Cluster Computing: A High-Performance Contender

Mark Baker, University of Portsmouth
Rajkumar Buyya, Monash University
Dan Hyde, Buknell University

Simplified and edited by Bilal Muddassir, University of Engineering and Technology for the Khwarzimic Science Society

When you first heard people speak of Piles of PCs, the first thing that came to mind may have been a cluttered computer room with processors, monitors, and snarls of cables all around. Collections of computers have undoubtedly become more sophisticated than in the early days of shared drives and modem connections. No matter what you call them - Clusters of Workstations (COW), Network of Workstations (NOW), Workstation Clusters (WCs), Clusters of PCs (CoPs) - clusters of computers are now filling the processing niche once occupied by more powerful stand-alone machines - minicomputers and single-box supercomputers.

In its simplest form, the computers in your office or your home that are connected to your local area network constitute a workstation cluster. In addition to the hardware, a workstation cluster also includes the middleware that allows the computers to act as a distributed or parallel system and the applications designed to run on it.

While a system based on low-end workstations and network technologies may not at first seem particularly useful, such systems have been the testbeds or a new computing paradigm: *high-performance and high-availability cluster computing*. This class of system is becoming increasingly commonplace; in fact most academic institutions and industries that use high-performance computing either already use or are thinking of using workstation clusters to run their most demanding applications. Even companies that can afford traditional supercomputers are becoming interested in commodity clusters.

Why the switch? For some, cluster-based systems provide a way to stretch their computing dollars, allowing the reuse of seemingly obsolete office or classroom systems. Others have found that a cluster of high-performance workstations can easily compete with the best supercomputers IBM or SGI have to offer. A company can download a few tools from a public Web site and order a collection of machines and network equipment to put together an 8-Gigaflops systems for around $50,000. Assembling a more powerful supercomputer would cost around $200,000.

TECHNOLOGIES, COMPONENTS AND APPLICATIONS

A cluster consists of all the components on any LAN with PCs or workstations: individual computers with their processors, memory, and disks; network interface cards; cabling; libraries; operating systems; middleware; tools and various other utilities. The architecture of clusters, however, can vary rather dramatically. At one end of the spectrum are the clusters based on commercial off-the-shelf components or put together from older systems, maybe originally used in offices or classrooms or homes of all our relatives. (For information on one of the very first COTS clusters, see the Beowulf project homepage, http://www.beowulf.org/). At the other end are proprietary clusters built around high-end Symmetric Multiprocessing (SMP) systems and custom network technologies. The physical configurations of clusters also vary widely; including anything from a bunch of PCs located in a campus junkyard to motherboards stored in custom racks in some computer science institute.

Of course, all these components are there to support

applications of one sort or another. Applications appropriate for clusters are not restricted to those traditionally run on supercomputers or other high-performance systems. The number and types of applications now using clusters are increasing all the time. Cluster-based systems support both high-performance parallel applications such as *computational chemistry, astrophysics, atomic structure calculations, finite analysis, biomedical image interpretation and computational fluid dynamics,* and commercials such as a *load-balanced high-performance Web server like HotBot* (http://www.hotbot.com), which uses a parallel Oracle database.

SUPERCLUSTER SYSTEMS AND ISSUES

The National Computing Science Alliance (http://www.ncsa.uiuc.edu/alliance/alliance/), an American academic partnership involving more than 50 US universities and research institutions, has built a supercomputer consisting of 192 nodes based on dual-processor Pentiums running Windows NT with Myrinet interconnects. A number of well-known scientific codes have been ported onto the NT supersluster. Its capability has been proven: Several applications attain a sustained performance of around 7 Gigaflops. This compares favorably to the performance reached on an SGI O2000.

However, as is the case with many professional environments, the differing technologies used in cluster computing can spark debate. For example, many disagree whether a cluster should be Ethernet - Fast or Gigabit - technology or specialized nonstandard intelligent network cards and protocols like those produced by Myrinet (http://www.myri.com) or SCI (http://www.scali.com). Experts in cluster computing circles, like those in other fields, also dispute the best operating system to use: commercial products such as Microsoft's Windows NT or freely available systems such as Linux.

The arguments over operating systems for cluster-based applications, like most arguments, are a combination of objective and subjective reasoning. To settle the OS debate more impartially, research teams working for NCSA have pitted a Linux-based supercluster against one running Windows NT. Housed at the University of New Mexico's Albuquerque High Performance Computing Center, the Linux cluster (http://www.arc.unm.edu/alta/) will run a range of computationally intensive applications to be compared to the proven power of the NCSA's NT cluster (http://www.ncsa.uiuc.edu/General/CC/ntcluster).

THE NEED FOR A NEW TASK FORCE

Recognizing the trend toward clusters for high-performance computing, the IEEE Computer Society has approved a Task Force on Cluster Computing (TFCC). You may ask, "What's so special about that? People have been using clusters of computers for years. It's a bit like Microsoft realizing that the Internet maybe a big thing!" *But think again*. The overwhelming number of cluster related projects and products appearing in the development arena and the commercial marketplace means that a focused group can help lead the international effort in cluster-based computing.

Why not place cluster-based activities under the umbrella of another related Technical Committee, such as the committees on the Internet, Supercomputing Applications, or Distributed or Parallel Processing? After all, what is cluster computing but a mixture of these disciplines? But cluster computing combines so many computing concepts and technologies that placing it under an existing banner would dilute the attentions of individuals interested in all the aspects that come together in this field.

With the advent of the TFCC, interested Computer Society members can participate in one focused group to champion the cause of cluster computing by sponsoring workshops, conferences, projects, and standards.

The TFCC has set up two mirrored Web sites: one in Australia (http://www.dgs.monash.edu.au/~rajkumar/tfcc/), and another in the UK (http://www.dcs.port.ac.uk/~mab/tfcc), as media for timely communication of its activities.

EDUCATIONAL RESOURCES

Although many educational institutions teach undergraduate and graduate students about the hardware and software that make up a cluster, few courses or programs concentrate on the wealth of technologies that constitute the complete cluster environment, from hardware to application development tools. In order to introduce cluster computing into the curricula of more university programmes , the Task Force on Cluster Computing has set up a Web site (http://www.coe.uncc.edu/~abw/parallel/links.html). This informative resource provides links to related journals, books, freely available software, projects from both academia and industry, white papers, and descriptions of hardware components. In addition, with our education donation program, we actively support academic faculty members around the world who are interested in introducing new cluster-based courses by providing sample curricular materials.

With the generous cooperation of leading publishers worldwide, we have arranged for the donation of some current books on cluster computing. While the books will be available for faculty members who request them, the TFCC has reserved 50 percent of donation to academic programmes in developing countries.

The titles available include

- *High Performance Cluster Computing: Architectures and Systems,* R. Buyya (ed.), Prentice Hall, 1999

- *High Performance Cluster Computing: Programming and Applications,* R. Buyya (ed.), Prentice Hall, 1999

- *In Search of Clusters,* 2nd ed., G.F. Pfiser, Prentice Hall, 1998

- *Metacomputing: Future Generation Computing Systems,* W. Gentzsch (ed.), Elsevier, 1999

- *Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers,* B. Wilkinson and C.M. Allen, Prentice Hall, 1998

CLUSTER COMPUTING AT THE KHWARZIMIC SCIENCE SOCIETY

At the *Khwarzimic Science Society*, a student group lead by *Mr. Jawad Mahmood* has taken up the study and implementation of cluster computing as the subject of their final year BE project in communication systems. At the moment, they plan to construct a low-end cluster using off-the-shelf Intel hardware and available LAN equipment. They can be of immense help if anyone finds interest in computer clusters. Mr. Jawad Mahmood can be contacted via The *Khwarzimic Science Society*. Details follow.

This document is intended to be a general introduction to the art and science of Beowulf clusters and the treasure of potential therein; other details about the *Khwarzimic Science Society's* project is delineated in some other document. Details are obtainable upon request.

EXTREME LINUX

Mr. Bilal Muddassir can be contacted via the *Khwarzimic Science Society* if any Member wants to taste the Extreme Linux CD. Extreme Linux (http://www.extremelinux.org) is a term coined to distinguish the Linux-based software layer that controls the cluster hardware. The CD has all the tools necessary for bringing a Linux-based cluster up and developing/running parallel applications on it. Linux is an operating system that was conceived by Mr. Linus Torvalds at the University of Helsinki as his study project around 1992. Due to a devoted effort of a large number of enthusiasts, code-freaks and professionals, It now stands as a robust alternative for networked environments. It is shipped with the GNU (http://www.gnu.org) General Public License that allows free use and distribution with the mentioned conditions.

Linux has proven to be an essential component in building clusters of PCs (pioneered by Beowulf), and its popularity is increasing in the world of scientific and high-performance computing. With a modular design and free source code that has been ported to several CPUs, the Linux kernel is also ideally suited for computer science research. Several companies have introduced Linux products to support powerful desktops and high-performance computing clusters. Dozens of universities and laboratories are using Linux for scientific computation and research. There are two Linux-based local area networks in our department available to the undergraduate and graduate students. Many companies are now beginning to market pre-configured Linux clusters using the latest Intel or DEC Alpha CPU. The marketplace is evolving. The Extreme Linux is there to help.

Bilal Muddassir
noorh@brain.net.pk
http://www.khwarzimic.org

Khwarzimic Science Society
Centre of Excellence in Solid State Physics, University of the Punjab, Lahore, Pakistan.

http://khwarzimic.org/beowulf
supercompute@khwarzimic.org