

Hellbender consists of 112 compute nodes, and 20 GPU nodes interconnected by an NDR Infiniband backbone through HDR Infiniband node connections. The compute nodes are configured with 2 AMD Epyc Milan processing chips with 64 cores each, 512GB of DDR 4 RAM, and a 1.6 TB NVME drive for local scratch. 17 of the GPU nodes are configured with 2 Intel Xeon Gold processing chips with 32 cores each, 256GB of DDR4 RAM, a 1.6TB NVME drive for local scratch, and 4 Nvidia A100 GPU co-processors with 80GB of on-board RAM each. 2 of the GPU nodes are configured with 52 core Intel processors, 2TB of DDR5 RAM, a 3.2 TB NVME, and 4 Nvidia H100 GPUs with 80GB of on-board RAM each. 1 of the GPU nodes is configured with 56 core Intel processors, 2TB of DDR5 RAM, a 3.2 TB NVME, and 8 Nvidia H100 GPUs with 80GB of on-board RAM each. In total there are 15,744 cpus and 84 GPUs. The NDR Infiniband backbone provides up to four hundred 400Gb/s of data throughput from point to point on the network, with an anticipated theoretical latency of less than 600 nanoseconds. Each node is attached to the backbone with an HDR Infiniband connection capable of providing 200Gb/s of data throughput to each node, with an anticipated theoretical latency of less than 600 nanoseconds.

Directly attached to Hellbender through the Infiniband backbone is the cornerstone of our Research Data Ecosystem (RDE) storage solution. An Arcastream Pixstor array providing 4PB of raw storage through a tiered spinning disk and NVME solution capable of providing 160GB/s of data throughput to clients. Along with this direct connection the Pixstor solution will have SMB and NFS gateways to provide connections to desktop clients across the newly redesigned research network. Along with the Pixstor solution, we also have a 4PB all flash VAST array as part of the RDE storage solution.

Hellbender is supported by the Research Support Solutions (RSS) team who provide system administration and software support.