


Background

- Legacy PSInet datacenter, now GW's Enterprise Hall on the Virginia Science and Technology Campus
- Original design was for 2kW per rack
- Colonial One HPC System runs at a sustained 25kW per rack:
 - Dell C8000 series hardware
 - Dual-socket Xeon + Infiniband Linux Cluster
 - Over 200 nodes

Challenge

- Plenty of CRAH/CRAC capacity; however airflow challenges lead to hot spots
- Excessive heat in the HPC cluster causing significant node failures
- Set point for the CRAH/CRAC is 72°F on return air, 50°F on supply under the raised floor
- Initially, rack “chimneys” were deployed on the HPC to shift exhaust air to above the ceiling
- System overheated during LINPACK benchmark testing & chimney fans were added
- Chimney fans could not keep up with the airflow demands



Need

- Keep Colonial One cool while adding additional racks
- Avoid worsening airflow/cooling issues
- Improve energy efficiency by raising temperature set point in the room

Solution

- Remove existing chimney and fan
- Add Coolcentric 30kW High Density Rear Door Heat Exchangers (RDHx) to five of the racks, Cooling Distribution Unit, and 8-port manifold
- Doors were installed, CDU commissioned, and secondary cooling loop filled and tested within 4 hours

Caveats

- Airflow within the racks needs to be managed, true for any solution
- Air inlet temp still ~60°F due to facility operations' concern
- Increasing the air inlet set point would allow for even greater energy efficiency gains

“Colonial One HPC Cluster is cooler resulting in an increase in speed.”

- Cluster failures reduced by 99%
- Hot Spots Eliminated