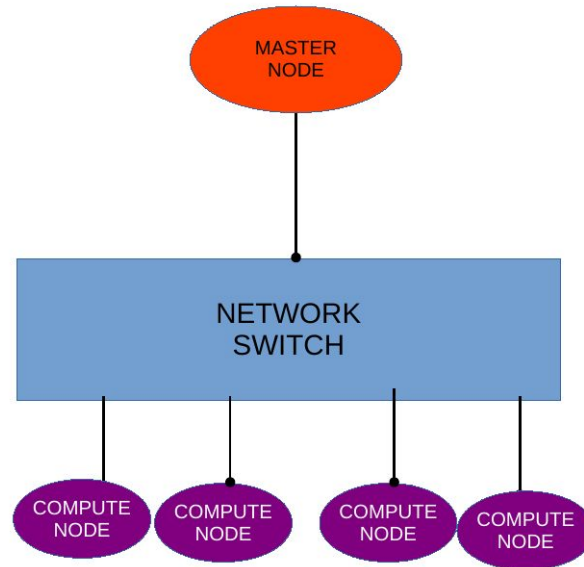# Beowulf Cluster for Research and Education

2018-08-24

—

Stefano C. Coronado & Brandon Rozek

University of Mary Washington Physics Department

Fredericksburg, VA

Overview

A Beowulf Cluster is a group of computers, individually known as nodes, that communicate with one other. One computer serves as the center of this cluster as a master node. The master node sends instructions to the others over a Local Area Network (LAN) to complete a given task. The hardware used were surplused Dell OptiPlex 990 models generously donated by the University of Mary Washington IT Department.

Stated Objectives

1. Build a Beowulf Cluster
2. Allow the use of multiple programming languages with their respective shells or compilers installed
3. Enable students and faculty to add programs to a run schedule
4. Create resources made to help students and faculty write or use programs for the cluster
5. Regulate power usage and cooling
6. Develop a modular cluster design, allowing for more nodes or computers to be added to the network

Achieved Objectives

1. Beowulf Cluster has been built
2. Multiple programming languages are installed
3. Modular cluster design was followed

Outcome

As a team, we have built a working cluster that can distribute the load of MPI scripts across the group of computers before revealing the output to the master computer. We are now looking at adding the extra graphics processors in our possession to expand opportunities of adding more applications to the cluster.

Future Direction

While some of our objectives were not fully achieved within the past semester, our team will continue to work closer towards achieving the following:

1. Allowing for students and faculty to be able to schedule execution of programs
   ○ We are working towards making the user interface cleaner to target users that are not experienced in the GNU/Linux operating system.
2. Resources to assist faculty and students to write programs.
   ○ Mostly done in the form of documentation compiled by the project.
   ○ Looking at other Linux vendors to use as a base for easier maintenance of the cluster.
3. Reduction of power consumption.
   ○ Experimentation operating system system calls to achieve this goal.

Summary of Major Expenditures

I. Data Storage

With last year's grant money, we have purchased many 16 GB Solid State Disks for each of our compute nodes to hold persistent data on the cluster such as, the deployed operating system that was installed to the SSD over PXE boot, similar to the Microsoft Deployment Toolkit used by IT Departments to image Windows PCs.

II.   Memory

A fast cluster needs modern memory modules rather than the older Dell OEM memory modules included with the donated Optiplexes'. We have decided to purchase Kingston HyperX DDR3 RAM modules to keep compatible with the motherboards that were donated to the cluster.

III.   Network Switch

The HP Enterprise Network Switch is used as a Local Area Network (LAN) with DHCP disabled, allowing the master computer to act as a DHCP server to deploy extra computing nodes to the array. This network switch assists in reducing network latency among the master and compute nodes.

IV.   USB Network Adapter

The USB Network Controller was purchased to follow through with what we have considered an experimental network design, where the master computer needed a second network adapter to communicate with the cluster, while the built-in controller could be used to physically connect users to our network switch.

V.   APC UPS Backup Battery

The backup battery is crucial for the cluster to continue operation in the event of a power outage. APC branded UPS Backup Batteries are used in server and telecom closets to allow for a graceful shutdown, and also work as a massive surge protector.