Competing in the world of science

Roy Dragseth, Senior Engineer, The HPC Project, University of Tromsø, discusses the acceleration of scientific progress through clusters...

On 10th November 2004, Norway's first teraflop computing facility was installed at the University of Tromsø. Due to increased activity among the research groups in Tromsø, there was an imminent need for more computing power. The requests for computational resources for the fourth quarter of 2004 showed a 500% increase in demand for the existing 22 processor Itanium 2 cluster.

In response, the University decided to acquire a new cluster that would be able to fulfil these demands and vendors on the regular campus-wide contract for IT equipment were invited to deliver tenders. Based on a performance evaluation of the main applications, the decision was made to acquire an Itanium 2-based cluster with 196 processors and a peak performance of one teraflop.

The idea behind clusters for high performance computing (HPC) is not new and is quite simple: instead of buying large, specialised supercomputers at enormous costs, a large number of ordinary servers are connected in such a way that they act as one large computer. In general, this appears as a very powerful computer at a fraction of the price of a supercomputer. The idea is very appealing, but there have been significant costs related to system administration, so when the first HPC cluster was installed in Tromsø three years ago, it was done with a strong focus on how to reduce these costs to the level of the supercomputer that was in use, and the search for available management tools started. The Rocks Cluster Distribution (Rocks) developed at San Diego Supercomputer Center emerged as the answer, as it provides a turnkey solution by taking care of all the tedious tasks related to the installation of the compute nodes, and the management of the users, software and batch job scheduler needed to make the cluster work. This allows the staff to spend their time on providing good user support for students and scientists with advanced computational needs. Since 2003, the HPC group in Tromsø has participated in the development of Rocks, and is responsible for the development and maintenance of one of the batch queuing systems supported.

The existing cluster had proved itself to be very robust, so the aim was not only to install the new compute nodes without taking down the existing cluster, but also to show the feasibility of clusters by broadcasting the installation live on the internet. The installation went very smoothly and the new nodes were ready for production less than four hours after arriving at the machine room, giving a 10-fold increase in the theoretical compute power from 100 gigaflops to over one teraflop.

In the power Grid analogy that is often used to describe the emerging Grid technology, the cluster will become a power station in the upcoming Nordic Grid collaboration, where scientists in the Nordic countries can benefit from the computing power it provides in a seamless manner.

The impact on the scientific progress has been much larger than anticipated and the response from the users has been overwhelmingly positive. Tasks that previously took weeks can now be completed in less than a day. Simulations can now be done that were previously not possible to execute due to few resources/long wait queues.

As a support function for scientific production, it is pleasant to see that the availability of a large computing facility has increased the competitiveness of the University of Tromsø in the world of science.

UNIVERSITY OF TROMSØ

Roy Dragseth
Senior Engineer, The HPC Project
royd@cc.uit.no
http://uit.no/itavd/HPC/