;login:
THE MAGAZINE OF USENIX & SAGE
August 2001 • Volume 26 • Number 5

Special Focus Issue:Clustering

Guest Editor: Joseph L. Kaiser

inside:

USENIX & SAGE
The Advanced Computing Systems Association &
The System Administrators Guild

special purpose clusters at the Cornell Theory Center

by Dave Lifka

Chief Technology Officer, Cornell Theory Center. Lifka has played a major role in the formation of the Advanced Cluster Computing Consortium, AC3, formed in conjunction with Dell, Microsoft, and Intel. CTC is focused on using Windows 2000 clusters for production quality computing.

lifka@tc.cornell.edu

Why CTC Moved to Windows Clusters

In August 1999, the Cornell Theory Center (CTC) moved from traditional UNIX HPC to Windows-based computing with the installation of a 256-processor Dell PowerEdge Windows cluster. CTC's Velocity cluster was the first production system of its kind within an academic environment. It was our goal to demonstrate that mid-range to high-end cluster computing systems with industry-standard components can provide the scalable, reliable resources needed for today's technical and business computing environment. We hoped to leverage exceptional economies of scale by integrating software and systems from desktop to HPC system. The CTC Velocity Complex now comprises an aggregate of 250 nodes, including several special purpose clusters.

Benefits of the Move

As a result of the move, CTC has experienced a very large increase in the number of users, especially new users attracted by the Windows HPC environment. We have also seen a dramatic shift in our cost-performance ratio, which is now at least one-fifth that for our prior traditional systems, especially considering the reduced costs of management and maintenance. This reduction in operating cost allows us to stay at the leading edge of the performance curve, which is critically important when the members of your user community need to maintain their competitive edge. At the same time, we have been able to provide extremely reliable resources to our users so that they can get their work done. (All of the nodes in the Velocity Complex are running at between 99.99% and 99.999% uptime.) In fact, our success in attracting new users helped drive the implementation of several special purpose clusters.

Special Purpose Clusters

A number of our users quickly realized that they were competing for a popular resource and decided to seek funding so that they could afford to purchase their own cluster systems. CTC builds and maintains these systems for the research groups in collaboration with their staffs, ensuring that the groups will benefit from our experience and be kept up-to-date on technical advances. At the same time, they have access to a custom environment that is tuned to their specific needs. Examples of special purpose clusters now implemented at CTC include:

Cornell Institute for Social and Economic Research

The Cornell Institute for Social and Economic Research (CISER) purchased a Dell cluster system that is installed and maintained at CTC. Cornell's social scientists interactively access a wide variety of statistical software packages for analyzing their research data or to gain access to CISER's nationally renowned data repository. They transitioned their users from a UNIX system to this Windows cluster in one year and now have more than 300 individuals on the system.

Computational Materials Institute

The Computational Materials Institute (CMI) at Cornell University is involved in a wide variety of research initiatives ranging from earthquake prediction to fracture simulations. CMI users were among the first to move to the initial CTC cluster, Velocity, and quickly found that they needed exclusive access to what is a shared resource. Their first strategy was to join with a few other large groups to drive the acquisition of a second production cluster, Velocity+. However, CMI researchers run 64-processor jobs 24x7; they needed a dedicated system to meet their demand. CMI recently received NSF CISE/RI funding for a dedicated 32-node, 64-processor cluster.

USDA-ARS Center for Agricultural Bioinformatics

The research emphasis of the USDA-ARS Center for Agricultural Bioinformatics at Cornell is on the exploitation of high-performance computing and communication technologies to solve practical problems in agricultural genomics. Initially, they required a cluster system that provided rapid turnaround for large-scale similarity searches on DNA sequences from agronomically important plant species. The output of these searches was used to populate a large data warehouse. This application required short queueing delays; the ability to handle very long jobs; fast access to large, shared file systems; and tight integration with the database server. Since this type of use is not well suited to a large, shared resource, USDA-ARS decided to fund the purchase of a dedicated 48-processor cluster that is managed and maintained at CTC. This cluster is now being used for the analysis of entire genomes of bacterial plant pathogens.

Conclusion

These research groups are finding that they get just what they need out of this working relationship – dedicated, up-to-date custom systems that are well supported and easily accessible. As a result, the CTC cluster complex now includes eight systems, each unique, some general purpose, some specially designed to meet the needs of a special user community. We learn more about clustering design and implementation with each project.

For more information:

CTC: *http://www.tc.cornell.edu*

CTC's Advanced Cluster Computing Consortium: *http://www.tc.cornell.edu/AC3/*

Windows HPC: *http://www.microsoft.com/WINDOWS2000/hpc*

Computational Clustering Technology Preview: *https://microsoft.order-1.com/cctp/*