Why PGS Made the Move from Clusters to Supercomputers

CASE STUDY : : PGS
MARINE SEISMIC COMPANY
LOCATED WORLDWIDE
HIGH-RES, 3-D MODELING

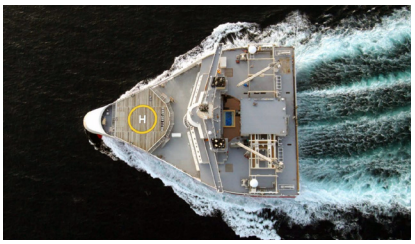

FACED WITH THE MOST COMPLEX SEISMIC SURVEY EVER COLLECTED, PGS TOOK A NEW APPROACH TO SYSTEM DESIGN… AND SUCCEEDED.

PGS BEFORE CRAY…

Incapable of processing Triton dataset

…AND AFTER

Industry-leading RTM at ~129M traces/min

Running more, larger jobs using more complex data and algorithms

Running individual jobs faster with higher-quality results

BUSINESS CHALLENGES

Increasing survey size and data density
Complexity of imaging algorithms
Tight turnaround and cost margins

CRITICAL COMPUTE NEEDS

Better system performance and throughput
Huge amount of memory
Strong interconnect
Ability to scale to problem size

BACKGROUND: THE TRITON SURVEY — When PGS launched the Triton survey in November 2013, they knew they'd end up with the largest seismic survey they'd ever collected. When they finished acquiring the data in August 2014, they had the most complex imaging challenge they'd ever faced.

The high-fold, long-offset, dual-sensor, broadband survey provided full azimuth data and covered approximately 10,000 square kilometers in a notoriously difficult-to-image area of the Gulf of Mexico.

Considered the most revolutionary and technologically advanced survey ever conducted in the region, Triton is impressive not just for its size but because with it, PGS deployed a breakthrough survey design intended to overcome the challenges of this structurally complex deepwater area.

Their specially designed acquisition template used five vessels in a combined wide-azimuth and simultaneous long offset (SLO) configuration, resulting in full azimuth coverage. The design, combined with PGS' acquisition technology GeoStreamer®, produced the massive amounts of data that formed the basis for the survey — and clearest subsurface images of the region yet.

However you look at it, the Triton program's resulting 660 terabyte dataset represented PGS' largest, most complex data processing and imaging challenge ever.

CRAY

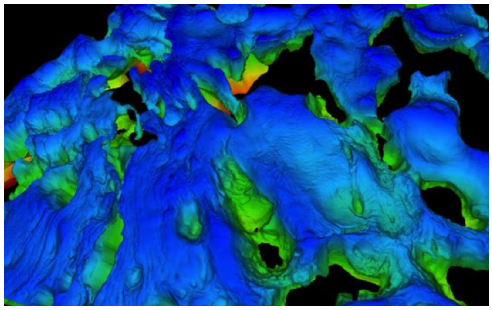

CHANGING SCIENTIFIC AND BUSINESS NEEDS ARE STRESSING COMPUTE INFRASTRUCTURES ACROSS THE SEISMIC INDUSTRY.

CHALLENGE: Producing Images from More Data with Less Time and Money — From a scientific perspective, seismic imaging is a highly complex task.

For PGS, it demands the production of accurate, clear images of varied geology at up to a depth of 16 kilometers from data with large velocity and density contrasts and incomplete boundary conditions using reverse time migration (RTM).

To solve the problem of how to process this data into an image, commensurately complex algorithms have evolved. Overall, the industry is struggling to handle the explosion of data volume, fidelity and algorithmic complexity.

From a business perspective, demand for energy is spiking at the same time that resources are getting harder to find and the oil and gas (O&G) industry as a whole is experiencing cutbacks. For PGS, a tightening market means they must collect data and produce images "faster, cheaper, better" by improving image quality and increasing productivity with more automation.

Together, these scientific and business needs are stressing compute infrastructures across the seismic industry.

PGS and companies like them have typically used clusters for their HPC needs, but this new era of massive data volume ushered in by the Triton survey and the overall tightening of margins led PGS to the realization that they couldn't meet any of these demands with their existing compute technology. They needed a "radically different" solution.

SOLUTION CHOICE

Cray® XC™ series supercomputer

Cray® Sonexion® storage system

SYSTEM DETAILS

24 cabinets, all CPU

600 TB of memory

Aries™ interconnect

TRITON SURVEY BY THE NUMBERS

660 TB dataset

Required RTM at 35 Hz, 16 km aperture, 16 km depth

6 azimuths, 46 angles for 276 angle volumes

Problem size: ~200 GB per shot

THE IDEAL SEISMIC IMAGING SUPERCOMPUTER

Cost effective for current applications

Fits existing technology pipeline

Enables implementation of new algorithms

Reduces development time

WWW.PGS.COM

SOLUTION: Why a Supercomputer, Not a Cluster

For decades, exploration seismology has been conducted in a very repetitive, linear process manner — one well-suited to cluster computing. But as survey complexity increases and data volume grows, the customary single-shot approach no longer produces accurate enough images.

In order to meet their business and scientific goals, PGS recognized they needed to shift their R&D approach. That shift included using all 3-D operators throughout the processing flow, moving toward full-wavefield and full-survey processing and imaging, and leveraging angle-domain image space for analysis, residual corrections and conditions.

All of these R&D drivers presented data management and processing challenges better suited to supercomputing. But at the same time that supercomputing solution needed to be cost effective for current applications, fit their technology pipeline, enable implementation (at scale) of new algorithms, reduce development time and keep up as volume and complexity continue to increase.

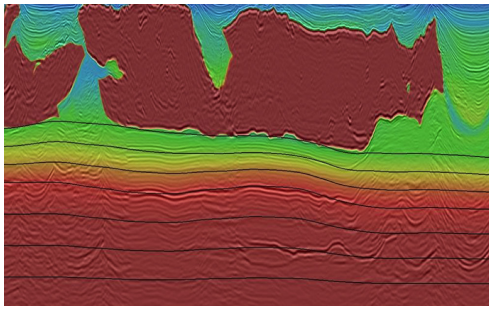

"WITH THE CRAY SUPERCOMPUTER, OUR IMAGING CAPABILITIES WILL LEAPFROG TO A WHOLE NEW LEVEL. WE ARE USING THIS TECHNOLOGY INVESTMENT TO SECURE OUR MARKET LEAD IN BROADBAND IMAGING AND POSITION OURSELVES FOR THE FUTURE."

Guillaume Cambois | Executive Vice President, Imaging & Engineering | PGS

ABOUT PGS

PGS is a marine seismic company that acquires high-resolution seismic data that is used to generate accurate images and 3-D models of the subsurface.

CRAY

Cray provides systems and solutions that help you solve your most difficult computing, storage and data analytics challenges. Our portfolio includes powerful supercomputers, optimized cluster systems, advanced storage systems and data analytics and discovery platforms.

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164
Tel: 206.701.2000
Fax: 206.701.2500
WWW.CRAY.COM

Enter Cray

Critical to PGS for their production imaging was a massively parallel system with a huge amount of memory, strong interconnect and the ability to scale to the problem size — while also fitting their production schedule, business goals and budget. They selected the Cray® XC™ series supercomputer with Cray® Sonexion® storage.

Of particular interest to PGS was Cray's system design, particularly the shared memory, Aries interconnect and software environment. The shared memory eliminates the time-consuming process of separating data and then putting it back together. The Aries™ interconnect provides cost-effective, scalable global bandwidth that significantly speeds processing. And the integrated software environment features a highly parallel MPI as well as the best compilers, debuggers, job schedulers, interprocessor communication capabilities and performance analysis tools.

With their new system, PGS went from being unable to process the Triton survey within their production deadline to industry-leading RTM processing capabilities at ~129M traces/min. What's more, their codes scaled and performed beyond the competition. They can now run more, larger jobs using more complex data and algorithms and run individual jobs faster with higher-quality results.

Expertise: The "Hidden" Cray Benefit

Cray has been honing its capabilities in the HPC marketplace for over 40 years. Today, it has a dedicated O&G team with an average of 30-plus years in the industry as well as industry-leading analysis, tuning, service and support. The company has successfully managed technical pathways for exploration companies around the world. What it means is that when you work with Cray, you get a partner. No other company in the market today can offer the same level of expertise and longevity in managing HPC deployments as Cray.

Looking Ahead

With the relentless turnaround, cost, algorithmic complexity and data volume challenges ahead, the last decade's status quo approach to building and deploying HPC systems won't meet the industry's needs going forward. Next-generation systems will need to include both capability- and capacity-oriented systems. Achieving scalability goals will depend on software development environments with a focus on parallelism. And finally, planning for next-generation HPC systems will require consideration from a holistic approach, designing, integrating and optimizing all aspects of the architecture as a complete system.

PGS solved these challenges by working with Cray to successfully move from cluster computing to supercomputing.