



Contribution ID: 12

Type: **Talk**

High Performance Data Access: the Hermes approach

Monday, 2 December 2019 11:00 (30 minutes)

High-performance computing (HPC) applications generate massive amounts of data. However, the performance improvement of disk-based storage systems has been much slower than that of memory, creating a significant I/O performance gap. To reduce the performance gap, storage subsystems are under extensive changes, adopting new technologies and adding more layers into the memory/storage hierarchy. With a deeper memory hierarchy, the data movement complexity of memory systems is increased significantly, making it harder to utilize the potential of the deep memory-storage hierarchy (DMSH) architecture. In this talk, we present the development of Hermes, an intelligent, multi-tiered, dynamic, and distributed I/O caching system that utilizes DMSH to significantly accelerate I/O performance. Hermes is a US NSF supported large software development project. It extends HPC I/O stacks to integrated memory and parallel I/O systems, extends the widely used Hierarchical Data Format (HDF) and HDF5 library to achieve application-aware optimization in a DMSH environment, and enhances caching systems to support vertical and horizontal non-inclusive caching in a distributed parallel I/O environment. We will introduce the Hermes' design and implementation; discuss its uniqueness and challenges; and present some initial implementation results.

Supported Student

Primary author: Prof. SUN, Xian-He

Presenter: Prof. SUN, Xian-He

Session Classification: HPC Technology

Track Classification: Storage and IO