



Contribution ID: 23

Type: **not specified**

Ceph Based Storage Systems at the RACF

Wednesday, 19 October 2016 14:50 (25 minutes)

We give a report on the status of Ceph based storage systems deployed at the RHIC & ATLAS Computing Facility (RACF) that are currently providing 1 PB of data storage capacity for the object store (with Amazon S3 compliant Rados Gateway front end), block storage (RBD), and shared file system (CephFS with dCache/GridFTP front-ends) layers of Ceph storage system. The hardware and software upgrades performed over the duration of the last year are reported, including the results of performance tuning for the Rados Gateway subsystem of the cluster in order to support the high concurrency (up to 24k simultaneous connections), high granularity (about 1-10 MB payloads per client session), and high bandwidth (up to 1 GB/s of aggregate bandwidth on the WAN) data transfers via Amazon S3 compatible API in order to match the growing requirements of the ATLAS Event Service. The results of boosting the performance of our Ceph clusters using the low latency PCIe NVMe SSD storage devices and the future plans for our Ceph based storage systems are also discussed.

Primary author: ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

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

Session Classification: Storage and Filesystems

Track Classification: Storage & Filesystems