

Panasas® ActiveStor® 14 Powers 3D Seismic Volume Interpretation

TerraSpark Geosciences designs software tools for energy exploration and production. Its flagship product, Insight Earth, enables visualization-guided 3D seismic volume interpretation of structure and stratigraphy, and integrated directional well-path planning.

SUMMARY

CUSTOMER

TerraSpark Geosciences

INDUSTRY

Energy

CHALLENGE

- To source a state-of-the-art, easy to use scale-out NAS solution that would provide top performance and maximum scalability
- To optimize Insight Earth performance
- Emulate tier-one customer interpretation environments

SOLUTION

- ActiveStor 14
- PanFS parallel file system

RESULT

- Processing times reduced from five days to 20 minutes
- 360x acceleration for a 1GB 3D seismic volume
- Streamlined workflows

TerraSpark had an opportunity to upgrade its development and test resources while equipping its engineers with powerful, state-of-the-art workstations and enterprise-grade network attached storage (NAS) in order to assure seamless cooperation with its technologically savvy customers.

THE CHALLENGE

According to TerraSpark Systems Administrator, Steve Dominguez, the company needed a high performance scale-out NAS solution that would be fast, easy to install and be essentially 'self-tuning' to effortlessly accommodate both very large and very small files simultaneously. Simple system management would also be a key consideration because taking the system offline for maintenance or management would simply not be an option as active production never stops and drilling dates must be met.

TerraSpark customers collect field data in complex workflows that take up to several months to complete. They deliver their files in the SEG-Y format and they can grow to 100GB. These files are typically compressed by a factor of four as they are pulled into workflows by geophysicists who analyze the data to assess potential drilling opportunities. Multiple versions of both very large and very small files may be generated and saved causing the number of project files to balloon into the tens of thousands. As a result, a single project might amount to 5TB of total content, and there might be dozens of projects over the course of a year.

THE SOLUTION

TerraSpark regularly surveys key customers to better understand their evolving requirements. The company recently sought input on its storage upgrade project. On the workstation side, customers described providing dual socket 16 core Xeon E5 systems running RHEL 5 with 256GB DDR III DRAM and the most powerful GPUs available to their own interpretation staff. On the storage side, it was a bit more complicated. "We received a wide range of input, so we looked at a lot of options," Dominguez said.

"Many storage vendor claims sounded pretty good until we dug into them a little bit. We eliminated many options pretty quickly, but with Panasas ActiveStor 14, the more we drilled down, the more we liked."

Steve Dominguez

Systems Administrator, TerraSpark Geosciences

"ActiveStor 14, in a 3+8 (three director blades and eight storage blades in a single chassis) configuration with 45TB of raw HDD capacity, 4.8TB of SSD capacity and 172GB of DDR III ECC DRAM cache, was an amazingly good fit for our mixed workload requirements. The integrated Panasas PanFS parallel file system delivers the performance and scalability we need.

It dynamically distributes and automatically load-balances data transfers across the architecture with small files and metadata served on integrated SSDs and large files on enterprise SATA disks. And the best thing is, I don't have to configure anything. It just happens. The Panasas DirectFlow™ parallel protocol eliminates I/O bottlenecks and saves me endless hours tuning NFS parameters to get the best performance on large file volumes," Dominguez said.

"Right now, our developers are driving a Ferrari on a farm road as network optimization is still a work in progress. Segments of "The Shark Tank," the TerraSpark developers' workspace, are still on 1GbE. It's a little bumpy at times, and slower than we could go otherwise," Dominguez asserted. "But even with this limitation, ActiveStor 14 helps. Native support for link aggregation enables us to gang multiple 1GbE links to the storage for additional network bandwidth, at least until we get 10GbE rolled out across the organization."

THE RESULT

In-house testing of ActiveStor 14 at TerraSpark has shown crucial algorithm processing times reduced from five days of compute time to 20 minutes – a 360x acceleration for a 1GB 3D seismic volume. Workflows are also streamlined so it takes fewer steps and just minutes to go from a 3D seismic amplitude volume to a complete set of fault surfaces.

"The 150GB/s throughput delivered by ActiveStor 14 is enough to keep large teams of interpreters happy and productive. The potential to scale capacity to 8PB while maintaining IOPS and throughput – this idea of 'linear scalability' – gives us confidence and peace-of-mind when we recommend Panasas to key customers," Dominguez said.

Performance, Scale and Ease-of-Use

"We chose Panasas because of its well-earned reputation among our most important customers as well as for its performance, scalability and ease-of-use features. When we polled our customers, they were already familiar with and – more importantly – very pleased with Panasas and the ActiveStor line. That made us feel like Panasas was a low risk play for us, and in this business, those don't come along very often."

The goal was to put customer-realistic tools in TerraSpark's development and test processes in order to optimize its Insight Earth software packages for real-world conditions, as well as employ best-in-class hardware for its development efforts. The company learned as part of its due diligence that Panasas and the ActiveStor line already had a strong reputation within oil and gas exploration and production environments.

"Panasas provided extraordinary support throughout our evaluation, proof-of-concept and deployment stages. The product is top notch and the service and support are exemplary. We have been well-served,"

Steve Dominguez
Systems Administrator, TerraSpark Geosciences

