



POWERWULF CLUSTERS: PHYSICS / ASTROPHYSICS

CASE STUDY

MASSIVELY PARALLEL COMPUTING FOR ENERGY & EARTHQUAKE RESEARCH.

"The AMD Opteron™ processor performance was clearly outstanding, so it was a simple decision from that point of view."

-SUE GRIESBACH, IT MANAGER,
PAULSSON GEOPHYSICAL SERVICES, INC.

PROFILE

Paulsson Geophysical Services, Inc. (P/GSI) Compute Cluster for 3D Seismic Analysis and Mapping

CHALLENGE

Processing large, multi-gigabyte data files to map underground conditions for oil industry exploration and earthquake monitoring. Increasing computing performance to reduce processing time and improve job turnaround.

SOLUTION

Paulsson Geophysical Services selected system integrator PSSC Labs to supply a PowerWulf cluster containing 47 Second-Generation AMD Opteron™ processors, equaling 94 processor cores, with more than 400GB of local disk space as well as multiple terabytes of DAS-connected storage.



POWERWULF CLUSTERS: PHYSICS / ASTROPHYSICS



CASE STUDY CONTINUED

IMPACT

The AMD Opteron™ processor-based PowerWulf cluster increased processing performance dramatically, delivering faster job turnaround and higher customer satisfaction, and helping to increase revenues.

ORGANIZATIONAL PROFILE

A spin-off from Chevron Research in 1997, Paulsson Geophysical Services, Inc. (P/GSI) acquires and processes geophysical data for major and independent oil companies. P/GSI (www.paulsson.com) is called in by oil companies to troubleshoot problems with extracting oil or gas from existing wells. The process (known as Massive 3D VSP®) involves lowering data collection devices into one or more wells to obtain seismic data needed for computing a high-resolution image of the subsurface geology. Energy-producing devices placed at known locations on the surface are used to generate sonic vibrations (like Sonar) whose echoes are then collected inside the wellbore, deep underground. The technique in general is called vertical seismic profiling, and is commonly referred to in the oil industry as VSP. The data are then returned to P/GSI's headquarters in Brea, California, where the company's geophysicists use a compute cluster to produce 3D maps of the subsurface. P/GSI is currently working in the Middle East, California, Canada, as well as other locations.

In addition, P/GSI uses its Massive 3D VSP® and 3D VSP data analysis services to perform microseismic studies. In such studies, there is no seismic energy source on the surface; the data collection devices are placed in one or more wellbores and simply "listen" to the earth for many days or weeks at a time. In these cases, the data collection gear monitors the sound of the earth's crust fracturing, listening for and recording small seismic events that are used to foretell larger earthquakes. P/GSI, in conjunction with a university consortium, completed a project last year to help map details of California's San Andreas Fault near the city of Parkfield in order to visualize how the fault is moving and to build better models that will aid future earthquake prediction.

CHALLENGE

P/GSI's VSP studies generate massive amounts of data. Oil field troubleshooting projects can produce from 10 to 25 gigabytes of data, while microseismic studies can easily produce 1 to 2 terabytes. The process of turning raw seismic data into reliable 3D images of the subsurface can take anywhere from one to three months. Since there is often a very large investment at stake, it is critical to its business that P/GSI provide accurate subsurface images to its clients in as short a timeframe as possible. Geophysicists are continually refining their models and parameters, working in an iterative fashion to generate the most accurate images. These sorts of data processing problems lend themselves to compute clusters where the processing can be broken into pieces, and each piece assigned to a cluster node.



POWERWULF CLUSTERS: PHYSICS / ASTROPHYSICS



CASE STUDY CONTINUED

Since 2000, P/GSI had worked with a Linux® Networx cluster that grew from an original ten nodes to 22 server nodes, all of which used dual-core Xeon processors. While initially more than adequate for generating seismic images, the cluster's performance was seriously lagging behind the state of the art by the time IT Manager Sue Griesbach started with P/GSI in 2004. "By the time I started with P/GSI, we knew that by upgrading to faster processors we would be able to use newer, more compute-intensive software to improve the resolution of the final images while simultaneously reducing the turnaround time for getting results to our clients," says Griesbach.

SOLUTION

To research a replacement for the aging compute cluster, Griesbach and her team talked with several major hardware vendors and system integrators, and were able to obtain test environments from many of them. "We run our own proprietary software on the cluster," she says, "so the standard benchmarking suites don't address our concerns very well. We wanted to test the potential solutions with jobs we would actually run."

In the end, servers based on AMD Opteron™ processors were the clear winners. "The AMD Opteron processing performance was clearly outstanding," Griesbach says, "so it was a pretty simple decision from that point of view."

In the spring of 2005, P/GSI purchased a 47-node PowerWulf cluster from PSSC Labs that features Second-Generation AMD Opteron processors. Each server node has 4GB RAM and two hard disks – one 80GB disk for the Linux operating system, and a second 400GB drive for scratch data storage during processing. Three of the server nodes have an additional 4GB of RAM. The head node uses two RAID1-configured hard drives that hold the operating system and custom applications. The head node and DAS storage provide the data warehouse for the cluster. The head node itself has 4.8TB of raw storage space, with an additional 2TB of attached storage. In addition, the PowerWulf cluster runs the latest version of PSSC Labs' CBeST (Complete Beowulf Software Toolkit) for management, monitoring, and maintenance of the cluster management software.



POWERWULF CLUSTERS: PHYSICS / ASTROPHYSICS



CASE STUDY CONTINUED

SOLUTION

The servers are interconnected via Gigabit Ethernet to the head node and to the DAS storage device via NFS. "Ours is a massively parallel computing environment, so each server does most of the work by itself without a lot of data transfer to other servers in the cluster," Griesbach says. "Gigabit Ethernet interconnect is fine for us."

Porting internal applications to the new cluster was not difficult, as most of the software development takes place on modern operating systems, even though the old cluster continued running under RedHat 7.2. As a result, moving the processing jobs to the new cluster was extremely smooth. "We were able to transition our production processing over to the new cluster with essentially no impact on our production schedule or customer deliveries," says Griesbach.

Since the spring of 2005, P/GSI has run all of its jobs on the new AMD Opteron processor-based PowerWulf cluster. Along with the server systems, P/GSI purchased a 3-year on-site maintenance contract from PSSC Labs. Since PSSC Labs is located just a few miles from P/GSI's office, PSSC has been able to offer same-day responses to requests for on-site service or support. "This has been invaluable to us as it's extremely important to minimize the cluster downtime and keep our production processing running efficiently and on budget," says Griesbach.

IMPACT

By nearly quadrupling the number of processor cores and delivering higher overall throughput via a performance-optimized platform, the PowerWulf cluster increased job processing performance at P/GSI dramatically. Jobs that formerly took six weeks to run can now be completed in much shorter times, and the processing of data from seismic monitoring projects that formerly took months can now be completed in weeks. While Griesbach attributes some of the performance improvements to ongoing tweaks to P/GSI's proprietary analysis software over time, the new cluster has been a major factor in the change.

"Obviously, the PowerWulf cluster has made a huge improvement in our job volume and customer satisfaction," says Griesbach. "Since we generally try to use the cluster on a serial basis and apply all our cluster resources to one job at a time, being able to turn jobs around so much more quickly means we can take in that many more jobs, or handle jobs that require an extraordinary amount of processing in far less time."



POWERWULF CLUSTERS: PHYSICS / ASTROPHYSICS



CASE STUDY CONTINUED

In addition to boosting job performance and throughput, CBeST has significantly reduced the management hassles, and faster processing has led to happier users and lower administrative overhead for Griesbach. "I don't have any complaints about cluster management anymore," Griesbach says.

Power and cooling requirements are a factor in any future decision to upgrade the cluster. Griesbach feels that the PowerWulf cluster's performance is meeting their needs today, while forthcoming Quad-Core AMD Opteron™ processor will give her options for scaling the cluster. "We plan upgrades of our computing resources in response to the needs of specific contracts, rather than as a strategic business decision," she says. "If we had the contracts in hand to warrant it, we could probably quadruple the number of cores in the cluster in our existing space to accommodate the new business, but we would also probably look at leasing the extra capacity to give us some flexibility."

Whatever the future holds, Griesbach knows that her decision to go with PSSC Labs and the AMD Opteron processor-based PowerWulf cluster has been the right one. It is delivering the massively parallel computing performance that P/GSI analysts need to create 3D underground maps, and PSSC's service team has been extremely responsive. "PSSC's proximity has really helped me," she says. "They have come out on the same day in several instances, and it's really made a difference for me. I'm really glad we took this route because I think it's really helped uptime for the whole cluster."

Whether it's monitoring a major earthquake fault to aid future predictions or improving oil and gas production with highly detailed 3D maps, P/GSI is offering sophisticated analytical services with rapid turnaround times. Today, the PowerWulf cluster and Second-Generation AMD Opteron processors are key factors in its success.