Subject: Checksumming files on tape
From: Bryan Hess <bhess@jlab.org>
Date: 7/6/20, 1:45 PM
To: Brad Sawatzky <brads@jlab.org>, Ole Hansen <ole@jlab.org>, Nathan Baltzell <baltzell@jlab.org>, Mark Ito <marki@jlab.org>
CC: David Lawrence <davidl@jlab.org>, Graham Heyes <heyes@jlab.org>, Jie Chen <chen@jlab.org>, Thomas Britton <tbritton@jlab.org>

Hello,

Chris Larrieu wrote the summary below about our plans to switch from MD5 to CRC32 checksumming on tape. Short version: It's getting hard to keep up with tape speed doing MD5 in band, and CRC32 has been sufficient. We are planning to implement this change on Tuesday, July 21st as part of the normal maintenance day cycle. Please let me know if you have any concerns or questions.

> *Starting in 2012 we started calculating MD5 hashes along with CRC checksums for every file placed into the tape library. For raw data writes the MD5 hash is calculated at the source machine and verified during tape write. For other writes it is only calculated while the data is being written to tape. Calculating an MD5 hash is computationally intensive; keeping up with higher data rates requires ever faster processors (the calculation is inherently serial, so is bound by clock speed and memory access rates). This has imposed a heavy CPU cost on some DAQ machines as well as the tape drive data movers.*
>
> *After further assessing the value of using MD5 hashes to verify data integrity and the associated cost in hardware requirements, we've come to the conclusion that CRC32 checksums are more than adequate for our purpose. They can be calculated essentially for free because they require very little CPU power. And they are reliable: the theoretical worst case performance of an n-bit cyclic redundancy check with respect to false positives is 2^-n, regardless of message size. This means that, using 32-bit CRC, we'll have one silently corrupted file out of every ~4 billion. It will be correct ~99.99999998% of the time.*
>
> *We have experienced several data corruption issues over the years. In every case the corruption was identified via CRC checksum disparity. We've never had a situation where CRC checksum verified but MD5 failed. We intend to disable MD5 hash calculations for future tape writes. This will be mostly invisible to users, with the exception that the 'md5' entry in /mss stub files will be absent for new files. We will need to update the jmirror software for the halls as well.*