Tape Challenge

Alessandra Forti (Manchester), Luca Mascetti (CERN), Maria Arsuaga Rios (CERN), Xin Zhao (BNL)

GDB, November 10th, 2021

Many thanks to :

- Experts from the WLCG ops team, FTS team, Rucio team, MONIT team
- Computing experts from all experiments
- Site experts from all T0/T1 sites

Outline

- WLCG Tape Challenge in October, 2021
- Results by experiments
- Site perspective
- What we learned
- Summary

Tape Challenge (1/4) : Goal

- WLCG Tape Challenge was at the second week of October, 2021. The goal is to validate the maximum tape bandwidth needed for reads and writes for Run3
  - This time we focused on T1 sites

- Survey on experiments and sites
  - Experiments expectations on T1 tape throughput for Run3 (table)
    - The values refer to delivered throughputs, not nominal tape bandwidth on the floor
  - All T0/T1 sites responded to WLCG survey, with information of their tape resources and configuration.

Overall T1s objectives for RUN3:

Indicate in this table, the bandwidths required for all T1s for reads and writes during data taking (DT) and right after data taking(A-DT).

| VO	Reads (DT) GB/s	Writes (DT) GB/s	Reads (A-DT) GB/s	Writes (A-DT) GB/s
ALICE	0	2.8	1.1	2.8
ATLAS	2.5	9.6	8.4	5.1
CMS	0.8	7.6	12.3	1.1
LHCb		11	3.38	
Total	2.5	24.78	25.18	8.3

Tape Challenge (2/4) : driven by experiments

- Each experiment ran its own tape tests during the week, driven by existing workflow and dataflow management systems
- "Shared Clock" among experiments (roughly)
  - Data-Taking(DT) mode in the first 2 days
  - After-Data-Taking (A-DT) mode in the last 3 days
- ATLAS and CMS ran tape write and tape read tests, in both DT and A-DT modes
- ALICE and LHCb ran only T0 export to T1s test, i.e. only tape write test
- All tape traffic during the week count toward the final throughput, ie. "activities=ALL" on the dashboard

Tape Challenge (3/4) : monitoring

- Central monitoring dashboard enhancement
  - Several knobs and plots were added to the Data Challenge dashboard, to distinguish tape traffic from disk traffic.
  - Covers experiments that manage data transfers via FTS
- ALICE has its own MonALISA monitoring
- Some sites provide monitoring links for site level tape activities
  - KIT, TRIUMF, NL-T1, CNAF

Tape Challenge (4/4) : tape throughputs over the week
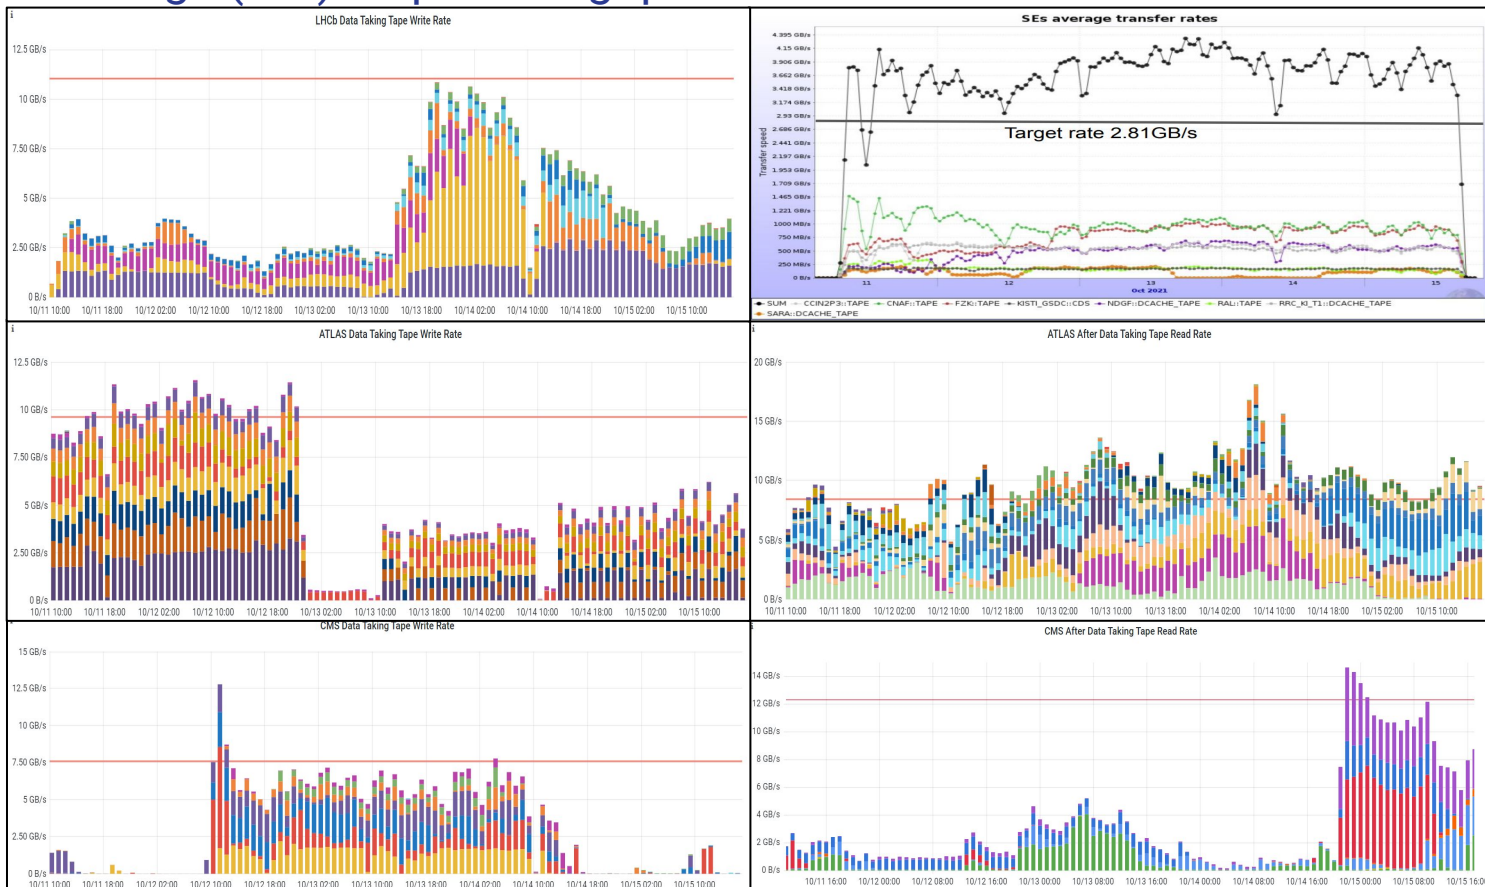
ALICE results

- Tested T0 export to T1 tape (DT and A-DT mode write test)
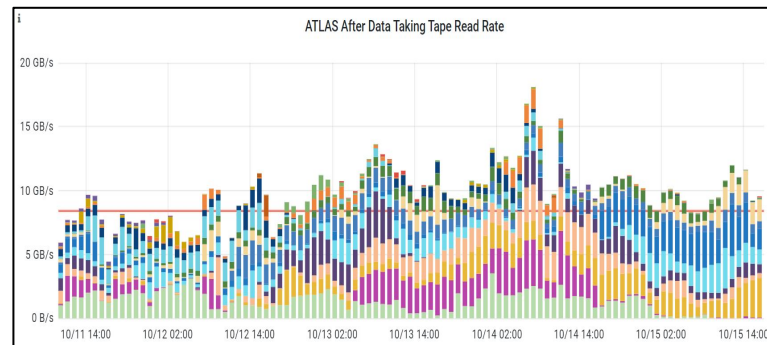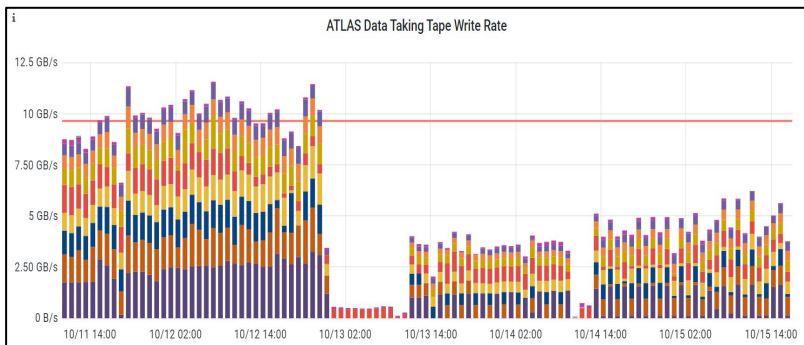- Smooth run, target rates achieved (exceeded), no particular issues found

SEs average transfer rates
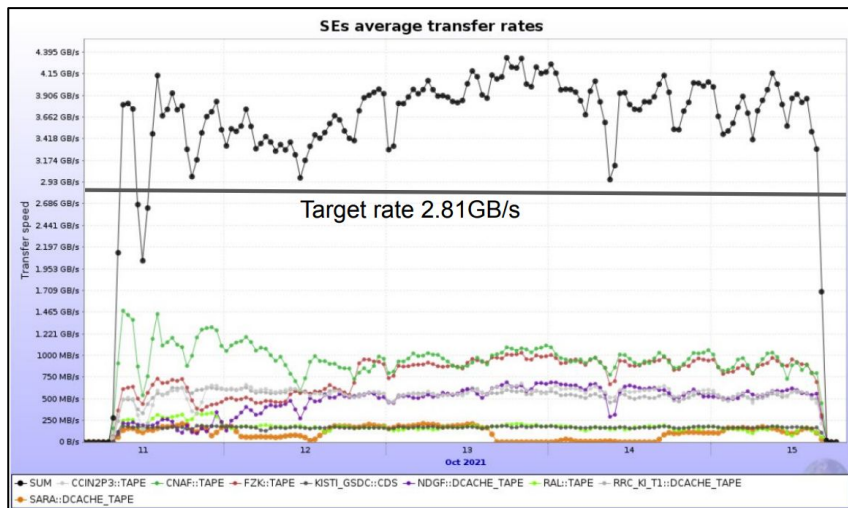
Target rate 2.81GB/s

T1 Centre	Target rate GB/s	Achieved rate GB/s
CNAF	0.8	0.94 (116%)
IN2P3	0.4	0.54 (130%)
KISTI	0.15	0.16 (106%)
GridKA	0.6	0.76 (123%)
NDGF	0.3	0.47 (144%)
NL-T1	0.08	0.1 (122%)
RRC-KI	0.4	0.53 (128%)
RAL	0.08	0.17 (172%)

Sum 2.81GB/s

ATLAS results

- Tests went well, following the plan
    - Staging test piggybacked on the Data Carousel traffic (real production requests from the ongoing reprocessing and derivation campaigns)
- Reached the overall target rates in both DT and A-DT modes
    - Overall staging rate exceeded target. Pressure on particular sites varied depending on the real production needs at the time

CMS results

- **Write test:**
  - Target write rate for prompt data+MC production (**7.6 GB/s**) not achieved during the week:
    - achieve target rate for data exports from T0 **(4.21 GB/s)**, despite FNAL having a large queue from the previous week
  - Write test T1_Disk → T1_Tape followed after.
    - Useful to measure max. tape-write capacity
    - Different sites tested at different time: time zone difference vs approval model
- **Read test:**
  - Staging test with containers of 300TBs
    - No specific target (wanted to stress the system)
    - Better performance for KIT, PIC, IN2P3 compared with Spring Tape tests

7.6 GB/s target is sum of estimated Prompt Reco (4.2 GB/s) and from MC production (3.4 GB/s) … but there was basically no production traffic
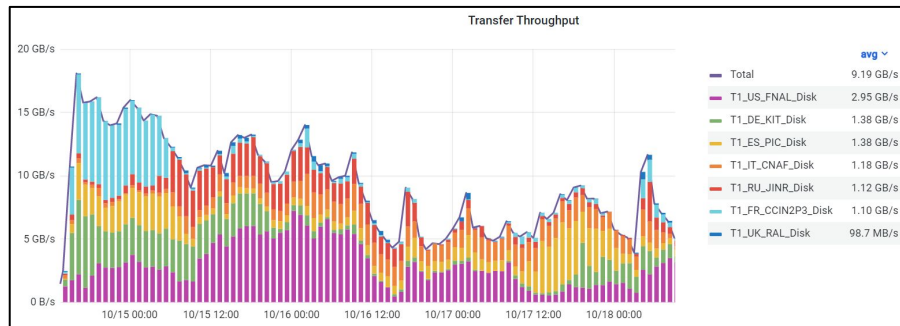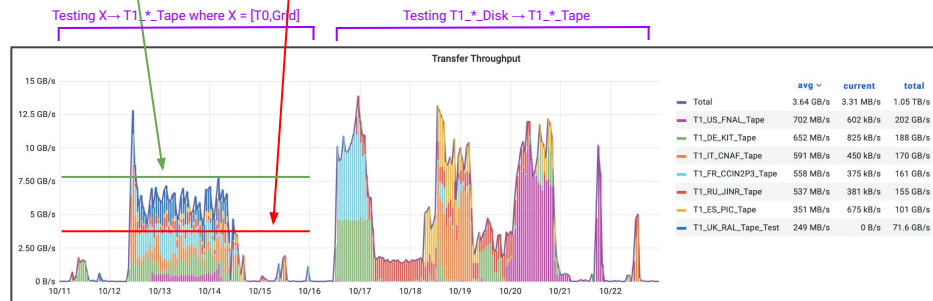
4.2 GB/s target for Prompt Reco (i.e., T0→ T1_Tapes)

Testing X→ T1_*_Tape where X = [T0,Grid]

Testing T1_*_Disk → T1_*_Tape

Transfer Throughput

	avg	current	total
Total	3.64 GB/s	3.31 MB/s	1.05 TB/s
T1_US_FNAL_Tape	702 MB/s	602 kB/s	202 GB/s
T1_DE_KIT_Tape	652 MB/s	825 kB/s	188 GB/s
T1_IT_CNAF_Tape	591 MB/s	450 kB/s	170 GB/s
T1_FR_CCIN2P3_Tape	558 MB/s	375 kB/s	161 GB/s
T1_RU_JINR_Tape	537 MB/s	381 kB/s	155 GB/s
T1_ES_PIC_Tape	351 MB/s	675 kB/s	101 GB/s
T1_UK_RAL_Tape_Test	249 MB/s	0 B/s	71.6 GB/s

Transfer Throughput

	avg
Total	9.19 GB/s
T1_US_FNAL_Disk	2.95 GB/s
T1_DE_KIT_Disk	1.38 GB/s
T1_ES_PIC_Disk	1.38 GB/s
T1_IT_CNAF_Disk	1.18 GB/s
T1_RU_JINR_Disk	1.12 GB/s
T1_FR_CCIN2P3_Disk	1.10 GB/s
T1_UK_RAL_Disk	98.7 MB/s

LHCb results

- Tested T0 export to T1 tape (DT mode write test)
- Slow start
  - FTS knobs not tuned to match the test scale
  - Not enough EOS Gridftp GWs
  - Misunderstanding of expected tape write rate by sites, leading to not enough tape drives allocated at one site
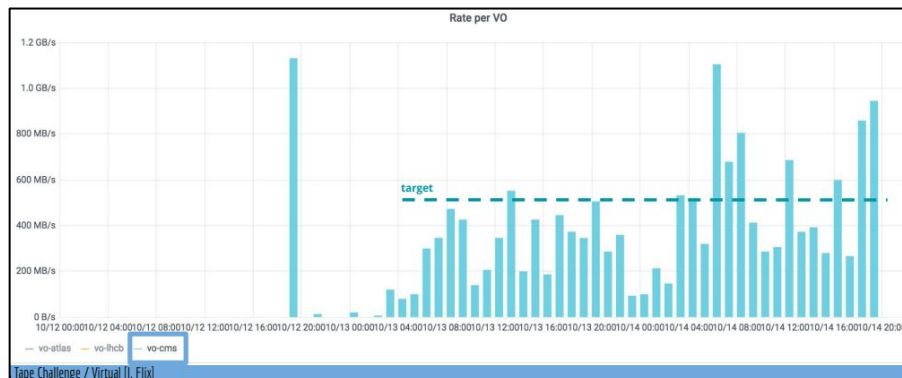- After fixing the above issues, tape write rate was much better, close to the target
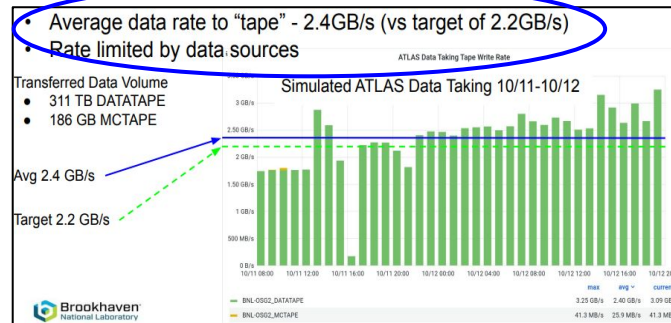

LHCb Data Taking Tape Write Rate
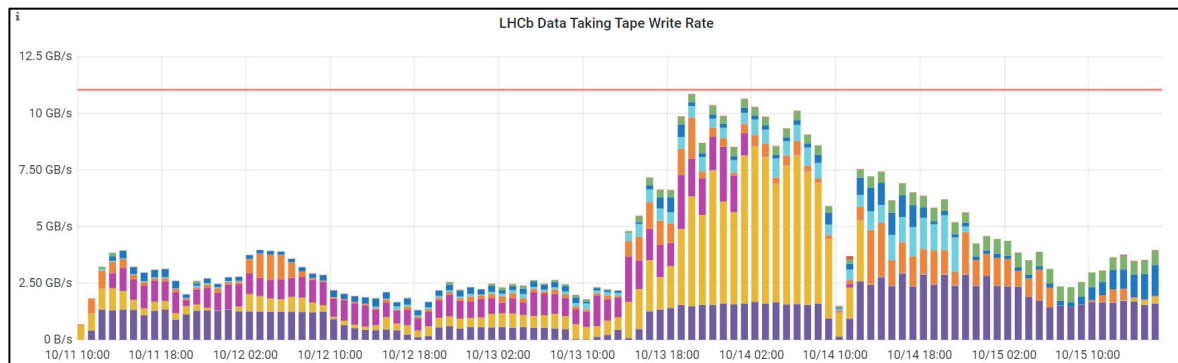
Site perspective (1/3)

- [Site post-mortem](#) still ongoing (next meeting tomorrow)
  - How the tape storage themselves performed in the challenge ?
- Sites are proactive in preparing for the tape challenge
  - Ongoing data migration to new tape library paused
  - Tape resources (drives/media) carefully allocated to meet the expected target rates from all experiments
  - Took this opportunity to test newly deployed tape system (e.g. RAL CTA instance)

Site perspective (2/3)

- Sites tape resources were not really stressed. Sites report meeting the target rates, if given enough external pressure (from experiments)

Site perspective (3/3)

- Some expected observations
  - Tape system has its own migration cycle
  - Shared resources among experiments
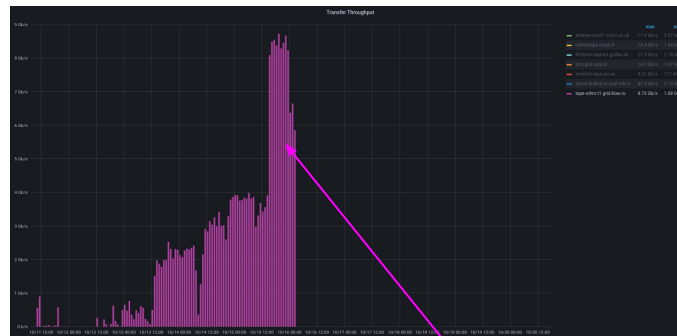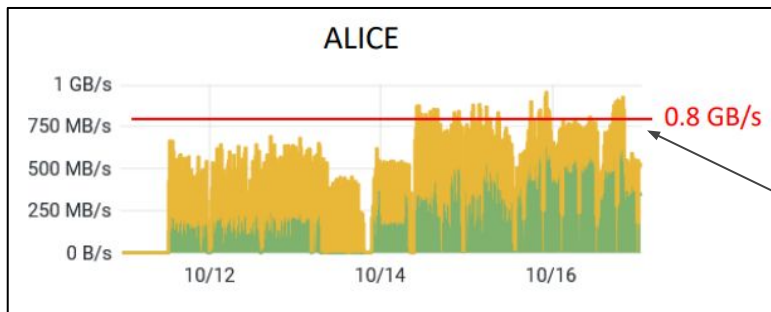
The incoming rate to disk buffer was higher (0.94GB/s), but tape system has its own migration cycle and rate

After ALICE was done writing, LHCb write rate doubled at a site

- More observations later …

What we learned (1/6)

- Orchestration of the tape challenge among all experiments
  - Each experiment planned and organized its own tests in the tape challenge, under a "shared clock".
  - The "shared clock" helped only partially to synchronize tests among experiments. Test didn't occur when all experiments peaked, either in tape write or in tape read, at the same time.
  - In future challenges, a proactive sharing of technical details during the preparation phase may help in orchestrating all experiments to move in full-sync.

What we learned (2/6)

- Is one week for a tape challenge too short ?
  - Not enough time for experiments and sites to tune their configuration, in order to reach the ideal test condition.
  - Leave room for retrying some parts of the test if needed
  - ALICE did a mock run before the official tape challenge, that appeared to have helped a lot
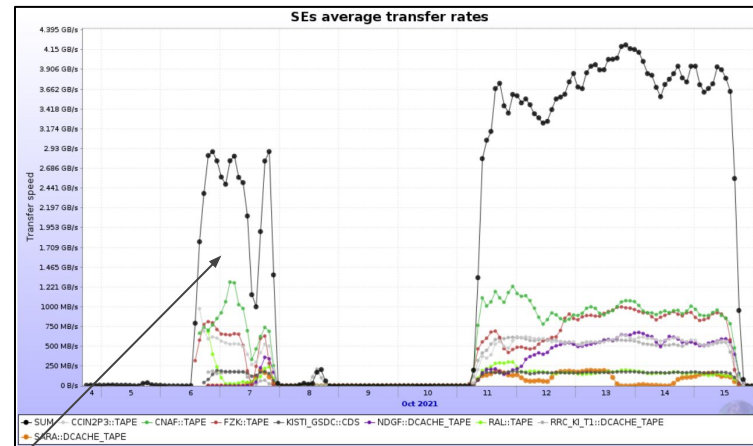
ALICE
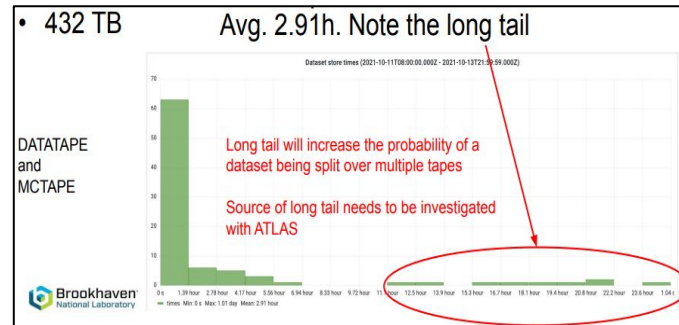mock run

What we learned (3/6)

- Site feedback on running a tape challenge
  - If fake data is to be used in tape write test, experiments should coordinate with sites, to keep the fake data separate from real data, e.g. in its own file families, to make the reclamation of tape space easier later
  - Experiments had better use different data samples for write and read tests, because cleaning data from disk buffer takes time, not very convenient for sites to do on the fly.

What we learned (4/6)

- Monitoring
    - Central monitoring
        - Very useful in showing the overall picture of how the available tape resources are being used by all experiments, and by what activities (read vs write) etc. These information are not available on individual experiment's dashboard.
        - needs to have a consistent representation of the experiments data (RSE names, activity names etc), also cover all traffic, including xrootd traffic.
    - Site monitoring
        - Very limited in this area, only few sites provided link.
        - Sites are not unwilling to provide local information. But sites are concerned about providing misleading information.
            - For example, tape resources are shared among LHC experiments, also with non-LHC experiments. Even for the same experiment, there are local traffic vs experiment-initiated traffic. While site monitoring not yet distinguish between them.
        - Some sites lack fancy visualization for sharing with the public.
        - Future improvements :
            - Sites should continue to try separate the different experiments and activity traffics on local monitoring (a good example at KIT)
            - WLCG to work on integrating site tape activities information into the central MONIT. e.g. revamp the WLCG tape metrics infrastructure ?

What we learned (5/6)

- Site feedback on optimal use of tape resources
  - Write data to tape the way it will be read back; read data from tape the way it is written
    - Group files on tape during migration
    - Reduce time spent in mounting/dismounting/positioning during recall
  - Reduce the dataset write windows and the dataset bringonline windows, coming to sites.
  - Bigger files leads to higher tape bandwidth utilization in both migration and staging
  - File aggregation helps with migration but hurts staging performance (unless the whole aggregate is recalled)
  - Collaboration between sites and experiments is needed for more optimal tape resources utilization.

- 432 TB   Avg. 2.91h. Note the long tail

DATATAPE and MCTAPE

Long tail will increase the probability of a dataset being split over multiple tapes

Source of long tail needs to be investigated with ATLAS

- CMS staging performance during A-DT better than ATLAS b/c
  - Bigger files (good for both migration and staging)
  - Less scattered across tapes
  - Less competing activities
  - Better data distribution across drive sets (more staging drives for CMS and underperforming drives for ATLAS)

What we learned (6/6)

- Several other recommendations
  - The current srm+gridftp protocol is used in this challenge. Gridftp will be deprecated in Run3, future challenges are expected to run with srm+https protocol.
  - Site staging profile, used by ATLAS Data Carousel, is found to be useful by sites. Other experiments may consider it as well :
    - It's easy to configure (already in CRIC)
    - If an experiment decides to adopt it, the workflow or dataflow management system needs to respect the profile limits and throttle staging requests to sites accordingly.

Summary

- A very fruitful exercise, first of its kind. Many people (central operations team, experiments, sites etc.) spent a lot of time and effort, obtained rates satisfactory (sometimes expressed with an embedded contingency).
- Both experiments and sites get a better idea what to expect in Run3 w.r.t. tape usage.
- What we learned from this tape challenge will benefit us all in future challenges, and Run3 and beyond.
- Target for a longer tape challenge (2 weeks) in the first quarter of next year, in combination with Tier0 commissioning tests.