

# ASIODS: An Application and I/O Collaboration System

Maxime Hugues<sup>1</sup>, Mike Moretti<sup>2</sup>, and Serge Petiton<sup>3</sup>

<sup>1</sup>*INRIA, Saclay, France*

<sup>2</sup>*DataDirect Networks, Sunnyvale, CA, USA*

<sup>3</sup>*LIFL, Villeneuve d'Ascq, France*

After the break of the 10 Pflops by the “K” computer, post-petascale era is few steps away and exascale machines are on the way. This continuing and insatiable appetite for computation power comes from the need to accelerate applications and solve larger, more complex problems to get a higher accuracy. However, building and operating of those machines are challenging for computer scientists by many aspects such as fault tolerance, energy consumption, programming, network and storage system performance. By focusing on storage, trends announce that more than hundred petabytes and 20TB/s of bandwidth would be necessary for applications, which would process and generate a large set of data.

However, storage performance is an ongoing problem for seismic imaging applications that generate and rely on many terabytes of data. Commonly, storage systems use caching techniques to improve performances. Then data are kept in RAM what avoid disk accesses and give a faster transfer. Nevertheless, this technique does not allow to optimize irregular I/O accesses. For that case, some applications anticipate data reading by prefetching data. Mainly users handle those operations because file system and application operate separately. For massive parallel computer, some researchers have obtained good performances by delegating I/O operations to avoid disk contentions and segregate computations from I/O.

Considering the coming of massive parallel computer and massive I/O, we propose a system that combines the advantage of I/O delegation and user expertise. Expertise describes task dependencies and I/O related through a graph that allows to make cooperation between application and I/O system. This has led to the design an asynchronous and smart I/O delegated system, named ASIODS.

This talk presents the main idea and design of ASIODS. First evaluation results are also presented and demonstrate a potential gain over 10% of execution time.