

Data Transfer Rates between ARCHER and JASMIN computers

William McGinty

February 10, 2017

1 Introduction

In large UM data projects such as SWAMMA, VERA and PRIMAVERA, it is necessary to archive large data sets while the model is running. In SWAMMA, the largest output streams for the 4km model were 45GB in netCDF format. PRIMAVERA is expected to produce larger files. Currently, the aim is to have a Rose/Cylc application do this at the end of an integration cycle.

The aim here is to measure the speed of transfer between ARCHER Research Data Facility (RDF) and JASMIN to support a decision on transfer method. There are two ways of accessing the RDF (dtn01 and dtn02), two ways of accessing JASMIN for transfers (jasmin-xfer1 and jasmin-xfer2) and it is possible to push or pull data from each. The jasmin-xfer2 computer requires priveleged access. Possible transfer programs are rsync, bbcp and GridFTP.

The method was to take a 9.9GB file¹ and send it to the other computer and time the transfer, then delete the file on the target computer. Repeat a number of times and collect basic statistics. The file is deleted to prevent clever transfer tools recognising that the file already exists and preventing transfer.

To get a sense of scale, at 100MB/s this would take 101.4 seconds to transfer.

There is a 10Gbit/sec “light path” connecting JASMIN Data transfer zone to the high performance data transfer computer jasmin-xfer2. This requires priveleged access. On the other hand, jasmin-xfer1 is connected to dtn02 via an optical private network (1Gbit/sec)², available to anyone with a JASMIN account, providing dedicated bandwidth for users. For comparision, the JANET backbone, shared by many users, runs at 100Gbit/sec with an interconnect capacity of 40Gbit/sec.

¹Actually 10,582,306,816 bytes

²<http://help.ceda.ac.uk/article/221-jasmin-external-connections>

1.1 Results

1.1.1 rsync

The results ($n=10$) of pushing the file from the RDF to jasmin-xfer2 is shown in Fig 1; the result of pulling the file from the RDF is shown in Fig 2. Pulling the file is faster than pushing it. There is not much difference between dtn01 and dtn02: the difference in median transfer time is swamped by the huge variability. A transfer time of 70 seconds corresponds to a rate of 145 MB/sec.

An rsync push to jasmin-xfer1 from dtn01 took 11,670 seconds (3.24 hours).

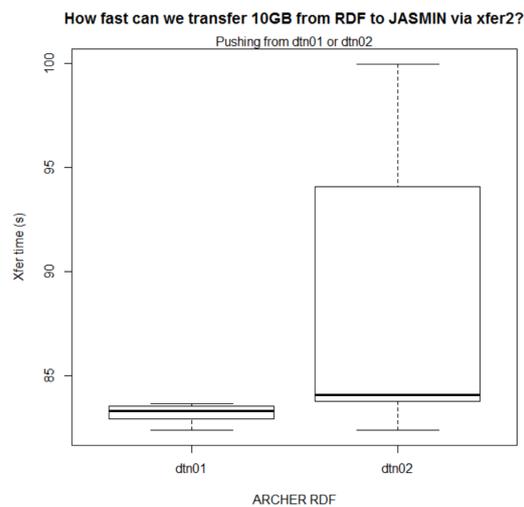


Figure 1: rsync push to RDF

1.1.2 bbcp

bbcp is a peer to peer transfer program which can use parallel streams. A bbcp program is required on each computer. In this work 8 and 16 streams were investigated.

Pulling the file from the RDF resulted in Fig 3 and Fig 4

The median transfer time is in the range 100 to 120 seconds. Again, there is large variability in the results, so that there is no real difference between dtn01 and dtn02 or between 8 and 16 streams.

Pushing from the RDF to jasmin-xfer2 results in Fig 5. Here the columns are labelled with either dtn01 or dtn02, the target jasmin and the streams.

Again, the variability swamps the differences. Transfers of around 100MB/sec are possible.

Pushing from the RDF to jasmin-xfer1 is much slower Fig 6. Transfers of around 36 MB/sec are possible.

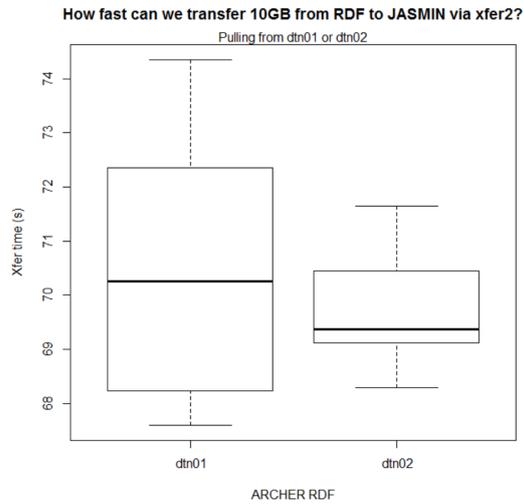


Figure 2: rsync pull from RDF

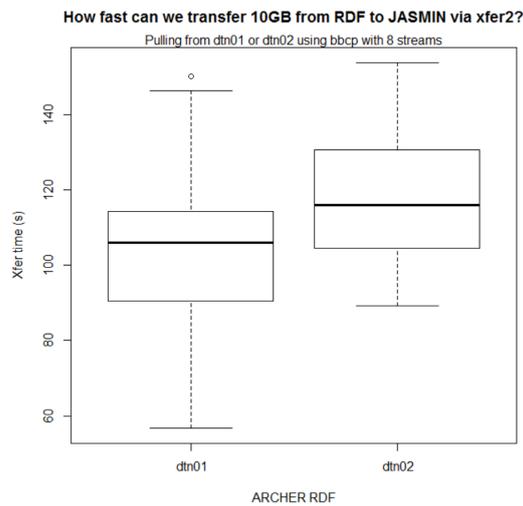


Figure 3: bbcpl pull from RDF 8 streams

1.1.3 GridFTP

GridFTP is a transfer program from Globus which can use parallel streams. In this work, 8 and 16 streams were studied.

Pushing to jasmin-xfer2 results in Fig 7. With median transfer times around 50 to 70 sec, the transfer rate is about 170 MB/sec.

Pushing to jasmin-xfer1 from dtn02 gives Fig 8.

However, notice that pushing from dtn01 to jasmin-xfer1 takes around 800 or 1600 seconds according to whether 16 or 8 streams are used.

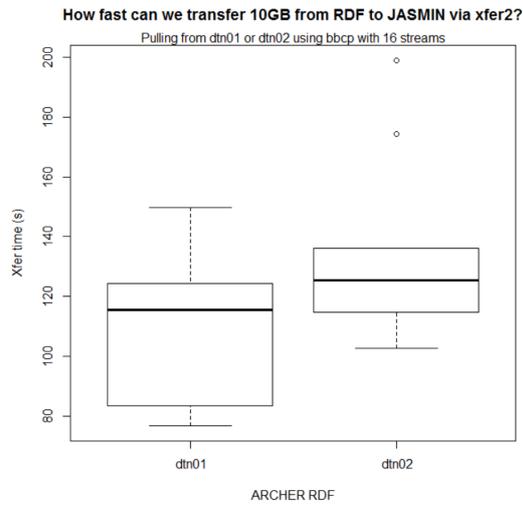


Figure 4: bbcp pull from RDF 16 streams

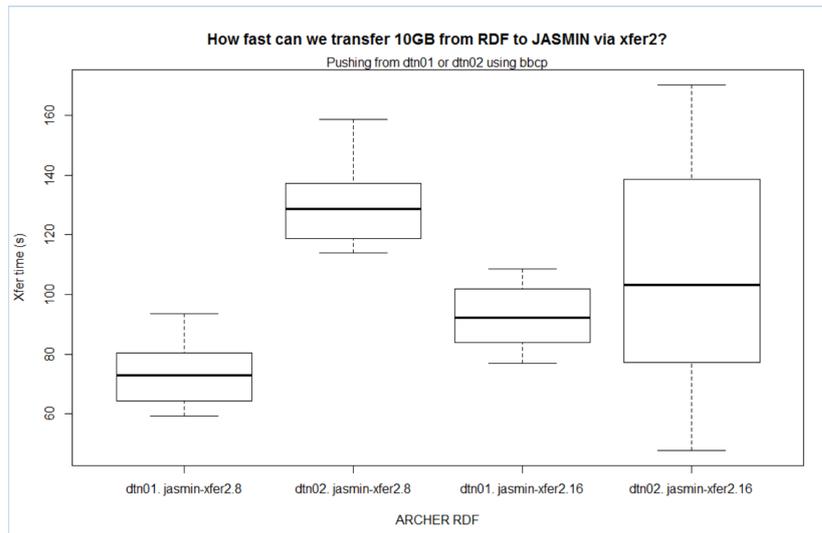


Figure 5: bbcp push from RDF

Pulling from RDF results in Fig 9.

Median transfer times range from 47 to 57 seconds, equating to a transfer rate of around 195 MB/sec.

On jasmin-xfer1, pulling with GridFTP and 8 streams results in transfer times between 250 and 430 seconds.

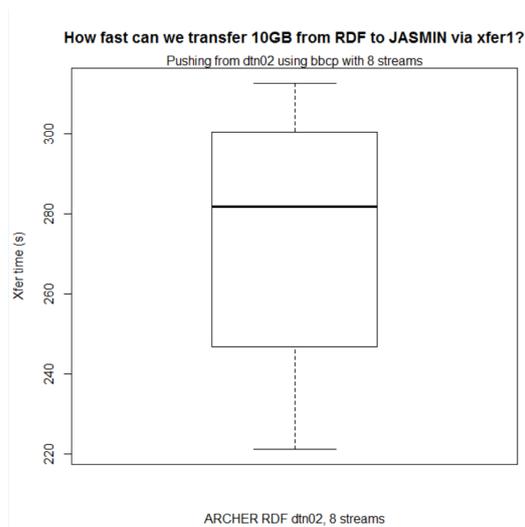


Figure 6: bbcp push from RDF

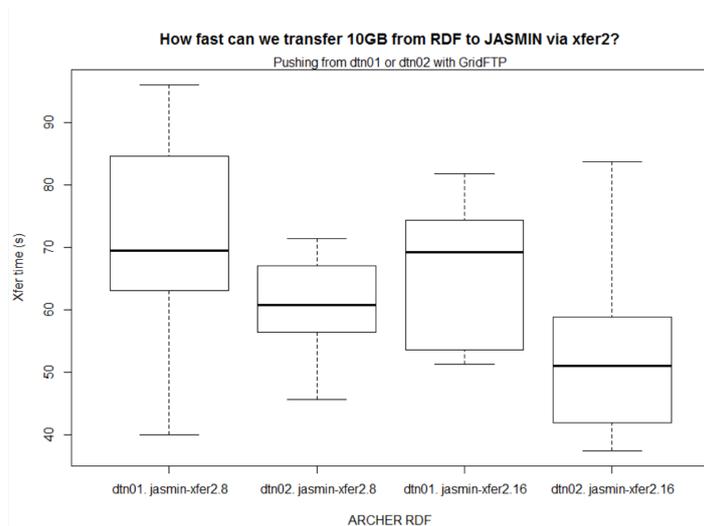


Figure 7: GridFTP push from RDF

1.2 Conclusions

The variation in the transfer rates is large due to user contention for the available bandwidths.

Transfer via jasmin-xfer2 gives the fastest transfer, but this requires privileged access.

Table 1 gives a rough idea of the range of median transfer rates observed. The differences between dtn01 and dtn02 are ignored.

The transfer rates for jasmin-xfer1 are very poor. This computer is open to all who have a JASMIN account, so the user contention for the available band width is high.

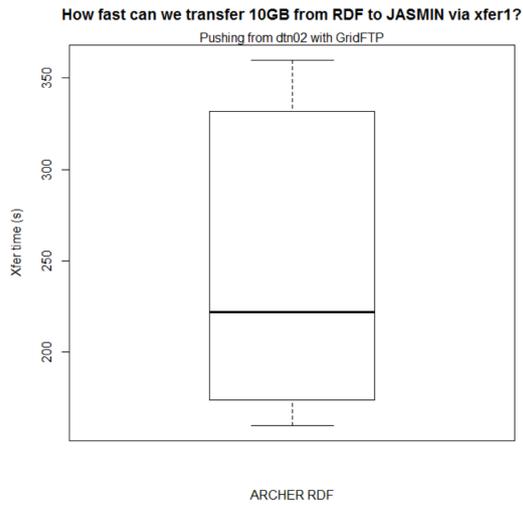


Figure 8: GridFTP push from RDF

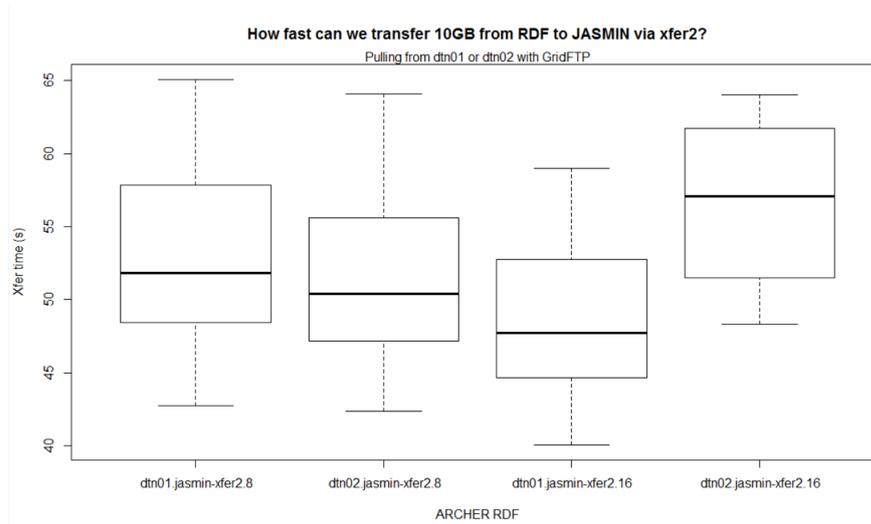


Figure 9: GridFTP pull from RDF

If it is necessary to use jasmin-xfer1, then the transfer should be from dtn02 and use GridFTP.

The best transfer rates are obtained by pushing from dtn02 to jasmin-xfer2 using GridFTP.

	jasmin-xfer1	jasmin-xfer2
rsync pull	0.8	145
rsync push	-	127
bbcp pull	-	82-95
bbcp push	36	80-150
GridFTP pull	24 - 40	175 -210
GridFTP push	46	145 - 200

Table 1: Median Data Transfer Rates in MB/s