

SeisSpace® and Parallel File Systems

Guidelines for setting up your seismic processing parallel file system

Cray's oil and gas performance engineering team, along with Taming Traces Consulting, profiled the I/O requirements for Landmark's SeisSpace® with a variety of workflow scenarios to produce a series of guidelines for parallel file system setup. Using the Cray® CS400™ cluster supercomputer and three separate storage platforms — the Cray® Sonexion® Lustre® parallel file system, IBM GL6 GPFS file system and DDN GS14KX™ GPFS file system — the team looked at how workflow performance would be affected under the following circumstances:

- Combining primary and secondary data on the same parallel file system
- Using GPFS instead of Lustre
- Modifying file system block size
- Using Lustre pools for primary and secondary separation

They ran hundreds of production-level read, write and read-then-write SeisSpace jobs with both GPFS and Lustre, varying the input for block size.



PRIMARY AND SECONDARY FILE SYSTEMS

Running multiple jobs with other non-SeisSpace contention on the system, the results showed that placing both primary and secondary file systems on the same parallel file system produced little performance reduction. This result will vary as the secondary user load increases, causing a possible delay in primary and in the use of the SeisSpace Navigator.



LUSTRE VERSUS GPFS

Both Lustre and GPFS perform equally well under production workflows and the load provided in testing. GPFS does have the tendency to provide better performance on edge case requirements and higher read loads.



OPTIMUM BLOCK SIZE

The newest versions of Lustre and GPFS work extremely well right out of the box. The most critical parameter to adjust is block size. Results showed 4 M to be optimum for a combined primary and secondary file system. As the file system block size is adjusted in a combined file system, 512 MB is where you begin to see I/O issues.



POOLING IN LUSTRE

Pooling in Lustre for the purpose of separating primary and secondary provides no performance improvement.

The bottom line

Keep things simple when designing your parallel file system for SeisSpace primary and secondary. The newest hardware and file system technology solves many of the old problems.

For more information: www.cray.com/energy