

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 23

Type: **oral presentation**

Current Status of the Ceph Based Storage Systems at the RACF

Thursday, April 16, 2015 11:30 AM (15 minutes)

Ceph based storage solutions are becoming increasingly popular within the HEP/NP community over the last few years. With the current status of the Ceph project, both its object storage and block storage layers are production ready on a large scale, and even the Ceph file system (CephFS) storage layer is rapidly getting to that state as well. This contribution contains a thorough review of various functionality, performance, and stability tests performed with all three (object storage, block storage and file system) levels of Ceph by using the resources of the RHIC and ATLAS Computing Facility (RACF) at Brookhaven National Laboratory (BNL) in 2012-2014 on various hardware platforms (including HP Moonshot) and with different networking solutions based on 10/40 Gbps Ethernet and IPoIB/4X FDR Infiniband interconnect technology. We also report the current status of a 1 PB scale (by usable capacity, taking into account the internal data replication factor 3x) Ceph based object storage system provided with Amazon S3 compliant RADOS Gateway interfaces that was built for the RACF in 2013, as well as the first results on its performance and the experience obtained while operating it in production over the last 8 months.

Primary authors: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); Dr ITO, Hironori (Brookhaven National Laboratory (US))

Co-authors: Mr HOLLOWELL, Christopher (Brookhaven National Laboratory (US)); Mr RAO, Tejas (Brookhaven National Laboratory (US)); Dr WONG, Tony (Brookhaven National Laboratory (US))

Presenter: Dr ITO, Hironori (Brookhaven National Laboratory (US))

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access