WLCG Update
Strategy Document & review

Ian Bird

LHCC Referee's meeting

CERN, 29th May 2018

Progress - 2018

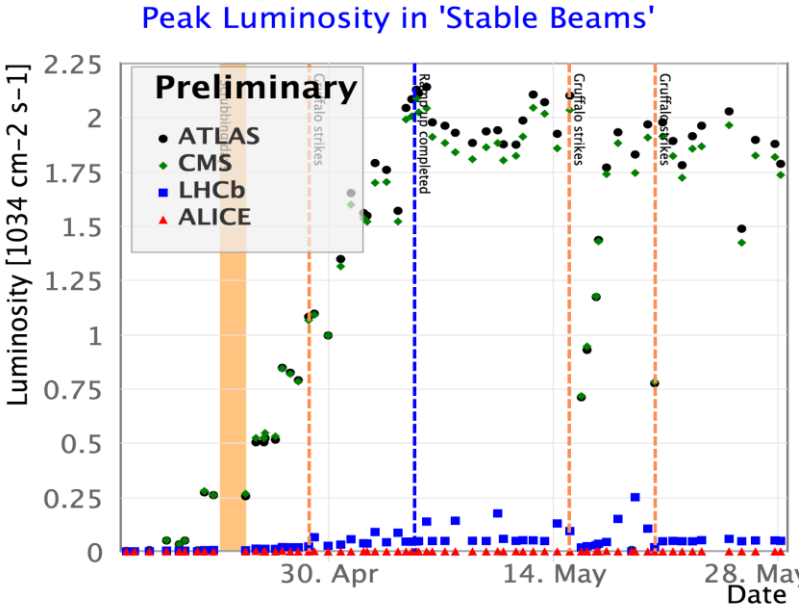

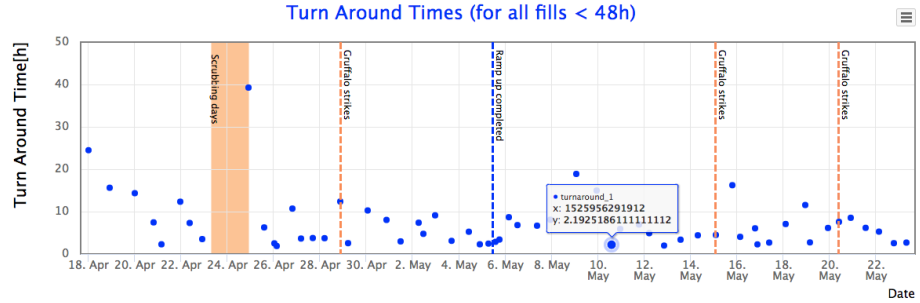

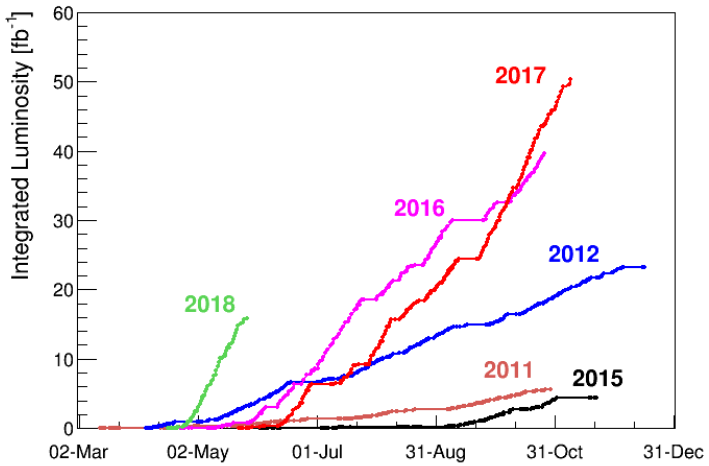

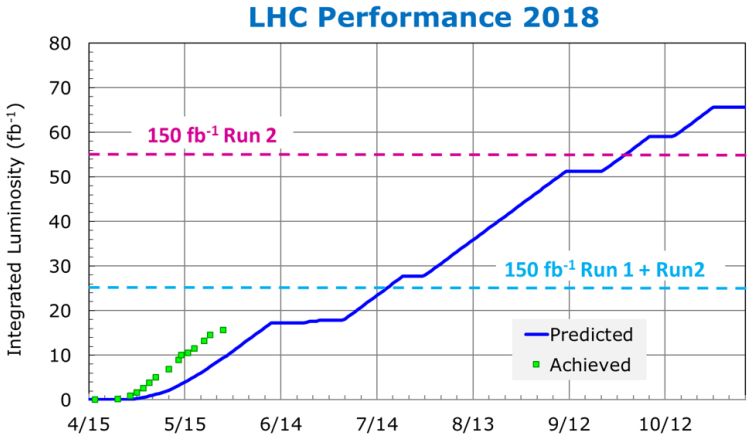

Data

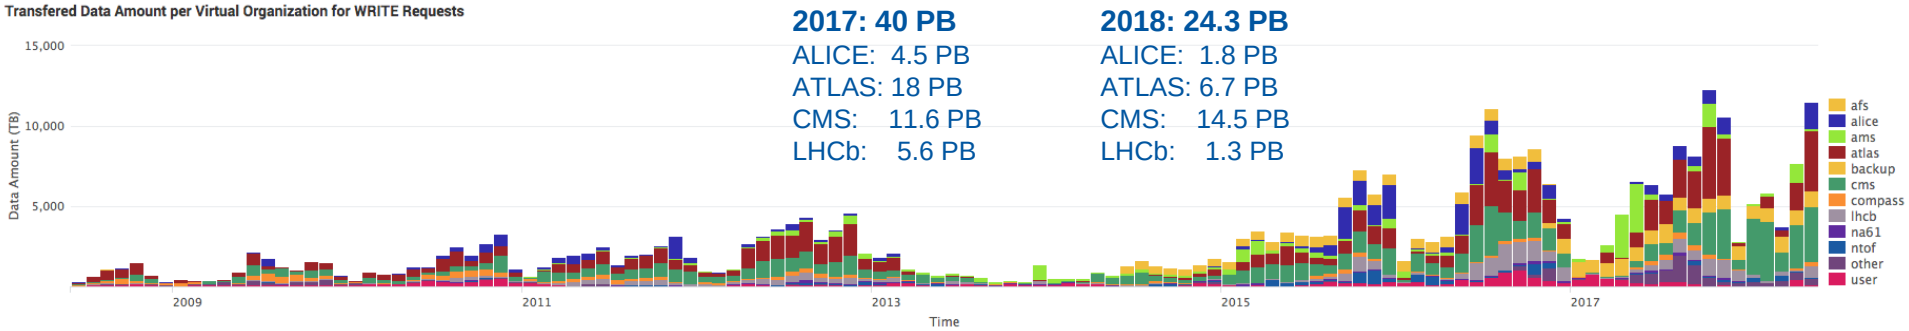

2017: 40 PB
ALICE: 4.5 PB
ATLAS: 18 PB
CMS: 11.6 PB
LHCb: 5.6 PB

2018: 24.3 PB
ALICE: 1.8 PB
ATLAS: 6.7 PB
CMS: 14.5 PB
LHCb: 1.3 PB

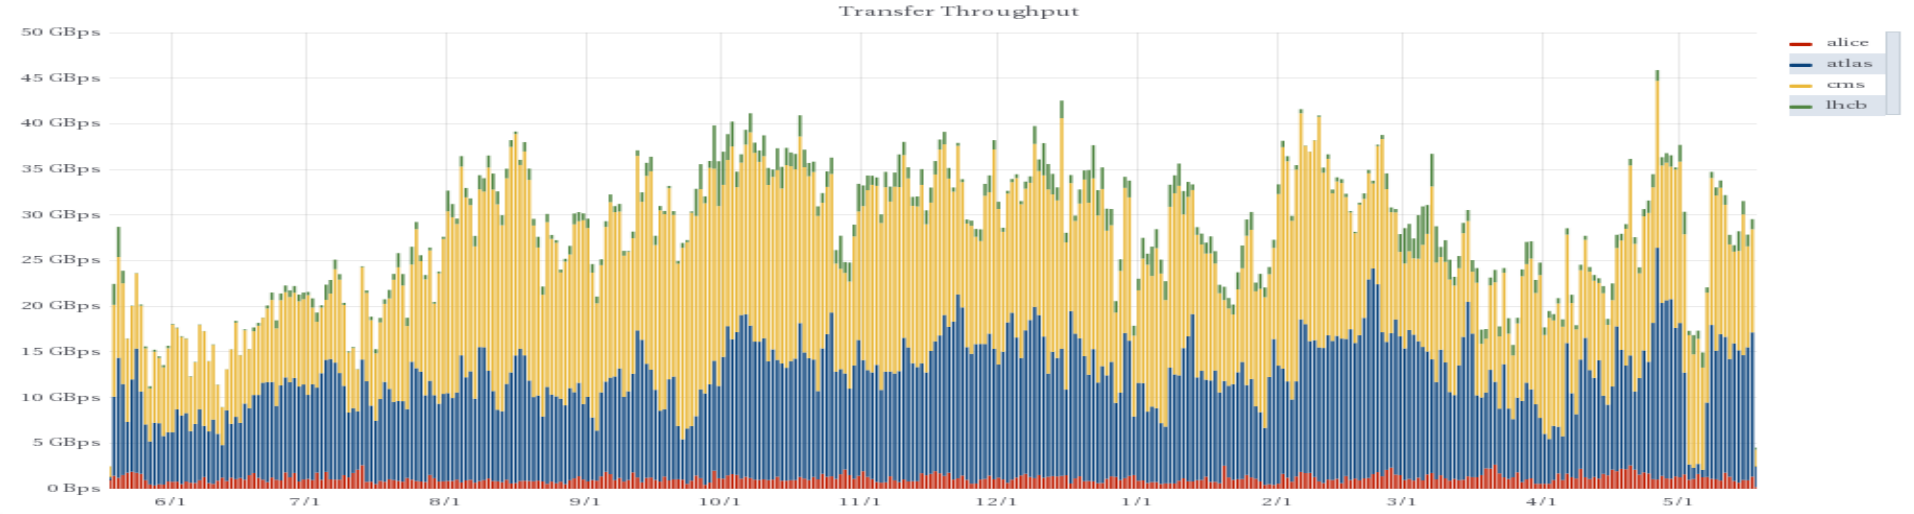

CPU Delivered

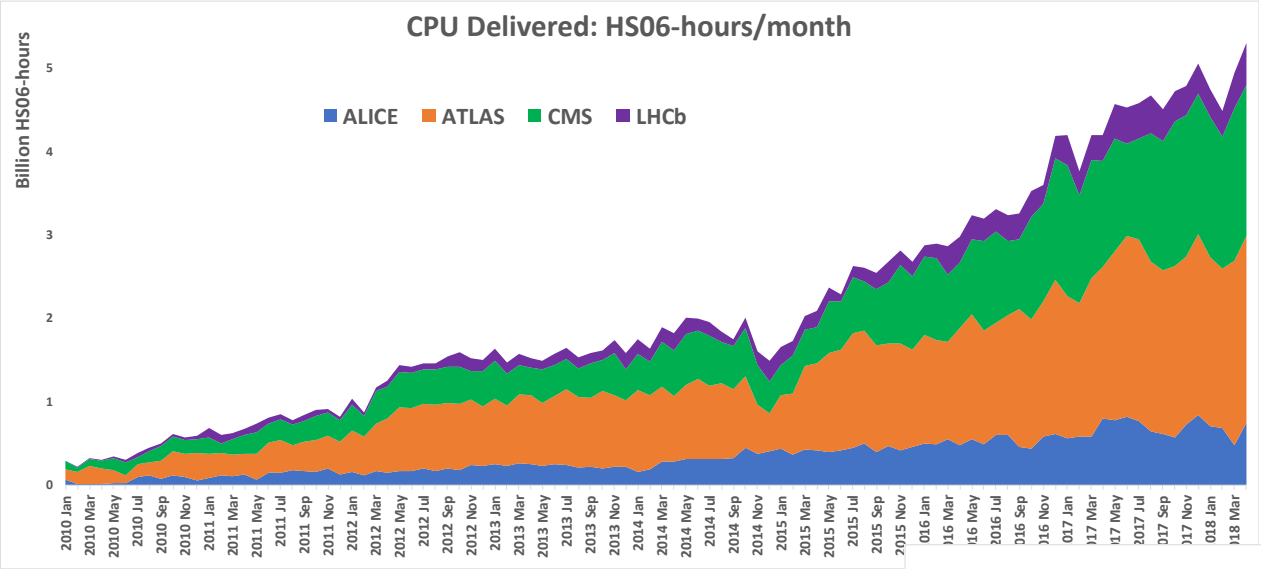

New peak: ~221 M HS06-days/month

~ 740 k cores continuous

(From sites that pledge)

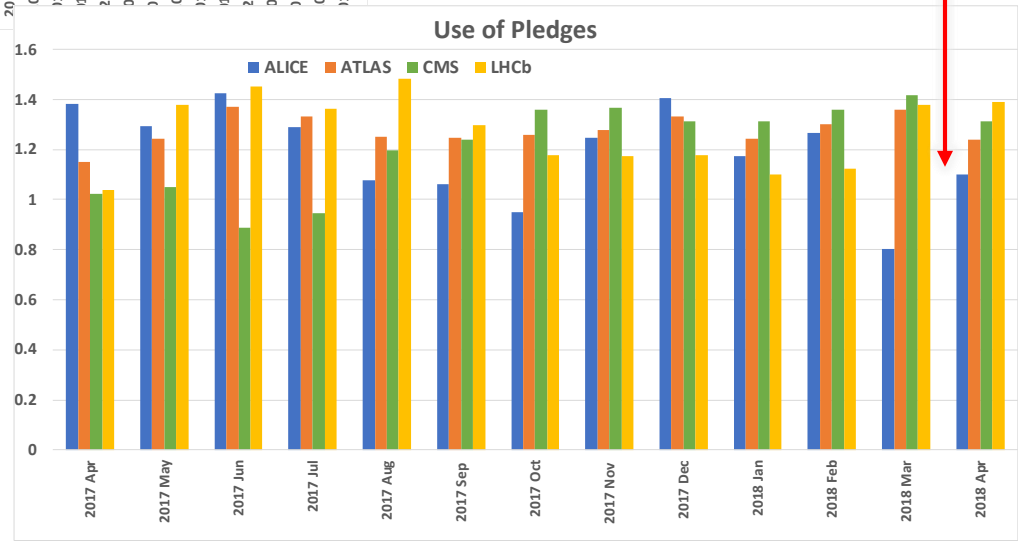

2018 Pledge situation

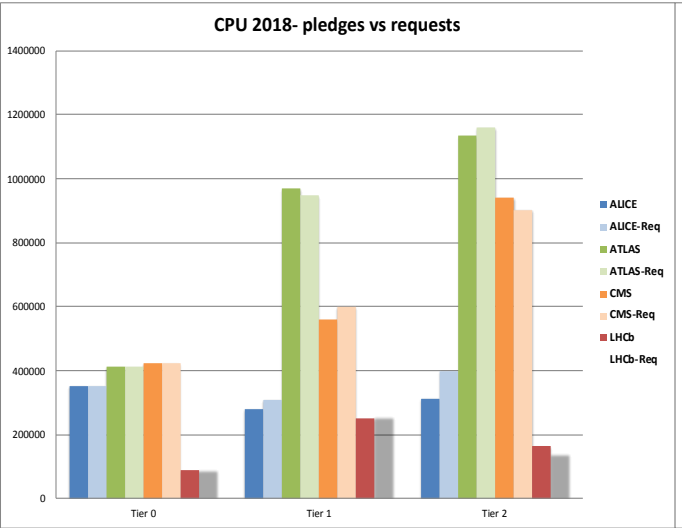

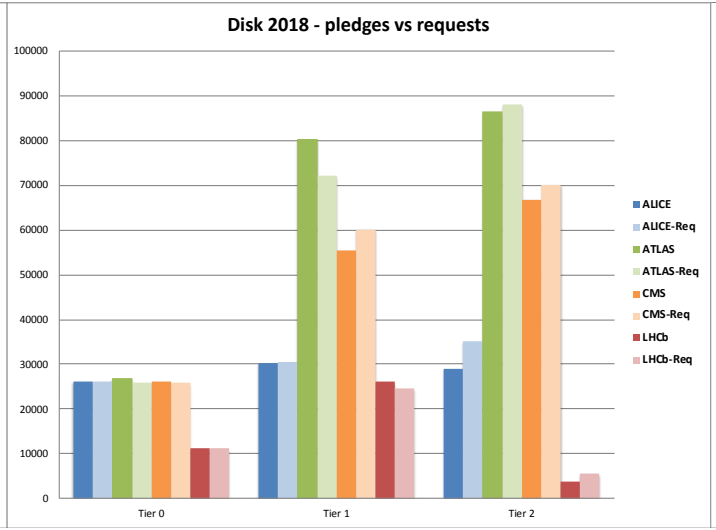

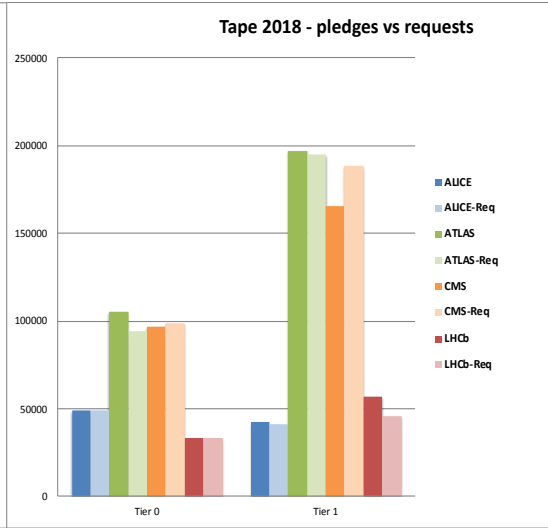

2018 pledges wrt requests:

As given in REBUS

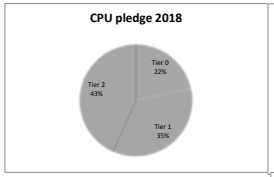

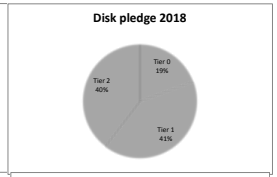

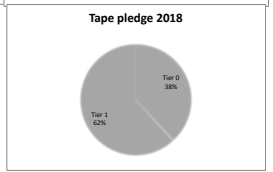

ESFRI Science Projects

HL-LHC	SKA
FAIR	CTA
KM3Net	JIVE-ERIC
ELT	EST
EURO-VO	EGO-VIRGO
(LSST)	(CERN,ESO)

ESCAPE

European Science Cluster of Astronomy & Particle physics
ESFRI research infrastructures

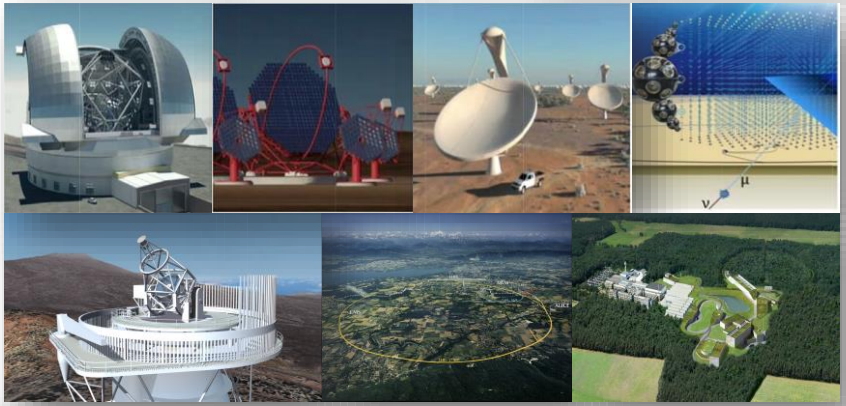

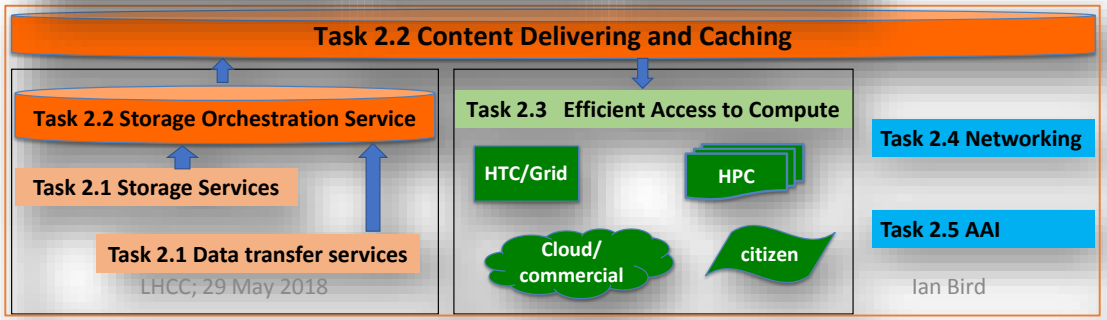

Goals:

Prototype an infrastructure for the EOSC that is adapted to the Exabyte-scale needs of the large ESFRI science projects.

Ensure that the science communities drive the development of the EOSC.

Has to address *FAIR* data management, long term preservation, open access, open science, and contribute to the EOSC catalogue of services.

Work Packages

WP2 – Data Infrastructure for Open Science
WP3 – Open-source scientific Software and Service Repository
WP4 – Connecting ESFRI projects to EOSC through VO framework
WP5 – ESFRI Science Analysis Platform

Data centres (funded in WP2)
CERN, INFN, DESY, GSI, Nikhef, SURFSara, RUG, CCIN2P3, PIC, LAPP, INAF

Run 3 running conditions – 1

- Following discussion with LHC operations
- Still many unknowns
  - E.g. experiment planned trigger rates are tbd
- Expected conditions:
  - 7 TeV per beam, gives small reduction in beam size
  - The main limitation is the heat load in the cryogenics
  - Expect BCMS filling scheme; 25ns
    - 2544/2556 bunches, $\beta^* = 27$cm
    - 1.3 x10^{11} protons/bunch
  - 2x10^{34} (could be a bit higher) is the limit due to the inner triplet cooling
    - This will not change in LS2
    - This is a pile up of ~60

BUT:

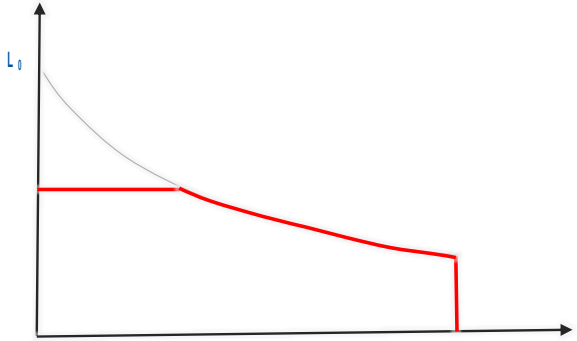

http://lhc-commissioning.web.cern.ch/lhc-commissioning/performance/Run-3-performance.htm

Parameter	BCMS	BCMS pushed a bit	Nominal - pushed	Comments
Energy [TeV]	7.0	7.0	7.0	
β* (1/2/5/8) [m]	0.3 / 10 / 0.3 / 3	0.3/ 10 / 0.3 / 3	0.3/ 10 / 0.3 / 3	Plus beta* levelling to 25 cm
Long-range separation [sigma] - assumed emittance	9.2 sigma - 2.5 um	9.2 sigma - 2.5 um	9.2 sigma - 2.5 um	
Initial Half X-angle (1/2/5/8) [µrad]	-160 / 120 / 160 / -150	-160 / 120 / 160 / -150	-205 / 120 / 205 / -150	Anti-levelled to 130 urad
Number of colliding bunches (1/5)	2592	2592	2748	BCMS - 240 bunches/injection from SPS
Bunch population	1.3e11	1.4e11	1.7e11*	* ruled out, initialy at least, by e-cloud heat load
Emittance into Stable Beams [µm]	2.5	2.6	3.0	
Bunch length [ns] - 4 sigma	1.1	1.1	1.1	
Virtual Luminosity (L0)	2.3e34	2.6e34	3.2e34	
Levelling time (hours)	2.0	3.8	7.9	
Luminosity per 12 hour fill (burn only)	0.65	0.7	0.8	
Luminosity lifetime (tauL) - end levelling	13 hours	14 hours	15 hours	Approx. - assuming burn only
Integrated/140 day year (fb-1)	65 - 70	70 - 75	85 - 90	NB Ballpark!

Summary – Run 3:

- Similar to 2018
- If the experiments luminosity level at a higher pile-up and for longer ➔
  - Potentially higher average pileup
  - Non-linear increase in CPU time
- Possibly less time between fills – more live time
- Overall the best estimate is 30% (50% conservatively) more resources needed than in 2018
  - But we have not seen 2018 yet
- For 2021: 1st year after LS2, could be only half-year live time but ramp up to optimal conditions rapidly
- Unknown:
  - Still need plans for experiment trigger rates
  - And plans for luminosity levelling

Resource evolution

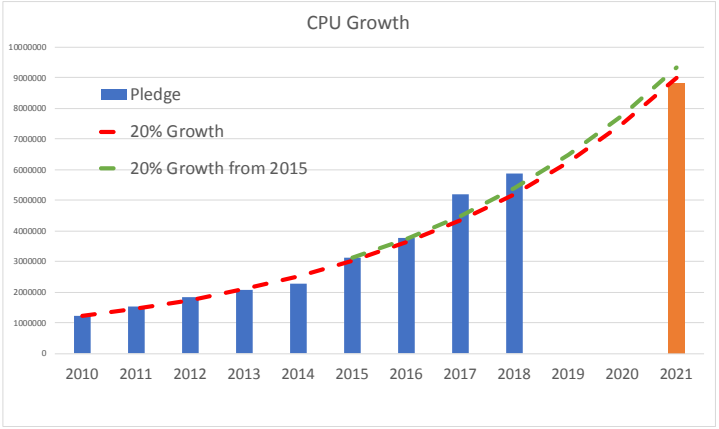

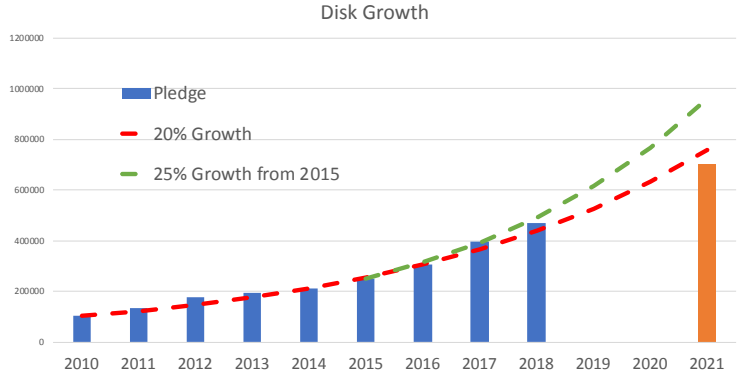

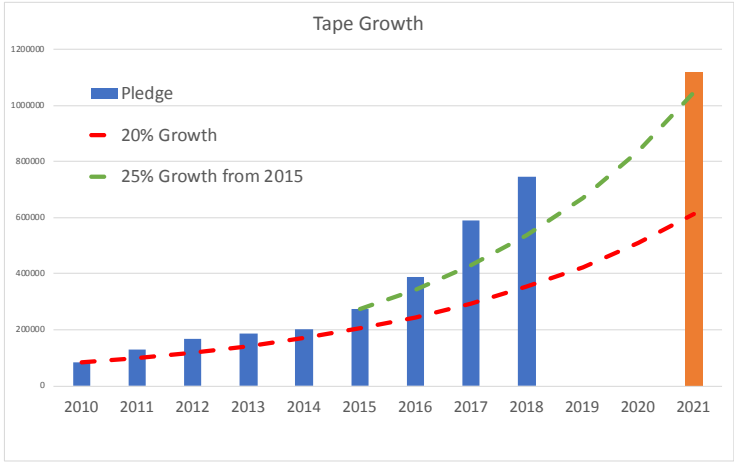

- 2010-2018 – pledges
- 2021 assume 1.5 x 2018

However …

- ALICE and LHCb are upgrading during LS2, so the expectations of their needs do not follow the assumptions in the previous slides:
  - LHCb:
    - luminosity and pileup increase by factor 5.
    - Major changes in computing model result in higher trigger rate and HLT output bandwidth.
    - LHCC milestone for computing model in Q3/2018, together with engineering TDR – currently under review
  - ALICE:
    - Factor 100 increase in readout rate (50 kHz)
    - Data volume increase mitigated by online reconstruction and raw data compression in new O2 facility
    - O2 TDR is approved; summary needs are:
    - Increases in 2021 wrt 2018: CPU: 48%, disk: 74%, tape 90%

Planning at CERN

- At CERN we are planning for procurements for 2021 as late as possible
- Budget constraints mean that we try to minimize purchases during 2019, 2020
- But we are pushed for additional resources during LS2
- We probably cannot satisfy all requests
- In addition logistics and infrastructure upgrades in the CC mean that we may lose some capacity during LS2

Summary

- Run 3 is very hard to plan for …
- 2018 is already at ~ nominal Run 3 conditions, but hints that LHC conditions will be pushed …
- Many unknowns from experiments – need some guidance on likely trigger rates, sustainable pile-up, etc.
  - Demands continue to increase – maximized trigger rates, parked data, much more HI data than foreseen, …
- 2021 is likely to be a short year, but 2022,23 could be very demanding on resources
  - We could reach limits of available budgets
  - This seems likely at CERN …
  - Strong hints remain that “constant budget” is the only realistic scenario
- What happens when we hit a resource limit?

Planning for HL-LHC
CWP & Strategy document

Naples workshop

- WLCG + HSF – follow up to CWP
  - ~200 attendees
- Key areas – various projects/wg's starting
  - Technology watch working group
  - Data Management R&D starting
  - Training group identified core skills as a prime first target
  - Common software libraries seriously discussed (VecCore, TrickTrack, Matriplex?)
  - Packaging group moving ahead with real tests
  - Software developers focus on performance and optimisation
  - Frameworks - take on the challenge of heterogeneity and organise workshop follow ups

Strategy document

Prioritize a program of work from the WLCG point of view:
A focus on HL-LHC, building on all of the background work provided in the CWP, and the experience of the past.

Themes

1. Software performance
2. Algorithmic improvements/changes
   - E.g. reco, fast MC, event generators
3. [illegible]s
4. [illegible]
5. Optimizing hardware costs

Demonstrate that we are in control of costs, while maximizing physics output

It defines an R&D program with rough timelines, organized in sections:

- The HL-LHC challenge, hardware trends and a cost model
- Computing Models
- Experiments Software
- System Performance and Efficiency
- Data and Processing Infrastructures
- Sustainability
- Data Preservation and Reuse

Timeline

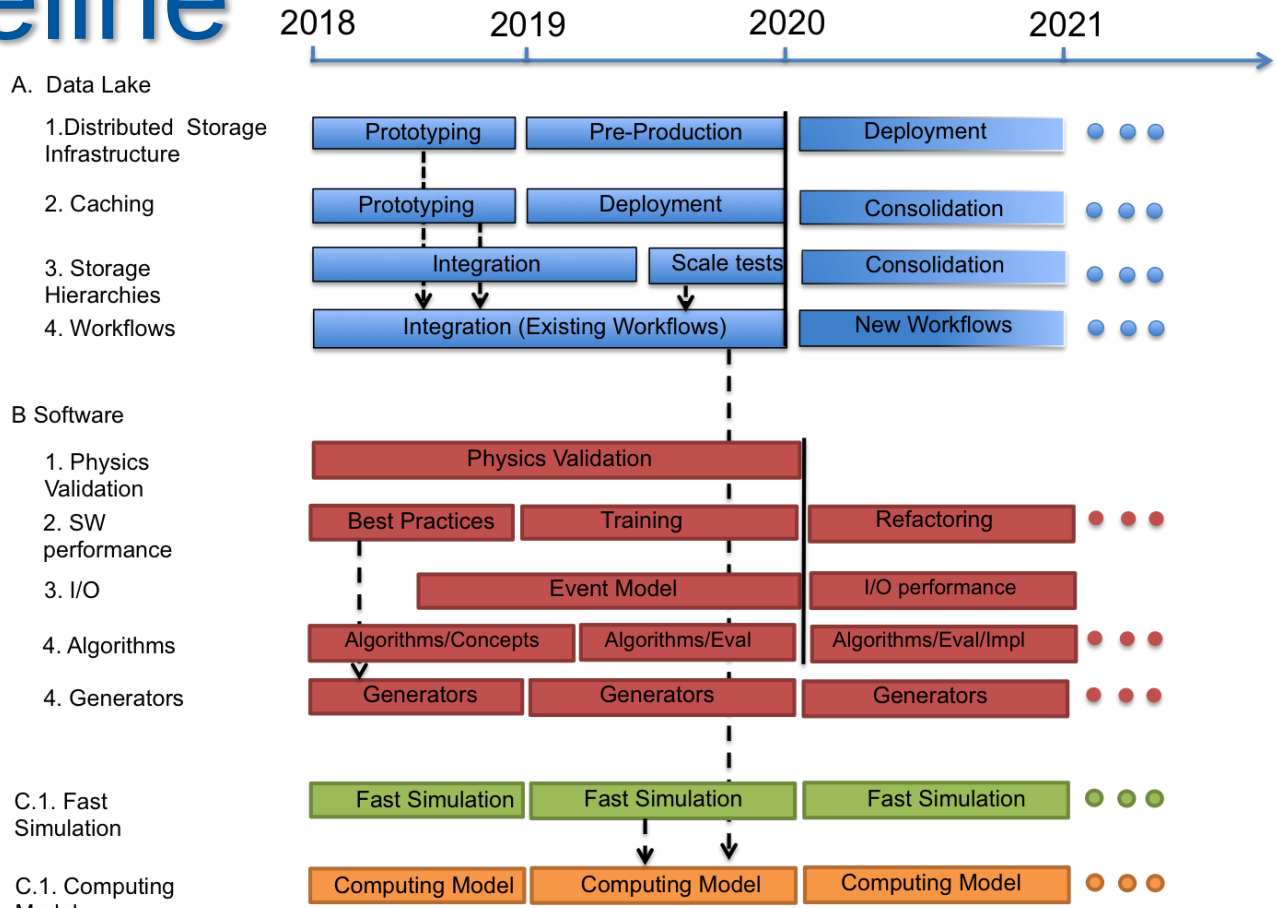
2018
2019
2020
2021
A. Data Lake
1.Distributed Storage Infrastructure
Prototyping
Pre-Production
Deployment
2. Caching
Prototyping
Deployment
Consolidation
3. Storage Hierarchies
Integration
Scale tests
Consolidation
4. Workflows
Integration (Existing Workflows)
New Workflows
B Software
1. Physics Validation
Physics Validation
2. SW performance
Best Practices
Training
Refactoring
3. I/O
Event Model
I/O performance
4. Algorithms
Algorithms/Concepts
Algorithms/Eval
Algorithms/Eval/Impl
4. Generators
Generators
Generators
Generators
C.1. Fast Simulation
Fast Simulation
Fast Simulation
Fast Simulation
C.1. Computing Model
Computing Model
Computing Model
Computing Model

Data and Compute Infrastructure

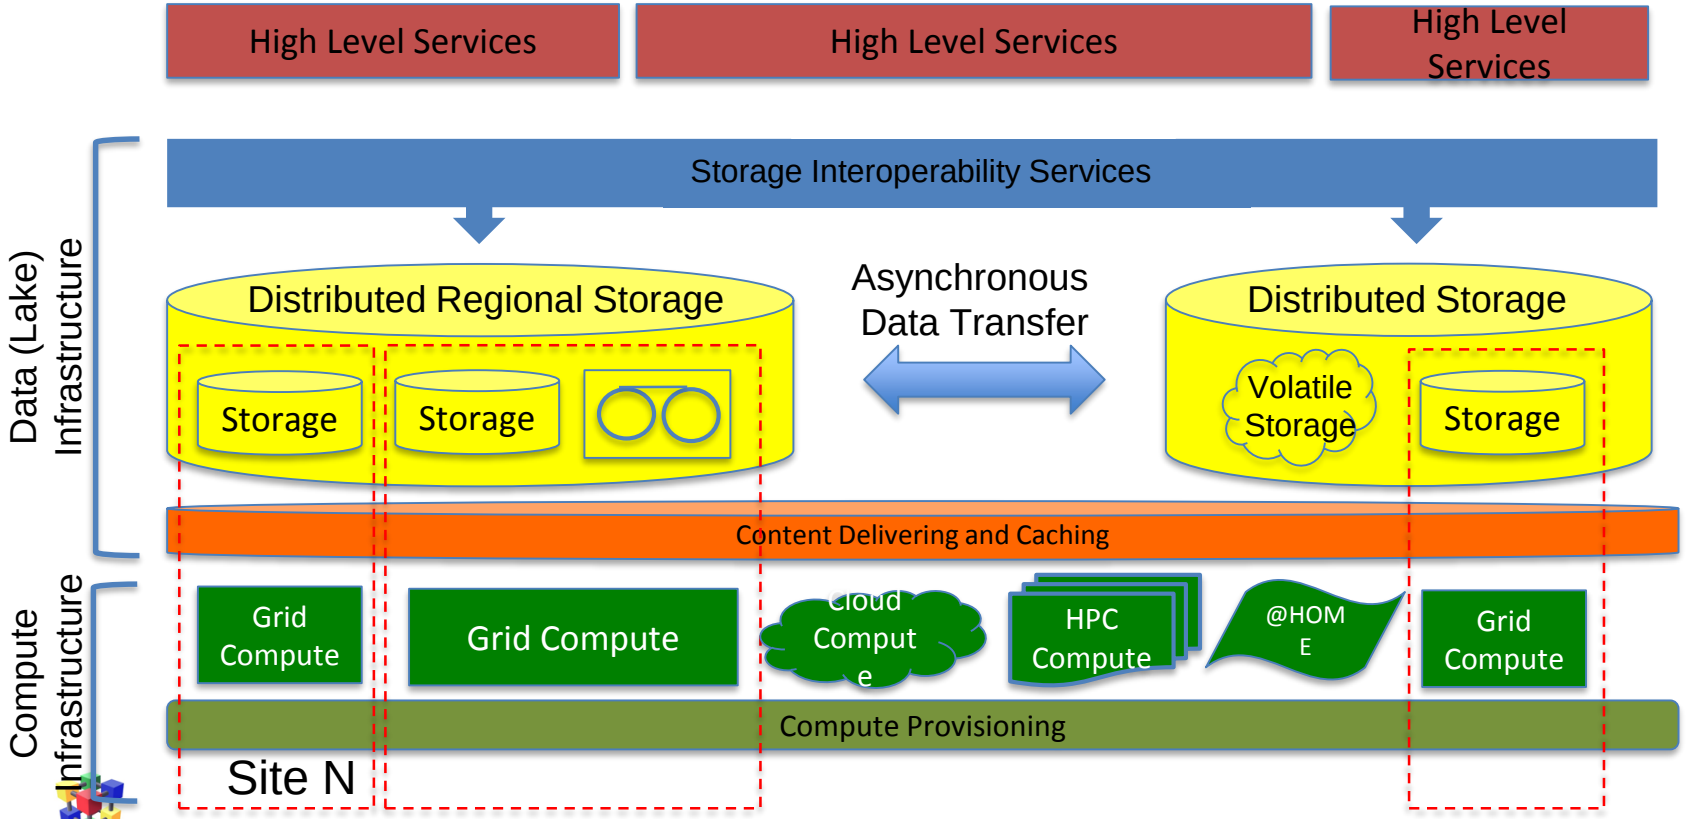

Priorities

- Setting up key projects to address the issues
  - NB. Several activities in hand in experiments, etc.
- We launched a WLCG DOMA (Data Organization, Management, Access) project to address the first
  - Create a forum to discuss ideas and present results
  - Track progress, review status, evolve the strategy, prioritize
  - Discussed in the WLCG/HSF workshop, agreed in WLCG management and Grid Deployment Board, kick off on June 4th and 5th
- We are discussing an equivalent initiative for Offline Software and the HSF is the natural umbrella for this
- Manpower is very short, particularly in software, also for the current core tasks
  - Need to leverage commonality at all costs, this is why the HSF plays a huge role here
  - Need to address the problem of recognition and career opportunity for people working on software related tasks

Review of HL-LHC Computing preparations

- As discussed in last LHCC
  - Should consider a review of the strategy
- We think a review will be extremely useful in a broader context:
- HL-LHC startup is ~2026/7 (optimistically)
  - TDR for computing in 2020 is perhaps not ideal – 6/7 years away from the need

➔ Have a serious review of the strategy in early 2019, a TDR slightly later (2022-3??), in concert with experiment computing TDR's (tbd).

➔ Use the review and (updated?) strategy to validate resourced projects with FA's

  ➔ But do not wait for this – start key activities now

A Review straw-man

- Scope: (tbd)
  - Comment on the appropriateness of the strategies to achieve an affordable and optimised computing model to maximize physics output for HL-LHC;
  - Are the ongoing work and proposed strategies realistic; and are there missing topics or opportunities?
  - Are the assumptions correct?
- It would also be useful to bring in all aspects of the problem of HL-LHC computing that need to all be contributing to the overall solution including aspects that don't traditionally regard themselves as "WLCG" like simulation and reconstruction. This would give a coherent overall picture.
- Needs 2-3 days
- High-level reviewers

➔ Needs some lead-time: likely timescale is thus early 2019

Review committee …

- Chair person: experienced but not directly in LHC programme
  - Amber Boehnlein (e.g.)
- Members: (this needs care)
  - Software expert – at least 1 senior person
  - Infrastructure expert
  - 2-3 high-level national representatives: involved in LHC computing and responsible for national structures
  - Rep from LHCC
  - Other?

Review structure

- Some introduction - overview of scale of the challenge, physics drivers, trigger rates, MC fraction etc. Experiment physics coordinator?

Each of the following should give prospects for improvement - performance factor, reduction factor, etc.

- Reconstruction
  - Improvements and prospects over coming 5 years - general view someone from common reconstruction activity
  - Experiment specific contributions and plans - ATLAS, CMS
- Simulation
  - Speed up of GEANT4 (vectorisation, parallelisation, re-engineering for performance, etc.) - Geant team
  - Fast simulation - Geant team
  - Explain plans for the evolution of GEANT
  - Full chain MC - ATLAS and CMS specific contributions
- Software performance and prospects in general - EP-SFT leading HSF activity, plus appropriate experts
  - portability (heterogenous architectures), I/O performance, EDM, etc
  - ROOT team should explain how they will help optimize I/O performance - what are plans?
  - Common activities (HSF) - parallelism, vectorisation, etc - how this will be managed
  - ATLAS and CMS - outlook for re-engineering core software?
- Prospects for reduction of data volumes - needs experiment specific contributions and plans
  - up-front/online processing (like LHCb plan to do in Run 3)
  - data formats - nanoAOD etc
  - use of virtual data
  - full chain MC
  - optimization of number of replicas - caching rather than storing, etc.
- Analysis evolution - who?
  - ROOT - what are future plans for analysis
  - Experiment outlook - what do they see as analysis needs - is ROOT sufficient?
  - Organized, local, cloud-based, etc. - Relative merits and costs
- Infrastructures - who? - Data-lake project + Rucio + experiments?
  - Data management ideas (data lake/DM project goals) - how much can be common
  - Workflow management - highly organized to allow use of tape vs disk
  - prospects for commonality - e.g. move of "~Rucio" into common layers, common workflows?
  - Hardware evolution outlook
  - Cost models?
  - Use of HPC (infrastructure level - software portability dealt with above)
- Event generators - someone very high-level (e.g. Ian Butterworth)
  - what are plans to re-engineer the code and support NLO, NNLO efficiently?
- Other topics?

Conclusions

- Very efficient and heavy use of WLCG during the winter stop, new peak usage reached
- Major incident at CNAF accommodated by other centres
- Resources and infrastructure in place for 2018
- Community White Paper published and WLCG Strategy document drafted –
  - R&D activities aimed at HL-LHC beginning