

Summary of Dual Phase activities during DC2 and further work

Denis Pugnere, Elisabetta Pennacchio, IPNL


IT-protodUNE coordination meeting, 17/05/2018

Data Challenge organization


- One goal of DC2 was to test the data transfer from the NP02 online storage servers at EHN1 to the CERN IT at different data bandwidths, going from the one foreseen for beam operation (3 Gbit/s averaged over the SPS beam cycle), up to the maximal one allowed by the link (40 Gbit/s).
- To simulate the raw data flow, a set of binary files corresponding to a total volume of 2 TB has been copied on the storage servers. Each file has a size of 3 GB, and could be copied more than once during tests.
- To copy "raw data" from EHN1 to the IT division a set of scripts based on the ones already used during operations of the 3x1x1 detector has been developed.
- The NP02 online storage facility is composed by 20 (+ 5 spares) DELL R510 servers including 72 TB of disk space per server. The system provides overall 1.44 PB of raw disk space with 20 servers operating in parallel, or 960 TB of total available disk space, by taking also into account RAID redundancy. Each server has 10 Gbit/s connectivity. The storage system operates under a distributed storage system based on a local EOS instance (Citrine version).
- We experienced problems with the third-party-copy protocol from the local EOS instance on the NP02 online storage cluster to the CERN central EOS system. Then, in order to generate the data volume for the Joint Data Challenge transfer test, instead than transferring from the local EOS of NP02 we then performed transfers from the local file systems of 5 machines out the 20 available servers.

Results of the transfer tests

- During the joint data challenge week we successfully tested file transfer from the NP02 online data storage at EHN1 to CERN EOS under various conditions, for a total transferred data volume of 430 TB.
- A data bandwidth of 20 Gbit/s (corresponding to the initial NP02 maximal link speed requirement) could be routinely achieved and sustained for several hours with peak at value 24 Gbit/s. Several tests under different conditions have been performed, with similar results.

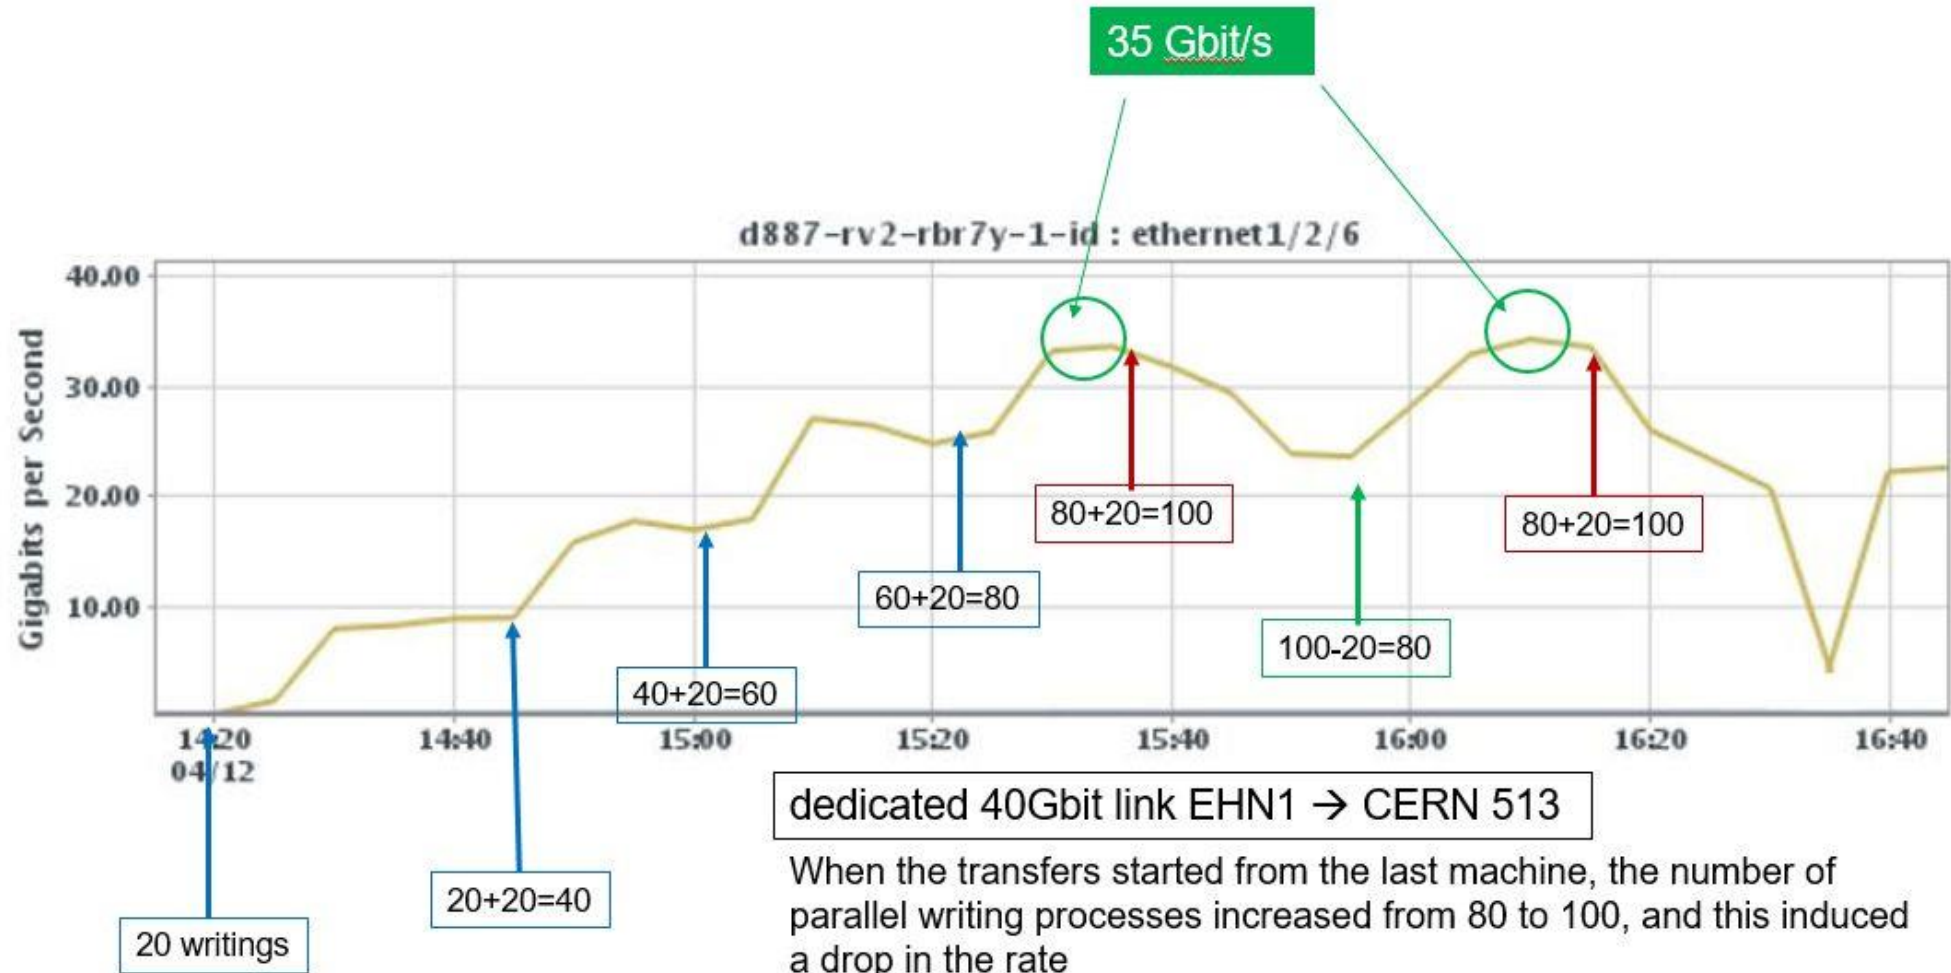


occupancy of the dedicated 40 Gbit/s link from NP02 to IT during 24 hours of the steady transfer test



occupancy of the dedicated 40 Gbit/s link from NP02 to IT during a shorter steady transfer test (~2 hours)

A data bandwidth of 35 Gbit/s, close to the actual NP02-IT link speed saturation, was achieved as well.



Problems fixed with the help of IT people

1. Errors during data transfer

data transfer errors occurred during the different tests represent 1% of the full statistics, and this value seems independent on the data transfer bandwidth.

50% of the errors was “strange”: the checksum values of the source and the target file are identical, both replica are available (checked with « eos fileinfo »). To understand and fix this issue, the asked the help of IT people.

The value of the environment variable XRD STREAMTIMEOUT (xrdcp timeout) is set to 1 minute by default. In some cases this value is too short to end successfully the copy process, and the error condition is raised. The value of the timeout has been increased to 5 minutes some more tests have been performed and the error disappeared.

2. Third-party-copy option

This issue has been solved this week by modifying the configuration of the storage nodes on the eospublic instance (IT side).

What's next

1. Repeat the transfer tests, at different data bandwidth, by copying files between the 2 EOS instance using the TPC option (the scripts already used have been modified accordingly).
2. Next week 9 servers Dell PowerEdge R610 support units will be delivered and installed at EHN1, in the daq room. They are going to be used as service machines.