An Innovative Parallel Cloud Storage System using OpenStack's Swift Object Store and Transformative Parallel I/O Approach

Hsing-Bung Chen,
Benjamin McClelland, HPC-5;
David Sherrill, HPC-3;
Alfred Torrez, HPC-1;
Parks Fields, HPC-5;
Pamela Smith, HPC-3

Cloud systems (computing, service, storage, architecture, and infrastructure) are very popular in the Information Technology (IT) community today. At LANL, we have worked on high-performance computing (HPC) systems for many years. The LANL parallel log file system (PLFS) has demonstrated its superior capability for the conversion of logical N-to-1 parallel I/O operations into physical N-to-N parallel I/O operations on HPC production systems. In this article, we describe the leveraging of the scaling capability of cloud object storage systems and the transformative parallel I/O feature (Fig. 1) of the LANL PLFS and the building of a parallel cloud storage system.

We used the Swift Object Store from OpenStack as our disk-based cloud storage system. OpenStack Object Storage (code-named Swift) is open-source software for creating redundant, scalable object storage using clusters of standardized servers to store petabytes of accessible data. It is not a file system or real-time data storage system, but rather a long-term storage system for a more permanent type of static data that can be retrieved, leveraged, and then updated if necessary.

We have applied a Fuse-based file system, S3QL, on top of the Switch object storage store. S3QL is an active Python/Fuse-based file system that can run on top of S3-type storage systems. These systems can transform the S3 storage system web service interface into a full-featured UNIX file system access interface. They serve as file system interfaces between the LANL PLFS and OpenStack's Switch Object Store. The proposed software architecture for this parallel cloud storage system is shown in Fig. 2.

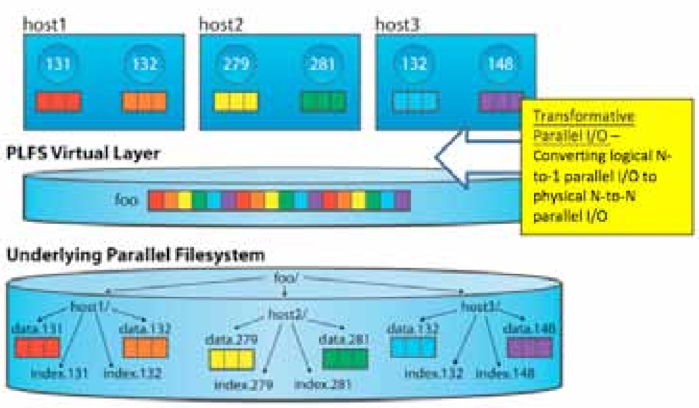

Fig. 1. Transformative parallel I/O.

We have successfully integrated the PLFS transformative parallel I/O features with the Openstack Swift Object Store. We also conducted various performance studies of I/O bandwidth and systems maintenance and management issues. We have demonstrated scaling performance results and justified performance issues. The PLFS parallel I/O scaling performance testing results are shown in Fig. 3.

As HPC systems move into the exascale computing era, the archiving exabytes of data becomes a real challenging issue. Simply buying more tapes and hard drives for storage is not a viable solution. We believe that merging advanced features from both HPC systems and Cloud systems is a promising direction.

In the future we would like to continue our research in several areas: (1) test with different data compressing algorithms; (2) examine the impact of cache size on CPU utilization, power consumption, and input/output bandwidth; (3) conduct performance testing on large-scale machines; (4) investigate other cloud storage systems (such as GlusterFS and CEPH file systems) and integrate them with the PLFS transformative parallel I/O feature; (5) enhance the S3QL with transformative parallel I/O functions; (6) apply the proposed Parallel Cloud Storage system architecture to Massive Array of Idle Disk (MAID) systems; and (7) implement dynamic power management features.

For more information contact Hsing-Bung Chen at hbchen@lanl.gov.

Funding Acknowledgments
DOE; National Science Foundation

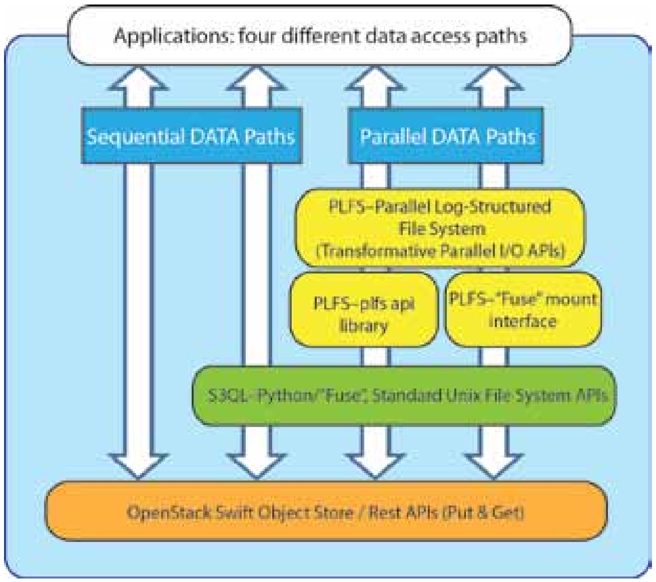

Fig. 2. Software architecture of the proposed Parallel Cloud Storage System.

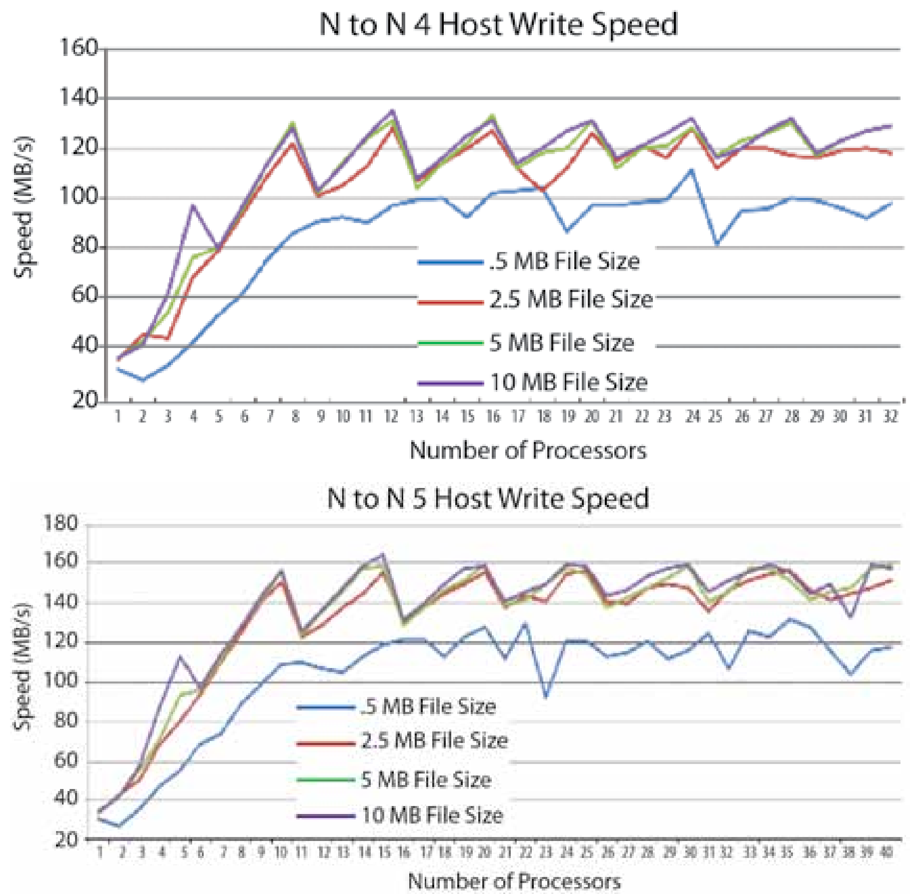

Fig. 3. PLFS parallel I/O scaling testing–multiple concurrent write session testing case.

Bibliography

Bent, J. et al., "PLFS: a Checkpoint Filesystem for Parallel Applications," *SC 09 Proceedings Conf High Performance Computing Networking, Storage and Analysis* (2009).

Manzanares A. et al., "The Power and Challenges of Transformative I/O," *IEEE Cluster 2012* (2012).

"OpenStack–Open Source Software for Building Private and Public Clouds," www.openstack.org

OpenStack Compute, *Starter Guide,* Openstack Document Essex Release (2012).

"S3QL–a Full-Featured File System for Online Data Storage," http://code.google.com/p/s3ql/, www.rath.org

Armburst, M. et al., "Above the Clouds: A Berkeley View of Cloud Computing," Technical Report UCB/EECS-2009-28, University of California, Berkeley (2009).

Xiaoming,G. et al., "A Survey of Cloud Storage Systems," CSCI-B649 Cloud Computng Project Report, Indiana University (2010).

Shankaran, N. and Sharma, R. "Cloud Storage Systesm-A Survey," CSCI-B649 Cloud Computing Project Report, Indiana University (2011).

Dewan, H. and R.C. Hansdah, "A Survey of Cloud Storage Facilities," *2011 IEEE World Congress on Services* (2011).