



Features

 [PRINT THIS PAGE NOW](#)

University of Utah: C-SAFE Uses Linux HPCC in Fire Research

2/3/2003

By Stephen Parker

A high performance computing cluster (HPCC) is an increasingly popular method for delivering computational power to CPU-hungry scientific applications. A cluster consists of several commodity personal computer systems (desktop or rackmount) connected with a commodity or special-purpose network. The low cost of these systems makes an attractive option for a wide range of scientific applications.

In 1997, the University of Utah created an alliance with the Department of Energy (D'E) Accelerated Simulation and Computing Program (ASCP), formerly ASCI, to form the Center for the Simulation of Accidental Fires and Explosions (C-SAFE, www.csafe.utah.edu). C-SAFE focuses specifically on providing state-of-the-art, science-based tools for the numerical simulation of accidental fires and explosions, especially within the context of the handling and storage of highly flammable materials. The primary objective of C-SAFE is to provide a software system comprising a Problem Solving Environment in which fundamental chemistry and engineering physics are fully coupled with non-linear solvers, optimization, computational steering, visualization, and experimental data verification. The availability of simulations using this system will help to better evaluate the risks and safety issues associated with fires and explosions. Our goal is to integrate and deliver a system that is validated and documented for practical application to accidents involving both hydrocarbon and energetic materials. Efforts of this nature require expertise from a wide variety of academic disciplines.

Additional Resources

Simulation of an explosive in a fire involves computing thermodynamics, fluid mechanics,

chemical reactions, and other physical processes at each of millions to billions of points in the simulation domain. These scenarios require significant computing resources to perform accurately. Simulations performed to date employed hundreds to thousands of processors for days at a time to simulate portions of this problem. Fortunately, we were also granted time on the D'E ASCP computers at Los Alamos, Lawrence Livermore and Sandia National Laboratories. These machines are some of the fastest in the world, and we received considerable allocations of CPU time as a part of the C-SAFE grant. However, sometimes these resources are not enough.

Recently, the University of Utah purchased a Linux cluster that will be used to augment the resources provided to us through the D'E. The cluster will be used for development and debugging of the simulation, parameter studies, and small- to medium-scale simulations. Large-scale computations will still require use of the larger D'E ASCP computing resources. Unlike the shared D'E resources, the cluster will be largely devoted to C-SAFE research 365 days a year.

This cluster consists of 128 Dell PowerEdge 2550 servers, each containing two Pentium 4 Xeon processors. The entire cluster contains 256 Gigabytes of main memory and over nine Terabytes of disk space. Each server (node) runs the Linux operating system, and is networked with a Gigabit Ethernet (currently being installed). This is a class of system often referred to as a Beowulf cluster, named after one of the first systems to demonstrate this concept. The cluster is listed on the November 2002 Top 500 list (www.top500.org) as the 89th fastest computer in the world.

Linux Advantage

C-SAFE chose to implement a Linux cluster for one primary reason: We could not afford a traditional supercomputer of the class and scale of those purchased by the D'E for the ASCP program. The cluster solution is based on commodity PC technology; each node is very similar to a personal computer. This leverages the mass-production of personal computers to provide a price/performance ratio unattainable by most other technologies.

Other advantages of this approach include leveraging a growing familiarity with the Linux operating system, and of course the fact that it is considered an "in vogue" thing to do in Computer Science circles. We also intend to utilize the cluster for real-time interaction with these long-running simulations, something that would be difficult to do on a larger shared resource.

However, this power does not come without caveat. First, the application must be able to perform well in this environment. A Linux cluster has communication latencies (both hardware and software) that are quite large relative to a traditional supercomputer. The C-SAFE software employs a novel approach to parallelism that helps mitigate the effects of these latencies, and we performed preliminary benchmarks that gave us confidence that we could achieve good performance on a machine of this size. Second, the installation and setup of a cluster can be quite complex and time-consuming.

Fortunately, Dell provided the physical installation of our cluster, which saved considerable time. Installation of operating systems on each of the 128 nodes, as well as installation and

setup of additional HPC software, consumed considerable additional time.

Numerous companies, both large and small, offer computational clusters of this type. The University of Utah chose Dell to implement our cluster because of a competitive price/performance ratio, company reputation, and features of the PowerEdge 2550. Dell provided physical installation and setup of the cluster, which immensely reduced the work required to deploy this machine. The cluster fills seven standard racks, stands about six feet tall, four feet deep and 14 feet long. It produces a considerable amount of heat and fan noise.

Lessons Learned

Those that are considering a large compute cluster should keep in mind several lessons that we have learned along the way (either from others or from the school of hard knocks).

First, start small; we leveraged expertise from building smaller clusters before we considered a project of this size. Second, do not underestimate the time involved in physical installation and software setup of a machine of this nature. Depending on the complexity of the required environment, this can take weeks to months. Third, planning is critical. Considerations must be taken for the power consumption, heat production and physical footprint of the machine.

Finally, the biggest lesson is that it can actually work; we are performing simulations that required computers that just five years ago cost nearly a factor of 100 more. This gives us hope that our largest simulations may be achievable on a relatively common platform in the not-too-distant future. Early access to this magnitude of resources from the D'E was critical to the development of these simulations. Once the hurdles are cleared, a high-performance cluster can provide a dramatic quantity of compute cycles for complex scientific computations such as those performed for C-SAFE research.

For information contact Steven Parker, Research Assistant Professor, University of Utah, at sparker@cs.utah.edu.

About the author: *Steven Parker is a Research Assistant Professor at the University of Utah.*

Proposals for articles and tips for news stories, as well as questions and comments about this publication, should be submitted to David Nagel, executive editor, at dnagel@1105media.com.