PEOPLE WHO GET IT

HIGH-PERFORMANCE COMPUTING IN OIL AND GAS

As energy exploration becomes increasingly challenging, oil and gas firms deploy ever more powerful computing and storage solutions to stay ahead.

Executive Summary

With more than 70 percent of the world's oil and gas production coming from mature fields, energy companies face a growing challenge with extracting natural resources under increasingly difficult conditions. To adapt, they are applying sophisticated survey techniques and cutting-edge science and technology to glean the location and character of gas and petroleum deposits.

Working from highly sophisticated seismic surveys, algorithms, accelerators, models and more, energy firms use high-performance computing (HPC) to process, analyze and visualize data. The amount of data is extreme. In fact, HPC data centers routinely host and manage petabytes (quadrillions of bytes) of data, with even more massive jobs on the horizon.

The task of processing and interpreting all this data falls to powerful, high-performance servers, workstations and other types of supercomputers. By understanding the challenge faced by the oil and gas sector, and applying the latest technology advancements to decision-making, companies can improve their opportunities for success.

SHARE THIS WHITE PAPER

Table of Contents

The Energy Challenge

HPC is defined by incredibly demanding missions. Yet, even when compared with daunting computing challenges posed by weather simulation, weapons research and genetic mapping, the hurdles posed by energy exploration are both unique and extreme.

Driven to discover fields of oil and natural gas that are increasingly more difficult to detect and access, companies in the energy sector must place ever-larger bets on the fields they lease and the horizontal well shafts they drill today. To manage these bets, which can cost billions of dollars, energy companies turn to HPC servers and workstations to capture, process and interpret the vast streams of data produced by detailed seismic surveys, simulations, models and more.

It takes serious hardware and software to leverage these data assets. Powerful, massively parallel server clusters, linked to advanced, multitier storage systems, drive the cutting-edge software used to process, interpret and visualize subsurface structures.

HPC systems address the challenge by bringing to bear unprecedented computing resources. Companies deploy racks of massively parallel servers, often with hundreds or even thousands of CPUs packaged in modular blades and attached to shared networked storage over fast InfiniBand or 10 Gigabit Ethernet links.

Intensive computing operations are streamlined by advanced accelerators or graphics processing units (GPUs), which are optimized to sprint through specific mathematical tasks at speeds that are orders of magnitude faster than general-purpose CPUs.

These arrays of machines deliver performance measured in teraflops (trillions of floating point operations per second) or even petaflops (quadrillions of floating point operations per second). By way of comparison, the most powerful supercomputer in the world — according to the November 2013 TOP500 rankings — is China's Tianhe-2, which produced 33.86 petaflops as measured by the Linpack Benchmark.

High-stakes Exploration

Companies in the energy industry operate under uniquely demanding, competitive conditions where speed is paramount. Analysis that takes months to complete can risk running past the lease term of prospective well sites. Further, delays in analysis at existing sites can produce delays that idle both equipment and personnel, thereby increasing operating costs.

In this high-stakes environment, where energy companies must place informed bets on the presence and character of prospective oil and gas deposits, the ultimate leverage is information. Energy companies conduct detailed, seismic surveys, deploying thousands of arrayed sensors to capture precise data about the seismic waves reflected back to them through subsurface structures. These surveys allow companies to bolster the efficiency of fracking and the unconventional horizontal drilling methods used today.

It is one thing to survey a field and capture the acoustic and seismic data that comes back. It's quite another to transform that raw data into actionable insight that enables oil and gas companies to make good decisions about where and how to drill.

The task of processing and interpreting survey data is incredibly complex, requiring the use of cutting-edge computing hardware and proprietary software. As a result, energy firms are as much engaged in cutting-edge geophysics and technical computing as they are in industrial drilling and production.

In fact, the oil and gas sector is uniquely dependent on high-performance computing, according to research data presented by IDC at the 2014 Oil and Gas HPC Workshop at Rice University in Houston. Overall, sales of HPC servers across all industry segments grew at a rate of 7.7 percent, to $11.1 billion from 2011 to 2012, and are expected to top $15 billion by 2017. In the oil and gas sector, sales of HPC servers are projected to grow even more rapidly, from $747 million in 2012 to $1.16 billion in 2017.

IDC projects that geosciences will be among the fastest-growing sectors tracked in its surveys. Figures from HPC analyst firm Intersect360 Research also project growth, showing that sales of all HPC products and services to the oil and gas sector is projected to grow between 6.5 percent and 7 percent in 2014.

HPC in Oil and Gas Industries

According to Halliburton, more than 70 percent of the world's current oil and gas production comes from mature fields, many of which are in the secondary or tertiary phases of production. The easy oil and gas has already been extracted, and sophisticated discovery and extraction techniques are needed to extend production. In many cases, raw survey data reanalysis and in-depth scenario analysis are performed to better understand the underlying structure and dynamics of aging fields.

To do so, companies rely heavily on complex, proprietary algorithms and code for more efficient exploration, producing higher-quality images that enable precise interpretation of subsurface data. Each generation of these algorithms yields increasingly effective interpretation, allowing exploration teams to locate otherwise hidden deposits and to extract a greater portion of the available oil and gas.

However, new algorithms often demand significantly more processing power and responsive storage, pushing the limits of HPC platforms. Oil and gas companies must process data at higher speeds to create models of ever-increasing fidelity.

That demand is clearly visible in the market. According to IDC, supercomputer deployments are dominating growth in the HPC sector. While sales of all HPC servers grew 7.7 percent from 2011 to 2012, sales of supercomputers over the same period grew by 30 percent. Supercomputer revenues accounted for $5.6 billion of the total $11 billion HPC server market in 2012. The oil and gas industry, along with the financial sector, is the premier commercial consumer of supercomputing systems and solutions.

Oil, Gas and Data

One area that sets the oil and gas sector apart is its consumption of data. Success in energy exploration ultimately boils down to the quality — and quantity — of raw data. The greater the volume and fidelity of the data, the better the chances that a targeted drill site or prospective deposit will pay out.

Companies have rapidly improved survey techniques and technologies, deploying higher numbers of advanced sensors to gain the most accurate possible picture of the underlying geology. Techniques such as 3D imaging and Wide-Azimuth (WAZ) surveys have multiplied the amount of data companies must capture and process.

Today, it is not unusual for a large exploration project to produce multiple petabytes of raw data, and for energy companies to possess a total data portfolio of tens or even hundreds of petabytes. Moving, processing and interpreting these vast stores of information have put oil and gas firms firmly at the leading edge of Big Data analytics and operations.

Spending on storage overall is growing fast — faster, in fact, than any other technology area at HPC sites, according to an Intersect360 survey. And IDC projects that HPC storage revenue will grow at an annual rate of 8.2 percent from 2012 to 2017, to $6 billion.

Big Data analytics are a target of investment across all sectors, according to the 2013 *High Performance Data Analysis Report* from IDC. It found that 67 percent of sites surveyed perform Big Data analysis on HPC systems, with an average of nearly one-third of available computing cycles devoted to the task.

Spending on high-performance data analysis (HPDA) servers will also continue to grow through 2017, according to IDC, from $743.8 million in 2012 to nearly $1.4 billion in 2017. HPDA storage, meanwhile, will approach $1 billion in revenue by 2017.

In the oil and gas sector, HPC budgets are shifting to account for Big Data and HPDA. An IDC survey, *HPC Market Update, HPC Trends In the Oil/Gas Sector and IDC's Top 10 Predictions for 2014,* found that companies in the oil and gas sector are budgeting an average of 12 percent of all HPC spending to support Big Data analytics — more than is spent on storage hardware, middleware or server maintenance. Only server hardware (30 percent) and application software (about 15 percent) consumed a greater share of budgeted spending.

The Human Resource

In this complex and competitive environment, energy companies face a stark challenge in hiring and retaining the highly skilled programmers, geoscientists, analysts and interpreters who are the backbone of any exploration effort. A Schlumberger Business Consulting study found that about 22,000 geoscientists and engineers will exit the workforce by 2015, as a generation of geoscientists and petroleum engineers hired before the mid-1980s reaches retirement age. Their departure is creating urgency in an industry in which leading companies have acknowledged project delays due to a lack of skilled employees.

The challenge is acute in the arena of high-performance computing or HPC-driven oil and gas exploration, where skilled programmers and scientists are needed to update existing efforts and develop new algorithms and computer code. Skilled developers are needed to update code assets to take advantage of the new capabilities of HPC systems.

For example, an HPC data center outfitted with server clusters incorporating NVIDIA Tesla graphics processing unit or GPU co-processors may not return full value in terms of computing performance unless existing algorithms and code have been tuned to work with NVIDIA's CUDA parallel computing architecture.

Even absent new algorithm and software development, maintenance tasks are critically important to oil and gas companies. Addison Snell, CEO at Intersect360 Research, says the problem is growing more acute, as companies struggle to find technical experts who operate at "the intersection of geoscience and computer science."

Consulting firm Bain in a 2013 brief singled out staffing and retention of technical personnel as a key factor for competitive success in the energy sector. With specialist skills in historically high demand ("more valued today than at any other time in the past two decades," according to the brief), Bain urged energy firms to develop clear staffing strategies to ensure success. Increasingly, the brief notes, large oil and gas companies are working to develop robust, differentiated core capabilities that will fuel ongoing growth.

Seismic Processing and Storage Infrastructure

Oil and gas companies acquire data by deploying arrays of sensors that record the reflection of seismic waves produced by sources such as high explosives, air guns, thumper trucks and seismic vibrators. Marine exploration of offshore deposits is often conducted by dragging arrays of sensors behind a boat, and firing air guns into the sea floor to produce acoustic reflections.

The amount of raw data captured is spiraling upward as high-resolution sensors become more numerous and more capable. WAZ surveys place numerous sensors across a wide area to develop effective imaging.

Offshore Surveying Using Towed Sensors

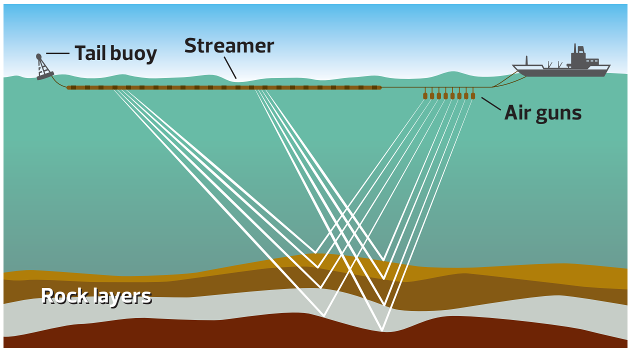

Sensor data is captured on tape or streamed to a mobile, onsite data center, which includes HPC servers and disk arrays. High availability and reliability are critical for data capture in the field, where even a brief outage or degradation can result in a loss of data that impairs analysis.

Traditionally, seismic data is offloaded from disk or tape copies to the corporate data center once the survey is complete. To save time, some companies are looking to preprocess data in the field, deploying HPC assets at the exploration site to perform quality control and achieve faster analysis.

Seismic processing continues at the corporate data center, consuming prodigious amounts of computing power across extremely large data sets. Sophisticated algorithms and transformations are executed against the data to provide velocity information, increase resolution and accurately locate the position of detected subsurface features.

The challenge is compounded by increasingly detailed survey techniques, which produce high-resolution, 3D and 4D (time-based) data sets that can be orders of magnitude larger than less sophisticated two-dimensional surveys. High-resolution arrays and higher density configurations combine to capture a fuller sampling of each seismic wave. Even a small-scale exploration program can generate terabytes of raw 3D data requiring substantial computing power.

The data challenge in oil and gas occurs at scales both large and small. Individual core samples, which are wrapped and secured before being pulled from drill sites, are scanned to produce highly detailed images of subsurface material. Each scanned cubic centimeter of a 150- or 200-foot-long core sample can produce 100 gigabytes of data, yielding vast amounts of information that must be stored, processed and managed.

The mixed nature of files managed by HPC systems in oil and gas exploration poses a challenge. Seismic data consists both of extremely large files and very large numbers of tiny files containing metadata.

Processing applications can also produce large volumes of small files. The result: Storage build-outs tuned to support sequential input-output (I/O) of large data sets can struggle with the mixed I/O environment. To manage this, storage systems employ powerful parallel file systems able to juggle multiple operations simultaneously, as well as sophisticated, tiered caching that increasingly leverages solid-state disks (SSDs) to reduce latency.

What has become clear is that seismic processing and analysis in oil and gas increasingly rely on the performance of storage infrastructure, which is advancing at a slower pace than that of silicon-based CPUs and GPUs.

For oil and gas companies engaged in exploration, storage hardware and file systems must be able to deliver gigabytes or even tens of gigabytes per second (Gbps) of throughput and hundreds of thousands or even millions of input/output operations per second (IOPS). Streamlined I/O is critical to efficient operation, as very large data sets must be made available to system memory and processing units (CPU and GPU/accelerators).

The Spider II file storage system, deployed at Oak Ridge National Laboratory, offers a case in point. The system is designed with 40 petabytes of raw capacity and is capable of ingesting, storing, processing and distributing research data at over a terabyte per second in a single Lustre File System.

Inside Seismic Processing

Seismic processing is a highly complex and data-dependent process that ultimately defines success or failure for oil and gas companies. Among the key goals, seismic processing and exploration seek to achieve the following:

- **Create models of ever-increasing fidelity:** Seismic processing is an iterative process. Earth models are applied to data to help synthesize clearer understanding of subsurface structures. That understanding helps inform more polished and detailed earth models, which are then applied to the practice of seismic processing, driving the effective interpretation of seismic data.
- **Arm companies with accurate information:** Data is at the heart of every oil and gas field project. The latest survey techniques have vastly increased the amount and fidelity of seismic data. They can also provide glimpses at subsurface structures that in the past remained hidden behind obstructing features such as salt domes.
- **Provide insight on where to drill and how much to bid:** Oil and gas companies are in the business of placing bets — huge, multimillion-dollar bets that can make or break a company. Energy firms rely on survey, modeling and computing techniques and technology to arm them with the information they need to make those bets pay off.
- **Make decisions fast for best ROI and profitability:** Rapid turnaround on survey analysis is critical when leases are set to expire or when seasonal weather threatens to hinder offshore and remote onshore drilling operations. HPC servers and workstations can shave valuable days or weeks off decision times and reduce operating costs.

Seismic Processing and Big Data

The challenge of seismic processing in the energy sector can be articulated in two words: Big Data. Big Data describes data sets that are so large and complex that they defy management with traditional relational database tools and software.

Big Data is often described in terms of three Vs: volume (the amount of data, or capacity), velocity (the flow of data in and out of systems, as well as changes that occur to data) and variety (the types of data, such as tables, unstructured files and large media objects). It also addresses the unique challenges of searching, manipulating, analyzing and visualizing these massive and changing stores of information.

In this regard, oil and gas companies conducting seismic processing are at the very front lines of the Big Data revolution. The volumes of data that energy exploration firms produce, manage and consume are extreme. Survey operations produce enormous, real-time streams of data, and significant costs can be incurred if that flow is interrupted.

Storage systems must support processing and visualization of all this data — at times millions of input/output operations per second (IOPS) and tens of gigabits per second of throughput. The performance demands placed on storage infrastructure and data-dependent applications and systems will only increase as the volume and richness of seismic data grows.

Managing and leveraging available data is another challenge, especially as oil and gas operations are increasingly instrumented. Real-time data captured, from drill sites, can be combined with information culled from scanned core samples to produce valuable insight beyond what is available from acoustic survey data. Combining these data flows, however, is a complex endeavor that requires thorough integration of software and systems.

The storage challenge continues as the workflow moves from processing to interpretation and analysis. Here, central and distributed storage must be marshaled to feed data to hundreds or even thousands of individual Windows- and Linux-based workstations. Throughput to these clients is measured in gigabytes per second, and loads are automatically balanced to ensure responsiveness.

Pushing the Boundaries of Scientific Discovery

The energy industry is a leading consumer of supercomputing resources in the commercial space, equaled only by the financial sector in its investment in cutting-edge processing, storage and network capabilities. Furthermore, a 2011 study by Intersect360 Research found that the oil and gas sector has the highest overall rate of internal software usage; companies in the energy sector produce and maintain more in-house applications and algorithms than any other commercial sector.

New techniques and algorithms, such as full-wavefield inversion (FWI), are enabled by increasingly powerful computing resources as well as more advanced and extensive sensor data capture. Today, HPC build-outs rely on powerful co-processors — such as the Intel Xeon Phi accelerator and NVIDIA Tesla GPU — to offload specific, mathematically intensive tasks, which then can be resolved far more rapidly than they could on a general-purpose CPU.

The importance of co-processing in HPC is reflected in 2013 market figures published by IDC. According to IDC's *Processors/Co-Processors/Accelerators Report* (one of six reports in the *IDC Worldwide HPC End-User Study*), the percentage of HPC sites employing co-processors or accelerators in HPC systems increased to 76.9 percent in 2013, up from 28.2 percent in 2011.

Intel Xeon Phi co-processors and NVIDIA GPUs were identified as the most widely deployed co-processors. The massive increases in the amount of data that needs to be processed, analyzed, stored and updated are driving these advances in computing power. Similarly, the advanced algorithms needed to handle this high-volume and high-velocity data place increased pressure on storage infrastructure.

The stakes continue to rise in the energy sector, as emerging techniques are developed to produce imagery of subsurface structures at ever higher fidelity. Time and depth migration techniques, for instance, take raw data gathered by acoustic sensors and apply complex mathematics to pinpoint the location of the subsurface features they detect. As computing and storage capabilities have improved, energy companies have moved to adopt more effective — and computing intensive — techniques for capturing subsurface images.

Among the techniques being used:

- **Kirchhoff migration:** A long-established acoustic method that traces rays and tries to infer travel times. In general, Kirchhoff migrations work best where the geology is layered, not complex.
- **Wave-equation migration:** A depth migration technique that handles multiple paths during wavefield extrapolation, especially for regions with salt bodies. It is superior to Kirchhoff migration in its ability to deliver accurate imaging.
- **Reverse-time migration (RTM):** Employed since 2009 on GPUs, RTM takes two passes through captured data, simulating the behavior of waves propagating both downward and upward through the earth. Complex wave models enable correlation between the two passes to yield a clearer image that reveals subsurface structures that would otherwise remain hidden.

- **Full-wavefield inversion:** A technique that develops high-resolution models of seismic data by iteratively comparing observed and modeled seismic waveforms. The repetitious FWI method has become affordable with the advent of GPU and co-processor acceleration.

The evolution of increasingly effective seismic imaging techniques is driven by gains in computing power. In a presentation at the 2014 Oil and Gas HPC Workshop at Rice University in Houston, Peter Breunig, general manager of Technology Management and Architecture at Chevron, showed the processing power required to drive contemporary data analytics in oil and gas.

In 2002, Kirchhoff migrations required roughly 1 teraflop of computing performance. Just two years later, that figure had leapt with the adoption of wave-equation migration. By 2010, two-pass, reverse-time migration was demanding 150 teraflops of performance. The emergence of acoustic, full-wavefield inversion represents another massive increase, demanding roughly 1.5 petaflops of processing power.

Computing Performance Required for Migration Techniques (in teraflops)

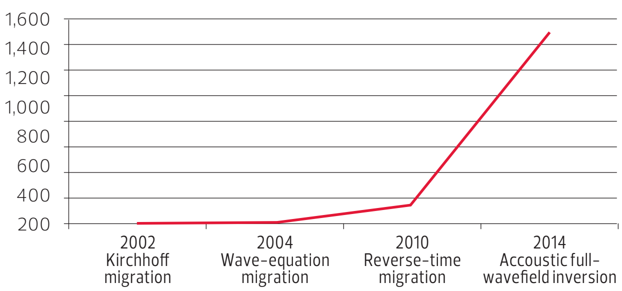

Source: Chevron Information Technology, March 2014

Beyond FWI, energy companies are looking toward the adoption of elastic migration, which simulates the subtle deformations of the earth as waves propagate through the subsurface structure. Further downfield, viscoelastic migrations, that account for the properties of rock in the survey, will again multiply the data and processing load.

Increases in computing power have also allowed oil and gas companies to expand the use of pre-stack migrations in analyzing seismic data. Advanced processing techniques such as pre-stack depth migration (PSDM) are particularly useful in resolving detailed, seismic imaging in areas of complex geology, such as salt domes, and allow interpreters to make much better judgments about prospective deposits. This capability comes at the cost of greatly increased stress on every component of the storage infrastructure.

The data-handling demands are extreme. Pre-stack migrations require working with every trace in a survey — often numbering more than one million pre-stack traces per square mile. By contrast, the number of traces can be 50 or even 100 times lower in post-stack migrations, which rely on shot data that has been stacked and combined to filter out noise and discern key geologic features.

Going Green: Power and Cooling

The vast growth of computing power and storage capacity has driven similar increases in the need for electrical power and cooling. With HPC deployments in oil and gas often accounting for thousands of computing nodes, and each of those often comprising multiple multi-core processors, the consumption of electricity and production of heat is massive.

Energy costs are a significant issue. Globally, more companies are using more electricity to power data centers than ever before. Data centers worldwide consumed 38 gigawatts of power in 2012, up 63 percent from 24GW in 2011, according to the DatacenterDynamics 2012 Global Census on data center trends. The rising energy demands pose a challenge for HPC data centers, which need a robust and reliable energy source.

Keeping systems cool is a serious challenge. Increasingly, vendors are turning to water-cooled solutions that are more efficient at moving waste heat away from system components. The IBM SuperMUC system, among the fastest supercomputers in the world, is cooled entirely by water piped through its structure. Thermally efficient hardware is a critical asset in HPC deployments. In the case of storage systems, which consume an increasing share of data center floorspace, densely configured storage can reduce cooling and power consumption by reducing the number of heat-producing storage controllers.

Gas and Oil: Sizing Up to HPC

While the HPC technology used to drive oil and gas exploration changes rapidly, the core concepts do not. IT management must align around highly scalable supercomputing architectures that can grow with demand.

Scale-out systems that allow organizations to add racks — or to add more hardware to existing racks — are important for keeping pace with fast-growing workloads. Reliability, manageability and flexibility are every bit as important as raw performance when it comes to cost-effective operation.

Notably, HPC deployments rely heavily on commercial components, which provide outstanding performance at far lower acquisition and management expense than those presented by specialty CPUs, GPUs and storage systems.

Likewise, industry standard InfiniBand and 10 Gigabit Ethernet links enable cost-effective, high-performance data access and I/O for storage systems.

Intersect360 Research found that nearly three-quarters of all deployed HPC systems are based on x86-compatible processors, with Intel processors appearing in nearly 80 percent of all new HPC systems shipped (AMD accounts for the remainder of x86 sales).

Among the things to consider in deploying an HPC build-out:

- Scalable supercomputing architectures that offer room for expansion
- High-density computing clusters, built on racks outfitted with multiple blade servers housed in standard-height (1U, 2U, 4U) enclosures and optimized both for top-end I/O and throughput, as well as for energy efficiency and cooling
- Powerful and energy-efficient, multi-core CPUs, such as Intel Xeon processors, supporting as many as 8 cores and 16 threads per socket
- Optimized on-board co-processors, such as Intel Xeon Phi accelerator chips and NVIDIA Tesla GPUs, to accelerate specific computation by an order of magnitude or more
- Massively scalable storage, equipped with intelligent storage controllers and integrated parallel file systems for peak responsiveness and throughput for both large and small file sets
- Robust interconnects, often based on industry-standard InfiniBand or 10 Gigabit Ethernet
- A robust physical plant, including ample electrical power and effective cooling to ensure reliable operation

HPC's Future

Sustained improvements in HPC performance and capabilities have fueled a drive toward a long-held goal — to produce the first exascale supercomputer capable of achieving 1 exaflops (1,000 petaflops per second) of sustained performance as measured by the Linpack Benchmark. With the Chinese Tianhe-2 currently tops among all supercomputers at 33.86 petaflops (or just 0.0339 exaflops), it's clear that an exascale supercomputer remains years (perhaps even decades) away.

Exascale supercomputers could help energy companies significantly reduce the time required to analyze and interpret seismic data. They also promise to open the door to more complex and effective techniques for processing and interpreting seismic data and modeling the characteristics of reservoirs. Elastic migration techniques, which account for the way seismic waves subtly deform the earth as they propagate, would be enabled by supercomputers in the exascale range.

Top Performing Supercomputers 2000–2013 (teraflops)

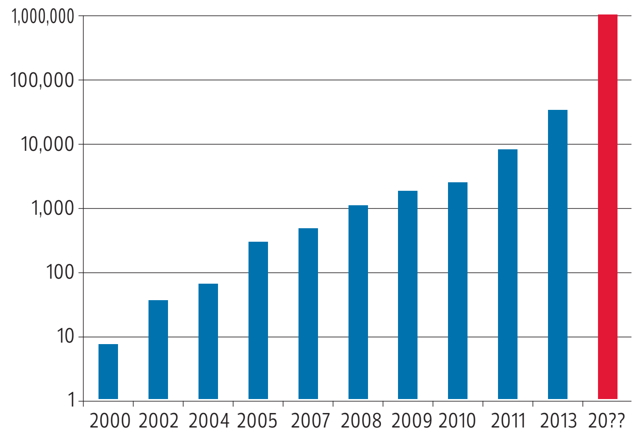

Source: *Visualizing the World's Supercomputers*, Peer 1 Hosting

Achieving exascale performance will require a significant overhaul of supercomputer architectures and software to manage bottlenecks that will otherwise hamper progress. While computing performance has increased impressively in recent years, memory bandwidth and capacity is not keeping pace with the increase, posing limits on future gains. Likewise, all system input-output — from chip to memory to I/O node to disk — will need re-engineering to enhance performance. Finally, developers face a significant challenge in managing failure in such a vast system. At exascale levels, it becomes exceedingly difficult to account for component failures and to ensure that system state and data can be preserved to protect against loss in case of failure.

The U.S. Department of Energy, which is supporting exascale development with grants in the supercomputing sector, describes an example of an exascale system as follows:

System performance	1,000 petaflops
System power	20 megawatts
System memory	64 petabytes
Storage	500–1000 petabytes
I/O aggregate bandwidth	60 terabytes per second

Source: *DOE Exascale Initiative*, September 13, 2013

Ultimately, exascale computing could change the nature of supercomputing. Traditionally, the most expensive aspect of supercomputing performance has been increasing the computing power. In an exascale machine, the opposite may be true. Massively parallel processing units will be capable of prodigious computation, but without significant innovation, the plumbing that connects disk, memory, network and other I/O systems may struggle to support them.

CDW: An HPC Partner that Gets IT

CDW understands high-performance computing and the technologies needed by energy companies. We offer a multipartner approach ensuring that oil and gas companies get the IT architecture they need.

Our account managers are dedicated to energy companies. They solve similar issues for industry players on a daily basis. Our solution architects offer expertise in designing customized solutions, while CDW advanced technology engineers assist with implementation and long-term management of those solutions.

CDW's high-performance computing offerings for the energy industry include technologies such as InfiniBand architecture, 10 Gigabit Ethernet switching, solid-state memory, high-availability/high-volume storage, high-performance servers, power management systems and more.

Projects begin with either an assessment or planning and design session, during which architects review system requirements and perform an evaluation of the existing environment to develop a comprehensive solution to meet specific project objectives.

In addition, offering a complete portfolio of products and services allows us to work on the entire energy-industry stack. Our areas of focus include:

- **Data Storage:** Let us work with you to devise a scalable solution to meet ever-increasing storage demands.
- **Servers:** We've formed alliances with strong technology brands to help your data center operate at peak performance – even while using limited resources.
- **Network Security:** Our brand partnerships help create the airtight, multilayer network defense you and your customers expect.
- **Telecommunications:** Discover how the latest telephony and unified communications technologies can greatly improve the efficiency and success of your firm.

To learn more about CDW's high-performance computing solutions, contact your CDW account manager, call 800.800.4239 or visit CDW.com/datacenter

HP Storage is positioned to provide an excellent value for customers needing increasing performance to support initiatives such as consolidation and virtualization.

CDW.com/hp

As an industry leader in networking, Cisco delivers a robust high-performance computing (HPC) solution to address customer needs. The broad Cisco solution portfolio delivers high speeds, low latencies, open standards and high system availability, allowing customers to deploy the right infrastructure for their applications.

CDW.com/cisco

IBM enterprise endpoint management software combines endpoint management and security into a single solution. It enables enterprise-wide real-time visibility so you can manage all your physical and virtual assets.

CDW.com/ibm

SHARE THIS WHITE PAPER

PEOPLE WHO GET IT™