Tape Recall Scheduling Optimization in the Storage System dCache

Wednesday 27 October 2021 09:25 (25 minutes)

Given the anticipated increase in the amount of scientific data, it is widely accepted that primarily disk based storage will become prohibitively expensive. Tape based storage, on the other hand, provides a viable and affordable solution for the ever increasing demand for storage space. Coupled with a disk caching layer that temporarily holds a small fraction of the total data volume to allow for low latency access, it turns tape based systems into active archival storage (write once, read many) that imposes additional demands on data flow optimization compared to traditional backup setups (write once, read never). In order to preserve the lifetime of tapes and minimize the inherently higher access latency, different tape usage strategies are being evaluated.

As an important disk storage system for scientific data that transparently handles tape access, dCache is making efforts to contribute to tape recall optimization by introducing a high-level tape recall request scheduling component within its SRM implementation. This presentation will include first experiences with this component on scientific data in a production environment.

Desired slot length

Speaker release

Yes

Primary authors: ROSSI, ALBERT (Fermi National Accelerator Laboratory); LITVINTSEV, Dmitry (Fermi National Accelerator Lab. (US)); CHITRAPU, Krishnaveni (National Supercomputer Center, Linköping University, Sweden); MORSCHEL, Lea (Deutsches Elektronen-Synchrotron DESY); SAHAKYAN, Marina; MILLAR, Paul; MEYER, Svenja (Deutsches Elektronen-Synchrotron DESY); MKRTCHYAN, Tigran (DESY); GARONNE, Vincent (Brookhaven National Laboratory (US))

Presenter: MORSCHEL, Lea (Deutsches Elektronen-Synchrotron DESY)

Session Classification: Storage & File Systems

Track Classification: Storage & Filesystems