

Tape Challenge '22

Xin Zhao on behalf of the Tape Challenge team

GDB, May 11th, 2022

Team effort –

- LHC experiments (including experiments Points experts)
- Central services, Tier0 and Tier1 experts
- WLCG operations
- Tape challenge coordination (Alessandra Forti/U. of Manchester, Luca Mascetti/CERN, Maria Arsuaga Rios/CERN, Xin Zhao/BNL)

This report is based on –

Experiment reports

- [ATLAS](#) (David South, David Cameron, Mario Lassnig, Xin Zhao)
- [CMS](#) (Garyfallia Paspalaki, Benedikt Maier, Fernando Garzon, Felipe Gomez, Conner Moore)
- [LHCb](#) (Christophe Haen)

Site reports

- [CERN](#) (Julien Leduc, Maria Arsuaga Rios, Steven Murray/CERN FTS team)
- [BNL](#) (Shigeki Misawa)
- [CC-IN2P3](#) (Aresh Vedaee)
- [CNAF](#) (Enrico Fattibene)
- [FNAL](#) (D. Mason for USCMS S&C)
- [KIT](#) (H. Musheghyan, A. Gottmann, P. Konstantinov, D. Lobontu, X. Mol, S. Pérez, A. Petzold, D. Ressmann)
- [PIC](#) (J. Flix, E. Acción, A. Pacheco, A. Pérez-Calero, E. Planas)
- [RAL](#) (Alastair Dewhurst for RAL Tier-1)

Outline

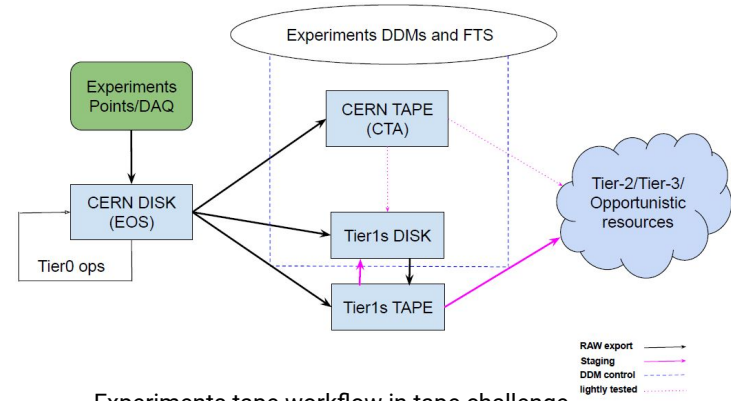
- Introduction
- Experiments perspective
- Central services
- Sites perspective
- Tape test operational experience and lessons learned
- Next steps and future improvements
- Summary

Introduction

- **Goal**
 - To validate the readiness of tape workflows and tape resources of LHC experiments for Run3.
- **Major differences from the 2021 tape challenge : get real !**
 - Full chain : pointX → CERN T0 → Tier1s
 - More complete and realistic tape scenarios covered
 - e.g. LHCb added tape read test
 - More Run3 hardware and software upgrades in place
 - e.g. WLCG finished migration to srm+https protocol for tape access
- **Time**
 - Three weeks in March (2022-03-07 ~ 2022-03-25)
 - For completeness, results below also cover some other tape tests that were carried out recently but outside of the above time window.
- **Participants**
 - Three LHC experiments : ATLAS, CMS and LHCb
 - ALICE skipped due to conflict with other commissioning activities. ALICE achieved Run3 goal in the last (2021) tape challenge.
 - Tier0 and majority of Tier-1 facilities
- **Monitoring**
 - [Central Data Challenge dashboard](#), [FTS dashboard](#), experiments DDM dashboards, site level monitoring
- **Communication channels**
 - Weekly tape challenge [meetings](#), [egroup](#), [slack channel](#), shared [google doc](#), GGUS tickets with all participant sites

Experiments perspective

- Experiments tested various major tape scenarios for Run3, in both DT (data-taking) and A-DT (after-data-taking) modes, on Tier0(include DAQ) and Tier1 sites.
 - Diagram to the right shows the major tape scenarios tested in tape challenge '22
- “Shared clock”
 - To synchronize the various tape tests among all experiments and sites, in the middle of many other ongoing commissioning activities, we reached a timeline with experiments (table to the right), at least two experiments running either DT or A-DT tests at the same time (table to the right).
- Below we will summarize the results for each experiments.



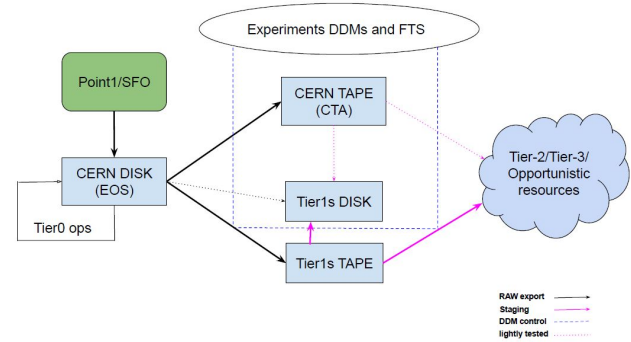
Experiments tape workflow in tape challenge

Experiments	Week 2 of March (03-07 ~ 03-11)	Week 3 of March (03-14 ~ 03-18)	Week 4 of March (03-21 ~ 03-25)
ATLAS		DT	A-DT
CMS	A-DT	DT	
LHCb		DT	A-DT

Timeline of experiments tests


Experiments perspective : ATLAS

- ATLAS ran both DT and A-DT mode tests
 - 2021 tape challenge demonstrated Tier1 tape resources meet the Run3 expectations in both DT and A-DT modes
 - This time the focus are 1) Tier0 scenarios and 2) peak export rate (3.5GB/s) with each Tier1
- Reached targets (again) on all fronts
 - Overshoot of tape write rate to Tier1s, follow up ([FTS-1782](#)) with FTS on throughput limit control




ATLAS tape workflow in tape challenge

Tier-1 sites	DT test (tape write)		A-DT test (tape read)		
	*Target (GB/s)	Real rate (GB/s)	Target rate (GB/s)	Max rate (GB/s)	*Stable rate (GB/s)
BNL	3.5	>8	1.9	6.8	5
CC-IN2P3	3.5	Up to 4	1.2	6.9	2.6
CNAF	3.5	>4	0.8	3.0	1.9
KIT	3.5	Up to 3.5	1.0	3.0	1.6
NL-T1	3.5	Up to 7	0.6	1.5	1.2
PIC	3.5	>6	0.3	2.1	1.2
RAL	3.5	Up to 3.5	1.2	2.1	1.7
*TRIUMF	3.5	Up to 4	0.8	2.8	1.6



P1/SFO to CERN EOS export rate



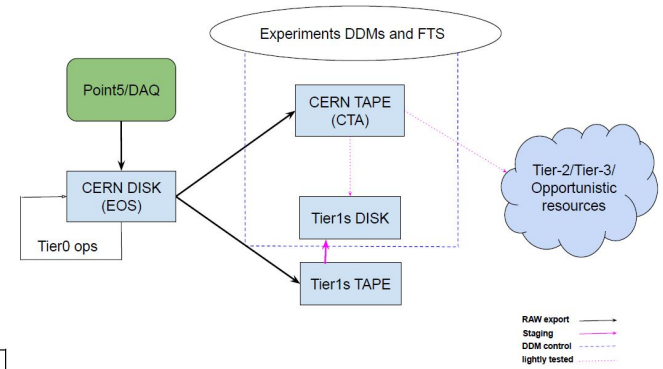
CERN EOS to CERN CTA export rate

* The DT target rate refers to the peak rate during a run (main stream RAW data)
 * stable rate refers to a rate sustained for >5 hours.
 * TRIUMF DT result was from a re-test after the tape challenge

Tier-1 DT and A-DT test results

Experiment perspective : CMS

- CMS ran both DT and A-DT tests
 - DT test used mock data produced on the fly from Point5 and transferred to EOS, more realistic Run3 condition. Replay mechanism used to transfer mock data from T0_Disk to T0_Tape/T1_Tape. Test designed to reach target rates at a given pressure (i.e. 5 days)
 - No specific target rates for the A-DT test, tried to stress the system. Reached smaller staging rates than tape challenge 2021, due to difference in test samples used.
- Overall a successful test, reached the Run3 (DT) target rates, which is calculated with large safety margin (50%)



CMS tape workflow in tape challenge

CMS site	March 2022 rates (GB/s)	Target Rates (GB/s)
CTA	5.72	3.2
KIT	0.59	0.29
PIC	0.58	0.15
Fermilab	1.97	0.73
RAL	0.8	0.29
IN2P3	0.59	0.29
JINR	0.50	0.55
CNAF	0.32	0.37

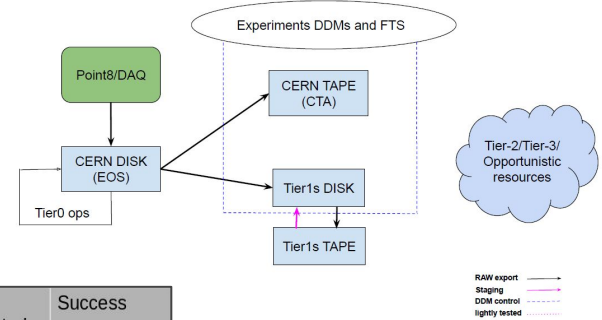
DT test results

CMS site	March 2022 average rates (GB/s)	March 2022 Highest rates (GB/s)	October 2021 average rates (GB/s)	Ratio 2021/2022 (average rates)
KIT	1.81	5.26	2.34	1.29
PIC	1.20	1.68	1.84	1.53
Fermilab	1.51 (1.57)	2.26	2.46	1.62
RAL	0.18 (0.21)	1.05	--	
JINR	0.705	3.31	1.56	2.2
CNAF	0.973	2.03	1.07	1.10
IN2P3	3.42	7.75	4.74	1.38

A-DT test results

Experiment perspective : LHCb

- LHCb ran both DT and A-DT tests
 - DT test mimic the Run3 RAW data flow (slightly different from ATLAS/CMS).
 - A-DT read back the data written during DT test
- Mixed results, more work needed with three sites (CNAF/NCBJ/RAL).



LHCb tape workflow in Tape challenge

Site	expected Speed (GB/s)	Max Speed (GB/s) disk/tape	Avg Speed (GB/s) disk/tape	Duration (hours) disk/tape	Success
CERN	11.00	NaN/13.13	NaN/11.31	NaN/38	Yes
CNAF	1.72	2.24/3.43	1.5/1.21	45/56	~Yes
GRIDKA	2.23	4.46/3.46	2.24/2.03	39/43	Yes
IN2P3	1.25	2.99/3.03	1.31/1.13	37/43	Yes
NCBJ	1.32	0.84/NaN	0.64/NaN	82/NaN	No
PIC	0.20	0.77/0.98	0.21/0.22	37/36	Yes
RAL	2.96	2.65/4.11	1.41/1.28	86/85	~No
RRCKI	0.25	0.57/0.30	0.29/0.13	22/49	Yes
SARA	1.07	3.16/2.23	1.12/1.04	37/40	Yes
Total Tier1s	11.00	11.00	7.65		

DT test results


Site	Expected Speed (GB/s)	Max Speed (GB/s)	Avg Speed (GB/s)	Duration real/expected	Success
CNAF	1.35	1.08	0.72	94/68	No
GRIDKA	1.36	1.99	0.66	133/87	Yes/No?
IN2P3	0.98	3.39	1.55	31/68	Yes
NCBJ	0.91				No
PIC	0.17	1.06	0.54	14/61	Yes
RAL	1.93	2.26	0.70	> 145 (cont)/82	No
RRCKI	0.21			>145 (cont)/61	No
SARA	0.74	1.4	0.5	84/76	Yes/No?

A-DT test results

*CNAF and GRIDKA passed in a re-test afterwards
 *RRCKI had pool manager issue, once fixed, test passed

Central Services :

- Rucio
 - Smooth run in general. Volume of the challenge was negligible to overall Rucio load (in terms of #files)
 - Running almost exclusively on Kubernetes
 - Earlier in the challenge (during CMS A-DT test), 3 out of 4 Rucio clusters got accidentally deleted by an (external) rogue maintenance script. Repaired by the rucio automated cluster recovery mechanism. ~1h downtime with limited impact (FTS queue already filled)
- FTS
 - FTS service was reliable during the tape challenge – Success!
 - 48M transfers, totally 63PB data volume during the tape challenge
 - Operational issues :
 - a configuration regression in the FTS LHCb instance (fixed)
 - Discrepancy in transfer throughput between FTS and DDM dashboards (fixed)
 - Some FTS web pages were confusing, JIRA tickets created to follow up.






Sites perspective : CERN Tier-0

- EOS (CERN disk service)
 - Smooth data collection in general, met the targets (table below)
 - Internal scheduling in EOS was modified and shown to work properly for all experiments
 - Hit by unexpected CMS T0 replay high load activity during the DT test ⇒ improvement on EOS I/O priority to ensure Point5 priority over the other CMS activities

Experiment	Raw Data to EOS 1.5 days	Experiment Tier0 Activity	EOS to CTA	Tier1s Export
ATLAS	8 GB/s	Read -> 2GB/s Write -> 2GB/s	10 GB/s	10 GB/s
CMS	10 GB/s	Parallel activity <ul style="list-style-type: none"> • P5 processing • Mock data processing 	10 GB/s from CMS Tier0 processing on real-time P5 data flow	From CMS Tier0 mock data processing
LHCb	None* * Tested separately reaching 20GB/s the 24 th Feb 2022	None	11GB/s	11 GB/s


Main objectives for the DT mode test



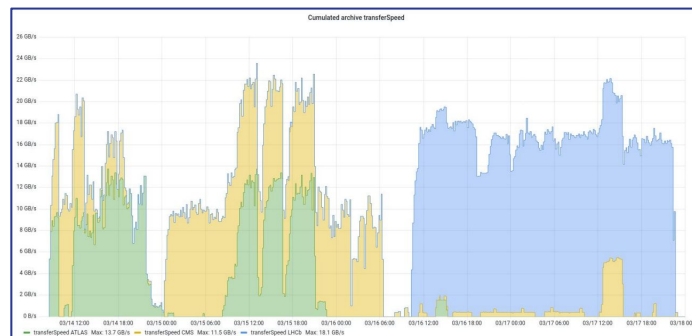
RAW data to EOS for ATLAS/CMS/LHCb

Sites perspective : CERN Tier-0

- CTA (CERN tape service)
 - Demonstrated predictable performance and very good tape write efficiency.
- CheckOnTape feature to make sure data is actually written on T0 tapes
 - Actual time spent on tape buffer depends on # of requests and upstream services. During the DT test, max time to tape for ATLAS files was 30m.



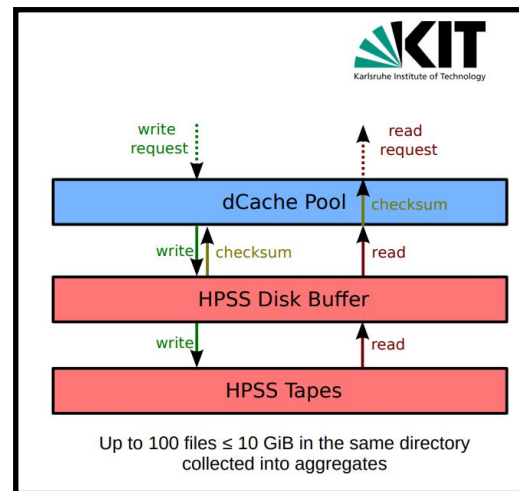
Incoming rate to CTA disk buffer (DT test)



Migration rate to CTA tape (DT test)

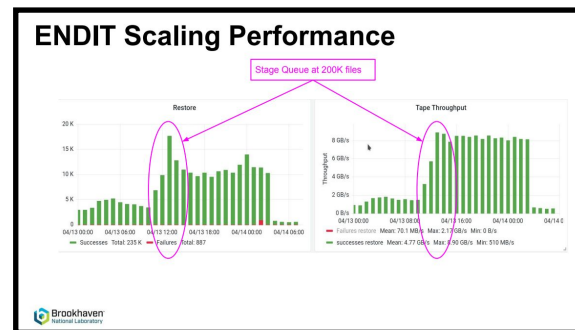

Sites perspective : Tier1 experiences & highlights (1/4)

- KIT – “smart writing” by “aggregates”
 - Two tape libraries in production now
 - Old tape library/TSM (ATLAS)
 - New tape library/HPSS (CMS & LHCb)
 - New setup in the new tape system tries to improve tape I/O utilization efficiency by grouping files into aggregates before writing to tape, and recalling files by aggregates
 - Results
 - The new setup improved the overall tape rate by more than factor of 2 per tape drive.
 - For tape reads, the new setup achieved max rate of ~400MB/s/drive (TS1160 drives)



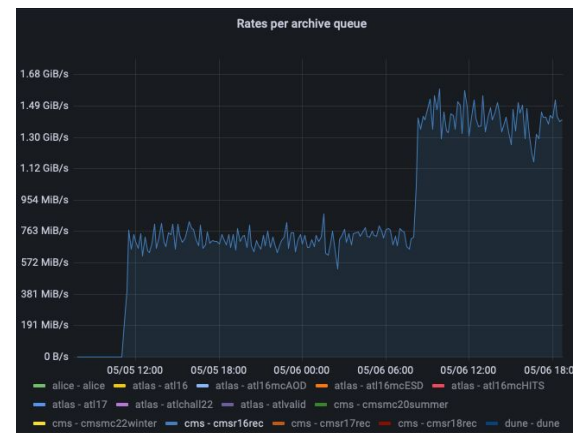
Sites perspective : Tier1 experiences & highlights (2/4)

- CNAF – bigger bulk staging requests
 - LHCb staging test didn't reach the target (1.35GB/s), due to frequent tape remounts. After increasing the limit of concurrent staging requests from 3.5k to 100k, new staging rate increased by a factor of 2, to 1.8GB/s (360MB/s/drive).
- BNL – ENDIT dCache HSM plugin
 - The ENDIT plugin was recently deployed at BNL (modified to work with BNL ERADAT HPSS batch system).
 - Dramatically improved scalability of HSM interface. A dedicated test shows the system works well with 200k concurrent staging requests now (previous limit is 80k).
 - Plan to extend ENDIT to tape writing as well for Run3.




Sites perspective : Tier1 experiences & highlights (3/4)

- RAL – improving tape write performance
 - Right before the tape challenge, RAL migrated to new Antares tape system (EOS+CTA), along with commissioning LTO-9 drives and media.
 - Initial test showed the write performance of LTO-9 drives only half of the expected rate.
 - Tuning the network ring buffer size vastly improved throughput
 - In a recent CMS production tape write activity, reached 375MB/s/drive rate with LTO-9 (400MB/s nominal rate)



Sites perspective : Tier1 experiences & highlights (4/4)

- TRIUMF – improving RAW export transfer rate
 - ATLAS DT test reached <1GB/s rate to TRIUMF tape buffer, far below the bandwidth limit. Investigation with site storage and network experts led to a configuration change in dCache to prioritize IPv6 over IPv4 in data transfer, which dramatically improved the throughput, reaching the target RAW export peak rate of 3.5GB/s.



- IN2P3 – system tunings for better performance
 - A single RFI0 server used for migration and staging between tape and tape buffer. This caused spike of failures in the ATLAS DT test, because the server was overloaded by concurrent CMS staging activities → now split into two RFI0 servers, one for reading with 1024 simultaneous connections and another for writing with 400 simultaneous connections.
 - LHCb A-DT test ran into a lot of timeout errors for bringonline requests, to fix this, the timeout limit in dCache is extended from 4h to 48h

Tape test operational experience and lessons learned

- About test sample
 - Fake data using in ATLAS and CMS DT tests are highly compressible, which is not the case in real production. This could artificially inflate tape migration performance as tape systems often compress data in archiving.
 - Some test data containers used in A-DT staging test were composed of files from different datasets, not a whole dataset staging. This could deflate the tape staging performance in the test, because in real production experiments usually trigger staging requests in units of dataset.
- Tape buffer was not completely cleaned up in some cases, before the A-DT staging tests. This affected measurement of the real tape performance.
- FTS limits not properly configured at the beginning of some tests, leading to bombarding sites with too many requests, causing timeout failures. The staging test had to be re-done later.

Next steps and future improvements (1/2)

- Experiments tape scenarios
 - ATLAS to control tape write rate to Tier-1s :
 - For RAW data export, the main stream comes out of Point1 at the peak rate of ~3.5GB/s, our goal is to avoid hammering any particular Tier1 with sustained peak rate streams. Development ongoing in Rucio to automate this process.
 - For non-time critical tape writes, try to keep the inbound data rate to Tier1 tape buffer under control. In the tape challenge, the FTS throughput limit settings were attempted but failed. Follow up with FTS team to improve this feature.
 - Another proposal is to export the small delayed streams of RAW data only at the end of a run, so that these RAW data arrive at Tier1 tape buffer within a much shorter time window, to help sites co-locating these datasets on tape.
 - LHCb DT mode is T0 DISK → T1 DISK → T1 tape. CNAF, unlike other Tier1s, has disk and tape buffer on the same hardware and FS. Discussion ongoing between LHCb and CNAF on how to simplify the current workflow.

Next steps and future improvements (2/2)

- Site infrastructure
 - RAL Antares – performance met expectations in general, tape writes is excellent.
 - More tuning and improvements are planned for tape recalls.
 - Webdav interface was brand new, significant improvements have been made since it entered production
 - Other near term improvements include upgrading to the latest EOS version and assigning IPv6 addresses to TPC proxy nodes.
 - BNL – continuously improve tape “smart writing”
 - Migration-by-directory HPSS feature in place. Still significant interleaving of dataset blocks observed on tape. Will look into adjusting the tape migration policy, to stop migration if there is insufficient volume of data to write.
 - Many sites run multiple tape libraries in parallel. One operational topic during Run3 will be how to balance usage between them, e.g. chronologically or in some other ways.
- As new tape libraries continue to arrive at several sites, e.g. IN2P3 and NDGF, commissioning tests are expected to happen with the relevant experiments. Dedicated tests can also be done after major improvements on existing tape systems.

Summary

- Tape challenge '22 is a successful test, a good opportunity for experiments and sites to validate Run3 tape workflows, end-to-end, in a more realistic condition.
- Results are satisfactory, which demonstrated (once again) that the current tape capacities meet the Run3 expectations.
- We learned a lot from this exercise. Many of the operational and service performance issues have been addressed. Areas for further improvements are identified, along with other longer term topics (e.g. better monitoring and more efficient use of tape resources), to continuously evolve our infrastructure to meet greater challenges towards HL-LHC.

Backup slides ...

ATLAS test results (Tier-1s)

Tier-1 sites	DT test (tape write)		A-DT test (tape read)		
	Target (GB/s)	Real rate (GB/s)	Target rate (GB/s)	Max rate (GB/s)	Stable rate (GB/s)
BNL	3.5	>8	1.9	6.8	<u>5</u>
CC-IN2P3	3.5	Up to 4	1.2	6.9	<u>2.6</u>
CNAF	3.5	>4	0.8	3.0	<u>1.9</u>
KIT	3.5	Up to 3.5	1.0	3.0	<u>1.6</u>
NL-T1	3.5	Up to 7	0.6	1.5	<u>1.2</u>
PIC	3.5	>6	0.3	2.1	<u>1.2</u>
RAL	3.5	Up to 3.5	1.2	2.1	<u>1.7</u>
TRIUMF	3.5	Up to 4	0.8	2.8	<u>1.6</u>

CMS Tier-1 staging test results

CMS site	March 2022 average rates (GB/s)	March 2022 Highest rates (GB/s)	October 2021 average rates (GB/s)	Ratio 2021/2022 (average rates)
<u>KIT</u>	1.81	5.26	2.34	1.29
<u>PIC</u>	1.20	1.68	1.84	1.53
<u>Fermilab</u>	1.51 (1.57)	2.26	2.46	1.62
<u>RAL</u>	0.18 (0.21)	1.05	--	
<u>JINR</u>	0.705	3.31	1.56	2.2
<u>CNAF</u>	0.973	2.03	1.07	1.10