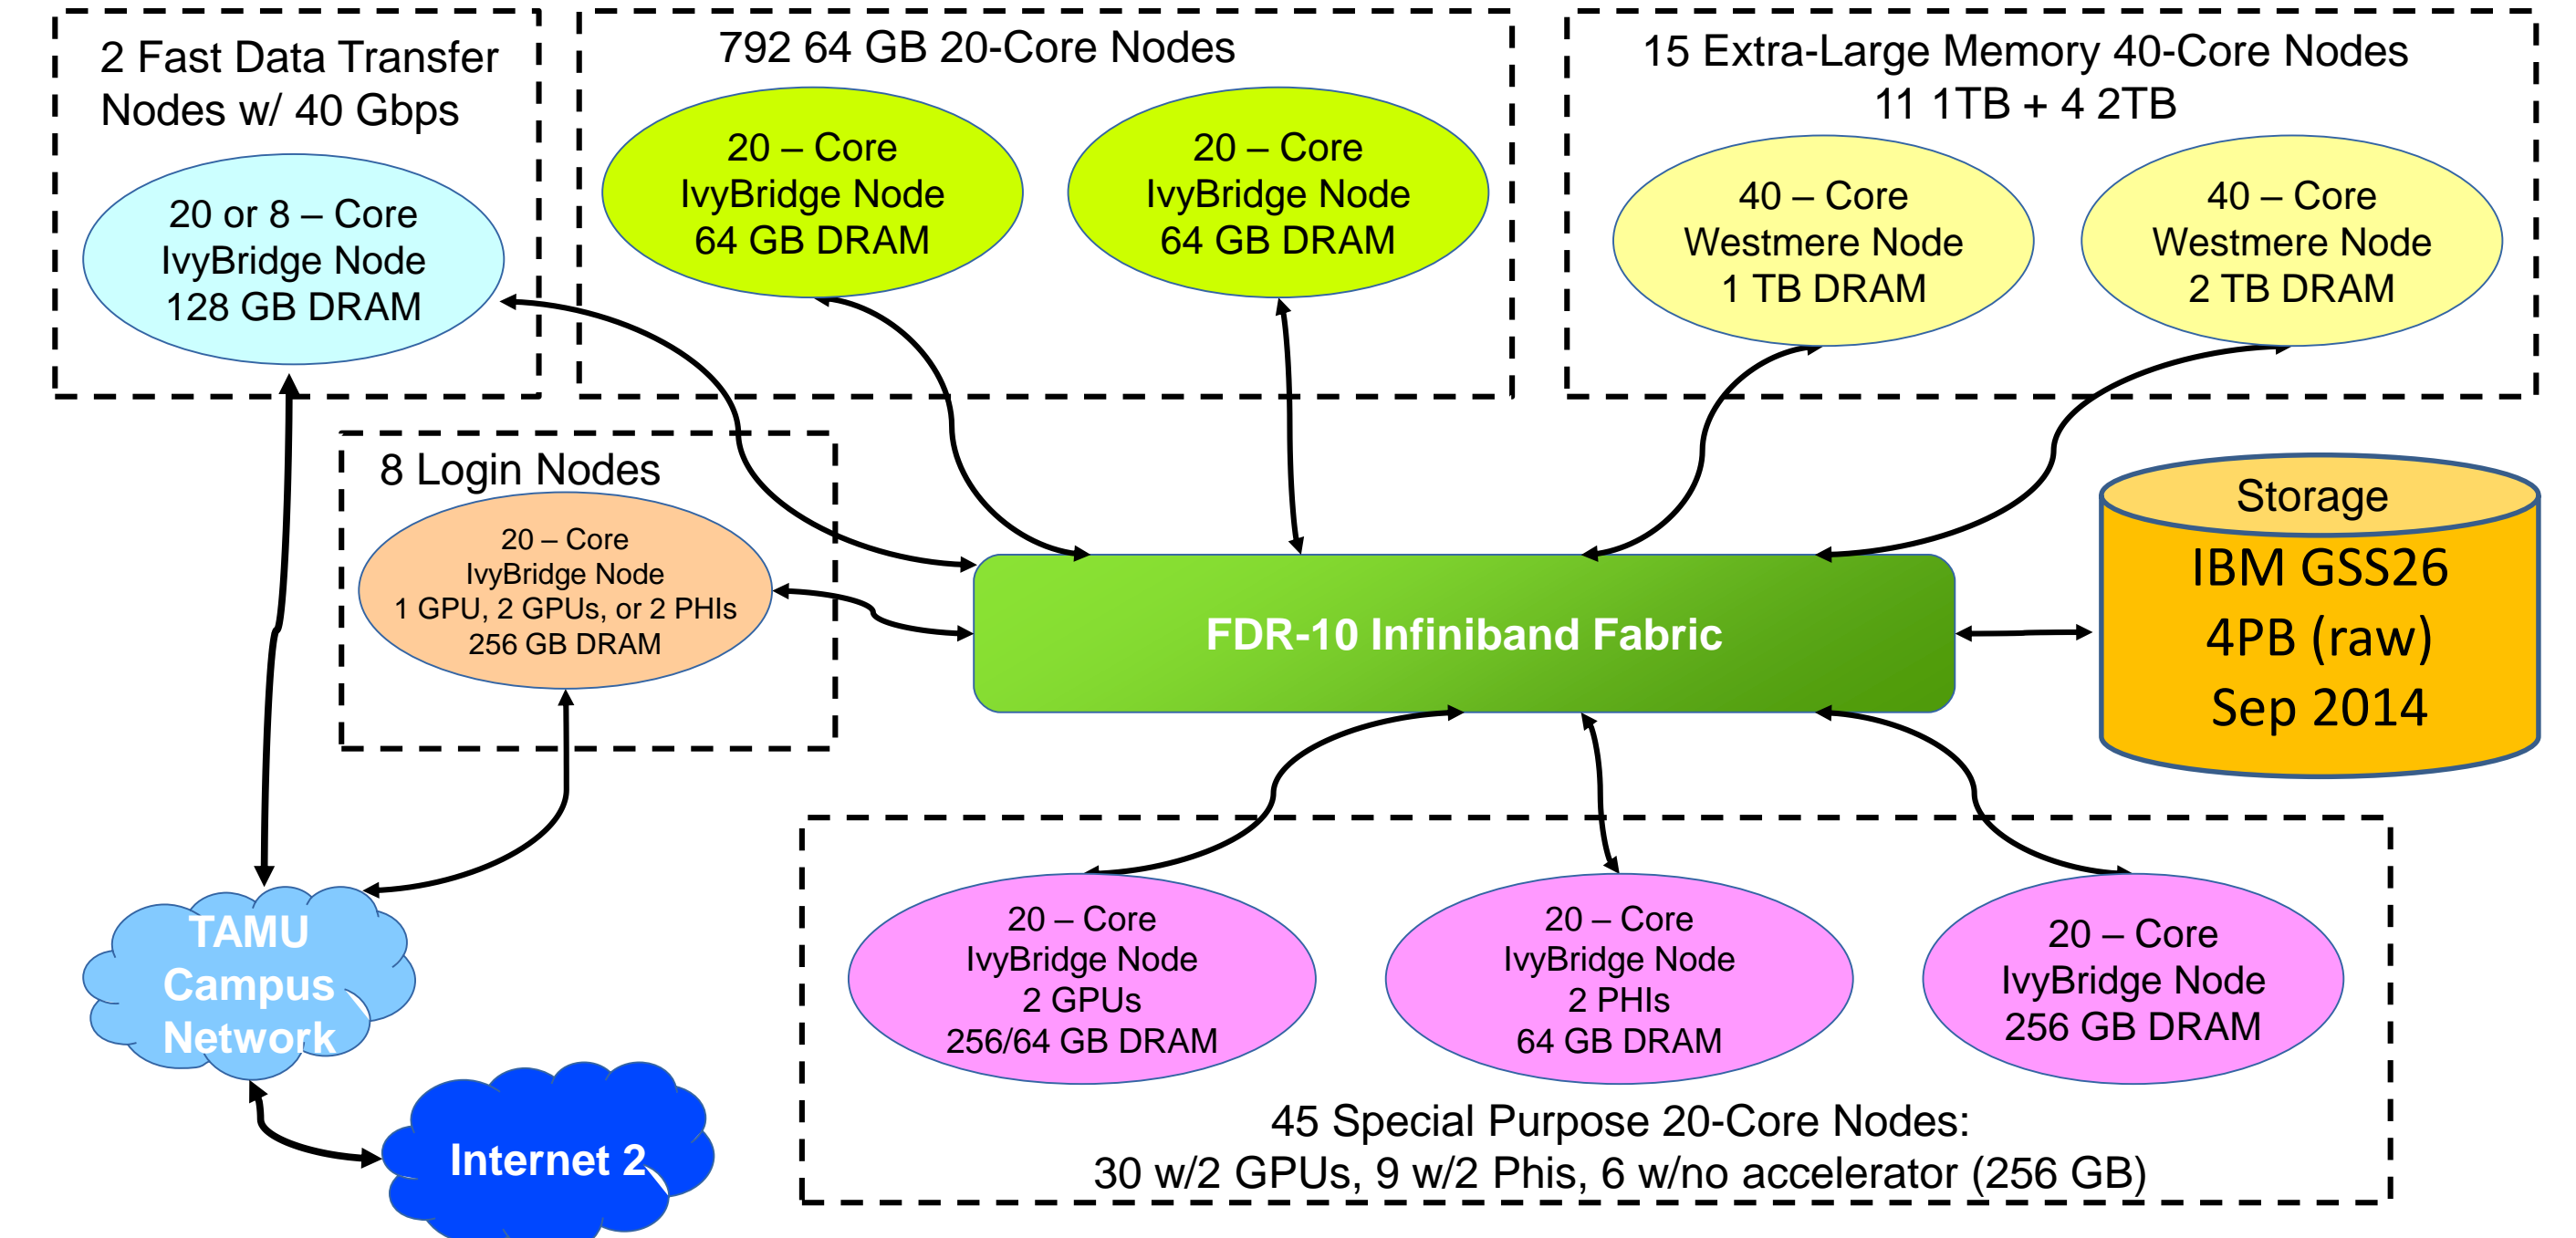


Ada: A 862-node IBM x86 Cluster



- A 17,528-core, 860-node cluster with:
 - 837 20-core compute nodes with two Intel 10-core 2.5GHz *IvyBridge* processors.
- Among these nodes,
 - 30 nodes have 2 GPUs (*K20*) each and 9 nodes have 2 *Phi* coprocessors.
 - 15 compute nodes are 1TB and 2TB memory, 4-processor SMPs with the Intel 10-core 2.26GHz *Westmere* processor.
- In addition, there are 8 20-core login nodes with two Intel 10-core 2.5GHz *IvyBridge* processors and 1 GPU, 2 GPUs, or 2 *Phi* coprocessors
- Nodes are interconnected with *FDR-10 InfiniBand* fabric in a two-level (core switch shown above in middle rack and leaf switches in each compute rack) fat-tree topology.
- 4 PB (raw) IBM Spectrum Scale (GPFS) file system via GSS26 appliance.
- Platform LSF as scheduler.
- 2 nodes with 40 Gbps link to Internet2 (Science DMZ)

Ada Schematic:



Curie: 72-node POWER7+ cluster



- Curie is an 1152-core IBM Power7+ cluster with nodes based on IBM's 64-bit 16-core Power7+ processors.
- In addition to the 70 compute nodes are 2 login nodes with 256GB of memory per node.
- Curie's file system (GPFS) and batch scheduler (LSF) are shared with Ada cluster.

Lonestar 5: 1252-node Cray XC40 cluster



- The latest in a series of Lonestar clusters hosted at TACC, Lonestar 5 is comprised of 1252 Cray XC40 nodes. Jointly funded by The University of Texas System, Texas A&M University, and Texas Tech University, it provides additional resources to Texas A&M researchers. Allocation requests are made through the HPRC request page.


ViDaL: 24-node Compliant Computing cluster

- Texas Virtual Data Library, or Texas ViDaL, a secure and cloud-based computing environment that will allow researchers from many academic disciplines to work with sensitive data from a variety of sources while fully protecting the confidentiality of the information.
- Types of Compliant Computing Available
 - Large RAM nodes: 4 nodes with 1.5TB Ram each
 - Regular RAM nodes: 16 nodes with 192GB Ram each
 - GPU nodes: 4 GPU nodes, each with 192 GB Ram and two NVIDIA V100 GPUs
- Pilot projects starting Jan 2019
- Project web site: <https://vidal.tamu.edu>

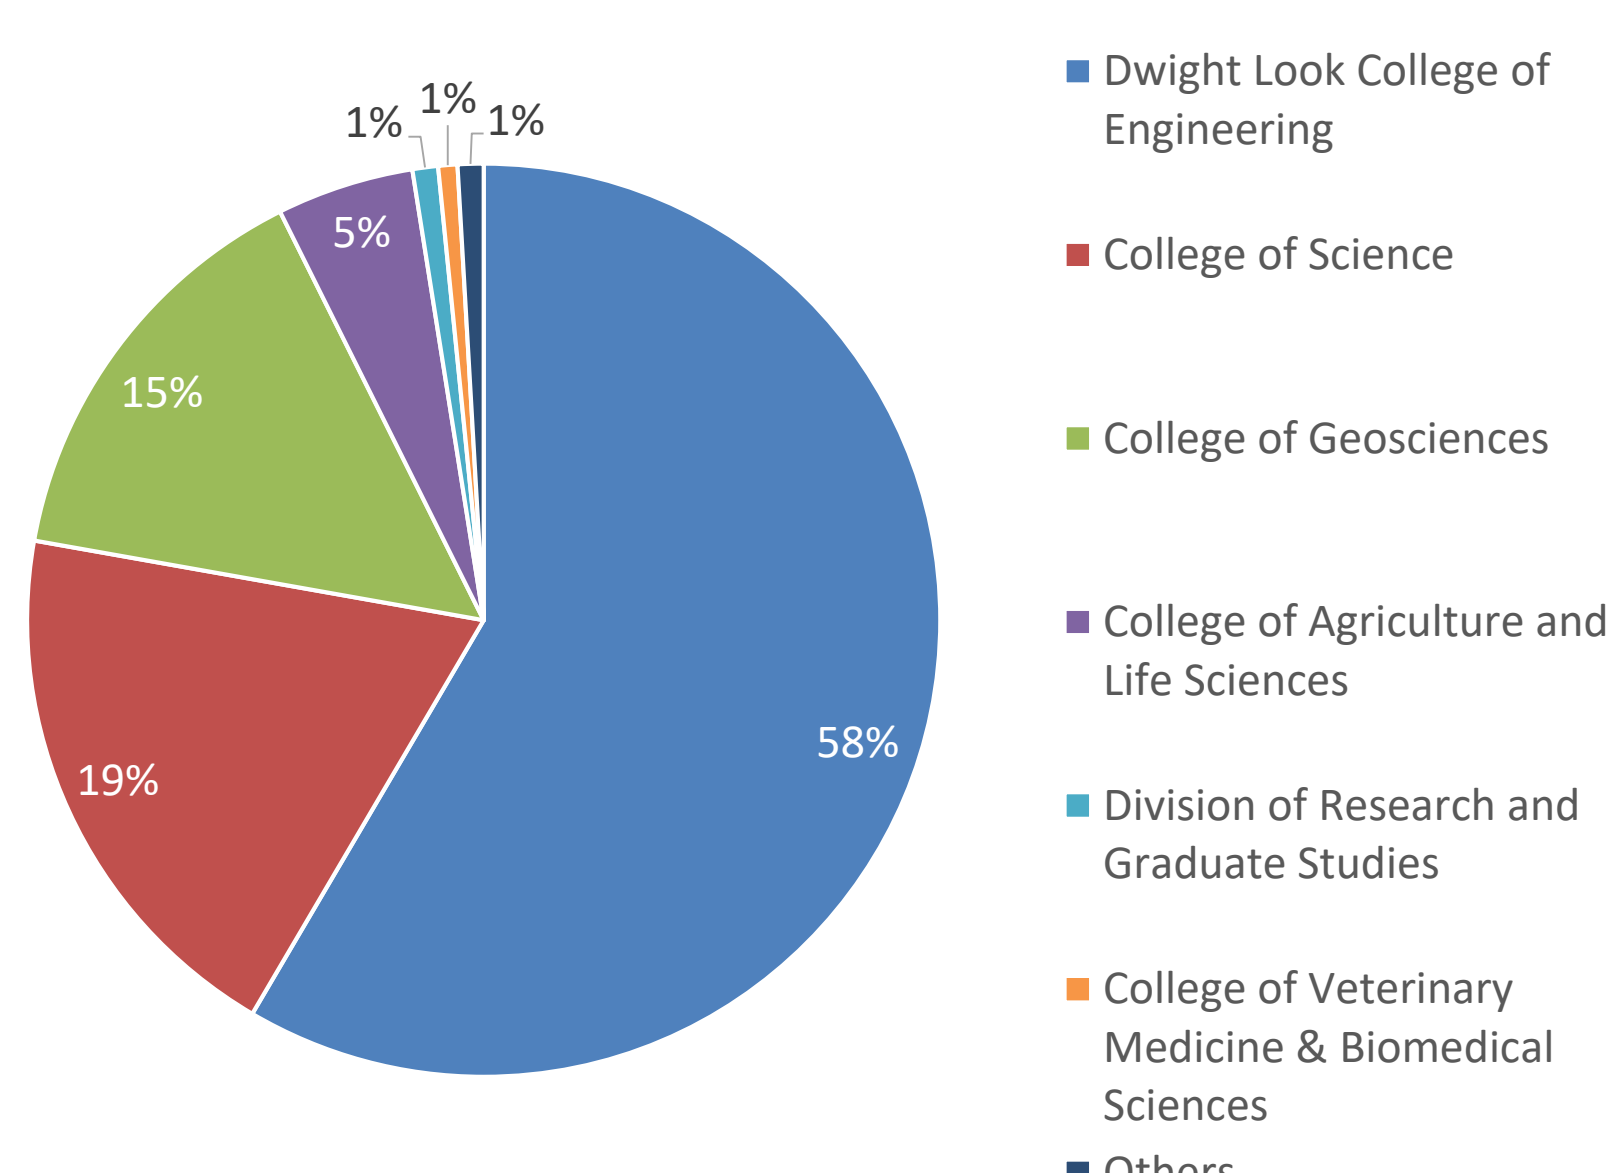


User profile:


FY2018 User Count per College



FY2018 Service Unit Usage per College



Research Highlight



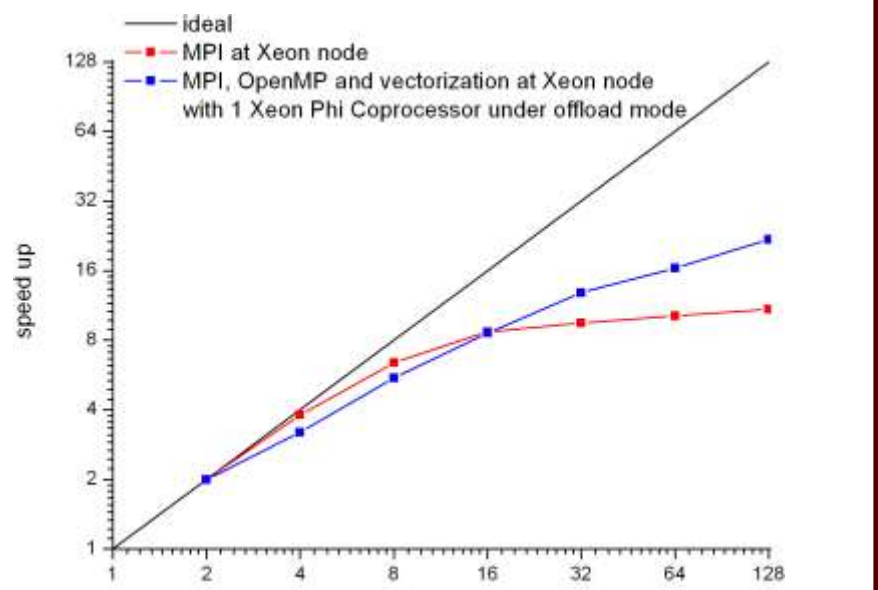

Cover picture: Teng and Hwang, Biophysical Journal, Volume 114, Issue 8, April 24, 2018

HPRC clusters hardware partners:



Intel Parallel Computing Center

HPRC became an Intel Parallel Computing Center in 2016 with primary focus on modernizing OpenFOAM to increase parallelism and scalability through optimizations that leverage Intel Xeon Phi processors.



Intel Parallel Computing Center

SC18 Student Cluster Competition



Six students from the College of Engineering, at the Texas A&M University will participate SC18 Student Cluster Competition, with hardware and software sponsored by Dell, Intel, Nvidia and Mellanox.

