



ALICE participation in DC24

L. Betev, M. Litmaath

Planned data rates

Centre	Target rate GB/s
CNAF	0.8
IN2P3	0.4
KISTI	0.2
GridKA	0.6
NDGF	0.3
NL-T1	0.1
RAL	0.1
<i>CERN</i>	<i>10</i>

- **ALICE target rates are defined for Run3/4 and remain unchanged**
- All rates tested and achieved during data challenges and subsequently in 2022/2023 (real data transfers)
- Rates coordinated with T0/T1s custodial storage experts
- Data to be transferred during DC24 - 34PB from 2023 Pb-Pb run
- Other WAN load from normal activities (full T0/1/2-T0/1/2 mesh)

Targets: 2.5GB/s (T1s) + 10GB/s (T0)

Rates rationale

- The rates to T0/T1s - proportional to the pledged resources of the corresponding centre
- Data volume share: $\frac{2}{3}$ to T0, $\frac{1}{3}$ to T1s
- Largest contribution are the Compressed Time Frames (CTF) - the ALICE equivalent of RAW data in Run3+
 - All data are located on the O2 disk buffer at CERN prior to transfer
 - File sizes: 10GB/CTF file + tail of calibration files associated to the CTFs
- Most intensive transfer period - after the end of Pb-Pb data taking for the corresponding year
 - Copy the data to tape in 3-4 months after completion of data taking

DC24 transfer methodology

- Transfer of real Pb-Pb data collected in 2023 (copy to custodial storage)
 - Total size to copy 34 PB
- No software changes in the ALICE DM system foreseen
 - Transfer tools - xrd3cp (TPC) + JAliEn transfer scheduler + ALICE tokens
 - Transfer tuning - number of streams per target, proportional to file size
- TPC monitoring (xrootd) is available and good enough
 - Data is exported from MonALISA to common monitoring system


Average achieved rates (in the period of DC24)

Centre	Target rate GB/s	Average achieved GB/s
CNAF	0.8	0.98 (+20%)
IN2P3	0.4	0.6 (+40%)
KISTI	0.2	0.25 (+22%)
GridKA	0.6	1.12 (+90%)
NDGF	0.3	0.35 (+15%)
NL-T1	0.1	0.25 (+150%)
RAL	0.1	0.58 (+500%)
<i>CERN</i>	<i>10</i>	<i>14.2 (+40%)</i>


- **The target rates were increased to most of the T1s on request from T1 experts to check the system capabilities**
- We did not observe any network instabilities or throughput issues
- All of the target rates were set by varying the number of parallel transfers

Averages: 3.75GB/s (T1s, +40%) + 14.2GB/s (T0, +40%)

Time evolution T0



Time evolution T1s

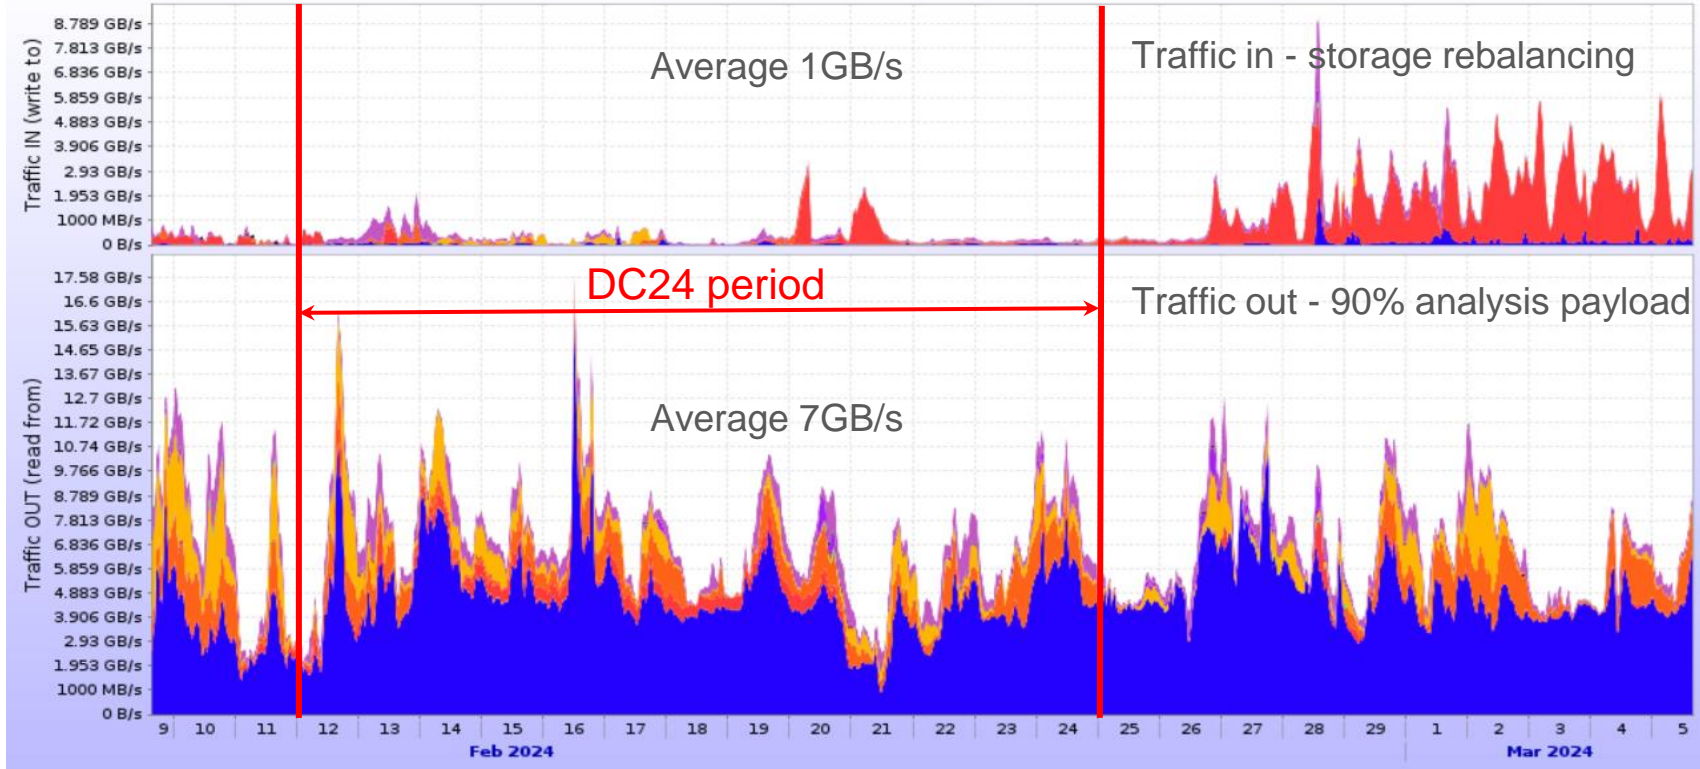


Incidents

- T0 rate - could not achieve 10GB/s in the first 3½ days
 - Reason - a limit of max active transfer threads was forgotten in the ALICE transfer system (self inflicted)
- 3h interruption of transfers to GridKA
 - Dead xrootd service on a disk buffer - repaired by restarting the service
- 24h interruption at CNAF
 - Too high rate observed on disk buffer, pinpointed to reads for md5sum calculation
 - Solved by adding SSD to the buffer
- All of the above are relatively trivial running issues, no structural or software problems identified

WAN traffic from data processing and analysis

WAN server traffic Matrix of 38x38 sites



Conclusions

- Used real data transfer in the exercise, i.e. tested the entire custodial storage chain - source, network, target storage and transfer software
- Target rates surpassed by 35% at T0 and 40% at T1s
- Steady transfers to all target storage elements, minimal intervention required in only two sites, zero network issues
- No interference observed from activities of other VOs, no effect on other activities within ALICE
- No errors during transfer
- **DC24 was successful for ALICE, many thanks to the sites and to the network!**