Muon g-2 Online FTS

Adam Lyon
March 2022

The Muon g-2 experiment chose the MIDAS[1] data acquisition system around 2012, before `artdaq` was a mature product. Therefore, the DAQ itself is not supported by SCD. The computing hardware is supported by the SCD SLAM group, and the mechanism for transferring data from the experiment in MC-1 to tape is supported by the SCD Scientific Data Services department. This mechanism is described in this document.

MIDAS collects data from each 700 μs muon *fill* and there are 12 fills per second. MIDAS writes a file about every 10 seconds. Each file is a *subrun* and is ~2 GB in size. Every few hours, MIDAS stops the current *run* and starts a new run. The subruns are thus grouped by run number. There may be several hundred subruns in a run corresponding to 500 GB – 1,000 GB of data.

For efficiency, the data from MIDAS are written to a SSD drive. There are two such drives, `/dataSSD1` and `/dataSSD2` and each drive is about 3 TB in size. Data for a run is written to one SSD drive and then MIDAS switches to the other SSD drive for the next run. As the next run is collected, the data from the previous run is copied to a larger spinning disk drive, `/data1` or `/data2`. Each of these larger disks are 28 TB in size. This configuration is depicted below. The large disks together can hold over three days worth of data from the experiment. There is a separate MIDAS system that collects the g-2 ring magnetic field information. Making a 2 GB file every 5 minutes, the data volume from this system is insignificant compared to the main MIDAS DAQ.

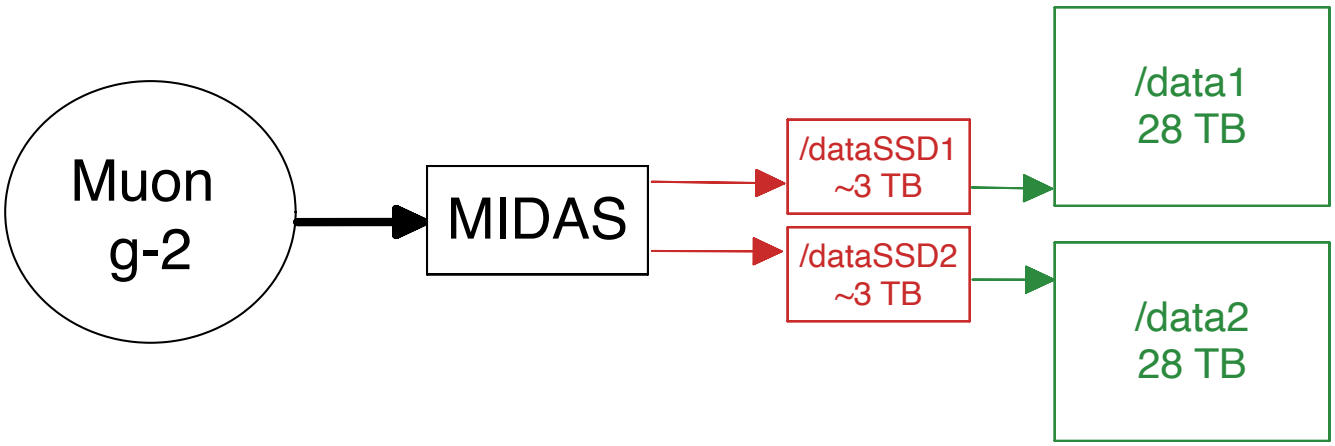

The Fermi File Transfer Service (FTS)[2] is a system for automatically and reliably transferring data from a source to a destination. It scans source directories at regular

intervals, notes new files that have appeared, runs a plugin for each new file to declare metadata to SAM, transfers each file to the destination in a reliable way, and if that destination is tape, it ensures that the file has a tape label in SAM. FTS may be configured to delete the file from the source once it is confirmed to be on tape. FTS runs on a machine on the MC-1 private network that also has a network port to the main Fermilab network for transferring files to the data center.

There are three main categories of files

- MIDAS data files
- Log files
- Nearline histogram files

The MIDAS data files are the subrun and field files and are the irreplaceable data from the experiment. The Nearline files are small histogram files with data quality information that are used for validation. Many of the nearline files and the log files go through the SCD Small File Aggregator.

The Files waiting for tape plot shows the four hour cycle of FTS checking for large files on tape (the little sawtooth pattern). The large rise and drops are the small nearline and log files. It can take FTS a long time (sometimes days) to confirm that a small file landed on tape.

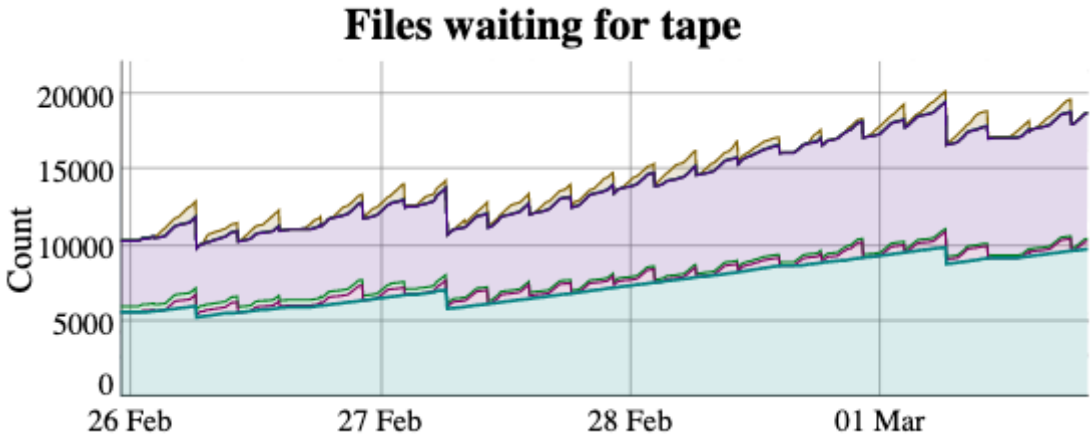

The destinations are areas in `/pnfs` that are tape backed. The MIDAS files go to an area with its own dCache write pool (54 TB - equivalent to the two large DAQ data disks) with a file family width of 6. See the figure below.

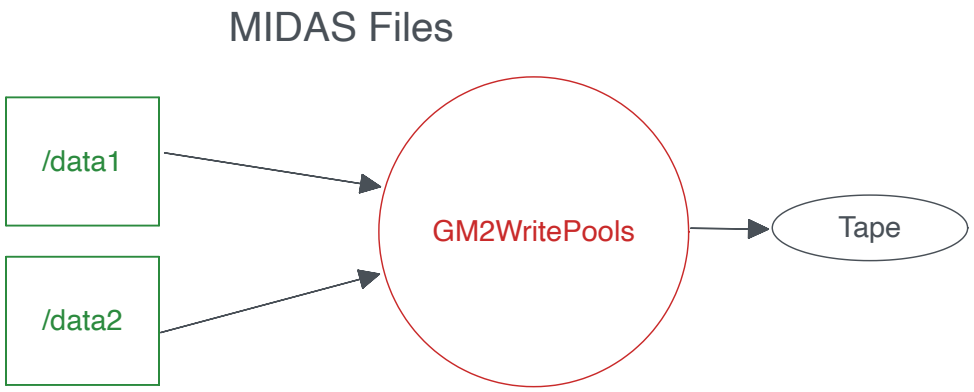

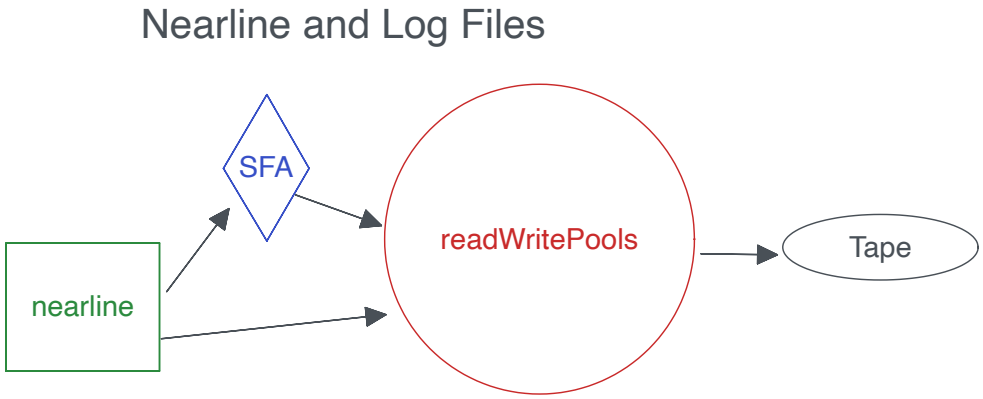

This separate write pool ensures that these essential files are written to tape in a timely manner even if the dCache system is busy. We could delete the files from the `/data1` and `/data2` DAQ disks once they are transferred to dCache (e.g. not wait for them to be written to tape). However, for extra redundancy we leave the files on the DAQ disks until the tape confirmation. If there were a long term tape or library problem (such as when both robot arms were broken), we could switch to immediate deletion and thus have enough storage for data corresponding to about a week.

---

1. https://midas.triumf.ca ↩
2. https://cdcvs.fnal.gov/redmine/projects/filetransferservice/wiki/Wiki ↩