WHITE PAPER
Intel® Xeon® Processors
Energy
High-Performance Computing

Shell Drills Down on Hundredfold Improvements in High-Performance Computing

Strategic collaboration, software optimization, and processor advances help Shell meet rising requirements

High-performance computing (HPC) is mission-critical for Shell. Having more HPC capacity helps us explore for and find hydrocarbons in increasingly more complex and difficult geologic environments, extract resources efficiently, and help meet the world's energy needs while minimizing environmental impact.

Shell's demand for HPC capacity is increasing dramatically. For our upstream exploration businesses, newer seismic acquisition techniques are generating larger and larger seismic datasets, which alongside new seismic algorithms are creating more and more demand for HPC capacity. In addition, dramatically more complex and larger reservoir models, combined with new reservoir simulator technologies, are increasing HPC demand. Shell's research and development (R&D) programs and other new emerging HPC opportunities also require next-generation HPC capabilities as well.

A key driver for Shell is the desire to deliver the required HPC capacity at the best-possible total cost of ownership (TCO). We deliver as much capacity as possible from Shell's global data centers in Houston and Amsterdam, and in general try to concentrate capacity in fewer locations. This approach helps us create scale benefits in cost and service levels and will help us transition to cloud-based solutions as they mature and become more affordable in the future.

Because of HPC's importance, Shell established an HPC Center of Excellence (CoE) in 2009, and set a goal of hundredfold improvements in application price/performance. Today, through rapid adoption of the latest technologies, an ongoing focus on software optimization, and close collaboration within Shell and with external leaders, we are well on our way to achieving that goal. In doing so, we are delivering technical successes, enabling significant results, and providing competitive advantages.

Scott Meyers
Shell Global Solutions US Inc. - Projects and Technology

Anu Krishnan
Shell Global Solutions US Inc. - Projects and Technology

Jeff Jackson
Shell Global Solutions US Inc. - Projects and Technology

Close Collaboration to Improve Performance

Our philosophy at Shell is that HPC is not only about compute, but also should include applications, data, high-speed storage, networks, and visualization technologies in order to deliver the most business value. Accordingly, the HPC CoE takes a comprehensive approach to meeting Shell's HPC requirements.

Our approach is highly collaborative. Within Shell, the HPC CoE works closely with:

- Innovation, Research, and Development (IRD), which brings the latest and best geophysical techniques
- Technical and Competitive IT (TaCIT), whose software engineers squeeze the last drop of performance out of techniques by heavily optimizing algorithms

Table of Contents

- Technical Landscape Services (TLS), where we receive the latest HPC infrastructure and services at the lowest possible cost and ensure that software and workflows perform optimally and reliably

Externally, we have established strategic collaborations with key suppliers such as Intel and with leading R&D academic leaders. Our collaboration with Intel gives us early access to the latest Intel® technologies and a longer-term view of next-generation microprocessor designs. Early in our collaboration with Intel, we created an R&D sandbox environment with 256 Intel® Xeon® processors 5500 series. We have since upgraded this environment and are using it to develop next-generation algorithms and optimize our key workloads. We have recently deployed the latest Intel Xeon processor E5 family, which provides a dramatic performance boost for most of our key applications.

Intel's consistency in delivering on its tick-tock model of architectural advances has allowed Shell to rapidly deploy each new generation of Intel Xeon processors into our environment and take maximum advantage of their performance advances.

Software Optimization

Our collaboration with Intel is not limited to hardware. Both Shell and Intel have made significant investments in engineer-to-engineer collaboration and training to optimize Shell's algorithms and applications for the Intel processors and to enable our codes to fully utilize rising core counts. We believe our optimization efforts are a key to our HPC success.

We utilize the Intel® C++ and Fortran compilers and Intel® Math Kernel Library, as well as tools such as Intel® VTune™ Performance Analyzer. We have provided training to a range of software developers and researchers. We have also built a small cadre of HPC experts who share best practices, allowing us to deliver even greater value from HPC optimization efforts.

Intel has supported our optimization efforts with its own expertise. Intel enterprise technology specialists work with Shell teams in the United States and the Netherlands. Software engineers from Intel's Software and Services Group have provided insights into new capabilities, and our experts have shared their thoughts on desired features. In one case, more than three dozen Shell and Intel software experts spent two full days digging deep into next-generation technology requirements and capabilities.

Infrastructure Choices for Performance and Throughput

HPC infrastructure at Shell includes the microprocessor, storage, network, and even visualization technologies used by the application portfolio. One key premise underpinning our HPC infrastructure is that we optimize on a combination of price and

"In the Gulf of Mexico we are typically looking at prospects in the ultra-deep subsurface underneath very complex salt. If you remember you are looking five to six miles under the surface of the Earth for a very small target, it's pretty miraculous that we can find these things."

David Schewitz,
Manager of Geophysics, Geomatics and Data Management for Upstream Americas, Shell

performance, so if a technology delivers a tenfold performance improvement, it must be less than 10 times more expensive.

Shell's primary HPC infrastructure is based on the Intel Xeon processor 5600 series, and we have already deployed the Intel Xeon processor E5 family in large numbers. The Intel Xeon processor E5 family delivers up to twice the floating-point performance of the previous generation along with double the I/O bandwidth.

We also evaluate innovative products with general-purpose graphical processing units (GPGPU) or field programmable gate array (FPGA) solutions based on their price-performance and overall TCO. We recently installed a cluster that combines Intel Xeon processors and NVIDIA* GPGPUs and are using it to test advanced visualization tools.

Results Highlights

Over the last three years, Shell's strategy has produced a number of technical successes and enabled significant results in our operations in the United States and internationally.

For example:

- Even before we deployed the Intel Xeon processor E5 family, the TaCIT team that supports Shell's custom-developed Dynamo* reservoir engineering model tripled its performance. Specific Dynamo workflows achieved significantly higher increases.
- Our software optimization work is producing improvements of two- to twelvefold depending on the workflow, with increases of eight- to sixtyfold based on access to a larger-capacity pool.
- We have more than doubled the HPC capacity in our largest data center between 2008 and 2011.
- The seismic technology team in IRD and the seismic applications team in TaCIT have deployed higher-performance versions of a number of critical applications, enabling the seismic processing teams to utilize the latest and greatest algorithms.

Intel's consistency in delivering on its tick-tock model of architectural advances has allowed Shell to rapidly deploy each new generation of Intel Xeon processors into our environment and take maximum advantage of their performance advances.

An Example of Business Impact

What do these improvements mean for Shell's business? David Schewitz, manager of geophysics, geomatics, and data management for Upstream Americas, recently provided an insight into how the advances in Shell's HPC platform have contributed to the success of subsurface seismic surveying activities in the Gulf of Mexico.

Last year Schewitz and his team were conducting pre-stack depth migrations using ocean-bottom survey data acquired in the Gulf of Mexico. They used a deep-penetration air gun to fire controlled concussion pulses into the water. These pulses traveled through the subsurface, reflected, and refracted off the geologic layers to produce sound waves. The team used special sensor nodes placed on the sea floor to digitally record the energy's travel time and amplitude strength after the trip through the subsurface.

Using HPC clusters based on the Intel Xeon processor 5600 series, Schewitz and his team analyzed the travel time and amplitude data and constructed a picture of the subsurface. Depending on the acquisition survey design, this method provided imaging in excess of 10 miles below the earth's surface.

"In the Gulf of Mexico we are typically looking at prospects in the ultra-deep subsurface underneath very complex salt," explains Schewitz. "If you remember you are looking five to six miles under the surface of the Earth for a very small target, it's pretty miraculous that we can find these things."

Schewitz says Shell's HPC resources were critical in storing and processing the data his team collected. "Without HPC none of these business results could have been obtained," he comments. "It underpins everything we do."

Processing and storing seismic data is expensive, but Schewitz says the costs in the last 10 years have dropped by a factor of more than two-and-a-half orders of magnitude. "These savings allow us to process more surveys, deploy new seismic technologies, run more projects simultaneously, and make each project larger and more dense," Schewitz adds. "Without the benefits of HPC, we could not do any of this work and certainly not to the scale that we have been."

More accurate seismic images, available through gathering multi-azimuth data, also enable significant cost reductions through better well placements and fewer development wells with higher recovery rates. The Mars-Ursa basin, where Shell acquired the ocean bottom seismic (OBS) survey, has many fields and prospects whose imaging is affected by salt. OBS wave-equation anisotropic imaging has already reduced risk on many of these prospects.

Things are heating up further this year as Schewitz's team acquires five different OBS surveys covering eight existing Shell fields. Shell will use this data to monitor appraisal drilling and to support 4-D work.

Conclusions

Shell is committed to a continuous effort to maximize the value of high-performance computing. With the intense collaboration and commitment of many organizations across Shell and Intel, the HPC CoE is on the path to achieving its goal of a hundredfold increase in application performance.

Meanwhile, we will continue to deliver every advantage possible to Shell and our customers by choosing leading-edge platform technologies, collaborating with industry and academic leaders, and optimizing our applications for maximum performance on each generation of Intel technology.

Find the solution that is right for your organization. Contact your Intel representative, visit **Business Success Stories for IT Managers**, or explore the **Intel IT Center**.

Printed in USA 1012/LJ/TDA/XX/PDF ♻ Please Recycle 328073-001US