OSC Updates

SUG

October 4, 2018

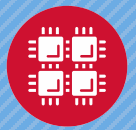

Agenda

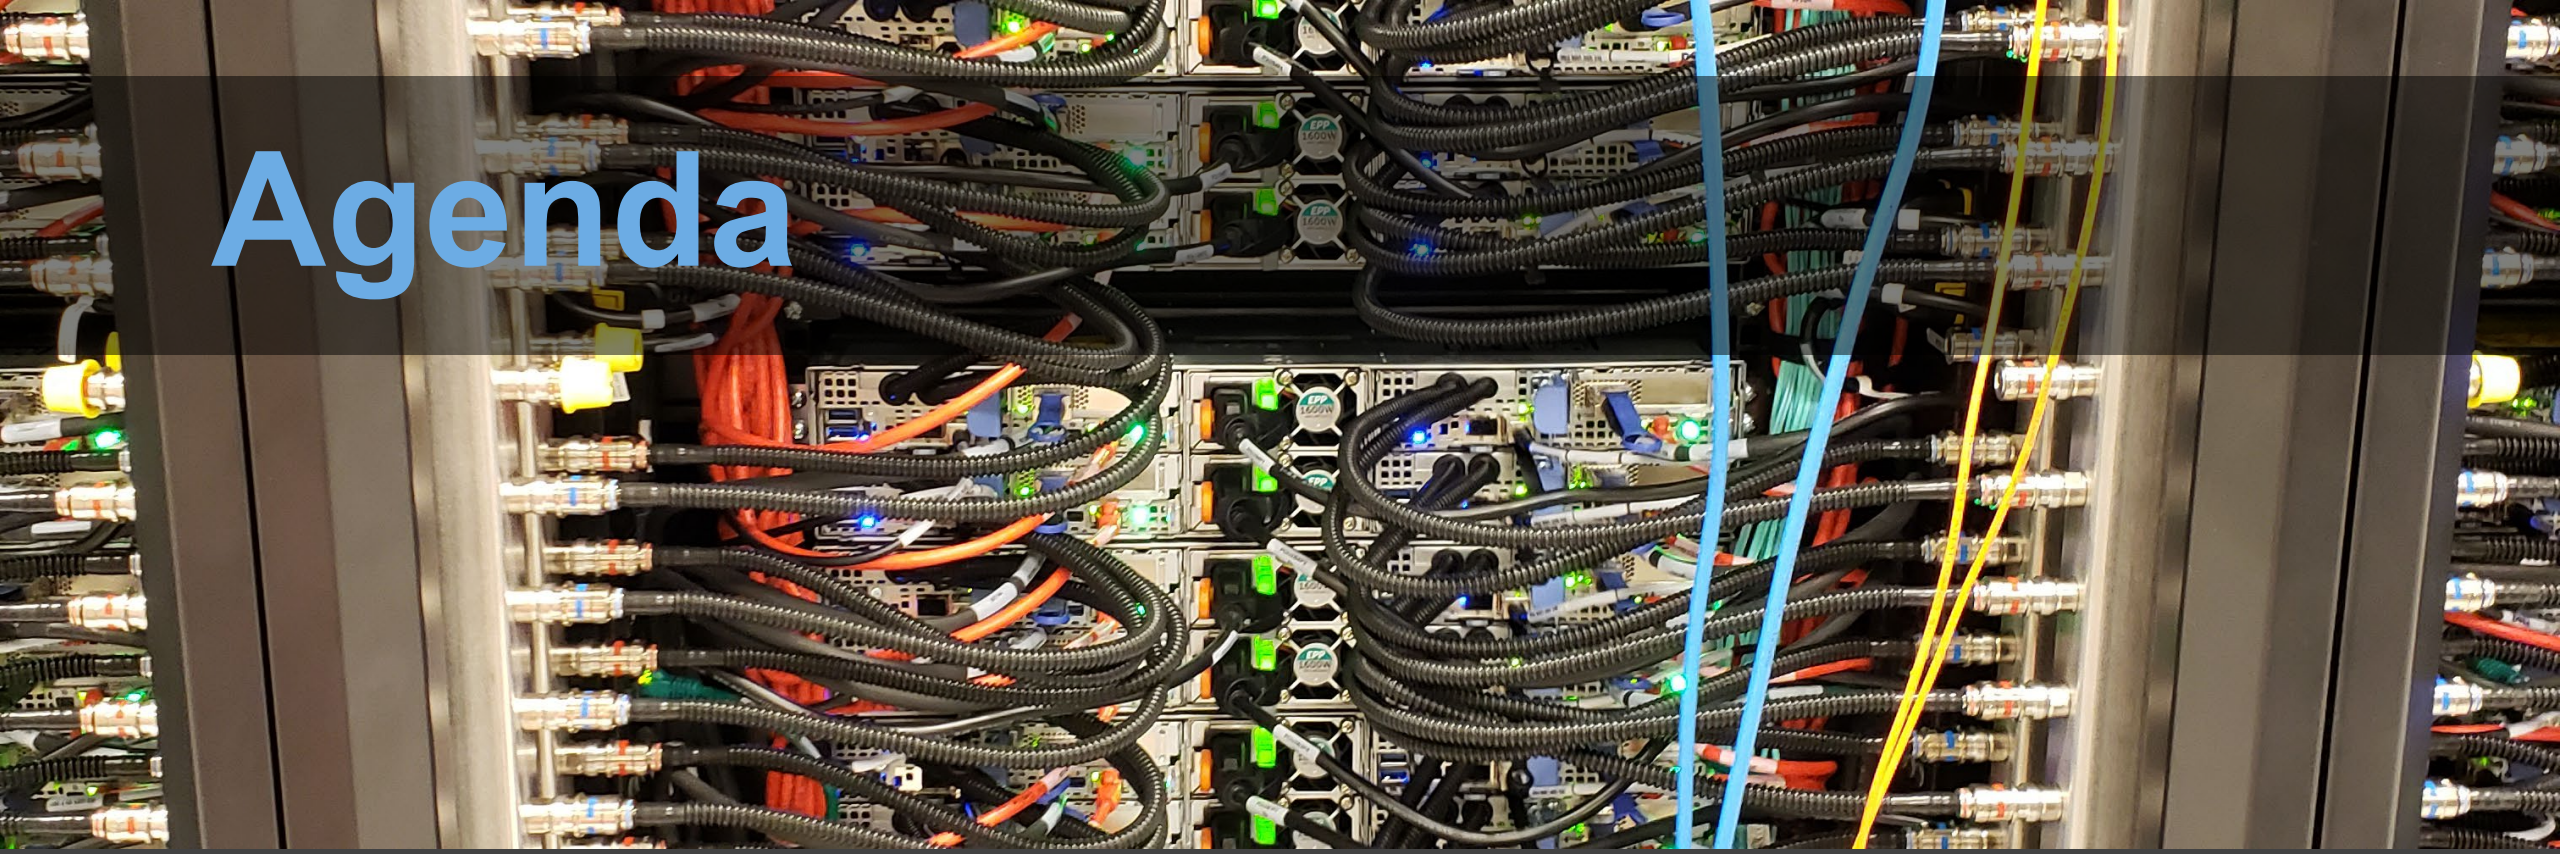

- Organizational Items
- Client Impact
- Services Overview
- Top Initiatives
- Opportunities
- Challenges

Organizational Items

David Hudak, Executive Director

Organizational Update

- FY19-20 capital budget allocation of $6.105M
  - Production infrastructure refresh
  - Protected Data Environment
  - Research Data Archive
- FY20-21 operating budget request underway
  - Asking to maintain current $4.388M/year
- Welcome new employees!
  - Morgan Rodgers (Web)
  - ZQ You (Scientific Apps)
  - Kyle Earley (Ops)
- Employee transitions
  - Janet Gregory retirement
  - Interface Lab transition to OSU

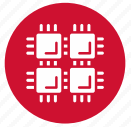

PEARC'18 Annual Conference

- "Practice & Experience in Advanced Research Computing" conference
- Attended by ~600 faculty and staff from academic supercomputer centers
- OSC had significant involvement

Type	Title
BOF	Open OnDemand – Present and Future Plans
BOF	Academic HPC center ROI calculations and cloud provider comparisons
BOF	Raising the Bar for High Quality HPC Learning Repositories
BOF	Supporting Student-Driven Research: Fostering Mentorship to Promote Student Success
Paper	Supporting parallel, interactive Jupyter, and RStudio in a scheduled HPC environment with Spark, and MPI frameworks using Open Ondemand
Paper	Teaching Data Science through Social Change
Paper	Scaling Puppet and Foreman in HPC
Poster	Code Optimization and Stabilization for a High-Resolution Terrain Generation Application
Poster	HPC Educational Programs for Middle School and High School Students
Poster	Scaling large parallel file system backups
Tutorial	Introduction to Python 3 and Jupyter Notebooks
Workshop	Workshop on Challenges to HPC Education and Training (CHET18)
Workshop	Student Modeling Challenge

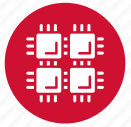

Client Impact

Brian Guilfoos, HPC Client Services Manager

Client Services

CY2017

23 academic institutions

48 companies

2,202 clients

256 awards made

23 training opportunities

461 trainees

604 projects served

33 courses used OSC

Usage by Field of Science

CY2017

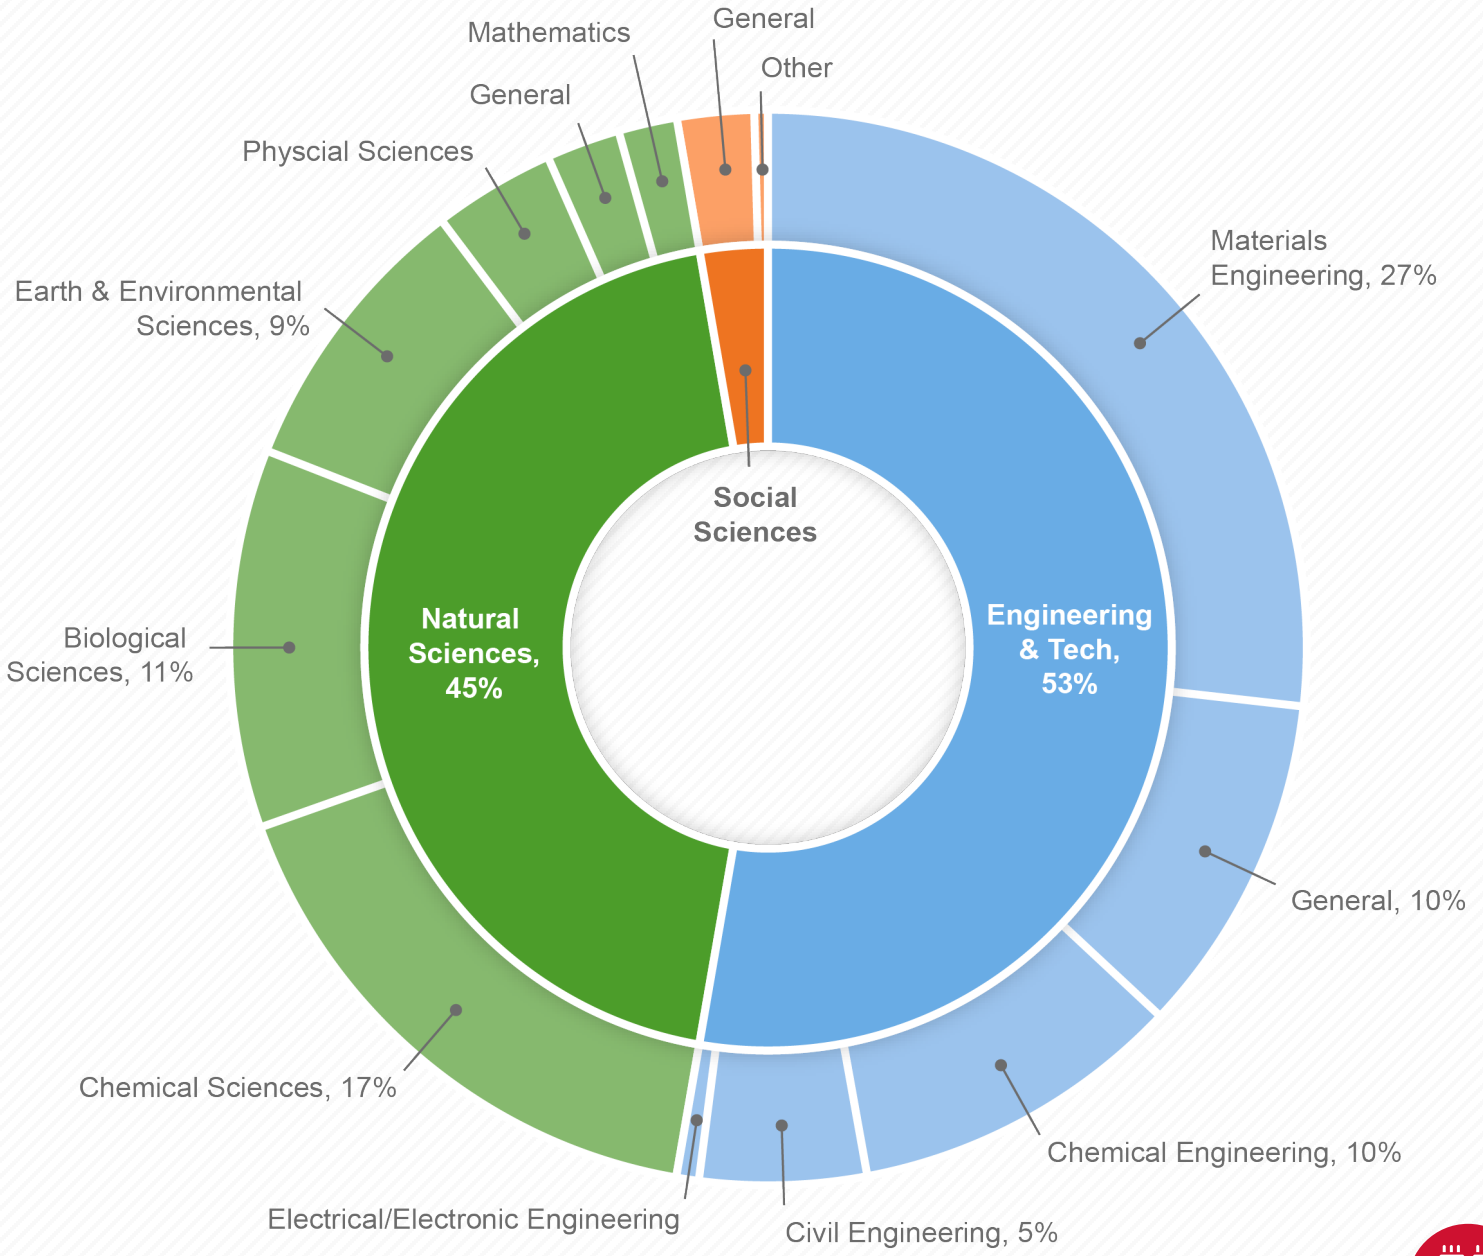

OSC Client Distribution by University (CY2017)

Institution	Active Users
Bowling Green State University	31
Bluffton University	7
Cedarville University	1
Cleveland State University	18
Case Western Reserve University	30
Kent State University	21
Kenyon College	1
Miami University	27
Oberlin College	2
Ohio Dominican University	1
Ohio Northern University	1
Ohio University	45
The Ohio State University	996
University of Akron	40
University of Cincinnati	162
University of Dayton	27
University of Toledo	22
Wittenberg University	1
Wright State University	11
Youngstown State University	9

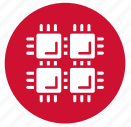

Academic Course Enrollment

CY2017

Department	
Bluffton Mathematics	9
OU Chemical & Biomolecular Engineering	10
OU Chemistry & Biochemistry	2
OSU Materials Science & Engineering	39
OSU Chemical & Biomolecular Engineering	19
OSU Materials Science & Engineering	16
OSU Materials Science & Engineering	6
OSU Computer Science & Engineering	30
OSU Computer Science & Engineering	21
OSU Evolution, Ecology, & Organismal Biology	18
OSU Computer Science & Engineering	3
OSU Evolution, Ecology, & Organismal Biology	12
OSU Biostatistics	3
OSU Chemistry & Biochemistry	28
OSU Computer Science & Engineering	34
OSU Computer Science & Engineering	75
OSU Geography	13

Department	
OSU Computer Science & Engineering	49
OSU Computer Science & Engineering	5
OSU Chemistry	400
OSU Chemistry	14
OSU Chemistry	10
OSU Chemistry	129
OSU Chemistry	105
OSU Chemistry	10
Akron Mechanical Engineering	12
UC Electrical Engineering & Computer Systems	18
UC Physics	23
UC Electrical Engineering & Computer Systems	63
UC Electrical Engineering & Computer Systems	60
UD Electrical & Computer Engineering	13
UD Electrical & Computer Engineering	14
Toledo Chemistry	16

1,279 total students

Web-based US Logins

CY2017

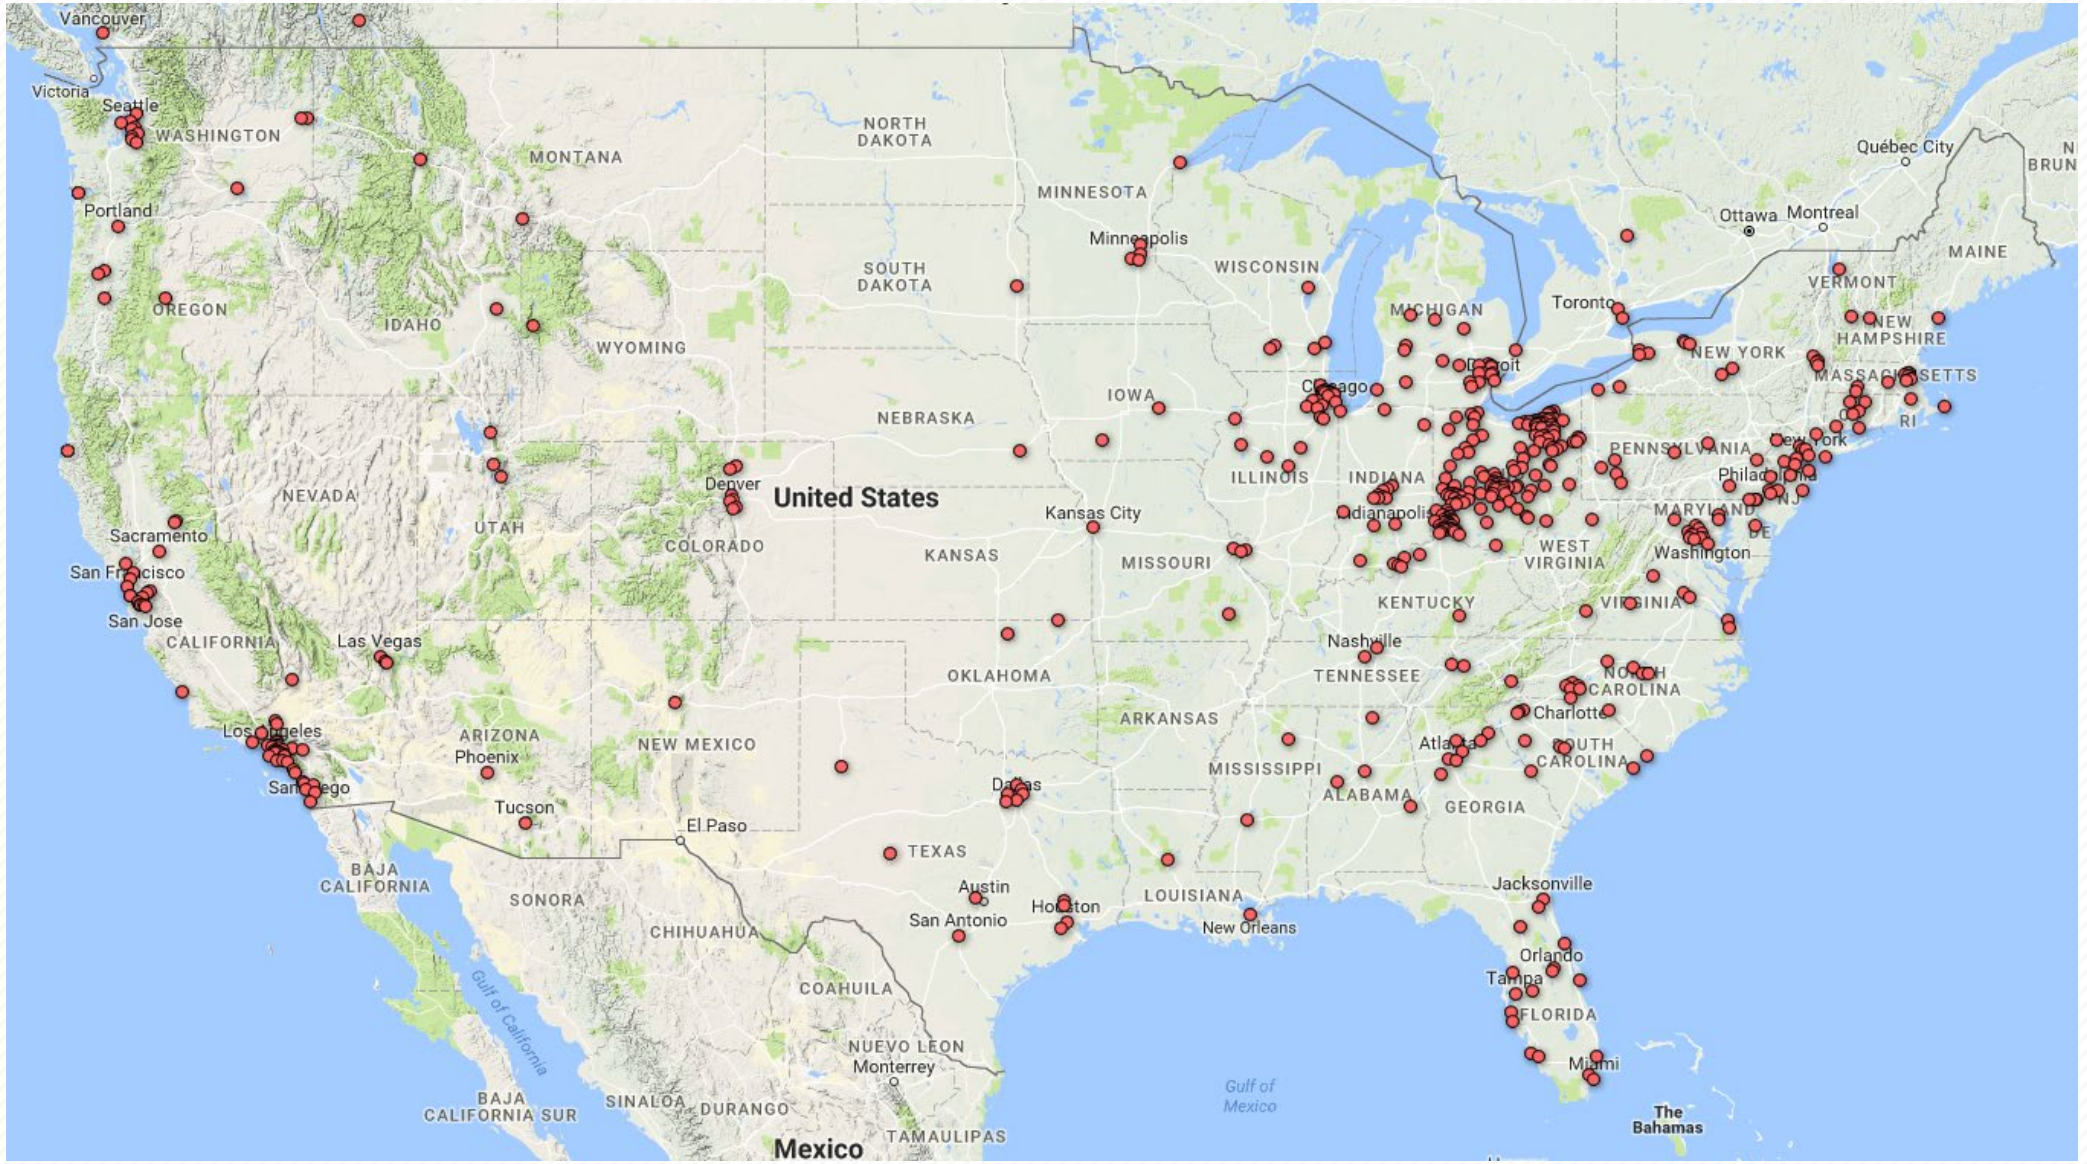

Example Client: Ian Howat

Title: “Automated, High Resolution Terrain Generation for XSEDE”

Funding Source: National Science Foundation (NSF)

Research: Establish a service for on-demand polar Digital Elevation Model (DEM) production and distribution utilizing the XSEDE High Performance Computing framework and the NSF-funded Polar Geospatial Center (PGC) data services

OSC Services: software engineering, code optimization and parallel software development

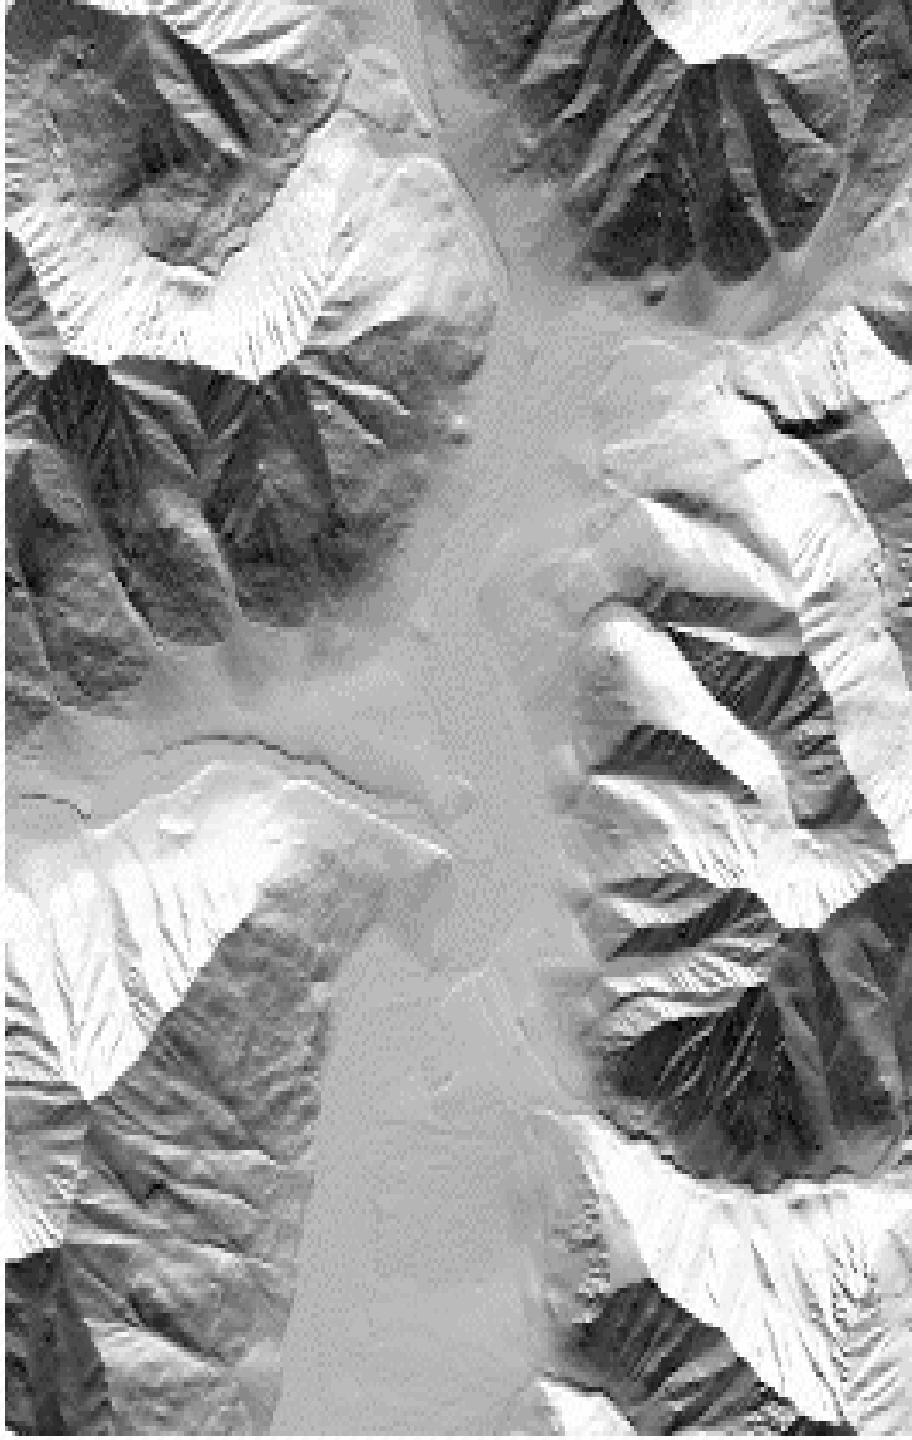

Training & Outreach Activities

Recent

- 5 classroom presentations on OSC services
- HPC Carpentry workshop at OSU/Physics
- OSC Big Data
- Several XSEDE events

Upcoming

- UC: Intro to OSC and Big Data workshops Oct 11th
- Wright State: Intro presentation, Oct 18th
- Wright State: hands-on workshop, Oct 30th
- Bowling Green: Nov 1st
- OSC: Big Data workshop, Nov 28th
- CWRU, UC, OU, and OSU all planned for the spring semester

Ongoing

- M-F 4-5PM in Pomerene Hall
- Every other Tuesday afternoon at Research Commons

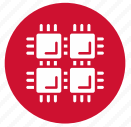

Summer Institute

si

- Two-week residential program for high school students
- 16-20 participants each year since 1989
- Students solve complex problems under mentorship of OSC faculty clients
- 2018 schools represented:
  - Scioto High School
  - Ottawa Hills High School
  - Olentangy Liberty High School
  - Olentangy High School
  - Indian Hill High School
  - New Albany High School
  - Upper Arlington High School
  - Columbus Academy
  - Phillips Exeter Academy
  - Dublin Jerome High School
  - Dublin Coffman High School
  - St. Xavier High School
  - Westlake High School
  - Hawken Upper School

Young Women's Summer Institute

- One-week residential program for middle school girls
- 16 participants each year since 2000
- Promotes STEM skills under mentorship of middle school teachers
- 2018 schools represented:
  - Immaculate Heart of Mary Cincinnati
  - Shanahan Middle School
  - Batavia Middle School
  - Sycamore Junior High School
  - Hilliard Station 6th Grade
  - Holy Angels
  - Sycamore Junior High School
  - Hilliard Tharp
  - Woodbury Elementary
  - Fredericktown Middle School
  - Columbus Academy
  - Karrer Middle School
  - Walnut Springs Middle School
  - Bexley Middle School
  - Dublin Coffman High School
  - Willard Grizzell Middle School

Unique Users

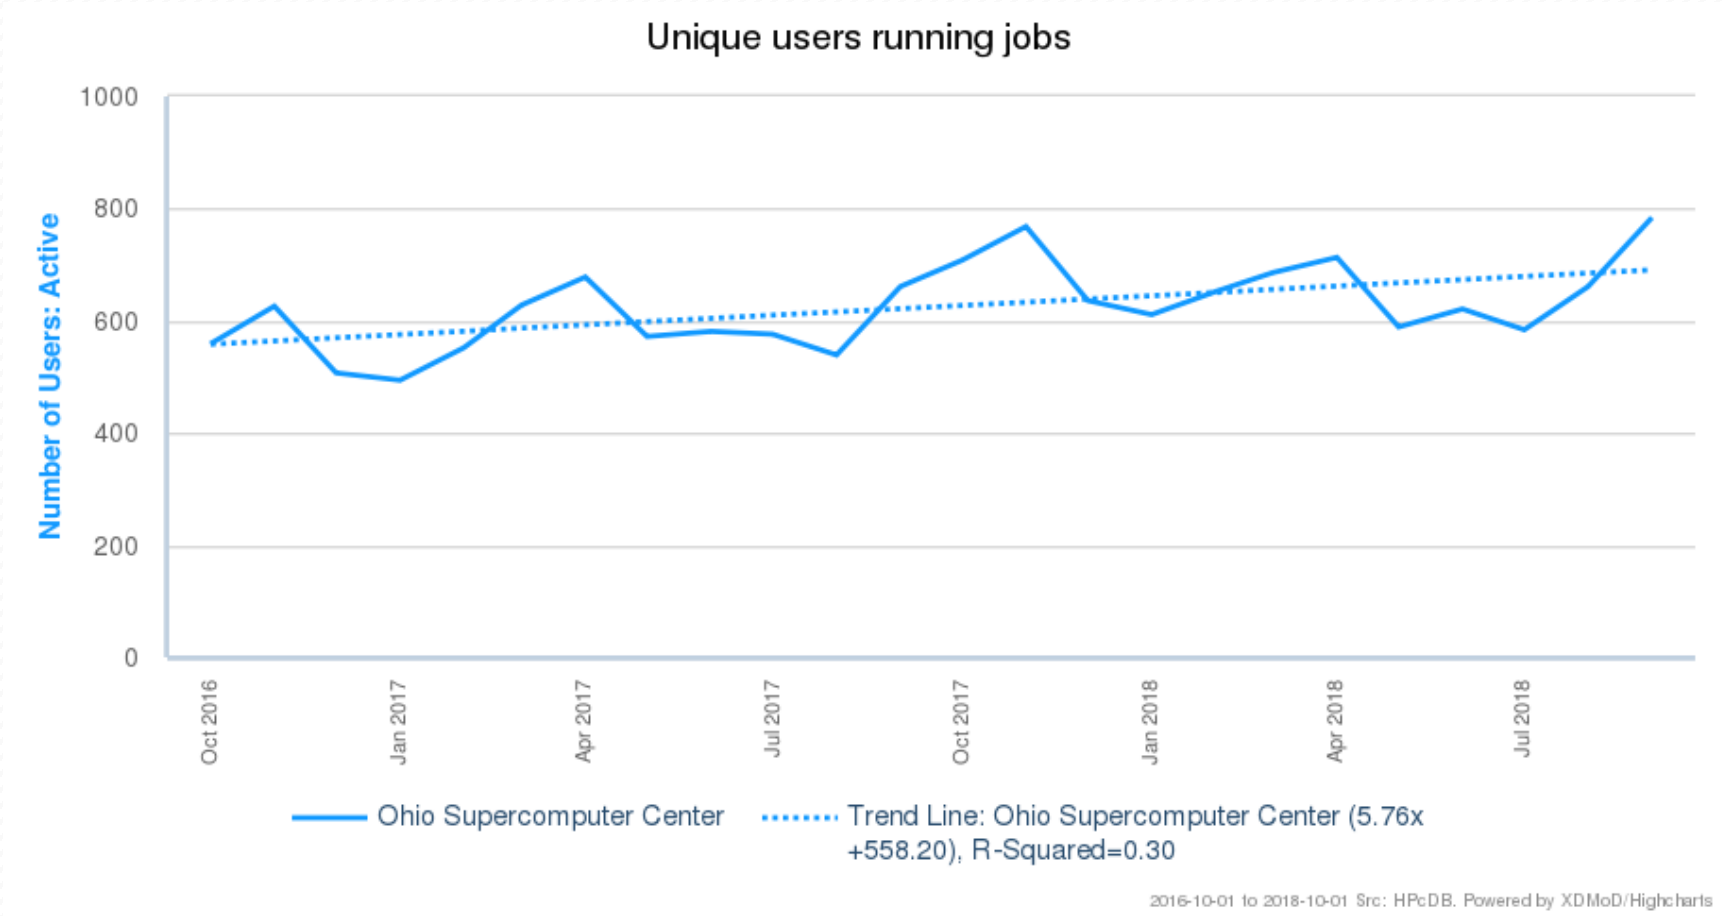

Active Client Service Support Load

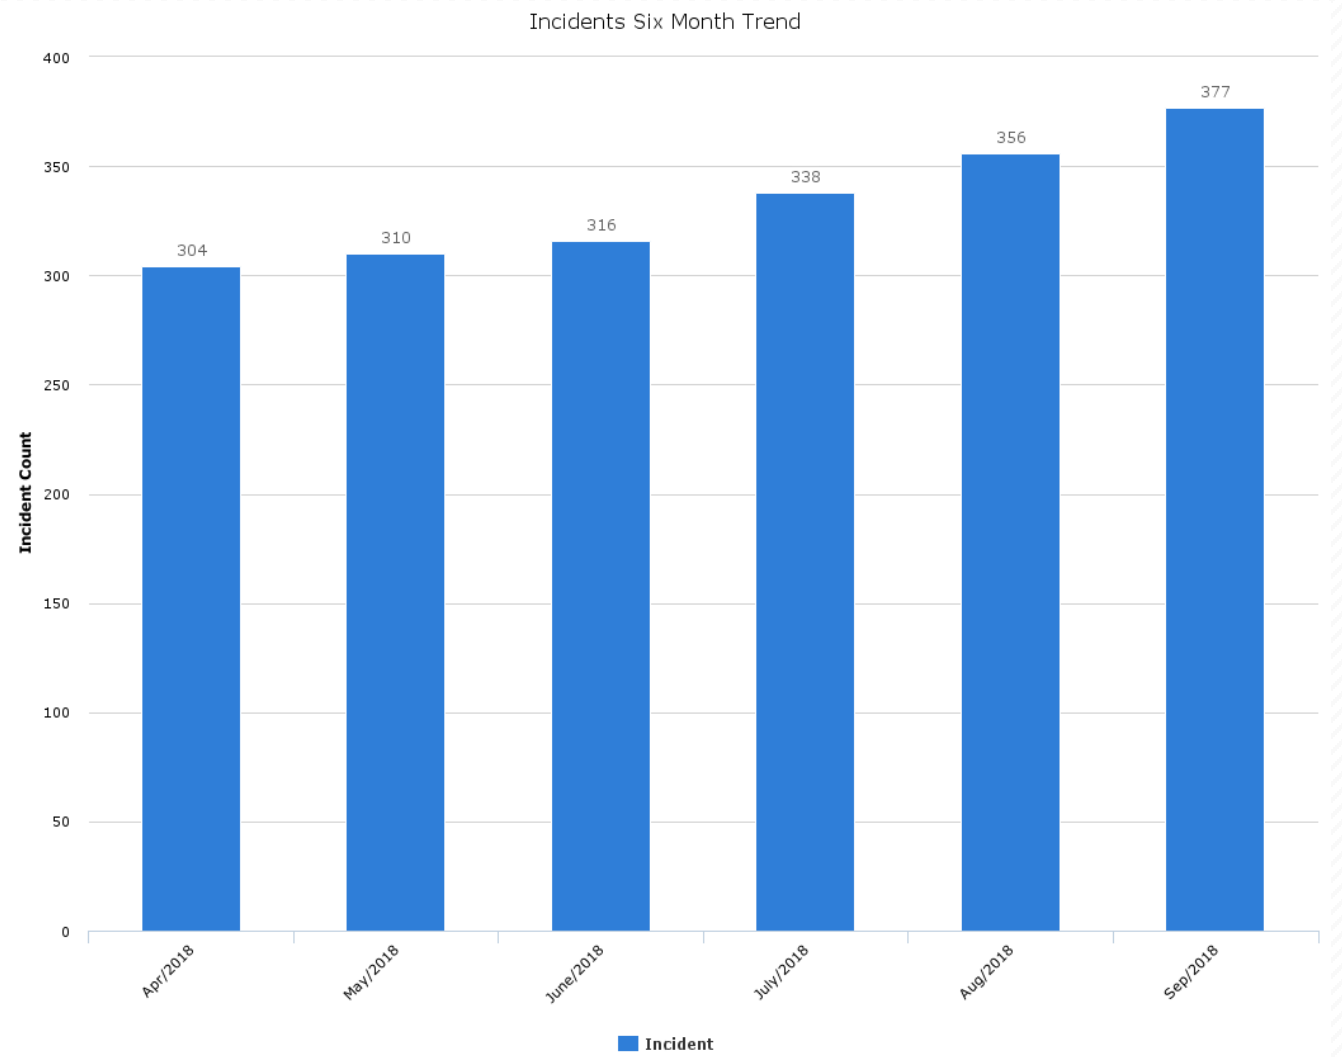

~300 tickets / month

~15 tickets / work day

Services Overview

Brian Guilfoos, HPC Client Services Manager

Service Catalog

Cluster Computing

A fully scalable center with mid-range machines to match those found at National Science Foundation centers and other national labs.

Research Data Storage

High-performance, large capacity data storage spaces along with others that are perfect for a wide variety of research data.

Education

High performance computing and networking resources come together to create an exciting and innovative teaching and research environment.

Web Software Development

Our expert web development team helps you create custom web interfaces to simplify the use of powerful HPC resources.

Scientific Software Development

Deep expertise in developing and deploying software that runs efficiently and correctly on large scale cluster computing platforms.

Production Capacity

CY2017

221,400,000+ core-hours consumed

78% average HPC system utilization

4,400,000+ computational jobs

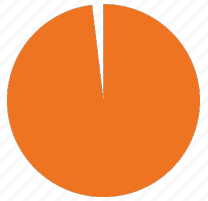

98% up-time

44% average storage system utilization

1.5 PB data stored

2 PB data transferred

79% jobs started within one hour

OSC Systems Core Hours

CY2017

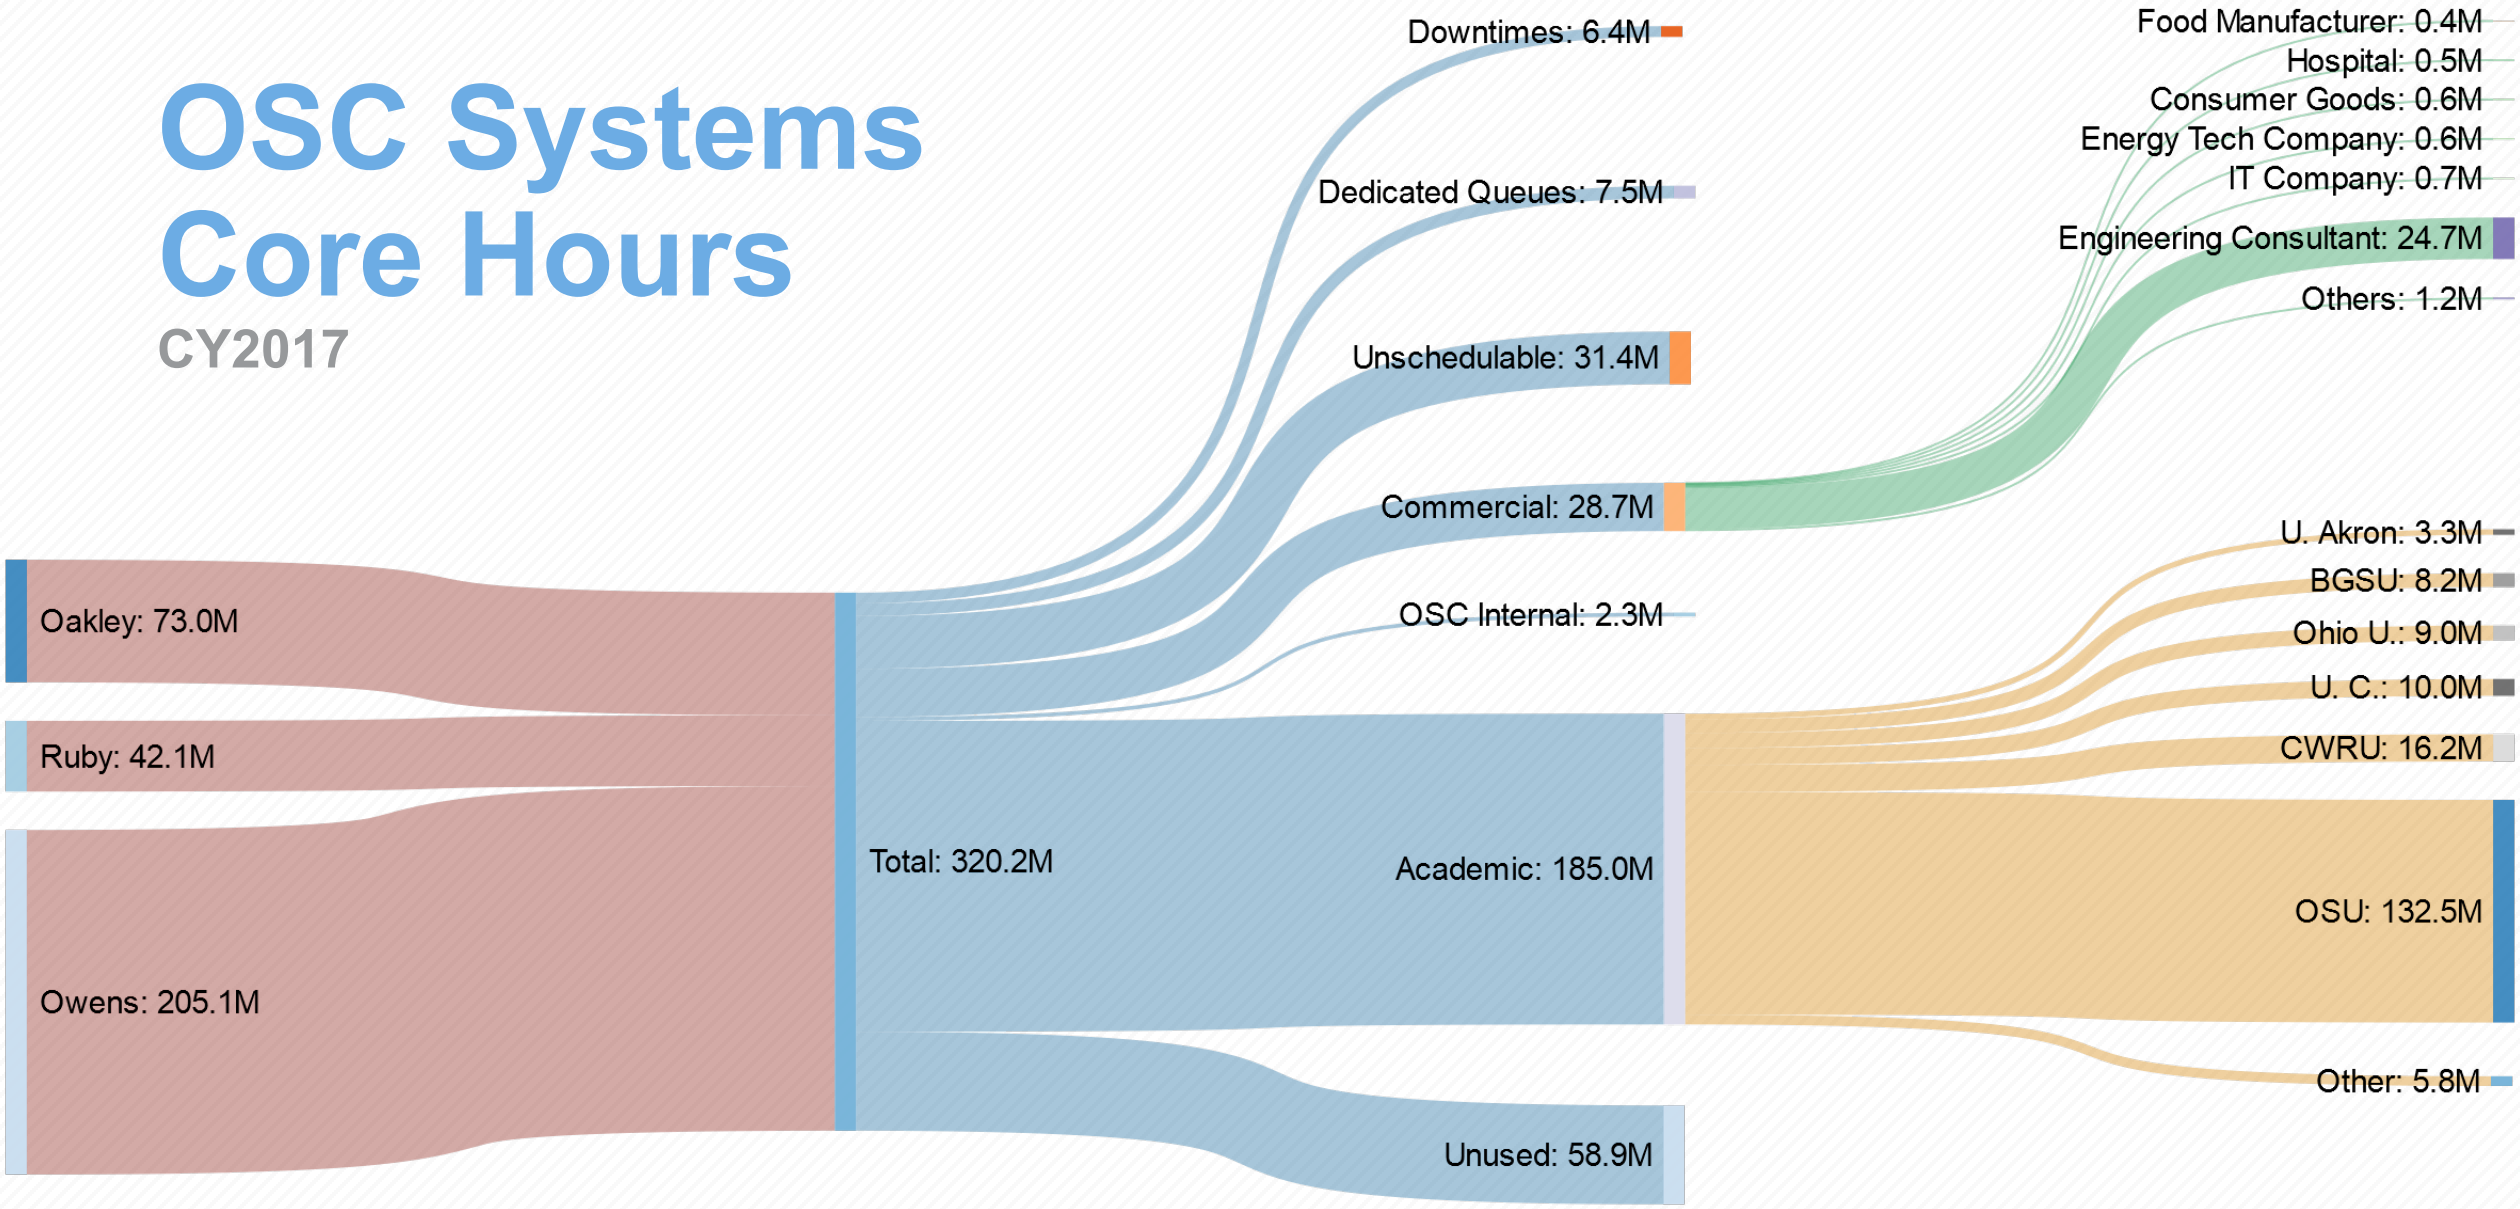

Project & Scratch Storage Utilization

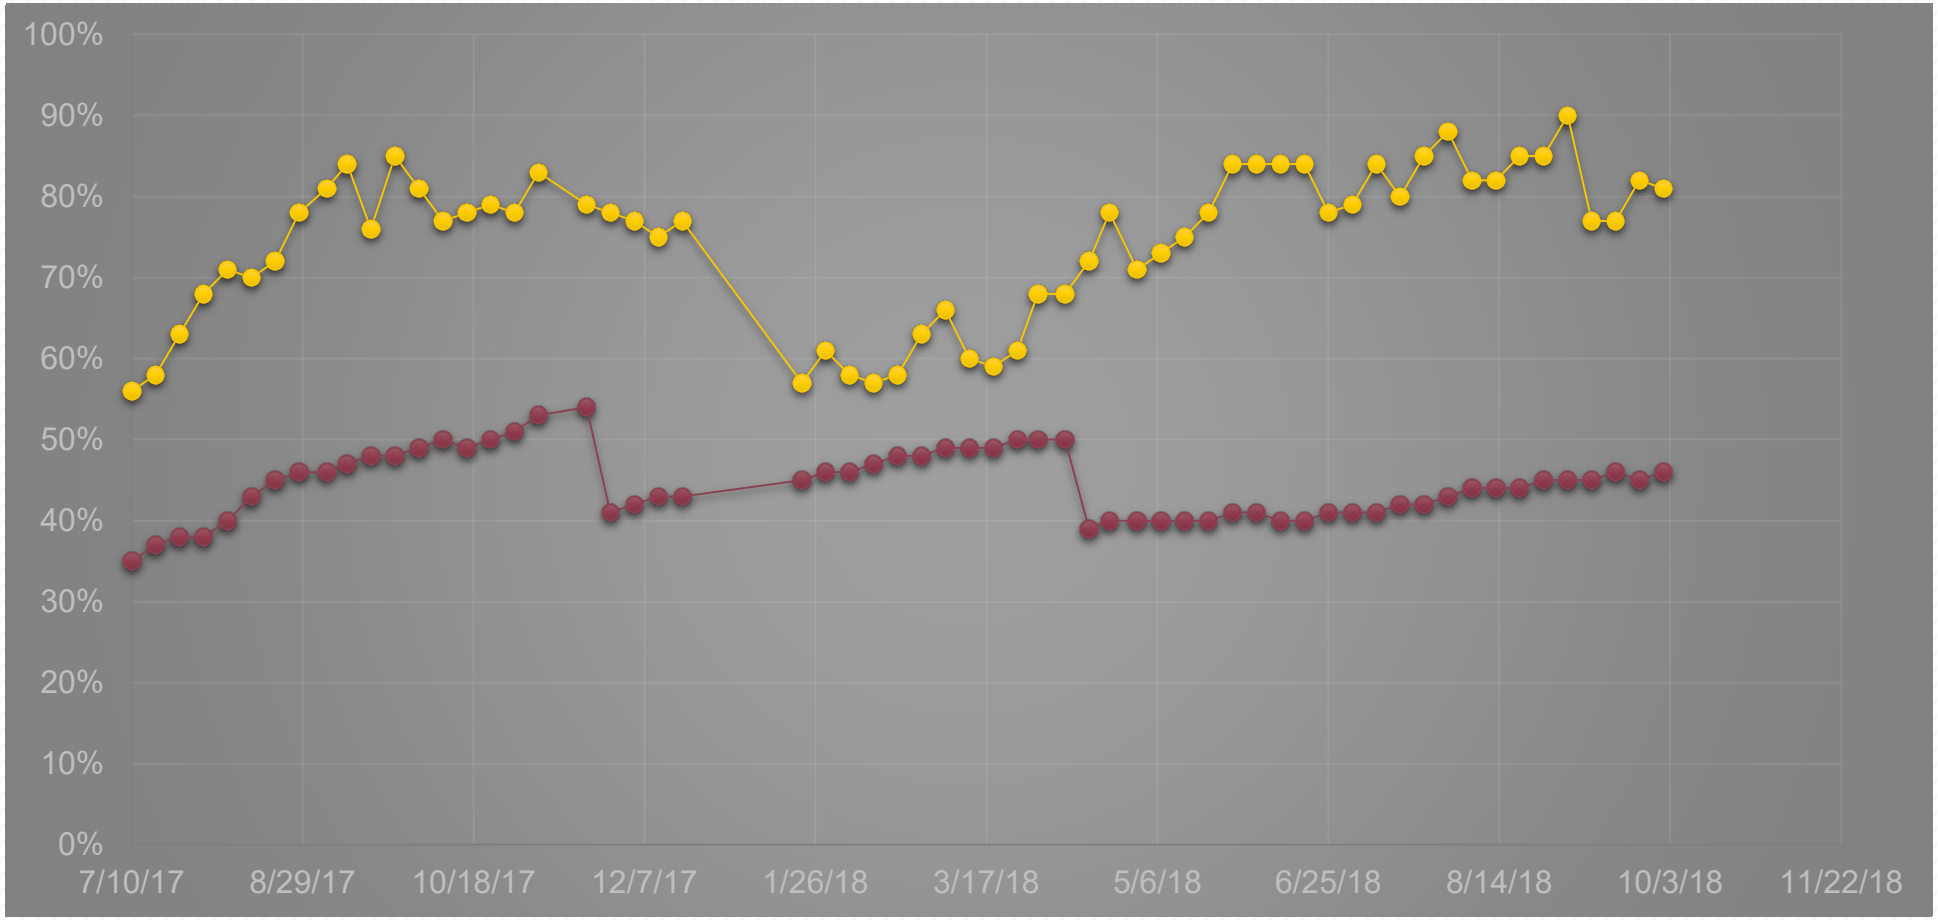

OnDemand Clients Jan'17-Aug'18

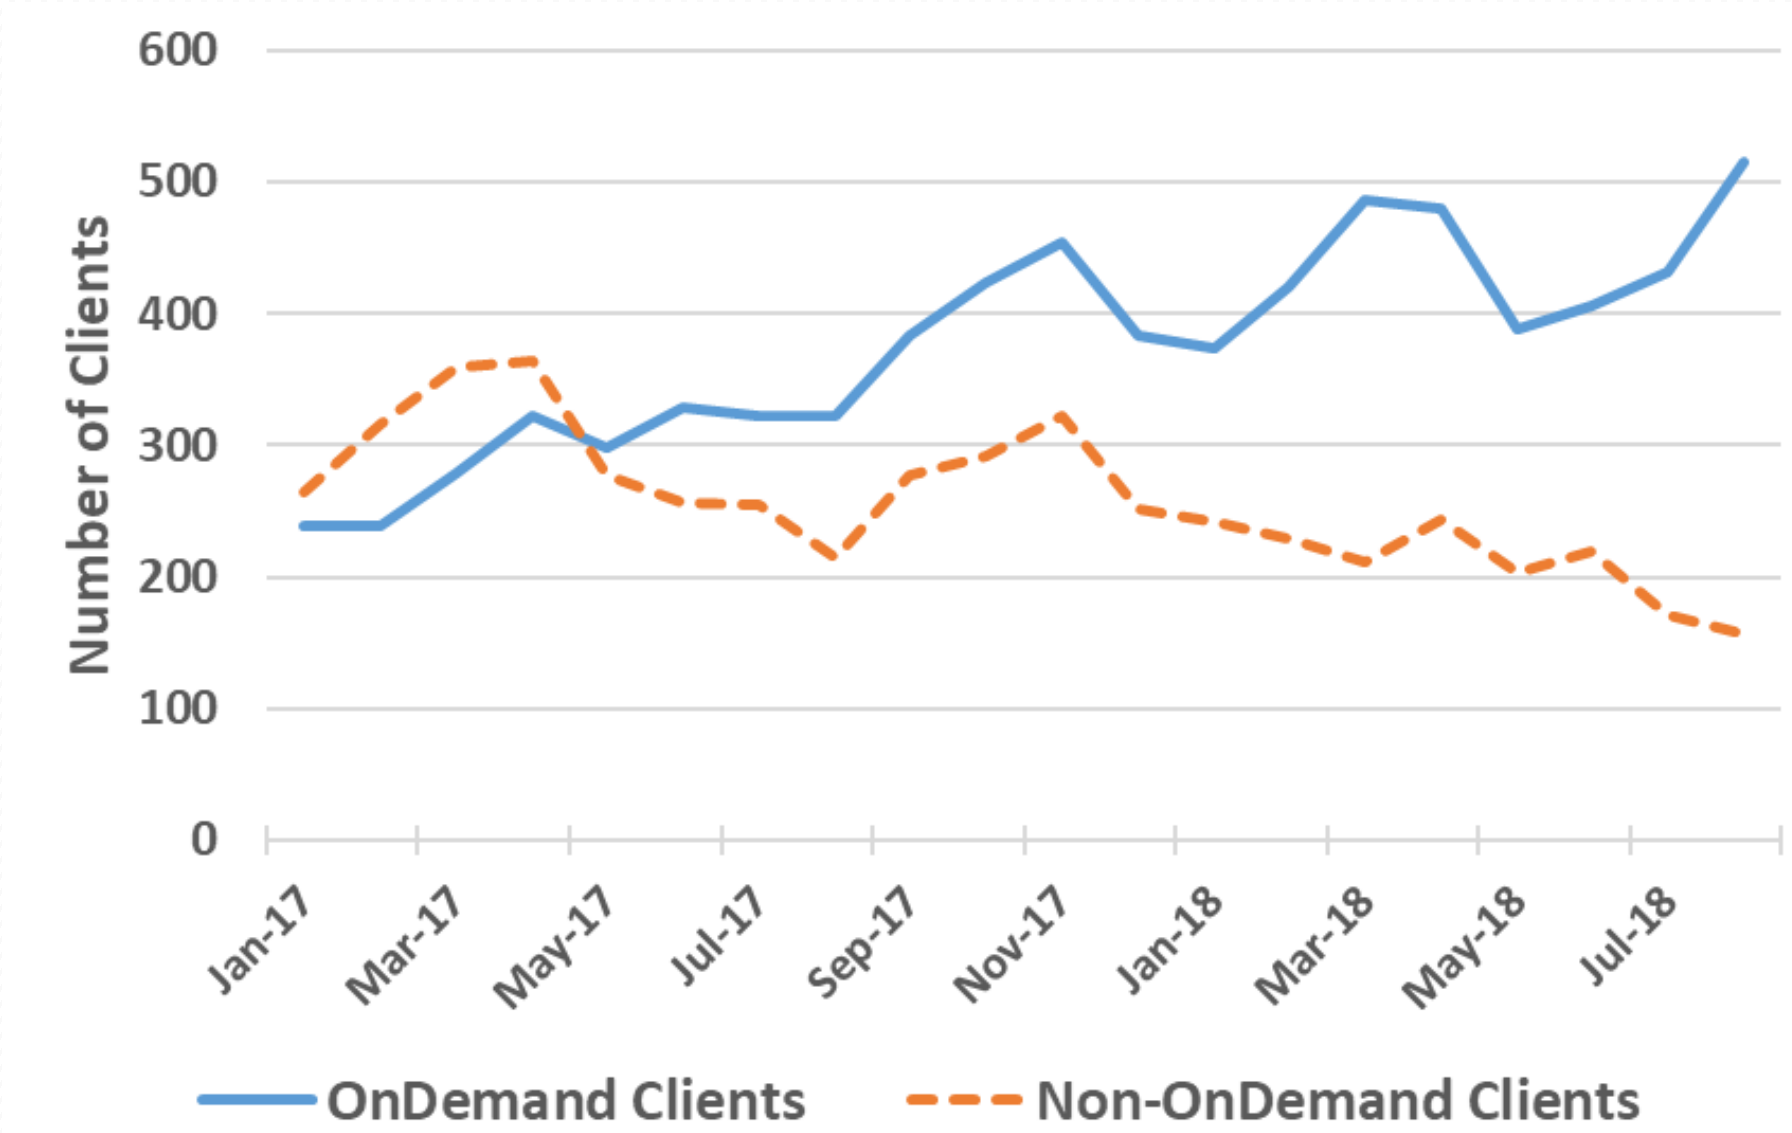

Open OnDemand Project

- New OSC OnDemand users start faster than ssh users: first login & job

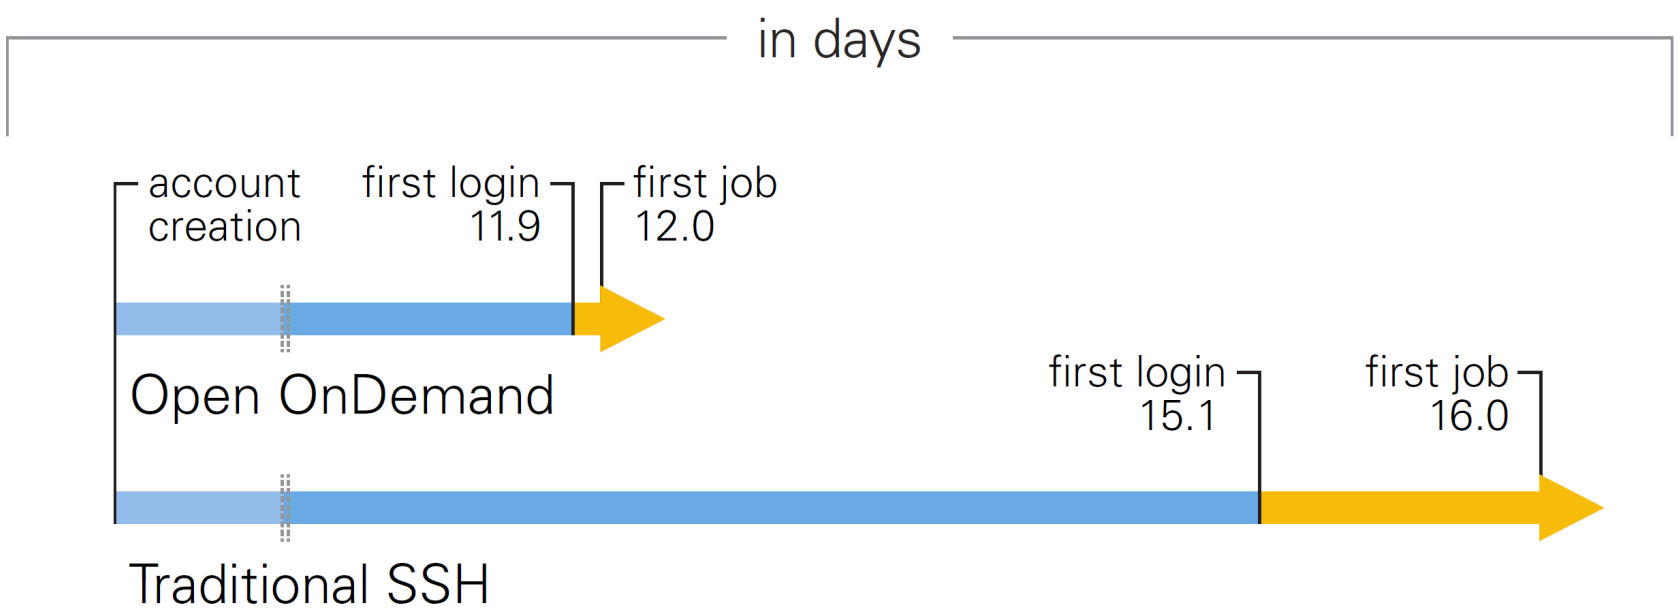

- Open OnDemand currently in use / evaluation at 30+ organizations
- Open OnDemand 2.0: NSF CSSI award, Jan'19 – Dec'22, $3.5M

Custom Portal Development

Recent Examples:

- OSC staff developed new OSU Data Commons service: datacommons.tdai.osu.edu
- R Shiny App server for OSU BMI department
- GoFly project CFD portal for TotalSim

- Contact us to discuss custom portals for your research!

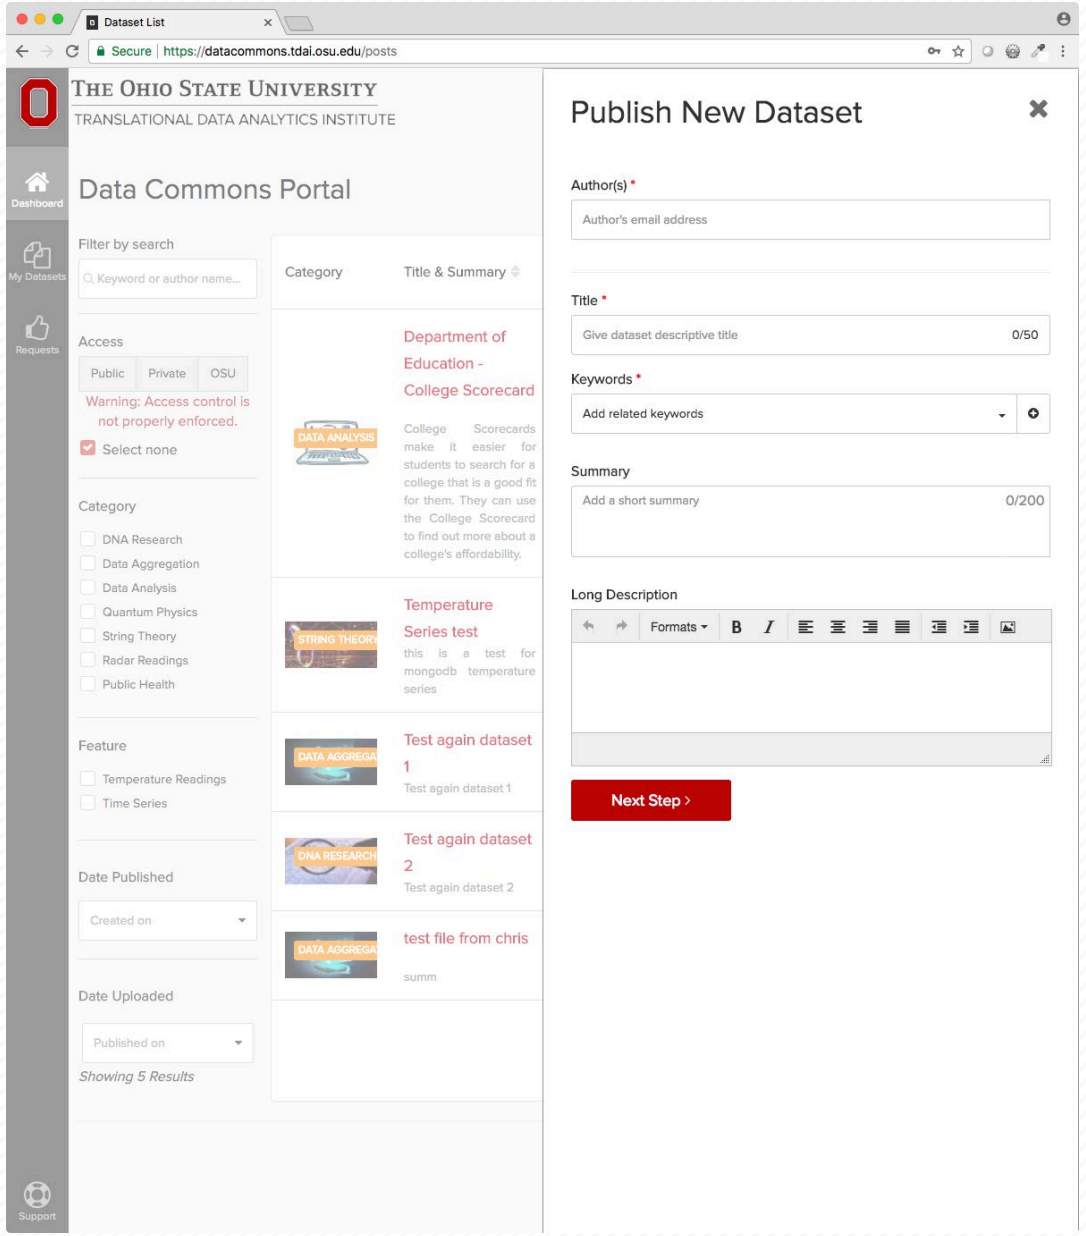

Top Initiatives

Doug Johnson, Chief Architect

Capital HPC Systems Projects

New HPC cluster “Pitzer”

- **Goals**
  - Complement existing systems
  - Replace Oakley with a petaflop class system

- **Timeline**
  - System delivery August 15, 2018
  - Full production November 2018
  - Oakley decommissioning Dec 2018

Pitzer Cluster

Characteristics relative to Oakley

- Delivers 8x the processing power (1,300 vs 154 TF)
- Costs 15% less ($4M vs $3.35M)
- Provides 25% more cores (10,560 vs 8,304)
- Has 2X the memory (70.6Tb vs 33.4TB)
- Uses 20% less power

Highlights

- 10,560 processor cores, ~1.3 petaflop peak
- Latest generation: SkyLake processors, 100Gb InfiniBand
- Warm water cooling supports high density, increased performance and efficiency

3 Types of Resources and Example Workload

- **Standard compute** (224 nodes) / Modeling and simulation for industry
- **GPU** (32 nodes) / Machine learning, artificial intelligence (AI)
- **Large memory** (4 nodes) / Genomics Sequence Assembly

System Status (2019)

SYSTEMS	Ruby	Owens	Pitzer
Date	2014	2016	2018
Cost	$1.5 million	$7 million	$3.35 million
Theoretical Perf.	~144 TF	~1600 TF	~1300 TF
Nodes	240	824	260
CPU Cores	4800	23392	10560
RAM	~15.3 TB	~120 TB	~ 70.6 TB
GPUs	20 NVIDIA Tesla K40	160 NVIDIA Pascal P100	64 NVIDIA Volta V100
	Total compute: ~3,044 TF		

STORAGE	Home	Project	Scratch	Tape Library
Capacity	0.8 PB	3.4 PB	1.1 PB	7+ PB
Current utilization Feb, 18	47%	48%	59%	47%

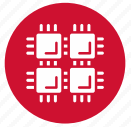

Capital Client Portal Projects

Replacement for my.osc.edu

- Friendly user testing starting August 17
- Scheduled go-live for all clients on October 23

New reports.osc.edu

- Provides OSC staff with robust client usage and billing reporting capabilities

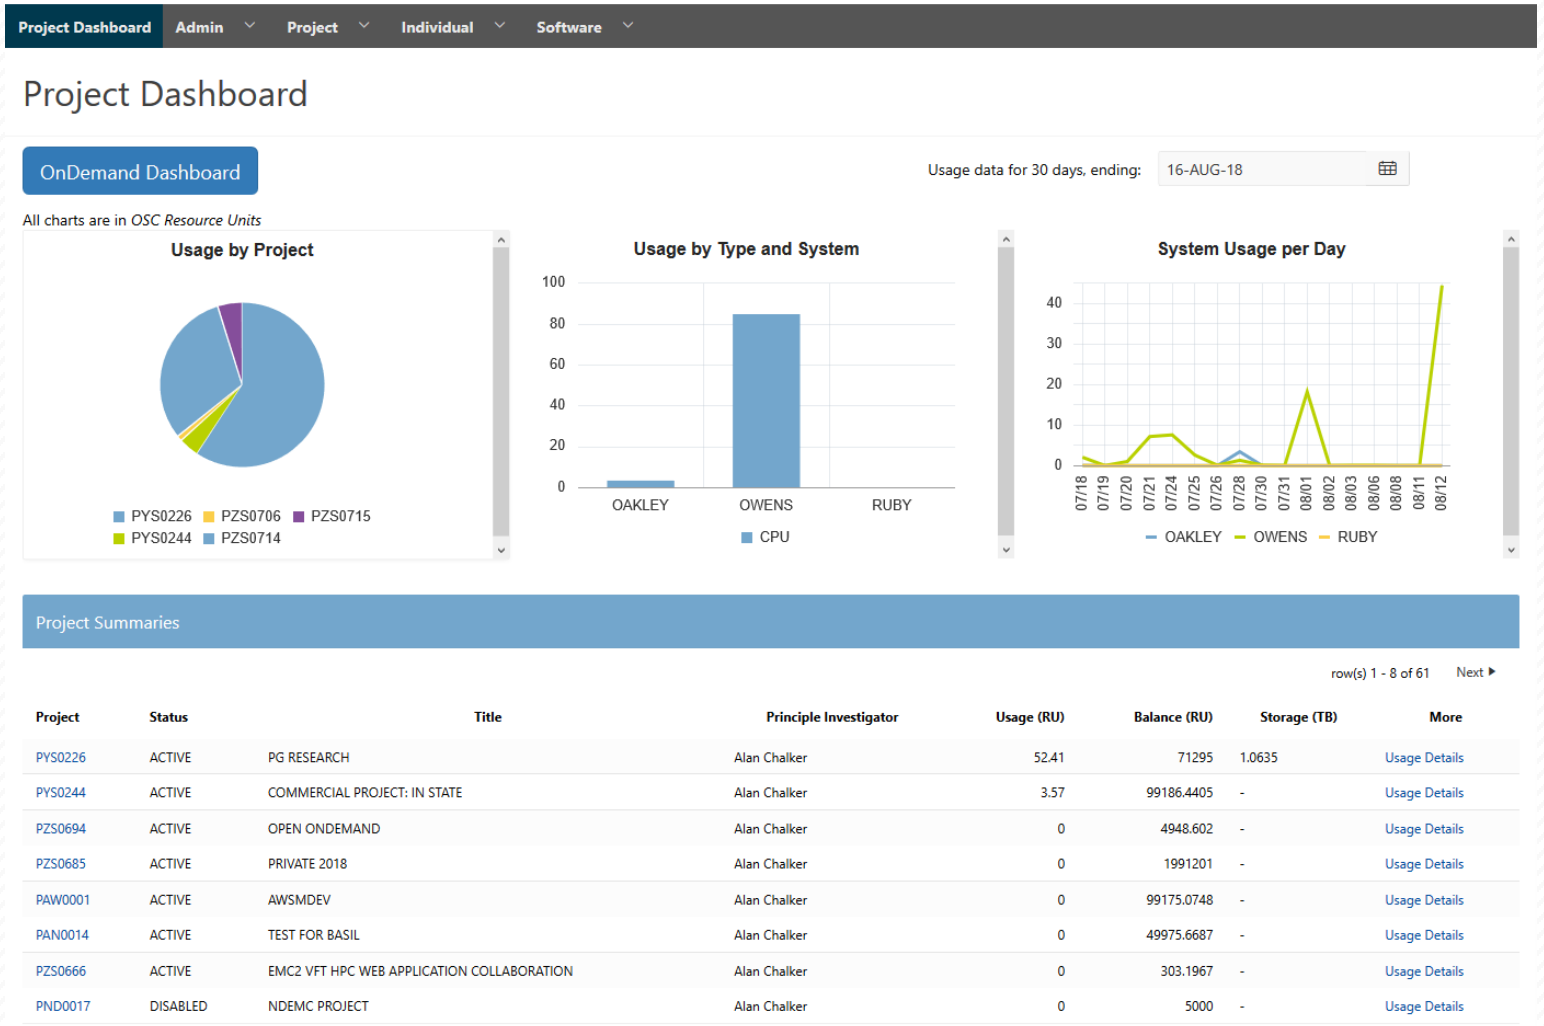

Capital Storage Projects

Upgrade tape library for backups capacity/performance, and future data archive project

- New library installed in December, 2017
- Data migration complete
- Scale Out Backup And Restore (SOBAR) implementation
- Backup servers, and disk storage pools upgrade 2nd half 2018 – Q1/2019

Research Data Archive

- Simplify data management plans for sponsored research
- Provide publishing and other abstraction capabilities
- Additional off-site copy of data for resiliency/availability

Project file system expansion

- Increase space for metadata, 2-3B files/directories (1B today)
- Slower tier of storage for infrequently accessed files

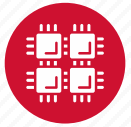

Capital Security Projects

Network firewall, and Ethernet network expansion for Pitzer

- Network expansion installed, final changes complete during October 23rd downtime
- Firewall hardware physically installed, deploy Q4/18 – Q1/19

Protected Data Environment

- Unique resource supporting HIPAA, ITAR, or other sensitive data sets
- Initial requirements gathering (OSU Wexner Medical Center, Nationwide Children's Hospital)
  - Provide standard OSC environment with security assurances, assistance with data management plans
  - Collect requirements for different services for new communities

Security Audits

- Annual OSU Information Security Standard (ISS) Audit completed September 28th
- New HPC security engineer started September 24th
- HIPAA risk assessment by external auditors in October/November, 2018

Other Capital Projects

Production infrastructure refresh

- “C20” HPC cluster (replaces the Ruby cluster)
- Storage upgrades
  - Additional fast tier storage
  - Other performance/capacity upgrades
  - Tape media capacity expansion
- Hardware to support protected data environment and research data archive
  - May include “Cloud” resources
- Other EOL hardware

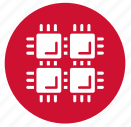

Opportunities

Alan Chalker, Director of Strategic Programs

New Research Fields
= New Client Needs

- Data Science
- Artificial Intelligence
- Machine Learning

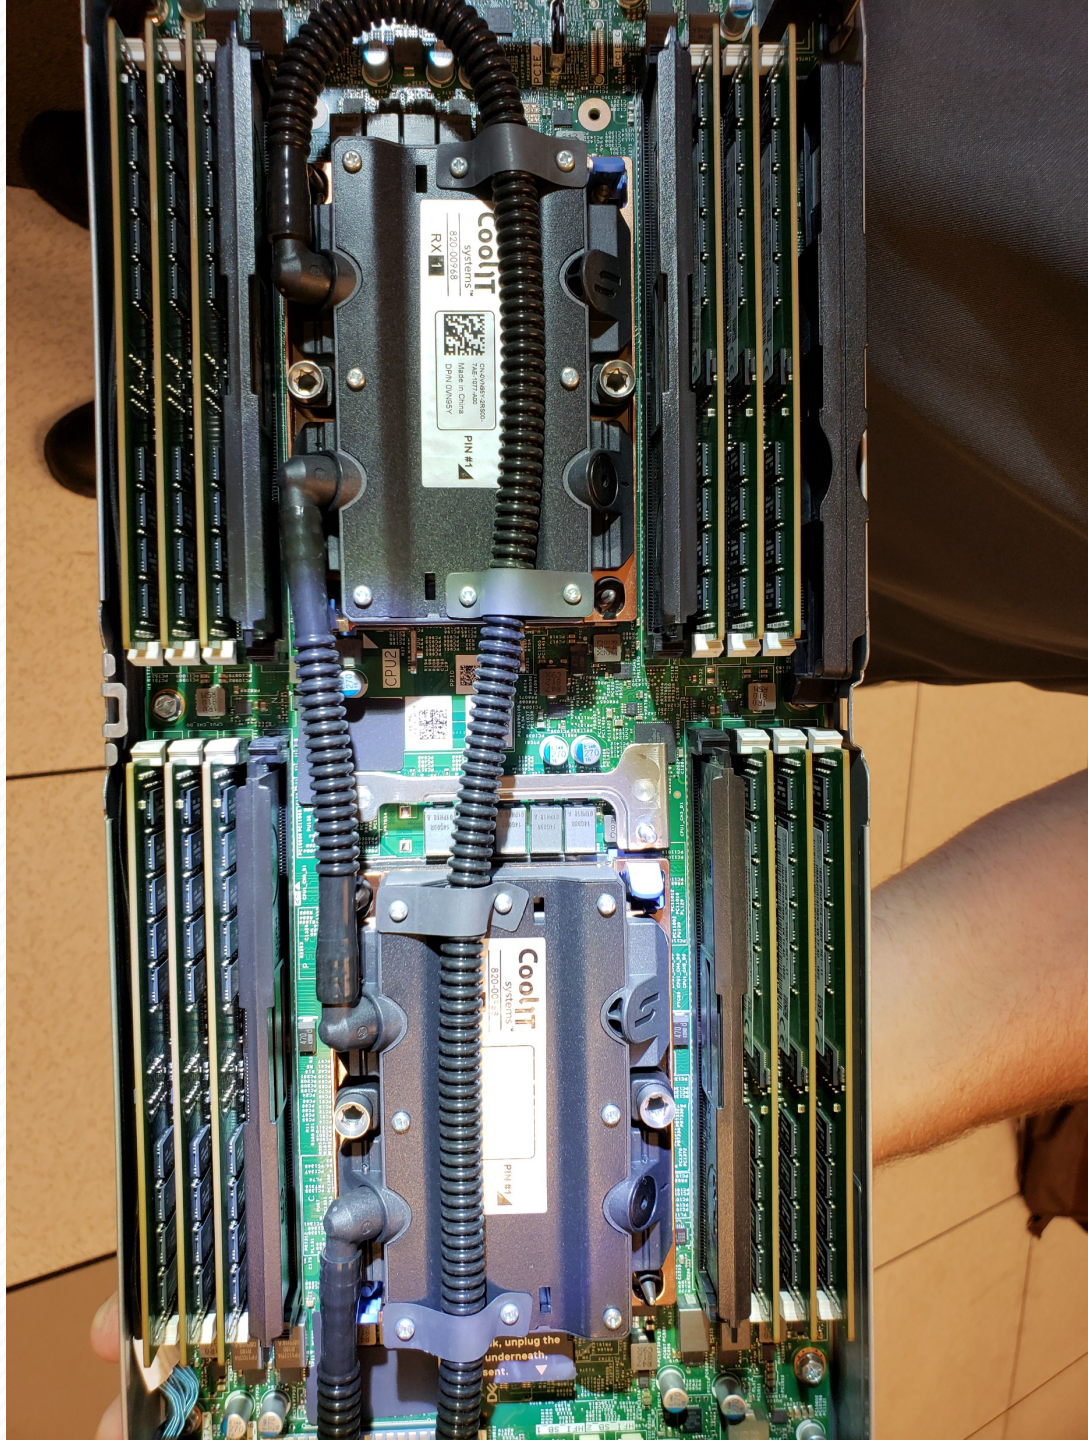

Integration with Cloud Computing Services

- Amazon AWS
- Microsoft Azure

Partnerships with Other HPC Centers

- Open OnDemand
- AFRL MSRC

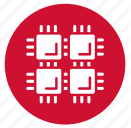

Challenges

David Hudak, Executive Director

Fee Structure Transition

FY19 Plan details

- FY19 rate of $0.075 / RU for cycles > 10K RUs per project; no storage charges
- Projected to provide the $1M needed to cover the OSC budget gap
- FY19 MOUs signed by the 6 biggest universities (OSU, BGSU, OU, UA, CWRU, UC: signing in process)
- OSC will be reaching out to other schools (e.g. Cleveland State) that have low expenses projected ($5K or more for the year).

Implementation

- Initial OSC Pricing Committee meeting March 12th, recent meeting on June 1st
- Implementation issues are being worked out, including an updated allocations process, regular institutional usage reporting, and refund policies
- Discussions are beginning with the OSC pricing committee regarding a revised model for FY20 that will likely include other charges / services

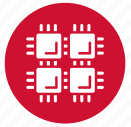

Academic Community Engagement

- Deepen outreach to Ohio universities
- Empowers local support staff to work directly with clients
- Projects created at UC, OSU, Miami, and CWRU

Organizational Development

Need to formalize and mature project management lifecycle