

THE AMERICAN OIL & GAS REPORTER[®]

NOVEMBER 2009

The "Better Business" Publication Serving the Exploration / Drilling / Production Industry

Hardware, Software Innovations *Ushering HPC Capabilities Into Oil And Gas Mainstream*

By Kari Johnson
Special Correspondent

As oil and gas executives know very well, any company engaged in the business of exploration, drilling and production is also an active participant in the technical computing business. The reason is simple: Computing has become ubiquitous to all facets of upstream operations. And now innovations in hardware and software technology are bringing into the oil and gas mainstream the kinds of high-performance computing capabilities that were once far out of the reach of most operators and most applications.

Although computing is at the core of virtually every work process and takes all imaginable forms—from wireless boxes passively recording seismic signals to remote terminal units self-regulating the operation of unmanned offshore platforms—exploration remains the most data- and computationally-intensive business function. Every day, oil and gas companies stake high-dollar investment decisions based on bits of insights gleaned from terabyte-sized data sets.

Reproduced for LSI Corporation with permission from The American Oil & Gas Reporter

With today's high-density, wide-azimuth seismic recording capabilities, even a relatively small-scale exploration program can generate enormous volumes of raw 3-D data that require massive amounts of compute power to acquire, store, manage, process, visualize and interpret, notes Larry Lunardi, vice president of geophysics at Chesapeake Energy Corporation.

"High-density data sets are becoming routine and much larger," he remarks. "The amount of data gathered per square mile and our expectations of it are both escalating."

Of course, the trend toward larger, high-definition data sets is nothing new, Lunardi points out. "It is really no different than when I started in this business years ago," he admits. "In those days, the big deal was to digitally migrate 2-D data. That brought existing mainframes to their knees. Today, we are designing seismic surveys we could not even dream about then, but we are still limited by compute capacity and it still takes too much time to process."

The goal of high-performance computing is to address the technical issues and limitations involved in analyzing and displaying large amounts of data, placing a burden on computing power and storage capacities. Lunardi says exploration presents a "moving target" for HPC, because as drilling prospects become more complex and geologic attributes become more subtle and less conventional, more data with more usable bandwidth are required. This, in turn, drives an ever-increasing need for more subsequent processing and analysis.

"In highly technical plays, where there are overturned beds, faulting and complex folding, a conventional seismic analysis would look like a television screen with nothing on it but static," Lunardi comments. "To have a better picture of the subsurface, we need the equivalent of high-definition in 3-D."


Obtaining a high-resolution seismic view comes at the cost of handling enormous amounts of data. "Where in the early days we used to get perhaps 10,000 prestack traces a square mile, now it is not unusual to gather 1 million traces," Lunardi relates.

Computing And Resource Plays

As a leading player in several North American unconventional natural gas resource plays, Chesapeake also puts HPC to work in its shales and tight sands, where Lunardi says 3-D imaging has to get down to the level necessary to detect fine-scale fracturing and anisotropy characteristics. "We use seismic to give us a structural road map," he explains. "It helps us identify dip changes and small faults, as well as natural fracture orientation."

To get to reservoir-quality determinations, however, Chesapeake must routinely perform inversions of prestack and post-stack data. Reservoir maps must be continuously updated. Because of the nature of resource plays, the company will drill a succession of wells that must be regularly reincorporated into the interpretation at the rate of every three to four months, according to Lunardi.

Using advanced prestack 3-D techniques, Chesapeake is beginning to extract detailed rock properties that Lunardi says allow interpreters to discriminate "good rock from bad rock" in the Marcellus, Haynesville and Fayetteville shales. The increasing focus on prestack data has dramatically increased data processing requirements. "With prestack, we have to deal with every trace,



With the advent of multicore chips, oil and gas firms will need new kinds of software to take advantage of multicore processing in high-performance data centers. To streamline computational workflows, software must be written to run in parallel on many processors at once, allowing multiple tasks to run simultaneously, rather than sequentially.

whereas post-stack reduces the trace count by 50 to 100 times," he states.

Chesapeake also is gathering multicomponent 3-D data in the Marcellus to determine if adding the shear-wave component to the compressional data will improve the understanding of fracture orientation and the existence of fracture "swarms" in shale plays, he adds. "However, depending on how the data are acquired, multicomponent seismic increases the total size of the data set by orders of magnitude," Lunardi says. "It requires a lot more processing and a lot more computing resources for the additional data contained in multicomponent surveys to be useful."

Fortunately, processing technology continues to evolve at a rapid pace, as advances in computing speed and power make it possible to run even the most complex algorithms in a timely and cost-effective manner, Lunardi goes on. "Initial processing routines and advanced algorithms are now fast enough to get results in weeks, making it possible to get a useful result within the lease schedule," he holds.

Thanks largely to the scalability and low-pricing points of clustered supercomputing systems, Lunardi says he sees a growing trend among independents to bring some advanced processing capabilities in-house to go along with the interpretation systems. For example, Chesapeake maintains an internal high-performance IBM system that it uses for processing and seismic data inversion, he adds, noting that Chesapeake relies on file servers for its data, which numbers in the terabytes.

"Storage has become a very big deal for us," Lunardi relates. "We now have more than 23 million acres worth of 3-D seismic, and those data just keep on growing and giving."

As technology evolves, Chesapeake can routinely revisit its data store and run new processing algorithms and interpretation tools to continually improve its understanding of the subsurface. "Whether its depth imaging, preprocessing filters, post-stack enhancements, coherence cubes or multiple inversion volumes, there are all kinds of options for reprocessing data," Lunardi concludes. "Right now, we simply are trying to keep up with

leasing in our shale plays. Our computing infrastructure is helping us stay ahead of the drill bit.”

Adhering To Moore’s Law

In keeping with Moore’s Law, users have come to expect computers to become more powerful and less costly each year. However, as computer makers increasingly turn to multicore chips to deliver performance, oil and gas firms will need new kinds of software to take advantage of that hardware in high-performance data centers, says Jan Odegard, executive director of Rice University’s Ken Kennedy Institute for Information Technology in Houston.

The institute’s computational scientists are working with a number of energy industry firms that want to avoid the “software bottleneck” that multicore processing will bring, according to Odegard.

“The Linux commodity cluster platform has been driven, until very recently, by Moore’s Law, and by the fact that until about 2004, the doubling of transistor density every 18-24 months roughly meant a doubling of performance every 18-24 months at ever-lower cost,” Odegard states. “While Moore’s Law is predicted to hold true until about 2020, we are not able to continue increasing the clock rate, and hence, performance is not ‘free.’ Now, instead of looking for increased performance from clock speed, the microprocessor industry is delivering an increasing number of computing cores with each new generation of processors. Effectively, this moves the quest for performance from a hardware approach to a software challenge.”

Odegard says the Kennedy Institute favors Linux-based commodity systems and acquired its first shared cluster in 2002. This past summer, it deployed a GreenBlade™ cluster from Appro to support rapidly growing computational needs in the area of information technology and computational science and engineering. In recent years, he says he has noticed energy companies paying increasing attention to the challenges of leveraging high-performance computing in their day-to-day operations.

To take advantage of HPC and to streamline computational workflows, software must be written in a way that allows it to run in parallel on many processors at once, Odegard explains. This allows multiple tasks to run simultaneously, rather than sequentially. While parallel processing is not new (it is the strategy programmers have used with supercomputers for years), he notes that the scale at which the industry needs to parallelize software is unprecedented.

“While some applications like seismic processing are, by nature, fairly easy to parallelize, others will take substantial effort, including developing new algorithms,” Odegard says.

With the proliferation of computing in all oil and gas technical disciplines, Odegard suggested that computational problem solving and some level of programming proficiency should be considered critical skill sets for any graduate aiming at a technical career in the oil and gas industry. “One of the key things I have discussed in my interaction with the industry is the rapidly growing need for workers who have the required computational skills to complement their domain specialties,” he states. “We already are seeing education and training programs focused on this, and many more are in the works across academia.”

Another major challenge facing the industry, Odegard observes, is that state-of-the-art processors are more power-hungry than ever. “We are seeing HPC data centers regularly drawing 3 megawatts-5 megawatts, and the trend is for more power, not less. A back-of-the-envelope calculation suggests that if current trends continue, data centers could soon need 100 megawatts of available power,” he remarks. “The cost of operating these kinds of facilities will break most budgets, so the current trend clearly is not sustainable from an environmental perspective or a business perspective.”

Since the demand for increased computational power continues to grow, the information technology industry is looking at some radical new approaches to delivering high-end performance with energy-stingy architectures. For example, Odegard says groups are experimenting with systems built on custom, low-power embedded processors—the same technology that allows long life on cell phone batteries.

“The idea is to get more computing per watt of energy used,” Odegard details. “Instead of having 100,000 or more compute cores, the new approach will require 1 million-10 million cores.”

Although there are major challenges in delivering such a massively parallel approach, efforts are underway. Green Flash, a project at Lawrence Berkeley National Laboratory, is one such project that Odegard says he is tracking. “Writing software to leverage performance from a system with 100,000 cores is hard, and leveraging the performance of 1 million-10 million cores will not be any easier,” he remarks. “We are right back at the software challenge, no matter how the industry evolves.”

Odegard says HPC requirements are expanding within oil and gas companies. “The energy industry began outsourcing a significant part of its computing operations years ago, but most companies maintain substantial in-house resources, and we have seen those grow over the past decade,” he states. “This internal capability allows companies to apply their own intellectual property to the final imaging stage. It is all about maintaining a business advantage over the competition.”



The trend toward higher-resolution survey designs with increased density has led to average channel counts more than tripling during the past decade at Dawson Geophysical, from 2,000-3,000 to more than 10,000 channels a crew. Here, Mike Weis, Dawson’s processing manager (left), and Sam Dobbs, vice president of data processing, discuss a data processing issue.

“While a substantial part of the in-house systems consist of commodity processing clusters, an increasing number are enhanced by accelerators, and most notably, general-purpose graphics processing units (GPGPUs),” Odegard continues. “One thing is for sure: The demand for computing infrastructure in the oil and gas industry will continue to grow as the industry faces increasingly challenging subsurface imaging conditions.”

Faster, Better, Cheaper

One primary reason for the growing demands on computing infrastructures is the fact that seismic surveys are being shot with higher-resolution arrays, higher-density configurations and collecting a fuller sampling of the seismic wave field. That is where companies like Dawson Geophysical enter the scene. “The demand on our business always has been higher-quality, cost-effective surveys in the shortest cycle time possible,” says President Steve Jumper. “The big push today is higher channel counts with greater surveying density, because better images and efficiencies ultimately come from increased channel counts.”

Only a decade ago, Jumper says Dawson Geophysical was averaging 2,000-3,000 channels a crew. Today, that number can be greater than 10,000. The increased density produces an effect on data processing requirements when it comes time to crunch the raw data. In addition, most 3-D data sets are now prestack time migrated, which also adds to the computational requirements. “I do not see any end to the increased intensity and size of data sets,” he offers.

That trend will only accelerate as the industry brings shear-wave energy into the imaging picture, Jumper goes on. “There are still issues in exactly how to acquire, process and understand the response from multicomponent data,” he says, noting that Dawson has acquired 12-15 multicomponent 3-D surveys over the past four years. “But multicomponent seismic will become more of a factor in our business going forward, and when it does, we are talking about increasing the sizes of the data sets by 300-900 percent.”

Phil Lathram, information technology manager at Dawson, adds, “The increased volumes and intensity of the data, combined with the mathematical algorithms associated with advanced processing techniques to improve final image quality, make compute power key. We have moved to a nodal system that ties multiple processors and provides huge disk space.”

The company has three dedicated U.S. processing centers to handle the data its crews acquire in the field. “On occasion, we also do some basic processing in the field as a check on data quality and survey parameters,” says Sam Dobbs, Dawson Geophysical’s vice president of data processing. “For field startups, for example, we may do a few hours of testing and some in-field processing.”

The company houses all of its data at a secure site in Midland, Tx., with remote links to its processing centers in Houston, Oklahoma City and Midland using high-speed T1 telephone lines with 3.0-megabit connections, Dobbs explains. Clusters of dual-processing central processing nodes are connected to a server environment for migration purposes. “Moving data out to a node and back is still an inefficient process, and it shows when a lot of data is being moved around,” he comments.

Dawson upgraded its cluster computing hardware two years



Brigham Exploration’s Chief Geophysicist Russell Pruitt (seated), Director of Geophysical Projects John Plappert (left) and IT Manager/Senior Geophysicist Dave Kirchhoff (right) analyze recent horizontal drilling results in the Bakken/Three Forks oil resource play in the Williston Basin using CGGVeritas Inc.’s multicient South Ross 3-D data set covering 115.5 square miles in Mountrail County, N.D.

ago to Intel® Xeon®quad pro processors, according to Lathram. “Our infrastructure is under continuous review,” he says. “We are particularly interested in technologies that provide high performance while reducing power consumption and space. It is important for us to watch for bottlenecks in the overall system. As hardware improves, the software needs to improve to take advantage of that. To maintain efficiency, you have to constantly attack whatever is the slowest component in the chain.”

Primary Tool

Brigham Exploration Company was founded in 1990 by geophysicist Bud Brigham, who now serves as the company’s president, chief executive officer and chairman of the board. Given his unique background, the company continues to place a heavy emphasis on geophysics and other advanced technologies in its exploration and development efforts. To date, the company has accumulated a library containing nearly 13,000 square miles of 3-D seismic data, reports John Plappert, Brigham’s director of geophysical projects.

“Throughout our history, 3-D seismic data has been the primary risk reducing tool used for prospect analysis and discovery,” he says.

That certainly remains the case today, as Brigham Exploration’s operational focus increasingly shifts from conventional to unconventional resource plays such as the Bakken oil shale in the Williston Basin, where Plappert says Brigham has amassed nearly 300,000 lease acres and is drilling horizontal wells with 10,000-foot vertical depths and 10,000-foot laterals.

In 2008, the company contracted with PGS Onshore Inc. to acquire 180 square miles of 3-D seismic data characterized by full-azimuth distribution and high fold to below the Bakken interval, Plappert says. “The goal was to derive information from the surface seismic data, such as anisotropy and inferred fracture patterns, which may influence the drilling and completion of our horizontal wells,” he comments.

Brigham has performed up to 28 frac stages when completing its long-lateral horizontals and achieved initial production rates approaching or exceeding 2,000 barrels of oil a well, Plappert updates. “Brigham always has been a pioneer. We were one of the first to advocate large-scale 3-D surveys, which provide a better overall view of the regional geologic framework, better economies of scale, and more opportunity for success,” he avers. “Our focus is driven by technology and we always are looking for ideas on the forefront.”

Throughout the industry’s history, technology has directly and indirectly led to substantial gains in recoverable resource estimates as the resulting initial frequency and size of new field discoveries increased, Plappert points out. The most obvious recent example is the opening of ultralow-permeability shales and tight sands that were considered noncommercial only a decade ago.

“There is definitely a link between technology and our understanding of available resources,” he holds. “When the industry moved from single-fold shot records to multifold 2-D, the number of large fields discovered increased. Moving from 2-D to 3-D dramatically increased success rates. Multicomponent 3-D, or perhaps some other advance in seismic technology, will no doubt yield another game change. That is why we are always keeping an eye on new and developing technology.”

Dave Kirchhoff, information technology manager and senior geophysicist at Brigham Exploration, says seismic interpretation is performed by teams of geophysicists and geologists who integrate advanced attribute analysis derived from prestack time migrated seismic data, as well as detailed evaluation of seismic gathers, with subsurface geology.

“Each explorationist has his own interpretation workstation running either Windows® or Linux, depending on the needs,” he details. “The workstations are outfitted with dual cores, high-end graphics and plenty of memory. Our focus is to maximize performance while staying aware of cost.”

Brigham’s conference rooms are equipped with dual projectors utilizing visualization software for collaborative peer review and management discussions, adds Kirchhoff. The company was also an early adopter of clustered storage. “We are looking for new technologies that better manage storage, like seismic-specific compression and deduplication (using matching logic to eliminate duplicate file records),” Kirchhoff concludes. “Interpretation teams create so many data volumes now, that terabytes of storage disappear in the blink of an eye.”

‘Intelligent’ Storage

When it comes to increasing speed and reducing overall cycle times, the storage system is a key piece of the puzzle. A well-tuned storage system can drive numerous efficiency improvements in HPC applications, according to Steve Hochberg, senior director of the worldwide HPC segment at LSI Corporation.

“High-performance computing demands that every component is as efficient as possible or it risks becoming a performance bottleneck,” he states. “With the volume of data involved in exploration activities, the opportunity to speed results using an HPC storage system is very real.”

A 40-acre seismic survey could generate more than 20 terabytes of data, Hochberg points out. Once added to the

storage system, all that data must be readily accessible for processing, interpretation and visualization, and multiple renderings of the same massive data sets sometimes have to be maintained on the storage system.

“It is important to consider all of these functions and balance the storage system,” says Hochberg. “Speed is of the essence, but writing data and reading data speeds are different. If a system is much faster at writing than reading, you will be able to quickly add data, but may have a bottleneck when it comes to analysis.”


The workloads created by seismic data processing requires a storage system that offers a balance between high-speed sequential movement of data and random IOPS (input/output operations per second) performance. Processing and interpretation require that the storage system be able to efficiently access random data from a variety of 3-D data and cultural data files. Visualization similarly pulls data from disparate content, Hochberg relates.

At the same time, storage system density is dramatically increasing. “Today, a single dense storage enclosure can provide data capacity of up to 480 terabytes, and that number will double by year’s end,” Hochberg says. “The increased storage density helps reduce floor space and power requirements, even as data volume increases. Improved storage density lowers power requirements by reducing the number of enclosures required to achieve similar storage capacity. Power supplies have also become more energy efficient.”

Combined with advanced capabilities for reliability, availability and scalability, modern storage systems are becoming what Hochberg calls “intelligent networks.” All components are redundant and hot swappable, with the system keeping track and rebuilding or rerouting drives and data as needed, he explains. “Intelligent storage systems can even predict the health of drives and recommend replacement before a drive failure,” Hochberg notes. “When the drive is replaced, its content is automatically rebuilt.”

Storage-related functions are migrating to storage controllers, which—like everything else in HPC—are getting faster. Cluster interconnects with up to 40-gigabits per second. InfiniBand® switched fabric communications architecture will replace fiber channel connections for even faster performance, he predicts.

“Over the next few years, I expect that much of the file



With inherent high-capacity parallel processing capabilities, graphics processing units developed initially for visualization and gaming applications can process large data sets and run complex algorithms more efficiently than central processing units.

system will run on the intelligent storage system,” Hochberg remarks. “Remote volume mirroring, replication and virtualization will all be storage controller-based.”

Harnessing GPU Power

The central processing unit has traditionally been at the heart of computational workloads. But Sumit Gupta, senior product manager for Tesla GPU computing at NVIDIA, says the focus for many computationally-intensive algorithms has shifted recently to graphics processing units.

“When there is a lot of math involved, tasks can be broken down and handled more efficiently by a GPU,” he explains. “Graphics processing units historically have been used to handle visual display, but the high-capacity parallel processing capabilities inherent in GPUs make them applicable to complex HPC.”

GPUs are designed for executing compute-intensive tasks while CPUs are best suited for sequential control tasks, Gupta says. “CPUs have four generals that do not communicate too well with each other, whereas a GPU has hundreds of soldiers that work collaboratively,” he remarks.

GPUs have been most prevalent in gaming devices for graphics manipulation. “Think of that action as putting pixels on the screen,” Gupta offers. “Increasingly, the GPU is used as a physics engine that does the computation for simulated motion. This goes way beyond graphics and makes games appear more realistic.”

Applied to exploration, GPUs can speed seismic processing by performing computations in parallel. “Hess Corp., for example, is using GPUs to run its finite difference and wave equation algorithms,” Gupta reports. “The company has reduced its computing footprint from a cluster of 500 CPUs to a GPU cluster of 32 nodes. The new infrastructure has 1/20th of the cost, consumes 1/27th of the power and takes up 1/31st of the space.”

In seismic processing, Gupta says, the GPU-to-CPU ratio can be about four-to-one today. To get fast performance, each GPU requires roughly 4 gigabytes of memory. Unlike a general purpose CPU, the GPU must be directly addressed by the software application. “To leverage the computational power of the GPU, software applications must have new code added to it,” he explains. “This code can be conditional, so that it takes advantage of GPUs if they are present, or otherwise runs sequentially on the CPU.”

Various seismic migration algorithms—including Kirchhoff, wave equation and reverse-time migrations—are well suited for computation using GPUs, according to Gupta. Other applications likely to take advantage of GPUs for computational power are seismic interpretation and well planning. And in reservoir simulation, the double precision support of GPUs can efficiently solve a series of simultaneous equations. In fact, he says the iterative solvers are a natural fit for the GPU.

“Since the GPU is fast and can be installed in a workstation, reservoir engineers are intrigued with the notion of having their own personal supercomputer to run more models as they iteratively converge on a solution without the hassles associated with waiting in long queues to access large CPU-based cluster farms,” Gupta concludes.

Scaling Up To HPC

Scalable supercomputing architectures and ultradense computer



Scalable supercomputing clusters with ultradense CPU/GPU dual-core processors give oil and gas companies high-performance capabilities with up to 1,000 teraflops of compute power at lower cost. Scalable “building block” architectures make it easier and more cost effective for smaller oil and gas companies to adopt HPC technology.

clusters that once were reserved for only the largest organizations are now accessible to independents of virtually any size to run seismic processing, interpretation and simulation, according to Maria McLaughlin, director of marketing at Appro.

“Oil and gas companies are really pushing high-performance computing,” she says. “The amounts of raw data needed to process oil and gas exploration create a huge demand for computing power. In addition, 3-D visualization analysis data sets have grown a lot in recent years, moving visualization centers from the desktop to GPU clusters.

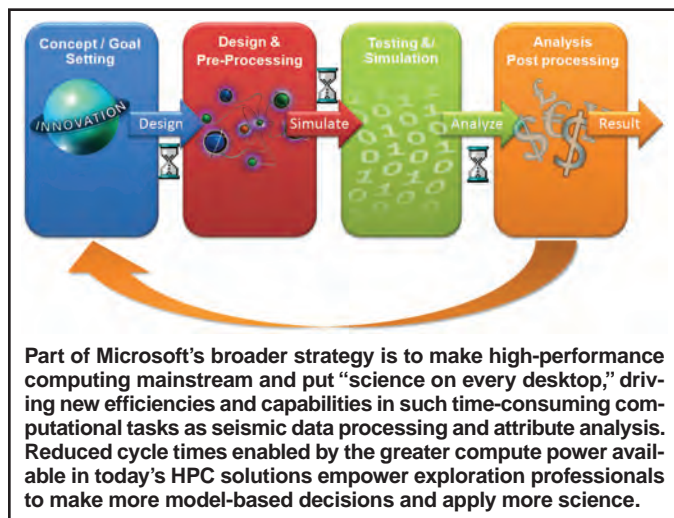
“With the need of performance and memory capacities, Appro’s clusters and supercomputers are ideal architectures, featuring systems that combine many cores with the latest x86 family of CPUs and GPUs based on NVIDIA® Tesla™ computing technologies, to give companies very high levels of performance at a much lower cost,” McLaughlin goes on.

Appro focuses exclusively on HPC with an emphasis on high-speed computations, according to McLaughlin. “The need for HPC is really pervasive,” she states. “Companies can now purchase HPC sized for their applications, making it very cost effective.”

One of Appro’s latest innovations is the new power-efficient and performance-optimized systems launched in the second quarter of 2009. These new platforms provide the building blocks of Appro’s scalable units (SUs) and supercomputing cluster architecture to deliver lower total cost of ownership with a quick path to a return on the investment. Appro’s SU building blocks enable volume purchase of multiple clusters to be deployed as standalone supercomputers for small-, medium- and large-sized HPC deployments across a broad range of vertical markets, McLaughlin explains.

“In this way, the customer can design a system to suit its immediate need and grow it over time as applications expand,” she states.

However, flexible management software is a key component of HPC, including remote server, network management and cluster management to optimize performance and identify prob-



lems, McLaughlin points out. For high-reliability systems, it is also important to have redundant components and network failover so that data are handled without interruption, she relates.

Appro Cluster Engine™ (ACE) management software is designed to provide plug-and-play for the Xtreme-X™ supercomputer series, meaning all network and server hardware is automatically discovered and monitored," McLaughlin says. "With its diskless booting of compute servers, ACE allows 10 to 10,000 compute nodes to boot quickly and simultaneously."

Individual clusters also can be independently configured and provisioned to enable workload and energy management with improved levels of control, she continues. Each application can see an optimized cluster environment, freeing the administrators to manage the global compute environment and energy consumption.

"General-purpose graphics processing units also can be incorporated for additional compute power," McLaughlin explains. "Appro will be launching a new GPGPU-based blade system in the first quarter of 2010. The system will hold up to 10 blades and can be combined with other blade servers into a cluster. Clusters can be combined to make larger systems, building to 'supercomputer' scale."

Enabling flexible system scalability over time makes it easier for smaller oil and gas companies to adopt HPC technology, she adds, stating, "With this approach, you simply build a cluster with all its component parts, management software, and then replicate it. In time, you can grow and scale it according to customers' technical computing requirements and budgets."

Science On The Desktop

With so much data streaming into oil and gas companies' technical departments around the clock, extracting business intelligence to make informed decisions can become a real challenge. To address this issue, Microsoft is taking high-performance computing mainstream, according to Craig Hodges, solutions director of the company's U.S. energy and chemical industry groups.

"Companies today have to do twice the work with half the staff," Hodges says. "Add to that the industry's graying workforce,

and there is a real need for a paradigm shift."


Microsoft's vision is "science on every desktop and HPC access for everyone," Hodges relates. At the 2009 Society of Exploration Geophysicists Annual Meeting in late October in Houston, one of Microsoft's partners supported this stance with a demonstration of 3-D visualization on the System 7™ graphical user interface-based operating system and Xbox® gaming technology.

"Everyone can have access to HPC today," he states. "And with advanced software and hardware solutions, oil and gas technical professionals can have an immersive, game-like experience when developing geologic interpretations."

Another key application—simulation—is supported through Windows® HPC Server 2008™, which can run in a cluster environment and is designed for high availability, reports Mike Sternesky, energy industry market development manager for Microsoft. If a node fails, the system immediately rolls over to the others. System management is greatly simplified, removing complexity that would otherwise deter use, he claims.

"Once the domain of only very large companies, HPC is going mainstream," Sternesky holds. "Low-cost software like Windows HPC Server 2008 enables faster parallelization while making it much easier to develop applications and submit jobs. Enabled by these technologies, companies can create new and powerful applications and leverage them for informed decision making."

In Microsoft's view, HPC should be part of an overall architecture that includes a rich collaborative environment with presence information, video conferencing, knowledge management and integrated telecommunications, Hodges goes on. "The architecture provides the tools that bring science to the desktop and removes the complexity of traditional HPC," he explains. "The collaborative environment facilitates rapid application development and integrates HPC applications into the overall organization." □



At the 2009 Society of Exploration Geophysicists Annual Meeting, Halliburton demonstrated the ability to run sophisticated 3-D visualizations of complex geological and geophysical data using the System 7™ graphical user interface-based operating system and Xbox® gaming technology.