



Next generation computing challenges for pre-stack depth migration and reservoir simulation

Bill Bartling, senior director, energy and sciences, Silicon Graphics provides an analysis of the shortcomings of today's computer systems and discusses where the future lies for computer-intensive operations in the E&P business

Oil and gas computational challenges, especially in seismic processing and reservoir simulation, have consistently outstripped the capabilities of even the most powerful computers. Continuing that trend, these same computational problems are now growing at a rate faster than Moore's Doubling Law. In truth, the very large problems we are trying to solve today have always been around. Adapting to our surroundings, we have scaled them back to match the limited abilities of our machines hoping that engineering breakthroughs would eventually compensate for those limitations.

Many breakthroughs have indeed come to pass, especially in microprocessor speed, power, pricing and form factor. But these revolutions exposed new limitations that in many cases cancelled out some or all of the benefits. Scientists and engineers have rushed to consume the promises of blazing fast gigahertz CPUs at highly competitive prices, only to find that I/O, bus speed, storage systems and code design prevented them from realizing their full potential. Idle cycles became the theme of the day.

The goal in working with the complexity of seismic and reservoir simulation data, is throughput: speed is merely a way to get there. The concept of using speed to deliver breakthrough throughput, and thus to deliver significantly faster computational times for much larger data models, continues to drive research both in computer and software design. The answer must lie in balanced systems, with each component optimized to do its part in delivering the revolution.

The past few years have ushered in a new computational paradigm: the cluster. This model is based on one of the most tantalizing promises of the Internet – a globally interconnected computational grid of inexpensive systems that, together, can potentially combine to create the most powerful supercomputer ever built. The SETI Grid (<http://setiathome.ssl.berkeley.edu/>) is an excellent example of such a system, but, of course, with inter-connects running at 28.8kb/sec in too many places, SETI@home highlights the importance of appropriate and adequate back plane bandwidth.

At the other end of the spectrum are supercomputers, which, in truth, are highly compact computational grids with extraordinarily fast interconnects. In spite of their slower

clock speed microprocessors, from their introduction they outran the fastest PCs due to efficient, stable and proven parallel operating systems and related optimizations that focused on the throughput of many joined processors, not just individual processor speed. Of course, the price point that comes with this sort of machine, while once easily justified, has come under increasing scrutiny from data processors, IT executives and P/L managers.

Key components to the success of supercomputers, in addition to their interconnect speeds, have been effective I/O and data delivery systems to keep the processors fed and working, the lack of which has been a striking shortcoming of PC architectures applied to computationally intensive tasks. Direct attached storage with fibre-channel connections has been the preferred solution to keep the data flowing to processors. New storage and delivery systems crush that paradigm, offering breakthroughs in data delivery via Storage Area Networks.

Real time computing and fusion of dynamic seismic and reservoir models

Even in the face of today's incredible data explosion, many companies are slow to adopt the latest tools. New algorithms for imaging, new software for visualization and interpretation, faster computers to crunch numbers and to display the results in collaborative, panoramic views, and exploration and reservoir models all remain drastically under-constrained. This directly results in nearly all exploration wells being failures and most reservoirs producing only a small percentage of oil in place. The appetite of the earth scientist and the petroleum engineer for new attributes, new views and new computational tools increases with every dry hole and every declining field. The business case for improving even a few percentage points in exploration success or reservoir recovery is staggering, fuelling the drive to real-time computing and to fusion of dynamic seismic and reservoir models.





Data Processing

3D seismic, especially 3D pre-stack depth migration (PSDM), has opened the doors to many oil fields around the world, serving as the keystone to discovery in geologically complex environs. Its concurrent application to existing assets is the engine to enhanced recovery. Especially in mid-life fields, PSDM infuses new life into assets that are otherwise inevitably headed to early obsolescence, transforming portfolio sows' ears into elegant silk purses. PSDM's advanced resolution, precision and accuracy comes via dramatically expanded data volumes and, by extension, substantially faster computational systems and clever software to transform the tens of terabytes of pre-stack data into robust, internally coherent, attribute-rich, easily understood images. Availability of ever-increasing data storage density with equally increasing access speeds feeds a variety of computing architectures, ranging from enormous vector machines to PCs.

Enter next-generation data. Tomorrow's prevalent 4D time-lapse seismic, always-recording passive seismic, on-demand reservoir simulation, real-time drilling data, real-time reservoir monitoring and real-time reservoir management are destined to work together to generate real-time decision systems. Data from such systems will dwarf that derived from today's practices, yet, seemingly paradoxically, the cost of storing, delivering, processing, visualizing and collaborating will drop. More importantly, the business benefits derived will be extraordinary, lengthening field life, increasing recovery factors, reducing operating costs and minimizing drilling failures. Data will be analyzed in collaborative decision support centres featuring always-on collaborative interfaces to share data, knowledge and models with experts, management and operations. These systems will be born of the confluence of extraordinary new data collection systems and powerful new computational systems, driven by businesses craving speed, flexibility, scalability and virtualization. The data volumes generated by these new sources, presented in forms that will deliver ever better business results, will require breakthroughs in computing, data storage/delivery and visualization.

Next-generation computing architectures: formulas for success

Moore's Law alone cannot address these needs. Keys to success will be revolutionary new architectures and systems that adapt to the needs of the business through scalability, flexibility and virtualization. These systems will leverage commodity processors, open source operating systems and software, and standards-based peripherals and data systems.

Formulas: scalability

Today's scalable systems come in many forms, clusters and grids being most popular, having taken the lead from symmetrical multi-processor (SMP) systems of only a few years

ago. Clusters arrived on the scene with many advantages, including very fast commodity microprocessors, inexpensive memory and open source operating systems. With them came a new set of emerging standards that now define peripherals and computing paradigms.

But access to commodity processors was far from free. With it came sometimes-unanticipated and unexpectedly significant new expenses: managing multitudes of inter-operating operating systems, each commanding only two or four processors; limitations of available memory for these processors; and, as a result in extreme cases, wholesale rewriting of applications. The plaintive cry from the field was the wish for the ease and reliability of programming of SMP systems and for the stability of programs running on single system images – if only this could be combined with the promise of the cluster. Scalability to thousands of processors has been established for these clusters. However, efficiency of their use, especially in memory-intensive applications, still leaves much room for improvement.

Enter-next generation computational systems, combining SMP and cluster technologies into a hybrid containing the advantages of both and eliminating the disadvantages of each. This new SuperCluster class of system is now on the market and is based on Intel's Itanium2 processor with cluster node sizes up to 64 processors, running RedHat binary compatible Linux, with single nodes carrying up to 128 GB of commodity memory. Because this system is based on time- and industry-hardened, proven SMP architectures, it features extraordinarily fast interconnects that enable another SuperCluster revolution, one in which the cluster can share memory across nodes. This marks a first in cluster computing and an achievement not matched by even gigabit cluster backbones. The released system currently scales to 512 processors (eight 64-processor cluster nodes), and, remarkably enough, the clustered processor count is doubling every fiscal quarter.

At first glance, these systems offer a distinct administrative advantage by eliminating thousands of operating system instances of classical clusters, minimizing the memory on the system to the total problem size (no need to maximize memory on each cluster node), and simplifying the adaptation of mpi-based software to the cluster. Digging deeper, software development, system administration, system cost, and system maintenance all become simplified and cheaper, while system stability and reliability are greatly enhanced. Recent benchmarking for PSDM and reservoir simulation codes show that this hybrid SuperCluster system out-performs traditional clusters and most SMP machines, especially those above eight processors, even those based on the same Intel Itanium2 microprocessor running standard Linux. When released, this system was labelled the fastest microprocessor-based computer ever built, winning many awards for innovation and performance.




In this architecture, the Linux operating system has been scaled up to a new peak of 64 processors within a single-system image. One of the most beneficial advantages is that the architecture obviates one of the huge disadvantages of the traditional cluster, the fact that each host must be provisioned with the maximum memory that any process may ever need. In addition, in that traditional cluster, there is the need for enough memory to run a copy of the operating system on each host. If a process needs more memory than has been provisioned on a single host, the code must be reworked to spread that load across multiple hosts, which today is only possible on this new system. With more processors per cluster node, a larger pool of memory and processing power is available to each individual process running on that node. For example, a 64-processor system running a single 64-processor Linux kernel carries 256 GB of memory available for a single process to use. As this system expands its processor count to over 1000 processors, the available memory to


each processor will increase to four terabytes.

The new system knits together component 'bricks' via an ultra-fast 6.4gb/sec interconnect backplane, allowing bricks that contain memory-only to be added to the system without the processors that would otherwise be present in those bricks. This allows the system to scale globally shared memory up to multiple terabytes for a single node, all without adding processors. No other Intel/Linux system and only one Unix-based system can do this. The advantages to software development and improved application performance are revolutionary.

This unprecedented capability gives application developers an extraordinarily large system to work with, so they can take full advantage of system resources and concentrate on the performance of their applications without the need to worry about node capacity or configuration limitations. Another major benefit is the ability to share user memory across the entire system with multiple hosts that are all con-



This seismic analysis image depicts four spatially synchronous seismic attribute qualities, along with a well attribute, from which can be derived a geobody, or oil reservoir. Credit: Rock Solid Images utilizing GeoProbe software by Magic Earth.





Data Processing

nected to the same interconnects. With this system, an application can approach work-sharing and parallelism problems with direct-memory access methods that until now have been available only on small SMP systems.

Formulas: flexibility

Today's problem is PSDM, tomorrow's will be reservoir simulation, the following day it will be operations optimization, and, the day after that, the decision model will need updating. The PSDM experts are in London, reservoir simulation experts in Aberdeen, operations optimization is done from the platform, and decision models are updated in Stavanger. Today, this dispersion problem is typically solved by establishing four redundant computer centres, quadrupling the cost of operations while each of the systems operates at less than 100% utilization. But this does not have to be.

The rate of change of business requires that systems have extreme degrees of flexibility. Assets are purchased and divested. Fields change in importance as the price of commodity or cost of capital changes. Bids are solicited for new blocks, and the winning, business-appropriate bid must be constrained by the best data processed in the best way – and processed in time to make a difference. Effective and timely optimization of product mix to meet market dynamics causes impressive improvements in the bottom line, all in real time.

Enter next-generation flexibility. There are three core capabilities for flexible systems: storage, maintenance and delivery of data; distributed and remote access to computational capacity; and distributed graphics, enabling always-on collaboration.

Storage, maintenance and delivery of data

Today's new systems connect data storage systems that are physically distributed across a wide region, up to 8000 km, into a common single-file system that all users on the WAN see as local data. This is connected at LAN speeds, giving each user, regardless of the desktop or server system, access to all data on the virtual data system. Thus, data are created, stored and resident in the location where they were generated, yet are accessible to all stakeholders and authorized users on the network, including partners and consultants. Based on advanced file system technology (CXFS), access is seamless to all users and operating systems, providing file size up to a terabyte and file system size up to a million petabytes. Leveraging fibre-optic connectivity, the file system is accessed at fibre-channel speeds across the WAN.

The big wins from this approach are reduction in storage, ease of backup, simplification of administration, and the biggest win of all, only one concurrent version. No longer do users ask, 'Is this the working version?' There is no longer a need to have multiple copies and no need to transfer data.

Coupled with these regionally distributed file systems are new hierarchical storage management (HSM) systems that rely on the spectrum of storage solutions (spindle speed, capacity, I/O subsystems, etc) to keep often-used data at the fingertips of interpreters while moving infrequently used data to slower, cheaper storage environments. Although the backplane of the HSM is slower, it remains online with speeds, ease of use and maintenance consistent with online data rather than off- or near-line systems. The entire environment operates RAID protocols, ensuring file safety and security, and, with its regional distribution, can operate as a fail-over or mirror.

Distributed and remote access to computational capacity

Grid computing has become most popular in university and research lab circles, but is not quite ready for prime-time industrial applications, where time is money. More affordable bandwidth in the future will change that equation but, for now, distributed and remote computing remains dependent upon moving data to the computer centre. With the new distributed, high-speed data storage and access systems described above, moving data to the computing centre is faster than ever. In time, geographically distributed computational grids of commercial capability will arrive, but will likely find as their backbone the large computing centres based on clusters and SuperClusters. Systems surviving to be part of the distributed grid will be those with architectural scalability and flexibility consistent with the goals of the grid.

Always-on collaboration

Asset teams are distributed according to the assets they manage. Producing oil is inherently a distributed task, and, with so many fields in distant and sometimes environmentally or politically hostile locations, remote management and long-range collaboration become increasingly important. Especially in offshore assets, where the cost of business is so high, consortia of companies requiring frequent interactions between technical and supervisory staff own the assets. The need to connect teams has always been high, but the high cost and low ease of connectivity have been unacceptable, causing coordinating meetings to be few and far between. Even more significantly, such meetings are largely a discussion of conclusions drawn from disparate data sets, necessitating many meetings simply to converge and rectify data sets.

Distributed graphics with always-on collaboration enables teams inside or between companies, be they partners or service providers, to synchronize and rectify data and views frequently, minimizing time spent during decision meetings on basic technical topics. In the production scenario, field offices collaborate with central offices, service providers and partners on such operations as geosteering



drilling wells, managing steam or water injection, and incorporating live data streams into decision models. Bringing the asset or exploration teams together around their common interface, the data model not only streamlines the work process, but also results in better well placement and superior field management.

Formulas: virtualization

Virtualization of the infrastructure is the key to flexibility, where rapid business changes are met with equally rapid changes in access to IT systems. Today, the solution is to put all of the IT equipment into a single large room with raised floors and centralized administration. The gatekeeper to remote access is the network, especially when large data is involved, fundamentally prohibiting any large computational projects. The greatest affordable bandwidth remains FedEx, which inspires most IT managers to adopt data replication/synchronization strategies as a cost-effective compromise.

The energy industry, majors and independents alike, require fundamental work process change to accommodate an exceptionally challenging technical demographic, an increasingly distributed work process, a step-function explosion of new data and demands for ever more competitive financial performance from investors. While existing systems are fast and affordable, they fall short of the high scalability and real-time flexibility that will be the fundamentals of the next generation of information systems.

The business thus presents a new challenge: to seamlessly and efficiently accommodate a highly decentralized workforce while retaining the cost benefits of a centralized company. These contrasting themes must be married in order to support advanced reservoir management and field optimization, along with outsource-heavy data management, computing and analysis. Both exploration and production will benefit from this marriage, simplifying and streamlining data management and access, interlocking the experience of regionally disparate workers, saturating investment and operational decisions with both data and knowledge, and delivering breakthroughs in cost containment and financial performance. Even more importantly, this new way of doing business will take advantage of experience-rich independent consultants and service company outsource businesses.

Virtualization will drive the next generation information system. Unlike many of today's systems, which ask businesses to adapt to the technology, virtualization will adapt to the needs of business. These systems must service the real rate of change of business, which is measured in months. IT departments locked into technologies not readily adaptable to changes in their business will find these obstacles measured on the bottom line.

Businesses which choose inflexible systems will have but

two choices: wait for their business to return to the capabilities of their information systems, or replace their systems. While neither choice offers near-term relief, only the second offers a path forward.

Virtualization is a simple concept. Data, storage, applications, computational servers, graphical servers, people, processes and services are distributed according to distribution of the business. As the business changes, so do the information systems that support them. The challenge is for systems to do all of this yet to adapt, respond and perform, with a single view to the enterprise. With such a system, any user sees a single file system, can access any or all computational resources with high-speed connects, utilizes real-time high-performance visualization, participates in realistic collaboration and can be provided instant service, regardless of location or desktop device.

Catalysts for change

Concepts do not solve business problems: products and solutions do. Some insightful companies working a cross-section of technical computing industries have laid the foundation for next-generation systems over the past several years and have introduced massively scalable Linux SuperClusters that deliver unprecedented computing power for pre-stack depth migration, reservoir simulation and other computationally intense applications. The catalyst to harvesting the bounty of these developments will be equally insightful companies taking advantage of this rich groundwork and applying this computing power to their most technically challenging yet most valuable business problems.

About the author

Bill Bartling is senior director of global energy solutions for SGI, responsible for SGI's strategy and position in the oil and gas industry. Prior to joining SGI in April 2001, he was president and CEO of SciFrame, a software company providing leading-edge collaborative solutions to the energy and pharmaceutical markets. He has also served as manager of technical computing for Occidental Oil and Gas, where he set strategy, architected solutions, and implemented E&P technology to pioneer the use of collaborative volume earth models in large-format visualization centres. Bartling also served as senior vice president, software engineering, for CogniSeis Development, where he oversaw the development of the company's flagship products, Focus, VoxelGeo, and GeoSec. Before joining CogniSeis, he spent 15 years in exploration, production, and research with Chevron Corporation, where he was involved in the use, testing, and design of computer interpretation and analysis products. Bartling holds an M.S. degree in geology from San Diego State University and a B.A. in biology from the University of California at Santa Barbara.

