

Improving Performance of Tape Restore Request Scheduling in the Storage System dCache

Tuesday, 18 May 2021 15:52 (13 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 evaluate its recall optimization potential and is introducing a proof-of-concept, high-level stage request scheduling component within its SRM implementation.

Primary authors: MORSCHEL, Lea (Deutsches Elektronen-Synchrotron DESY); Ms CHITRAPU, Krishnaveni (National Supercomputer Centre, Linköping University); GARONNE, Vincent (University of Oslo (NO)); LITVINSEV, Dmitry (Fermi National Accelerator Lab. (US)); Ms MEYER, Svenja (Deutsches Elektronen-Synchrotron DESY); MILLAR, Paul; MKRTCHYAN, Tigran (DESY); ROSSI, ALBERT (Fermi National Accelerator Laboratory); SAHAKYAN, Marina

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

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities