21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



21st International Conference on Computing in High Energy and Nuclear Physics CHEP2015 Okinawa Japan: April 13 - 17, 2015

Contribution ID: 2

Type: oral presentation

Mean PB to Failure – Initial results from a long-term study of disk storage patterns at the RACF

Tuesday 14 April 2015 16:30 (15 minutes)

The RACF (RHIC-ATLAS Computing Facility) has operated a large, multi-purpose dedicated computing facility since the mid-1990's, serving a worldwide, geographically diverse scientific community that is a major contributor to various HEPN projects. A central component of the RACF is the Linux-based worker node cluster that is used for both computing and data storage purposes. It currently has nearly 50,000 computing cores and over 23 PB of storage capacity distributed over 12,000+ (non-SSD) disk drives.

The majority of the 12,000+ disk drives provides a cost-effective solution for dCache/xRootd-managed storage, and a key concern is the reliability of this solution over the lifetime of the hardware, particularly as the number of disk drives and the storage capacity of individual drives grow. We report initial results of a long-term study to measure lifetime PB read/written to disk drives in the worker node cluster. We discuss the historical disk drive mortality rate, disk drive manufacturers' published MPBTF (Mean PB to Failure) data and how they are correlated to our results.

The results helps the RACF understand the productivity and reliability of its storage solutions and has implications for other highly-available storage systems (NFS, GPFS, CVMFS, etc) with large I/O requirements.

Primary author: Dr WONG, Tony (Brookhaven National Laboratory)

Co-authors: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); HOLLOWELL, Christopher (Brookhaven National Laboratory); CARAMARCU, Costin (Brookhaven National Laboratory (US)); Mr RAO, Tejas (Brookhaven National Laboratory); STRECKER-KELLOGG, William (Brookhaven National Lab)

Presenter: HOLLOWELL, Christopher (Brookhaven National Laboratory)

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access