

Tutorial A: Design and Analysis of High Performance Clusters

Rob Pennington, David A. Bader, Barney Maccabe, Patricia Kovatch, and Stephen L. Scott

Abstract

The National Computational Science Alliance (the Alliance) has created several production superclusters for scientists and researchers to run a variety of parallel applications. The goal of this tutorial is to bring together and share the latest information on the state of high-end commodity clusters. The Alliance and others are active in developing and benchmarking applications on both Linux and NT clusters. Presenters are all active in the development and management of these systems.

The goal of these clusters is to provide easy-to-use high performance computing systems at reasonable prices. Superclusters are large scale clusters built from commodity parts and high performance interconnects.

We will discuss the details on the design, implementation and management of these systems. System monitoring and management tools are becoming more sophisticated and will be demonstrated. A wide variety of applications and community codes run on these superclusters. We will examine several of these applications and include details on porting and application development tools on both NT and Linux. We will also discuss how to use tools to tune the system and applications for optimal performance. Within the Alliance, the Virtual Machine Room (VMR) provides a sophisticated interface to the national computational grid infrastructure. These Alliance clusters are linked together through the VMR, allowing easy and uniform access to these resources. Using this interface, scientists and researchers can submit and monitor their jobs through a common interface from geographically distributed locations.

Many research and development efforts are enhancing superclusters. Several of these trends such as IA-64 and parallel file systems will be discussed.