

# NDGF tape test

2022-06-13 to 2022-06-18

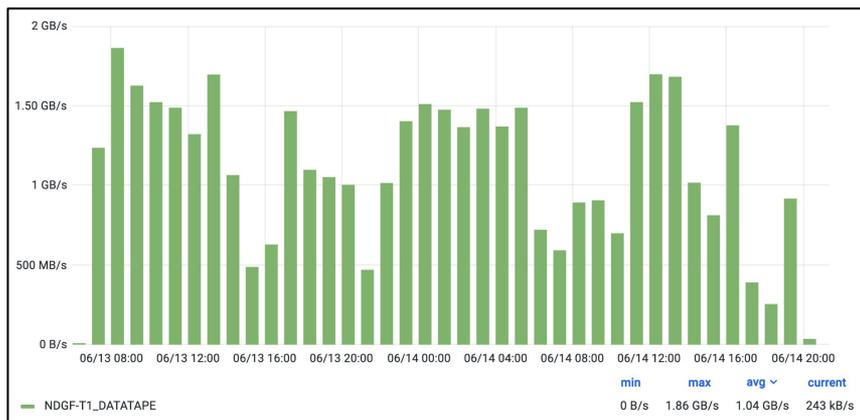
# The test

- DT test
  - Test transfer ran from Monday 2022-06-13 7am ~ Tuesday 2022-06-14 7pm (UTC)
- A-DT test
  - Ran from Thursday 2022-06-16 8am (UTC) ~
- GGUS communication ticket

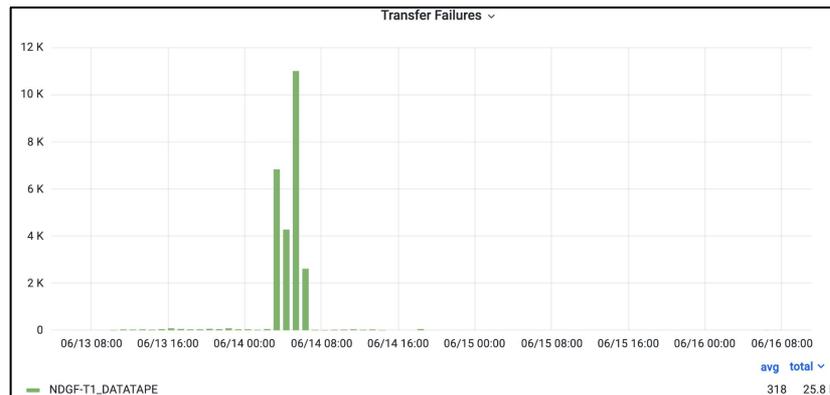
[https://ggus.eu/index.php?mode=ticket\\_info&ticket\\_id=157733](https://ggus.eu/index.php?mode=ticket_info&ticket_id=157733)

# DT test (1)

- Peak export rate from CERN EOS to NDGF DATATAPE
  - One main stream fake dataset used, total volume ~150TB, highly compressible though
- Spike of failures
  - DESTINATION SRM\_PUT\_TURL error on the turl request : [SE][PrepareToPut][SRM\_EXCEED\_ALLOCATION] Space associated with the space token 95216 is not enough to hold SURL.
  - Site comment: The space reservation 95216 ran out of space when one write pool was filling up due to slow flushing, while the other was just a third full. Technically there was space for writes, but not within the reservation. Sizing the reservation is tricky for us, because we actively manage the tape write pool group depending on which tape systems have free space (sometimes only one), and per spec we're not allowed to have a larger space reservation than we can guarantee space for writes.



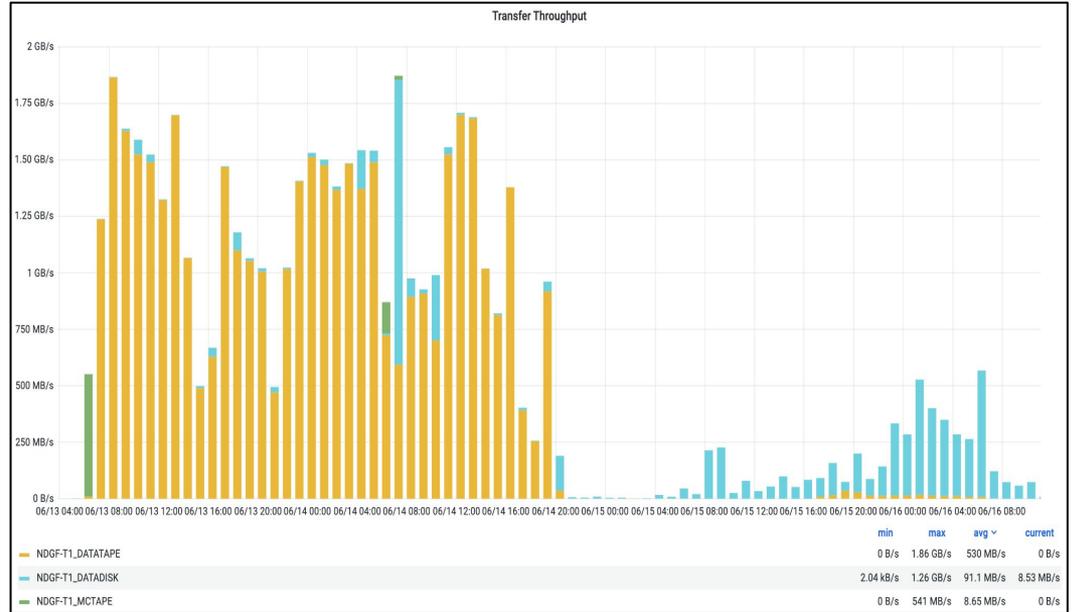
Transfer rate from CERN EOS to NDGF DATATAPE (disk buffer)



Failures during the tape write test (same time window)

# DT test (2)

- Avg rate : 1.1GB/s vs the real main stream rate of 3.5GB/s expected in Run3.
  - limited by 20Gbps network bandwidth on the NDGF side. Upgrade to 100Gbps soon ?
  - Data will be buffered longer at CERN if the export rate to T1 is less than the incoming data rate from DAQ during a run.
  - Plot shows there are some other traffics from CERN to NDGF at the same time, not much though

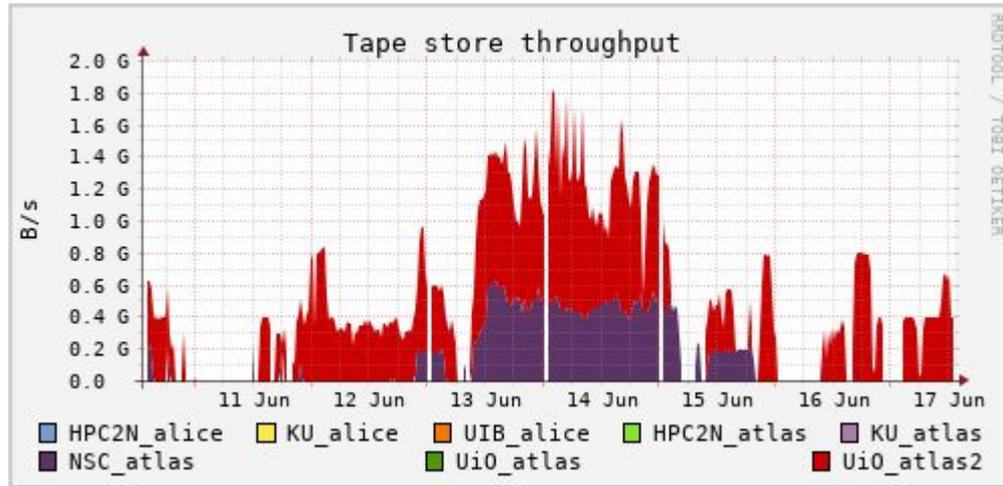


There were other traffic from CERN to NDGF during the same time window

# Site view of DT test

Week timescale, no surprises, and the CERN networking was the limiting factor.

The incoming data outside of the test period is mostly transfers from old UiO library to the new one



# A-DT test (1) : test sample

data\_test:data\_test.00414018.physics\_bulk.daq.RAW

+-----+-----+-----+-----+-----+

Total files : 31478  
Total size : 147 TB

data16\_13TeV:data16\_13TeV.00304198.physics\_Main.daq.RAW

+-----+-----+-----+-----+-----+

Total files : 9911  
Total size : 22.8 TB

data16\_13TeV:data16\_13TeV.00304409.physics\_Main.daq.RAW

+-----+-----+-----+-----+-----+

Total files : 6904  
Total size : 16.4 TB

mc16\_13TeV:mc16\_13TeV.361106.PowhegPythia8EvtGen\_AZNLOCTEQ6L1\_Zee.recon.ESD.e3601\_s3170\_r12660\_tid25676792\_00

+-----+-----+-----+-----+-----+

Total files : 500  
Total size : 3.4 TB

mc16\_13TeV:mc16\_13TeV.361107.PowhegPythia8EvtGen\_AZNLOCTEQ6L1\_Zmumu.recon.ESD.e3601\_s3170\_r12638\_tid25649553\_00

+-----+-----+-----+-----+-----+

Total files : 500  
Total size : 3.4 TB

mc16\_13TeV:mc16\_13TeV.361107.PowhegPythia8EvtGen\_AZNLOCTEQ6L1\_Zmumu.recon.ESD.e3601\_s3170\_r12659\_tid25676628\_00

+-----+-----+-----+-----+-----+

Total files : 500  
Total size : 3.4 TB

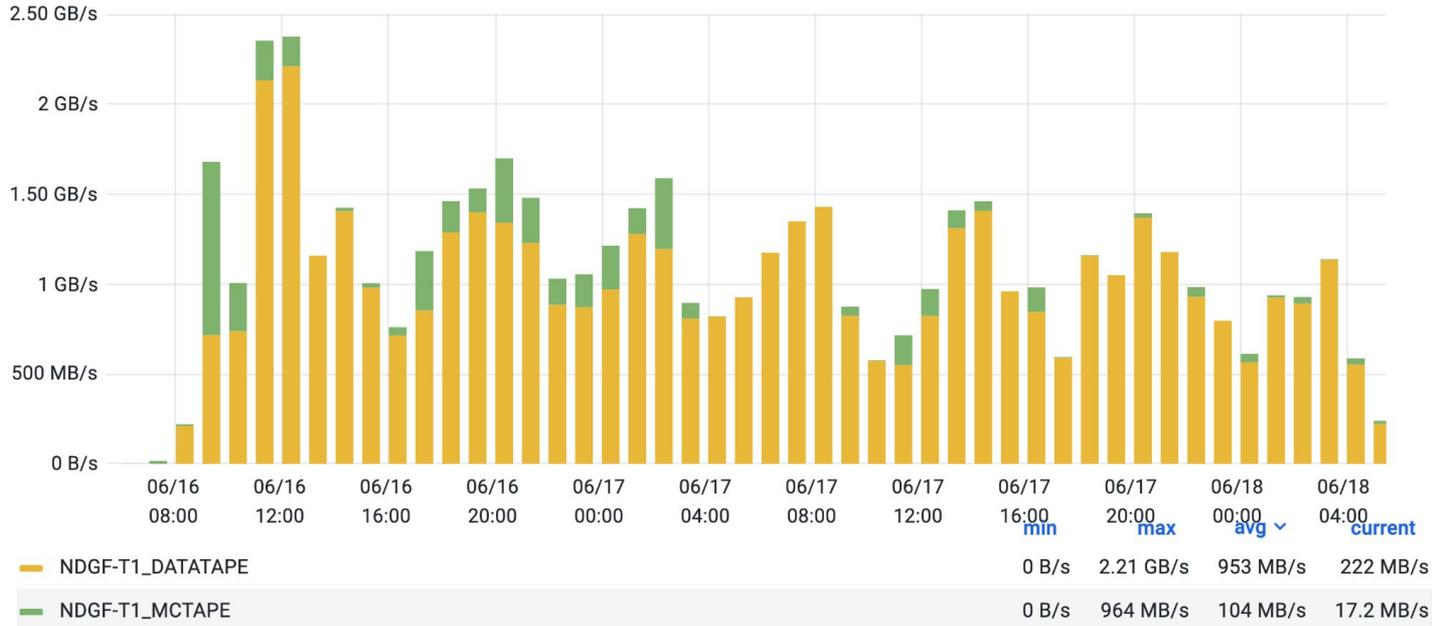
mc16\_13TeV:mc16\_13TeV.364701.Pythia8EvtGen\_A14NNPDF23LO\_jetjet\_JZ1WithSW.recon.ESD.e7142\_e5984\_s3170\_r12670\_tid25731743\_00

+-----+-----+-----+-----+-----+

Total files : 250  
Total size : 2.4 TB

# A-DT test (2) : throughput from DDM dashboard

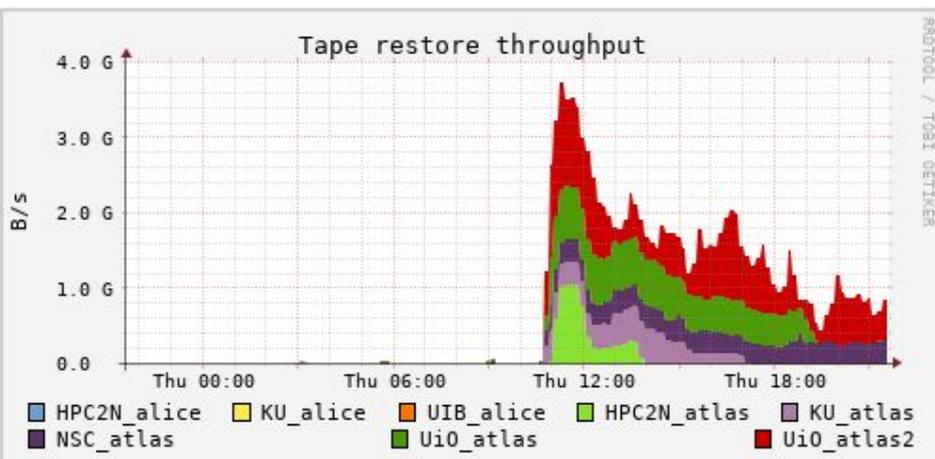
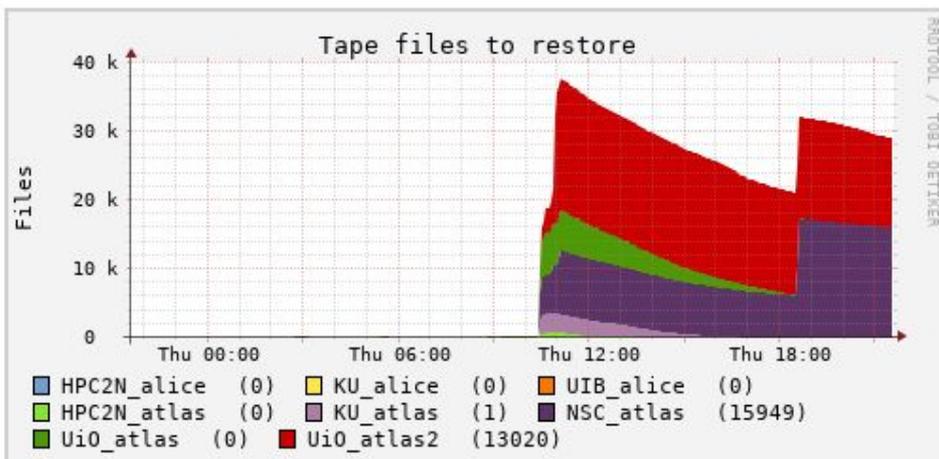
- Target avg staging rate for NDGF (A-DT mode) : 0.5GB/s



# A-DT test, graphs from site monitoring

Big peak early on when there was parallelism to read data from all five tape libraries with ATLAS data

## Tape restore statistics



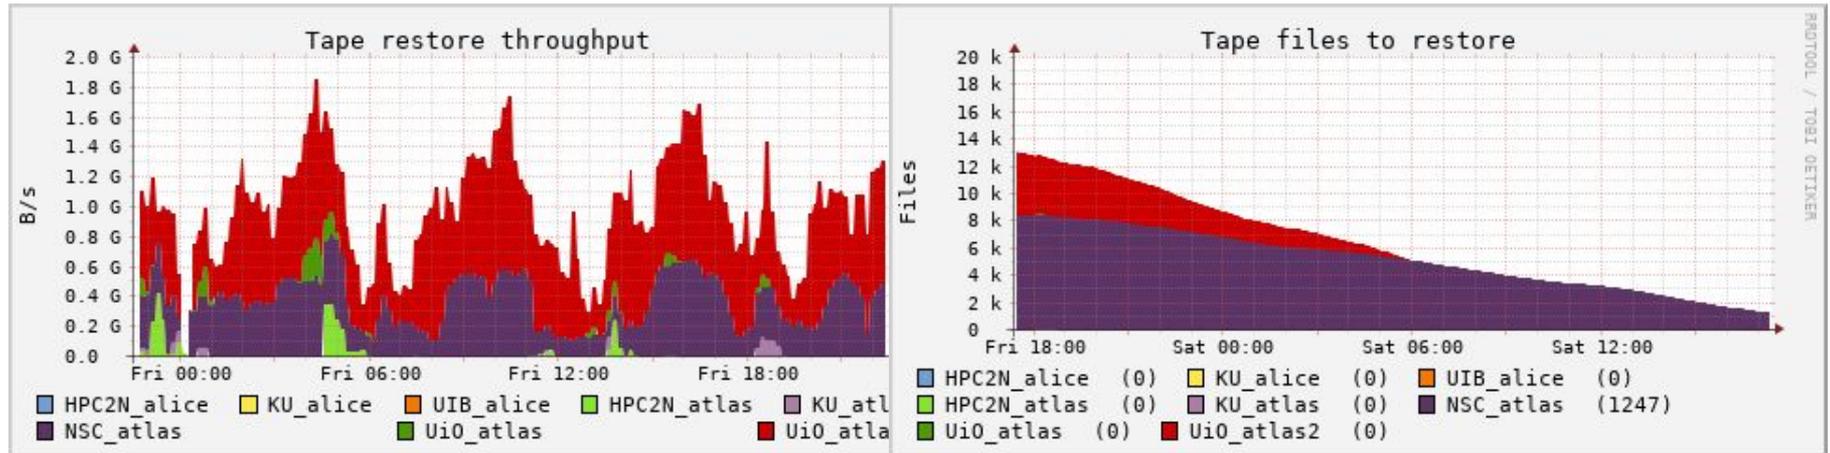
# A-DT test, graphs from site monitoring

Throughput then waned as fewer tape libraries could participate in restores

In one way the test was worse than real traffic: the UiO read pool was almost full with requests for the compressible test data stored on one tape, so for a long while only one tape drive was working

Day two throughput

Day three restore queue



NDGF-internal discussion, Xin: Please don't present this to a wide audience (ATLAS tape experts internally is fine)

NSC tape has this weird fluctuation. During this entire period there were two full tape restores running on tapes likely with compressible test data. Networking issue between TSM server and pool?

