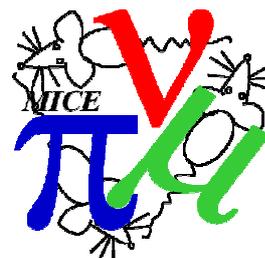


Moving Files from the Online Buffer to micestore – User's Guide



J.J. Nebrensky (Brunel University, Uxbridge UB8 3PH, UK)

1 Introduction

- The Event Builder(s) will produce a series of ~250 MB files per run saved on the Online Buffer (currently either miceraid4 or miceraid5 at any one time).
- The files are named $x.y$ where x is the run number and y starts at 000 and increments to 001 , 002 , etc. The top level of organisation is a single directory per MICE step with a new subdirectory created every hundredth run; thus data from run 987 will be stored in `.../MICE/StepS/00900/00987.0nn`.
- The output from the Online applications are arranged in a hierarchy corresponding to the global namespace used for the data. The files are named `On<app>.00987.*`, e.g. `OnMon.00987.root`, `OnScalers.00987.txt` etc. for run 987. Henry should be notified when a new application is added - the applications currently handled are Monitoring and Reconstruction.
- Both the DAQ and Online data are initially stored in a variety of locations. Therefore, the datamover must rename and rearrange the files into the required hierarchies on the Buffer and the Online Reconstruction machines respectively.
- The DAQ data and Online applications output (if available) are all then combined into a single tarball per run, stored on a separate RAID system (*micestore*). This process is known as the “File Compactor”.
- Both the internal file moving and the tarball creation are done using a pre-prepared shell script – the `moveFiles` script – which must be run at a suitable opportunity.
- The Grid Transfer Box NFS-mounts the relevant area on micestore and uploads the tarballs to the Grid automatically soon after their creation.
- The Online Buffer also NFS-mounts the relevant OnRec machine, in order to include the various Monitoring, etc. output into the tarballs.
- The `moveFiles` script will normally be triggered at the end of each DAQ run by the Run Control application. It may also be run by hand, e.g. to catch up a backlog after any problems.

For more details, see MICE Notes 247, 252, 255 or 264, and eLog entries #821, #834 and #3138.

2 Installation/Cold Starting

The present configuration on MICE is that the Event Builder (DATE GDC) and Online Buffer are run on a single machine, which may be either of miceraid4 or miceraid5.

The files `fileCompactorWrapEPICS.sh` and `moveFiles-manual.sh` must exist in the “daq” user’s home directory on miceraid4 and miceraid5. As shell scripts, they should have execute permissions set (755).

The scripts should have been automatically fetched from bazaar when the machine was configured.

The `fileCompactorWrapEPICS.sh` script is called remotely by Run Control, and *should never be run by hand*.

In order for the script to access all required data, the Online Buffer (miceraid4 or miceraid5) must be able to see the relevant files systems on both micestore (output tarballs) and the Online PC (OnRec01). This is implemented using NFS – see eLog entry #3138. It is therefore necessary to check that the required NFS-mounts exist:

On the Online Buffer in use (miceraid4 or miceraid5);

- Display the current list of mounted filesystems:

```
ls /storage/for/MICE
ls /storage/for/Online
mount
```

- Check that lines exist in the response starting with either micestore and miconrec01, or their IP numbers (as listed in MICEmine).

If not, log on *as root* on the Online Buffer and check that the `autofs` service is running:

```
service autofs restart
```

3 Preparation for shift

3.1 Warm Start

If any DAQ machines have been rebooted since the previous shift, check NFS mounts as above.

3.2 Function Test

Not really necessary – you could start and stop the DAQ to create a small data file, and run the manual script (as below) if expecting problems.

4 Running Instructions

The File Compactor is triggered automatically at the end of each DAQ run within Run Control. Progress is shown at the bottom of the Run Control window: when each run is complete the File Compactor will cycle through “Sleeping”, “Compacting Data”, and finally “Compacted RunNumber OK” (fig. 1). Each invocation by Run Control applies to that specific run only. The File Compactor does not interact with the Configuration Database (CDB).

Occasionally the DAQ output may be less than 1 megabyte in size; below this threshold it is deemed too small to preserve (as there was likely a DAQ problem) and the script will report “Compacted RunNumber OK - tarball too small” – this is **not** an error with the File Compactor. Note also that if Run Control fails to fully carry out its end-of-run process then the File Compactor obviously will obviously not be called for that run.

In the event of a problem, an error message will be displayed and an alert raised in the Alarm Handler. The shifters do not need to interrupt data-taking to investigate (beyond making an eLog entry and notifying experts) – they should be able to use Run Control to start the next run. Data that is not successfully written into a tarball is merely delayed from distribution and will form a backlog on the Online Buffer.

If the File Compactor encounters certain errors that may lead to data corruption it will set semaphores to prevent the script being run; future invocations will then report “Error - Can't run due to earlier failure”. This will persist pending expert attention.

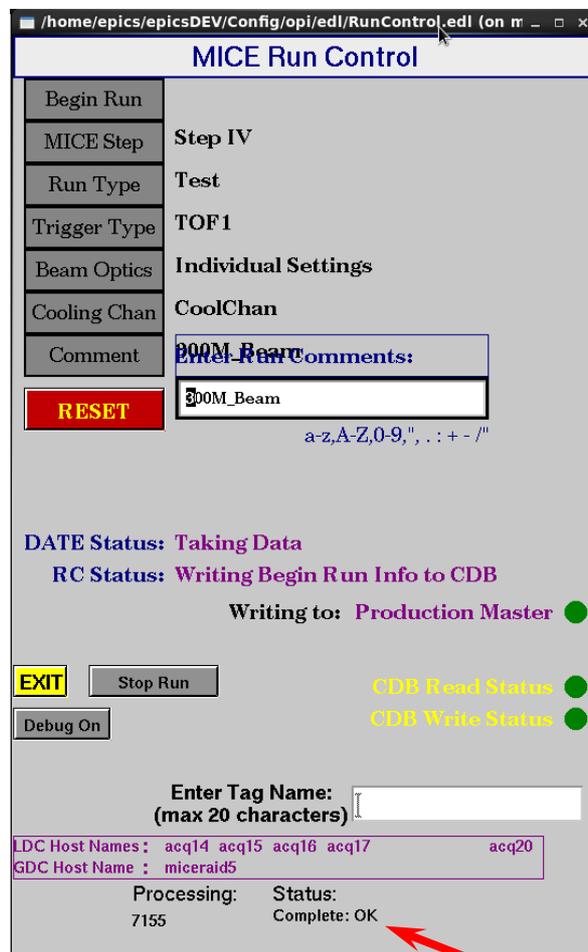


Figure 1: Run Control panel with File Compactor status indicated

If requested by the expert or to catch up with the backlog, it is possible to run the script by hand. To run the script manually:

Ensure DATE has been stopped. You must NEVER run the moveFiles script whilst data-taking is in progress.

1. Check which DAQ machine has the Event Builder (either miceraid4 or miceraid5). If unsure, consult the DAQ expert.
2. Log onto a terminal on the Event Builder (miceraid4 or 5)
 - Username: **daq**; Password: see hint in MLCR
3. Check that the Online Monitoring ROOT GUI has been restarted after the last run to be moved.
 - Note: this is necessary in order to produce the correct Online Monitoring output that is saved with the data.
4. Hit the “Start Processes” button on the DAQ control to close out the processes of the last run to be moved. Do NOT then press “Start”.
5. Check the old directory to see which runs need to be moved:
 - `ls /data/mice/*.000`
6. Run the script giving the number of the first and last runs in the set to be compacted:
 - `moveFiles-manual.sh 1181 1185`
7. You must give two numbers, but they can be the same (for a single run). There is some error checking, for simpler mistakes. The script will pause while displaying

```
Moving files for runs 1181 to 1185 inclusive
```

```
Please make sure the Online Monitoring has been
restarted from run 1185
```

- During this time, quickly check that the run numbers are correct and that there isn't a typo of e.g. O for 0. Also make sure that the DAQ and the Online Monitoring ROOT GUI have been restarted.
 - If necessary or if you are unsure, it is safe to stop the script at this point with ^C.
8. The script will then display

```
Checking run 1181
```

 - During this time, it will checksum the files, move them to the required place and then validate they are correct.

- You should see a series of 1181.000 OK messages.
 - Do not stop the script at this point. It should not permanently lose data, but may leave the run files in an inconsistent state with some in the new space and some in the old.
 - If the script has been stopped, re-run it starting from the run **after** where it stopped.
9. If the script encounters a problem it will stop showing any error messages from the command that failed
 10. If the problem looks to be minor retry with the next run as above, else for serious or repeated errors do not use the script and contact MICE Data Manager (*micedataman @stfc.ac.uk*).
 11. Once the script has finished, it will return to the command prompt.

5 End of Shift

In principle you should watch to see that the manual script finishes without errors. In practice, there may not be enough time for it to run between shift end and departure so the output should be checked the next morning.

6 Shutdown

Not applicable – the script will stop on its own.

7 Emergency Shutdown

If something goes badly wrong or the Online Buffer urgently needs rebooting, you can stop the script with ^C – there shouldn't be any data loss but the data files will be left in an inconsistent state. Contact the MICE Data Manager (micedataman@stfc.ac.uk) to check things over before running the script again.

8 Troubleshooting

If the manual script encounters a problem it will stop, showing any error messages from the command that failed.

If the automated script encounters a problem it will return an error code and message via EPICS. Script activity is copied to a log file `/tmp/DFC-RunNumber.log` on the Online Buffer.

- If the problem appears minor, check the NFS mounts as per section 2. Also check the general state of the machines – i.e. no network failures or full disks. Then retry with the next run as above.
- For serious or repeated errors do NOT use the script. Contact the MICE Data Manager (micedataman@stfc.ac.uk), currently Henry.

9 Documented Procedures

Not applicable.