

Dynamic Hierarchical Storage

Proposed by: Jean-Thomas Acquaviva, DDN France

Description of the problem

With the arrival of new technologies, namely SSDs and Storage Class Memories, the need for hierarchies in storage systems has dramatically increased. As a consequence several technologies of burst buffers have arisen in the last year. While extremely efficient to tolerate the overhead induced by most workloads, these schemes do not fully leverage the ability to store data directly within the compute node.

Current implementation of burst buffer still relies on the concept of delegation servers and to not leverage the ability for every compute node to store locally data on a persistent non volatile storage. The purpose of this study is to propose new distributed algorithms able to allow a set of compute nodes to interact cooperatively with the burst buffer layer in order to minimize the network traffic and most importantly to guarantee a scalability able to cope with Exascale class systems.

Proposed Work

- Bibliographic survey on current burst buffer implementations and challenges.
- Bibliographic survey on distributed algorithms for peer to peer data structure sharing (DHT)
- Design and implementation of a communication protocol to allow compute node to share information
- Proof of this protocol using synchronous product and other classic theoretical tools
- Validation of the prototype using existing software stack and performance analysis tools

A metric of success will be the acceptance of a publication in a scientific conference.

Skills required

- Distributed algorithms,
- Knowledge of C, or willingness to learn it;
- Knowledge of MPI