

Research Computing Advisory Committee

Minutes Apr 14, 2025 (taken by Erik Deumens)

Present: Nikolay Bliznyuk, Chris Barnes, Philip Chang, Clint Collins, Gogce Crynen, Erik Deumens, Richard Hennig, Gail Keenan, Rafa Munoz-Carpena, Massoud Rouhizadeh, Kuttichantran Subramaniam

Updates

- HiPerGator 4th Generation refresh update

- Classes this Spring have been taught on the NVIDIA DGX Cloud that we arranged as part of the upgrade process with NVIDIA. That has gone without issues.
- A few researchers with appropriate workflows have also used the NVIDIA DGX Cloud. We will keep using the NVIDIA DGX Cloud until June 30.
- **April** 32 DGX B200 nodes have been handed over to UF at the end of March. UFIT RC staff are now installing RHEL 9.5 and running tests through the end of April
- **May** Users will be given access to the DGX B200 nodes starting May. This will be “early user” access. That means that work can be interrupted at short notice, unlike the typical 4 to 6 weeks’ notice during production status of HiPerGator. A small group of complex projects will receive priority support; this is because most likely fixing their problems will fix the problems of many in this early access stage.
- **May** The Lenovo CPU and L4 GPU nodes will be delivered as well as the new Blue storage system. The new nodes will also be in early user mode. The new storage will be tested, data will be copied. At some point the new Blue will become the production system.
- **June-July** The remaining 80 DGX A100 nodes will be returned to trigger delivery of the remaining 31 DGX B200 nodes. These nodes will then be wired, tested by NVIDIA, and tested by UFIT RC.
- **August** The whole DGX B200 system will run HPL and AI benchmarks to claim UF’s status on the performance scales.

Discussion

- The question was asked about the HPG prices.
 - The CPU core and storage process will stay the same
 - The process for the GPU cards will be set by the cost of the NVIDIA L4 GPU, which is \$2,800 for 5 years hardware life; this makes the price of the NVIDIA B200 extremely highly subsidized, just as the NVIDIA A100 was.
- The operating system of HiPerGator will be upgraded from RHEL 8.10 to RHEL 9.5 during the coming months, because that is what the new DGX B200 require, and we want HiPerGator to be homogenous.
- Orange storage will be upgraded in 2026.

Next meeting will be on June 2, 2025 from 2:00 – 2:30 pm. The meeting will be Zoom only.