Developments at CERN & WLCG in the context of Data Lakes and Caches

Mario Lassnig
with content from many colleagues

HL-LHC Computing Roadmap 1/2

HEP Software Foundation Community Whitepaper

Bottom-up exercise	Identify the areas of work to address the HEP challenges of the 2020s

WLCG strategy toward HL-LHC Whitepaper

Top-down	High-level prioritization of the whitepaper, for the LHC needs

The LHCC review series of HL-LHC computing

Multistep process	Tracking the progress towards HL-LHC
May 2020	Review of ATLAS and CMS plans, Data Management (DOMA)
	Offline software, WLCG collaboration and infrastructure
November 2021	Update from ATLAS and CMS
	Common Software Activities (Generators, Simulation, Foundation, Analysis, DOMA)

HL-LHC Computing Roadmap 2/2

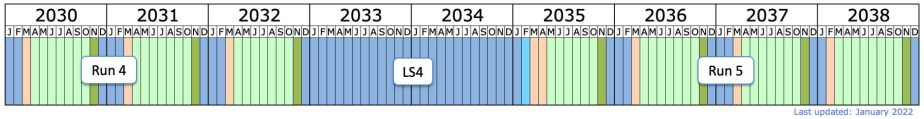

Gap between available and needed resources

Revised HL-LHC schedule

Key DOMA topics identified

- Token-based authentication & authorisation
- Software-Defined Networks
- Revisiting experiment dataflows

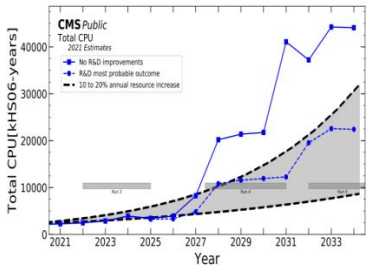

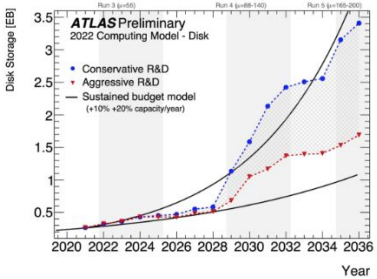

Leverage R&D projects to fill resource gaps

- ESCAPE / EOSC DataLake-as-a-Service
- Experiment-based R&Ds: Data Carousel, Virtual Placement, …
- Collaborations with commercial clouds
- Cross-experiment collaborations on infrastructure & computing, esp. with astronomy, neutrino, and photons

Data Challenge planning

Collection of DC#21 feedback available: https://indico.cern.ch/event/1026399/

Start planning next one (DC#24)

- Considering a 3x increase of required throughput
- Considering recent updates to HL-LHC timeline

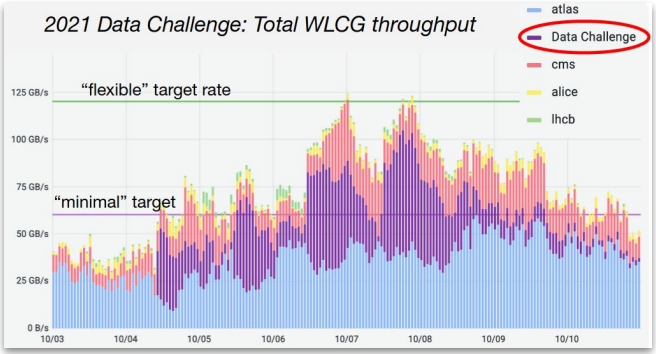

Networking will play a central role in HL-LHC as enabler for HEP computing

- Support the core functions of WLCG (data acquisition/archival/processing)
- Provide more flexibility to the computing models, allowing to optimise
- 4.8 Tbps of total network capacity needed for HL-LHC

Full stack packet marking

- Work being organized by the RNTWG Packet Marking subgroup
- Passing more elaborate metadata from Rucio to FTS to the network layer to improve our monitoring and understanding of flows

Network orchestration

- Sequence of limited-scope tests for network function virtualisation
- The Network Orchestration WG is planning to describe options and a possible model that will require iteration
- Rucio/FTS will need to participate in this longer-term activity

Software-Defined Networks: NOTED

Self-aware network load-balancing

- Automatically discovers large transfers
- Integrated with Rucio & FTS

Offload from LHCOPN to LHCONE

- Achieves additional bandwidth
- Can exploit unused links

Demonstration driven by ATLAS Rucio

- CERN, TRIUMF, KIT
- Preparing for SC'22

See next week's network meeting

- Full presentation

Software-Defined Networks: SENSE

Allow priority traffic for priority datasets

- Ongoing work, driven by CMS Rucio

Make Rucio capable of using SENSE

- To schedule transfers on the network
- Automatically configure VPNs
- Enforce a given path for set of transfers
- Implement QoS to prioritize at the DTN level

Fine-grain managed transfers allows

- Fine-grain monitoring
- Within a well identified network channel

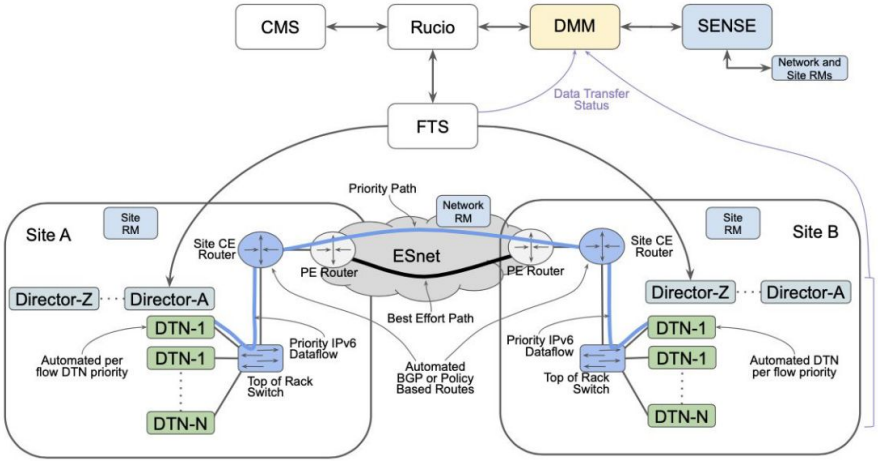

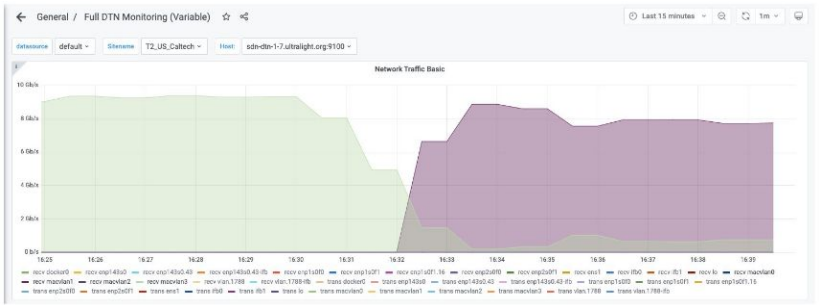

Caching examples: Virtual Placement, SoCal Cache

VP: Cache-aware brokering of ATLAS jobs

- Strategically deployed 6 small XCaches (12TB)
- Implemented via Rucio geo-awareness

SoCal: Federated XCache at the CMS infrastructure-level cache with 2.6 PB capacity

Very encouraging results

- Cache hit rate ~80% delivery of data from cache instead origin
- Reduce the traffic by factor 2+ during normal uses

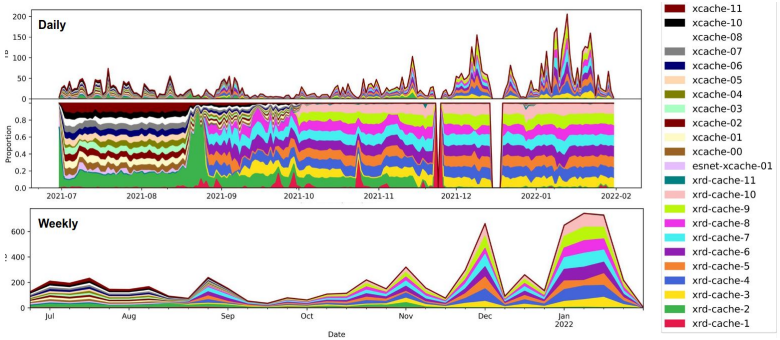

ESCAPE Data-Lake-as-a-Service 1/2

Prototype an infrastructure adapted to exabyte-scale **future needs** of large science projects

- Ensure sciences drive the development of EOSC
- Address FAIR data management principles

Data Lake as modular ecosystem of services and tools

- Shaped around the ESCAPE scientific communities
- Federated data management and access solution
- Heterogeneous resources

Hiding complexity and providing transparent access to data

- Layer for orchestration of resources as entry point for sciences
- Content delivery and caching layer

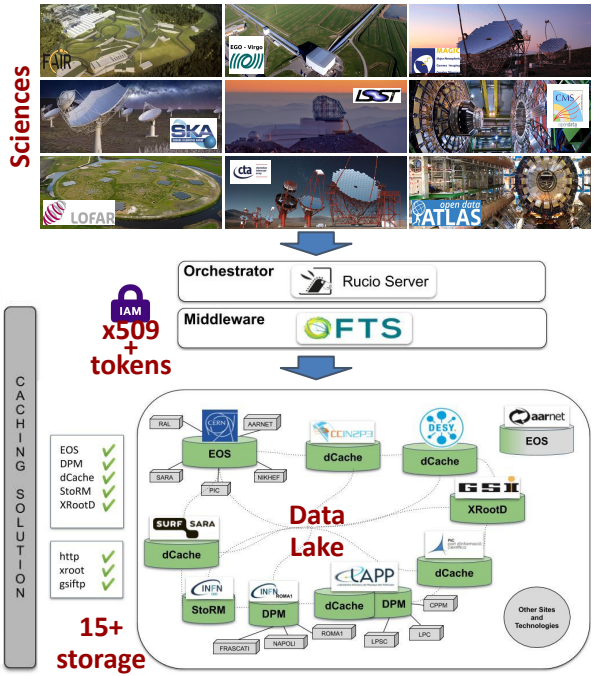

ESCAPE Data-Lake-as-a-Service 2/2

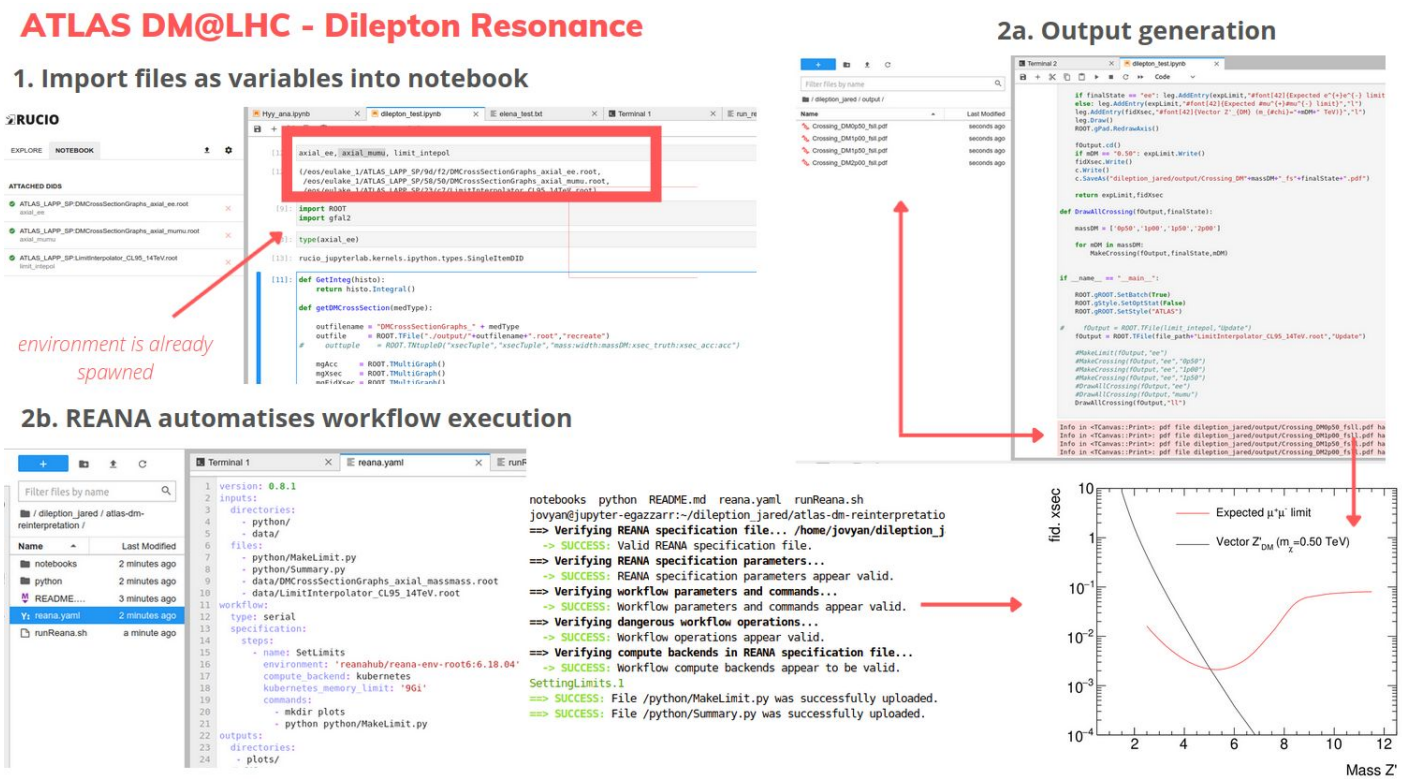

Backup

WLCG Data Challenges

Major WLCG objectives: 2 scenarios

Export of RAW data from CERN to the T1s

Data reprocessing

Two-yearly "steps" until HL-LHC

Improve through R&D programmes

2020 estimation total

4.8 Tbps of total network capacity

ATLAS & CMS 400 Gbps flat

ALICE & LHCb 100 Gbps flat

x2 to absorb expected bursts

x2 overprovisioning for operations

T1	%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
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	0	40	0	0	800	0	0	800	1600
(atlantic link)					1250	0	0	1250	2500
Sum	100	100	100	100	4000	400	410	4810	9620

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	800	1600	400	240	120	40
(atlantic link)	1250	2500	630	380	190	60
Sum	4810	9620	2430	1450	730	240

ESCAPE Data-Lake-as-a-Service