

Plans for the next WLCG Data Challenges

Katy Ellis / STFC & Shawn McKee / University of Michigan

WLCG DOMA <https://indico.cern.ch/event/1439566/>

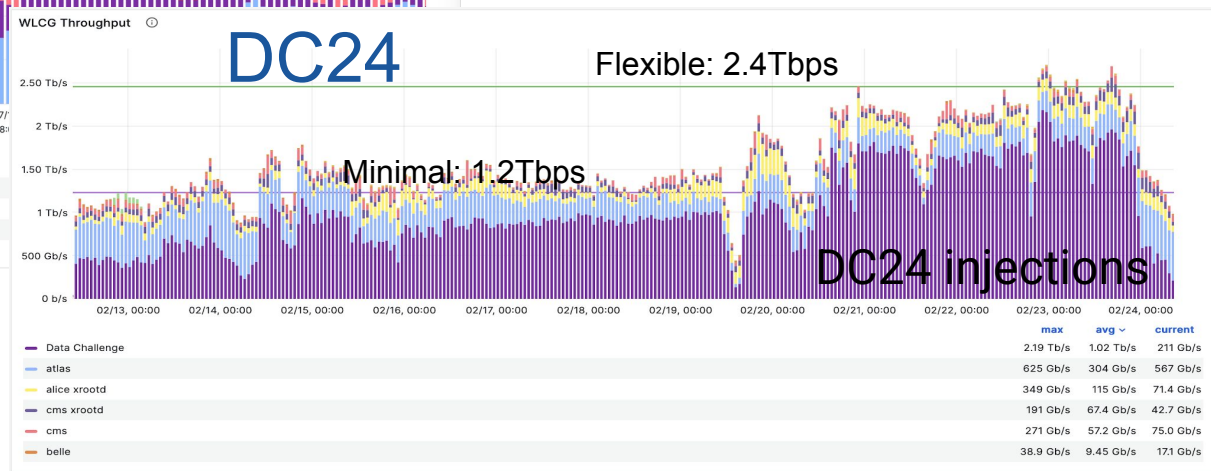
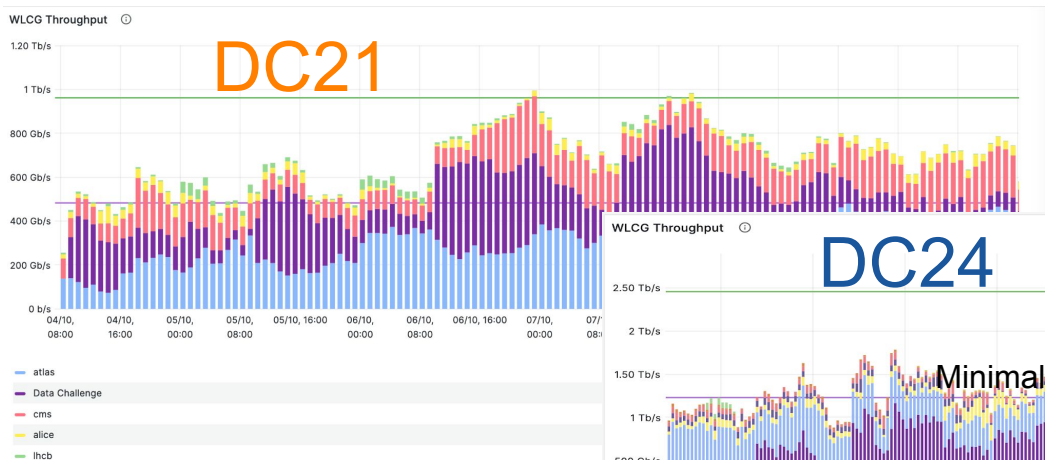
March 19, 2025



WLCG Data Challenges (1)

Previous data challenges (DC21 and DC24) were useful exercises

- a. Lots was learnt by experiments, middleware providers, and sites
- b. Network provision or performance were not identified as bottlenecks



c. Improved monitoring

WLCG Data Challenges (2)

DC27


- a. 50% HL-LHC aka DC27 (~Feb 2027?)
 - b. Confirmed by WLCG MB for early 2027
 - c. Our biggest challenge yet..?
-
- 25% of the ‘flexible’ target was already hugely challenging
 - What will have changed in the 3 years since DC24?
 - Still 3 years before the start of HL-LHC
 - So sites may not have HL-LHC hardware in place, or be paying for HL-LHC network connection
 - Maybe existing kit will be sufficient...?
 - New FTS version?
 - Tokens
 - Possible change in approach?
 - Re-assess the goals around the flexible model?



WLCG Data Challenges (3)


Mini data challenges

- a. Keep up momentum in the 3 year gap between major challenges
- b. Continue to try to improve our existing infrastructure
- c. Capacity and Capability mini challenges are in progress
 - i. Capacity challenges test the existing technology
 - ii. Capability challenges test new technologies that might be used in future
 - iii. Mini challenges should be limited in scope
- d. Have a look at <https://twiki.cern.ch/twiki/bin/view/LCG/DomaMiniChallenges> and the GoogleCalendar linked from there
 - i. The US and UK have already run a first set of capacity challenges
 - ii. Capability challenges have also started - e.g. jumbo frames (see Maria's talk, next)
- e. Please contact us if you'd like to join or already planning to make similar tests outside of the already existing (on-going) tests - see next slide



Where are we now?

DC24 dashboard (25% of HL-LHC), last 30 days, representative



Run-3 was 'supposed'
to be ~10% of HL-LHC

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

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?

Capability testers

We continue to solicit advocates and mini-challenge planners/operators.

- See [Google folder](#)
- If a capability you are interested in is not there, add it!
- **Please participate.**

US Goal is to continue running various capability mini-challenges leading up to a capacity mini-challenge in June-July 2025 (needs VO buy-in)

- We would like to see some beneficial capabilities (like Jumbo frames) part of many sites' production infrastructure by then.
- Are the R&E networks interesting in participating in any way?

In general, we plan to continue both types of mini-challenges (capability and capacity) through the last WLCG Data challenge in 2029.

Mini-Challenge Components

The important topics to plan for and discuss for mini-challenges:

WHO? We cannot do a mini-challenge without participants. The guiding principle so far has been “bottom up” organization, where advocates of technologies, sites or experiments initiate the mini-challenges.

WHAT? We need to understand the items to test and measure. Will the mini-challenge focus on simply capacity/load-testing between select sites or will it also test capabilities at some scale?

WHEN? Scheduling can be critical, since we need the “WHO” able to run the challenge and we may want to co-schedule with other mini-challenges while avoiding critical time periods for the experiments.

Known Issues

Do we have the needed monitoring to properly understand each mini-challenge?

- What is missing or incorrect?
- For technology focused mini-challenges, what new monitoring is needed?
- Will the WLCG Monitoring Task-force be responsible for this?

Can we easily run the mini-challenges?

- The `dc_inject.py` (from Mario) and Hiro's load-testing scripts likely need improvements to enable easy use
- How to converge-on and then create the toolkit? (Who, When)

Questions / Discussion

I hope we can continue fruitful discussions about mini-challenges and here are some “seed” questions/topics we have been exploring:

Are there new advocates willing to participate (including bringing new capabilities to test)?

What load-testing (**capacity**) mini-challenges are needed (specific sites, regional, targeting possible shared network “choke” points)?

How best to share information and scheduling? (Google docs, WLCG DOMA calendar, ...)

How to make progress on the tooling and infrastructure that enables mini-challenges?

The Floor is Open for Discussion...

Background Material


Here are some resources we know about:

Presentations

- [WLCG Data Challenge 2024 \(DC24\) Status and Plans Related to ATLAS DDM](#) (Jun 2023)
- [DC24 Planning and Near Term Activities](#) (Jul 2023)
- [USATLAS Data Challenge 2024 Take-aways](#) (Feb 2024)
- [Medium to Long Term Network Plans for ATLAS and CMS](#) (Mar 2024)
- [DC24 Network Activities & Results](#) (May 2024)

Some Google Docs

- [WLCG/DOMA Data Challenge 2024: Final Report](#)
- [USATLAS Milestones/MiniChallenges for Next WLCG Data Challenge in 2024](#)
- [Planning Mini-Challenges for US ATLAS Facilities and Distributed Computing](#)
- [NOTES: USATLAS Facility Status and Evolution Discussion](#)



Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- [OSG: NSF MPS-1148698](#)
- [IRIS-HEP: NSF OAC-1836650](#)



Open Science Grid



WLCG
Worldwide LHC Computing Grid



Backup Slides Follow

ESnet Monitoring for WLCG Data Challenge

ESnet created a very nice [monitoring dashboard](#)

LHC Data Challenge / LHC Data Challenge Overview

Last 6 hours


LHC Data Challenge Overview

Menu: Overview | Interfaces | Sites | Transatlantic | LHCOPN

This dashboard shows an overview of statistics relevant to the LHC data challenge. It contains a combination of SNMP and flow statistics from ESnet's Stardust measurement system. Use the navigation menu above this text or links in the data below to move to other dashboards that provide different views of the data.

SNMP Statistics

Top 10 Interfaces by Incoming Rate (SNMP)

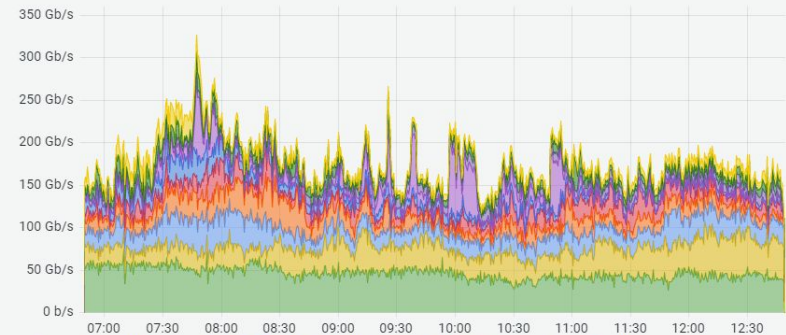


Legend for Top 10 Interfaces by Incoming Rate (SNMP):

- star-cr55:to_fnal-site-lhcone
- newy32aaa-cr6:mit_se-115
- cern-513-cr5:to_lhcone_geant
- chic-cr55:to_lhcone_uchicago-b_v4
- kans-cr5:to_lhcone_unl
- bnl515-cr6:bnl_se-105
- star-cr55:to_lhcone_aglt2-a
- bnl515b-cr6:bnl_se-101
- chic-cr55:aglt2_se-591
- pnwg-cr5:to_lhcone_canet

Top Interfaces by Incoming Volume (SNMP)

Top 10 Interfaces by Outgoing Rate (SNMP)



Legend for Top 10 Interfaces by Outgoing Rate (SNMP):

- wash-cr5:to_lhcone_internet2
- cern-513-cr5:to_lhcone_geant
- kans-cr5:to_lhcone_unl
- amst-cr5:to_lhcone_geant
- star-cr55:to_lhcone_in-gpop-b
- chic-cr55:to_lhcone_uchicago-b_v4
- bnl515b-cr6:bnl_se-101
- atla-cr5:to_lhcone_fir-a
- star-cr55:to_lhcone_aglt2-a
- chic-cr55:to_lhcone_uwmadison-b

Top Interfaces by Outgoing Volume (SNMP)