OSC Fall Update

- Budget Trends
- SOCC Charges
- Business Model Proposal
- Discussion

Budget Trends – Revenue Analysis

OSC Revenue Categories:

- State Subsidy
- Sponsored Research Funding
- User Fees
  - Commercial Usage
  - Condo (Dedicated) Usage

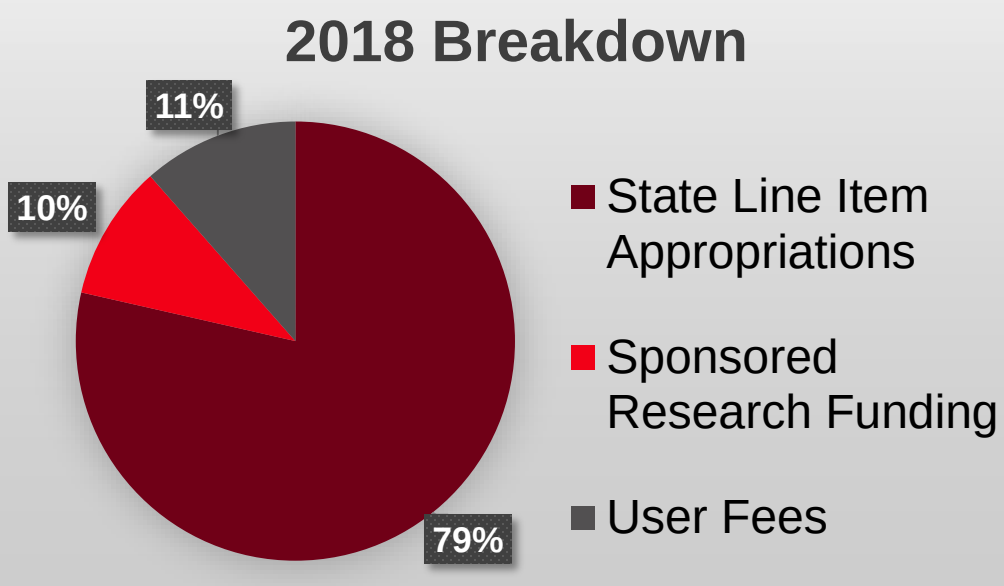

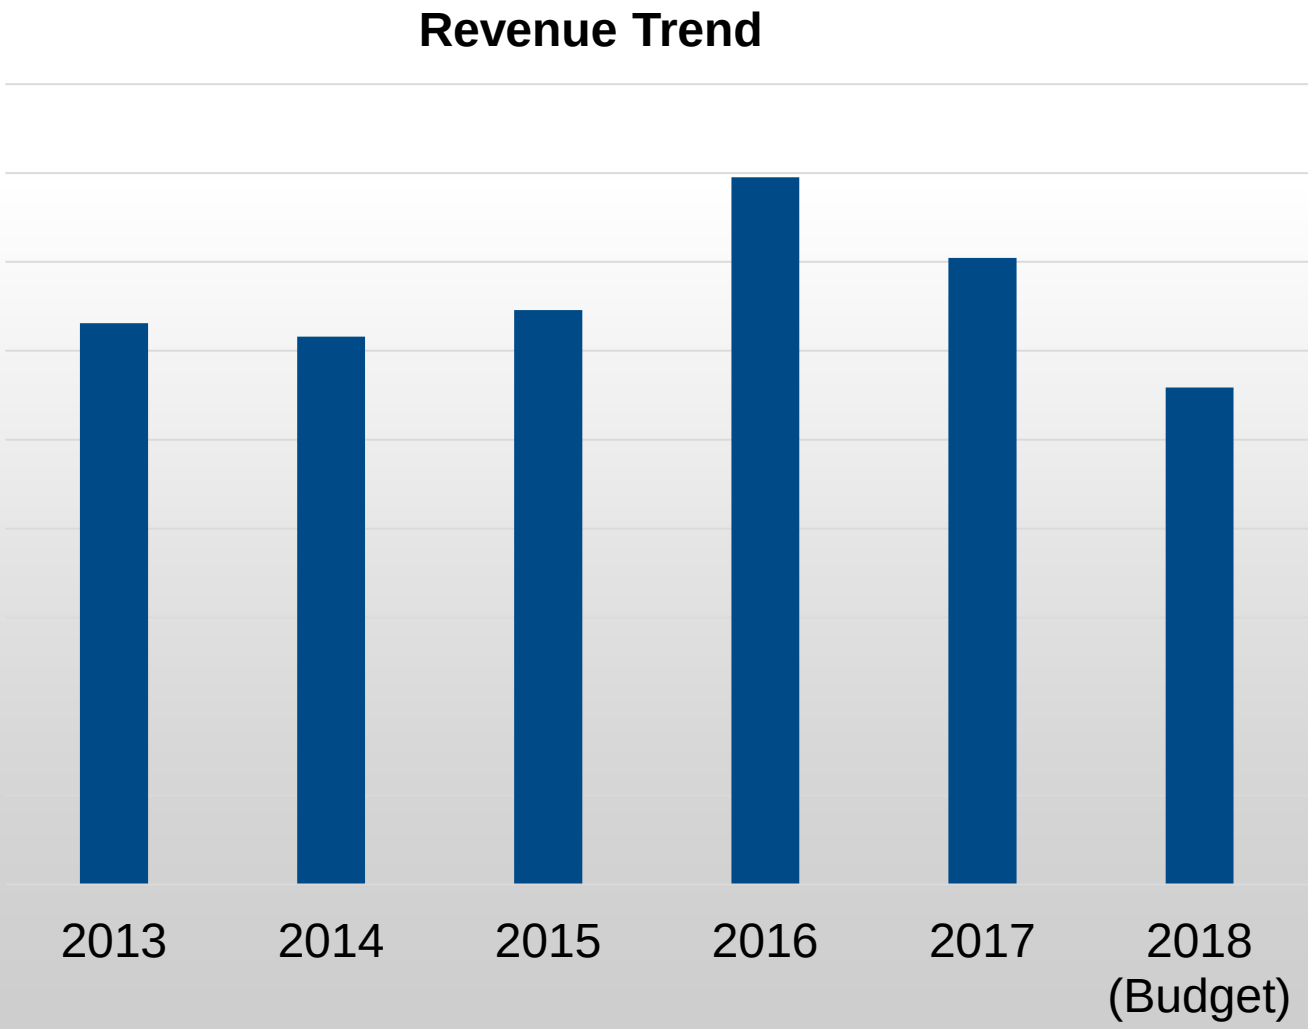

Budget Trends – Expense Analysis

OSC Expense Categories:

- Direct Costs
- Indirect Costs
- Personnel

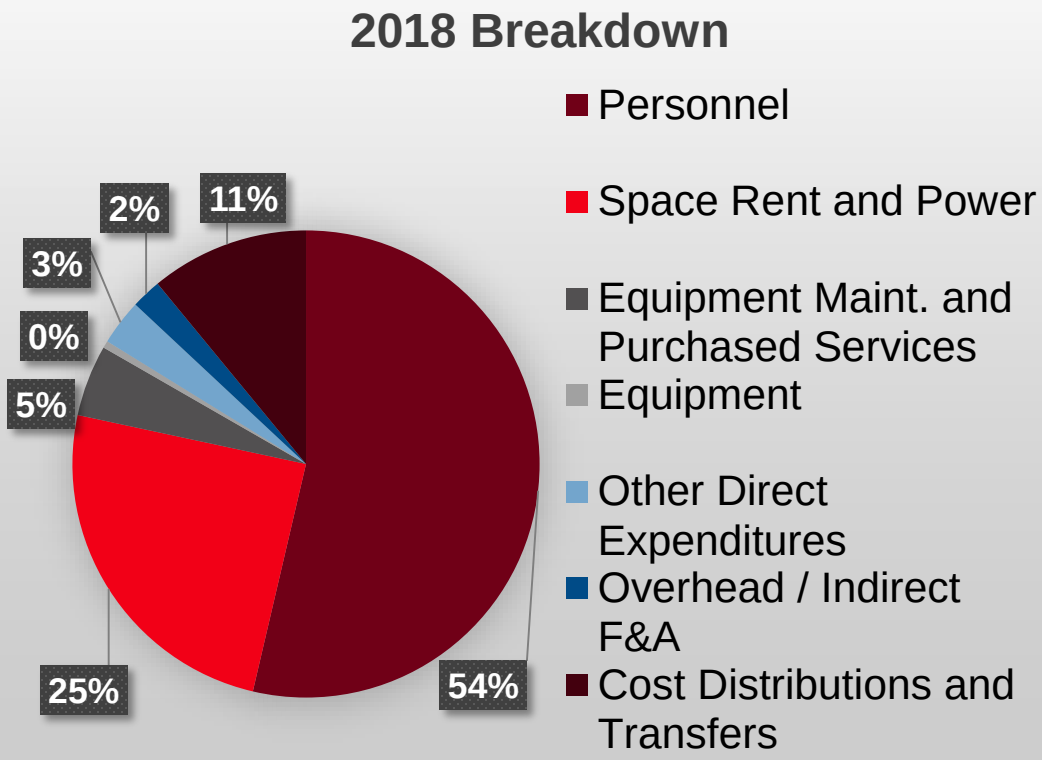

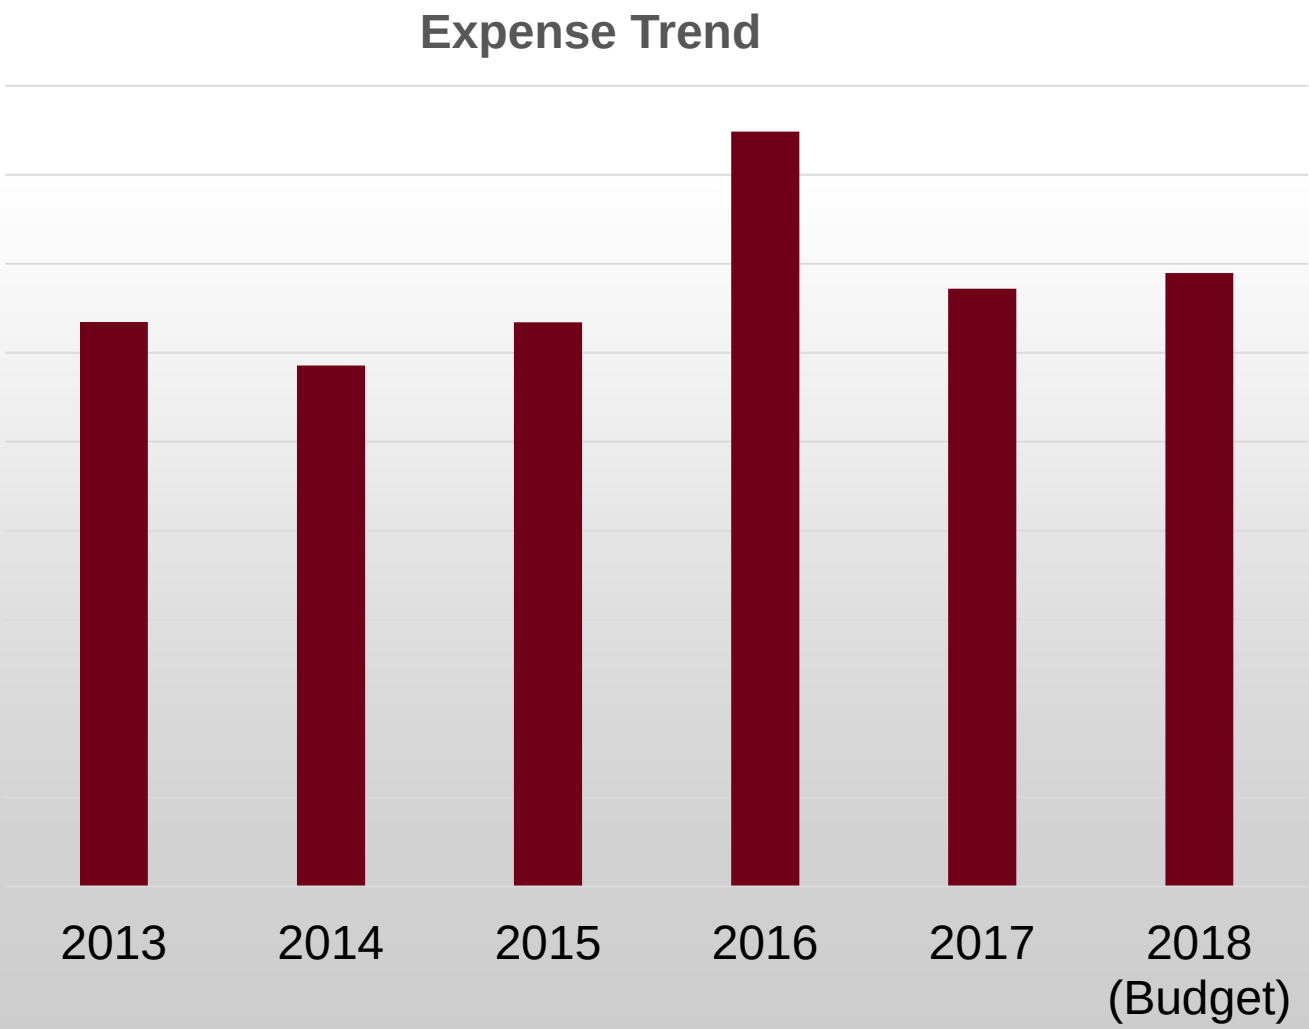

Budget Trends – Overall

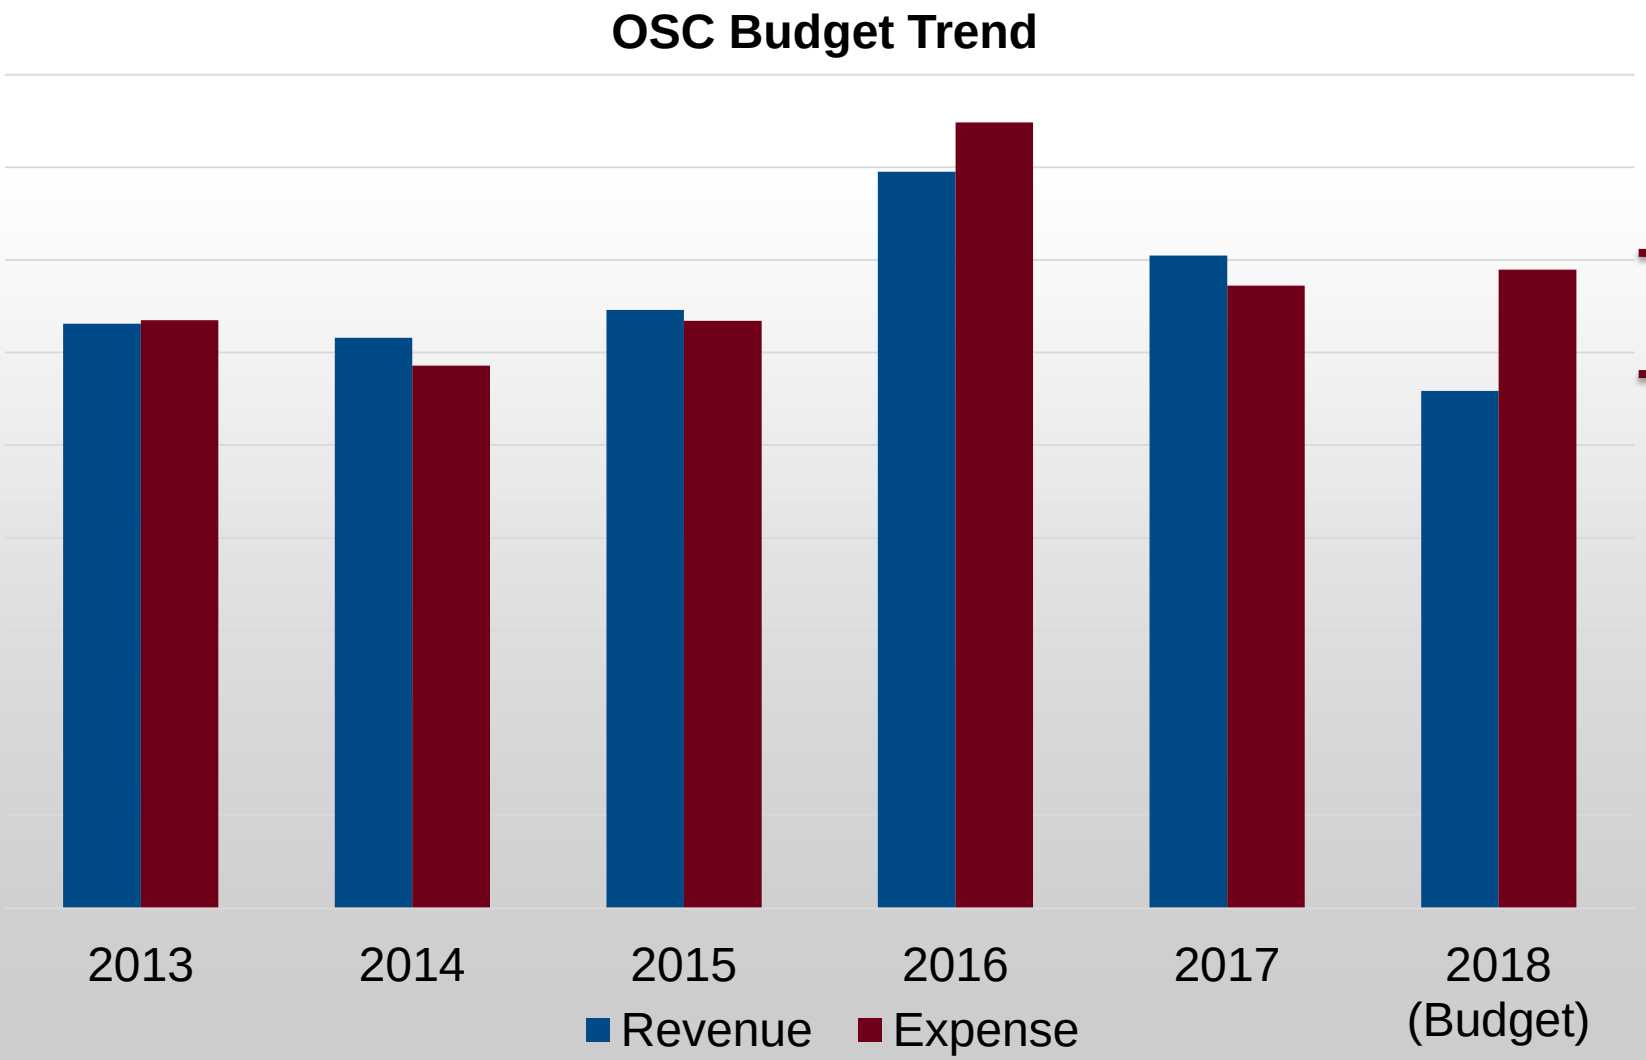

Projected funding gap is a result of:

Decreasing Funding
Increasing Expenses

Impact of SOCC Charges

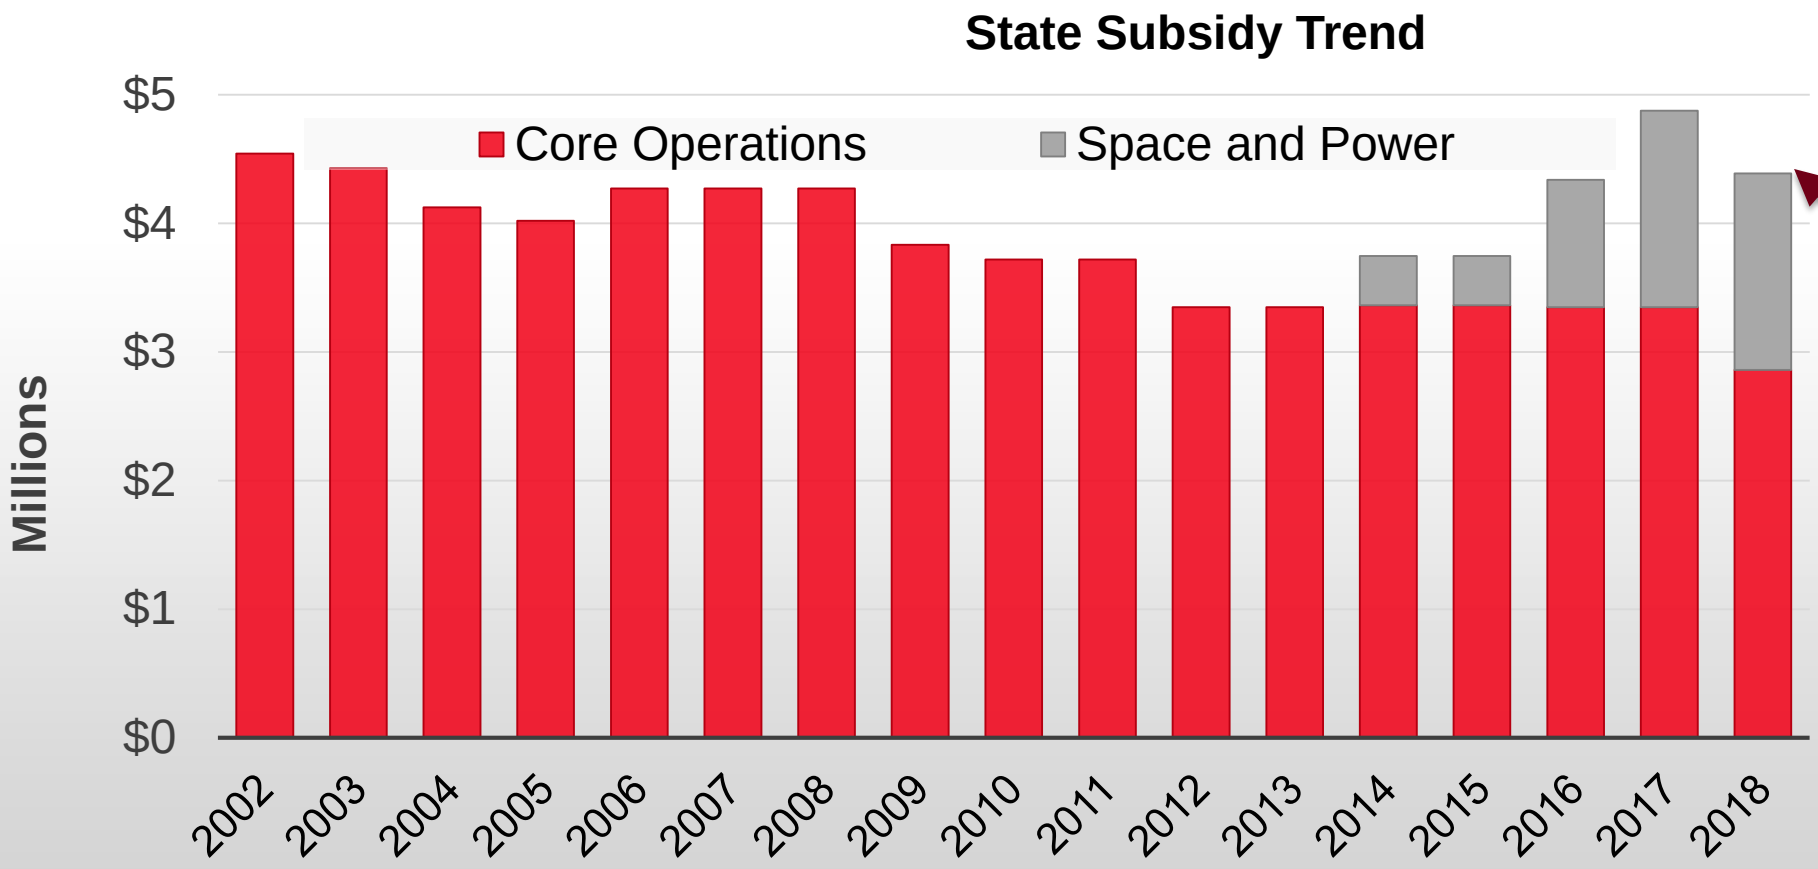

10% Decrease in FY18

- Data center costs fully transferred to OSC by 2017
- Permanent, substantial addition to OSC's expenses
- From 2008 to 2018, OSC lost nearly $1.4M for core operations

OSC Business Model Proposal

- OSC is an initiative of the Ohio Department of Higher Education (ODHE)
- Decision was made to ask universities for support
  - Need a single policy to apply to all Ohio universities
- ODHE has experience with fee models for services (e.g. OARnet, OhioLink), wants to do the same thing with OSC
  - Faculty chargeback policies up to each university

OSC Business Model Proposal Features

- Maintain subsidized access model
  - Subject to peer review by OSC Allocations Committee
  - Fully subsidize a portion of ALL awards
- Charge fees based on compute (RU) usage
  - Options for institutions to pay centrally to lower per-RU fees
  - Establish predictable costs over multiple years
- Create an advisory model similar to OARnet
  - User/Institutional-based Finance Committee
  - Each fall, Committee will develop pricing for coming fiscal year

We are doing everything we can to constrain costs and will continue to offer heavily subsidized services that are well below market rates.

Business Model Proposal Annual Process

Step 1: Project resource usage

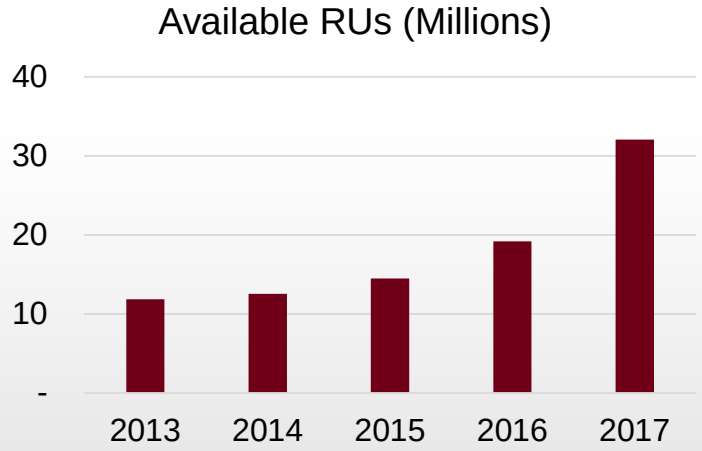

Step 2: Quantify recovery target

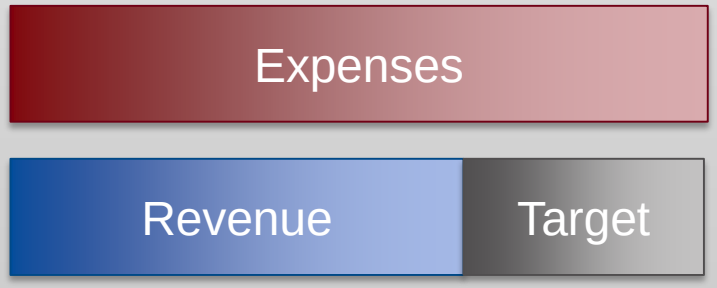

Step 3: Subsidize 1st 10K RU for all clients

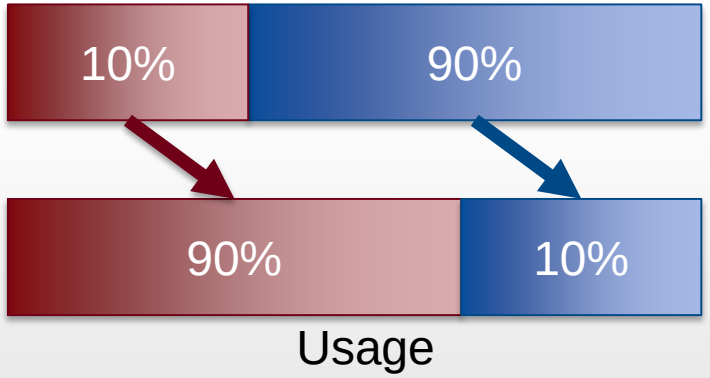

Step 4: Divide target by usage

RUs	Unit Cost
<10,000	No Cost
>10,000	$0.17/RU

DRAFT COSTS

Step 5: Incentivize central payment

Payment	Discount	Unit Cost
$10,000	20%	$0.136/RU
$30,000	40%	$0.102/RU
$60,000	60%	$0.068/RU
$350,000	80%	$0.034/RU

Step 6: Calculate cost/university

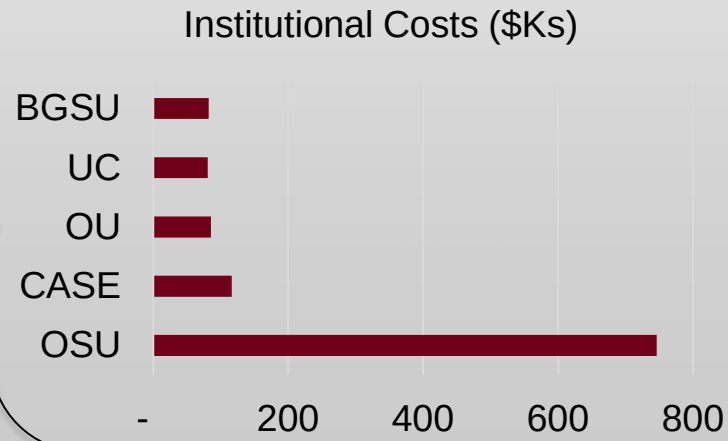

Alternatives to OSC for Example 100K RUs / Year Client

Local Cluster

- 1 Owens node = ~22K RUs / year
- 100K RUs = ~5 nodes
- Typical node upfront cost = ~$5,500
- Typical node annual operating cost = ~$1,900
- 5 node cluster cost = ~$28K upfront, ~$10K/year
- 5 node cluster five-year lifetime cost = $78K
- Additional costs for:
  - system administration services
  - software licenses
  - data storage and backup services

Amazon Cloud Services

- 1 Owens node ~= Amazon r4.8xlarge instance
- 100K RUs = ~5 instances
- Typical instance upfront annual cost = ~$12K
- 5 instance annual cost (paid upfront) = ~$60K
- 5 instance five-year lifetime cost = ~$300K
- Additional costs for:
  - data transfer
  - system administration services
  - software licenses
  - data storage and backup services

These costs are hard to quantify generically since each client has specific unique needs, although they are covered by OSC today.

Comparison to OSC for Example 100K RUs / Year Client

Institution	Annual Central Payment	RU Rate	Annual Cost for 100K RUs
Local Cluster	$0	$0	$15,600
Amazon	$0	$0	$60,000
Ohio State University	$350,000	$0.034	$3,060
Case Western Reserve University	$60,000	$0.068	$6,120
Ohio University	$30,000	$0.102	$9,180
University of Cincinnati	$30,000	$0.102	$9,180
Bowling Green State University	$10,000	$0.136	$12,240
All Other Institutions	$0	$0.170	$15,300

Does NOT include:

- System administration
- Software licenses
- Data storage
- Backup services

These are based upon projected institutional usage and lowest overall impact to each institution.

Client Feedback to Date

Representative quotes:

- "I would have to downgrade my research program to fit within the computation that I can afford, which is basically nothing at the rates that you are proposing."
- "Funding rates are plummeting, and the amount per grant is dropping in relative terms…If costs are pushed onto the grants that mainly support the salaries of young talent, the obvious outcome is that the workforce must be reduced."

Summary of comments:

- Paying for compute from direct grant funds hurts ability to hire students
- Concern over cost containment
- Concern over long-term cost predictability
- Desired flexibility to spend funding when available
- Explore charging for other services (e.g. storage, GPUs, big-data, HIPPA)

Status and Next Steps

- Discussions with provosts from 6 institutions this summer
- DHE and OSC will continue to engage universities to discuss the proposed model and consider adjustments
- OSC and Finance Committee will be looking at budget and usage projections
- Begin charging in FY19 (July 1, 2018) with the goal of reaching $1.5M in funding by FY20

Questions?

Campus Visits

- Bowling Green State University (January)
- University of Cincinnati (April)
- Ohio State University (March, June)
  - Consultation hours at Research Commons every other Tuesday

Upcoming Training and Events

- Introduction to Supercomputing / Big Data at UC: October 10th
- Big Data at OSC: October 26th

See https://www.osc.edu/events

OSC Supercomputers + Storage

Systems

	Oakley	Ruby	Owens
Date	2012	2014	2016
Cost	$4 million	$1.5 million	$7 million
Theoretical Performance	~154 TF	~144 TF	~1600 TF
Nodes	692	240	824
CPU Cores	8304	4800	23392
RAM	~33.4 TB	~15.3 TB	~120 TB
GPUs	128 NVIDIA Tesla M2070	20 NVIDIA Tesla K40	160 NVIDIA Pascal P100
	Total compute: ~1900 TF		

Storage

	Home	Project	Scratch	Tape Library (backup & archive)
Capacity	.8 PB	3.4 PB	1.1 PB	5+ PB
Bandwidth	10 Gbps	40 Gbps	100 Gbps	3.5 Gbps

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 GPU Rollout and Adoption

- 160 Nvidia P100 GPUs were made available to clients in the beginning of April 2017

- Usage:
  - ~50% Molecular dynamics (e.g. Amber, Gromacs, Namd, Lammps)
  - ~50% Machine learning / neural networks (e.g. Tensorflow, Caffe, Torch)

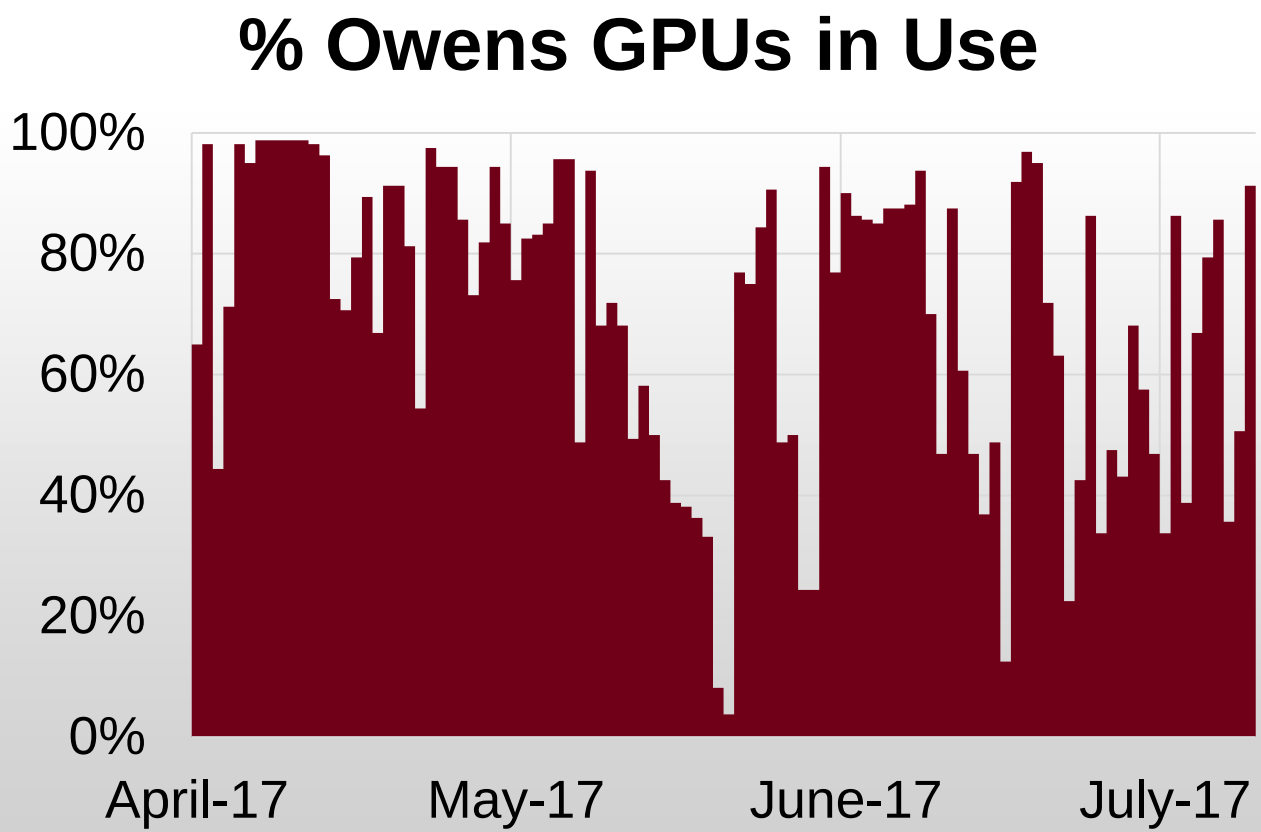

System Monitoring

Client Portal Project

- A user friendly client portal that is a “one-stop” shop for registration, project management, profile updates, etc.

- Status:
  - Internal user testing underway, external client testing beginning shortly
  - Development timeline expected to finish first week of October
  - AweSim website in process of being transferred to internal hosting

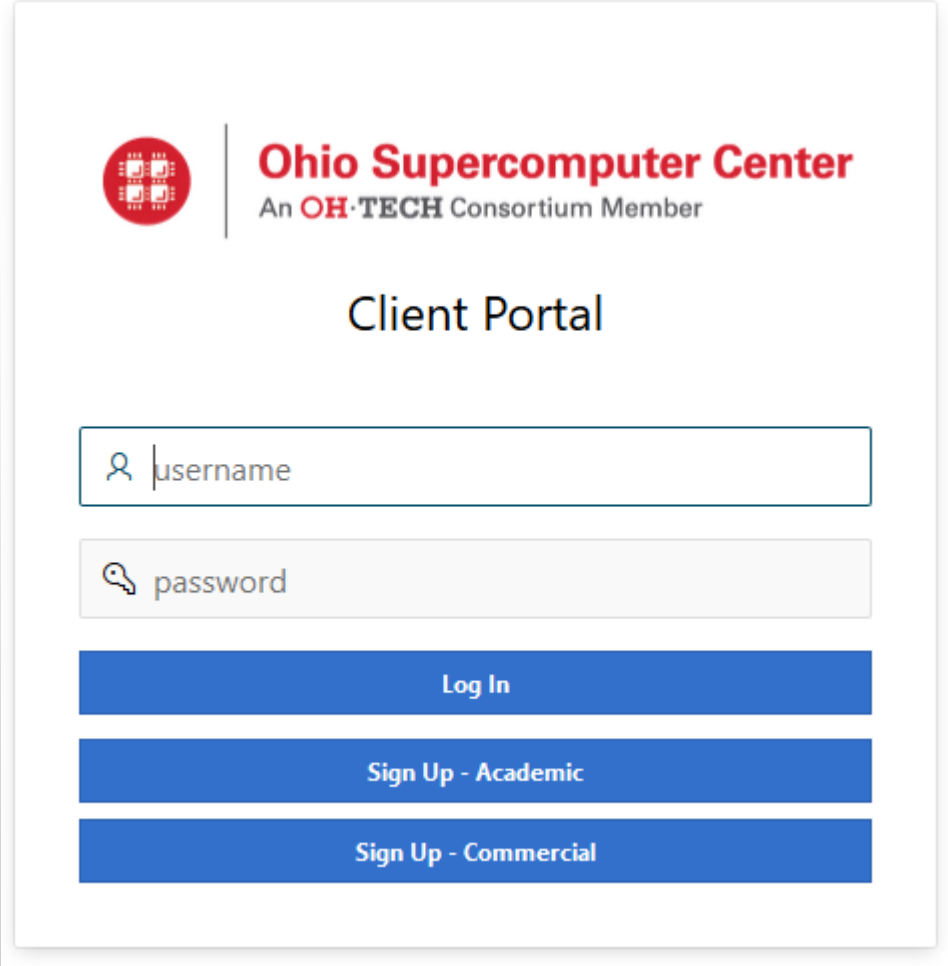

FY17-18 Roadmap of OSC Resources

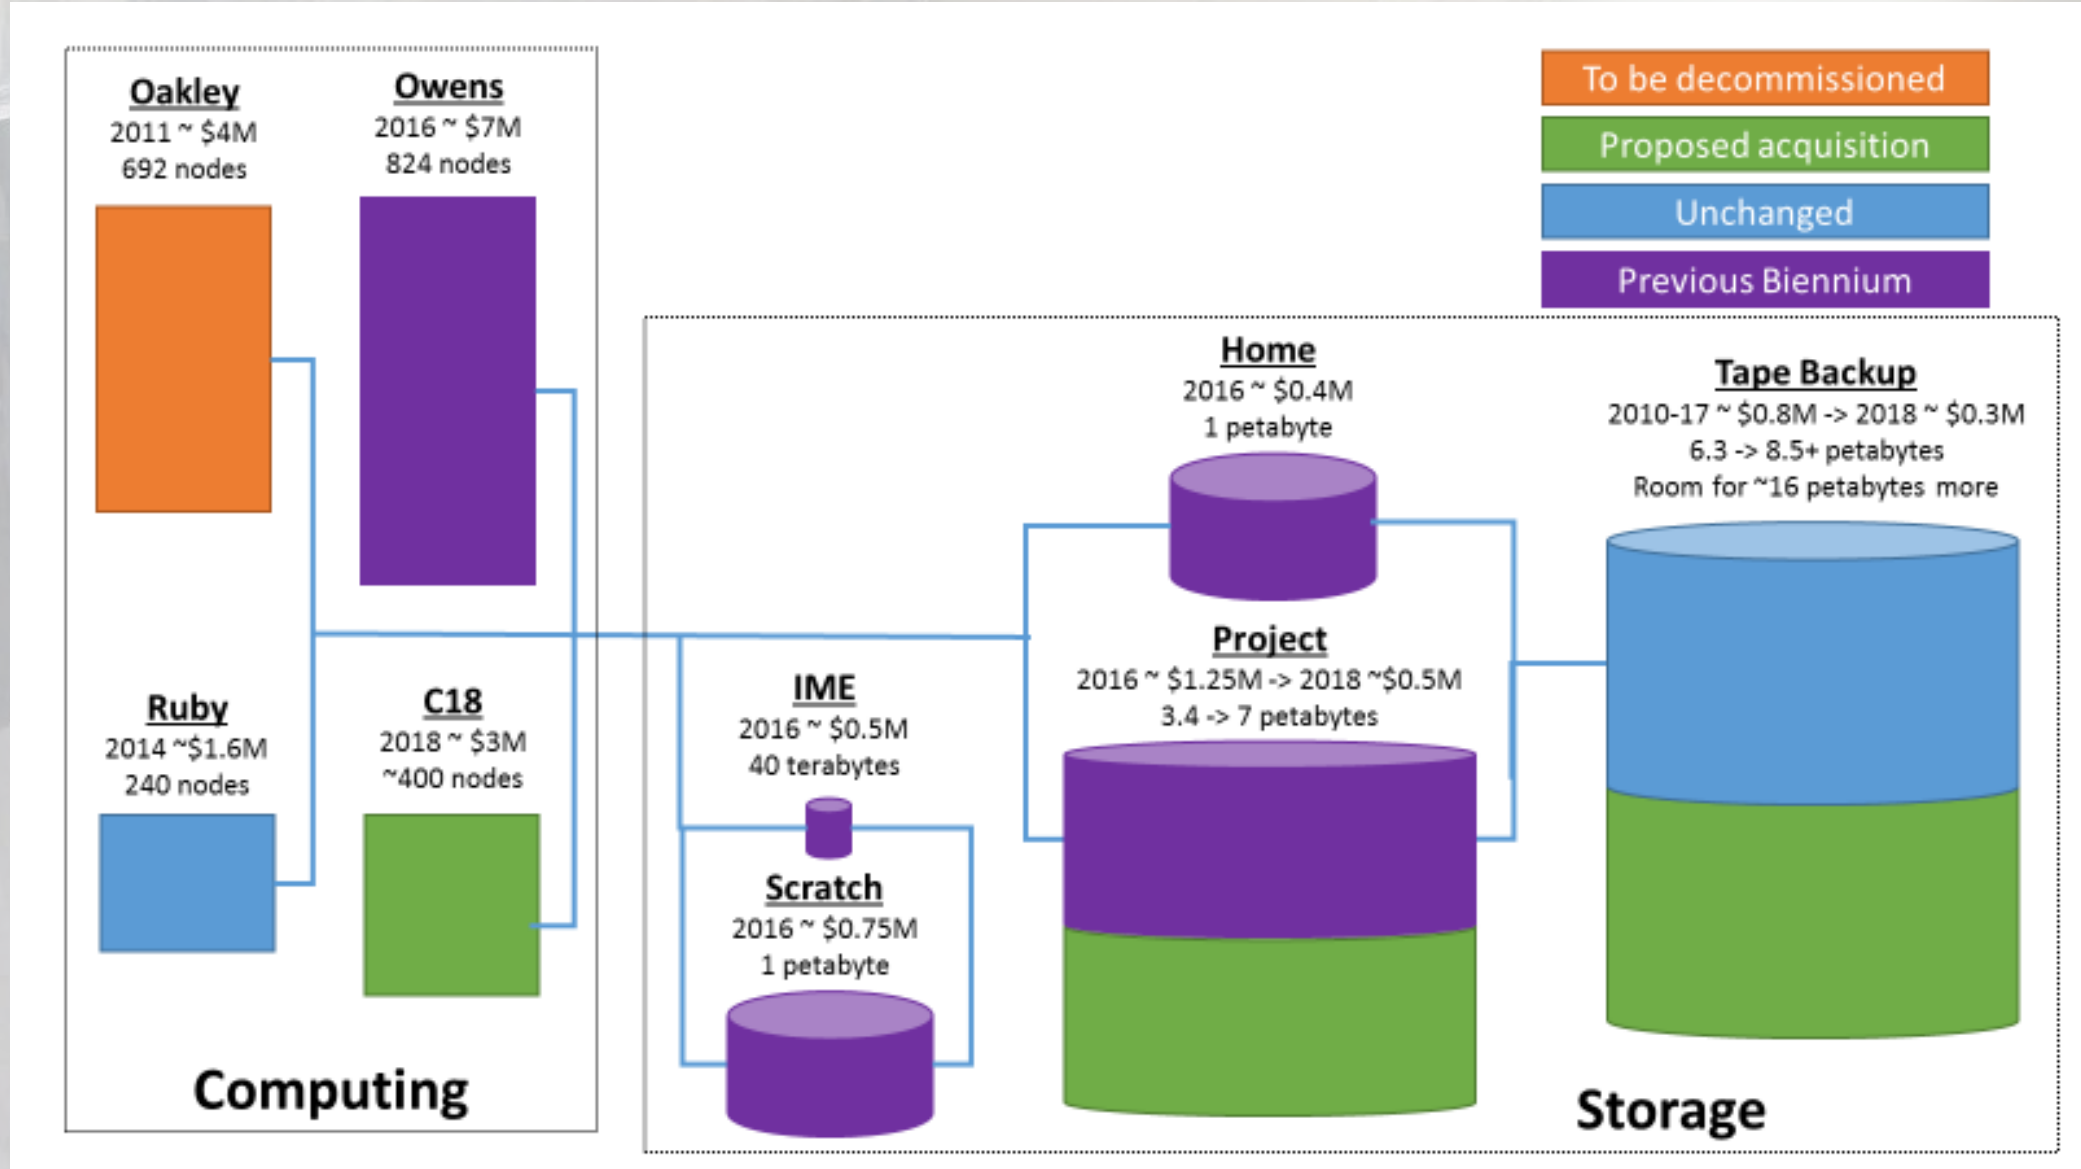

New Cluster (C18) Vision / Timeline

- Vision
  - Complement Owens
  - Dense compute component
  - GPU computer component
  - Big data component
- Timeline
  - RFI issued Sept 25
  - RFI due Oct 22
  - RFP issued Nov 2
  - RFP due Dec 15
  - Facilities updates May
  - System delivery June