

Combustion Lab Discovers Cost Effective Way To Turn Clusters into More Research

Harnessing Departmental Resources in Cluster Technology Increases Compute Power, Saves Money



Products:

- » Altus 1000E Opteron™ server
- » Altus 3200 Opteron™ server
- » Scyld Beowulf® Series 29 Linux clustering software

Key Benefits:

- » Cost-effective Compute-Intensive Solution
- » Eliminated Reliance on Other Facilities
- » Easy to Set Up, Minimal Maintenance Required

Advice from Sandia:

1. Do your homework as to how your specific applications perform on the clusters being evaluated.
2. Reliability is key. Cluster technology isn't perfect and depending on which products you chose, you may need supplemental cluster management expertise in-house.
3. Buy in logical increments. Anticipate what you'll need in the future especially if you work with uncertain budgets.

“The minimal intervention required to manage a Scyld cluster significantly enhanced our productivity but in a cost effective manner.”

– Joe Oefelein, Senior Member, Technical Staff,
Sandia National Laboratory's Combustion Research Facility

Summary:

Even though Sandia National Laboratory's Combustion Research Facility (CRF) was doing science for the Department of Energy (DOE) with real world pocket-book impact, they were often limited by a lack of supercomputer resources required to conduct numerical simulations and analyze data.

Knowing they had a tight budget and needed massive compute power, Sandia's CRF team and Penguin Computing put their heads together and harnessed the power of several Penguin Altus Opteron™ servers using Scyld Software's Beowulf® cluster software.

This single departmental cluster gave them over five million compute hours per year and a five fold increase in performance. The result was a dramatic increase in the amount of research they could complete. It also saved CRF \$150,000 a year in administration overhead. Being brilliant scientists, they were smart enough to invest that “found money” in more compute nodes so they can crank out even more research!

The Challenge:

The DOE funds Sandia to find high efficiency, low emission solutions to complex combustion problems. As part of this mission, the CRF conducts complex simulations and analysis of turbulent reacting flows to study interactions between fluid dynamics and combustion chemistry that affect the performance and emissions of combustion devices.

Simulations done by CRF are so complex that a typical baseline case takes one to two weeks. More sophisticated jobs need up to eight weeks of compute time. Nothing in-house was powerful enough to do that sort of work so CRF had to rely on off-site supercomputer grants to complete their research.

But because demand for supercomputing time always exceeds supply, the CRF had to compete for resources at the handful of supercomputing centers nationwide against other government facilities, academic researchers, and other applicants. In the end, CRF only got a small fraction of the total compute hours per year needed to perform the required calculations.

The The Solution:

Two members of CRF's technical staff, Senior Member Joe Oefelein and Distinguished Member Jackie Chen, realized high performance technology had evolved to the point where CRF might be able to cost-effectively create a departmental scale Linux cluster that would allow them to run many calculations. Supercomputer time could be reserved for larger simulations that require substantial system support.

“We purchased hardware from Penguin Computing [because] their system engineering appeared to be the most robust,” Oefelein said. “We chose Scyld Beowulf because it is easier to use [than other options] and offered us a turn-key solution. The Scyld interface emulates a workstation.”

Within a day of arrival, the new Penguin/Scyld system was in place. It turned out that the biggest implementation challenge to CRF was the time the staff needed to get the computer lab prepared and the power in place. Since then, the cluster has been in continuous operation.

Combustion Lab Discovers Cost Effective Way To Turn Clusters into More Research

Harnessing Departmental Resources in Cluster Technology Increases Compute Power, Saves Money

The Solution (cont.):

Their new departmental cluster gave them over five million compute hours per year and a five time increase in performance. The result was a dramatic increase in the amount of research they could complete.

“It’s been clear as we’ve emerged from the shake-down phase that [the new server and cluster is] performing as promised. The time I need to manage the cluster is truly minimal.”

- said Oefelein

The Scyld Beowulf clustering software chosen by CRF also offers dramatically simple ease of deployment and manageability. This means three CRF Principal Investigators can access the Scyld Beowulf cluster from their workstations just by using an informal queue to manage shared use of the cluster. A single process ID space for the entire cluster on the master node also means that the cluster seems like one computer so CRF

can even scale up incremental without redesigning or administrative effort. The ease of use of the cluster also saved CRF \$150,000 a year in administration overhead.

“Cluster technology provides an affordable system with fairly significant capability dedicated to our problems,” concluded Oefelein. “We expect to buy more clusters. Other groups at Sandia also recognize this and are currently evaluating Linux clusters.”

The Installation

Penguin Computing created a 72 node Linux cluster with 144 processors for CRF. The Altus master node, back up battery system, GigaBit Ethernet switch and Infiniband are on one rack. Two other racks house Altus compute nodes with motherboards and communication hardware.

CRF invested in Infiniband because they knew that they would likely add another 72 nodes later if the strong performance they expected materialized. The experiment results were even better than expected so CRF is adding more nodes.



About Sandia National Laboratory Combustion Research Facility

The CRF is an internationally recognized Department of Energy Office of Science user facility. The CRF is home to about 100 scientists, engineers, and technologists who conduct basic and applied research aimed at improving our nation’s ability to use and control combustion processes. The need for a thorough and basic understanding of combustion and combustion-related processes lies at the heart of CRF research.

The CRF is an Office of Science user facility for broad-based research in energy science and technology. Using the facility’s unique laser diagnostic capabilities, staff researchers and visiting investigators explore fundamental chemical reactivity and dynamics problems, as well as conduct applied studies that support industry’s needs in areas such as engines and materials processing.

Sandia National Laboratory

» <http://www.ca.sandia.gov/CRF>

About Penguin Computing and Scyld Software

Penguin Computing is the leading innovator of highly scalable, powerful Linux cluster, server and workstation platforms based on open standards hardware and software. The company has focused 100% on Linux since its inception and as a result, its solutions are robust and reliable, with a price/performance that’s hard to beat. Scyld Software, a subsidiary of Penguin Computing, provides the industry leading Linux clustering software, Scyld Beowulf™, whose superior ease of use dramatically simplifies the deployment and management of Linux clusters. The company’s extensive customer base includes Fortune 1000 companies, government agencies and educational institutions. Founded in 1998, Penguin Computing is headquartered in San Francisco, California. For more information, please visit:

Penguin Computing

» <http://www.penguincomputing.com>

Scyld Software

» <http://www.scyld.com>