

Mini data challenges

WLCG DOMA session in WLCG/HSF workshop

IJCLab, Paris

08 May 2025


Katy Ellis (STFC, RAL), Johannes Elmsheuser (BNL)

- DC27
 - **50% HL-LHC**
 - Confirmed by WLCG MB for early 2027
 - Our biggest challenge yet ?
 - Note 25% of the “flexible” target in D24 was already a **significant** challenge !
 - What will have changed in the 3 years since DC24 ?
 - Still 3 years before the start of the HL-LHC
 - Sites may not have HL-LHC hardware in place or be paying for the HL-LHC network connections, but we expect some incremental improvements
 - Maybe existing setup will be sufficient ... ?
 - New FTS version ?
 - Tokens ?
 - Possible change in approach ?
 - Re-assess the goals around the flexible model ?

- Mini data challenges
 - Keep up momentum in the 3 year gap between major challenges
 - Continue to try to improve our existing infrastructure
 - Test individual components to avoid getting lost in the noise of the major DCs which are testing all at once
 - Capacity and Capability mini challenges continue
 - Capacity challenges test the existing technology
 - Capability challenges test new technologies that might be used in future
 - Mini challenges should be limited in scope
 - USATLAS/USCMS capability test suggestions at:
<https://drive.google.com/drive/u/0/folders/1Af7hWa0Zm30EuqsV1PbekSjb--gXAsVG>
 - Have a look at <https://twiki.cern.ch/twiki/bin/view/LCG/DomaMiniChallenges> and the GoogleCalendar linked from there
 - Please contact us if you'd like to join or already planning to make similar tests outside of the existing (on-going) tests - we need volunteers!

Where are we now ?

DC24 dashboard (25% of HL-LHC), 60 days in Mar-May 2025, representative



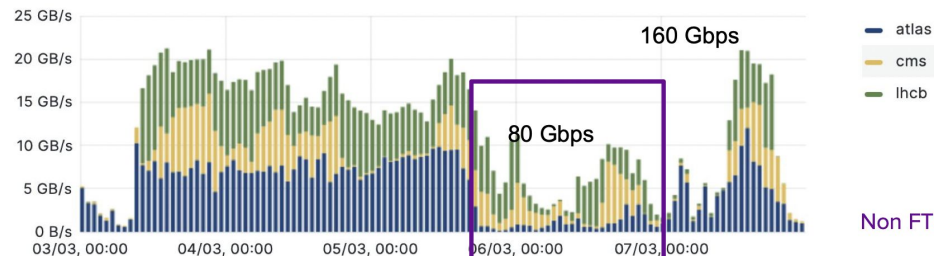
Note: this plot does not contain e.g. XRootD activity by CMS/ALICE nor any other non-FTS traffic

Several presentations with results from US and UK at DOMA general meetings in the past months:

- March 2025: UK mini DC
- Feb 2025: Jumbo frames CERN to US and UK, UK mini DC
- Jan 2025: US and UK mini DC

FTS vs total (IN to RAL)

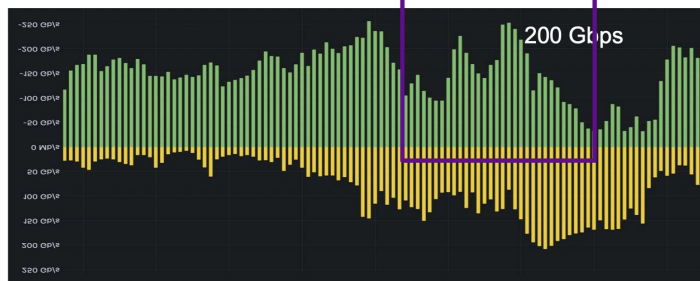
Transfer Throughput



Non FTS traffic to RAL.

CMS AAA jobs pulling data over root

New LHCb workflow pulling data over https



Alessandra Forti



CMS Testing plan with Jumbo frames January 2025

Request: WLCG monitoring
with white background

CMS Rucio: EOSPILOT →
FNAL

Data per test:

± 126TB reads

- FTS settings 2000/2000 min/
max

Key Observations:

- **20-25% faster transfers** with
Jumbo Frames for Rucio Activity.
- **Significant performance boost**,
especially over longer distances.

Standard Frames — MTU 1500



Jumbo Frames — MTU 9000



Conclusion: Jumbo Frames improve long-distance transfers (Geneva-Chicago, RTT 130ms) with 20-25% more bandwidth.

Maria Arsuaga-Rios

Maria Arsuaga-Rios - WLCG Doma 19/02/2025

11

Next WLCG DOMA mini data challenges

- Next tentatively planned mini challenges:
 - Tape testing in the UK
 - Strategy should be different to disk
 - Large file (30 GB) transfer tests in USCMS
 - From David Bouvet: France is planning a multi-VO tests at the Tier2s and also involving the Tier1 - similar to UK/US tests from the last months - more details after LCG-France workshop in June
 - From the Storm developers: new release with focus on improved read performance will be released soon - run a mini-DC if things look as expected
 - Others ?

Bottom up approach - engagement from regions to organize and execute these tests needed with support from existing experts

Taking stock before DC27

- DC24 was successful, but was it too ambitious?
 - **No** - we identified multiple bottlenecks which were better found sooner rather than later
- Does DC27 automatically need to be double DC24?
 - Not necessarily
 - We plan to recalculate our needs for HL-LHC based on the latest inputs
- Was too much pressure put on Tier 1s?
- Should DC27 involve tape?
- How to include throughput by direct transfer?
 - E.g. ALICE XRootD, CMS AAA, others
- There are many future unknowns
 - Will computing models change (Analysis Facilities, Data Lakes, etc.)?
 - If jobs run faster and storage isn't scaled up, will data have to move faster?
- Will work on updated “throughput tables” together with stakeholders in coming months
- *Don't buy hardware just for DC27*

