Optimizing performance: Object storage for seismic applications

GARTH GIBSON, Panasas, Fremont, California, U.S.

Political instability at many points of the world ... rising gasoline prices ... limited natural gas and oil reserves. These are just a few of the challenges facing the oil and gas industry as it attempts to keep pace with the energy demands of an increasingly developed world. In this quest to supply the globe, it's no surprise that a growing number of energy companies rely upon technological advances to increase their efficiency and productivity. Nowhere is this more evident than seismic data acquisition and processing capabilities, where 4D seismic and other advanced imaging techniques are giving geophysicists more opportunities to accurately analyze the earth's subsurface to make predictions on where and when to drill.

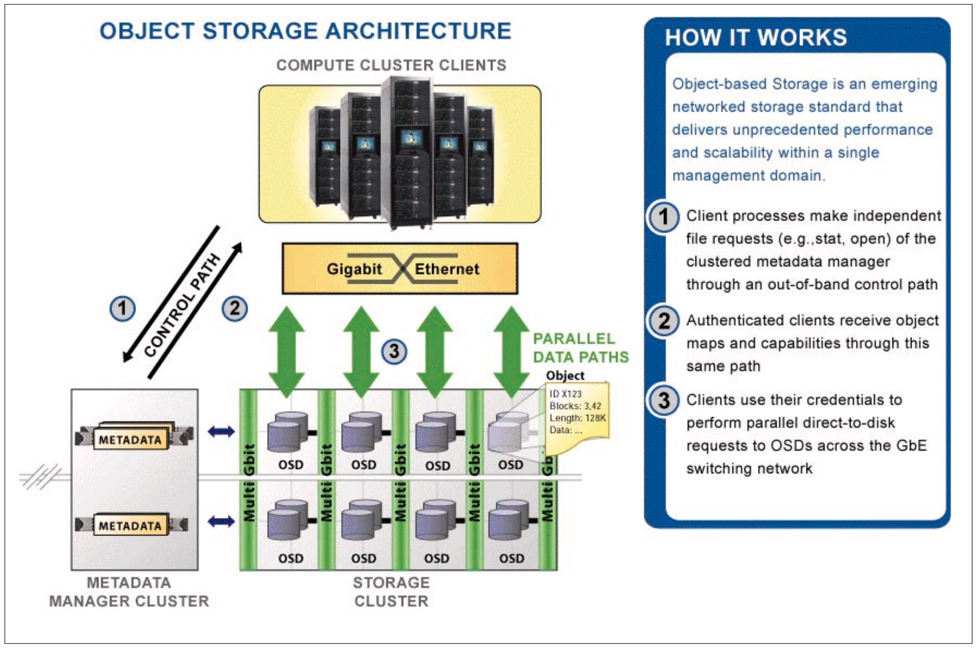

Figure 1. *Object storage architecture.*

Complementing these recent developments has been the advent of Linux compute clusters, which give affordable and extremely powerful computing capabilities to handle the demands of these advanced imaging techniques. These clusters have forever changed the high-performance computing landscape but the storage capabilities have not kept up with the cluster's demands for massive data bandwidth. But today, an object-based storage architecture offers a solution in which the storage system's scalability and price-performance can match that of the Linux clusters.

Current network storage systems have been simply incapable of providing the data throughput needed to keep Linux clusters operating efficiently. Each of the two major types is distinguished by its command sets. First is the SCSI block I/O command set, used by storage area networks (SANs), which provides high random I/O and data throughput performance via direct access to the data at the level of the disk drive or fiber channel. Second, network attached storage (NAS) systems use NFS or CIFS command sets for accessing data with the benefit that multiple nodes can access the data as the metadata (describes where the data exists) on the media is shared.

In order to get the benefits of both high performance and data sharing that Linux clusters demand, a fundamentally new storage design is required that supports systems which provide both the performance benefits of direct access to disk and the ease of administration provided by shared files and metadata. Panasas believes that new storage system design is the object-based storage architecture.

A new storage frontier. Object storage is designed to provide better utilization for large datasets and virtually unlimited growth in capacity and bandwidth. Unlike conventional storage systems, data are managed as large virtual objects. An object is a combination of application (file) data-seismic imaging information, in the case of oil and gas exploration--and storage attributes (metadata) that define the data. Managing data as objects, as opposed to traditional storage blocks, means that files can be divided into separate pieces and distributed across the storage media, known as object-based storage devices (OSDs). So just as the Linux compute clusters spread the work evenly across compute nodes for parallel processing, the object-based storage architecture allows data to be spread across OSDs for parallel access. It's massively parallel processing on the front end, matched by massively parallel storage on the back end.

This architecture (Figure 1) delivers substantial benefits. By separating the control path and the data path, file system and metadata management capabilities are moved to the nodes in the Linux cluster, giving those nodes direct access to storage devices. This allows OSDs to autonomously serve data to the end users and radically improve data throughput by creating parallel data paths. Instead of pumping all information through one path, which creates major bottlenecks as data size and number of nodes increases, Linux cluster nodes can securely read and write data objects in parallel to all OSDs in the storage cluster. Consider this analogy: a massive oil field relies on just one pump to access all of the oil within a reservoir. Now imagine adding 10 or even 100 pumps with little or no additional cost; obviously, reserves would reach market much faster, just as data reach the client much faster.

Object-based storage means the compute cluster now has parallel and direct access to all data spread across the OSDs within the shared storage. The large volume of seismic data is therefore accessed in one simple step by the Linux cluster for computation. While the seismic and imaging data may still require weeks for adequate processing, the object model of storage drastically improves the amount, speed and movement of data between storage and compute clusters.

Seeing more, computing more: A case study. One company utilizing the object storage architecture is TGS Imaging, formerly NuTec Energy, which recently implemented the Panasas Active Scale Storage Cluster to store and retrieve 75 terabytes of data. TGS uses massively parallel supercomputing and imaging solutions in its processing and needed a storage solution that could perform in a single namespace and handle the hundreds of clients connected to that system. The TGS Linux clusters were easily meeting the computational requirements of their advanced imaging techniques, but the equally ravenous appetite for high-performance storage and data access was not being met. The company had been mounting NFS to standard direct attached disks that would sit behind the nodes in their cluster, but this solution had limited bandwidth and scaling difficulty in relation to the number of clients. The company explored an SAN solution, but while offering some performance improvements, integration and system management proved too complex and costly.

Ultimately, TGS chose an object storage architecture built to complement Linux clusters. It combines a distributed file system with intelligent hardware to deliver a high-performance, cost-effective, integrated storage solution. TGS Imaging's cluster nodes, which are allocated on a continual basis to deliver the necessary compute power to complete a job, are now supported by 75 terabytes of storage. The architecture allows each node in the Linux cluster to store and retrieve data in a massively parallel fashion.

The result? TGS has achieved more than 10 times improvement in performance over their previous NFS solution. The company can analyze seismic data sets faster, more efficiently and more accurately. In addition, the system can be scaled to give TGS Imaging nearly unlimited growth and the ability to manage multiple compute processes.

"We've been able to achieve an order of magnitude improvement in performance," said Tony Katz, IT manager at TGS Imaging. "With other products, we were forced to make trade-offs, but with the object system, we were able to get everything we needed and more."

Energy companies face increasing challenges in meeting an unending demand for oil and gas. Technology continues to play a pivotal role in the race to provide valuable resources, and energy companies are now leveraging an object storage architecture in tandem with their Linux compute clusters. This combination allows them to run more advanced seismic images through highly accurate computational models, greatly reducing the time it takes to find promising oil fields, while fully utilizing existing reserves. TLE