

# Data and Network challenges in preparation of HL-LHC

PRACE-CERN-GÉANT-SKAO kick-off workshop

29<sup>th</sup> of September 2020

[edoardo.martelli@cern.ch](mailto:edoardo.martelli@cern.ch)

# Data production in HL-LHC

CMS and ATLAS will produce ~350PB of raw data per year, running for 100 days/year

**The traffic from CERN to all the T1s will be ~400Gbps for 100 days per year per experiment**, as it will export in quasi-real time.

- e.g. FNAL (40% of resources for CMS) will import at ~150 Gbps

- e.g. BNL (25% of resources for ATLAS) will import at ~100 Gbps

Therefore it is estimated the need for **3-4Tbps of network capacity from CERN to the T1s** by the time of HL-LHC, of which **1-1.5Tbps will be needed across the Atlantic** to cover the needs of ATLAS and CMS

# Data and Network challenges

Network challenges could consist in **demonstrating the capability to transfer an increasing volume of data over the next years to reach the production transfer target, sustained for a few days, by the start of HL-LHC in 2027.**

We could foresee **milestones as 15% of the target 2021, 35% in 2023, 60% in 2025 and 100% in 2027.**

This could be adjusted based on the growth plan of the NRENS.

# T1-T2: Reprocessing at T2s

The data at the T1 needs to be staged from tape and exported to the T2s for processing

**The target is to be able to reprocess 100% of the data collected in the year and stored at a specific T1 in less than three months.**

The data could be streamed directly to the processing centres or buffered at the T1 and transferred in a burst. This has different implications on the storage needs at T1s and T2s, the balance with CPUs and the network needs.

**A T1 will need to commission its capability to stream an aggregated 1Tbps to the T2s.** The 1Tbps T1 egress capacity is the target for 2027 for a 40% T1 serving only one experiment (e.g. FNAL). The targets for the other T1s can be derived from there.

# Reprocessing at HPCs

HPC will also be used for reprocessing


Tier1s will need to stage the data to be reprocessed at the HPC premises

The use case where an HPC would provide an allocation of 5k nodes (128 cores each) for many days capable to process 10kHz of events, implies **demonstrating the capability to stream 1Tbps of data into a HPC in 2027**

Intermediate targets should be defined for the coming years

# European challenge: T0 to T1s


- For Tier1s in Europe several links could be aggregated, leveraging the large availability of dark-fibres
- Target: 1Tbps



E.g. soon available 400Gbps test link between CERN and NL-T1

# T1 challenge: reprocessing at T2s




- T1s can aggregate existing and available bandwidth
- Target: 1Tbps



**WLCG T1**

**LHCONE**


**NRENs connectivity**



**WLCG T2s**

# HPC challenge: reprocessing

- Connect HPC centres to R&E networks at large bandwidth, possibly to LHCONE
- Target: 1Tbps



WLCG T1

LHCONE

NRENs connectivity




Beware:

- HPC centres not always equipped with large Internet access
- HPC may not use Ethernet, would require media converters



# Connectivity options for HPC centres

- HPC pulls data from the Tier1
- Staging at local storage at HPC centre
- Use of DTNs (Data Transfer Nodes) and caches



**WLCG T1**



**HPC centre**

# What's next

- Involve Experiments and Storage-Transfer service managers
- Define set of meaningful data challenges
- Assess the aggregated network capacity of existing and upcoming storage
- Agree on the schedule of the challenges

*Questions or comments?*