DSRI AI Activity Group Organizational Meeting 4/23/2026

Attendance: Kelley Branson, Jim Cox, Kacie Reed, Andrew Kuppinger, Ronnie Banerjee, Ravi Mangal, Marcelo Melo, Jesse Burkhardt, Jude Bayham, Mendy Smith, Anping Chen, Santiago Di Pietro, Bruce Draper, Tim Hansen, Adam Preston, Ben Shaby, Tyson Claffey, Bill Carpenter, Brandon Shanks, Matt Koslovsky, Rob Paton, Michael Kirby, Melissa Reynolds

Welcome

- Michael Kirby introduction and brief overview of "why activity group"
    - Foster things, collaborations, facilities
    - Expertise in this room to give feedback to university of leadership
    - Respond to university opportunities and develop proposals
- Take back to depart to show what we are doing, and make this more connected across campus
- Forming a Campus-Wide Community
    - Forming cooperation and bridging gaps

Where we are headed

- Moving HPC to Energy Institute (synergy with energy and computing)
- Working on getting level 2 (second floor) of HPC infrastructure as well. There is room to expand on HPC at Energy Institute
- Future HPC opportunity and resources also coming with UWYO via joint grant with DSRI
- Aiming to make it possible for researchers at CSU to do big data research in AI with energy efficient computing
- Aiming for Fall 2026 for launch of Bighorn (name pending approval) at Energy Institute (replacing Riviera)

Previous AI meeting

- Keynote speaker
- 2 panels of AI faculty
- If you missed out, and want to hear about future DSRI activities, email Brandon Shanks to be added to the emailing list and/or check the website for updates

Sustainable HPC

- Faculty can purchase or lease compute and storage nodes in a parallel file system
- Service Unit (SU) chargeback model

- Balance operating costs with usage revenues (non-profit)

Service Units

- Definition: A service unit is one CPU core hour
- A 2TB high memory node CPU is the equivalent of 4 SUs (total: 1024 SUs)
- One A100 GPU core is the equivalent of 100 SUs (total: 1400+448 SUs)
- One GH200 GPU core is the equivalent of 200 SUs (total: 200+72 SUs)
- Total (3144 SUs)
- Total (27,541,400 Sus/year)

Proposed fee structure

- Installation $1000 per node
- Yearly Fee $1000 per U per node
    - These fees can be deducted by scavenging units
- Storage $100/TB/ Year parallel file system (compute side)
- Working on purchasing more storage

Discussion on HPC expansion and charge back

- Working on incorporating all related costs into charge
- Purchasing more than just a yearly amount?
- Do you purchase in set units of storage? Or can you purchase the exact amount you would need.
- Can purchase a node and NOT open it up for other uses or to be scavenged
- What happens if hardware fails?
- When will researchers have numbers (to put this in grants they are writing currently)
    - Numbers coming very soon for faculty to use
- What if the system is not at 100% utilization? For a year or period of time?
    - Working in new area with “bighorn” and will use the first years as learning experience and transition period
- Benefits of having a system on site
- Could you bring storage? Will you charge for rack space?
- What about security levels for different data?
    - Trying to get answers now from Deans and have a clear idea by July 1
- Will there be a free tier of access to Bighorn?
    - Each SU must be accounted for, but that cost might not be directly charged to you. We are exploring opportunities and options for different ways colleges and units can pay for their researchers to run on HPC.
- Free time for new users?

- If you have a price, will this increase if the system is under utilized? This could make a spiral making the system more expensive and less users
- How do you give an opportunity to users who are newer to HPC or not experts on a command line?
- What about a classroom / workshops to bridge gaps?
- Microsoft teams for this group to continue the conversation