

SOLUTION BRIEF

OIL & GAS – SWIMMING IN THE SEA OF SEISMIC DATA



Oil & Gas Storage Needs

- Extreme Scalability
- High Durability
- I/O performance
- Seamless integration
- Self healing
- Single-pane management
- Very low TCO



OPPORTUNITY / CHALLENGE

Big data is the lifeblood of the oil & gas industry. It is a simple fact that “quality equals volume” for example: new 3D seismic imagery requires 10X the amount of storage capacity of standard seismic imagery. The question all HPC managers are dealing with is “How do I balance the need to spend money on the latest interpretation and visualization applications that help drive revenue versus the need to expand storage capacity to store the massive amounts of data?” The logical answer to this question is to spend less on storage and more on things that help drive revenue, but how?

TRADITIONAL STORAGE NOT SUITED FOR TODAY’S BIG SEISMIC DATA

Storage efficiency, from both a capacity and performance perspective, is key to delivering a solution that works with today’s big seismic data. Once data is stored it must be accessible for processing, interpretation and visualization and the storage system must have read/write performance to handle the transfer of large data sets to and from the HPC cluster or to end users or other applications. Finally, data loss is not an option for oil & gas HPC environments, so high data durability is a must. Traditional storage options utilizing RAID-based storage and tape backup no longer meet the challenge. The main reason is that as drive sizes increase within RAID sets, so do the rebuild times during failure, with some rebuilds taking as long as a week causing severe performance problems. In addition, standard tape backup processes are no longer capable of handling the massive amounts of data, are increasingly cumbersome and do not provide the necessary reliability.

SOLUTION

The new generation of object storage solutions like Himalaya were built specifically to solve this big energy exploration data challenge. Its modular architecture and fully abstracted software stack delivers unbreakable durability, infinite, scalability, and extreme efficiency at a lower cost. Optimized for Intel-based commercial-off-the-shelf hardware Himalaya is designed to take full advantage of the latest Intel* Xeon* Processors with performance scaling linearly with each new controller in the system. Himalaya is ideal for creating a common repository for the vast amount of seismic data and can consolidate several tiers of storage into a single tier.

HIMALAYA™

HIMALAYA™

Reduced CapEx

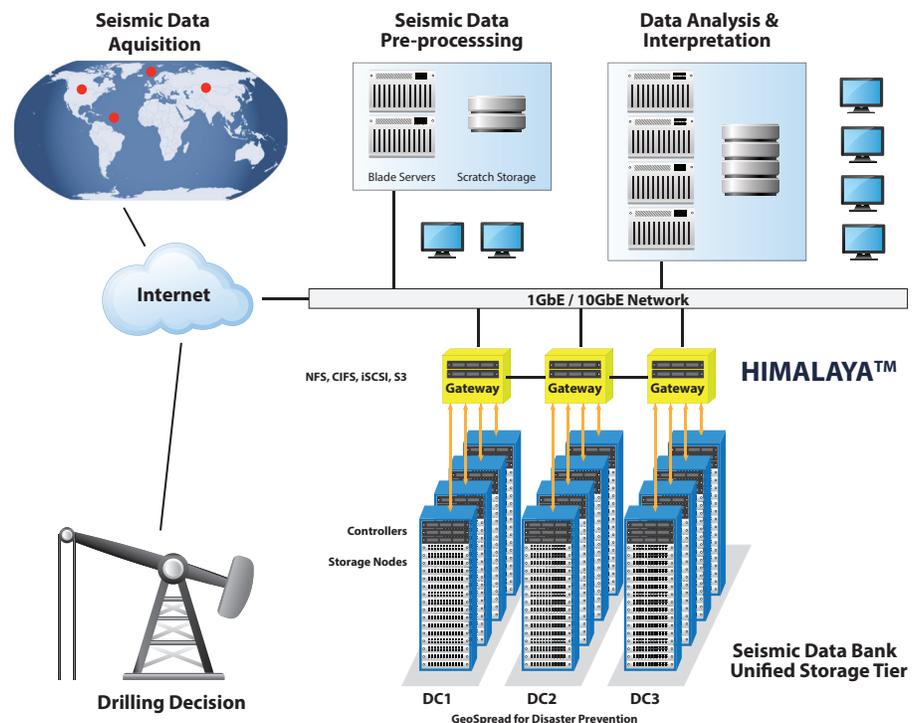
- COTS hardware optimized for Intel® Xeon® Processors
- Reduced storage overhead
- Lower software cost

Reduced OpEx

- High availability
- Single system view
- Longer system useful life
- Fail in place model
- Automation
- Remote monitoring
- High density storage
- Low power consumption

A global namespace allows applications and users to access key data no matter where they are located. Integrating with existing applications and environments is easy with support for NFS/CIFS or iSCSI protocols via certified technology partner gateways. Patented BitSpread® technology protects data with up to fifteen 9s data durability and handles 19 simultaneous drive failures without losing data or access. GeoSpread offers disaster prevention by spreading data across multiple sites using 65% less capacity than RAID 6 and replication. System data remains protected and accessible even if an entire site becomes unavailable. Hardware components include Intel® Xeon® processor-based controllers and storage nodes, along with a 10GbE/1GbE fabric. Increasing system capacity or performance can be done independently by simply adding more storage or controller nodes.

Patented BitDynamics® continuously monitors data integrity. When an issue is detected, it quickly and automatically self heals with virtually no impact to system performance. Failed drives can remain in place to be addressed during a planned maintenance interval vs. disruptive ad hoc drive replacements. Below is an example of a large oil & gas seismic data storage infrastructure.



Contact Amplidata

North America: 408-433-1600

EMEA: +32 9 324 25 90

Email: info@amplidata.com

CONCLUSION

With Himalaya, seismic IT operations can dramatically reduce the overall cost of data storage while creating opportunity for new levels of operational efficiency and work-flow improvements. Optimized for Intel-based COTS hardware, scaling capacity and performance over time is easy and cost effective. Silos of seismic data can be consolidated into a large central repository putting more data on-line and accessible by applications and users no matter where they reside. Traditional backup solutions, massive data migrations and long RAID rebuilds become things of the past. Himalaya makes storing and managing the large and growing seismic data sets more efficient, allowing organizations to free up critical resources so they can be applied to high value activities that move the company forward.

AmpliStor, BitSpread and BitDynamics are registered trademarks of Amplidata. *Other Brands and Names are the property of their respective owners.