# Data Grand Challenge – Looking toward DC24

Brian Bockelman

# Data Grand Challenge

- We expect the HL-LHC will require a ~20x increase in data rates compared to the beginning of IRIS-HEP v1.0.
  - Recently it's been "all bad news": requirements appear to be creeping upward as we gain a better understanding of the HL-LHC environment.
- The community has defined a series of WLCG-coordinated "Data Challenges", milestones to show progress. DCXX aims:
  - Hit data transfer rate milestones (right)
  - Technology maturity (e.g., tokens)
- Each value of XX has its own goals. Will focus on DC24 today.

| **Year** | **Minimal** (Gbps, %) | **Flexible** (Gbps) |
|---|---|---|
| 2021 ✅ | 480, 10% ✅ | 960 ✅ |
| ~~2023~~ 2024 | ~~1,440, 30%~~ 25% | ~~2,880~~ |
| ~~2025~~ 2026 | 2,880, 60% | 5,760 |
| ~~2027~~ 2028 | 4,800, 100% | 9,600 |

Note: above shows both original plans from 2020 and how they are evolving!

# DC21

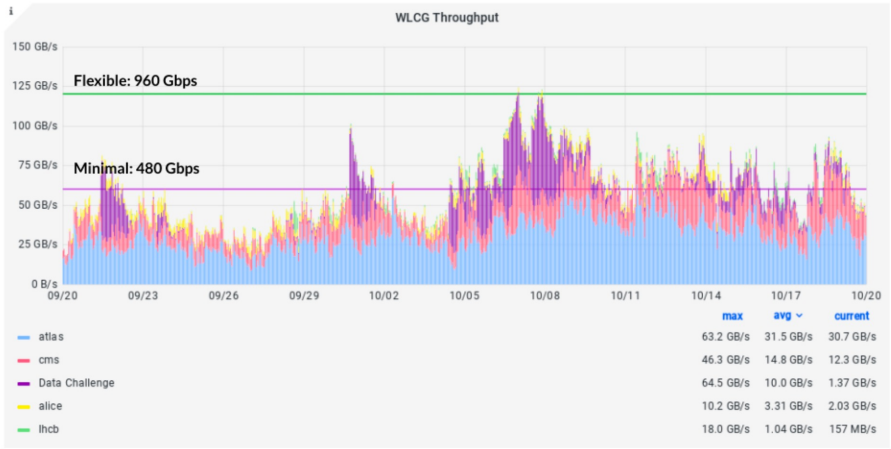


Figure 1 - Mock DC1 22/09/2021; Mock DC2 01/10/2021; Network Challenge (DC) 04-10/10/2021; Tape Challenge 11-19/10/2021.

Figure reproduced from

https://zenodo.org/record/5767913

The first data challenge, DC21, was highly successful in that it

- Demonstrated the coordination and tooling necessary to run a data challenge,
- Showed technology readiness for HTTP-TPC, *and*
- Met its data rate milestones (peaking at 960Gbps).

DC21 was not a particular stretch in data rates – in the same neighborhood of what is needed for Run3 – but the first two were quite significant!

# Collecting Inputs for DC24

- We are in the process of collecting inputs from experiments and technology providers on their plans for DC24.
- Experiments include those beyond LHC: Belle II, JUNO,  DUNE. Networks are a shared resource -- its useful to have as broad participation as possible!
- Example technology demos (organized outside IRIS-HEP) proposed:
  - ALTO project's integration with FTS/Rucio for fine-grained network tuning.
  - Exploring the performance of BBR-based TCP congestion control (CERN networking).
  - Use of packet pacing to improve overall throughput (ESNet, USATLAS).

# Example Input - CMS

Starting point: Minimal model
T0 export to T1s:

| RSE | Total space (TB) | Proportional share | Rate (GB/s) |
|---|---|---|---|
| T1_DE_KIT_Disk | 8530 | 0.106 | 3.277 |
| T1_ES_PIC_Disk | 4100 | 0.051 | 1.575 |
| T1_FR_CCIN2P3_Disk | 8000 | 0.099 | 3.074 |
| T1_IT_CNAF_Disk | 10000 | 0.124 | 3.842 |
| T1_RU_JINR_Disk | 10600 | 0.131 | 4.073 |
| T1_UK_RAL_Disk | 7252 | 0.090 | 2.786 |
| T1_US_FNAL_Disk | 32200 | 0.399 | 12.372 |
| | | | 31.000 |

31GB/s is approximately 250Gbps.

TBD: Rates for MC production traffic and other traffic to and from T1s and T2s. Probably equivalent traffic as above for export from T1 to T2, and again the same for T2 to T1. Total traffic would be around 750Gbps.

- N.B. – Giving CMS as an example because they've posted their initial plans already.
  - Hoping for more LHC experiments this week!
- In DC21, we took base load from the experiment's production activities and injected added traffic from a WLCG-run Rucio instance.
  - CMS plan is to use load generation tool within their Rucio instance.
  - Will start doing preliminary tests in early November.

# Technology Demos – WLCG BDT Group

- The WLCG Bulk Data Transfer (BDT) group is pushing to have end-to-end token-based authorization in use during DC24.
- **Goals**: 5% of sites have token-based transfers performed during DC24. Target to have success rates with tokens comparable to overall success rates.
  - My commentary: Given existing progress in the US and elsewhere, I think the 5% threshold is conservative. CMS has SAM/ETF test coverage for token functionality and is in the middle of a deployment campaign.
- Currently-proposed milestones leading to DC24:
  - October: Rucio deletions with tokens
  - October: Accept updated JSON submission format (Rucio data distribution depends on this new interface)
  - November (DC24 Workshop): demonstrate Rucio deletions, basic FTS transfers with tokens using new submission format
  - January 2023: mini-challenge: stress test for FTS implementation

# Technology Demos – Rucio/SENSE integration

- Rucio/SENSE aims to provide Rucio with capabilities to request network services via SENSE to: a) improve accountability, b) increase predictability, and c) isolate and prioritize transfer requests.
  - This project uses dedicated Rucio and XRootD instances so it does not interfere with Production systems.
- **Goals**: Demonstration of Rucio-initiated priority paths between 3 different pairs of sites
  - UCSD <- FNAL (100Gbps)
  - Caltech -> UNL (100Gbps)
  - UCSD -> Caltech (400Gbps)
- Milestones:
  - Demonstrate above matrix of transfers by November 2023.

# Upcoming Advertisement – Data Challenge Workshop 2023 @ CERN

- The WLCG is organizing a Data Challenge Workshop this November 9-10. Stated Objectives:
  - Discuss the latest developments of the DOMA R&D program,
  - First results from the ramp-up challenges, and
  - to finalize the plan for DC24.
- In the same week: Pre-GDB (Tuesday) on tape-storage and a GDB (Wednesday).
- In other words,
  - Far off enough we can wrap up various summer projects.
  - Soon enough that we can do course corrections before the data challenge.
  - This is an opportunity to help refine the scope and make concrete plans.

https://indico.cern.ch/event/1307338

Hosted by WLCG.

# Questions?

This project is supported by the National Science Foundation under Cooperative Agreements OAC-1836650 and PHY-2323298. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.