

www.datadirectnet.com



IFP Advances Oil and Gas Research Using DataDirect Networks Storage Array

www.datadirectnet.com

Application

IFP provides leading-edge research for oil, gas, and alternative fuels. To concurrently support multiple compute-intensive modeling and simulation applications in this highly competitive industry, they needed a high-performance storage solution to speed up their processing time for critical research studies.

Challenge

To eliminate storage bottlenecks while providing a high-performance storage pool that supports massive parallel processing and sustained reads and writes with the reliability and scalability to support over a thousand multi-core processors.

Solution

DataDirect Networks' S2A storage array delivers superior performance for multiple simulation and modeling applications and hundreds of multi-core processors doing simultaneous reads and writes, while providing the scalability required to reliably support over a thousand processors.

IFP Advances Oil and Gas Research Using DataDirect Networks Storage Array

IFP is an independent center for industrial research and development, training and information for sustainable development in the field of energy, transport and environment. IFP is headquartered in Rueil-Malmaison, near Paris, France and its activities cover all aspects of oil, gas and their derivatives.

To innovate and conduct its research programs, IFP anticipates market needs and transfers the results of its work to industry, contributing to the development and technological competitiveness of the oil, petroleum service and supply, gas, and automobile industries. Hundreds of highly specialized engineers and technicians form the Institute's body of expertise.

Research and development at IFP embraces every technical field related to oil and natural gas, including exploration, drilling, production, refining, petrochemicals, internal combustion engines, and the intelligent use of energy, together with the associated aspects of environmental protection.

Representing more than 50 professions, including geology, geochemistry, geophysics, reservoir engineering, applied mechanics, physics, chemistry, engines and fuels specialists, paired with computer science and applied mathematics, IFP's multidisciplinary staff of nearly 1800 forms a unique body of specialists and an unparalleled network of expertise.

“DataDirect Networks is the best solution for providing high-performance I/O.”

— Stéphane Requena,
Project Manager
of IFP's HPC facilities

Complex Research Demands High-Performance Storage

To successfully support the work of its researchers, IFP has used high-performance computing (HPC) since the 1960s. In recent

years, the HPC facility implemented the use of Linux-based clusters, and now has 700 processor cores working to perform the computationally intensive tasks of IFP's many simulation applications. The Institute develops complex simulation and modeling applications using both commercially available code and its own sophisticated algorithms, resulting in an average of 40 to 50 different applications being used by more than 100 engineers on the high-performance cluster at any one time.



One of the vitally important area of IFP's research involves oil reservoir simulation. As oil reserves dwindle, oil companies seek ways to improve oil recovery from their reservoirs, and simulation software is an essential part of this process. IFP's reservoir simulation software maximizes existing oil and gas reserves by identifying ways to extract more oil from existing reservoirs, which can extend the life of an oil field in production by years. It uses voluminous historical and reference data combined with mathematically complex algorithms to simulate and model reservoir properties, performing successive simulations to produce the best possible scenario in order to reduce drilling uncertainties.

Another research area of great importance is innovative car engine combustion systems. IFP works with auto makers to develop new engines that optimize fuel consumption and minimize pollution by improving combustion and reducing exhaust.

Stéphane Requena, project manager of IFP's HPC facilities, explained, “The engine simulation application is so large and, with over a hundred multi-core processors working on it, was taking so long to run that without upgrading to a world-class storage solution, it wasn't going to be competitively feasible for us to use.”

Both applications are incredibly compute-intensive, performing massive sequential writes and multi-processor parallel I/O and creating a heavy I/O burden for the cluster's storage system.

www.datadirectnet.com



Finding oil and gas deposits and modeling the most efficient method of extracting the precious resource demands sophisticated computer modeling enabled by DataDirect Networks' high performance S2A storage systems.

"Before we upgraded our storage system, a single car engine simulation would have to write the output data to disk, then read it back in, which took close to twice as long as it does today," outlined Requena. "A study requires many, many simulations to be run before it is completed, so in order to be competitive, we had to find a much higher performance storage solution. I/O performance is very important and we can't afford to reduce our overall efficiency because of a storage subsystem bottleneck."

Because IFP's existing storage solution had limited performance with parallel I/O, IFP's research applications were being impeded by slow access and the excessive amounts of time required for completion of processing.

DataDirect Networks S2A Storage Solution

IFP's HPC team upgraded their storage system using DataDirect Networks S2A (Silicon Storage Architecture) solution, eliminating the bottlenecks they had been experiencing and achieving a six times improvement in performance.

Requena illustrated, "We made huge progress in terms of performance with our reservoir simulation software because we were able to switch on the parallel I/O system with the DataDirect Networks array. Our previous performance was around 50MBps, but we're now seeing 300-400MBps to a single client. Our parallel I/O benchmarks are running at 1.1 to 1.2GBps." He continued, "The S2A has much better I/O scalability than our previous system."

The cluster computer consists of a variety of processor types using IBM's GPFS parallel file system, with the S2A storage array connected via InfiniBand. The S2A system utilizes 60 300GB Fibre Channel drives for nearly 14 Terabytes of usable storage. The simplified architecture enabled by the S2A array has provided the additional benefit of administration and management cost savings

"Because our applications use so many multi-core processors, it's important to have a balanced system. The bandwidth and in-

terconnect of the storage system are key," commented Requena.


Requena's team was pleased with the painless addition of the DataDirect Networks' storage, explaining, "It took less than one day to get everything up and running."

While the addition of the S2A array has eliminated the cluster's bottlenecks and significantly improved performance, a single engine simulation still takes over a week to complete using more than 100 processors. Consequently, IFP plans to add 600 more processors. To support the additional computing requirements, they are also planning to add an additional 14TB of storage to the S2A array.

Requena cited the S2A's ability to scale while continuing to provide outstanding performance under this massive computing load as a decisive factor in IFP's selection. He also pointed out that the S2A's fault-tolerant architecture with inherent zero-time failover was an important for reliability.

In addition to the cluster computer, the HPC team is planning to connect their visualization nodes and 10 to 15 graphics workstations to the InfiniBand switch to take advantage of the DataDirect Networks storage pool.

Summing up his choice of DataDirect Networks, Requena said, "We had known about DataDirect Networks for a few years. It was our obvious choice for the upgrade."



3D structural model used in oil reservoir simulation

DataDirect Networks is the leading provider of open, scalable storage systems for performance and capacity driven applications. DataDirect's S2A (Silicon Storage Architecture) appliance enables modern applications such as video streaming, content delivery, modeling and simulation, backup and archiving, cluster and supercomputing, and real-time collaborative workflows, that are driving the explosive demand for storage performance and capacity. DataDirect's S2A technology and solutions solve today's most challenging storage requirements, including providing shared, high-speed access to a common pool of data, minimizing data center footprints and storage costs for massive archives, reducing simulation computational times, and capturing and serving massive amounts of digital content.

Major corporations, supercomputing centers and rich media organizations, including AOL, Ascent Media, Boeing, CINECA, CGGVeritas, CNN, Disney, Federal Reserve Board, Ford, Hess, Kodak Gallery, Lawrence Livermore National Laboratories, NASA Ames, RIOT, Sandia National Laboratories, Sony, Technical University Dresden, Technicolor, Time Warner, Thomson, Trinity College Dublin and Universal, utilize DataDirect high performance, high capacity solutions.

DataDirect
N E T W O R K S
Performance. Capacity. Innovation.

9351 Deering Avenue . Chatsworth . California 91311
phone +1.800.TERABYTE (837.2298) . fax +1.818.700.7601
sales@datadirectnet.com
www.datadirectnet.com