

# DSRI HPC/AI Activity Group Meeting

April 23, 2026

# Agenda

- 1 Welcome!
- 2 Who, what, why, when is the Activity Group?
- 3 From Riviera to Bighorn
- 4 Sustainability
- 5 Questions

## Vision

**To aid society through the discovery of knowledge from data.**

## Mission

**To support and enhance the Data Science research environment and facilitate faculty collaborations that stimulate the development of new analytical tools and applications.**

## Goals

**To harness multidisciplinary expertise across campus, build research teams, and cultivate the discipline of data discovery.**

# Who, What, Why, When is the Activity Group?

Who: You! A collection of CSU stakeholders.

# Who, What, Why, When is the Activity Group?

Who: You! A collection of CSU stakeholders.

What: A group to provide expertise, vision and feedback to university leadership on matters relating to HPC and AI, and lead proposals.

# Who, What, Why, When is the Activity Group?

Who: You! A collection of CSU stakeholders.

What: A group to provide expertise, vision and feedback to university leadership on matters relating to HPC and AI, and lead proposals.

Why: form a campus wide community to foster excellence in HPC/AI

# Who, What, Why, When is the Activity Group?

Who: You! A collection of CSU stakeholders.

What: A group to provide expertise, vision and feedback to university leadership on matters relating to HPC and AI, and lead proposals.

Why: form a campus wide community to foster excellence in HPC/AI

When: monthly meetings with additional activities as planned by the activity group





## Past:

- AI/HPC symposium 2025
- Keynote speaker and panel session 2026
- CC\* lunches
- list-serve

## Future:

- AI/HPC advisory group to OVPR
- Monthly meetings
- additional activities?

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

Spring 2022: Order Riviera equipment with OVPR support

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

Spring 2022: Order Riviera equipment with OVPR support

Fall 2023: Riviera running in mainsite

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

Spring 2022: Order Riviera equipment with OVPR support

Fall 2023: Riviera running in mainsite

Spring 2024: Awarded CC\* cyberinfrastructure planning grant

Fall 2024: Joint NSF MRI proposal with University of Wyoming awarded \$4M

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

Spring 2022: Order Riviera equipment with OVPR support

Fall 2023: Riviera running in mainsite

Spring 2024: Awarded CC\* cyberinfrastructure planning grant

Fall 2024: Joint NSF MRI proposal with University of Wyoming awarded \$4M

Fall 2025: Decision to expand HPC at new facility in the Energy Institute: Energy efficiency, Big Data and HPC/AI

# From Riviera to Bighorn

Spring 2020: Data Science Open forum on computing: Should HPC@CSU be top 500?

Spring 2022: Order Riviera equipment with OVPR support

Fall 2023: Riviera running in mainsite

Spring 2024: Awarded CC\* cyberinfrastructure planning grant

Fall 2024: Joint NSF MRI proposal with University of Wyoming awarded \$4M

Fall 2025: Decision to expand HPC at new facility in the Energy Institute: Energy efficiency, Big Data and HPC/AI

Fall 2026: Anticipated launch of Big Horn

- Faculty can purchase or lease compute and storage nodes
- Service unit (SU) chargeback model
- Balance operating costs with usage revenues (non-profit)

# System Overview

- $N_{\text{CPU}} = 22$  is the number of nodes with 128 CPUs. (2816 cores)
- $N_{\text{GPU}} = 3.5$  is the number of nodes with 4 A100 GPUs. (14 gpu core, 448 cpu core)
- $N_{\text{HM}} = 1$  is the number of high-memory nodes. (256 cores)
- $N_{\text{GH}} = 22$  is the number of Grace Hopper superchips. (1 gpu core 72 cpu core)
- Data transfer node

# 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)

These SU equivalences are similar to other institutions.

Total: 3144 SUs

Total: 27,541,440 SUs/year

# Purchasing Compute Node

Proposed fee structure:

- Installation: \$1000 per node\*
- Yearly fees: \$1000 per U per node\*

These fees can be deducted from node scavenging revenue if available (yearly fee equivalent to 10% of node usage).

# Scavenging Income From Faculty Purchased Systems

- $S_{\text{CPU}} = 0$  nodes with 128 CPUs.
- $S_{\text{GPU}} = 1$  nodes with 4 A100 GPUs.
- $S_{\text{HM}} = 1$  high-memory nodes.
- $S_{\text{GH}} = 0$  Grace Hopper superchips.

## Renting:

- \$100/TB/Year parallel file system (compute side)

## Purchasing:

- t.b.d.

# Questions

?