

dCache on steroids - delegated storage solutions

Monday 10 October 2016 14:00 (15 minutes)

For over a decade, dCache.ORG has provided robust software that is used at more than 80 Universities and research institutes around the world, allowing these sites to provide reliable storage services for the WLCG experiments and many other scientific communities. The flexible architecture of dCache allows running it in a wide variety of configurations and platforms - from all-in-one Raspberry-Pi up to hundreds of nodes in multi-petabyte infrastructures.

Due to lack of managed storage at the time, dCache implemented data placement, replication and data integrity directly. Today, many alternatives are available: S3, GlusterFS, CEPH and others. While such systems position themselves as scalable storage systems, they can not be used by many scientific communities out of the box. The absence of specific authentication and authorization mechanisms, the use of product specific protocols and the lack of namespace are some of reasons that prevent wide-scale adoption of these alternatives.

Most of these limitations are already solved by dCache. By delegating low level storage management functionality to the above mentioned new systems and providing the missing layer through dCache, we provide a system which combines the benefits of both worlds - industry standard storage building blocks with the access protocols and authentication required by scientific communities.

In this presentation, we focus on CEPH, a popular software for clustered storage that supports file, block and object interfaces. CEPH is often used in modern computing centres, for example as a backend to OpenStack services. We will show prototypes of dCache running with a CEPH backend and discuss the benefits and limitations of such an approach. We will also outline the roadmap for supporting 'delegated storage' within the dCache releases.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Object stores

Primary Keyword (Mandatory)

Storage systems

Primary authors: BEHRMANN, Gerd (NDGF); FUHRMANN, Patrick (DESY); MILLAR, Paul; MKRTCHYAN, Tigran

Presenter: MKRTCHYAN, Tigran

Session Classification: Track 4: Data Handling

Track Classification: Track 4: Data Handling