



SBEUC

Simulation Based Engineering User Center

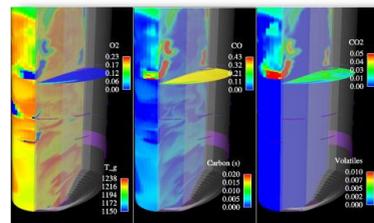
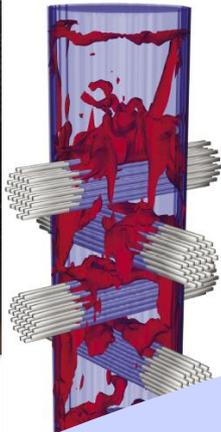
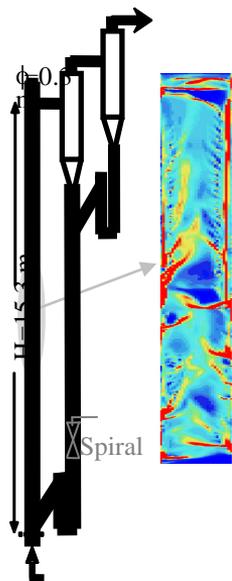
Chris Guenther,
Director, Computational Science Division

Agenda

- A brief look back
- What is SBEUC
- Modular Data Center
- HPCEE
Hardware/Software
- Visualization & Collaboration
- RUA Access



A Look Back at NETL's Computational Resources

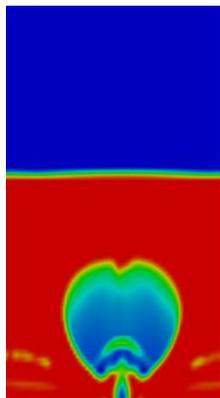


SBEUC

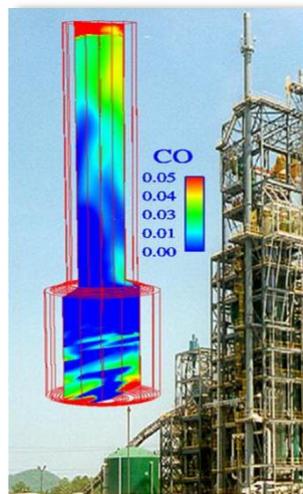
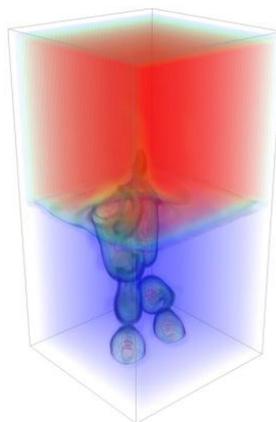


2000
NETL's First Computer Cluster
24 Cores
7.2 GFLOPS

1998
SGI-Octane
8 cores
300 MFLOPS

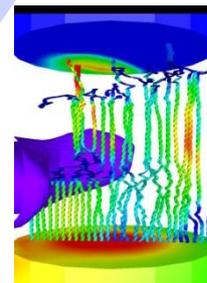


2002 Pittsburgh
Supercomputing Center
3000 cores, 6 TFLOPS



2005 Clusters at all
NETL sites
834 Cores
3 TFLOPS

2011
Clusters at all
NETL sites
7932 Cores
84.3 TFLOPS



The SBEUC Project

**ARRA Project 4000.600.007.001
Dept of Energy - Fossil Energy**

This project is funded by the American Recovery and Reinvestment Act (ARRA). The primary focus of is to provide a multi-laboratory collaborative tool to perform computations for the completion of the Carbon Capture and Storage Simulation Initiative (CCSSI) program.

\$20,000,000 ARRA funding

\$2,500,000 NETL funding



SBEUC Mission

- **Provide advanced high performance computing capabilities to accelerate progress in NETL programs to meet DOE's Fossil Energy mission**
- **The SBEUC was built specifically to facilitate collaborative computational research with an applied energy focus between researchers at all the three NETL sites and with the partner Regional University Alliance (RUA).**

Elements of the SBEUC

- **High Performance Computer for Energy and Environment (HPCEE)**
- **Modular Data Center (MDC) – The structure that houses, powers and cools the HPCEE**
- **Visualization Centers – located at all NETL sites and collaborative workstations at all the RUA member sites**
- **High speed connectivity in/out of the MDC**

Modular Data Center

- **NETL studies showed lower costs, quicker installation and higher energy efficiency from Modular Data Centers**
- **18 proposals from 11 vendors were received**
- **Focus was on performance (TFLOPS) and power usage effectiveness (PUE) for a fixed price.**
- **Bid went to Silicon Graphics International (SGI)**

NETL Morgantown, WV, July 2012



MDC Enclosure Arrives in 7 Shipments



***Prefabricated
power supply
for MDC being
lifted (65,000 lbs)
into place***



Power Delivery

350' of 4180 volt cable feeds a stepdown transformer.

The HPCEE has a 1 Megawatt power Requirement

5 times more efficient than current clusters (TFLOPS per watt)



**MDC Assembly
required only
2 weeks**



**Inside one of the
MDC hot aisle prior
to equipment
installation**



99 Pallets in 6 Truckloads



Installing IT Equipment

All servers were assembled and tested by SGI and Aeolus in Wisconsin

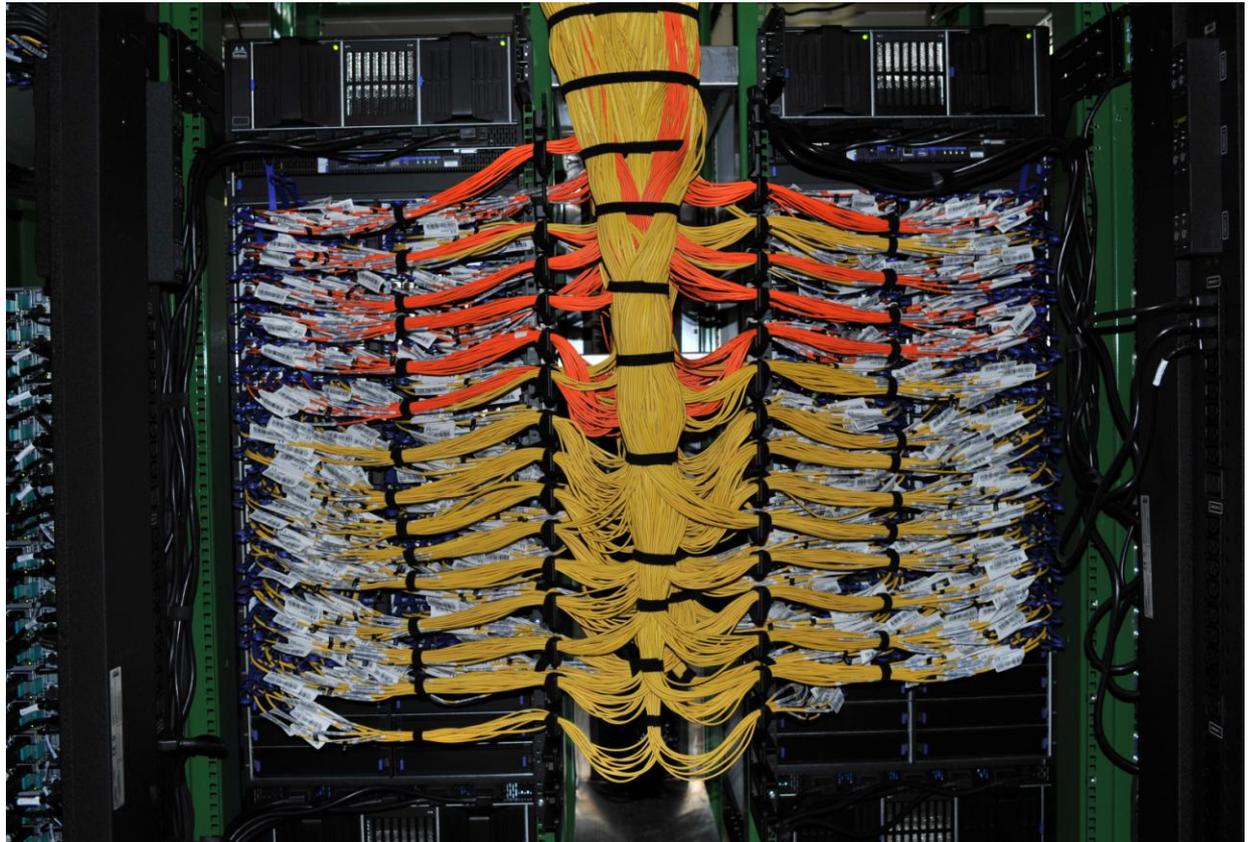
All servers and switches shipped to Morgantown for installation into Modular Data Center



Compute Node Hub

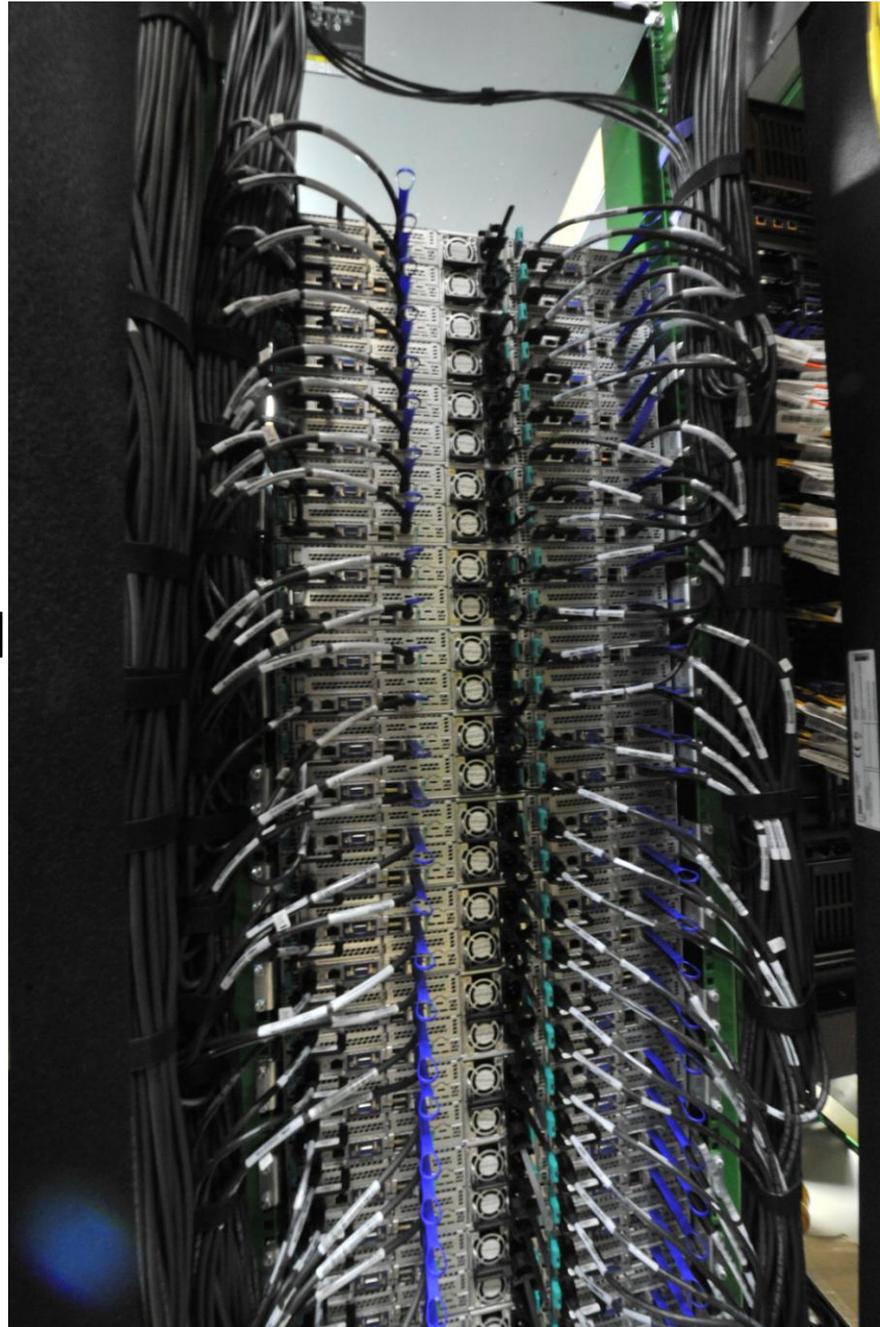
Fiber optic cable connections to a pair of Mellanox 648 port IB switches

The Trunk of the Fat Tree Topology



A Single Rack of Compute Nodes from the hot aisle

**Airflow is from the cold
aisle through the
computers to the
hot aisle and
exhausted**



Modular Data Center

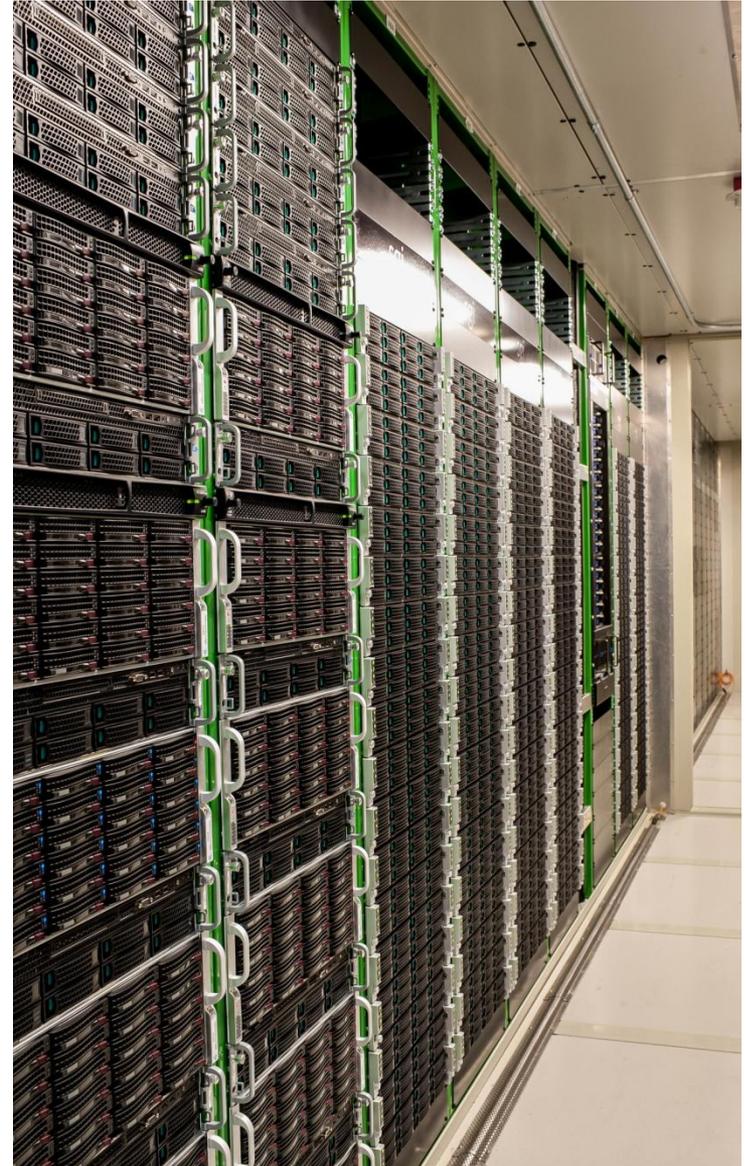


Modular Data Center

MDC

- **Cost effective: 50% lower cost based on traditional data centers**
- **The Modular Datacenter (MDC) represents the cutting-edge of efficient datacenter design with a PUE of 1.06**
 - Uses free-air cooling for most of the year
 - Supplemented with evaporative cooling for hot days
 - Advanced system management
 - Tracks energy utilization, air temperature and humidity.
 - Adjusts fan speeds, louver opening, and air recirculation

$$PUE = \frac{\text{Total Power Entering Data Center}}{\text{Power to Run Computer Infrastructure}}$$



MDC: Efficiency By Design

Two Measures of Computational Efficiency

- **MFLOPS/watt**

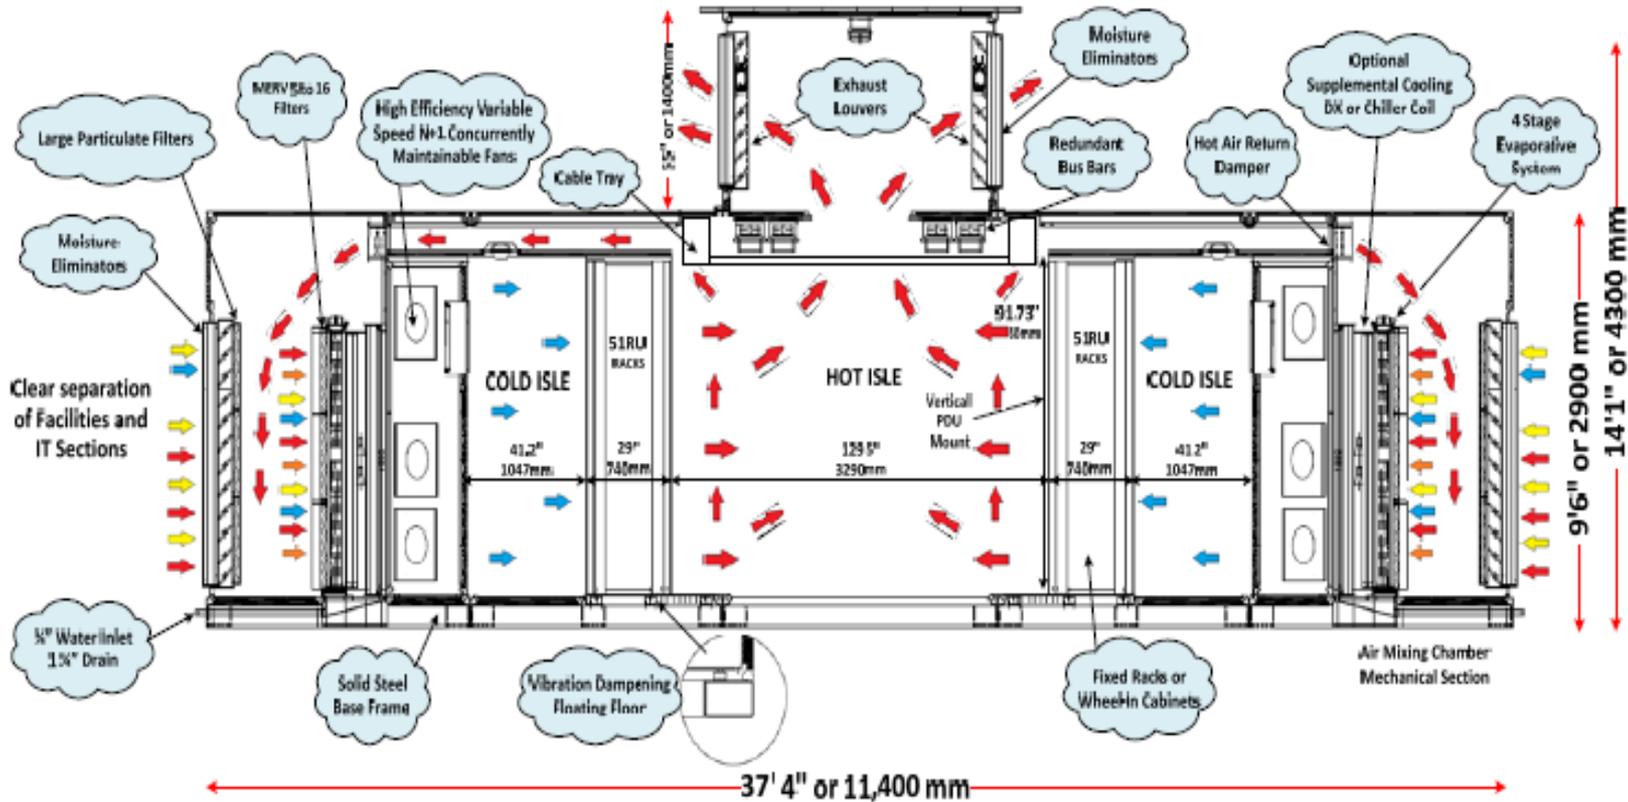
- Existing NETL cluster average = 250
- HPCEE = 1300 **5X IMPROVEMENT!**

- **Power Utilization Efficiency (PUE)**

Total Data Center Power / IT Power

- Current NETL aggregate = 1.6
- HPCEE = 1.06 **WORLD CLASS PERFORMANCE!**

HPCEE Cooling



- **Four Stage Evaporative Cooling**
- **No Mechanical Refrigeration**

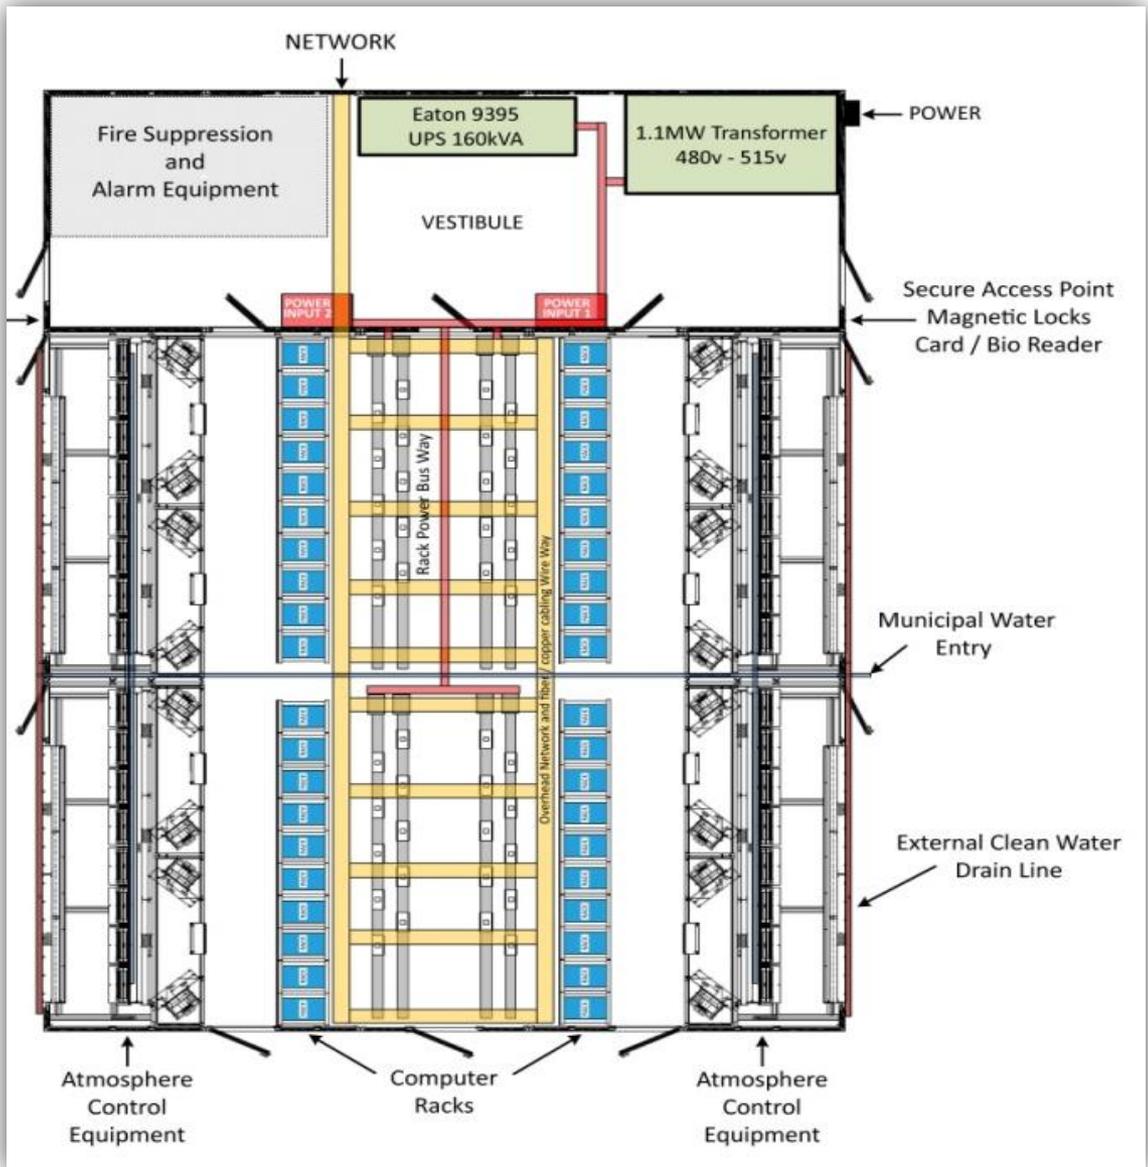
MDC Interior

Elements of the MDC

Vestibule – power, power back-up, fire suppression and alarm equipment

Mixing Chambers
Air Handling – filtration, humidification and fans

Cold Aisle
Computing Equipment
Hot Aisle



Compute Nodes

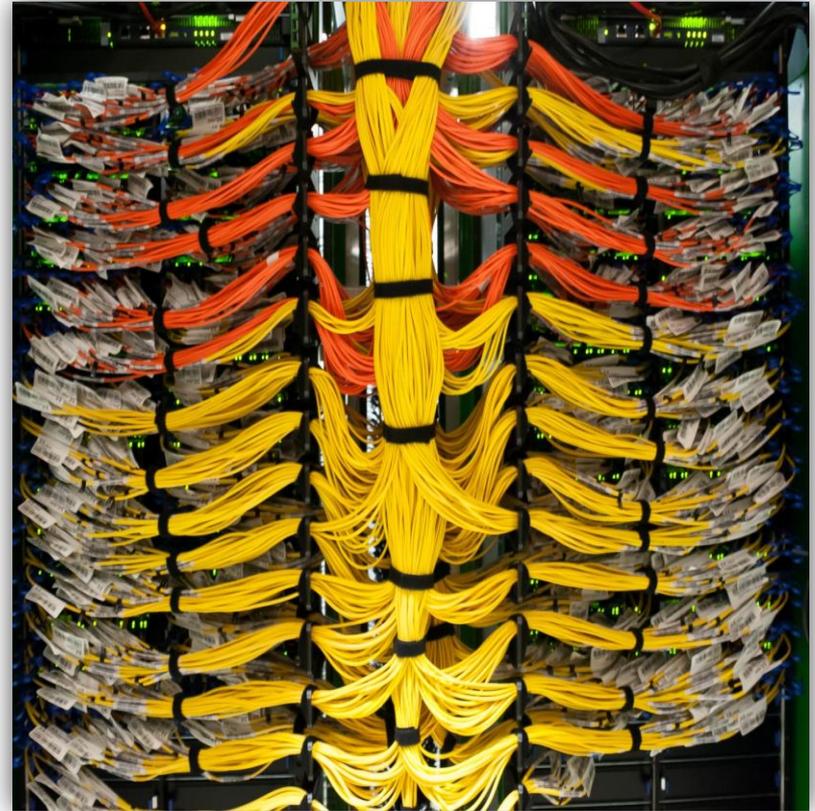
*378 chassis containing four nodes
(1512 total)*

*Each node has two 8-core 2.6 GHz Intel
Sandy Bridge CPUs (24,192 total)*



Connectivity

- Each node is equipped with a 40-Gbps QDR Infiniband network interface in an optimized network topology: a full-bisection bandwidth via Fat-Tree topology. Provides 40Gbps links among all 1512 nodes
- This results in a 82% efficiency



Storage

- **Total of 9 petabytes of disk storage**
- **1 petabyte of primary disk storage attached to the compute nodes by Infiniband**
- **Storage is mirrored by an identical 1 petabyte array in B39 at NETL Morgantown**

Software

- **The HPC EE systems run Linux as the sole operating system.**
- **OpenSUSE 11.4 is the current base distribution with specially compiled kernels to support parallel processing.**
- **Application software is supplied by the researcher.**
- **Currently, MFiX, ANSYS/FLUENT, Barracuda, VASP, OpenFOAM and other packages are being used on the system.**
- **Post-processing/Vis packages, Ensight, Paraview, VisIt**

Interacting with the HPCEE

Primary mechanism for interfacing with the HPCEE is remote login from user workstations.

- **Access via VirtualGL/TurboVNC software that provides a full 3D accelerated graphics interface remotely and makes user logins persistent, *providing the same active login in moving from workstation to workstation, to laptops or personal systems, or even to the large scale SBEUC special purpose displays.***

Visualization

- **SBEUC Visualization Centers are located at all three NETL sites (Morgantown, Pittsburgh Morgantown**
- **Tile displays, large screen LCD's and collaborative workstations**
- **High speed access to high memory visualization nodes for visualization and post-processing**

NETL Morgantown Visualization Center



Visualization Spec's

HPCEE has dedicated visualization hardware specifically designed to parse and render large data sets from high resolution/massively parallel simulations

- Six dedicated visualization servers, each with:
 - 12 2.9 GHz Intel Sandy Bridge cores,
 - 256 GB main memory
- Four Nvidia Tesla C2090 graphics cards
 - Total of 2048 GPU cores and 24 GB of GPU RAM
- plus a large dedicated local RAID6 scratch space and is directly connected to the one petabyte Infiniband storage network.

NETL Morgantown Collaborative Area

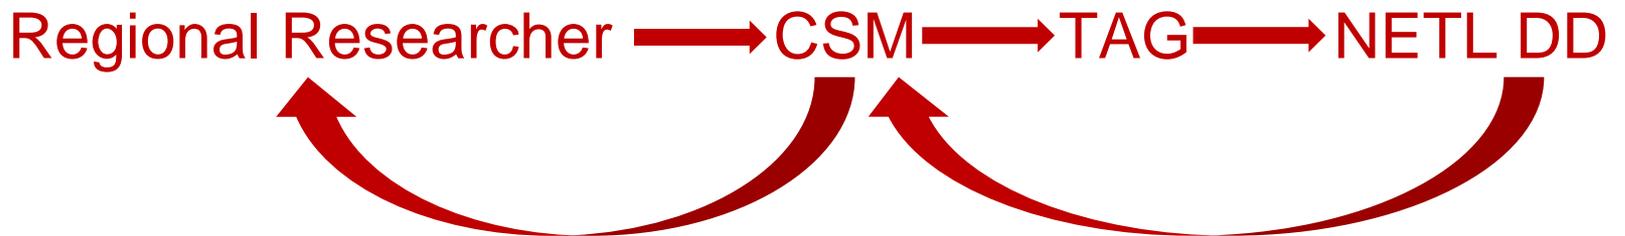


Where are we today

- Requirements Complete – December 2012
- Contract Signed -April 2012
- Testing Components - March 2012
- Delivery to NETL - July 2012
- Turnover to URS - October 2012
- HP LINPACK completed - November 2012 (**HPCEE Number 55 in the world!**)
- Turnover to DOE - January 2013
- User Migration - February 2013
- Current NETL/RUA users migrated April 2013
- Outdated NETL clusters shutdown – Spring 2013
- Resource available to RUA - Spring 2013

User Access

- Current NETL/RUA researchers are being migrated to the new system
- Spring 2013 regional researchers may apply for access to the HPCEE and supporting equipment.
- Access request form via NETL ORD website
- Access will be managed in accordance with operating procedures approved by the ORD Computational Systems Manager (CSM).



Thank You!

Questions can be directed to

Chris Guenther

Chris.Guenther@netl.doe.gov