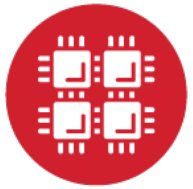

Ohio Supercomputer Center

An OH·TECH Consortium Member

OSC Spring 2017

Brian Guilfoos

Doug Johnson

April 2017 SUG General Meeting

General Agenda

- Organizational Update
- New Services
- Hardware Futures
- Committee Reports

Client Services CY2016

26 Ohio-based universities

41 companies

1,358 clients

218 trainees

8 training opportunities

218 awards made

533 projects served

31 academic courses used OSC's supercomputers

Production Capacity CY2016

115,000,000+ core-hours consumed

3,400,000+ computational jobs

1,101 TB data stored

97.9% up-time (target: 96%)

469 Active Projects CY2016

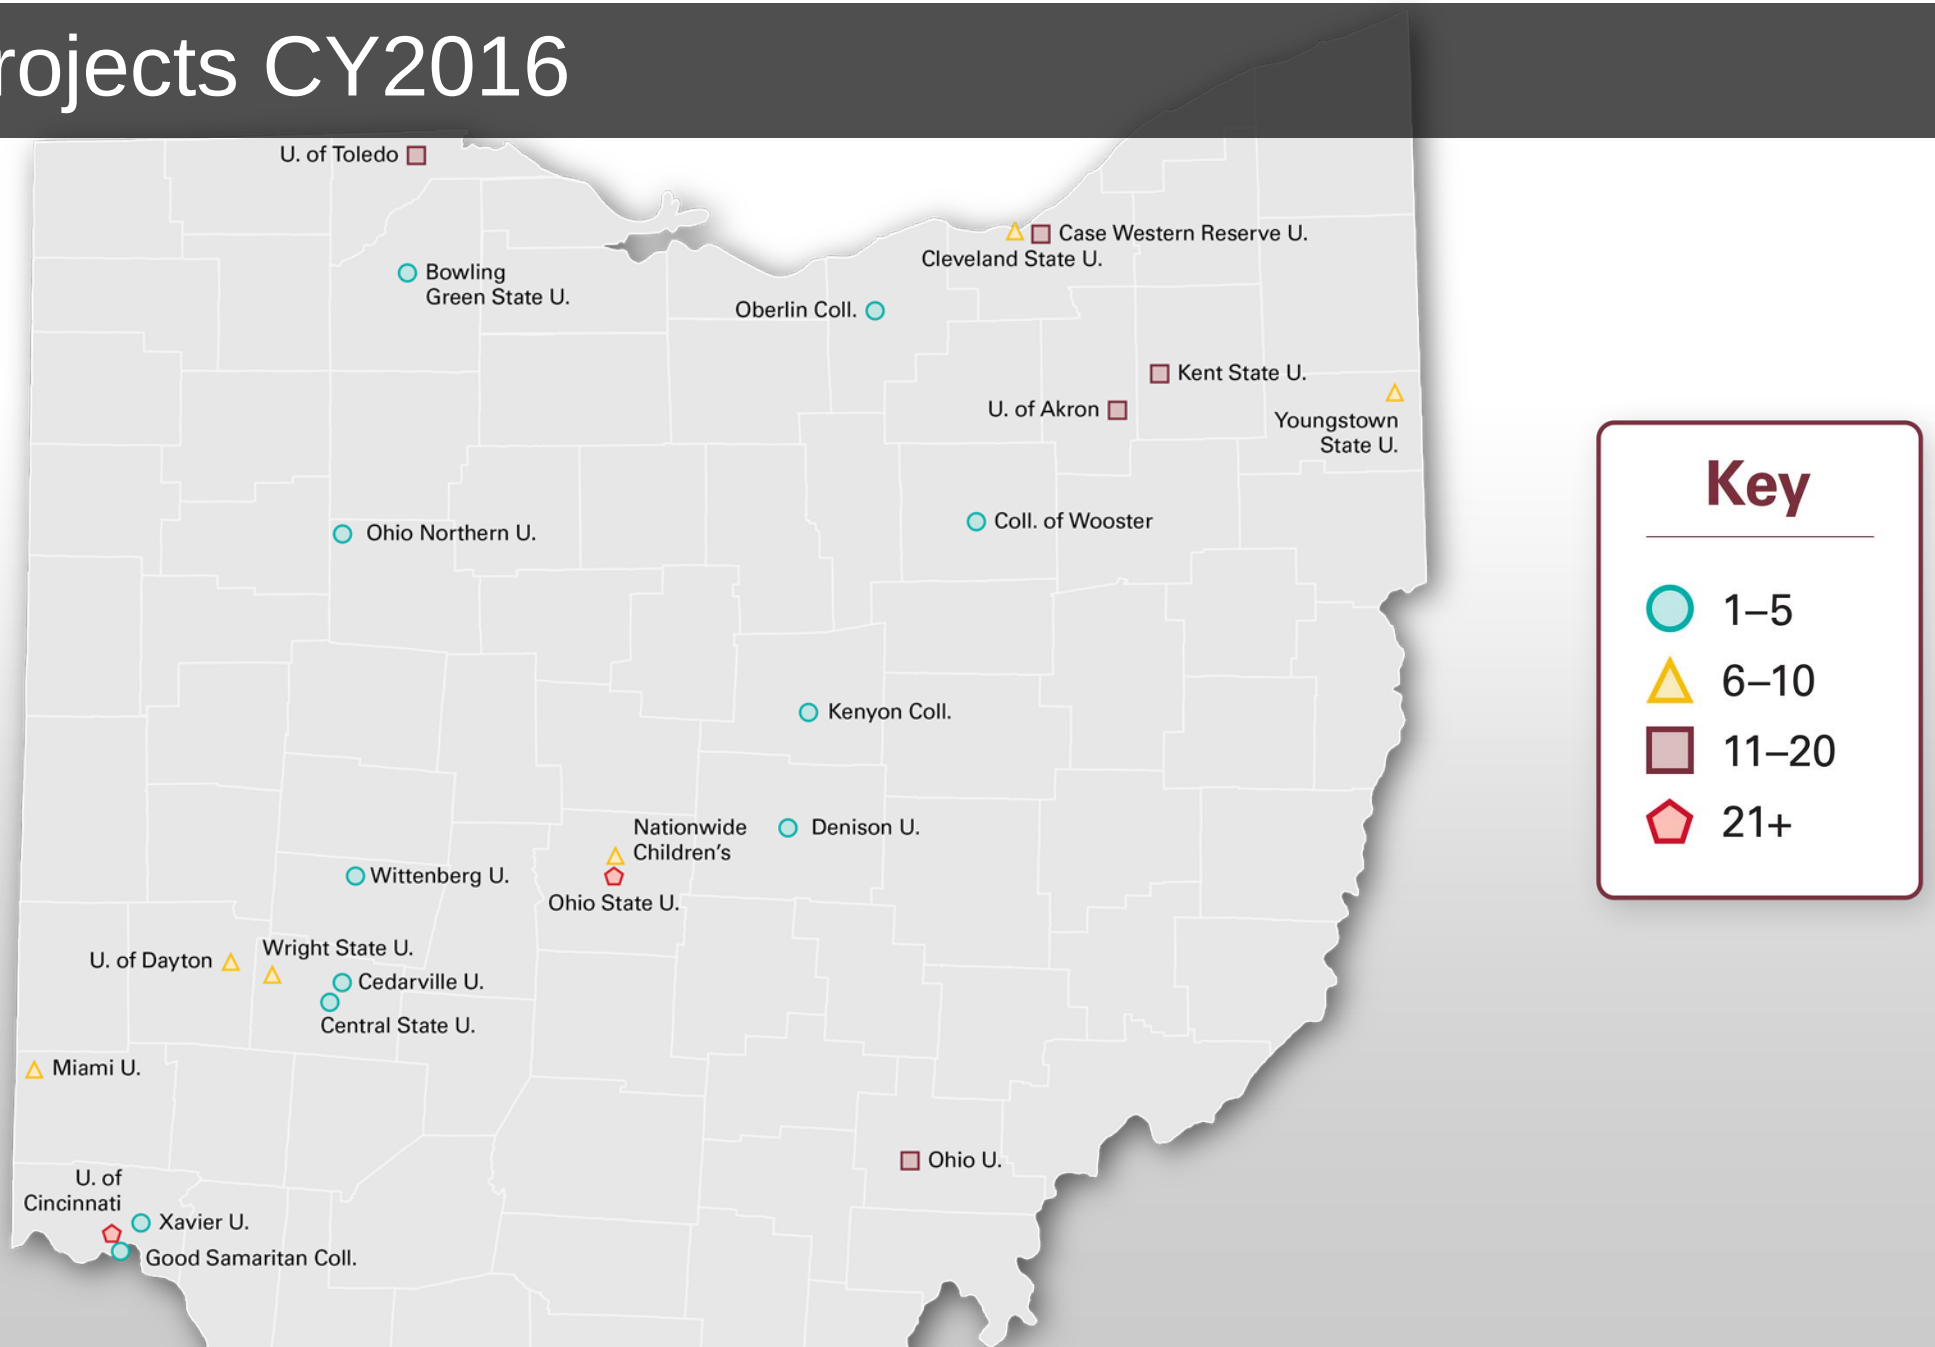

Ohio Academic Resource Network (OARnet) Update

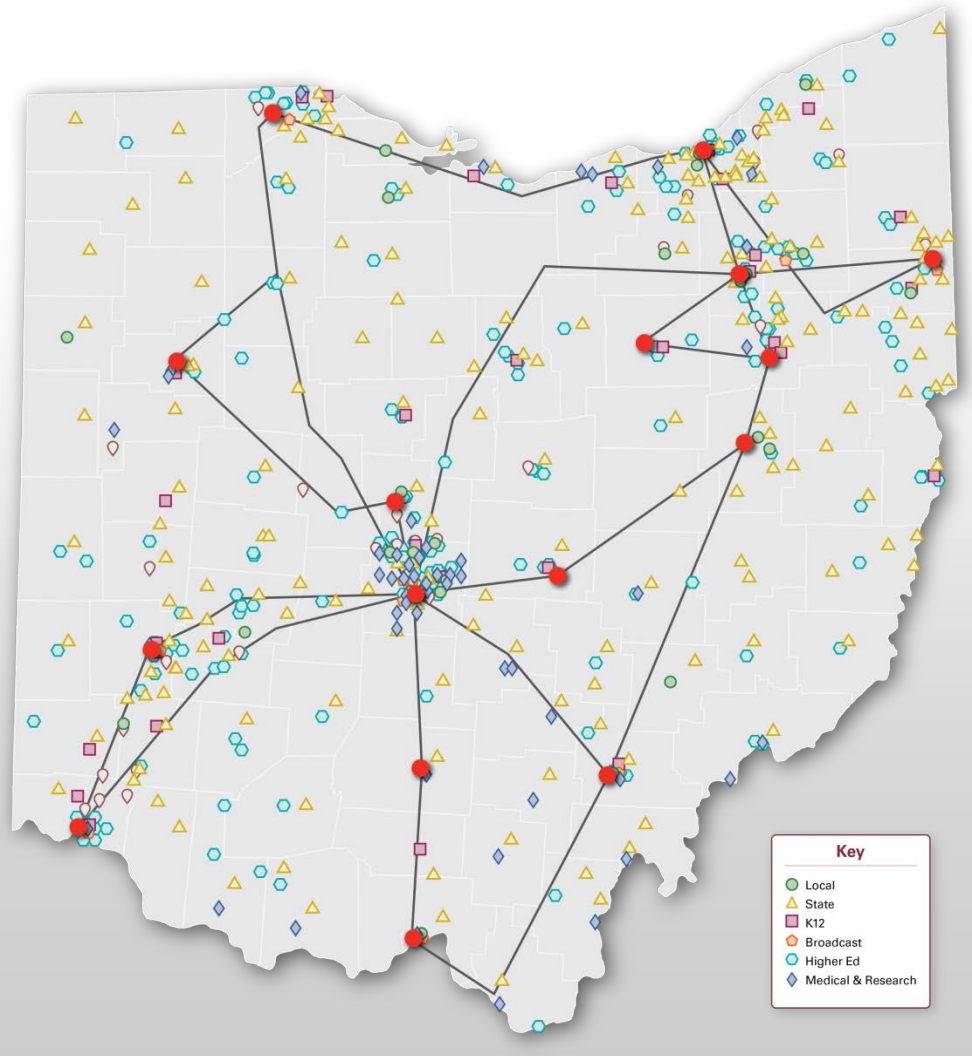

OARnet 100 Gigabet/second network backbone with connected partners

26
K-12 Schools

15
Broadcast Stations

91
Higher Education

35
Local Entities

65
Healthcare Facilities

543
State of Ohio Agency Sites

6
Research Facilities

2,240+ **miles** of high-speed network backbone

$8 **million** savings to Ohio education

$47+ **million** savings through virtualization program

300+ **campuses** of college, university & branch partners

700+ **schools** in public K–12 sector

New Services

- Owens in full production (Dedication ceremony last week!)
- Expansion of GPU services with new NVIDIA P100s
- Expansion of data analytics services
- Interactive applications via the web

Owens Dedication

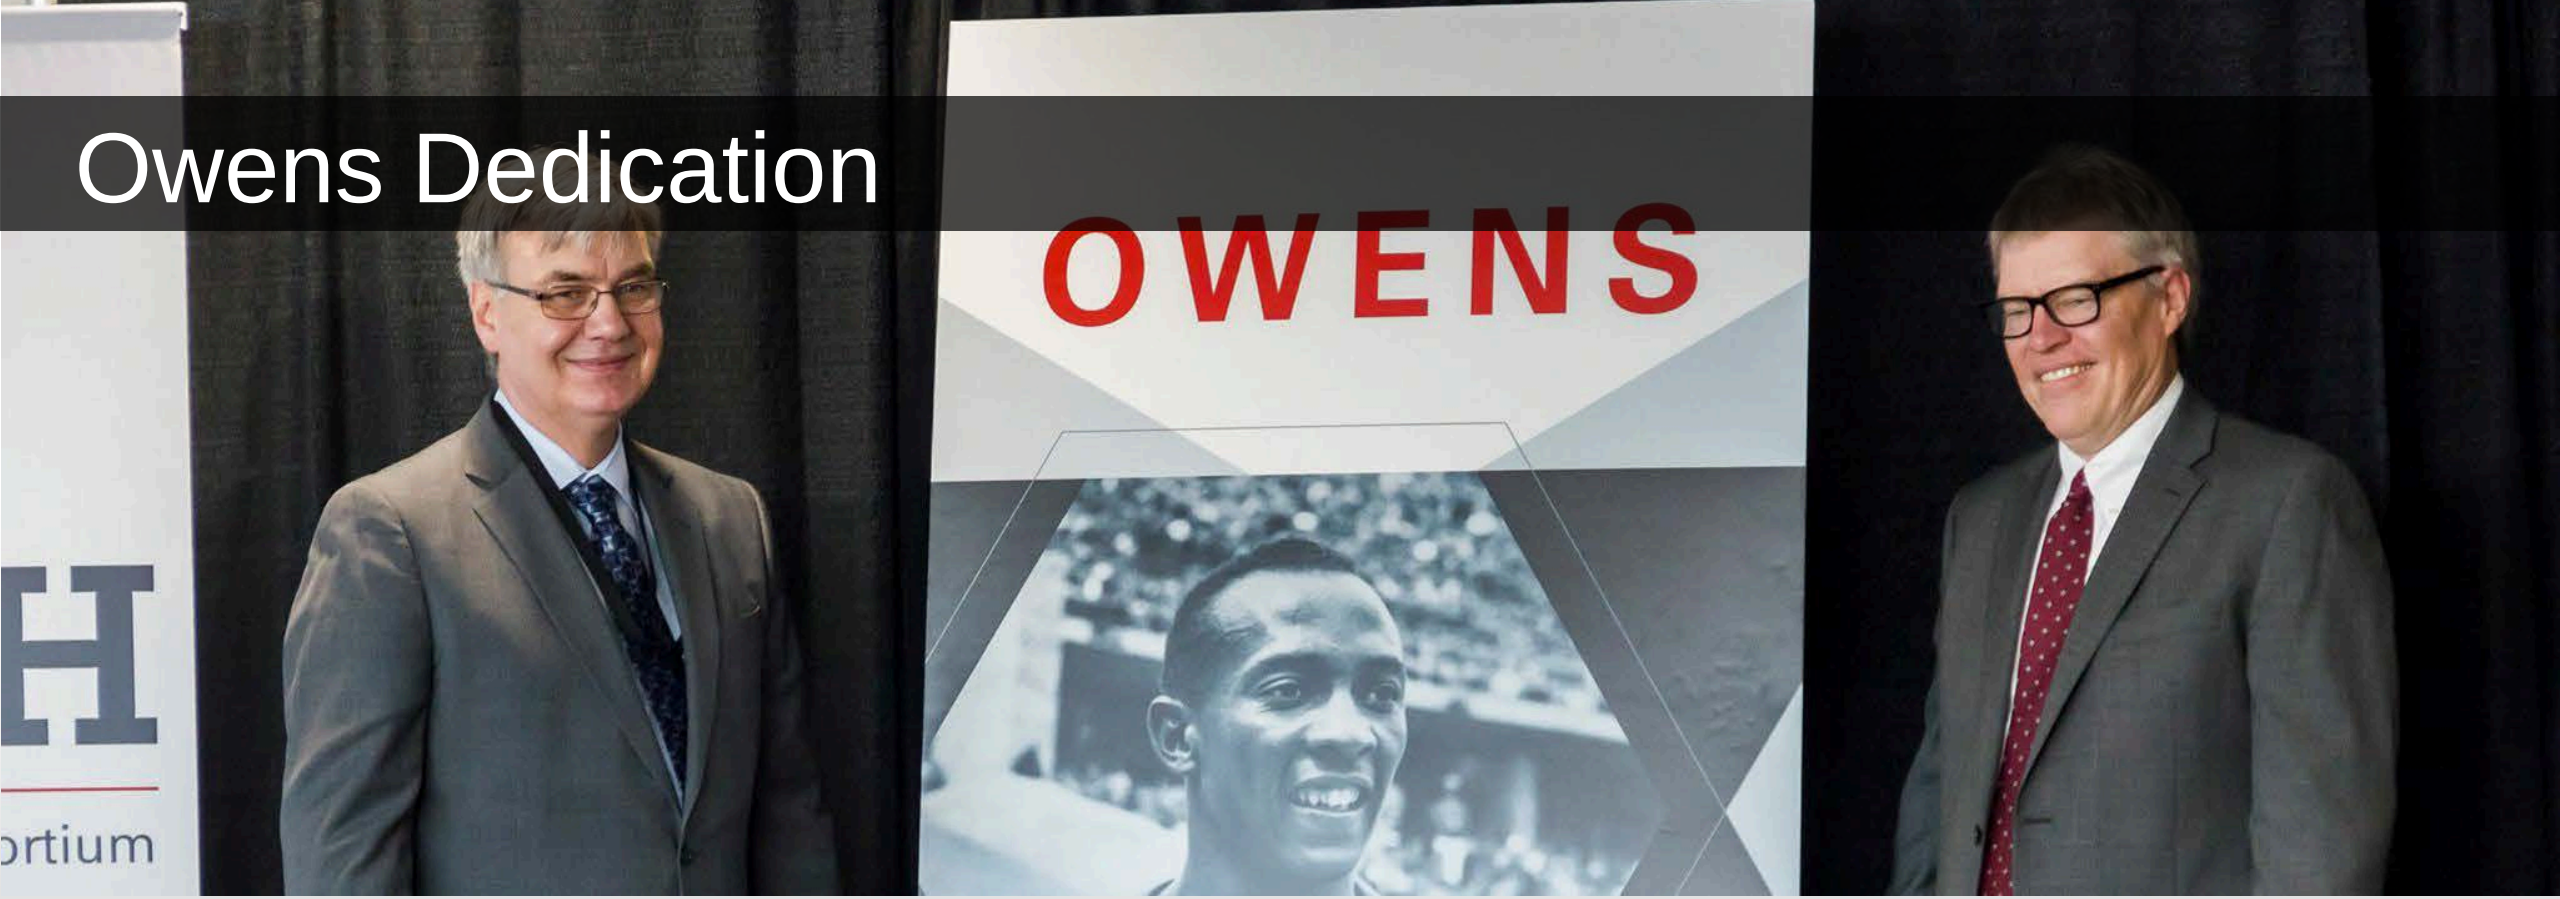

- March 29th, 2017
- Well attended, with representatives from major partners, vendors, and R1 institutions in Ohio

OSC Supercomputer and Storage Services

#202 on the Top 500

	Owens (2016)	Ruby (2014)	Oakley (2012)
Theoretical Performance (TF)	~860	~144	~154
# Nodes	824	240	692
# CPU Cores	23,392	4,800	8,304
Total Memory (TB)	~120	~15.3	~33.4
Memory per Core (GB)	4.5	3.2	4
Interconnect Fabric (IB)	EDR	FDR/EN	QDR

	Capacity (PB)	Bandwidth (GB/s)
Home Storage	0.8	10
Project Storage	3.4	40
Scratch Storage	1.1	100
Tape Library (backup & archive)	5+	3.5

Owens Node Configurations “side-by-side” Comparison

Node Type	Compute	GPGPU	Data Analytics
Node Count	648	160	16
Core Count	28	28	48
Core Type	Broadwell	Broadwell	Haswell
Memory	128 GB	128 GB	1500 GB
Disk	1 TB	1 TB	20 TB
GPU	N/A	P100	None

Owens GPGPU Installation

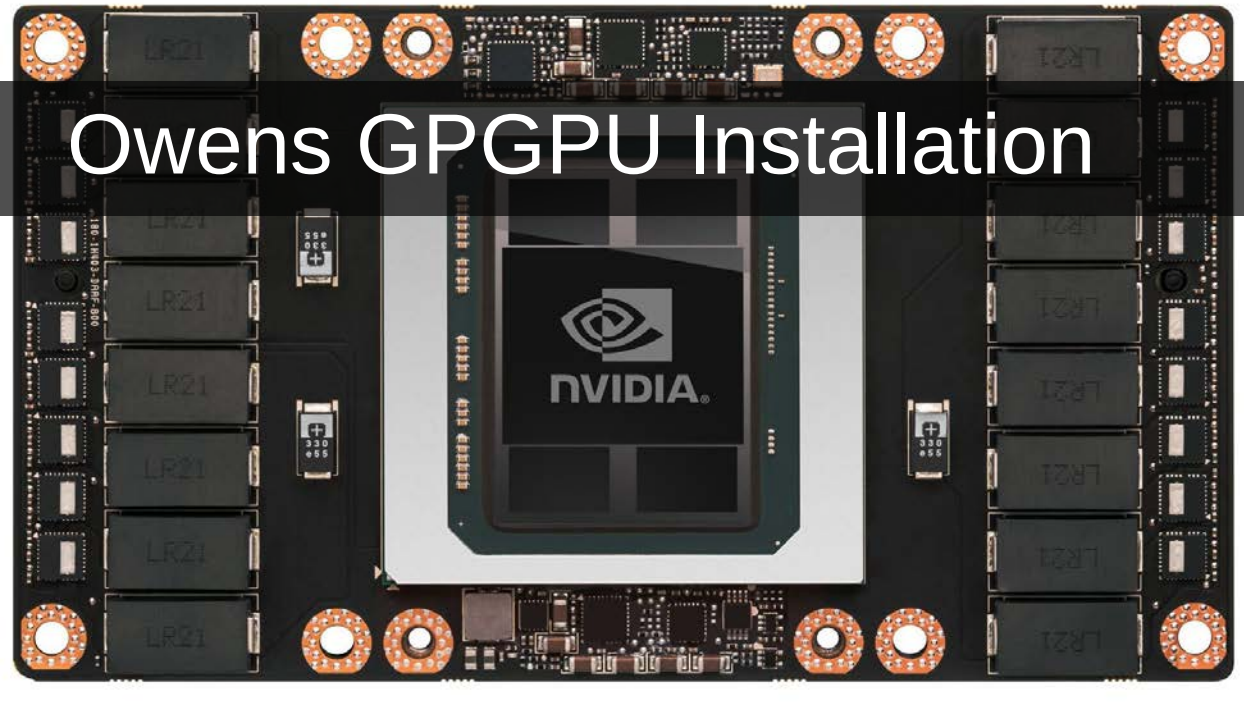

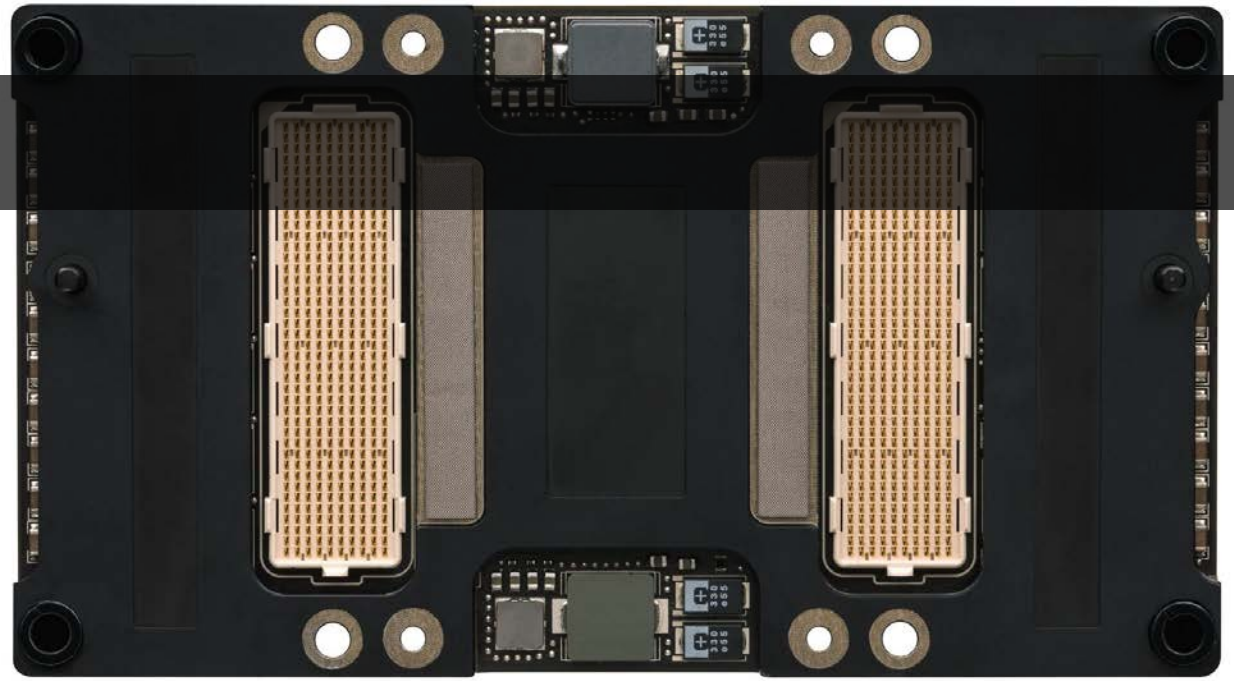

- Tradeoffs: very large number of compute cores, high bandwidth memory
- Model: NVIDIA “Pascal” P100
- Purchase Price: $770K
- Quantity: 160
- Expected Performance: ~750TF (will make Owens ~1.6PF)
- Customer availability now!

GPGPU Example Client Use Cases

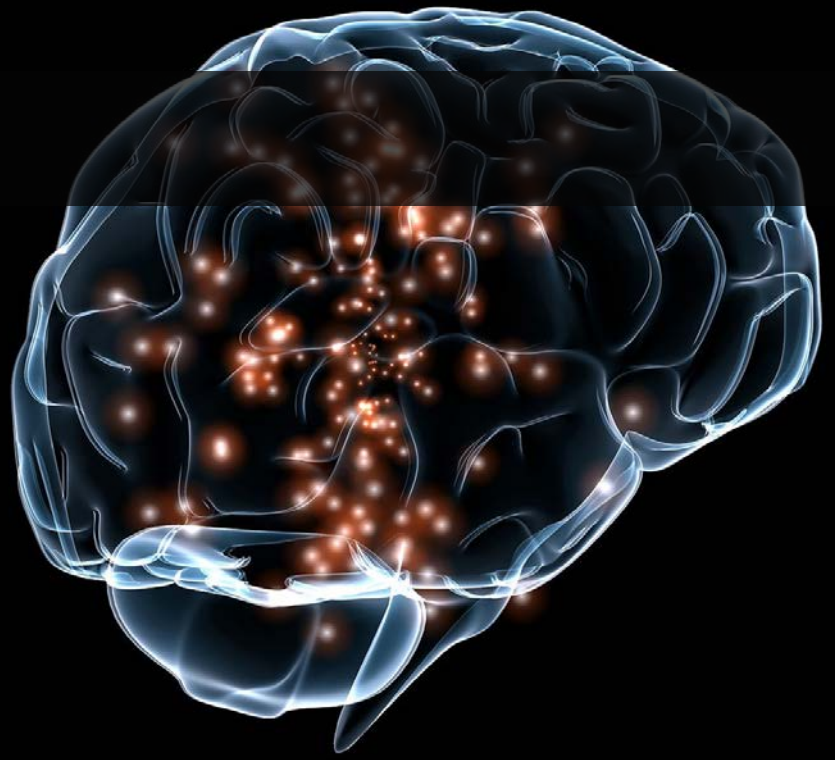

- Molecular Dynamics (MD) Simulations
  - 3X - 7X faster than CPU
  - Materials Science, Biochemistry, Chemistry, Biophysics
  - Software: NAMD, LAMMPS, AMBER, GROMACS
- Machine Learning/Deep Learning
  - 4X - 10x faster for “training” than CPU
  - Wide range of disciplines
  - Software : Caffe, TensorFlow, Torch

Owens Data Analytics Nodes

- Tradeoffs: very large memory, increased core count, large local storage
- Quantity: 16
- Cores: 48 / node (Intel Haswell)
- Memory: 1.5TB / node
- Local Disk: 24TB

Data Analytics Use Cases and Services

Statistics	MYSQL	SPARK
Job data vs user	1.30 hour	5 sec
CPU vs Account	1.25 hour	5 sec
Walltime vs user	1.40 hour	5 sec

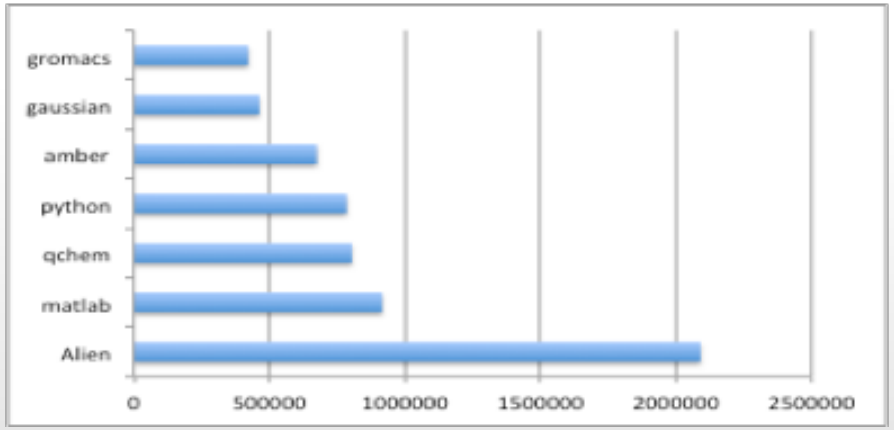

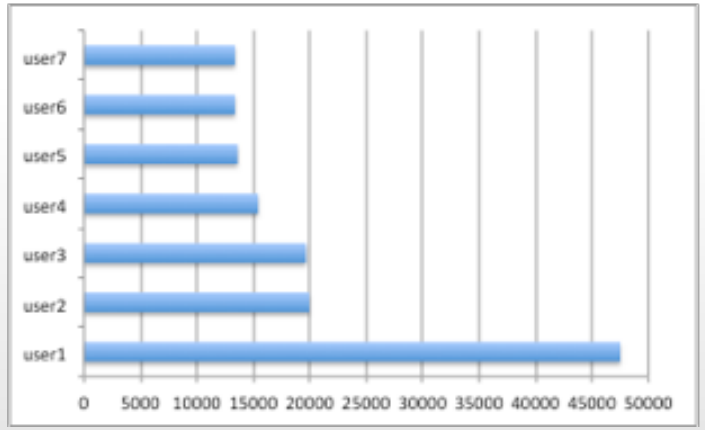

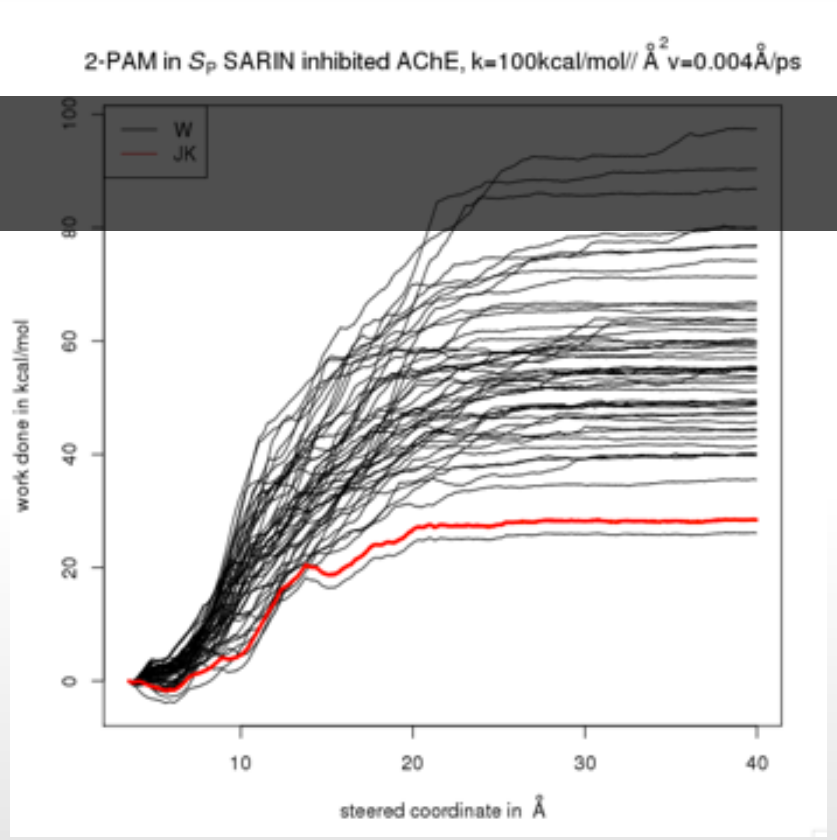

- Analytics on OSC Job data
  - Complex queries on historical job data
  - More than 700x faster than MYSQL query of same data
  - Software: Apache SPARK, PySpark
- Analysis of Simulation results
  - Large data sets from suite of simulation runs
  - Biochemistry/Bioinformatics
  - Software : VMD, R
- Other services: Hadoop, Statistical and mathematical software, high performance storage

Interactive Applications via Web Browser

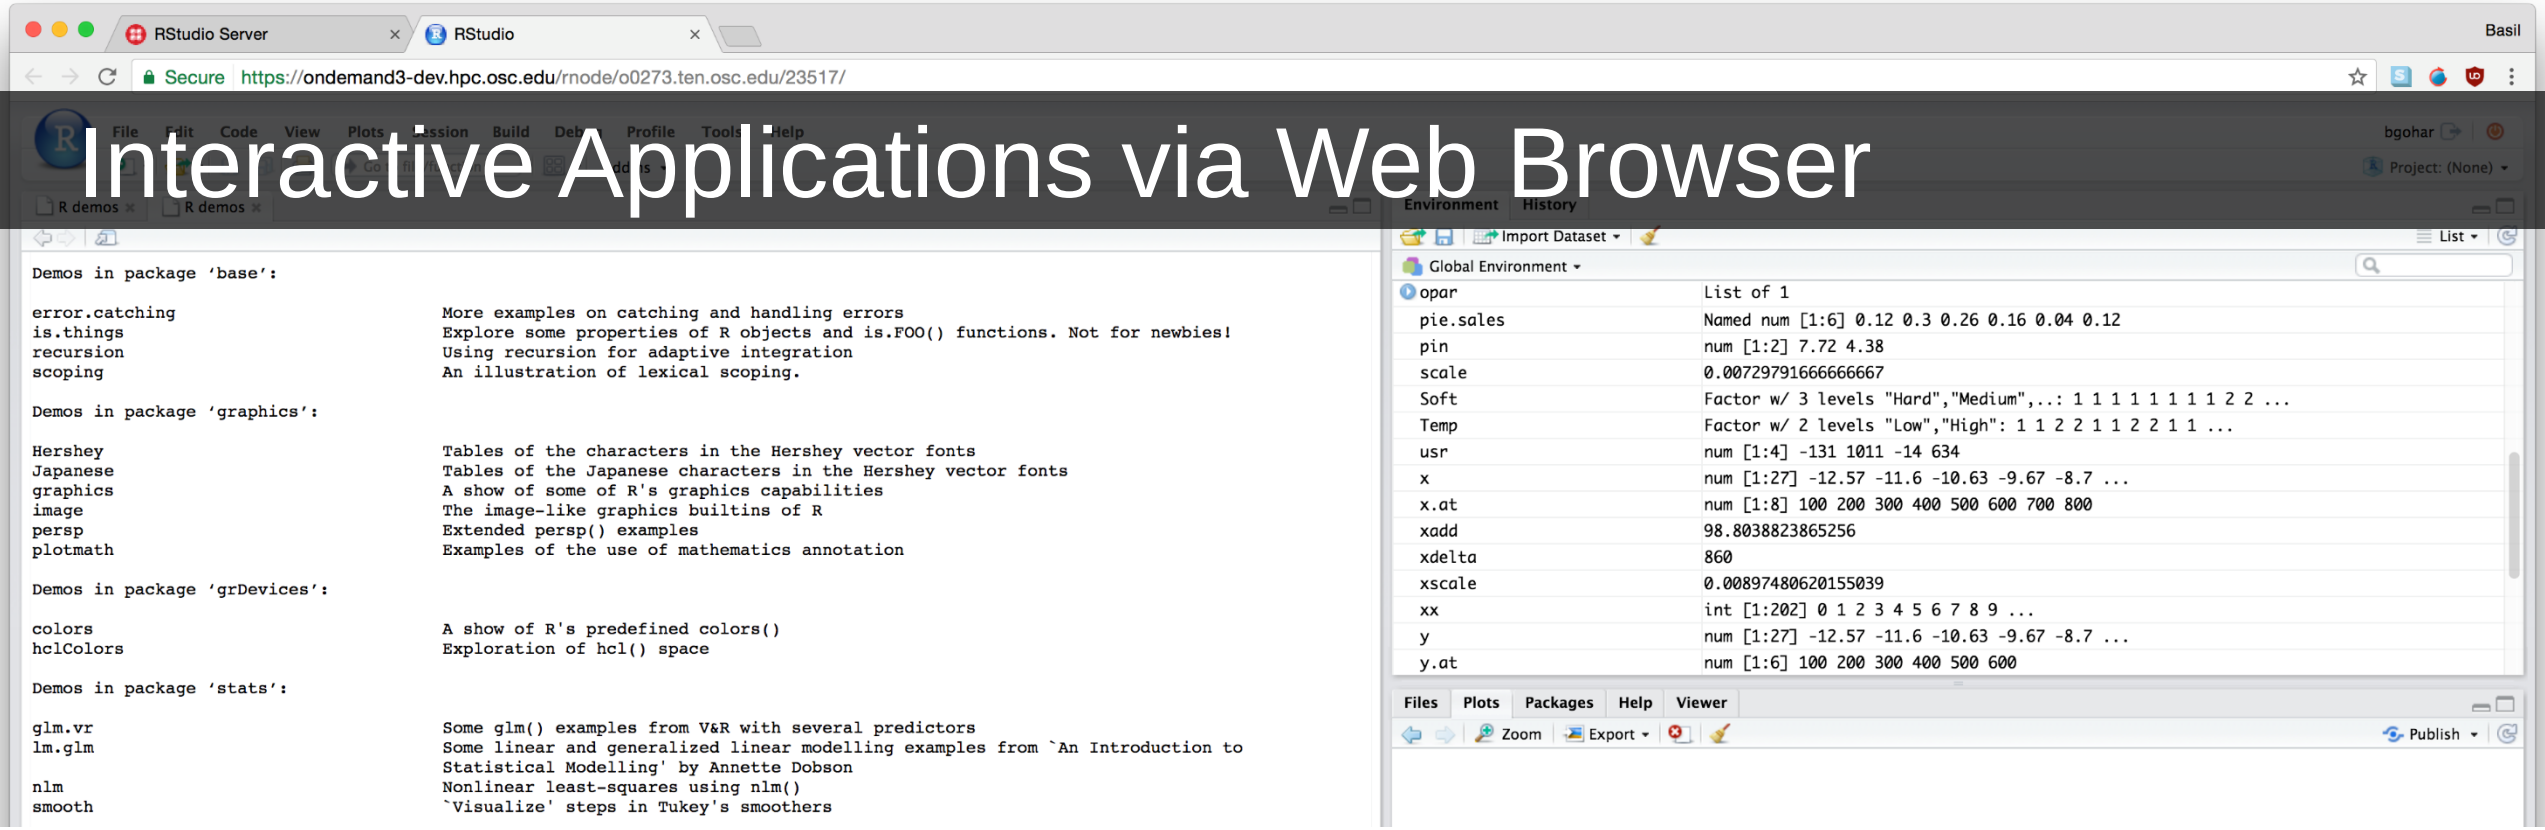

- New capability not available at other supercomputer centers
- Accessible via a web browser with a few clicks through OSC OnDemand
- High performance computing live via a dedicated HPC node(s) (vs. local laptop)
- Currently in Beta testing: Rstudio, Jupyter Notebook for python, MATLAB

DDN Infinite Memory Engine (IME)

- “Burst Buffer” for /fs/scratch file system
- NVMe SSD based storage (same hardware as storage arrays, no spinning media)
- Logically sits between compute nodes, and file system
- Acts as write-back/read cache, or temporary storage
  - Additional tier in storage hierarchy
  - Can smooth peak demand on file system
  - Better suited for small, or unaligned writes than parallel file system

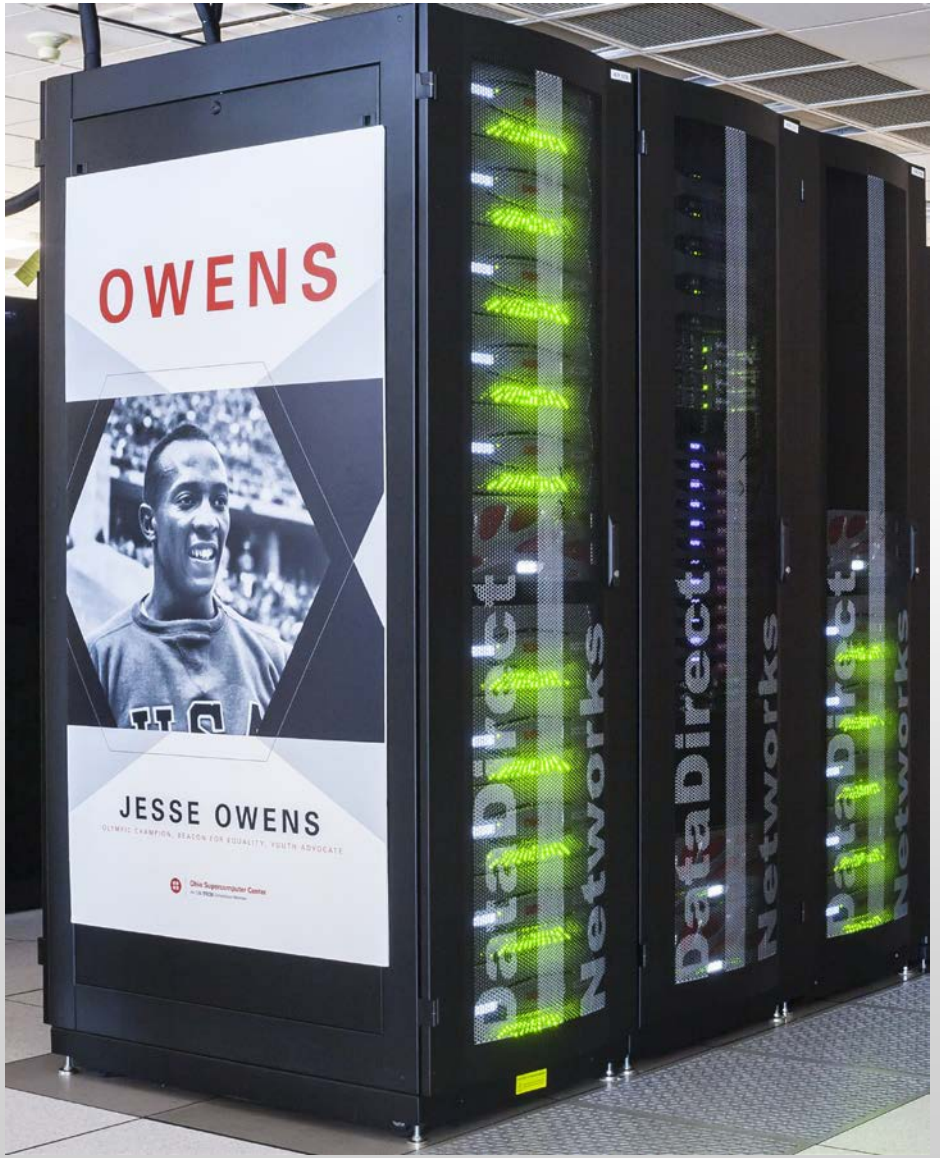

DDN IME Performance, and Status

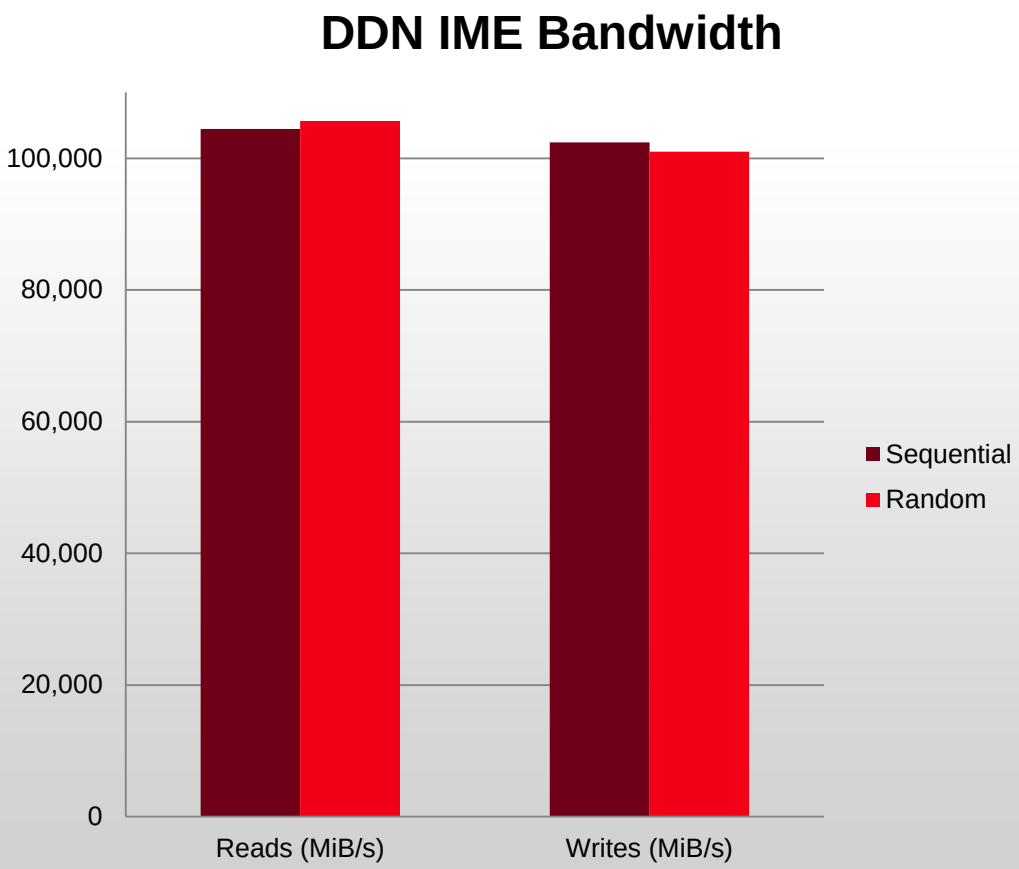

- Capacity: ~40TB
  - Only ½ disk slots populated
- Methods for access
  - POSIX interface, /ime/scratch instead of /fs/scratch
  - Native API
  - MPI-IO (NetCDF, HDF5, etc)
- Data location management not completely automatic
- Still in testing, friendly user availability soon

Hardware Futures

- Compute
  - Oakley decommissioning, and replacement
- Storage
  - Performance and capacity upgrades for backups
  - Infrastructure storage upgrade
  - Project storage expansion, additional tier(s) when needed
- Network
  - Upgrade to 40Gb uplink to OARnet

Upcoming Events

- OSC Workshop: Computing Services to Accelerate Research and Innovation: Thursday, April 13th @UC
- OSC Workshop: Big Data at OSC: Intro to Hadoop and Spark at OSC: Thursday, April 13th @ UC
- XSEDE Workshop: MPI: Tuesday April 18th & Thursday April 19th
- Scratch Policy public comment period closes: Friday, April 28th
- Client Survey currently open
- Office Hours at OSU's Research Commons (alternating Tuesdays) – in person or remote
- OSC 30th Anniversary: TBD (Fall)

Ohio Supercomputer Center

An OH·TECH Consortium Member

Committee Reports

Allocations Committee

- Allocations:
  - 7.7M+ RUs allocated
  - Reviewed 215 applications
    - 25 discovery-level
    - 14 major-level
    - 28 standard-level
    - 5 emerita
    - 25 classroom
    - 118 startup
  - 22 institutions allocated RUs

- Annual Allocations (for CY16):
  - 4.9M+ RUs allocated
  - 7 institutions

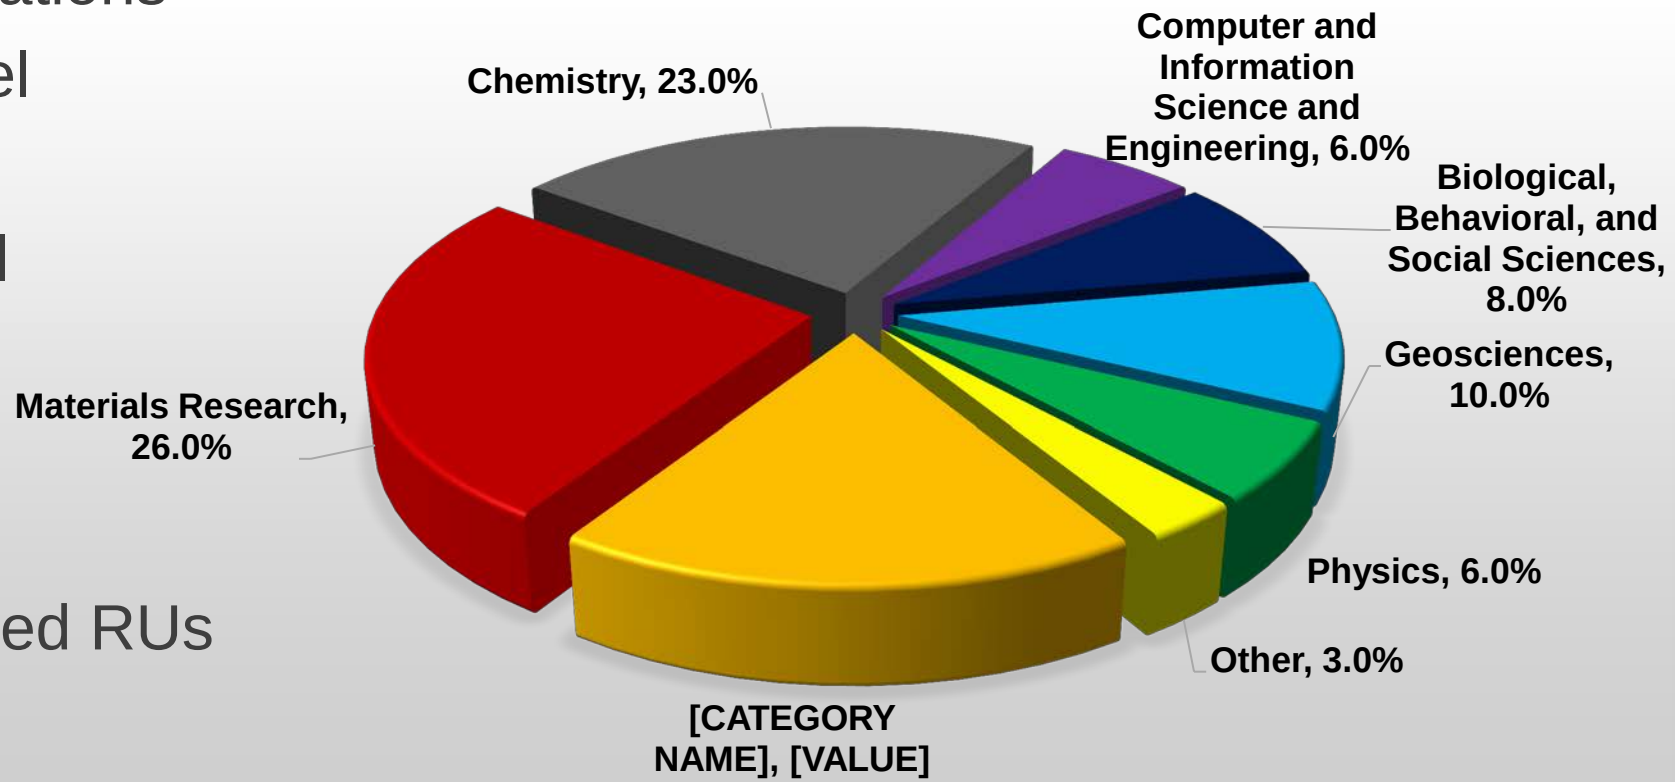

Resource Units consumed by Field of Science, CY16

Hardware Committee

John Heimaster, Committee Chair

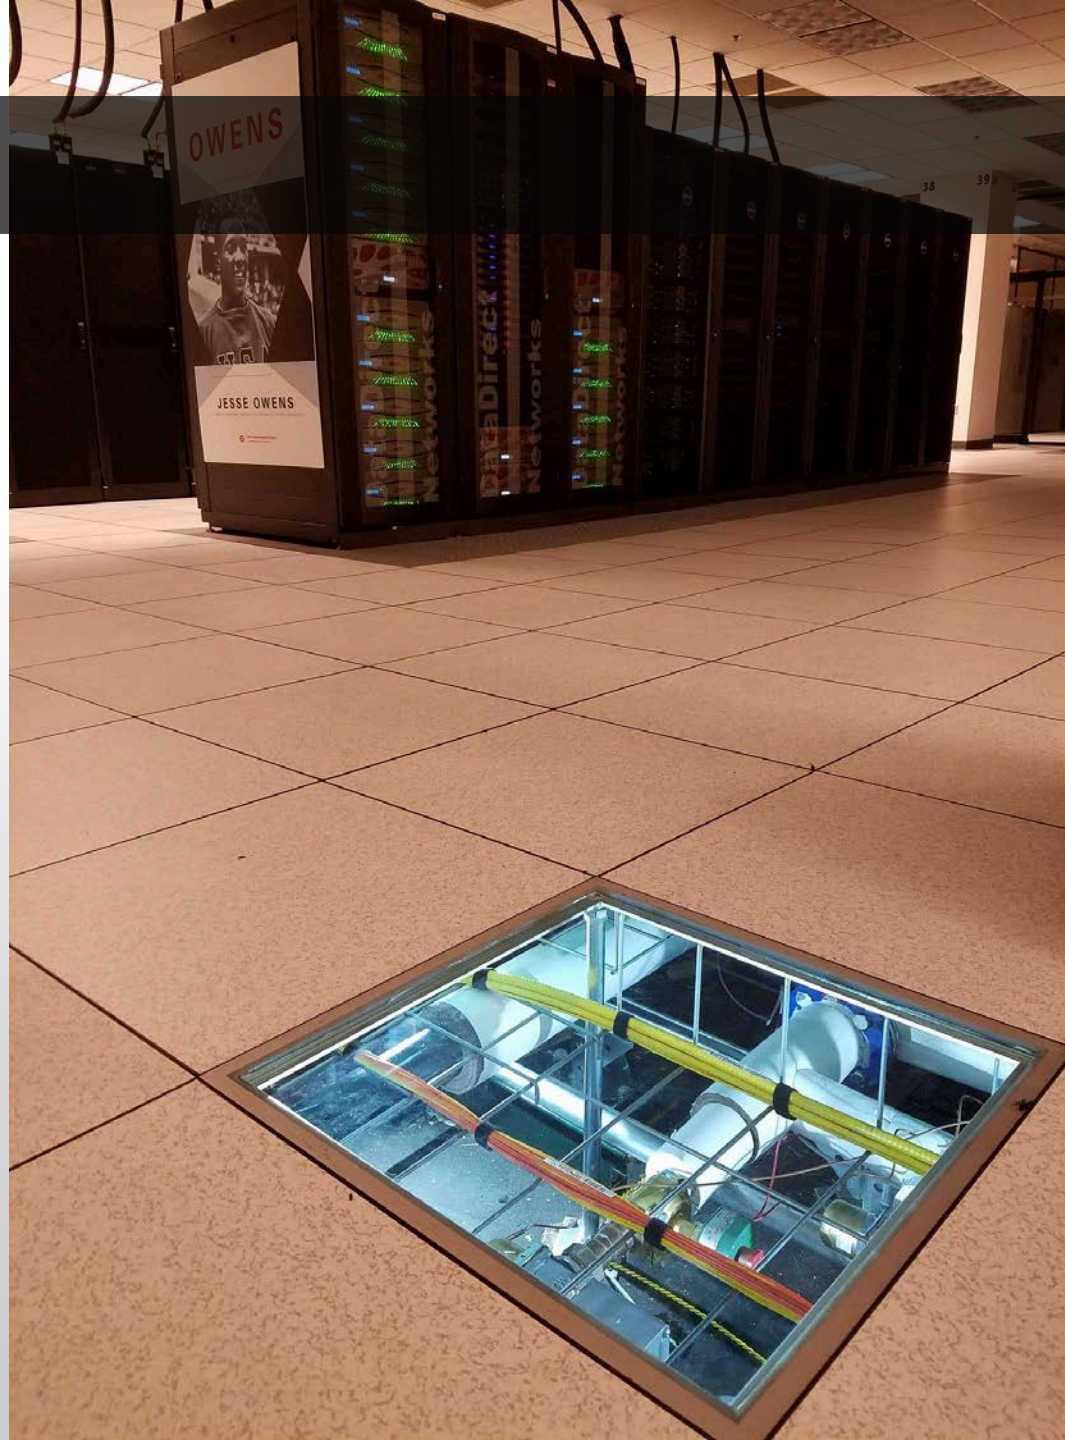

Software Committee

- **New Purchase Discussion**
  - Comsol Server: provide non-OSU users
  - Debugger: Totalview vs. DDT
- **Third party hosting**
  - Matlab
    - Any academic users in Ohio can use Matlab on OSC as part of our license.
- **Renewal (Since Oct 2016 meeting)**
  - Abaqus, pgi, CSD, Gaussian, Intel Cluster (Capital), MDCS, Turbomole
- **Upcoming (during 2017 calendar year)**
  - Discontinue?: CSD (Cambridge Structural Database)
  - Capital purchase: Totalview, Ansys
  - Regular renewal: Amber, Comsol, Q-chem, Star-CCM+, Schrodinger, Allinea, abaqus, pgi, Gaussian