Cluster Construction to Advance Analytics

Sawyer Bergeron, Joseph Carrara, Griffin Leclerc, Scott Valcourt

Computer Science Department, University of New Hampshire, Durham, NH 03824

Introduction

The high performance computation community has developed numerous ways to construct computer clusters. A variety of options are available to construct computational clusters, yet some community members are unaware of emerging technologies and continue to use conventional methods. This project's aim was to design and construct a low cost, high performance computing (HPC) cluster for use by University of New Hampshire researchers looking to complete experiments involving intense computational components. These experiments can include simulations of city transit systems, cross-model parameter estimation in epidemiology, or simulation of atmospheric conditions, among others. Additionally, our cluster construction documentation will provides the HPC community with an alternative perspective on cluster construction and implementation.

Methodology

The project cluster utilizes six rack mounted servers, or nodes, reclaimed by the project's sponsor. These nodes were provisioned with a modern Linux operating system and networked together to form the project's computational cluster. Slurm, a versatile job scheduling technology, has been installed on all of the nodes. HPC6 acts as the task distribution node, which determines how resources are managed and which worker nodes receive tasks pertaining to any given computation, or job. A sample job was written, in which the first 100,000 digits of Pi were calculated, assembled and returned to the user. This job was used to measure and compare the completion times of an average personal computer, our cluster, and another cluster made available to us by the Extreme Science and Engineering Discovery Environment (XSEDE).

Data

Using the sample job described in the methodology section, the performance of our cluster was compared with other systems. We performed this comparison with our cluster in two separate configurations. The configuration titled "6 Node Project Cluster" utilized all six of the cluster's nodes to complete the sample job (HPC1-HPC6). The second configuration, "4 Node Project Cluster", executed the sample job across the cluster's three more computationally advanced nodes (HPC1-HPC3) and compiled the results on a fourth, less robust node (HPC6). Finally, we configured our job to run over 4 nodes of the XSEDE cluster.
All machines and configurations were able to complete the sample job.

Cluster Implementation

Physical Cluster

Full-scale project cluster view

Technologies

XSEDE
Extreme Science and Engineering Discovery Environment

slurm
workload manager

Graphics depicting organizations whose technologies were utilized during this project

Cluster Workflow

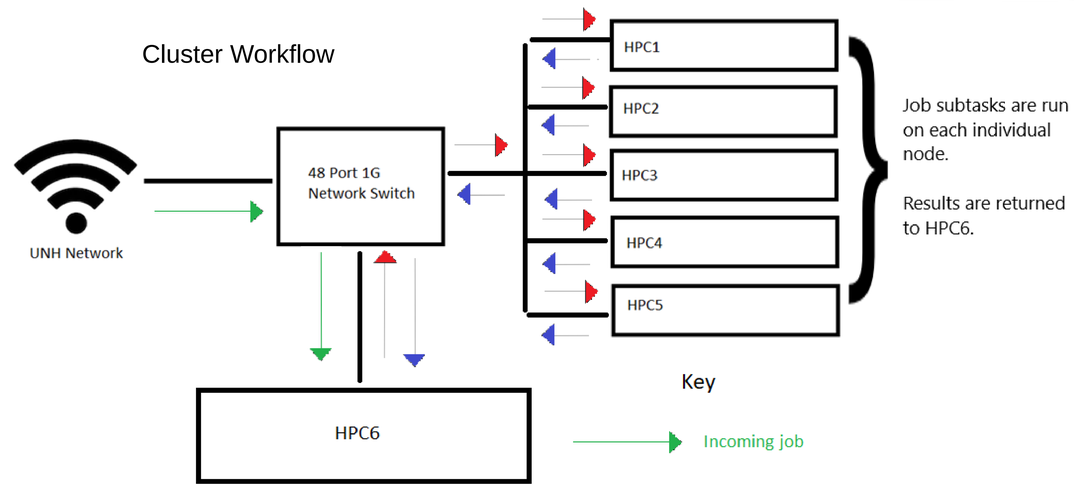

A job is sumbitted to HPC6.

That job is divided into subtasks that are transmitted to nodes.

Results are gathered and stored on HPC6 for later retreval.

Benchmarking

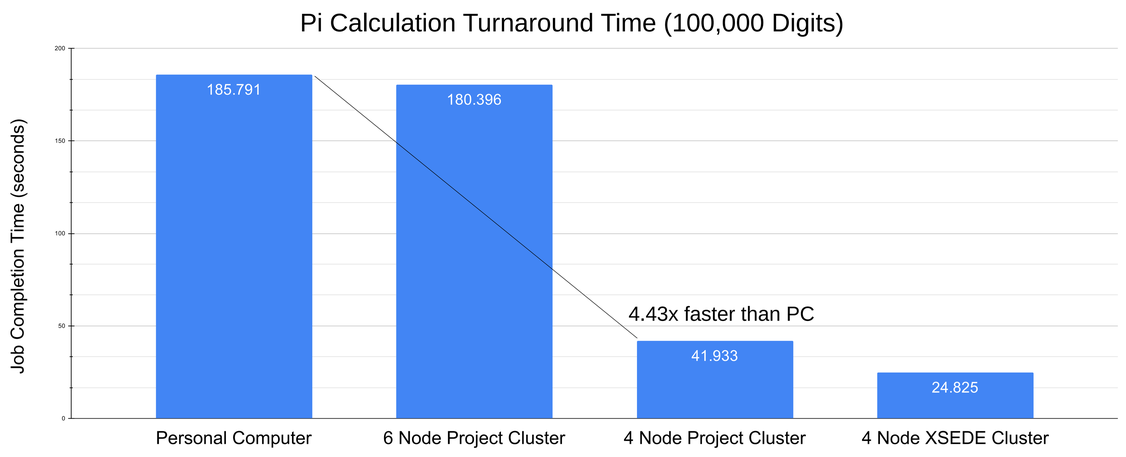

Results

See the Benchmarking section for performance results. A fully functional cluster has been constructed with all nodes networked as described in the Cluster Workflow diagram. All nodes are able to communicate with one another and with other networked devices. The cluster computation job schedule software, Slurm, has been installed and configured to distribute load evenly across all of the computationally advanced worker nodes (HPC1-HPC3). This cluster is currently available to UNH students and faculty for use in computationally intensive research. Documentation detailing the cluster implementation has been provided to the project sponsor for the benefit of external users interested in this technology and submitted to the Technical Program Committee for evaluation in advanced research computing at this year's PEARC21 conference.

Conclusions

The project demonstrates it is possible to construct high performance computer clusters, using reclaimed technology, with respectable performance at minimal cost. The project's 4 node cluster configuration completed the sample job 4.43 times faster than the personal computer and only 1.7 times slower than the XSEDE cluster.
In addition, the project cluster completed the sample job four times faster when configured to use fewer nodes. This is due to the fact that some of our machines completed computations at a slower rate due to less robust hardware components.
These findings were counterintuitive as our initial hypothesis was additional nodes would decrease job turnaround time, regardless of hardware capabilities.

Acknowledgments

Project Sponsor: Professor Scott Valcourt
s.valcourt@northeastern.edu

Project Advisor: Professor Matthew Plumlee
matthew.plumlee@unh.edu

This work is partially funded by the National Science Foundation under grant OAC-1659377

References

Cluster Construction Documentation:
https://docs.google.com/document/d/11NsXENrCb9wPGNAvC20MZJTGAEeyAcQIFLU0yfDZDxI/edit?usp=sharing

Slurm Documentation:
https://slurm.schedmd.com/documentation.html