WLCG Status

LHCC Computing Referees meeting

I. Bird, S. Campana (CERN)

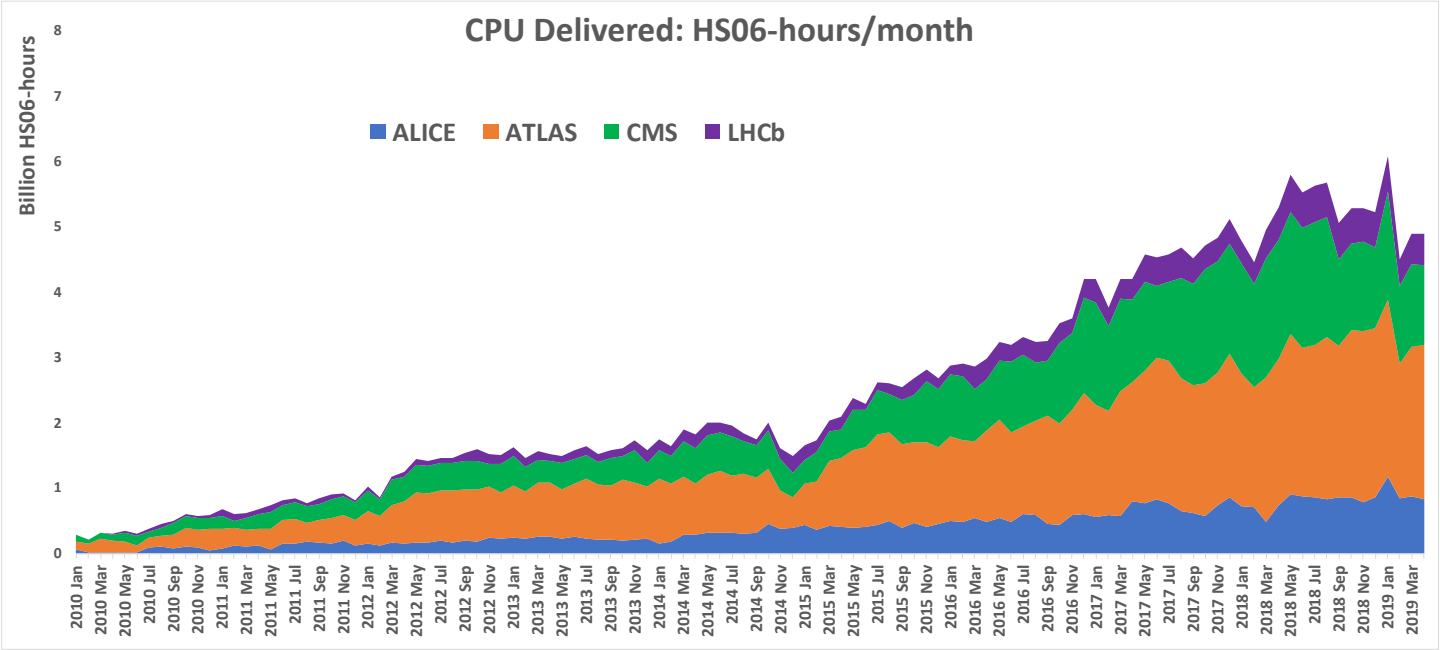

Pledged resources continue being fully used and some opportunistic ones continue being available. Not including most of HPCs/Clouds/HLTs

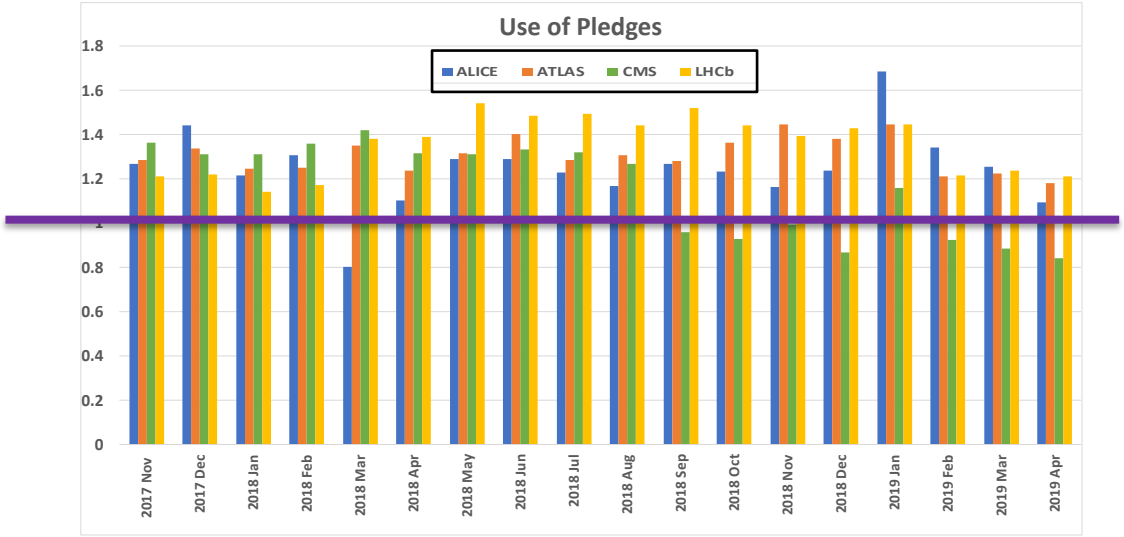

LHC Run-3

- There are still many unknowns on the schedule and the running parameters, probably more than one year ago
  - https://lhc-commissioning.web.cern.ch/lhc-commissioning/performance/Run-3-performance.htm
- The CRSG is asking firm numbers for Run-3 at the October RRB
- We would like to ask the LHCC to take this as an action

Time in Stable Beams

- Assume a physics efficiency 50%.
- 150 days with 50% of the time in Stable Beams - 6.5 million seconds

Virtual Luminosity (L0)	
Levelling time (hours)	

General comments:

- Usual caveats very much apply
- 2021 - recommissioning after LS2 for operation at 7 TeV - short year

ALL-2 The approved physics program for LHCb and the proposed program for ALICE will require significant increases starting already in 2021. The C-RSG recommends that the ATLAS and CMS collaborations review for the Fall 2019 scrutiny their tape storage requirements.

ALL-3 Recommends that a common set of planning assumptions (e.g., instantaneous and integrated luminosity) for Run 3 by adopted by the experiments. These should include estimates of upper and lower performance to make estimates of computing resources for the 2021 and future years.

ALL-6 The C-RSG requests that as part of resource assessments for Run 3 the experiments provide a strategy to address the case where running conditions for the experiment (e.g., pile-up, luminosity or other effects) are 20% larger than anticipated. For this mitigation strategy the experiments should assume no increase in pledged disk or CPU.

Computing road toward HL-LHC

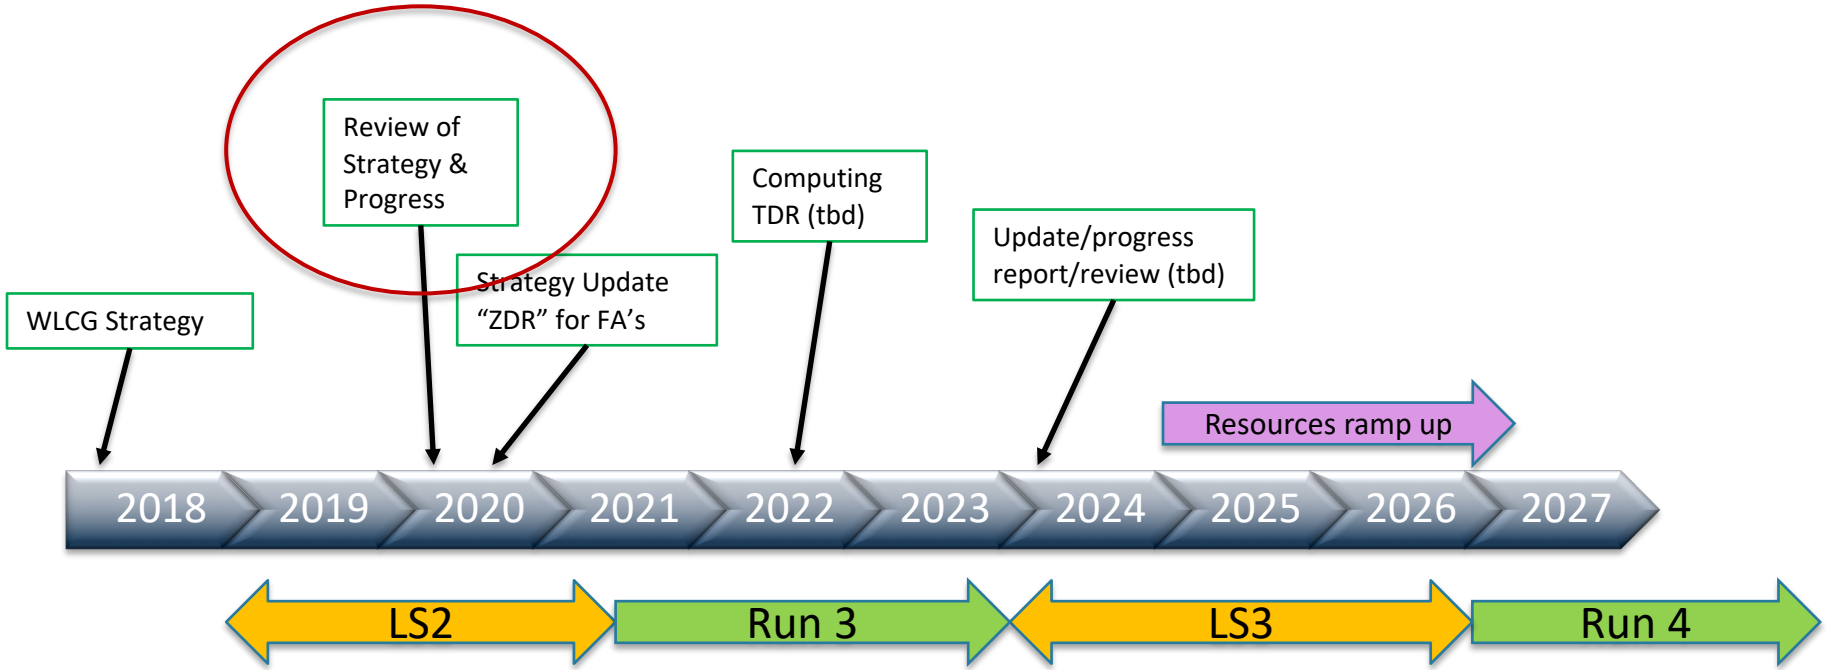

We propose a review of the strategy and the progress made in Spring 2020. We suggest to start the preparation soon (identify dates, possible reviewers, charge). We presented some suggestions in the past which we are happy to re-iterate

Computing road toward HL-LHC

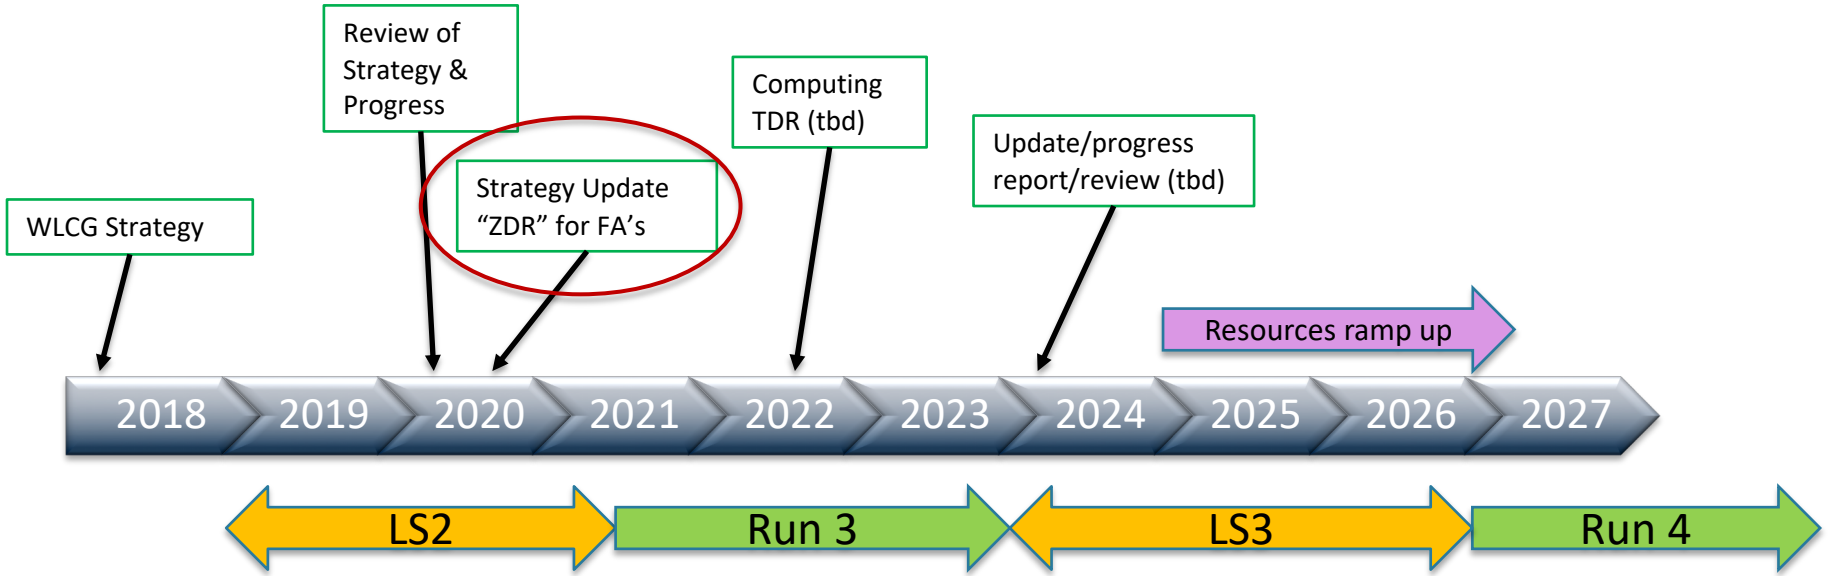

The Funding Agencies are asking for indications about HL-LHC resource needs and progress

This "ZDR" (or "CDR") can be done quickly after the review (Q3 2020). One ATLAS and one CMS document, one common/complementary WLCG one. Statements from LHCb and Alice

Computing road toward HL-LHC

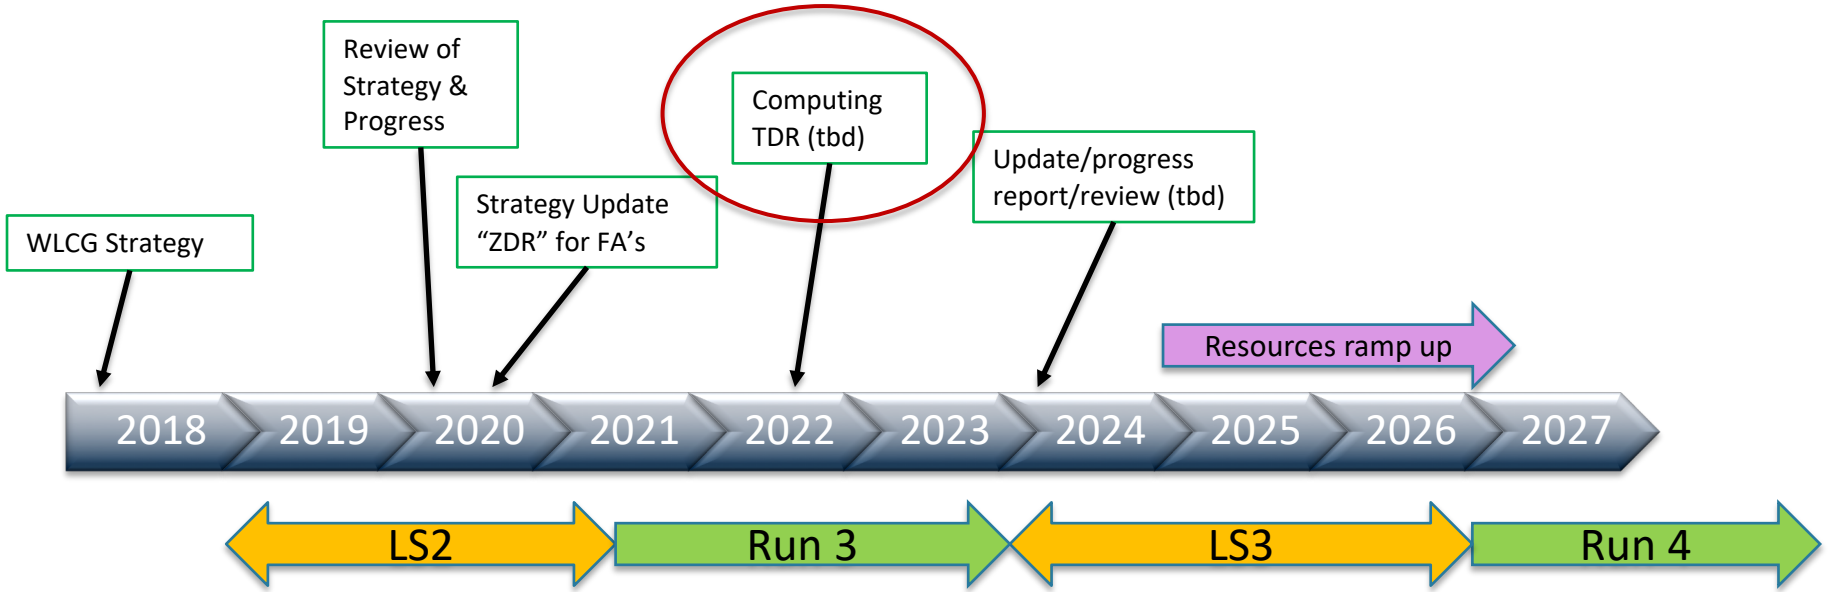

A Computing TDR will be meaningful not before middle 2022, with an update/progress/report/review after 18 months. Ready to ramp-up resources before Run-4

Schedule of TDR and what follows is subject to the LHC schedule

Concerns from hardware trends

- Cost of hardware decreasing exponentially vs time but not as steeply as before
- In general, trends driven by market (revenues) rather than technology

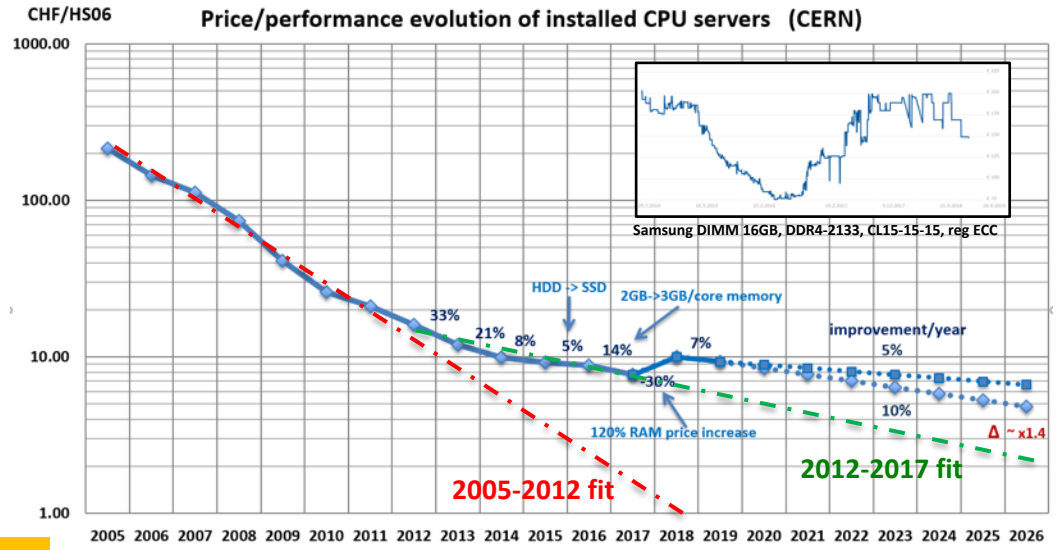

General loss of long term predictability

B. Panzer @ Scientific Computing Forum

Trends (not only cost) really change both with the country and the technical solutions adopted. 15% - 20% growth in many countries is over optimistic. A dedicated discussion at the next WLCG Overview Board:

“”””Round Table is to invite Funding Agencies to comment on the current view on budget outlooks and purchasing power. How do the agencies intend to fulfil the computing needs of LHC Run 3 and possibly beyond? What hardware do they intend to procure? How do they intend to address the potential case of the requirements surpassing the resources by 20%, 50%, 100%?”””

DOMA update

- Third Party Copy with xroot and HTTP protocols
  - Currently testing 33 endpoints, 6 storage technologies.
  - A dedicated Rucio instance manages transfers and monitoring
  - Improvements in the protocols themselves, documentation, deployment
  - Testing token-based auth as added value
- Finishing implementation and prototype phase. Reaching early deployment phase

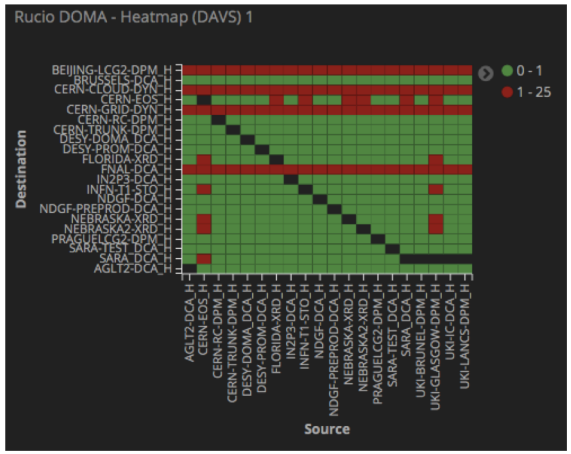

All documented in TPC Twiki

Extensive report at

HSF/WLCG/OSG workshop

DOMA update

- Data Access, Content Delivery, Caching, Latency Hiding: bridging CPUs and data
  - A huge amount of performance measurements, technologies evaluation, workflows studies
  - A strawman model proposed for the data analysis use case
  - Xcache emerges as the most promising caching and latency hiding technology
- Solutions are being prototyped by sites and experiment.

Extensive report a at HSF/WLCG/OSG workshop

All documented in ACCESS Twiki

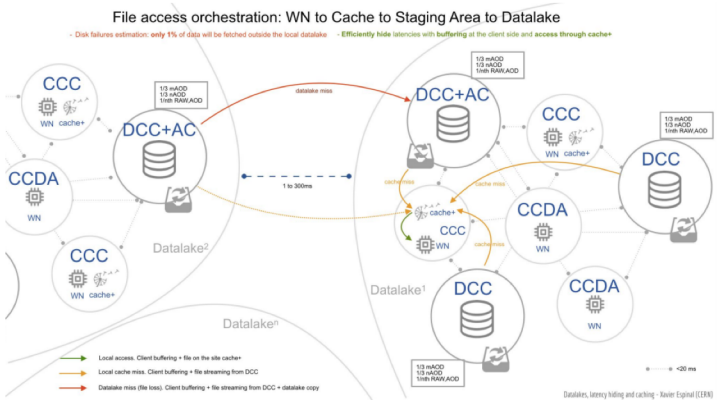

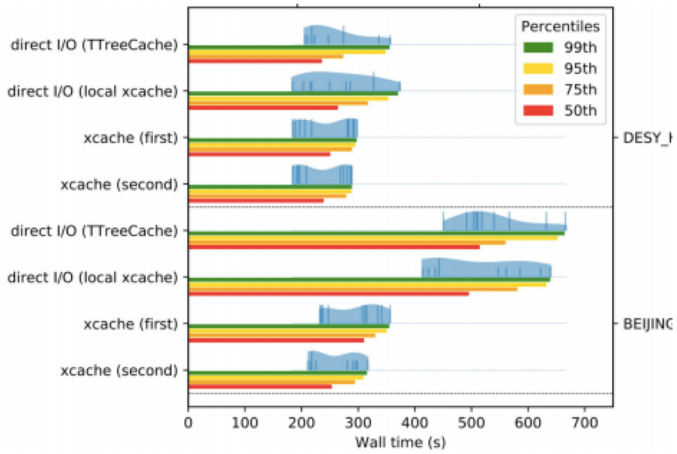

DOMA update

- Quality of Service
  - Describes how different classes of storage behave. It links needs and cost. Potentially a big hardware “saver”
  - Surveying use cases and storage capabilities at sites. Together with other WLCG operations related aspects
  - Understanding the integration of QoS classes with Rucio

Again, extensive report a at HSF/WLCG/OSG workshop

HPC cross experiment activities

- Technical meeting held in May, organized by the experiments
- The experiments would prefer continuing with HPCs as opportunistic resources. At the same time, they understand FAs might present them HPCs as pledged capacity and they need to prepare for this.
- Use cases for HPCs differ with the experiment
  - Focus on Full Simulation for ATLAS and LHCb, all workflows (predominant is Reco) for CMS. Impact of generators on compute needs differs as well..
- Two aspects being considered
  - Computing issues: workload and Data management on leadership class facilities
  - Software issues: portability to non X86, usage of accelerators

HPC elephants

From the discussion summary:

> We experiments think that while on our side we are committed to do our best, we should be facilitated by the FAs in these aspects. The best approach is to have HEP experts to be part of the definition process of both architectures and policies, having hence the HEP use case as a first-class citizen for HPCs. Sadly, this is not happening today, with HEP entering the game only when HPC machines have already been built and put into production.

Leveraging GPU capacity is particularly challenging

1. Evolve experimental frameworks to offload code to accelerators for Reco. HSF frameworks WG
2. Opportunities in generators (again, commonality x exp)
3. GPUs for Geant4 would be an opportunity but complicated
4. ML offers an opportunity (training). How much training do we need?

Any of this will require significant time (several years) and effort (many FTEs).
We need the support of Funding Agencies to address those aspects

Other WLCG news

Coordination with Worldwide LHC Computing Grid

- WLCG undergoing a transition in the organizational structure towards a Scientific Computing Infrastructure Steering Committee
- Approved middleware and interfaces (CREAM, ARC, HTCondorCE, dCache,EOS, SRM, etc)
- Tickets (GGUS), registry (GOCDB), and accounting (APEL) from EGI

- DUNE is now an **observing member** of the WLCG Management Board - announced at the HSF/OSG/WLCG Workshop at JLab this spring
- will participate in Global Deployment Board (GBD) for infrastructure planning
- DUNE has already started to participate at an operational level through GGUS tickets - report and aid in identification of site specific issues (UK and French sites currently)
- Hoping to have a GDB meeting in the US later this year with focus on DUNE-relevant areas
- **Technical contributions** made with development and evaluation of services (not focused on getting everything into production operations)
  - Edinburgh + FNAL: development work on RUCIO for Data Management
  - RAL-PPD + Manchester: exploring DIRAC capabilities for Workload Management (FNAL kx509 certs and file catalog filed w/ metadata from SAM)
  - RAL-PPD: Running small MC simulation on WLCG and RAL SE integration

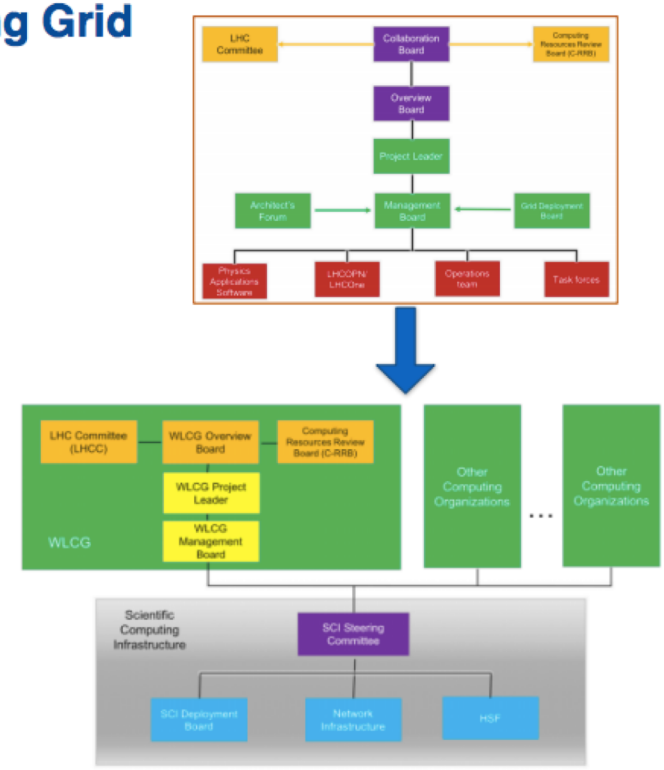

3

Fermilab DUNE

From the computing overview talk at the DUNE collaboration meeting

Highlights from the ESPP

One computing plenary talk. Parallel track on Computing & Instrumentation: one morning on computing, including a panel discussion

Many synergies identified between Computing and Instrumentation

- Difficulties in recognition and in offering career opportunities
- Importance of training on technical aspects
- Innovation needs strong support, maintaining a proper mixture between R&D and operations

Strong desire in collaboration across HEP and with big data experiment in the future

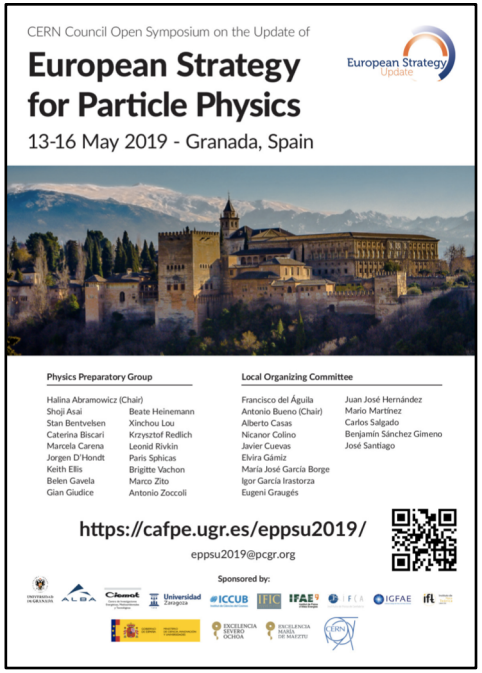

Highlights from the ESPP

Cost of both detector AND computing should be considered together when designing new projects

Data Management Plans (preservation, open access..) should be thought upfront in new projects

From the summary talk:

- Need to find way that continues to foster innovation while developing shared and common projects that have a long-life span and typically an Open Source Model.

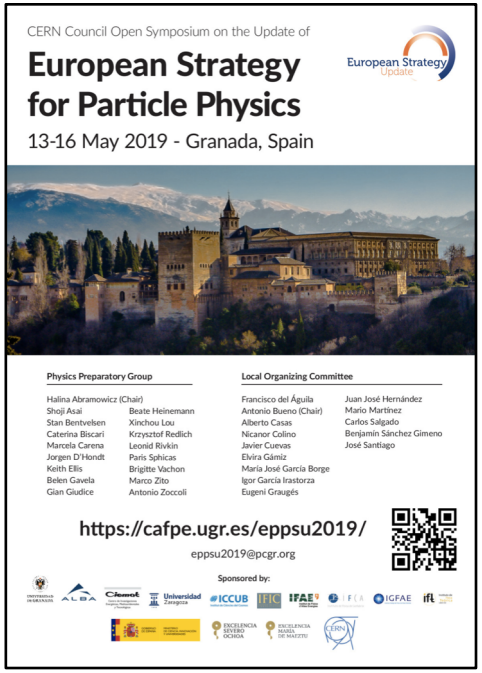

Key Messages

1. We need the help of the LHCC in defining the Run-3 running parameters

2. We propose a timeline for the HL-LHC computing strategy review and TDR. We are happy to contribute organizing the process

3. We continue working with the Funding Agencies in preparing for HL-LHC. Upfront investments in expert personnel is needed for the success of the project. A supportive message of the LHCC in this direction would be very helpful

4. A lot of progress in many WLCG areas (DOMA, AAI). Operations are smooth. WLCG infrastructure opening to other experiments

5. We feel the messages we passed for computing at ESPP in Granada were very well received. We should ensure they are well documented