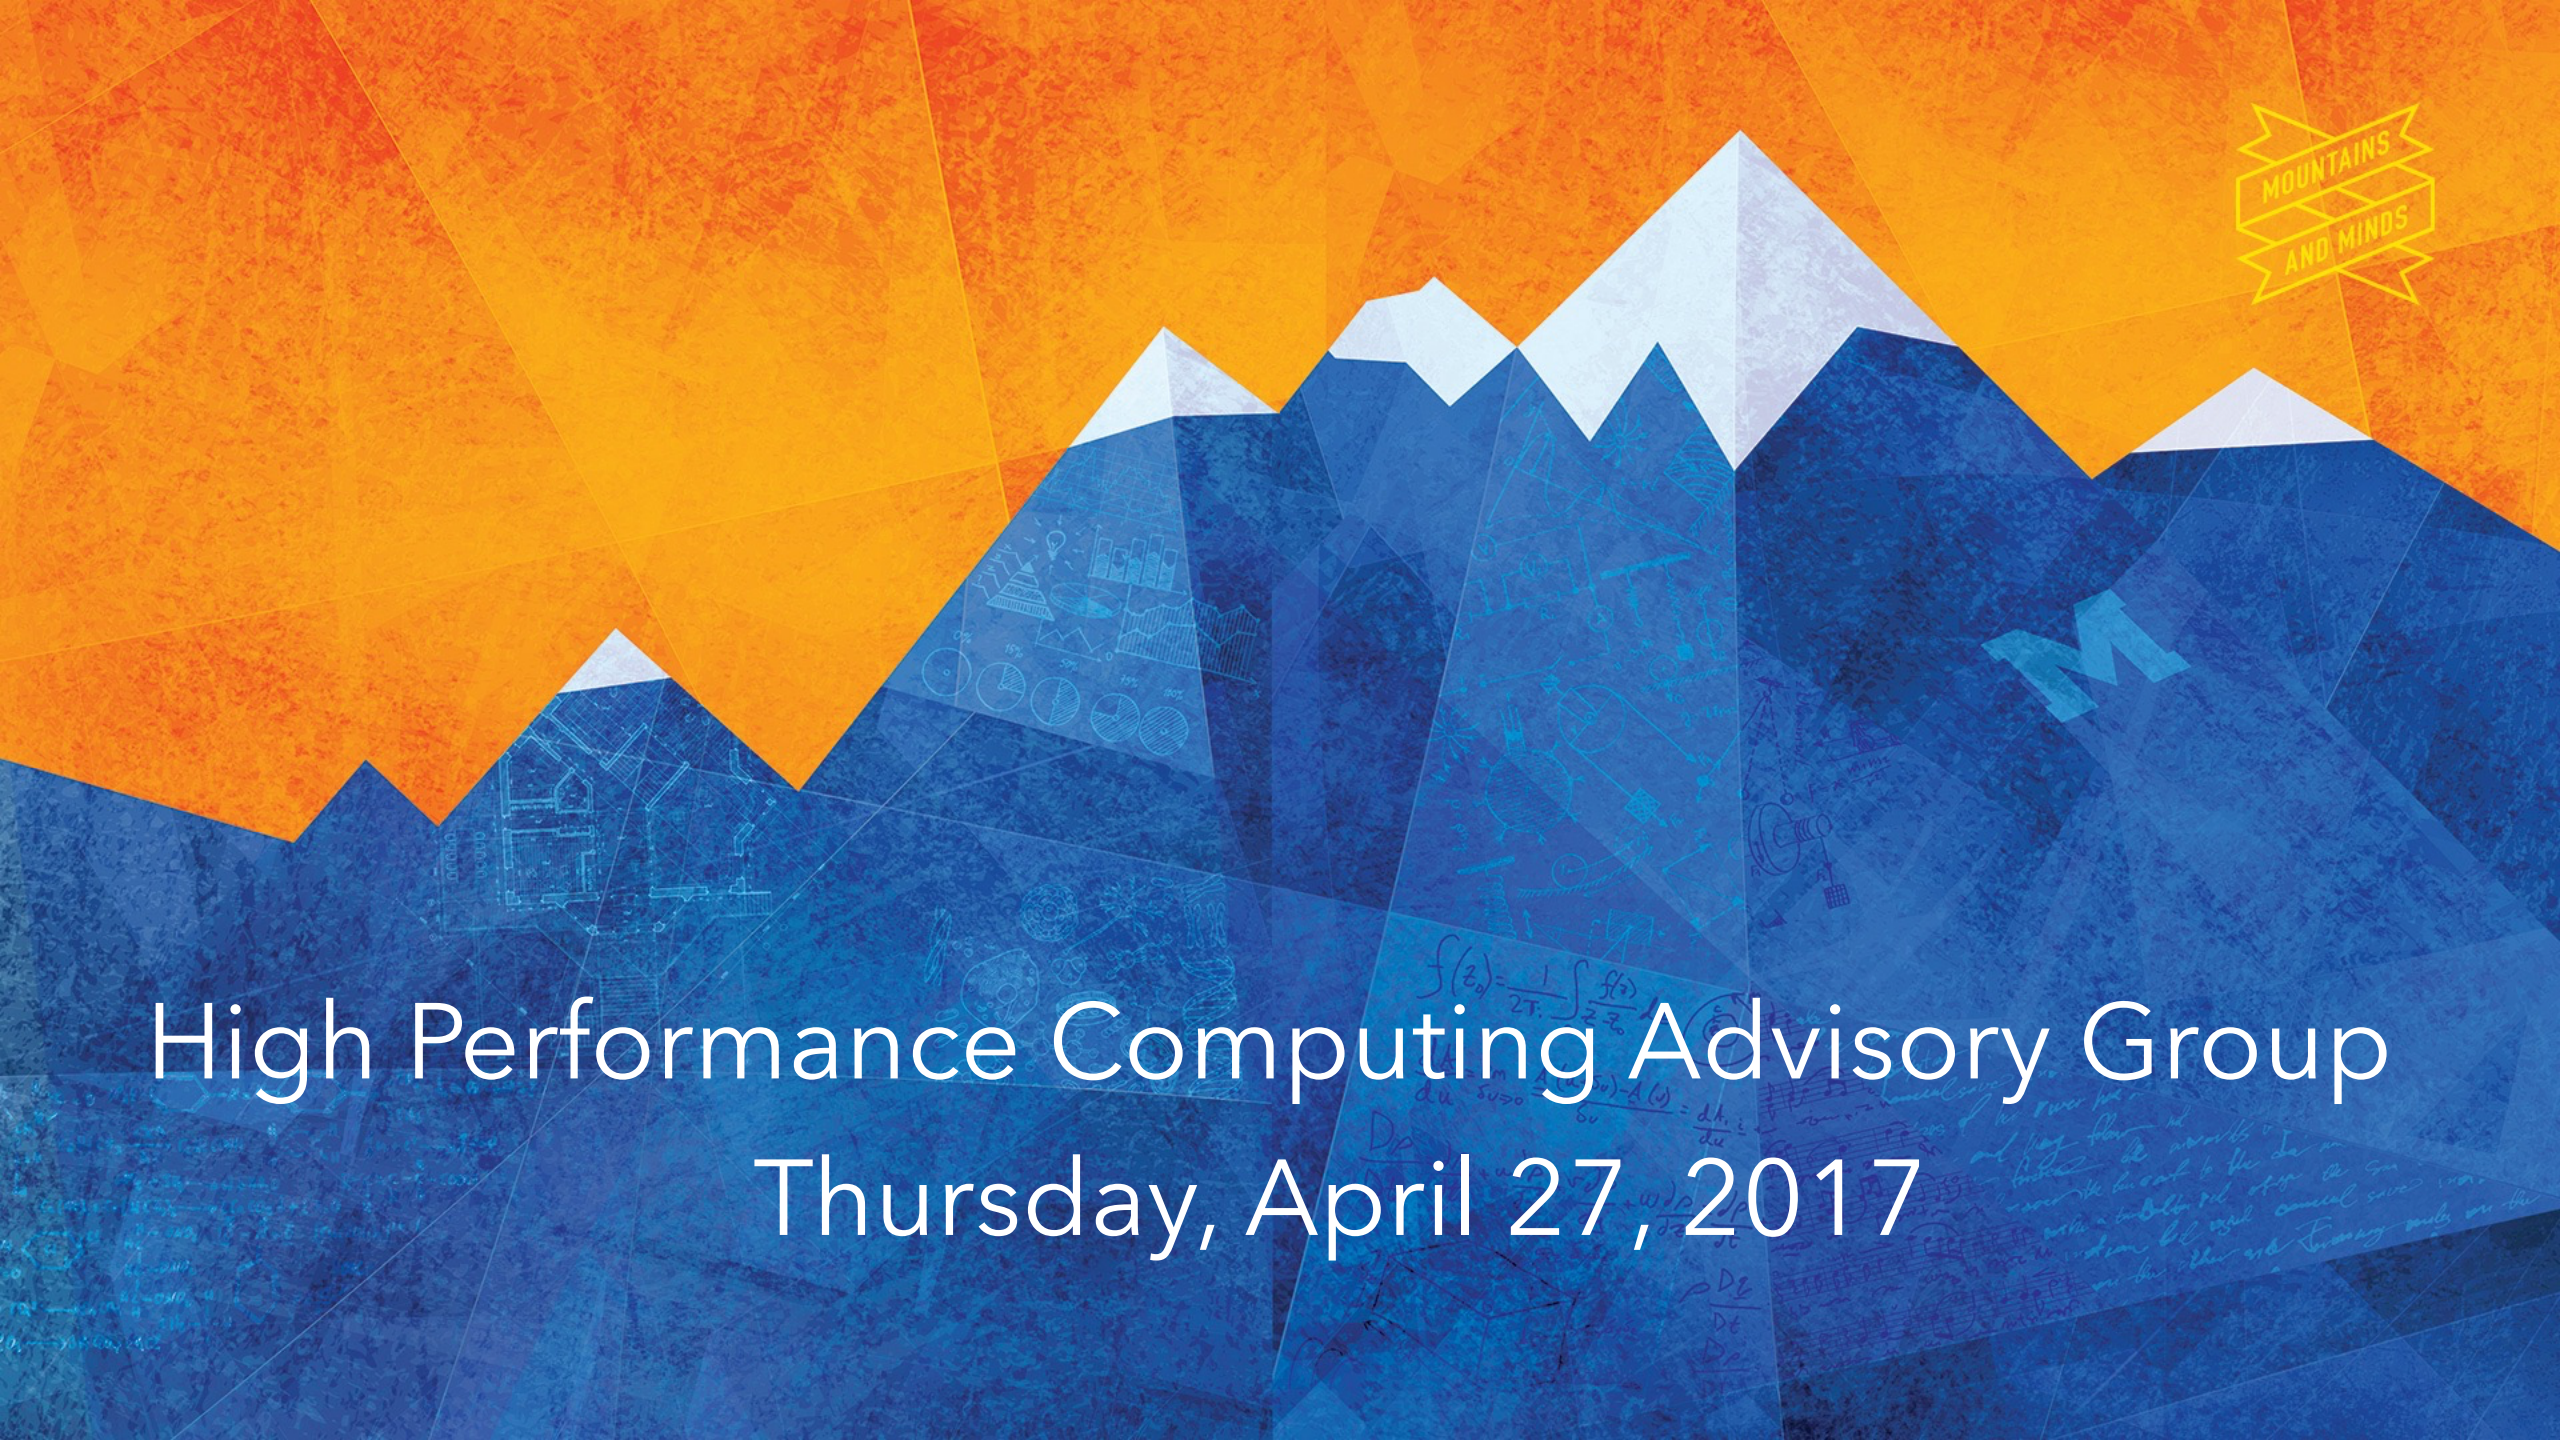

High Performance Computing Advisory Group

Thursday, April 27, 2017

AGENDA

1. Technical Issues
   a. Filesystem Issues
   b. Orphaned Process Issue
   c. CPU Performance Issue
2. CFAC Classroom HPC Acquisition
3. Increased Job Wait Time
   a. XDMoD: Job Wait Time
   b. Addressing Increased Wait Time
4. XSEDE Resources
   a. XSEDE Pioneers
   b. XSEDE Resources Comparisons
5. Faculty Candidate Topics
6. CONDA – a better way to manage software packages?
7. ALCES FLIGHT

MONTANA STATE UNIVERSITY

Lustre Filesystem Issues

- Incident Summary
  - Write failures (3/11 – 3/12)
  - Slow reads/writes (3/29 – 3/30)
- Impact
  - ~48 hours of outage, ~48 hours degraded.
  - No at-rest data loss, some loss from failed jobs
  - Some jobs unaffected
- Investigation
  - Intel: LSI RAID Driver
  - No NTP causing time drift
  - Also: HA failover – ping only
- Resolution
  - Servers Rebooted
  - NTP started on all nodes
  - Additional monitoring added
    - Splunk for log analysis
    - NetData for system stats

Orphaned Jobs

- Incident Summary
  - User reported jobs running slowly and/or crashing
- Impact
  - Unexpected CPU/RAM usage causing jobs to fail
  - 7 compute nodes affected
- Investigation
  - Non-SLURM processes running on compute nodes
  - Slurm failing to kill some jobs that go overtime
- Resolution
  - Worked 1 on 1 with researcher
  - Trained to check check jobs for rogue processes and clean up
  - Github issues posted, still in-progress

CPU Performance Problems

- Incident Summary
  - User reported jobs running very slowly
- Impact
  - Some compute nodes fail to speed up and also hang with CPUs stuck
- Investigation
  - Compute nodes idle at 1.2Ghz for energy savings
  - Process bumps the speed to 2.4Ghz when under load
  - Process hanging at 100% CPU
- Resolution
  - Set all CPUs to performance mode (no idle speed)
  - No reboot required, change applied hot to the cluster

CFAC Classroom HPC Acquisition

- Robert Szilagyi, CHMY 513 (Computational Chemistry)
  - Description excerpt: “comprehensive overview of computational chemistry methods”
- Continues previous CHMY 591 course
- Offered Fall 17 & Spring 17
- *Hyalite Installation Timeline*
  - May: Procure Nodes
    - 4 or 8 standard compute nodes, depending on budget
  - June: Install Nodes, create Classroom queue
  - July: Testing with Faculty of the Classroom queue and scripts

MONTANA STATE UNIVERSITY

XDMoD: Increase in Job Wait Time

- Average Wait Time is going up (last 90 days)
  - defq: 13.16 hours (12 hour target)
  - priority: 3.5 hours (2 hour target)

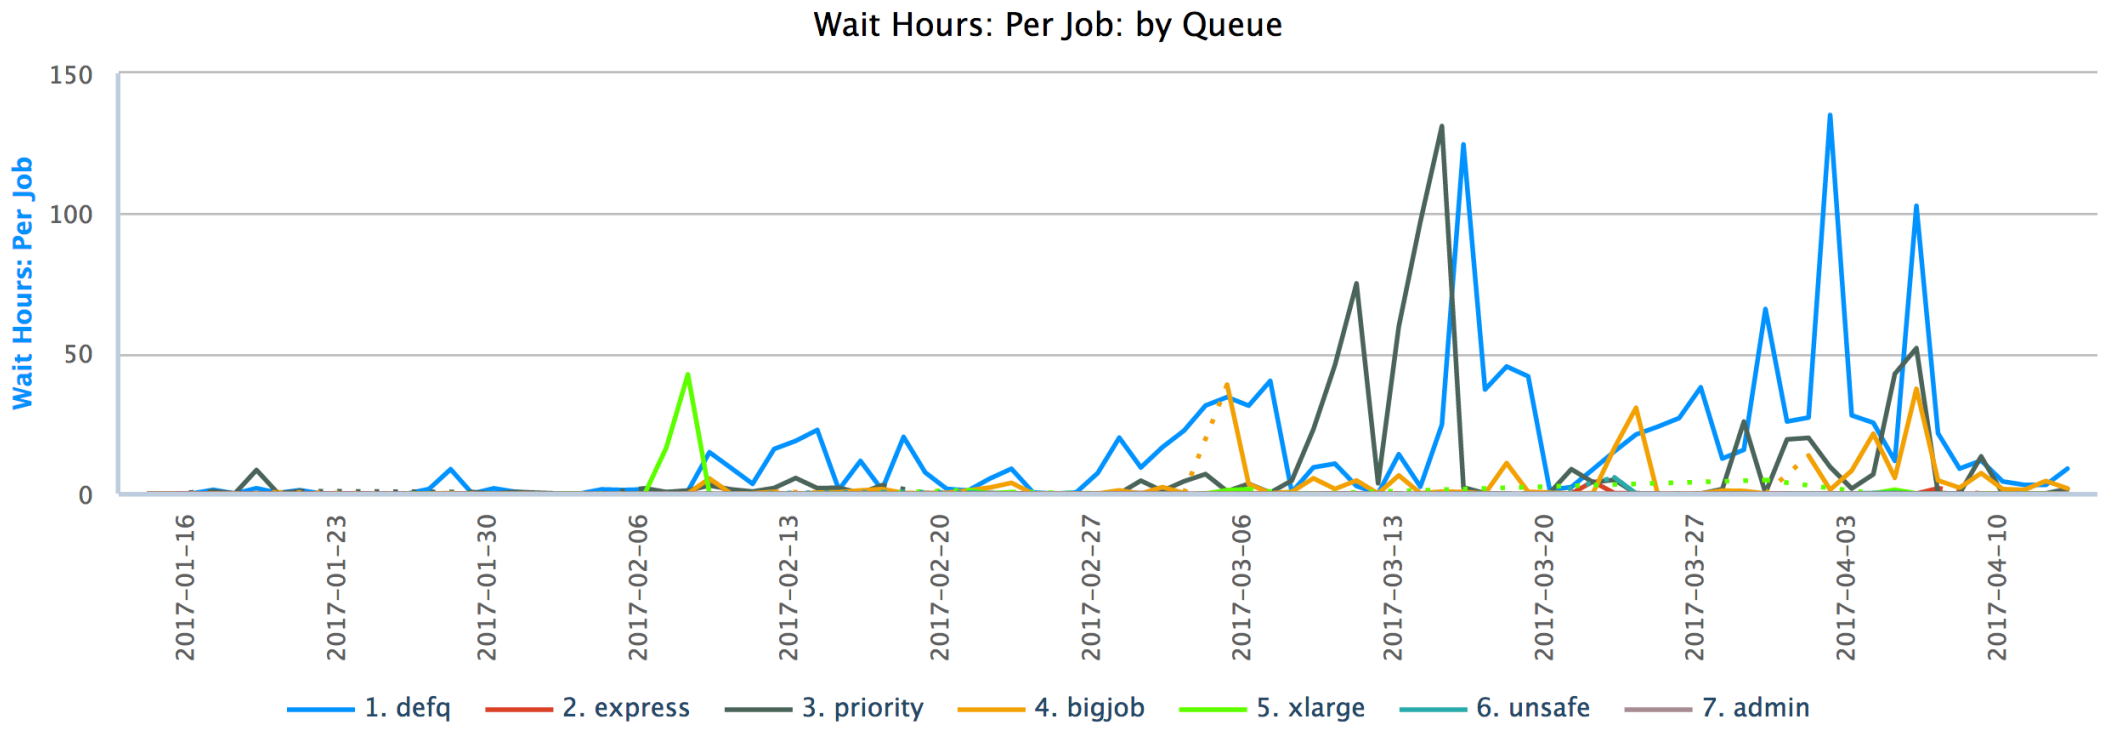

MONTANA STATE UNIVERSITY

Addressing Hyalite Job Wait Time

3 Ways to Reduce Wait Time:

1. Shorter Jobs (lower max runtime)
2. Add more Nodes (buy more)
3. Reduce Users (migrate users to XSEDE)

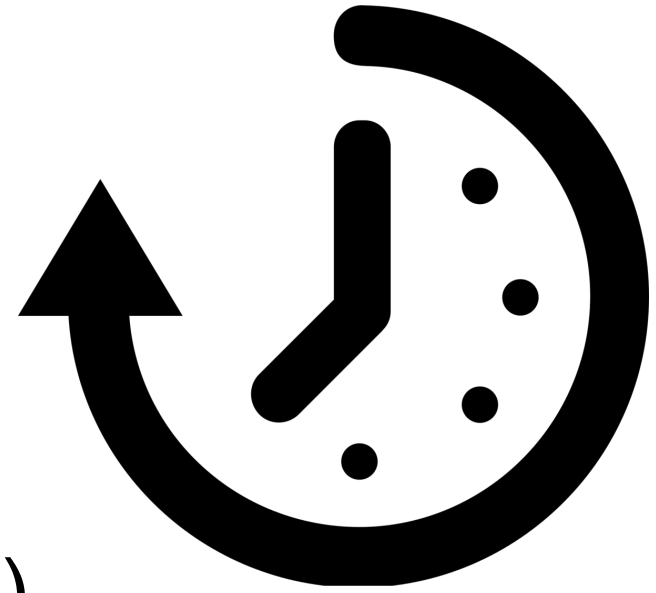

XSEDE Pioneers!

- Two birds
  - Reduces load on Hyalite
  - More benefit for researcher

- XSEDE Pioneers
  - Need a few heavy users
  - RCi Support Commitment
    - Getting connected
    - Adding/Compiling needed software
    - Porting SLURM scripts
    - Assistance writing Startup Allocation Proposals

XSEDE

Extreme Science and Engineering Discovery Environment

- *How much bigger than Hyalite?*
  - *Stampede: 6400 nodes*
  - *Comet: 1984 nodes*
  - *Bridges: 2632 TB of Memory*
  - *Jetstream 640 nodes*

MONTANA STATE UNIVERSITY

XSEDE Resource Comparisons

More comparisons of XSEDE resources:

Name	Organization	CPUs	Nodes	CPUs/ node	Peak Tflops	Mem/ node (GB)	Storage per node (GB)	Scheduler	Timelimit	Jobs (Monday)	Mean wait time (Monday)	Median wait time (Monday)	Jobs (last month)	Mean wait time (last month)
Stampede	TACC	102,400	6,400	16-32	9,600	32-1000	50	SLURM	1-10 days	1,300	6h 45m	2h 45m	48,320	17h 15m
Comet	SDSC	47,616	1,984	24-64	2,000	128-1500		SLURM	2-14 days	5,450	1h 30m	15s	424,380	2h 40m
Gordon	SDSC	16,384	1,024	16	341	64		Torque	2 days				17,710	2h 45m
Xstream	SRCC (Stanford)	1,300	65	20	1,001	256	450	SLURM	2-7 days				12,880	28m
Bridges	PSC (Pittsburgh)	21,056	752	28	895	128-3096	8	SLURM	2-14 days	700	5h	50m	63,240	6h 45m
SuperMIC	CCT (LSU)	7,200	360	20	925	64	500	Torque	3 days				4,390	8h 05m
Hyalite	MSU	1,080	65	32-40	18	64-256	1000	SLURM	1-infinite days	460	1h 50m	2h 25m	7,540	15h 50m

XSEDE

Extreme Science and Engineering Discovery Environment

MONTANA STATE UNIVERSITY

Faculty Candidates on HPC

- Met Faculty Candidates in Math and ME
- Infrastructure Questions
  - Is there on-campus HPC and storage?
  - Some interest in GPU computing
- Availability Questions
  - Who can use it?
  - What is the cost?
  - Difference between contributor and guest?
  - Is it hard/time consuming to get access to resources?
  - Is there training and support?

- MSU Resources
  - Excited about unlimited Box.com
  - Interested in XSEDE
  - Positive about cluster availability and cost
  - Would like more on-site research storage options.
- Users were Experienced
  - Familiar with HPC clusters and concepts

MONTANA STATE UNIVERSITY

CONDA – alternative package manager

- CONDA
  - Originally a package manager for python
  - Now manages many environments including: Perl & R
- Possibly a more powerful "**module**" command alternative; manages the environment in users home directory
- More portable (modules are not very portable, CONDA runs everywhere)
- Supports separate virtual environments, each with distinct package versions (if desired)
- Many domain-specific repositories available (mostly community-driven, and including bioinformatics)
- Deployed with success in Young and Budak labs

ALCES FLIGHT

Alces Flight is a commercial product that allows users to create virtual compute clusters of any size and capacity. A dashboard is available to monitor the cluster in real time, and a marketplace (Alces Flight Gridware) makes it straightforward to deploy scientific software packages on its nodes.

- Online platform that facilitates the deployment of custom, on-demand compute clusters
- Based on Amazon Web Services
- Clusters can adapt to compute loads, by adding or removing nodes automatically
- Clusters can take advantage of AWS spot pricing, dramatically reducing cost (by up to 80%)
- Clusters can be provisioned with open-source scientific software packages, from a list of tested packages (Alces Flight Gridware)

alcesflight

Alces Flight: http://alces-flight.com/
Alces Flight Gridware: https://gridware.alces-flight.com/software

MONTANA STATE UNIVERSITY