

Best Practices in Accessing Tape-Resident Data in HPSS

Tuesday 10 July 2018 16:45 (15 minutes)

Tape is an excellent choice for archival storage because of the capacity, cost per GB and long retention intervals, but its main drawback is the slow access time due to the nature of sequential medium. Modern enterprise tape drives now support Recommended Access Ordering (RAO), which is designed to improve recall/retrieval times.

BNL's mass storage system currently holds more than 100 PB of data on tapes, managed by HPSS. Starting with HPSS version 7.5.1, a new feature called "Tape Order Recall (TOR)" has been introduced. It uses the RAO mechanism for improving the access times over sequential media claiming a performance improvement by 30% to 60%.

Prior to HPSS 7.5.1, we have been using an in-house developed scheduling software, called ERADAT. ERADAT accesses files based on the order of the logical position of each file and consequently has demonstrated great performance at BNL. We have compared the new TOR method to the method of using logical position under different conditions such as number of access requests.

In this presentation we will demonstrate a series of tests, which indicate how effective the TOR (RAO) is under different scenarios and what are the best methods in restoring data from tape storage under different conditions.

Authors: Dr CHE, Guangwei (Brookhaven National Lab); YU, David (Brookhaven National Laboratory (US))

Co-authors: Mr CHOU, Tim (Brookhaven National Lab); Mr NOVAKOV, Ognian (Brookhaven National Lab)

Presenter: YU, David (Brookhaven National Laboratory (US))

Session Classification: Posters

Track Classification: Track 4 - Data Handling