Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810, www.cabotpartners.com

# BRIEF

## Productive High Performance Computing for Upstream Petroleum with the IBM Cluster Solution Powered by the Intel Xeon Processor 5500 Series

Sponsored by IBM

**Srini Chari, Ph.D., MBA**

**April, 2009**

mailto:chari@cabotpartners.com

### Introduction

Finding oil forty years ago was mainly hit or miss. Explorers drilled ten "dry holes" for every success. Today, these odds have quintupled largely due to the impressive advances in sensors and remote monitoring and visualization, along with high-performance computing technologies and algorithms that have enabled multidimensional seismic imaging and reservoir modeling. Three dimensional (3-D) seismic imaging is now mainstream practice in the industry for exploration. Furthermore, seismic imaging coupled with very accurate reservoir modeling helps engineers better optimize the number and placement of wells in a reservoir or even on an entire basin composed of multiple reservoirs by "virtual drilling" i.e. without actually drilling. This computational capability is immensely valuable especially since it costs about $60M to drill a new deepwater well.

Using IBM System x clusters powered by the latest Intel Xeon processor 5500 series will enable upstream petroleum engineers to make significant improvements in the understanding and the solution of some very challenging problems in exploration and production. For competitive advantage, the industry has developed sophisticated multidimensional seismic imaging and very accurate multi-component, multi-scale, reservoir models to further increase the probability of finding and recovering oil. These innovative computational solutions are more imperative in the years ahead as exploration must increasingly occur in environmentally hostile regions such as in deep water, wide basins with salt mounds, Siberia, and Alaska. This shift has entailed a continuing investment by upstream petroleum engineers, geoscientists, mathematicians, and computer scientists to develop new algorithms and enhance applications for parallel computing environments. This investment will be protected as newer and more powerful cluster systems become available with continuing innovations from IBM and Intel. As before, the payoff will far outweigh this investment as energy costs continue to escalate.

Today's petroleum companies are challenged to continually improve hydrocarbon recovery strategies, predict the nature of the fields, rapidly bridge the gap between technical and business decision making, and eliminate system engineering delays and guesswork. This requires a robust high-performance environment – powerful and reliable cluster solutions from IBM are designed to provide exactly that.

IBM System x cluster[1] solutions powered by the Intel Xeon processor 5500[2] series are delivered along with leading reservoir modeling and seismic imaging applications through both direct and business partner channels. These integrated and optimized solutions deliver significant value differentiators with up to three times the performance over the previous generation of cluster systems, improved productivity, and lower energy and total cost of ownership. Oil and gas companies will be able to deploy and use these high performance applications more easily and significantly enhance productivity.

---

[1] IBM System x and BladeCenter HPC cluster solutions,
http://www-03.ibm.com/systems/x/solutions/infrastructure/departmentalclusters/

[2] Intel Xeon Processor 5000 Sequence Based Server Platforms Performance summary,
http://www.intel.com/performance/server/xeon/summary.htm?iid=perf_server_lhn+dp_sum

## Petroleum Exploration and Production Computing Trends and Challenges

Seismic imaging and reservoir simulation are routinely used in conjunction with sensor information of fluid and rock properties at given locations to significantly enhance the success of finding oil and to optimize its recovery. More accurate imaging and modeling techniques increase the probablility of hydrocarbon discovery and exploitation, extend the life of current reservoirs, improve the risk-reward ratio of locating new hydrocarbon sources, and reduce the overall cost of exploration and production. The availability of flexible and cost-effective cluster computing solutions with optimized versions of major upstream petroleum applications has further energized increased adoption among smaller oil and gas service providers who provide high value added services to the petroleum industry.

Seismic imaging

As exploration increasingly occurs in environmentally hostile regions with complex geology on wider basins and much deeper in the ground, seismic imaging problems have become very large and information needs to be aggregated from billions of points from seismic surveys. Various algorithms – migration types - have been used to solve these wave equations that describe the passage of sound waves through rock. In the past, poststack migration was very common. Prestack migration takes about 100 times longer than poststack migration but can exploit cluster computers more effectively. Prestack migration retains amplitude variation with offset and phase changes which are very useful for subsequent analysis. Production specialists increasingly use several algorithms concurrently to analyze the data, which drives up the computing requirements even further. But these concurrent analyses are very effective on clusters as each analysis can be done on a separate partition. Also – as the price of oil continues to escalate - revisiting existing reservoir basins and analyzing older seismic data with newer computational algorithms can reveal potential opportunities for production.

Reservoir simulation

Reservoir simulation involves tackling a wide range of complexity in modeling flow physics in or around complex shapes and geometries. Sophisticated reservoir simulation applications integrate knowledge gained from geophysical, geological, and drilling and production operations, and predict dynamic reservoir behavior under a wide range of conditions. Dynamic Data Driven Applications and Simulations (DDDAS) is becoming a new paradigm of computing, one involving symbiotic feedback in which reservoir simulation and sensor field data interact in real time to dramatically improve the fidelity of the analysis, its accuracy, and reliability. Integrating reservoir modeling with 4-D sensor data enables the ongoing management and optimization of large reservoir basins and the production in traditionally difficult terrains.

Challenges and opportunities

Many production reservoir simulation and seismic imaging applications have underlying numerical algorithms that scale very well on large clusters, and are being increasingly used by reservoir engineers and asset management teams to make well-informed and speedy decisions. However, high performance computing (HPC) application performance is not solely tied to parallel capability. It is also closely related to file I/O throughput and to other components of the computing environment. Conventional “white box” computing clusters just can not offer the integrated high-performance I/O solution and architectural flexibility often needed for production seismic analysis and reservoir simulation.

Furthermore, white box systems simply cannot meet all the business objectives for smaller oil and gas organizations that are often limited in skills and resources to manage and deploy distributed computing solutions. Specifically, they fail when it comes to working to reduce total cost of cluster ownership, reverse the challenges of managing disparate servers, execute numerous applications concurrently, and achieve more flexible resource management. To reduce the total cost of ownership, a flexible, balanced cluster environment with an optimized implementation of major upstream petroleum applications is needed.

## IBM Cluster Solutions Tailored to Maximize Total Value of Ownership

IBM has developed a portfolio of robust cluster servers built on IBM X-Architecture that enable a customer to jumpstart the production of applications or products faster than the competition. This is realized through research, experience, innovation, and support, delivered and proven over the years.

The latest Intel Xeon processor 5500 series incorporates dynamic scalability, Turbo Boost technology, improvised Hyper-Threading, and scalable shared memory based on QuickPath technology for up to 3.5 times the peak performance over the previous generation Intel micro-architecture. IBM HPC clusters can integrate new IBM System x servers using the Intel enhancements to provide strong new levels of performance and functionality.

Enhanced by the capabilities of intelligent management tools, leadership offerings in energy-efficiency, consolidation and virtualization, security, and remote maintenance, the data center gets the much needed optimization and cost reduction that is critical for business growth. Multiple layers of redundancy, memory protection, high availability tools, in System x servers, are a part of the IBM system design.

HPC offerings in System x from IBM consist of internet-scale computing iDataPlex, BladeCenter, Cluster 1350, and Storage servers and can include integrated systems powered by the Intel Xeon processor 5500 series.

Intel Xeon processor 5500 series

This processor architecture design draws on the benefits of hafnium-based Intel 45nm high-k metal gate silicon technology, for parallel processing performance. Intel's QuickPath Interconnect (QPI) delivers substantial increase in bandwidth from a scalable shared memory by incorporating an integrated DDR3 memory controller onto the processor die in lieu of the previously bottlenecking Front-Side Bus (FSB). The processor can run two threads per core simultaneously with Intel Hyper-Threading technology which equates to 8 threads per 2 socket server. Multi-level shared cache reduces latency to frequently used data thereby improving performance and efficiency significantly.

Intel Turbo Boost Technology increases performance of both multi-threaded and single threaded workloads. This technology is activated when the Operating System requests the highest processor performance state. The maximum frequency of Intel Turbo Boost Technology is dependent on the number of active cores. While the length of time the processor spends in the Intel Turbo Boost Technology state depends on the workload and operating environment, it provides the performance a user needs, when and where the user or application needs it.

For a given workload, the number of active cores, estimated electricity consumption, or processor temperature can set the upper limit of Intel Turbo Boost technology. When the processor is operating below these limits and the user's workload demands additional performance, the processor frequency will dynamically increase by 133 MHz on short and regular intervals until the upper limit is met or the maximum possible upside for the number of active cores is reached. Conversely, when any of the limits are reached or exceeded, the processor frequency will automatically decrease by 133 MHz until the processor is again operating within its limits with the stated frequency as a lower bound. For example, a database query, which doesn't use any of the processor's floating point silicon can take advantage of that thermal headroom and increase the frequency of all 4 cores.

**Exceptional performance improvement for petroleum applications with the Intel Xeon processor 5500 series**

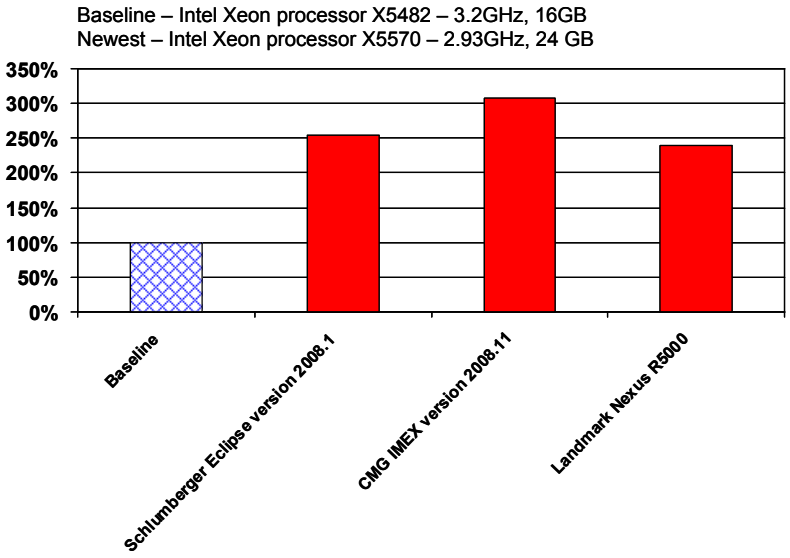


- All results shown have Hyper-Threading and Turbo Boost disabled
- Data obtained from the Intel Corporation http://www.intel.com/performance/server/xeon/hpc.htm

Early benchmarks on the newest Intel Xeon processor 5500 series for a wide range of upstream petroleum applications record up to three times the performance over the previous generation Intel architecture. Memory bandwidth sensitive algorithms like reservoir simulation, reverse time and wave equation migration benefit most from the newest Intel Xeon processor 5500 series. This is particularly beneficial for petroleum companies as these applications are often the ones that are very accurate, stress the IT infrastructure, and are most bottlenecked by CPU performance.

IBM System x iDataPlex

The System x iDataPlex internet-scale computing server solution from IBM is uniquely positioned to help enterprise clients overcome compute density and energy efficiency constraints. The iDataPlex supports massive scale-out data centers and high performance computing solutions.

The iDataPlex is a half-depth server solution, optimized both mechanically and component-wise for maximum efficiency with power and cooling. It is an industry-standard based server platform designed to minimize utilization of data center floor space (doubles the compute density/square foot compared to a 1U rack), and power and cooling infrastructure (20% less cooling and 63% less fan power). An easily maintainable solution with individually serviceable servers, front access hard drives/cabling, and common tools across the System x portfolio for management at the node, rack, or data center level, it is configurable for customer-specific compute, storage, or I/O needs and delivered pre-configured for rapid deployment.

Large scale reservoir simulation problems with over 3-5 million unknowns scale very well on the iDataPlex. Interdisciplinary problems with multiple physics and/or complex, coupled, problems in reservoir optimization over an entire basin and parametric studies iterating over thousands of input parameters perform extremely well on the high performance scalable iDataPlex system. Also, hundreds of

parametric simulations can be done concurrently on multiple iDataPlex partitions. This significantly reduces the time to results for large-scale reservoir optimization studies.

IBM BladeCenter

By integrating servers, storage and networking, IBM BladeCenter is helping oil and gas companies sweep aside complexity. Its wide application solution blades, truly efficient chassis, and open design, are packed into an answer to today's data center challenge of straggling racks and overheated server rooms.

IBM System x Cluster 1350

IBM System x Cluster 1350 reduces deployment time for Linux clusters and Windows clusters, benefits from power and cooling through IBM Systems innovation, and offers integrated global hardware support. Oil and gas clients can speed up installation of an HPC cluster, simplify its management and support, and reduce mean time to payback.

IBM Storage servers

IBM storage servers come at versatile ranges of entry level disk storage servers, mid-range disk storage systems—the IBM TotalStorage Enterprise Storage Server (ESS) and the IBM TotalStorage DS4000 series, the Network attached storage or N series products, and IBM's TotalStorage SAN solutions.

Innovative software environment

Together IBM and Intel drive toward simplicity in the software design to take advantage of system features that deliver high performance at consistent reliability and security. The programming and administration environment is based on familiar programming languages, libraries, job management tools, and parallel file systems. Reservoir engineers greatly benefit from these innovative software components without facing a steep learning curve.

IBM System x clusters are highly optimized to reduce the deployment time for Linux clusters and Windows clusters. A Linux HPC Cluster is equipped with computing power at low-cost, collective intelligence of open standards, portability, flexibility and high availability. Likewise, the Windows HPC Server 2008 on IBM System x clusters and the BladeCenter servers provide reliable, cost-effective HPC cluster solutions. Windows HPC Server 2008 is designed to simplify the deployment, administration and management of the client's entire system with scalable performance, an easy-to-use scheduler and a new management interface that facilitates a familiar Windows environment so that supercomputing is more accessible.

Additional HPC cluster software available on IBM clusters includes the General Parallel File System (GPFS) for Linux and Tivoli Workload Scheduler LoadLeveler for Linux. GPFS is the top performing cluster-wide file system providing superior scalability and high reliability. Tivoli Workload Scheduler LoadLeveler is a job scheduler designed to maximize resource utilization and job throughput to get the most out of the available resources. IBM clusters also support other workload management solutions available from partners. These integrated software tools sustain highly productive environments with thousands of server nodes running large parallel computational workloads typical in seismic imaging and reservoir simulation.

IBM differentiation for reservoir simulation and seismic imaging

IBM Cluster Solutions powered by the Intel Xeon processor 5500 series are delivered jointly with leading application providers in order to leverage new and existing IBM channels – both direct and through business partners. These certified, optimized, and customizable solutions aggregate and optimize a portfolio of

hardware, software, and services components from IBM, Intel, Microsoft, major reservoir simulation and seismic imaging application providers, and other business partners. Some of the key IBM and Intel value differentiators for petroleum E&P companies are to:

- Increase the total value of ownership and improve asset decisions
- Offer access to pre-engineered and optimized cluster configurations based on sizing expertise for a range of leading reservoir simulation and seismic imaging application solutions
- Reduce risk, time, and costs associated with cluster installation and deployment
- Provide deep HPC expertise to enable, significantly improve, and optimize applications workloads
- Deliver an affordable, scalable, and flexible environment that can be tailored to a customer-specific environment with superior technology components from
    - Intel – Multi-core Xeon processor 5500 series
    - The open source community – Linux, Message Passing Interface (MPI)
    - Microsoft – Windows HPC Server 2008
    - IBM unique technology components include:
        - Enhanced networking connectivity with power accelerator technologies
        - Superior thermal and energy management with Cool Blue
        - Systems Management for clusters
        - High performance I/O systems – IBM General Parallel File System (GPFS) and Memory mapped IO (MIO)

The IBM Cluster solution including iDataPlex with several hundreds of multi-core Intel Xeon nodes has been deployed for seismic imaging and reservoir simulation at oil and gas companies. These customers and others can now gain a large performance boost from this new generation of IBM System x clusters powered by the Intel Xeon processor 5500 series. IBM provides access to a wide range of HPC systems through the Computing on Demand (CoD) cloud center. Oil and gas applications developers and prospective users can test out and get in-depth experience on these systems in a cost-effective manner.

Using an IBM Cluster solution maximizes the total value of ownership for production upstream petroleum environments and enables faster time to results and delivers substantial returns on investment for oil and gas companies.