# BUILDING ONE OF THE WORLD'S LARGEST WINDOWS-BASED HPC CLUSTERS

By Scot A. Schultz
Charles Rubin

**When the University of Nebraska at Omaha's Holland Computing Center teamed up with Dell to build a high-performance computing (HPC) cluster, the result was a flexible dual-boot system of Dell™ PowerEdge™ servers with dual-core AMD Opteron™ processors—and one of the world's largest and most powerful Microsoft® Windows® OS–based clusters.**

**Related Categories:**

AMD

AMD Opteron processors

Case study

Dell PowerEdge servers

High-performance computing (HPC)

Microsoft

Visit DELL.COM/PowerSolutions for the complete category index.

Over a year ago, Dell engineers set out to design and build one of the world's largest Microsoft Windows–based high-performance computing (HPC) clusters at the University of Nebraska at Omaha (UNO). Funded with a budget of US$20 million and located in a 2,000-square-foot glass-walled enclosure in the Holland Computing Center (HCC) at the Peter Kiewit Institute, the cluster represents a major milestone in establishing the stature of UNO as a leading U.S. research center. The HCC has already attracted cluster users from the U.S. Department of Defense, Gallup, Microsoft, and the Milliman actuarial firm, and provides a major IT resource for UNO faculty and students as well as research scientists from around the United States. The HCC cluster was ranked 43rd on the TOP500 list of the world's most powerful supercomputers in November 2007.[1]

## DESIGNING THE CLUSTER

The three-person HCC staff began by selecting Dell as the master project manager. As part of the design phase, the team looked at other major HPC centers in the United States and received input from local and national business organizations with previous supercomputing experience, and determined that the proposed cluster must be both collaborative and able to scale to meet several general requirements:

- It had to support both Microsoft Windows and Linux® operating systems to accommodate enterprise users, university and government researchers, and students.
- It had to be one of the most powerful clusters in the world.
- It had to fit inside a 2,000-square-foot facility in a region where outside temperatures can range from –20°F to over 100°F.
- It had to include an access and management system that could provision dedicated nodes for several simultaneous users with minimal staff intervention.
- It had to be able to run TOP500 benchmark tests.

American Power Conversion (APC), Cisco, Dell, Force10 Networks, Microsoft, and other vendors were eager to provide equipment and support for this ambitious project to help ensure that the installation went seamlessly. Dell specified 1,150 PowerEdge

[1] For more information on the TOP500 list, visit www.top500.org.

SC1435 1U servers as the compute nodes, with a PowerEdge 6950 server as the master node (see Figures 1 and 2). The cluster connects to two dedicated networks: a 10 Gigabit Ethernet management network that uses a Force10 TeraScale switching fabric, and a Cisco-based InfiniBand network to deliver high bandwidth and low latency for server-to-server interconnects.

Each compute node was initially configured with dual-core AMD Opteron 2220 SE processors at 2.8 GHz, 8 GB of RAM, and 80 GB of local storage. The HCC selected these AMD Opteron processors because of the server motherboards' plug compatibility with the new generation of quad-core AMD Opteron processors, which enables the HCC to upgrade to quad-core processors in the future simply by installing them in place of the dual-core processors and then updating the server BIOSs—without needing to change power and cooling requirements. While the dual-core processors could deliver 21.5 TFLOPS of performance, the HCC estimates that an upgrade to quad-core processors could provide up to a 320 percent performance increase, in addition to raising the cluster's teraflop rating to more than 60 TFLOPS.

***Figure 2.*** *The HCC cluster is based on 1,150 Dell PowerEdge SC1435 server nodes and a PowerEdge 6950 master node*

For data storage, the cluster links to a Panasas storage and high-speed file system with 150 TB of space over Fibre Channel links designed to provide throughput of 3.5 Gbps (see Figure 3). The system supports both a Microsoft Windows Compute Cluster Server 2003 environment and a Linux environment. The HCC can automatically provision OS images to specific nodes, and can dedicate specific nodes to specific users as needed.

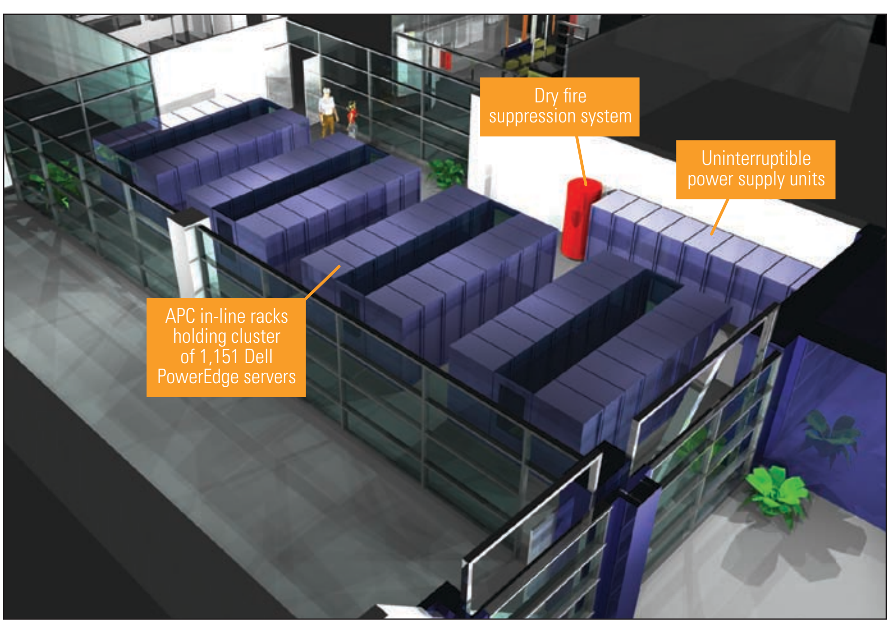


***Figure 1.*** *The powerful dual-boot HCC cluster is housed in a 2,000-square-foot glass-walled enclosure*

To help meet the physical demands of the installation in terms of footprint, space, and power and cooling, the team configured the cluster in four modular pods. The design uses a hot aisle/cold aisle configuration, where the pods hood the hot aisles and recycle the warm exhaust air through an in-row self-contained water-cooled system designed and provided by APC. The system uses two redundant chillers that can provide more than 200 tons of cooling. When outside temperatures drop below 20°F, a dry air cooler turns on and shuts down the chillers to help reduce energy use.

## MEETING DIVERSE DEMANDS

The HCC debuted on the TOP500 list in November 2007, placing within the top 50 of the world's largest supercomputing-class systems at number 43. The configuration and testing processes have been ongoing since the initial installation. One of the most important discoveries related to daily operations has been the demand for Microsoft Windows applications; because the HCC team's experience had been primarily with Linux-based clusters, they had not anticipated the level of demand for Windows applications that the cluster has received.

Initially, the HCC had planned a zero-touch automated deployment of the Windows Compute Cluster Server 2003 OS. However, the Remote Installation Services (RIS) tool used by this OS did not support zero-touch deployment on systems that had already been partitioned using Linux. To work around this limitation, the HCC employed a light-touch deployment, which consisted of an attended network service boot on each node that polled the RIS server for the preconfigured Windows image and allowed the side-by-side deployment.

Microsoft played an active role in the success of the cluster, which is one of the largest Windows-based clusters in the world. The company provided on-site assistance and training, helping the HCC staff overcome early challenges. The collaboration on a project of this large scale not only helped the HCC provide Windows users with comprehensive access to the facility and its available resources, but also enabled Microsoft to refine many of its own best practices and product features.

"The HCC and its collaborators plan to continue pushing the boundaries of HPC performance based on the Dell focus on best-of-breed standards and cost-effective components."

## PLANNING FOR THE FUTURE

As a result of the cluster's success, the HCC plans to move forward with further planning and development, and to increase its internal staffing to help accommodate the large demand from the business community as well as university and research requirements. In addition to the Department of Defense, Gallup, Microsoft, and Milliman, the cluster supports genetics researchers at the Henry Doorly Zoo, researchers studying protein strands at the University of Nebraska Medical Center, and both undergraduate and graduate students at UNO and the University of Nebraska at Lincoln.

The HCC, Dell, AMD, Microsoft, and a range of collaborative partners have continued working together since the cluster's original implementation. In December, the team performed an initial upgrade of 250 nodes to quad-core AMD Opteron 2356 processors at 2.3 GHz, anticipating an upgrade of the entire system in time for the next TOP500 test period. For the Dell consulting team, the experience continues to advance the state of computing performance while making that performance available to a worldwide community. The HCC and its collaborators plan to continue pushing the boundaries of HPC performance based on the Dell focus on best-of-breed standards and cost-effective components—helping expand the reach of HPC and deliver it to a growing set of users including enterprises, university and government researchers, and students.

*Figure 3.* Panasas storage provides the HCC cluster with over 150 TB of space over Fibre Channel links

**Scot A. Schultz** is a senior strategic alliances manager for HPC at AMD, where his responsibilities include understanding computational and operational requirements from commercial vertical markets as diverse as bioinformatics, manufacturing, and oil and gas, as well as managing partnerships with key technology partners to create solutions based on those requirements. He also works with numerous software and hardware partners to optimize their technologies, software, and drivers to deliver superior performance on AMD-based architectures, and serves in various industry organizations including the HyperTransport Technology Consortium and the OpenFabrics Alliance.

**Charles Rubin** is a freelance technology writer focused on network systems. He is the author of 15 books about computer hardware and software, and has B.A. and M.A. degrees in English from San Francisco State University.

MORE ONLINE
DELL.COM/PowerSolutions

**QUICK LINK**

**Dell HPC cluster solutions:**
DELL.COM/HPCC