

# WLCG Data Challenge 24 Summary


WLCG Workshop  
2024-05-17

Christoph Wissing (DESY), Mario Lassnig (CERN)  
On behalf of the DC24 community



# DC24 in one plot

WLCG Throughput ⓘ



- Data Challenge
- atlas
- alice xrootd
- cms xrootd
- cms
- belle

**DC24 met the (main) goals [25% of HL-LHC]:**  
- Achieved full throughput of minimal model (1st week)  
- Had to push hard for flexible target (2nd week)


	max	avg	current
Data Challenge	2.19 Tb/s	1.02 Tb/s	211 Gb/s
atlas	625 Gb/s	304 Gb/s	567 Gb/s
alice xrootd	349 Gb/s	115 Gb/s	71.4 Gb/s
cms xrootd	191 Gb/s	67.4 Gb/s	42.7 Gb/s
cms	271 Gb/s	57.2 Gb/s	75.0 Gb/s
belle	38.9 Gb/s	9.45 Gb/s	17.1 Gb/s

- The minimal target was reached without problems
  - Too easy?
  - No! There were a number of pre-challenges, particularly targeting T0 and T1s
  - Intense preparation work across the community over several months
- Flexible target "just" met, a few things (almost) broke, a number of issues
- Still a success?
  - Yes!
  - Reminder: Production continued without interruption
  - Challenges are there to uncover weak areas, scaling limitations, insufficient concepts
    - FTS & Rucio running at unprecedented scales
    - First large scale using of token based authentication for transfers
  - Monitoring greatly improved, addressing the recommendations of DC21
    - WLCG WAN site monitoring
    - New XRootD monitoring
- Sites considered the challenge as very useful
  - No major problems, some localised WAN and/or IO saturation

# Token based Authentication




- Token infrastructure became ready for DC24 just in time
  - FTS pre-release with token support
  - Rucio with base set of features for ATLAS and CMS
    - Dedicated token caching
  - Deployment campaign to prepare storage elements
- About half of the transferred DC injected traffic via token authentication
  - Very high load on IAM for LHCb
    - Used 1 token per transfer
  - ATLAS switched tokens off at the end of 2nd week
    - Refresh was very expensive for FTS
  - Valuable experiences gained with token usage at production scale
- Need a process to agree on a reasonable approach how to use tokens in a sustained way
  - Token lifetime and granularity of scopes vs. load on central IAM service




# FTS operating at unprecedented scales

- Particularly FTS ATLAS instance survived thanks to permanent baby sitting by FTS team
  - Database surgery in production
  - Increase of hardware resources
- Improved understanding of current FTS scaling
  - Optimizer cycle needed several hours
  - FTS treats all links with the same activity with equal priority
  - Various ideas for more complex storage back pressure algorithms could be investigated
- FTS team started to iterate developments items and related priorities with stakeholders of the community
- Need a process to agree on development priorities
- First official FTS release with token support this spring

DC24 file transfers per FTS instance per hour




DC24 data volume transferred per hour



Plots show DC injected 'activity' only  
Parallel ongoing production not included!

- WLCG DCs should also (scale) test new technologies
  - Deployment can vary depending on level of matureness
- Some technical topics addressed in the context of DC24
  - Measures to improve monitoring
    - Site based network monitoring
      - Captures all traffic
    - Network flow marking
      - with *SciTags* and *UDP Fireflies*
  - Software Defined Networking (SDN)
    - NOTED
    - SENSE/Rucio
  - Low level network stack
    - BBR vs Cubic tests for congestion control



- **Aftermath of DC24**

- Have a set of 'lessons learned', e.g.,
  - What are realistic options for using tokens?
  - Behaviour of tools & components in unprecedented territory?
- Set priorities of for ongoing developments
  - VO & community specific tools such as Rucio, FTS
  - Storage middleware
  - Network equipment

- **Planning of next DC**

- So far nothing is set except the global target of **about 50%** of expected HL-LHC throughput
- Dates
  - Likely in 2026 (or even later?)
  - Almost for sure in LS3, which makes scheduling much easier for LHC experiments
- Participating experiments
  - LHC experiments, hopefully again Belle-2 and DUNE
  - Interest (already expressed during DC24) by JUNO, SKA, Neutrino experiments in Japan
- Experience shows that planning needs to start early: **1 year is *not enough***

- Incredible effort by everyone involved, a big thanks to y'all!
  - **87 pages** of detailed information, **53 authors!**
  - **Experiments** (LHC + DUNE + Belle II)
  - **Middleware** (FTS, Rucio, Monitoring, IAM)
  - **Networks** (Features, Infrastructure, SDNs)
  - **Sites** (16 sites incl. T1s & T2s)
- Final report available for comments at [https://docs.google.com/document/d/1k8-P-aSbBOoJ0XoFWqKSawX6rQZvxJPw\\_xvqNAMhtx4/edit?usp=sharing](https://docs.google.com/document/d/1k8-P-aSbBOoJ0XoFWqKSawX6rQZvxJPw_xvqNAMhtx4/edit?usp=sharing)
  - Please check your sections!
- Will upload to Zenodo on Monday May 27th
  - Get your comments in before that!
  - We will add a one-page executive summary/abstract
- We still think there is value in sending a ~10p article to CSBS



## WLCG/DOMA Data Challenge 2024

### Final Report

**Authors:** Arora Aashay, Agostini Federica, Arsuaga Rios Maria, Balcas Justas, Balci Berk, Barisits Martin, Benjamin Doug, Betev Latchezar, Carder Dale, Chauhan Rahul, Christidis Dimitrios, Dart Eli, Davila Diego, Dewhurst Alastair, Ellis Katy, Fernández Casaní Álvaro, Flix Molina Josep, Forti Alessandra, Garonne Vincent, Garrido Borja, Giacomini Francesco, Glushkov Ivan, Haen Christophe, Hoeft Bruno, Lassnig Mario, Lehman Tom, Litmaath Maarten, Lukaszczk Mark, Manrique Andres, Mascetti Luca, McKee Shawn, Miccoli Roberta, Morganti Lucia, Murray Steven, Musheghyan Haykuhi, Nappi Antonio, Ozturk Hasan, Pacheco Pages Andres, Paparrigopoulos Panos, Pardi Silvio, Paspalaki Garyfallia, Patrascoiu Mihai, Perez-Calero Yzquierdo Antonio, Robinson Kate, Rogovskiy Alexander, Sapunenko Vladimir, Shah Asif, Timm Steven, Vianello Enrico, Vokac Petr, Wissing Christoph, Yang Xi, Zani Stefano

**Editors:** Lassnig Mario, Wissing Christoph



- Network capacity of today is fine for DC24 rates
- Network capacity is actively being upgraded for "DC29"/HL-LHC rates
  - Need to understand concrete demands for DC26 (~50%)
  - When to purchase network equipment for HL-LHC?
- Token progress accelerated, needs a consolidation of approach now
- Most sites could cope with DC24 load
  - Some reports of IO overload though
- Monitoring progress very much appreciated
  - Some discrepancies remain to be investigated
  - New XRootD monitoring requires validation and full roll-out
- Some DC tools still to be improved for future DCs
  - Smooth out injections even more
  - Deletion performance is crucial to maintain throughput
- Ramp-up challenges were crucial to the success

On to Data Challenge 2026 !



DC26 Bridge