Effective Aggregation of Idle Computing Resources for Cluster Computing

by Bruno Richard and Philippe Augerat

Modern research and industry projects require extensive computational power, which is conveniently provided by PC clusters. However, typical use of these clusters is irregular and shows some usage peaks. Researchers at Icatis are developing the ComputeMode™ software, which takes advantage of idle PCs on the corporate network and aggregates them to the cluster during usage peaks in order to reduce its load.

Computing clusters can be found in lots of companies and institutions today. Researchers and engineers have high data-processing needs and use them to distribute large jobs to a set of homogeneous machines. The Linux operating system is a de facto standard for cluster management, providing easy administration, good performance and a wide software support basis. Moreover, the set of available libraries and tools makes Linux a good choice for scientific applications.

Past projects have been focusing on how to aggregate user workstations from the enterprise network to the clusters. In this way, the company can take advantage of the availability of the processing power of user PCs. However, the real world shows that most corporate users are running Microsoft Windows™, making it difficult to aggregate user machines to the corporate cluster, which is based on Linux. Other approaches such as SETI@home or XtremWeb use a centralized distribution of computing tasks to Internet machines, but do not offer the smoothness and ease of use of a Linux cluster.

Icatis is developing the ComputeMode™ software suite, which smoothly handles this specific issue and aggregates user machines to the corporate cluster. A server is installed on the customer premises and keeps track of user PCs running Windows. During cluster usage peaks, a number of idle user machines can be aggregated to the cluster. This is done through a transparent switch of the PC to a secondary, protected mode from which it proceeds into a network boot from the ComputeMode™ server, taking advantage of the PXE protocol. This patented technology provides several benefits, such as the full isolation of the PCs' hard disks: these are not accessible while the PCs are dedicated to cluster computing. The OS and system configuration of a computing PC are also the same as a PC from the cluster, hence providing homogeneity and easing administration.

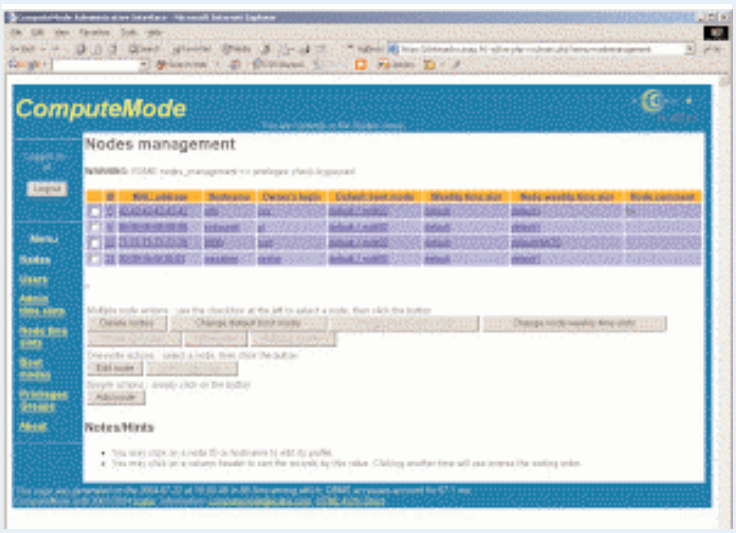

A ComputeMode screen shot.

The system is designed to be very transparent to PC owners, and the machines are only used at times when the PC is likely to be idle (nights, weekends and during business trips or vacation). If the user returns unexpectedly while his/her PC is doing a computation, the user can claim the PC back, and it restores to the state in which the user left it within one minute. This includes the user session, open files, on-going programs and the desktop.

On the administration side, ComputeMode™ offers a Web administration interface to register/unregister machines to the system, handle specific system image parameters and the usage schedules for the machines (this can be done automatically), and check usage logs. The Job Management System (JMS) administration for the cluster shows additional machines in the computing pool, and priorities can be adjusted using the standard JMS configuration tools.

Users of the cluster extended through ComputeMode™ do not see much difference when ComputeMode™ is installed. Job management is done in the standard way through the JMS. The only noticeable difference is the boost in reactivity that can be expected when the cluster is heavily loaded. In such cases the PCs that ComputeMode™ aggregates to the cluster provide some extra computational power and processing occurs faster for the user.

Icatis is a young company, having been created in January 2004 after several years of investigation and refinement of its offer. It was successful on the commercial side: a contract has already been signed with a major oil and gas company and in June 2004 Icatis was elected a laureate in the 'Innovation-

Development' national contest and won a prize from the French Agency for Innovation (http://www.anvar.fr/).

Most Icatis researchers have been working within the ID-IMAG Laboratory (http://www-id.imag.fr/), the Apache project run by INRIA, CNRS, IMAG, and UJF. ID is a French public research laboratory, which for the past twenty years has been researching concepts, algorithms and tools for high-performance, parallel and distributed computing. Successful experiments include the development of a supercomputer from standard hardware components such as those that might be found in a typical large company. An unusual supercomputer built from 225 standard PCs (733 MHz, 256 MB RAM, Ethernet 100) entered the TOP500 contest (http://www.top500.org/) in May 2001 and was ranked 385th worldwide for supercomputing. Other successful experiments such as I-Cluster and NFSp, as well as the Ka-Tools developed in close partnership with Mandrakesoft (http://www.mandrakesoft.com/) for the MandrakeClustering product (CLIC project from ID-IMAG), have built some sound technical bases for Icatis.

Icatis benefits from a strong experience in cluster computing and the Linux System, and has strong links with the high-performance computing community.

Some Icatis customers have already evaluated a ComputeMode™ prototype on their own premises. It has shown good results for peak usage absorption on an application for seismic terrain exploration, with each job running several hundreds or thousands of tasks. The full product will be released from QA in December 2004. Among other features, future versions of ComputeMode will have wider grid capabilities, such as inter-Cluster load balancing, and multiple administration roles/domains. At the same time, Icatis is working on a high-end visualization cluster.

Link:
http://www.icatis.com/

Please contact:
Philippe Augerat, Icatis CEO,
Grenoble, France
Tel: +33 6 76 27 27 92
E-mail: philippe.augerat@icatis.com

The European Learning Grid Infrastructure Integrated Project

by Pierluigi Ritrovato and Matteo Gaeta

A semantic Grid for human learning is the vision behind the European ELeGI Integrated Project for the implementation of future learning scenarios based on ubiquitous, collaborative, experiential-based and contextualized learning through the design, implementation and validation of the Learning Grid.

In recent years, teaching and learning practices have been mainly based on the information transfer paradigm. This focuses on content, and on the teacher as the key authoritative figure who provides information. Teachers' efforts have generally been devoted to finding the best ways of presenting content in order to transmit information to learners. Unfortunately the current generation of 'e-Learning solutions' has adopted the rather narrow pedagogic paradigm of 'information transfer'. This occurred simply because it is an easy way to use the Web's basic facilities. Failures, such as massive drop-out rates, are usually explained by a lack of staff awareness in the use of the Web, rather than critical reflection on the limits of this approach.

So, the question remains – how do we provide better access while maintaining or improving quality of learning through the use of ICT? The aim of ELeGI (European Learning Grid Infrastructure), an EU-funded Integrated Project (23 partners from nine EU countries), is to promote effective human learning by demonstrating and supporting a learning paradigm shift from the current idea of information transfer to one that focuses on the learner and on the construction of knowledge using experiential and collaborative learning approaches in a contextualized, personalised and ubiquitous way.

In our vision, the learner has an active and central role in the learning process. Learning activities are aimed at aiding the construction of knowledge and skills in the learner, rather than the memorisation of information. In keeping the learner at the centre of the learning process, personalisation (exploiting the existing learner's capability and skills) and individualisation (creating and adapting learning paths according to learner's preferences) become relevant aspects to be supported by technologies through the creation of the right context.

The Learning Grid

This new vision has two strong implications with respect to technology: first, teaching and learning will move toward the form of services, and second, the creation of virtual communities that provide services (eg contents, tutoring, searching for learners for sharing experiences etc) will be encouraged. Technology must be selected according to these implications. ELeGI will adopt a service-oriented model, which is deeply intertwined with the use of semantic tagging and aligns itself with the global community, helping to develop trends such as OGSA (Open Grid Services