

Contribution ID: 14

Type: **not specified**

CephFS in an HTC cluster and VMs on Ceph RBD with TRIM and differential backups in Bonn

Tuesday, 17 September 2019 16:50 (25 minutes)

CephFS is used as the shared file system of the HTC cluster for physicists of various fields in Bonn since beginning of 2018.

The cluster uses IP over InfiniBand. High performance for sequential reads is achieved even though erasure coding and on-the-fly compression are employed.

CephFS is complemented by a CernVM-FS for software packages and containers which come with many small files.

Operational experience with CephFS and exporting it via NFS Ganesha to users' desktop machines, upgrade experiences, and design decisions e.g. concerning the quota setup will be presented.

Additionally, Ceph RBD is used as backend for a libvirt/KVM based virtualisation infrastructure operated by two institutes replicated across multiple buildings.

Backups are performed via regular snapshots which allows for differential backups using open-source tools to an external backup storage.

Via file system trimming through VirtIO-SCSI and compression of the backups, significant storage is saved.

Writeback caching allows to achieve sufficient performance. The system has been tested for resilience in various possible failure scenarios.

Primary author: FREYERMUTH, Oliver (University of Bonn (DE))

Co-author: WIENEMANN, Peter (University of Bonn (DE))

Presenter: FREYERMUTH, Oliver (University of Bonn (DE))