WLCG Data Challenge 24
- A first & preliminary look -

DCMS Meeting
2024-02-29

Christoph Wissing (DESY)

LHC & High Lumi LHC

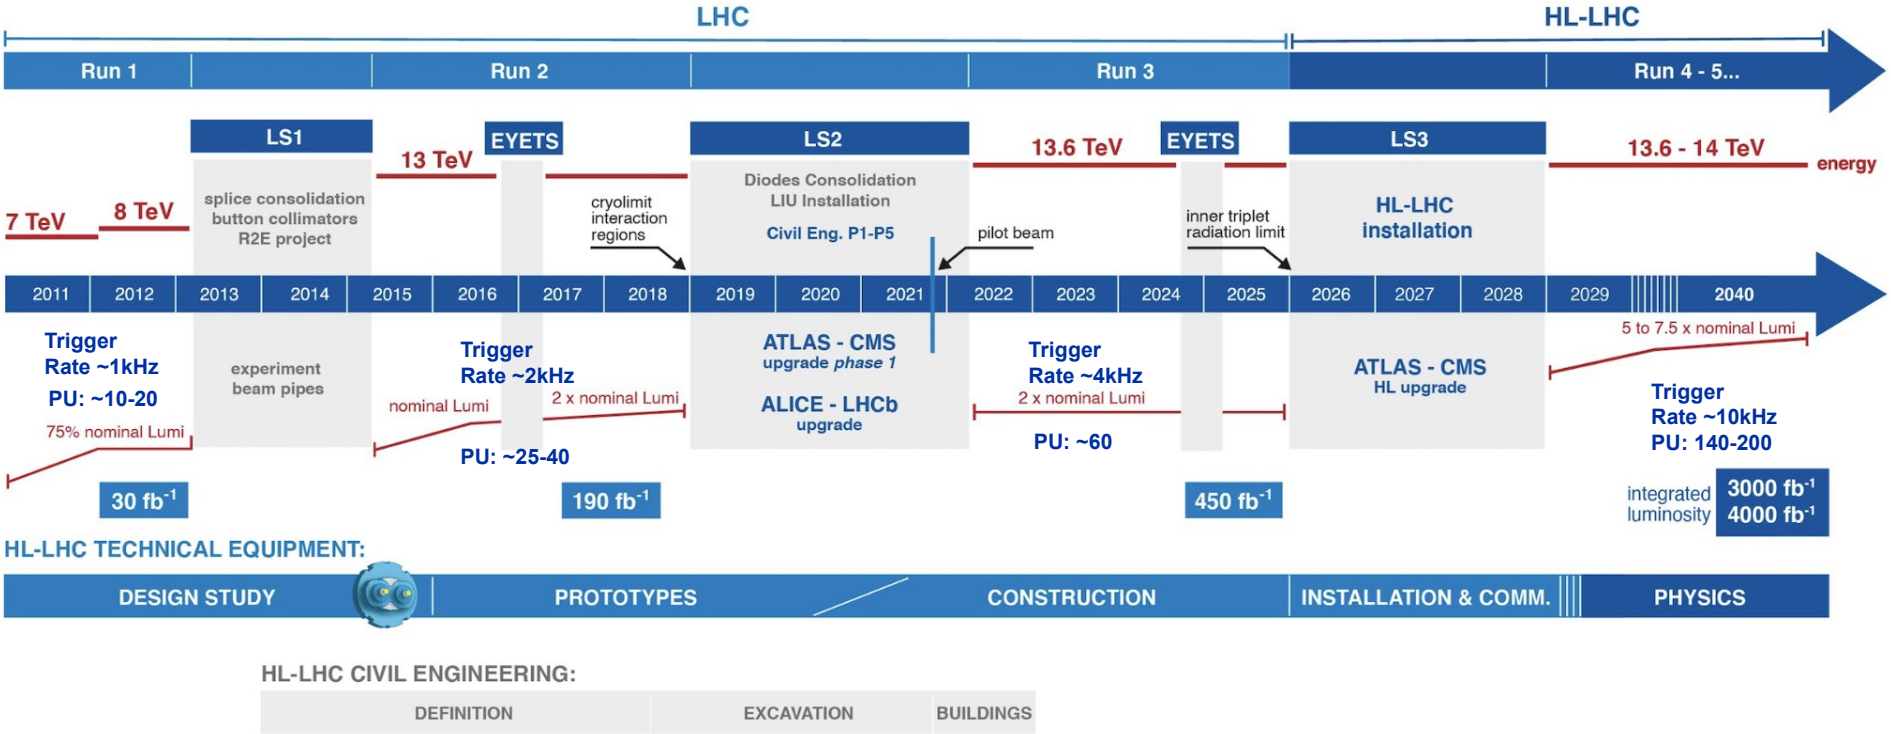

Figure adopted from:
Zerlauth, Markus & Bruning, Oliver. (2024). Status and prospects of the HL-LHC project. DOI; 615. 10.22323/1.449.0615.

Data Challenges for HL-LHC

- WLCG has been mandated to execute data challenges (DC) for HL-LHC
  - Demonstrate readiness for expected HL-LHC data rates by a series of challenges
  - Increasing volume/rates
  - Increase complexity (e.g. additional technology)
  - A data challenge roughly every two years

- DOMA is the coordination and execution platform
  - Data Organization Management & Access
    - Forum across all LHC experiments to address **technical** needs and challenges
  - For the DCs find agreements across the LHC experiments and beyond
    - Suited dates
    - Reasonable targets
    - Functionalities
  - Help in orchestration

- Dates and high level goals always approved by WLCG Management Board

Recap of (initial) modelling & resulting rates for HL-LHC

ATLAS & CMS T0 to T1 per experiment

350PB RAW per year, taken and distributed during typical LHC uptime of 7M seconds

- 50GB/s or 400Gbps

Another 100Gbps estimated for prompt reconstruction data tiers (AOD, other derived output)

1Tbps for CMS and ATLAS summed

WLCG data challenges for HL-LHC - 2021 planning
https://zenodo.org/records/5532452

ALICE & LHCb T0 Export

100 Gbps per experiment estimated from Run-3 rates

Minimal Model

Sum (ATLAS,ALICE,CMS,LHCb)*2(for bursts)*2(*overprovisioning*) = **4.8Tbps for the expected HL-LHC bandwidth needs**

Flexible Model

Assumes reading of data from above for reprocessing/reconstruction in 3 months (about 7M seconds)
Means doubling the Minimal Model: **9.6Tbps for the expected HL-LHC bandwidth needs**
However data flows primarily from the T1s to T2s and T1s!

Data Challenges target: **50% filling of expected** HL-LHC bandwidth needs

DC21 - 10% of HL-LHC Troughput

However, we managed to fill 100% of the (minimal) DC21 target!

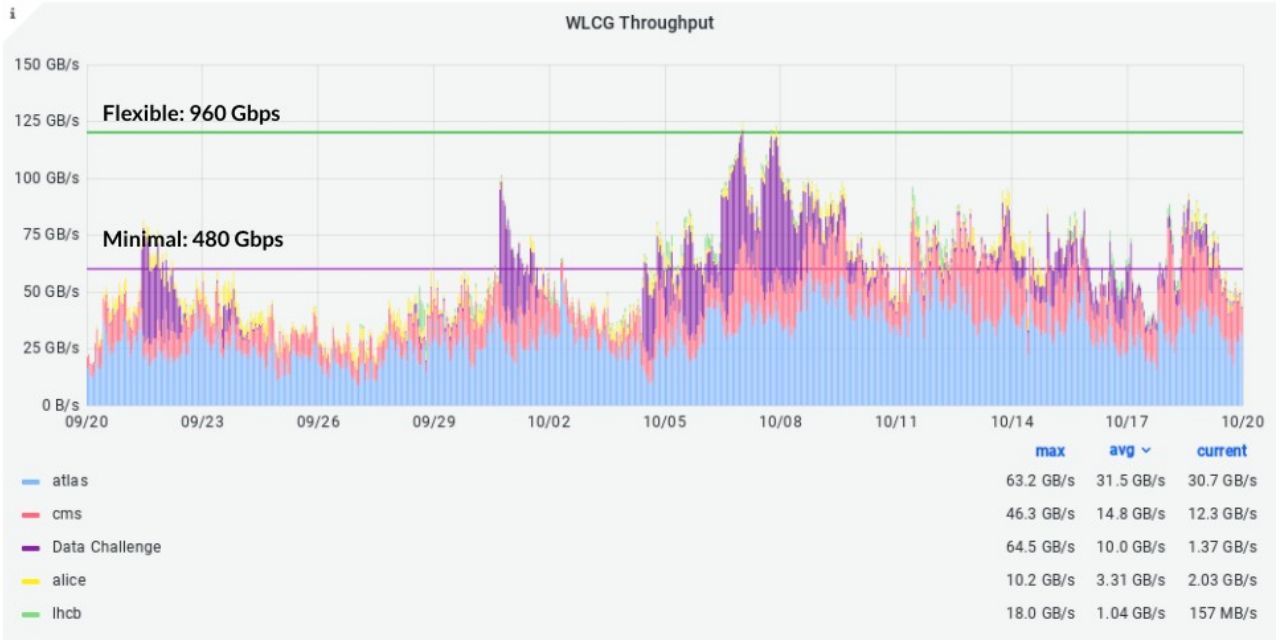

Network Data Challenges 2021 wrap-up and recommendations

https://zenodo.org/records/5767913

Planning of DC24

- Overall target: **25%** of HL-LHC throughput
  - Slightly lowered from originally 30% due to delayed start of HL-LHC
- Long way to towards the DC24 program
  - Agreement on dates:
    2 weeks before beam operation in 2024
  - Full transfers from disk to disk,
    Not just network traffic
  - Experiments had room to optimize their set of exercises
    - ALICE and LHCb involved tapes,
      ATLAS and CMS decided not to
  - Preparation of monitoring
  - Regular preparation started one year before
    - DOMA general meetings
    - Dedicated workshop in Nov 2023

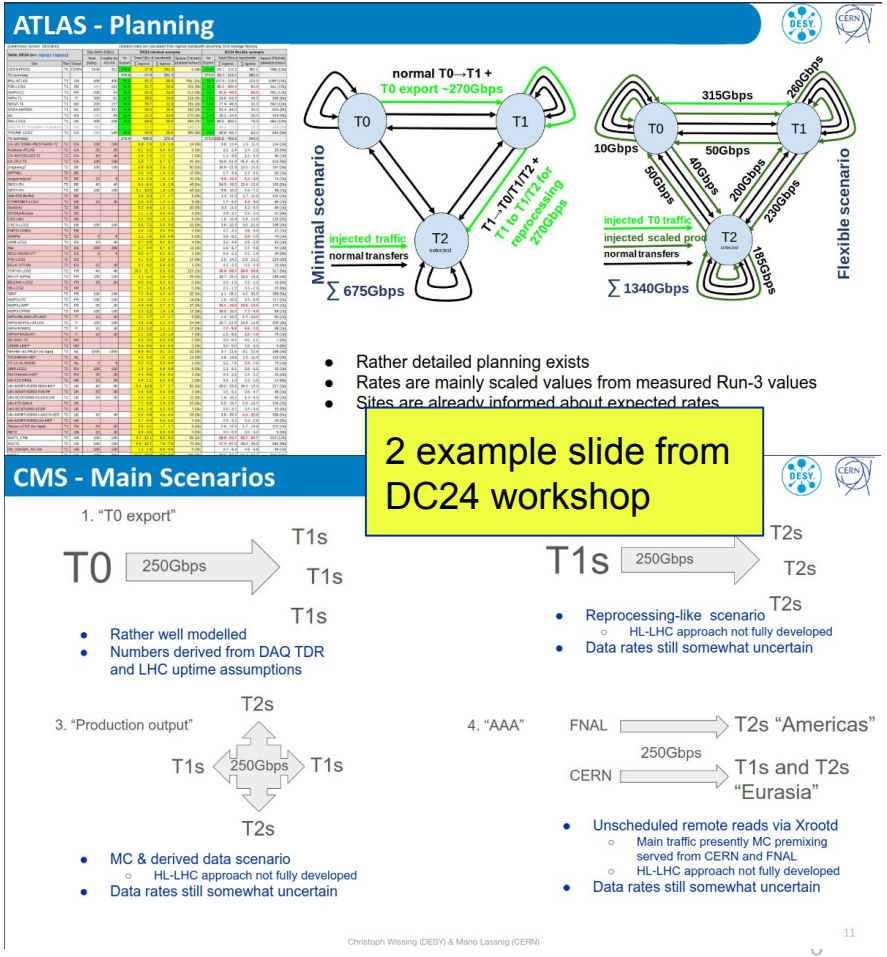

DC24 in one Plot

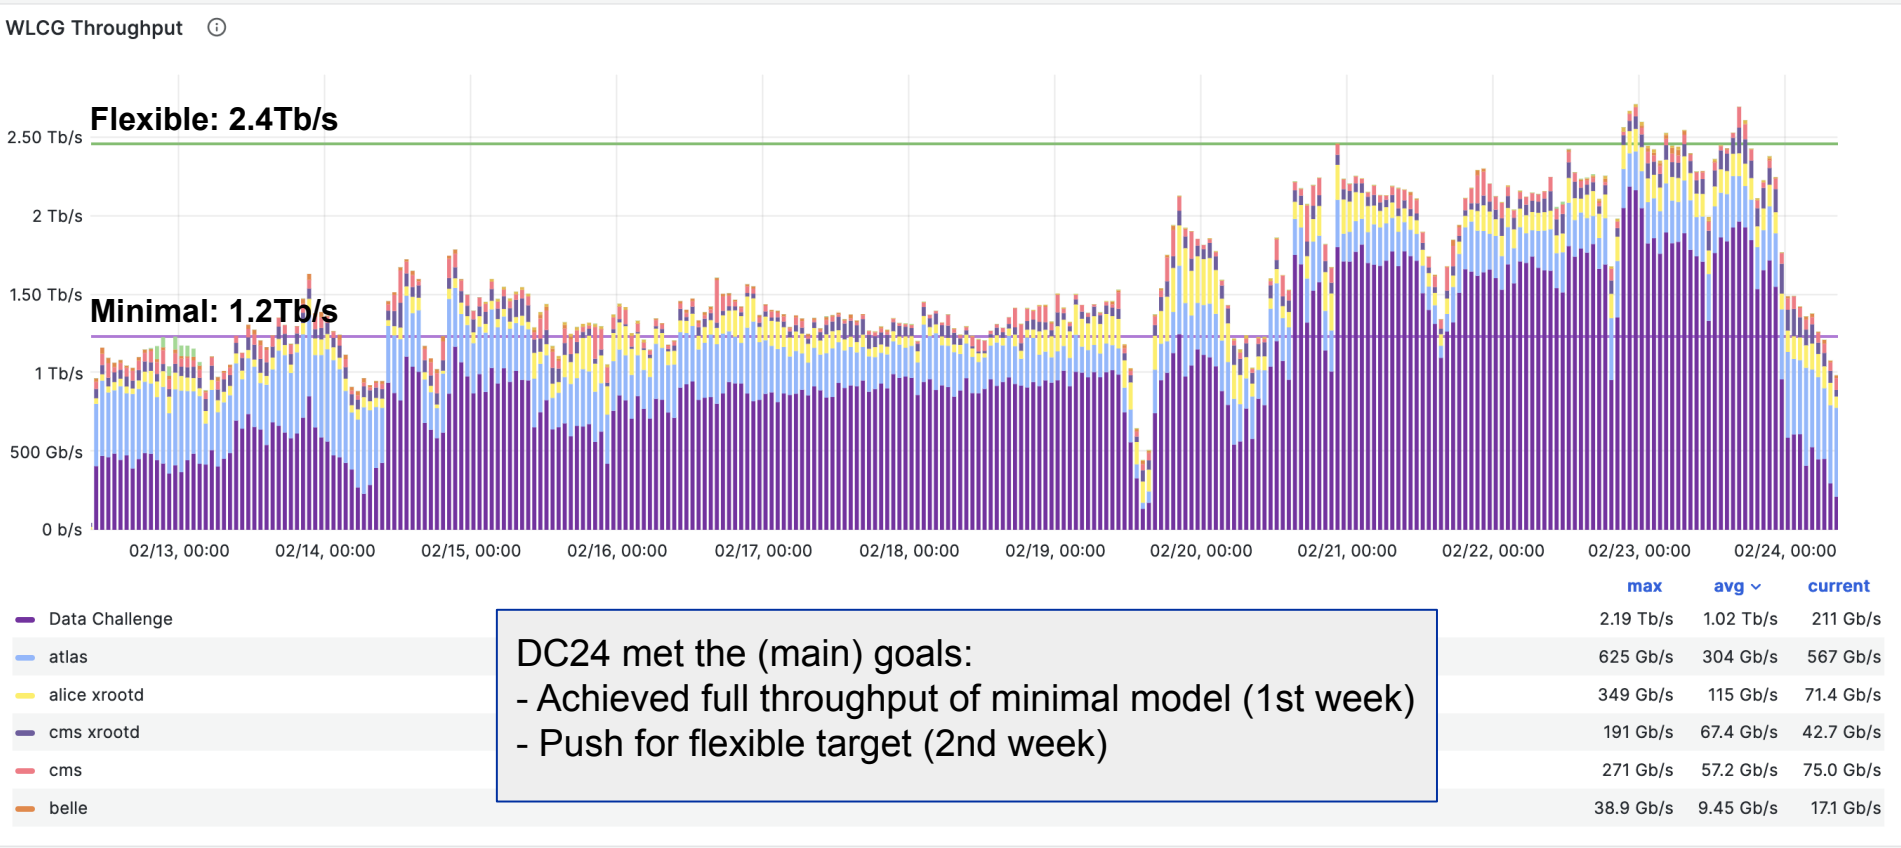

DC24 met the (main) goals:
- Achieved full throughput of minimal model (1st week)
- Push for flexible target (2nd week)

Beyond throughput

- 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
  - Token based authentication
    - About half of the throughput authenticated via tokens
  - 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
    - Jumbo frames
    - BBRv2, BBRv3 TCP stacks

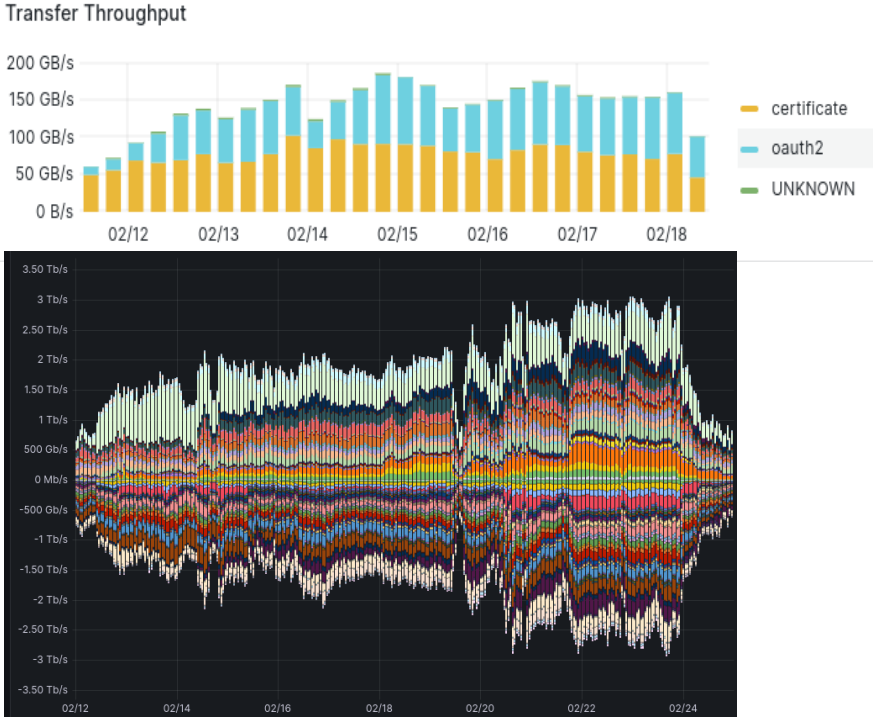

After the Challenge is before the next Challenge

- Aftermath of DC24
  - Derive 'lessons learned'
    - What went well, where were bottlenecks, organizational improvements …
  - Set priorities of for ongoing developments
    - VO & community specific tools, e.g. Rucio, FTS,
    - Storage middleware
    - Network equipment
- Planning of next DC
  - So far nothing is set except the global target of **about 60%** 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, likely 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 before, at least)

Some Random preliminary Observations & Remarks

- There are other bottlenecks than network bandwidth
  - Maintenance of DC injections was challenging
    - FTS instances got pushed to their limits, particular the ATLAS one
    - Keeping up with deletions is not trivial, systems not designed for best scaling here
  - It needs time before a complex system reacts to parameter changes
    - The parameter space is huge
    - Not many attempts to re-adjust (a very few per day)
  - A number of CMS sites asked for more (than planned) traffic to exercise their WAN connectivity
- German sites
  - Did in general rather well during the challenge
  - DESY 100GBit WAN link close to max during 2nd week (ingress & egress)

What AI thinks we are doing …

Bing Image Creator: "Worldwide LHC Computing Grid, Data Challenge Workshop, Happy Mood"

Bing Image Creator: "Worldwide LHC Computing Grid, Data Challenge Workshop, Serious Mood"