

Managing a Large UNIX Cluster

Customer

JET / UKAEA

Introduction

The JET Facility in Oxfordshire is home to the world's largest fusion experiment. Roughly thirty times a day, heavy isotopes of hydrogen are raised to 200 million °C in order to emulate the thermonuclear reactions at the centre of the sun.

The ultimate goal for this research is to provide an energy source that is plentiful, safe and inexpensive. However, there are many technical difficulties to be overcome before this Holy Grail can be reached. Over 50,000 experiments have been run at the JET Facility, which have produced over 10 Terabytes of data. This data must be processed, analysed, and compared with theoretical predictions, in order for the scientists to progress.

The major computing platform for this analysis is a cluster of UNIX nodes referred to as JAC. This cluster enables the scientists to analyse and process the data with IDL, Matlab, "C" and FORTRAN programs. Many of these codes are extraordinarily CPU intensive, often running for days and sometimes weeks.

A Tessella consultant has been the Systems Manager of the JAC cluster since 1996. At that time, the cluster consisted of fifteen workstations, and was managed by a staff of two System Administrators. The cluster has since expanded significantly to its present configuration of over 90 nodes running the Linux, AIX and Solaris flavours of UNIX. It is now used by over 200 active users who can logon locally from office PCs, or remotely from various laboratories around the world.

Another improvement enables the cluster to run as a "Beowulf" supercomputer. This allows programs, tailored to use the Message Passing Interface (MPI) and Parallel Virtual Machines (PVM) systems, to be executed in parallel on a large number of different nodes. Such an approach dramatically reduces the run-time of the codes, and is particularly beneficial to Monte-Carlo simulation codes which model millions of independent particles in motion, as different subsets of the particles can be analysed on different nodes.

The challenge of running a large cluster

The management of such a cluster presents major challenges. Running three distinct operating systems poses problems in presenting a homogenous environment for the users, and also makes the system more difficult to administer. The large number of nodes increases the possibility of hardware and software problems. Finally, the technical support given to the large number of users also creates a significant workload.

Perhaps the biggest challenge is to successfully manage all of these problems with the minimum of downtime, and without an increase in the number of administrators.

Configurations

The individual nodes must be configured so that the users have a consistent interface, and so that the cluster can be managed efficiently. It is also important that the cluster is configured in a scaleable manner, so that new nodes can be incorporated at a later time. To this end, most common system information is shared across all the nodes using the Network Information Services (NIS). The file systems containing common applications and user files are mounted across all nodes from a server. This gives each node equal access to all data, without having to replicate files.

Load Balancing

The ninety nodes can deliver a considerable amount of computer power, but only if the

workload is spread evenly over the cluster. If a large number of users log on to a single node, their response times will deteriorate until the node runs out of some critical resource.

The work done by the cluster is evenly spread by balancing the number of users logged on interactively to each node, and providing a facility whereby a 'batch job' can be run on a node with freer resources.

Problem Management

UNIX systems can suffer from a variety of problems. File systems can fill up, memory can deplete, processes can hang, and remote disks can be inaccessible. Any one of these or other problems, if left unattended, could render a node useless, or impact the whole cluster. Most of these problems, if detected early, could easily be resolved. However, early detection is not always easy because of the large number of nodes. It was therefore found necessary to automate the monitoring of potential problems, and in some cases, automating their correction.

The first step in detecting problems is to measure the "vital life signs" of each of the nodes. This includes recording the processor activity, memory usage, and network and disk activity. Only after the general behaviour of the cluster is understood, is it possible to detect the early signs of problems.

A suite of programs was developed to measure these life signs. They have to deal with the different interfaces in the three operating systems, but store their results in a common format. These results are processed by report programs, and also by programs which search for unusual system behaviour, such as high system loads or the processor spending excessive times in the "wait" state. Reports of any unusual activity are emailed to the administrators. If the problem is sufficiently severe, the system will page the administrators, informing them of the need to take urgent, corrective action.

Occasional problems still occur, and have a significant impact on the users. All problems of this type are formally recorded and reviewed. Their likelihood of reoccurring is assessed, and if necessary, changes are made to the existing procedures.

The most time-consuming aspect of the management of cluster is the provision of technical support to the users. This involves answering a variety of questions, ranging from simple UNIX queries to more complex problems involving programming, data access and networking. Particular efforts have been made to reduce the time spent in this area without sacrificing the quality of the support. These include the writing of manuals, web pages and Frequently Asked Questions (FAQs).

Conclusion

By using a range of techniques, the 90 workstations in the UNIX cluster have been able to provide an excellent high availability service to the growing user population using only two administrators. This, combined with the use of low cost Intel systems and public domain Linux based software, has led to significant cost savings whilst maintaining enviable service levels.