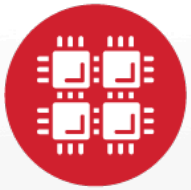

Ohio Supercomputer Center

An OH·TECH Consortium Member

July 2014 HPC Tech Talk

Agenda

- Mission
- WebEX tips
- Overview of service updates
- Attendee-driven discussion
- "Tech Notes" (30 minutes) – OSC Roadmap
- Slides are available at
  - http://www.osc.edu/tech_talks

Mission

- We want to better engage the daily users of the system
- Provide another avenue for the community to raise issues and talk about unmet needs.
- This event is for you! Please ask questions, make comments, and provide feedback as to how these can be improved to better serve you.
- https://www.surveymonkey.com/s/W3Y2JLZ

WebEX tips

- You can use the “Raise hand” icon to ask a question or make a comment (this will notify us so that we can acknowledge you)
- You may also use the Q&A section to ask questions.
- Please mute your microphones when not talking, to avoid feedback and interference noises.

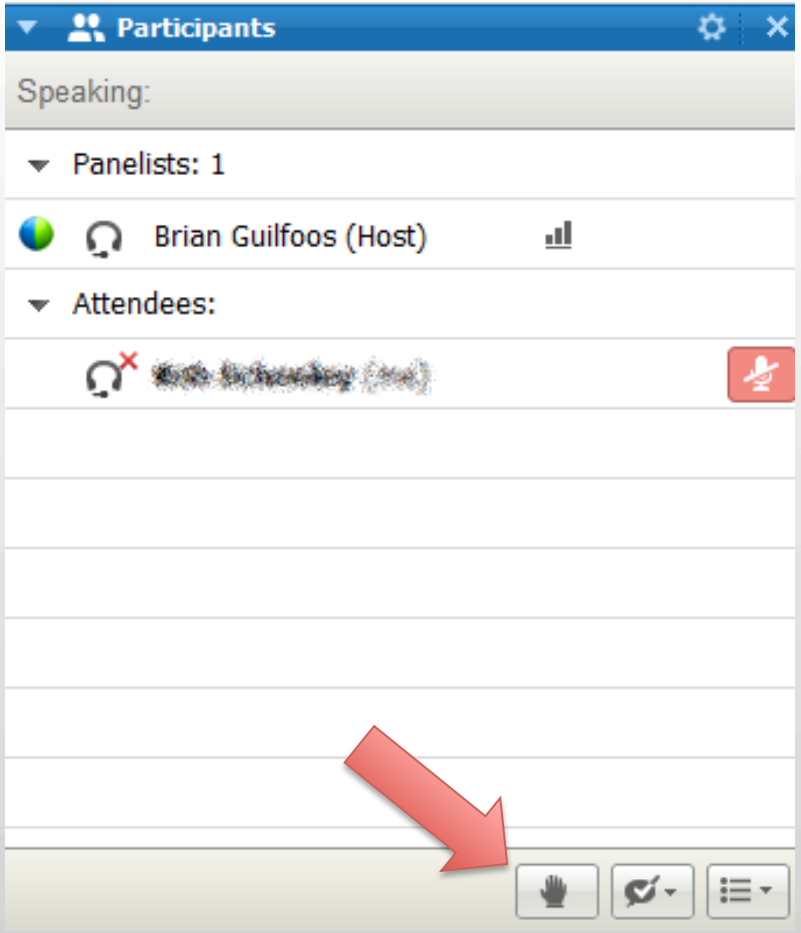

Introducing MyOSC!

- We have launched a new service called MyOSC
  - https://my.osc.edu
- Currently allows changing password, email address, and shell.
- Can regain access to your account if you forgot your password.
- More features in development!
- These functions have been removed from ARMSTRONG

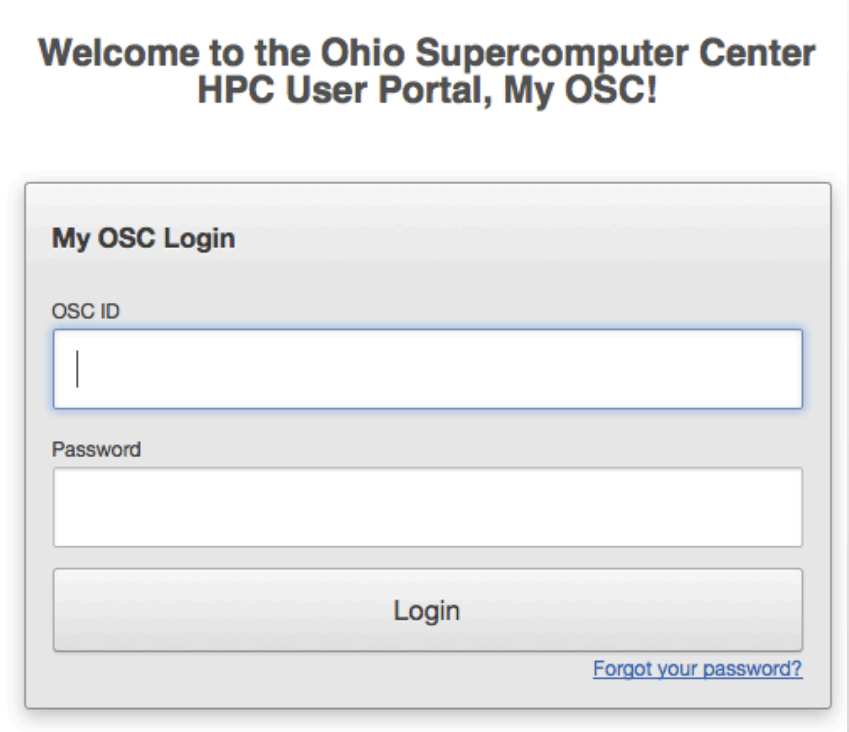

“Changelog” now available on osc.edu

- Significant configuration changes will be reported here
  - http://www.osc.edu/supercomputing/changelog
- Can view individual years (“…changelog/2014”) or filter by HPC.
- Will be embedding the changelog in the sidebar on some pages.

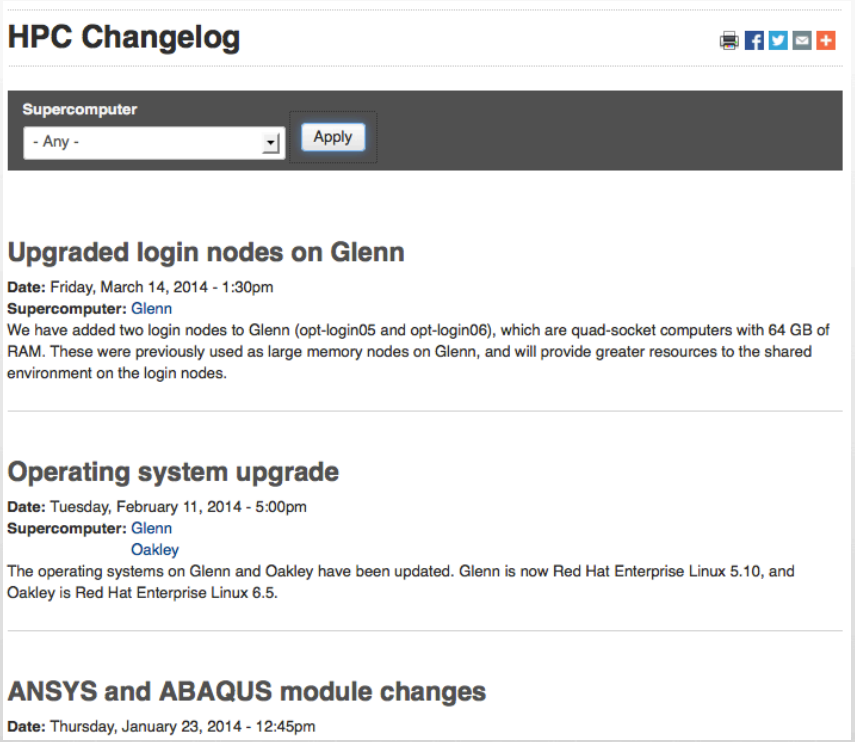

SUG Software Focus Groups

- We'd like your input on decisions regarding software installed for general use (via modules) on our systems
- To indicate interest, please visit http://goo.gl/Dfsfqp
  - Use the access code "OSCSUG"
- Establishing several initial focus groups:
  - Bio informatics/Bio Sciences
  - Fluid Dynamics
  - Structural mechanics
  - Quantum Chemistry/ Materials
  - Physics
  - Atmospheric and Environmental Modeling

AweSim Apps

- We are looking for ideas for the AweSim app store, and for developers for beta testing.
- Please contact us if you are interested.
- http://www.awesim.org/

July Downtime

- We had a two-stage downtime to conduct maintenance on our infrastructure
- Experienced some difficulties with the Lustre upgrades, which resulted in the filesystem being unavailable for the rest of that week.
- Recently, we've still been seeing intermittent performance issues.
- Please contact OSC Help whenever you notice an unexpected change in behavior. Users occasionally find problems we did not encounter in testing, or issues not caught by our automated monitoring

Glenn is still in production!

- Long waits on Oakley, but often no wait at all on Glenn
- We can provide assistance in migrating your jobs to Glenn, or helping you decide if you will benefit from switching.
- If there is anything preventing you from using Glenn (for example, missing software), please let OSC Help know.

Ruby cluster available for general use late this summer

- Ruby is a small 8 node cluster that has been used for research purposes
- We will be buying 240 nodes for the cluster this summer (4800 cores)
- Ruby will much larger than Glenn and nearly as large as Oakley (in peak performance)
- Large number of newer NVIDIA accelerators and Intel Xeon Phi accelerators (20 nodes of each)
- We will be retiring a portion of Glenn to free up the power necessary to expand Ruby.

Charging policy change implemented

- Serial jobs on Oakley are charged a number of cores proportional to requested memory
  - 4GB of RAM per core on Oakley compute nodes
  - nodes=1:ppn=1:mem=12GB would be charged for 3 cores
- Memory use limited to the amount requested, or the implied amount (4GB * ppn)
- No impact for whole-node jobs (including parallel jobs)
- More details: http://www.osc.edu/memcharging

Known Issues

- You cannot check your quota on GPFS. No workaround; the quota numbers reported when you log in are calculated once per day.
- If you block popups on Safari, some OnDemand features will silently fail. We are working on both a workaround and a long-term fix.
- Look at the bottom of http://www.osc.edu/supercomputing for up to date issue reporting.

Discussion period

- Floor open to questions of presented material, or for the community to raise other issues to discuss.

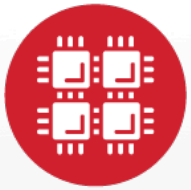

Ohio Supercomputer Center

An OH·TECH Consortium Member

OSC Roadmap Update

Dave Hudak, Doug Johnson, Brian Guilfoos, Karen Tomko

OSC Focus Areas 2014-2015

- OSC business model and service catalog
- “Condo” Pilot Project
- Ruby Cluster deployment
- FY15 capital budget and system acquisition

OSC Business Model and Service Catalog

- Goals
  - Streamline clients' access to our services
  - Demonstrate value of services to clients and stakeholders
  - Create/customize services to meet client needs
  - Enhance OSC sustainability
- Outcomes
  - Create clearly explained, well-defined service cost sheet
  - Succinct value proposition for clients
  - Show value of services to clients and their administration
  - Show value (ROI) of state subsidy

OSC Service Catalog

- Majority of current clients fully subsidized by state
- Increasing requests for more and extraordinary capacity funded by research projects and startup packages
  - Clients want option to augment our capacity with their funds
- Two primary mechanisms: Lease or Purchase (“Condo”)

Services	Lease	Purchase
Compute	Per core-hour	Per node
Storage (Project)	Per TB/month	Per disk
Archive (Tape)	Per TB/month	Per tape

“Condo” Pilot Project

- Faculty from Physics and Mechanical Engineering at OSU purchasing nodes in Ruby, our new cluster
  - Combining purchases: $640K worth of hardware for $550K
- Value proposition for OSC
  - Condo owners vacating allocations on Oakley, freeing it for other clients
  - Backfill on condo nodes offsets OSC workload
- Defining business case for expansion of condo services
  - Recovering some costs while leveraging subsidy for centralized infrastructure

Ruby* Cluster

- New cluster coming to OSC late summer
- Transitional cluster with following goals
  - Increase capacity to lessen load on Oakley
  - House the Condo Pilot Project
  - Expanded platform for accelerator development
- $1.6M purchase; ~30% funded by condo owners

*Ruby Dee is an actress, poet, playwright, screenwriter, journalist and activist. She was born in Cleveland.

Ruby, Oakley & Glenn: System Configurations

	Ruby System (2014)	Oakley System (2012)	Glenn System (2009, Phase II)
Theoretical Peak Performance	96 TF + 28.6 TF (GPU) + 20 TF (Xeon Phi) ~144 TF	88.6 TFs + 65.5 TFs (GPU) ~154 TF	53TFs + 6 TFs (GPU) ~60 TF
# Nodes	240	692	650
# CPU Sockets	480	1384	2632
# CPU Cores	4800	8304	5300
# / Kind of Accelerators	20 nVidia Tesla K40 20 Xeon Phi 5110p	128 nVidia M2070s	20 nVidia Quadro Plex 2000 S4's
Total Memory	~15.3 TB	~33.4 TB	24 TBs
Memory per Node	64 GB	48 GBs	24 GBs
Memory per Core	3.2 GBs	4 GBs	3 GBs
Interconnect	FDR/EN IB	QDR IB	DDR IB

Continued Growth In Accelerators

- Power efficiency is compelling
- Adoption is a concern
- Chemistry codes have been updated with GPU-capable versions on Oakley
- Will be working via the Software Committee to identify GPU-capable packages of interest to researchers

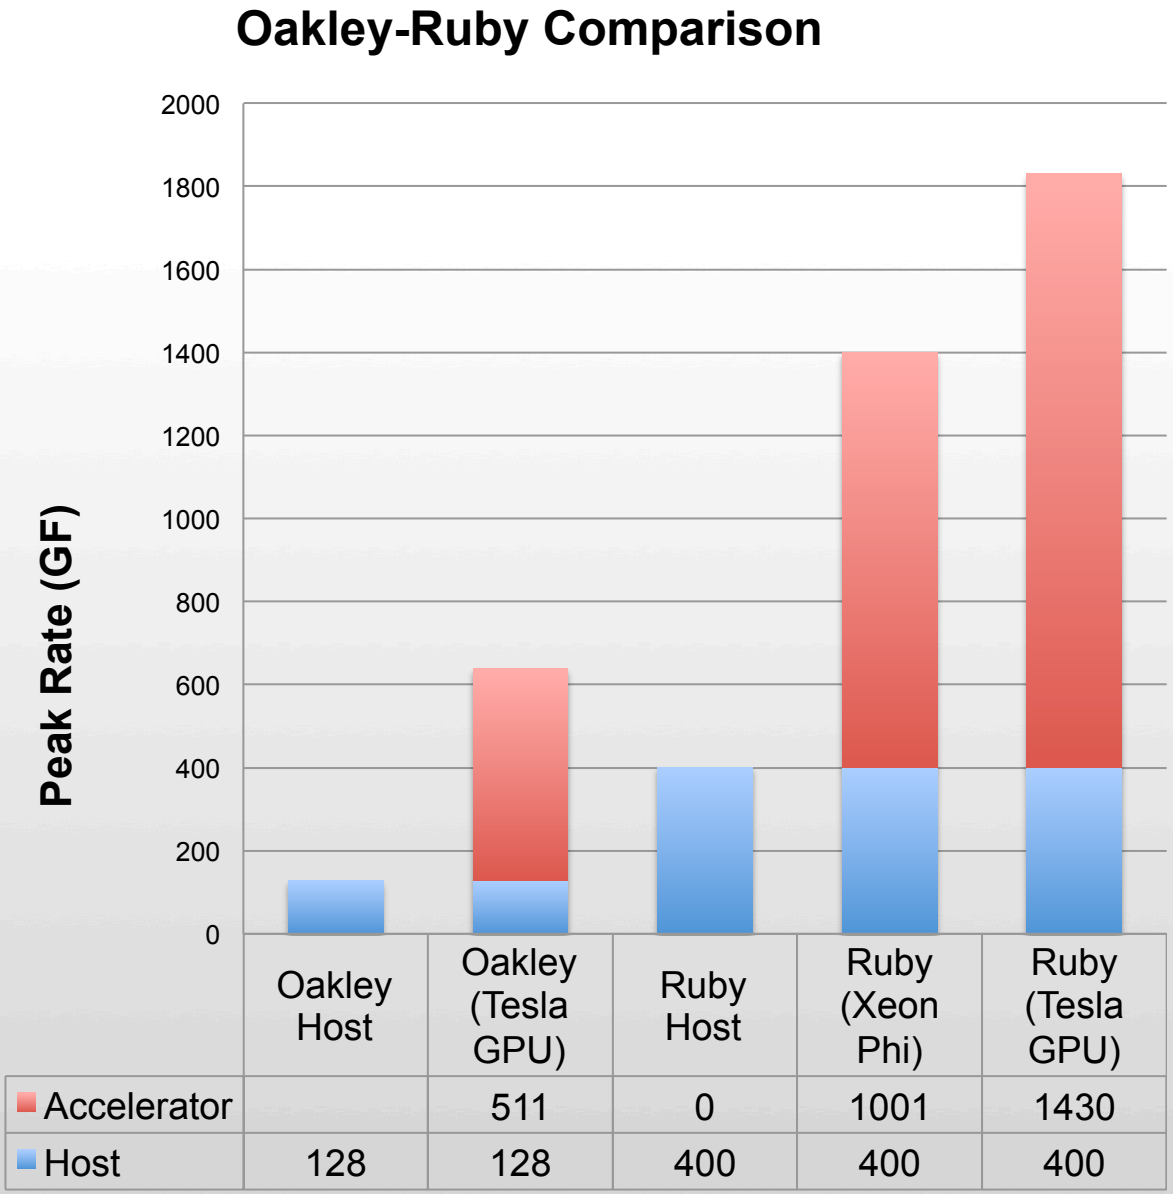

FY15 Capital Budget and System Acquisition

- Request for Information (RFI) responses from all major vendors, including IBM, Cray, HP and Dell
- Based on responses, our strategy:
  - A new cluster to quadruple capacity while maintaining compatibility, in early to mid 2015
  - Put Ohio in the Top 5 (likely) or Top 10 (definitely) of academic systems in the United States
  - Align purchases to major technology releases
  - Fit into current infrastructure without major facility modifications (i.e. power or cooling changes)
- $12M allocated in state capital budget

Roadmap Opportunities

- Promising near-term technologies
  - Intel Haswell processors
  - EDR InfiniBand (100 Gbps)
- Refresh strategy
  - What can we do to improve performance two years down the road?
    - Replace Haswell CPUs with Broadwell?
    - Replace Kepler GPUs with Maxwell?
    - Add new Knights Landing (Xeon Phi) systems?

Acquisition Timeframe

- RFP Publication starts process
- RFP Responses – 2 months
- Vendor Selection – 1 month
- Controlling Board approval and P.O. to vendor – 3-4 months
- Delivery and Installation – 2 months
- …in total, 8-9 month process

How Can You Participate?

- Many hardware options for purchase:
  - Node configuration (processors, RAM/core)
  - Interconnect selection and configuration
  - GPU and/or Xeon Phi?
  - Large memory
  - Storage

How Can You Participate?

- We'd like common needs for research communities
  - What can't you do at OSC?
  - Important benchmark examples?
  - Data intensive/analytics (Hadoop?)
- How do we collect input?
  - Survey: breadth of input
  - Interviews: depth of input
- It's more than just hardware
  - Policies, services, software?