

# Oil and Energy Case Study

## Highlights

### Oil and Energy

#### Challenges

- High Concurrency
- Large Capacity Storage
- High Availability

#### Scala Storage Solutions

- Dynamic increase capacity and performance
- No single point of failure
- Automatic failure detection
- Fast data recovery

#### Best TCO

- Pay as you grow
- Single management interface
- Predictable scale up capacity and performance.

#### Real-Life Cases

- Support 500+ blade servers
- Over PB of data
- Outperform competitor

## Scala Storage Scale Out Solution helps Oilfield make Better Decisions and Resolved Storage Limitation

### Background:

At an undisclosed top 10 oilfield location, which produces 40 million to 50 million tons of oil annually, Scala Storage Scale-Out Solution has been tapped to greatly improve performance and drive its oil and gas seismic processing business into the future. Advancements in the collection of seismic data with unprecedented levels of detail is important information that has overwhelmed its current storage systems, usually Fibre Channel SAN with a legacy file system such as HDS with StorNEXT or IBM hardware with GPFS. Scala Storage has dramatically reduced the time and resources required to manage data and has allowed the customer to refocus vital resources towards interpretation and delivering results to its management.

The oil and energy business is not simply about providing seismic data, but rather delivering detailed analyses that can be used to make more resourceful drilling decisions. By choosing Scala Storage, the ability to quickly organize and analyze seismic data was enabled and accurate exploration decisions were made based upon:

- Leveraging vast archive of seismic data to compute new lines of analyses
- Reduced analysis iteration cycle from four weeks to one week

Among many requirements of a scale-out solution was one that could quickly add on new storage nodes without system downtime, meet the heavy usage demands of a 24x7 operating environment, and unify disconnected file volumes into a single, fully redundant file system.

### Business Challenges

- The need to accurately and quickly analyze massive amounts of data
- The concurrency connection of thousands of blade servers generates pressure to the storage systems.
- Traditional storage systems (FC SAN) are expensive and are unable to meet data requirements.
- High capacity, high availability, and high recovery process requirement.

### Solutions

- Phase I: Scala Storage, 200 TB in 2009
- Phase II: Scala Storage, 1.5 PB in 2012

### Results

- Disconnected data silos into a single file system.
- Reduced sample analysis workload cycle from 24 hours to 6 hours.
- Phase 1 consisted of a no downtime replacement of all original 750 GB Disks with new 2TB disks, resulting to zero downtime over 6 years.
- Phase II provided over 12 GB/s aggregated bandwidth via 10 GbE, which shortened the same operation workload shortened from 4 days to just 10 hours.



## Customer Comments

“Scala Storage resolved previous bottlenecking of bandwidth issues when thousands of computing nodes read and write together. At the same time performance from original has improved 5X more. With preventive failure detection and auto self-healing, critical information is protected without impact of performance.”

“Scala Storage not only provides high bandwidth to resolves high read and write concurrency issues, but also allow dynamic upgrade performances without impact on productions.”

## Solutions Comparison

### Testing Environments

Testing software: Omega Geophysical Data Processing Software

Hardware: 500 computing nodes in a blade clustering, connected with GbE, using Omega

Resolved issues	FC SAN + StorNEXT	Scala Storage
<b>High Concurrency Support</b>	At high concurrency situations, FC SAN Controller CPU is 100% utilized, creating bottleneck.	MetaData and storage nodes are in a cluster structure, and all processing requirements are divided to each node, which provides the best performance.
<b>Failure recovery</b>	Hard drive rebuild takes over 10 hours depending on data size.  Performance is degraded during rebuilding.	90% full 3TB hard drives only take 20 minutes to rebuild.  Preventive failure detection and auto healing provides non-degraded performance.
<b>Performance</b>	Read Write Performance is limited by the storage controller, bottlenecking at the controllers hardware.	Clustered, distributed model provides combined performance from each of the storage nodes. Performance linearly expanded by adding more storage nodes.

## About Scala Storage

Scala Storage has redefined storage solutions to solve these challenges with Scala Storage Scale Out System, Scala File Systems and Scala Manager. These storage solution implementations are designed for enterprise, private and public cloud computing, and across a variety of industries. Scala Storage is scalable from hundreds of terabytes all the way up to petabytes, while still providing predictable increasing linear performance and high availability storage structure. The Scala Storage family delivers extreme capacity and performance, and significantly reduces an organization's capital expense, operational expense and total cost of ownership.