


WLCG perspective

S. Campana (CERN) @ ATCF2022


We are in LHC Run-3

Data written in the CERN tape storage per month (since 2008)



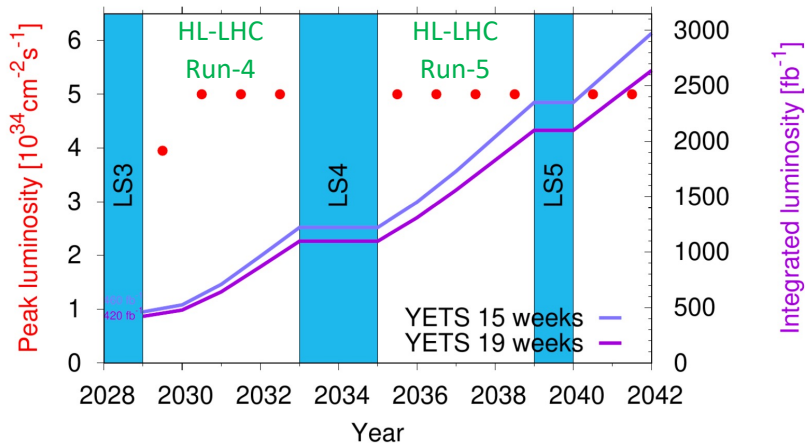
Run-3 is a major challenge for LHCb and ALICE (new detectors, new computing models)
Run-3 brings novel challenges also for ATLAS and CMS (high trigger rates)

WLCG processing



The WLCG processing does not know about LHC shutdowns. The level of resources in use continue increasing.

The HL-LHC challenge




The challenge is about

- Higher data volume
- Higher event complexity

In the context of

- Constrained funding
- Sustainability concerns
- Global politics
- Security and trust

Driven by ATLAS and CMS, though ALICE and LHCb will also be a challenge, particularly Run-5



Run-3/4/5, pile-up estimated to average 54/140/200

The HL-LHC challenge statement at ECFA 2016



In a nutshell:

- Projections assume constant funding every year for LHC computing
- Technology improvements will bring ~20% more resources every year

And this was the initial conclusion

ECFA2016

Estimates of resource needs for HL-LHC



Storage
Raw 2016: 50 PB → 2027: 600 PB
Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU
x60 from 2016

Technology at ~20%/year will bring x6-10 in 10-11 years

=> x10 above what is realistic to expect from technology with constant cost





Simone.Campana@cern.ch - ECFA2016


3/10/2016

4

ATLAS and CMS needs for HL-LHC



- The gap between available and needed resources is filling up, assuming the main R&D activities are successful. There are still large uncertainties



Networks' role in HL-LHC

Networking will play a central role in HL-LHC as enabler for HEP computing

- Support the core functions of WLCG (data acquisition/archival/processing)
- Provide more flexibility to the computing models, allowing to optimise

WLCG continues engaging with Funding Agencies and NRENs to ensure that enough capacity is made available and the LHC traffic does not get segregated below a critical level.

Several R&Ds were launched to study how to better leverage the network resources in the data and processing infrastructures for HL-LHC

- Regularly discussed at the LHCONe/LHCOPN [meetings](#)

WLCG network data challenges

The data challenges are an incremental **process** to prepare for the HL-LHC network needs, through a regular dialog between the network providers, the experiments and the facilities.

We identified the main use cases at HL-LHC in terms of network use (RAW data export and reprocessing), for the 4 LHC experiments

We estimated the network needs including contingency and considering different scenarios (“baseline” form minimal use and “flexible” to exploit the full capabilities)

We set metrics and intermediate targets to be progressively challenged

The challenges offer the possibility to bring in production many network R&D activities

The data challenge planning document can be found at [this link](#)

Bandwith Needs

	%ATLAS	%CMS	% Alice	% LHCb	ATLAS+CMS Network Needs (Gbps) Minimal Scenario in 2027	Alice Network Needs (Gbps) Minimal Scenario in 2027	LHCb Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Flexible Scenario in 2027
T1									
CA-TRIUMF	10	0	0	0	200	0	0	200	400
DE-KIT	12	10	21	17	450	80	70	600	1200
ES-PIC	4	5	0	4	180	0	20	200	400
FR-CCIN2P3	13	10	14	15	450	60	60	570	1140
IT-INFN-CNAF	9	15	26	24	480	110	100	690	1380
KR-KISTI-GSDC	0	0	12	0	0	50	0	50	100
NDGF	6	0	8	0	110	30	0	140	280
NL-T1	7	0	3	8	140	10	30	180	360
NRC-KI-T1	3	0	13	5	50	50	20	120	240
UK-T1-RAL	15	10	3	27	490	10	110	610	1220
RU-JINR-T1	0	10	0	0	200	0	0	200	400
US-T1-BNL	23	0	0	0	450	0	0	450	900
US-FNAL-CMS (atlantic link)	0	40	0	0	800	0	0	800	1600
					1250	0	0	1250	2500
Sum	100	100	100	100	4000	400	410	4810	9620

This table presents the bandwidth needs per T1 (ingress from CERN and egress to the T2s)

For Asia the ingress from CERN makes sense, the egress to T2s does not really

Target Rates


T1	LHC Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Flexible Scenario in 2027	Data Challenge target 2027 (Gbps)	Data Challenge target 2025 (Gbps)	Data Challenge target 2023 (Gbps)	Data Challenge target 2021 (Gbps)
CA-TRIUMF	200	400	100	60	30	10
DE-KIT	600	1200	300	180	90	30
ES-PIC	200	400	100	60	30	10
FR-CCIN2P3	570	1140	290	170	90	30
IT-INFN-CNAF	690	1380	350	210	100	30
KR-KISTI-GSDC	50	100	30	20	10	0
NDGF	140	280	70	40	20	10
NL-T1	180	360	90	50	30	10
NRC-KI-T1	120	240	60	40	20	10
UK-T1-RAL	610	1220	310	180	90	30
RU-JINR-T1	200	400	100	60	30	10
US-T1-BNL	450	900	230	140	70	20
US-FNAL-CMS (atlantic link)	800 1250	1600 2500	400 630	240 380	120 190	40 60
Sum	4810	9620	2430	1450	730	240

Same as the previous slide, but setting targets for the data challenge

2021 network data challenges

We executed data challenges testing the WLCG network and archive storages in preparation for Run-3 (Oct 2021)

The data challenges are an end-to-end test including storage, protocols, networks, data management services (e.g. Rucio, FTS)



Targets were met:

- the nominal transfer rate was sustained
- the peak transfer rate were reached


2021 network data challenges

The data challenges are not just about throughput but also functionality

In 2021 they provided an opportunity to commission new features that are now in use during Run-3.

- For example the HTTP protocol (replacing gridFTP) for asynchronous transfers

Transfer Rate per protocol

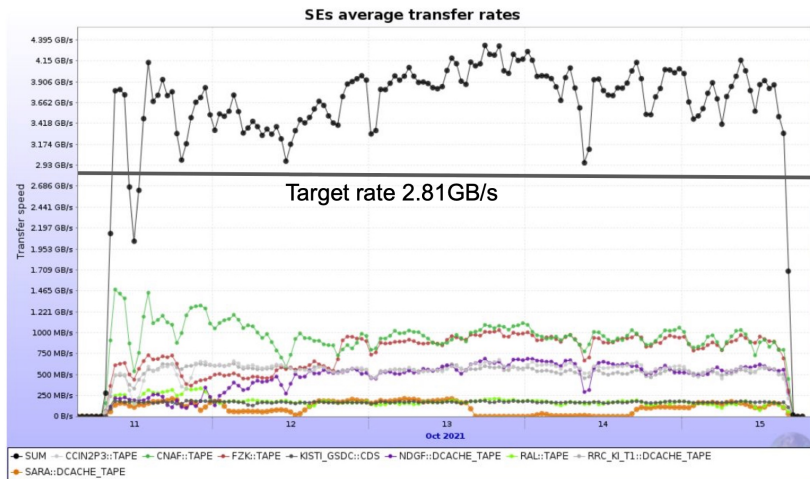


Traffic mostly through HTTP (**RED**)

Commissioning for Run-3: tape systems


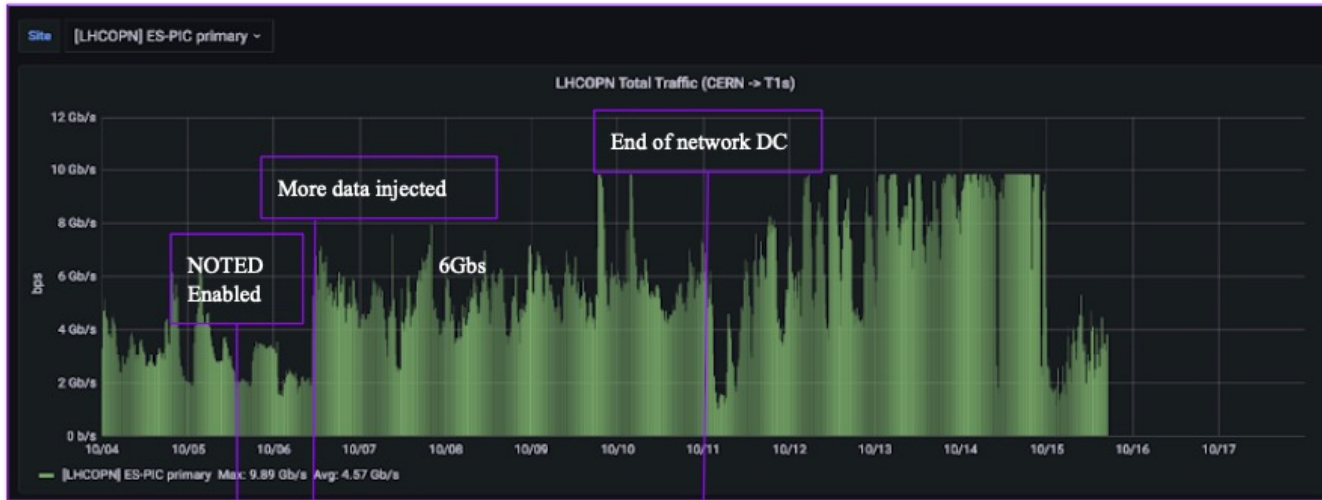
Tape Challenges in preparation for Run-3 were recommended by the RRB and the LHCC

Target rates differ for Data Taking (DT) and After Data Taking (A-DT). All relevant info [here](#). Post Mortem [here](#)



VO	Reads (DT) GB/s	Writes (DT) GB/s	Reads (A-DT) GB/s	Writes (A-DT) GB/s
ALICE	0	2.8	1.1	2.8
ATLAS	2.5	9.6	8.4	5.1
CMS	0.8	7.6	12.3	1.1
LHCb		4.38	3.38	
Total	2.5	24.78	25.18	8.3

CERN-PIC NOTED



NOTED is a Software Defined Network R&D project to share network traffic between different paths

Enabled during the data challenge between CERN and PIC

When the 6 Gbps LHCOPN link saturated, NOTED added the LHCONE link to complement it.

10Gbps target reached

WLCG data challenges

The 2023 data challenges will likely be in ~Q1 2024

The target rates (30% of the HL-LHC traffic) seem too aggressive given the shift of the HL-LHC schedule


- A 20% to 25% target seems more realistic and will not require a considerable hardware investment for the sites

30% to be reviewed

Specific targets for T2s will need to be discussed for different regions. **We should discuss what makes sense for ATCF**

We would like to have a joint challenge with some of the partners (e.g. DUNE, Belle-2). Under discussion

	LHC Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Flexible Scenario in 2027	Data Challenge target 2027 (Gbps)	Data Challenge target 2025 (Gbps)	Data Challenge target 2023 (Gbps)	Data Challenge target 2021 (Gbps)
T1						
CA-TRIUMF	200	400	100	60	30	10
DE-KIT	600	1200	300	180	90	30
ES-PIC	200	400	100	60	30	10
FR-CCIN2P3	570	1140	290	170	90	30
IT-INFN-CNAF	690	1380	350	210	100	30
KR-KISTI-GSDC	50	100	30	20	10	0
NDGF	140	280	70	40	20	10
NL-T1	180	360	90	50	30	10
NRC-KI-T1	120	240	60	40	20	10
UK-T1-RAL	610	1220	310	180	90	30
RU-JINR-T1	200	400	100	60	30	10
US-T1-BNL	450	900	230	140	70	20
US-FNAL-CMS	800	1600	400	240	120	40
(atlantic link)	1250	2500	630	380	190	60
Sum	4810	9620	2430	1450	730	240



WLCG data challenges

Data challenges are not just about bandwidth. They are an end-to-end test involving the full middleware stack. They offer the opportunity to commission new features at scale

The 2024 data challenges should demonstrate the use of tokens for bulk data transfer

Progress needs to be made in the next year:

- Support of tokens in storage (see right)
- Support of tokens in FTS
- Support of tokens in the experiment data management systems
- Strengthen the performance of IAM (token issuer)

Different level of maturity for the items above

Statistics by Tag	Pass	Fail	Pass / Fail / Skip
critical	369	39	
not-critical	8	26	
se-cern-eos	20	6	
se-cnaf-amnesiac-storm	24	2	
se-florida-xrootd	23	3	
se-florida-xrootd-redis	23	3	
se-fnal-dcache	26	0	
se-infn-t1-xfer-storm	24	2	
se-nebraska-xrootd	20	6	
se-nebraska-xrootd-redis	18	8	
se-prague-dcache	14	12	
se-prague-xrootd	24	2	
se-prometheus-dcache	26	0	
se-ral-test-xrootd	22	4	
se-ubonn-xrootd	24	2	
se-ucsd-xrootd	23	3	
se-ucsd-xrootd-redis	22	4	
se-wisconsin-xrootd	22	4	
se-wisconsin-xrootd-redis	22	4	

WLCG Workshop



1st WLCG workshop in person after the start of the pandemic. Jointly organised with the Rucio community workshop

A very good attendance (80 people on site and ~30 remote)

Discussion items: experiment inputs, DOMA, HPCs/Clouds, tokens, facilities (with a focus on energy). Collaboration Board meeting, discussing the review of the WLCG MoU annexes

WLCG partners and collaborators

Many communities collaborate with WLCG and share some of the same technologies and services

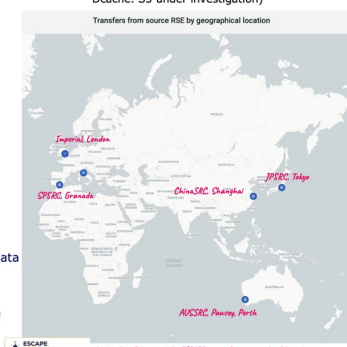
In some cases, we share the same resources

Summary and future timeline

- DUNE successfully utilizing resources at WLCG sites
- Looking forward to addressing challenges
 - fuller Rucio integration (see talks on Thu)
 - developing new workflows and workflow




Progress - SRC Rucio prototype



Not a production instance! Development and testing



Operations model for data
Rucio stress tests, test framework
Initial SRC efforts to increase
assignment
to SRC use cases.
based system!



hop - Nov 2022




Belle II Network


100G Global Ring
runned by SINET



LHCOPN Optical
infrastructure that can
be used without
jeopardizing resources



LHCONE L3 VPN
Connecting all the major
Data Centres



Fostering this collaboration is essential for the sustainability of the WLCG infrastructure



Conclusions

- Run-3 is ongoing and we have seen only the beginning of it
- In this presentation I focused on the medium to long term WLCG planning and the preparation for HL-LHC . That plan foresees gradual improvements and new features being demonstrated early in production
- The 2024 data challenges are an opportunity to continue assessing the progress at the facilities and the technologies. The ACTF is the right forum to organise the participation of Asian centers to the challenges
- The collaboration between WLCG and other experiments and sciences will be a benefit for the facilities and needs to be pursued