HA-OSCAR

Alyssa Leangsuksun
The University of Texas at Austin

Research Alliance in Math and Science
Computer Science and Mathematics Division
Mentors: Stephen L. Scott, Ph.D., John Mugler, Thomas Naughton

http://www.csm.ornl.gov/Internships/abstracts/AlyssaLeangsuksun.pdf

Introduction

High Performance Computing (HPC) is a significant enabler in the process of Scientific Discovery. Today, cluster computers are the low-cost computing workhorses in such environments; however, node failures can greatly impede progress. The acknowledgement of this problem has led to the development of High-Availability Open Source Cluster Application Resources (HA-OSCAR) whose goal is to strive towards eliminating single points of failure, ultimately resulting in reduced downtime in computer clusters.

Research Methods

The first method in researching high-availability (HA) in a cluster computing environment is the construction of a cluster from commodity workstation-class computers. This will be accomplished using the OSCAR and HA-OSCAR packages. An evaluation of high-availability Linux will follow, as to consider how the HA-OSCAR environment may leverage its features.

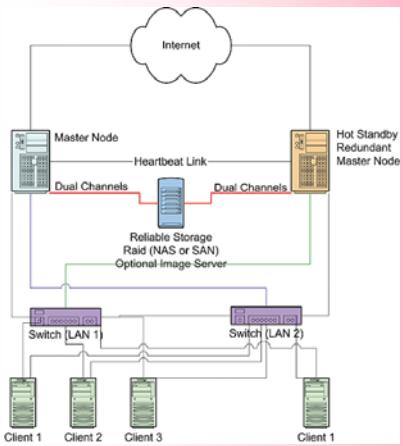

Project Objectives

- Gain understanding of cluster and high-availability computing techniques
- Construct HA cluster
- Evaluate LinuxHA
- Migrate promising techniques to HA OSCAR

Results

	LinuxHA	HA OSCAR
HPC Capability	Can only support two servers	Two head nodes, multiple compute nodes
Heartbeat Mechanism	Yes	Yes
Failover	Automatic	Does self-checking first
Failure Prediction	N/A	Proof of Concept- Trend Analysis
Planned Downtime	N/A	Work in progress

Conclusion

At the commencement of this project, it was expected that HA OSCAR would be able to leverage promising techniques from LinuxHA; however, HA OSCAR's features not only match, but surpass those of its predecessor. It is able to support two head nodes and multiple compute nodes -- an essential component in a project requiring ultra-scale computations. In addition, HA OSCAR does a self-checking before performing a blind failover, first attempting to resolve problems through checkpoints. Progressive and future works include failure prediction and planned downtime, two preemptive features which will be vital in the significant reduction of downtime in a cluster computing environment.

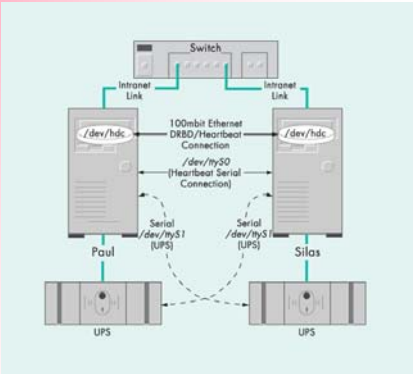

The Research Alliance in Math and Science program is sponsored by the Mathematical, Information, and Computational Sciences Division, Office of Advanced Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725. This work has been authored by a contractor of the U.S. Government, accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

This research used resources of the Center for Computational Sciences at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

The author would like to thank her mentor Dr. Stephen Scott and his research associates John Mugler and Thomas Naughton for their tireless efforts in facilitating the project. She would also like to express her gratitude to Debbie McCoy & the Office of Science for providing such an amazing and beneficial opportunity. Finally, she would like to acknowledge Cheryl Hamby, Cynthia Latham, Lora Wolfe and the rest of the supportive staff at ORNL for their assistance to the Research Alliance in Math and Science program.

OAK RIDGE NATIONAL LABORATORY
U.S. DEPARTMENT OF ENERGY