



Contribution ID: 39

Type: **not specified**

Ceph Based Storage Systems for RACF

Thursday, 16 October 2014 10:00 (30 minutes)

Ceph based storage solutions are becoming increasingly popular within the HEP/NP community over the last few years. With the current status of Ceph project, both object storage and block storage layers are production ready on a large scale, and the Ceph file system storage layer (CephFS) 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 RACF computing resources in 2012-2014 on various hardware platforms (including HP Moonshot) and with different networking solutions (10/40 GbE and IPoIB/4X FDR Infiniband based). We also report the status of commissioning a large scale (1 PB of usable capacity, 4.0k HDDs behind the RAID arrays by design) Ceph based object storage system provided with AMZ/S3 compliant RadosGW interfaces which is currently being finalized within RACF, and the early performance results obtained with it.

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

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

Presenter: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

Session Classification: Storage and Filesystems

Track Classification: Storage & Filesystems