



OIL AND GAS

SOLUTIONS BRIEF | APRIL 2007

ABSTRACT

The oil and gas industry's main needs are to quickly find and access oil and gas resources and to do this in a cost-effective manner. Seismic imaging and modeling continue to shorten the business cycle and improve the drilling success rates in the oil and gas industry. At the core of this discipline are the increasingly effective tools running on powerful Linux cluster computers that are available to geophysicists. Together, these tools and Linux clusters place increasing demands on today's storage systems. This solutions note provides detailed information on how the Panasas® Parallel Storage Cluster is uniquely designed to meet the demands of today's seismic imaging techniques. The Panasas storage architecture provides breakthrough scalability and performance that delivers accelerated time to results. Furthermore, the Panasas storage solution establishes a single shared storage infrastructure for the entire spectrum of the oil and gas data processing flow. As such, the Panasas storage solution simplifies administration, lowering the total cost of ownership (TCO) with a "compute in place" capability that reduces overhead and costs associated with data copying and replication. The Panasas Parallel Storage Cluster provides virtually unlimited growth – in capacity and bandwidth – for the most aggressive imaging and modeling applications.



OIL AND GAS INDUSTRY NEEDS AND CHALLENGES

Oil and gas businesses have two very basic needs and challenges:

- quickly find and access oil and gas resources
- reduce total cost of operations.

To find oil and gas quickly, these businesses are dependent on data – acquiring it, processing it, analyzing and interpreting it, doing simulation and visualization of reservoirs, and drilling in the right place before going into full production.

Some businesses are outsourcing companies that process and sell information; others go into actual production. The entire oil and gas industry as a whole, however, has a need to acquire data and refine it. This is essential in making good decisions on where to look for oil and gas resources.

The data processing flow is quite complicated and involves a lot of computation, including work that refines images of what is under the surface of the earth. When undertaking data acquisition, salt deposits disrupt the waves that sensors pick up, and the acquired data has to be processed for correction. Oil and gas applications, both home-grown and commercial packages, perform this kind of processing. The faster a company can move through the entire data processing cycle, the better its chances are for business success.

Another key to success is keeping the total cost of ownership (TCO) low, in particular, for IT resources. In addition to IT hardware and software acquisition costs, reducing TCO also depends on streamlining processes and reducing overhead. When a unified infrastructure can be used throughout all phases of data manipulation, the TCO can be reduced.

OIL AND GAS DATA PROCESSING FLOW

Oil and gas businesses deal with a massive amount of data. Today’s 4D modeling technology adds a time element, and increases the amount of processing power necessary to compute. By taking numerous surveys and overlaying them with time-dependent information, well logs and production histories, geophysical scientists are processing as much as 100 terabytes (TB) of data at one time. This data is moved through a processing flow that includes seismic analysis and modeling, analysis, and visualization.

In a simplified view, the data flow includes the following components:

- seismic analysis
- modeling, analysis, and visualization.

Seismic Analysis

Seismic analysis begins with data acquisition, the generation and recording of seismic data, using seismometers and geophones on land, in wells, or in the sea. It is followed by seismic processing to suppress noise, enhance signals and migrate seismic “events” (which occur at subsurface structure transitions) to the appropriate location in space. Seismic processing provides more accurate substructure models.

Modeling, Analysis, and Visualization

The modeling, analysis, and visualization part of the data processing flow include reservoir modeling. A reservoir is a subsurface body of rock with characteristics (porosity and permeability) that make it suitable for storing and transmitting fluids – including oil and gas. Reservoirs are characterized (modeled) and simulated, and then visualization techniques are used to analyze various extraction or drilling alternatives.

Compartmentalized IT Resources

There is a tendency to compartmentalize IT resources for the various stages of oil and gas data processing. Instead of a unified compute and storage infrastructure, there may be silos of infrastructure for the various data processing stages. While this approach may have near-term benefits, the bigger picture that involves TCO may not be factored in. Ideally, both compute and storage resources need to be evaluated for overall value with performance, scalability, leveraged infrastructure, and overall management overhead in mind.

THE COMPUTING FACTOR

In order for oil and gas companies to shorten their business cycles by getting to the oil and gas fast, they need to have the computational power to move data through the processing flow. The best solution is to use compute clusters, in particular, Linux clusters. With Linux clusters, customers get value in terms of performance and price. The power of hundreds to thousands of multiprocessor nodes allows parallel processing of data and accelerates times to results.

Along with computational power, Linux clusters have a ravenous appetite for high-performance storage and data access. Without efficient access to numerous seismic surveys, the newly updated computing resources are under-utilized and the benefits of 4D models cannot be realized. Up to this point, traditional networked storage systems have been simply incapable of providing the data throughput needed to keep these Linux clusters operating efficiently. With the capital investment of the cluster in place, customers need to keep it busy, and that means keeping it well-fed with data. Cluster utilization is a key factor in keeping IT costs low.

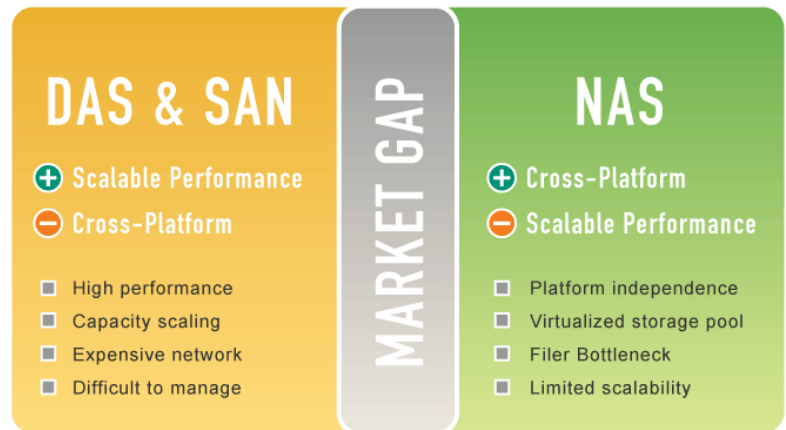
In addition to utilization, it is also important to allow the cluster to keep making progress on the jobs that it is processing. Checkpointing is an important part of this. The computing workload is spread out for parallel processing amongst the compute nodes. Checkpointing allows for the various completed parts of the job to be consolidated and “synched up” for the next step of computing. Checkpointing can also be thought of as a snapshot so that if there is a failure in the system, the overall job does not have to be started from scratch again. The job can be started and resumed from a checkpoint. If the cluster demand for data can be satisfied and checkpointing can be done efficiently, then the customer will get the most return from the cluster investment.

The cluster interdependency with storage is then obvious. Oil and gas data processing jobs have very long compute cycles that may last for weeks. During those cycles, there may be negligible I/O activity. Then suddenly, there may be a huge burst of I/O activity, and the storage system must be able to handle that demand, which may be for multiple GB/sec from the aggregate of applications running on entire set of cluster nodes. The datasets involved are large, requiring from a few hundred GBs to TBs of storage. There are typically several passes with large bursts of I/O. At the start of the job, a lot of data may be read; during processing, there may be several passes with a lot of write activity. As such, it is very important that the storage system have the ability to scale performance as capacity increases. Without an adequate storage solution, the cluster will not be able to make progress, and will have to stand idle while waiting for an I/O bottleneck to be cleared.

THE STORAGE FACTOR

As the right kind of storage system is essential to success in the oil and gas, let us examine the various storage approaches and architectures. Each has its own tradeoffs and drawbacks. Sometimes it is not the approach that is the problem; it is the storage architecture or implementation that results in serious tradeoffs that are shown below for Direct-Attached Storage (DAS) and Storage Area Networks (SAN) versus Network-Attached Storage (NAS).

Of course, reliability is foundational; if the storage system fails, the cluster will not have access to data. Sharing is important to span the various stages of oil and gas data processing flow and to reduce overhead. Sharing data across multiple clusters enables leveraged use of storage resources and valuable data. Performance is what enables companies to accelerate time to results. Scalability includes being able to add capacity in concert with performance, so that growth does not have the drawback of building isolated islands of storage. Finally, manageability is essential so that the IT staff overhead is kept low.



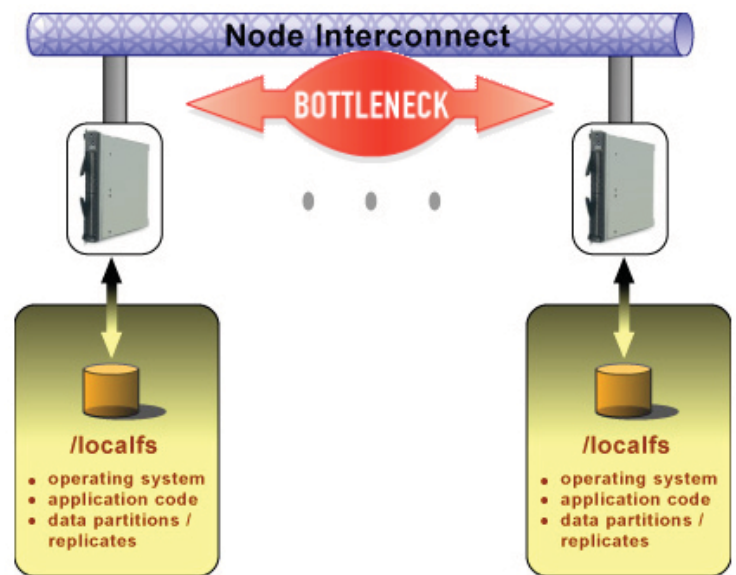
The storage challenges have to do with choosing where to store data in view of potential tradeoffs. It is important to map the storage to each and every phase of the data processing flow. This is especially critical to the oil and gas sector. Let us explore why by examining the traditional approaches to storage and the market gap that they leave.

Direct-Attached Storage

Direct-Attached Storage uses local disks for each compute node. While disks have specific physical and logical association with nodes and there can be good aggregate bandwidth, there is no RAID protection. This results in poor reliability. A local disk failure essentially is the same as a node outage.

Data cannot be shared as there is no central storage pool, so in order for nodes to get access data that is not local, data copying or replication is required. The data has to go over the node interconnect, which then becomes a performance bottleneck.

With DAS on a Linux cluster, there is a lack of scalability and flexibility. To add storage capacity, you have to install local disks to the compute nodes. If you add storage capacity



to one set of nodes in the cluster, you need to consider adding the same incremental capacity to the rest of the nodes. Otherwise, your system will not be balanced, and copying and replication may run into capacity limits. In fact, this presents a manageability issue as well; in order to add more storage capacity, you may reach the point where you need to add more compute nodes and network connections. This tie-in to compute and networking infrastructure proves to be very inconvenient. Furthermore, because the storage is disjointed and not in a centrally shared pool, checkpointing overhead increases as all local disks have to be “synched up” with one another. A better alternative would be a centralized storage pool where the entire storage system could be checkpointed with the equivalent of a system-wide checkpoint “snapshot.”

Storage Area Network

A Storage Area Network does provide a shared pool of storage. It uses a dedicated high-performance network (usually Fibre Channel) to connect the cluster to the storage systems. Logical volumes are assigned to the various nodes. Data is accessed by way of a file system that is managed from the cluster nodes. As the data is stored in logical volumes managed by specific nodes, in order to have wide access across the cluster, nodes may have to handle I/O requests from other nodes.


Reliability is improved over local storage as storage systems with RAID can be included in the SAN. However, there are still problems with data sharing. SANs actually allow sharing of storage, not so much sharing of data. The storage pool has storage systems that partition their capacity for specific nodes. Sharing requires the nodes to act as file servers or I/O nodes. There is potentially still a need for copying and replication of data. This would be the case when a node would need more control and access to data that belongs to another node. If a lot of data sharing is required, the node interconnect is still a performance bottleneck. Remember that in addition to I/O traffic between nodes, the interconnect is used to coordinate the overall activity of the aggregate cluster.

In terms of scalability and flexibility, SAN infrastructure still comes with a price premium, and managing the various copies of data is not trivial. It is much like a checkpointing issue across the storage pool requiring to keep track of the various versions of data and managing them.

Network-Attached Storage


Network-Attached Storage allows for both shared storage and shared data. It uses a file system that is managed by the storage system, which typically has a file server (filer head). The filer head can be thought of as an I/O node that serves data from its disks to the cluster compute nodes. It fulfills the I/O requests from the cluster by using standard protocols such as NFS and CIFS.

As with the SAN approach, NAS provides RAID protection for better reliability and availability of data. Data sharing is easily done as the file system in the filer and all the nodes can access the same volume. However, the filer head has to handle all I/O traffic for that volume. Therefore, sometimes replication may be necessary in order to scale and service a much-shared volume.



The filer head is a performance bottleneck under heavy I/O load. This presents a lack of scalability and flexibility. In order to scale performance, additional filers have to be introduced into the configuration. With that comes manageability issues; a filer may be convenient and suitable for a small business department, but a farm of filers servicing a high-performance compute cluster is another story. The manageability issues quickly rise and overhead goes up. In addition, if replication is used, it will require management to keep track of the various copies, and also to keep them synchronized.

The idea of NAS still leaves something to be desired for shared storage. In fact, in real life implementations, shared storage in general often deteriorates again into “islands of shared storage.”



Shared Storage Challenges

The idea of shared storage for compute clusters is intended to greatly streamline processing and simplify administrative tasks. Ideally, data partitions can be made available globally, with no replicates required. This greatly simplifies data management, especially for datasets that change periodically. Applications can be written to be agnostic of data location and can thus run anywhere in the cluster. This affords greater flexibility in running applications on clusters with varying configurations. In this scenario, no data staging or de-staging is required, as all partitioned and global datasets are accessed from shared storage. This greatly reduces the chance of error due to improper file propagation (data consistency issues) and eliminates time consuming data movement activities.

Unfortunately, the traditional architectures for shared storage for clusters result in “islands of shared storage.” Because of the multi-volume (and mount point) structure of these systems, most organizations restrict a single application’s datasets to a single storage system or NAS head. This greatly limits the performance (both in terms of bandwidth and IO ops/sec) that a single application can achieve, and means that applications running in a “shared storage” environment will suffer performance degradation due to resource contention. Moreover, it means that total system utilization is poor (as only one server’s resources would be used at a time). Attempting to distribute application datasets across multiple storage systems or filers for capacity and load balancing is a difficult task, one that requires constant attention in an operational setting. This is, at best, costly and often unfeasible for most organizations.

Finally, this is a very inefficient model for managing storage costs. One has to trade off administrative concerns for performance concerns, for example, maintaining related files on one storage system versus spreading them across multiple systems for performance. Capacity and bandwidth are added to the overall system only by adding new systems/NAS filer heads. This often requires manual load and capacity re-balancing of the storage “cluster” which results in application downtime. Additionally, this procurement model encourages large quantum capacity expansion, causing the customer to pay for storage in large chunks ahead of their need. As such, purchase decisions often focus on minimum cost for a particular phase of the data processing flow (e.g., for seismic processing) rather than on what is best in the big picture of a shared infrastructure and overall TCO.

According to Gartner, 80% of storage TCO comes from ongoing management activities. While performance is important for oil and gas businesses, it is important to remember that ease of management is very important, especially at large storage capacity levels.

SAN does not allow for true data sharing. While NAS products were initially viewed as easy to manage, due to scaling limitations and the need to purchase and manage multiple systems, what was once perceived as a benefit of NAS has now become a management liability.

PANASAS ACTIVESCALE® OPERATING ENVIRONMENT


The Panasas ActiveScale Operating Environment delivers an architecture that has been built from the ground up to address many of the management issues plaguing oil and gas customers today. It allows for a single namespace that can be shared across the cluster and does dynamic load balancing. The idea is to provide a single storage infrastructure that can be fully leveraged and utilized for the entire oil and gas data processing flow.

The key differentiation of the Panasas storage infrastructure lies in the software architecture. Panasas is leading the development of a new wave of intelligent storage, called the object-based storage architecture. The core principle of this architecture is that data is managed in large virtual objects and not as small blocks of bytes.

This architecture allows the Panasas Parallel Storage Cluster to breakthrough the scalability barrier with performance that scales in concert with capacity. The scaling of performance with capacity is almost linear; Panasas does this with an architecture that uses Panasas DirectorBlade® modules to scale and manage metadata growth and Panasas StorageBlade® modules to scale and manage capacity growth. The Panasas storage system shelves scale up in networking infrastructure as well, with each shelf having a 16-port switch blade for enterprise-class scalability. This is truly a breakthrough storage technology that is ideally suited for Linux clusters and specifically for the oil and gas industry.

In addition, the Panasas DirectFLOW® client access software allows for maximum performance with parallel I/O to each compute cluster node. Before starting, all nodes in the cluster are empowered with a small installable file system called the DirectFLOW client to enable direct communication between clients and StorageBlade modules.

The Panasas Object Storage architecture divides files into objects; this is especially effective for large files. An object is a logical unit of storage. It can be accessed with file-like methods such as open, close, read, and write. Objects have associated application data and attributes and metadata. Each object has attributes assigned to it and is designed to be managed, grown, shrunk, deleted, and dynamically distributed and redistributed across physical media. As these objects are

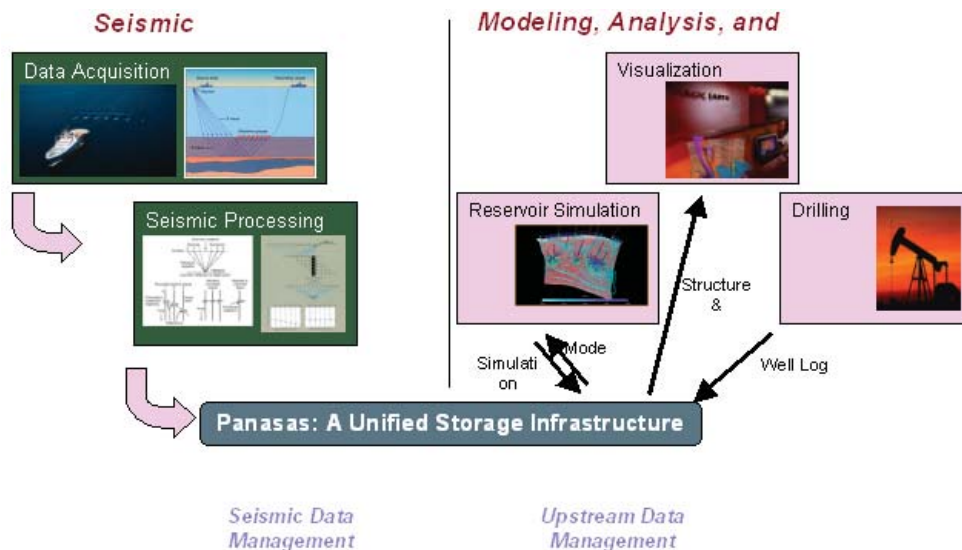


distributed across the Panasas StorageBlade modules, they can be accessed in parallel by the cluster compute nodes. However, the management of the objects is done by the DirectorBlade modules. The intelligence within each object allows the system to offload work from the traditional filer head or file server to the actual storage.

Object storage enables two primary storage breakthroughs. First, since the system is able to offload work to the storage instead of going through a central filer head or manager; the system is able to deliver parallel performance directly from disk. Further, since each object is injected with attributes, it can be managed intelligently.

DirectFLOW Protocol

Once the DirectFLOW client is installed in all the cluster nodes, a simple three-step process is required to initiate to all direct data transfers. First, the DirectFLOW client makes a request to the DirectorBlade module to get and/or store data files. Next, the DirectorBlade module authenticates the user and request, then obtains the latest map of all applicable objects across the StorageBlade modules and sends a virtual copy of the map to the client. Finally, with authentication and a virtual map, the clients are able to directly request, obtain, and store data on StorageBlade modules. Data is streamed off each StorageBlade module in parallel. This concurrency eliminates the bottleneck of traditional, monolithic storage systems and delivers an order of magnitude performance throughput improvement versus traditional storage solutions.



The number of data streams is limited only by the number of StorageBlade modules the file sits on and the number of clients (or clusters nodes) in the compute cluster. Specifically, if you have one large file, you can have multiple StorageBlade modules and multiple clients reading and writing that data concurrently.

With the Panasas breakthrough architecture and technology, the scalability barrier is shattered. Instead of islands of storage, you can have one storage infrastructure that spans the entire spectrum of your oil and gas processing flow: the Panasas Unified Storage Infrastructure. With this, you will avoid the overhead of islands of storage in a fragmented set of infrastructure silos.

PANASAS VALUE PROPOSITIONS

Going back to the needs and challenges of the oil and gas industry, the goal is to find oil and gas and do it quickly. It is a competitive market, and the faster a business can locate and access oil and gas resources, the more successful it will be. With the Panasas storage architecture and its massively parallel I/O and industry-leading performance, you have the very best technology to speed your Linux compute clusters to the completion of your data processing flow. The bottom line is that you will be able to shorten your business cycle because the Panasas storage architecture has none of the bottlenecks of traditional storage systems. In fact, you scale your performance as you increase your data and storage capacity.

With a centralized, unified storage infrastructure, you will reduce your IT overhead with simplified storage management.

Finally, it is not just the price that you pay for storage that matters; it is your overall total cost of ownership. Without having to purchase additional storage for data copying and replication and without having to manage data versions and synchronization, you come out ahead with Panasas. You reduce your true cost by being able to leverage the Panasas storage infrastructure to “compute in place” rather than have to copy data from one island of storage to another.

CONCLUSION

To recap, the oil and gas industry depends on Linux compute clusters because of their computational power, their value, their economies of scale, and their ability to contribute to accelerated time to results.

Panasas is uniquely designed for Linux clusters, and will be able to meet the demanding I/O loads that these Linux clusters generate. The Panasas storage solution scales performance and capacity in concert to shorten the oil and gas business cycle. It provides a single unified storage infrastructure for the entire data processing flow of an oil and gas enterprise. Finally, it delivers the best overall storage value and reduces TCO.

Storage purchase decisions need to be made with the big picture in mind. Stopgap measures focused solely on localized criteria such as price per unit storage often detract from investing in storage that will serve an entire oil and gas data processing flow.

The Panasas storage solution is a perfect fit for oil and gas as it was intended to meet the needs of an I/O profile that exactly matches that found in the oil and gas industry. The Panasas Storage Cluster is designed for advanced seismic imaging environments that demand a highly scalable, yet easily manageable storage network.

As such, the Panasas parallel I/O architecture presents a business advantage to oil and gas businesses.

Accelerating Time to Results™

6520 Kaiser Drive Fremont, California 94555 Phone: 1-888-PANASAS Fax: 510-608-4798 www.panasas.com
1-888-PANASAS (US & Canada) 00 (800) PANASAS2 (UK & France) 00 (800) 787-702 (Italy) +001 (510) 608-7790 (All Other Countries)

©2007 Panasas, Inc. All rights reserved. Panasas, the Panasas logo, and DirectFlow are trademarks or registered trademarks of Panasas, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.

Information supplied by Panasas, Inc. is believed to be accurate and reliable at the time of publication, but Panasas, Inc. assumes no responsibility for any errors that may appear in this document. Panasas, Inc. reserves the right, without notice, to make changes in product design, specifications and prices. Information is subject to change without notice.