

WLCG Outlook to Run 3 and HL-LHC

Ian Bird

Scientific Computing Forum

CERN, 2nd October 2019



2/10/2019

Ian Bird

1

Run 3 guidance – June 2019

For 2021 the baseline assumption for the proton run is as presented in the March 6 LMC:

- ❑ bunch intensities ramping linearly up from 0 to 1.4×10^{11} ppb over the year with limited availability of the injectors/LHC resulting in **only 20% machine efficiency**.
- ❑ Note that LHCb is expected to be leveled at 1×10^{33} for most of the time as the bunch intensity will not be sufficient to reach 2×10^{33} for the duration of a full fill.
- ❑ For contingency planning, the machine efficiency can be assumed to reach the normal value of 50%. This results in the following luminosity envelope:

| | Baseline | Upper limit |
|-----------|----------|-------------|
| ATLAS/CMS | 17/fb | 42/fb |
| LHCb | 3/fb | 7/fb |
| ALICE | 36/pb | 90/pb |

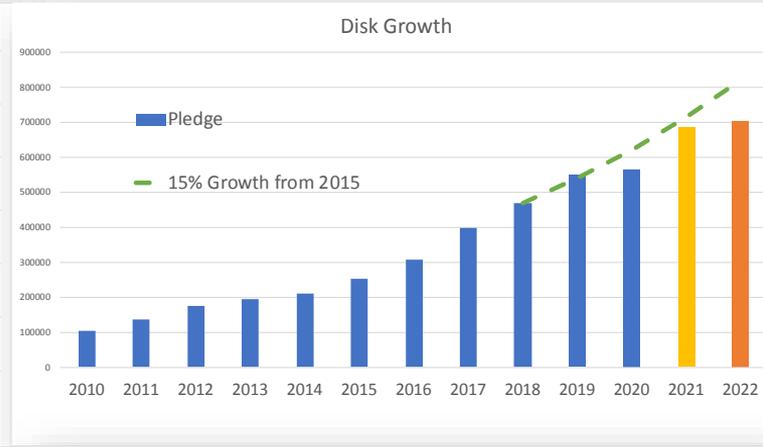
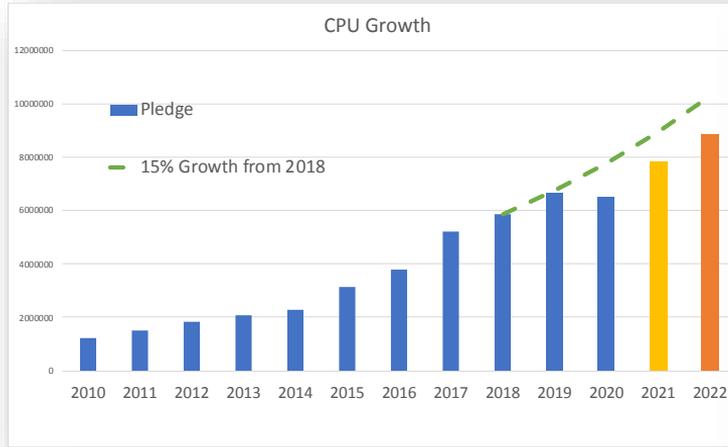
- ❑ We would like to emphasize that the upper limit is for **contingency planning only (i.e. raw data tape storage), not physics planning**.
- ❑ Furthermore the baseline numbers are likely to shift and should not be considered as luminosity targets yet. Updated numbers are likely to come towards the end of the year.
- ❑ For the PbPb period, it is assumed to be a full production year for now, which means more $>2/nb$ for ATLAS, ALICE and CMS.

Run 3 : best guess

- Given information from LHC team, following June LHCC request:
 - Agree baseline for 2021 is 17 fb^{-1}
 - With contingency plan (for tape) is up to 42 fb^{-1}
 - No additional (needed!) input on likely conditions was given
 - So agreement is to assume levelling is for 4 hours per fill
 - Assume a full HI run
- Likely situation as far as we understand:
 - 2021 is a very low data test run, resources same as 2018 for pp;
 - However, full HI run is likely – will need some level additional resources
 - 2022 is a full year with a resource level of 1.5×2018
 - Moderate (20%) growth rates for 2023 (and 2024)

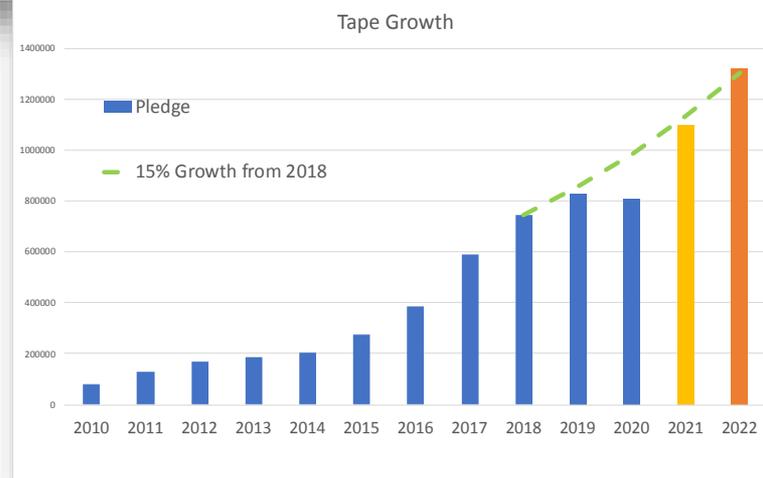
Resource evolution

- Actual installed
- 2021 request to C-RSG
- 2022 = 1.5*2018

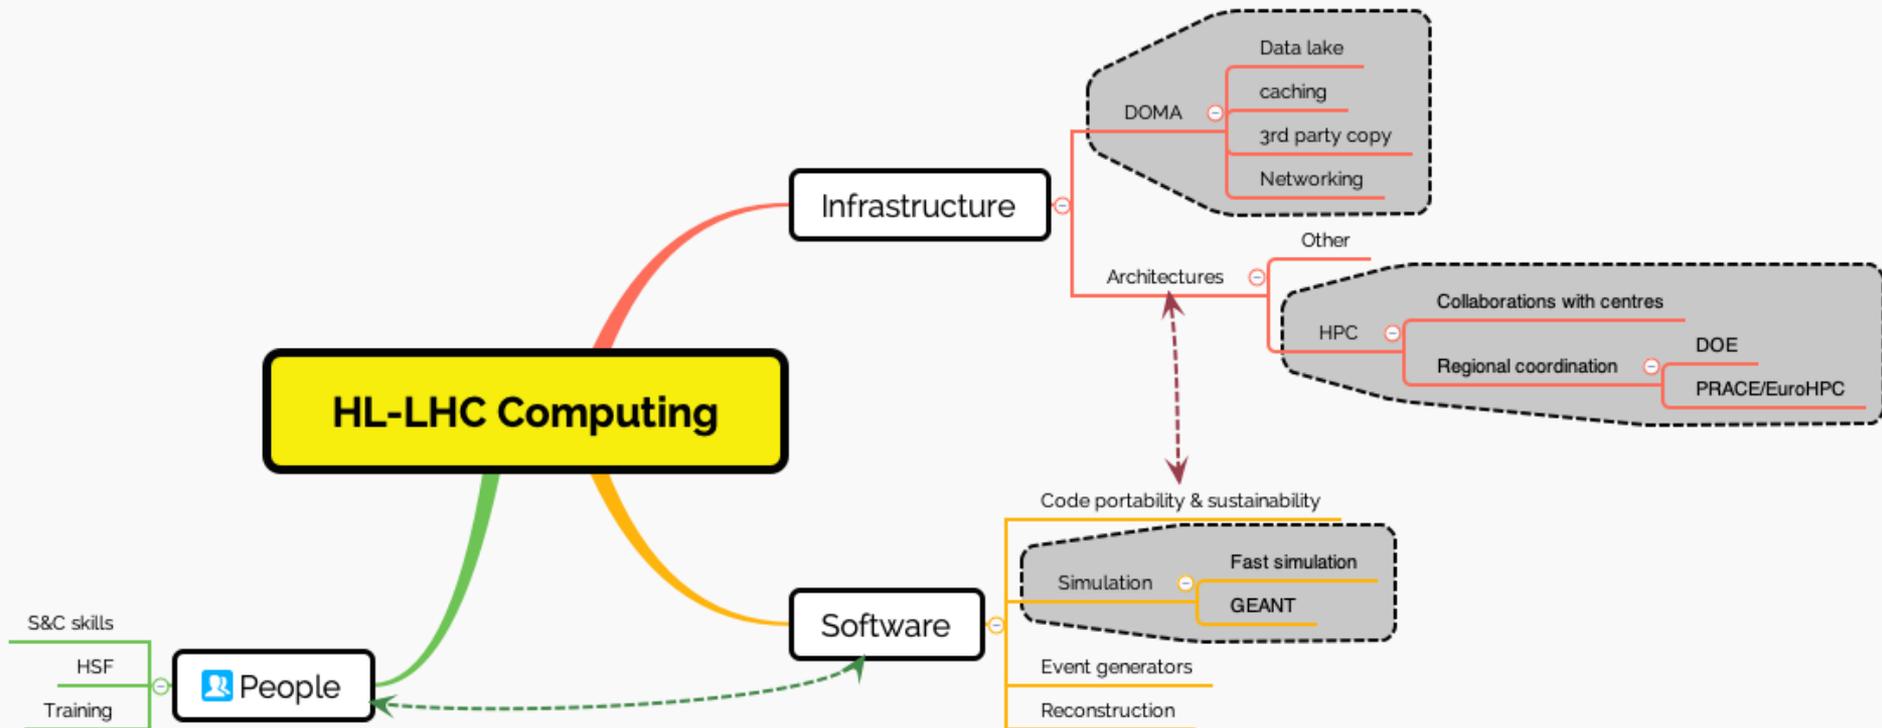


NB: Run 3 probably manageable overall, *but* constant budget growth until Run 4 is essential for HL-LHC

NB2: 15% maybe an optimistic assumption – many indications that flat budget is much worse (~ 10%)



Towards HL-LHC



DOMA in a nutshell

DOMA project

(Data Organization, Management, Access)

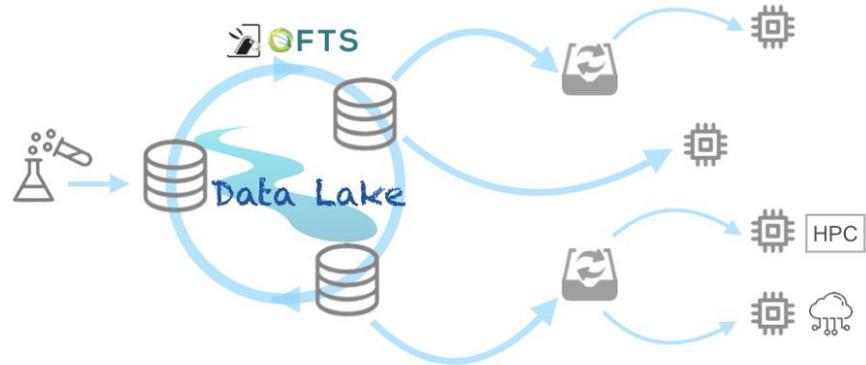
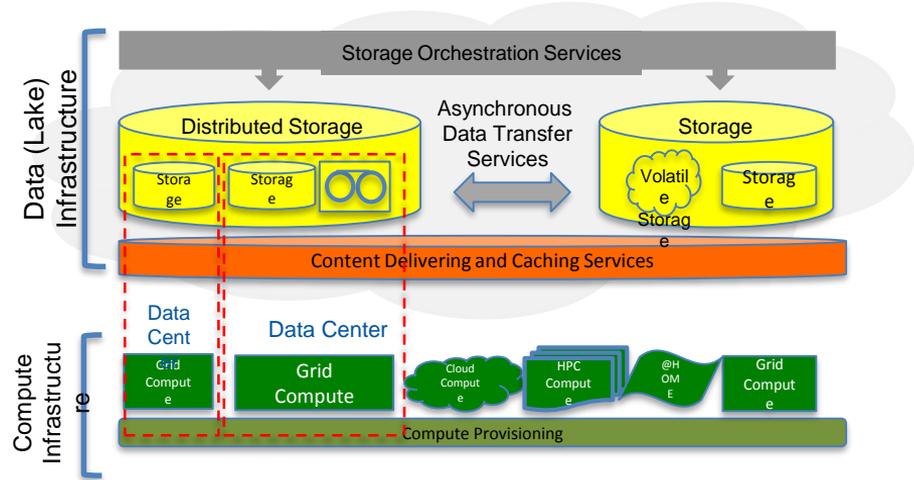
A set of R&D activities evaluating components and techniques to build a common HEP data cloud

Three Working Groups

- ACCESS for Content Delivery and Caching
- TPC for Third Party Copy
- QoS for storage Quality of Service

And many activities, reporting regularly

<https://twiki.cern.ch/twiki/bin/view/LCG/DomaActivities>



Data Storage Data Manager Data Mover Data Stream Data Cache Data Processing

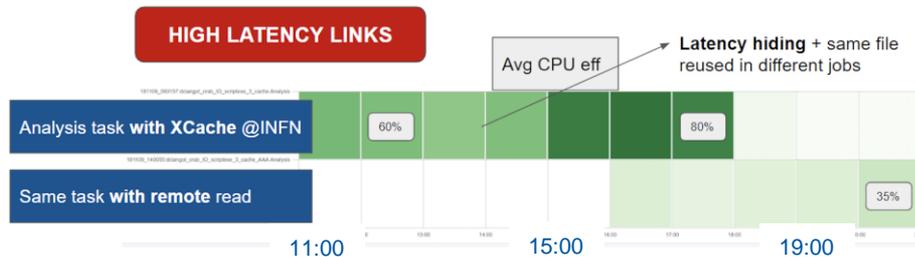
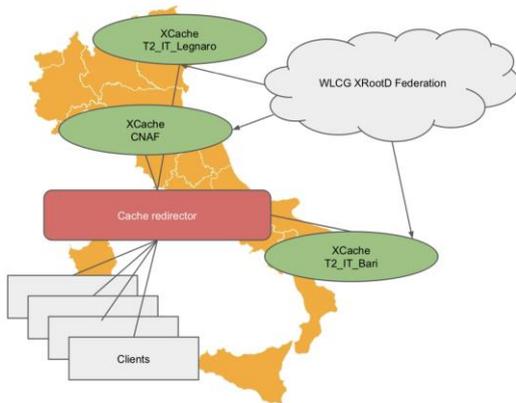


2/10/2019

From Simone Campana @ LHCC 10/09/19

ACCESS: caching layer prototype

A distributed caching system in INFN



From Simone Campana @ LHCC 10/09/19

TPC

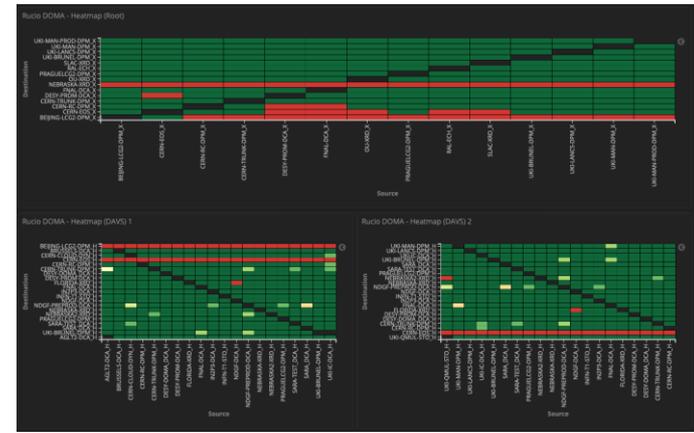
Goal: commission non-gridFTP protocols for asynchronous data transfer (Third Party Copy)

- Phase-2 (deadline June 2019): all sites providing > 3PB of storage to WLCG should provide a non-gridFTP endpoint in production

Functional and Stress testing



Capable to fill available bandwidth



Point-point functional testing

- Phase-3 (Dec 2019): all sites to have a non-gridFTP endpoint

NB: some features needed for TPC are available only in recent versions of storage

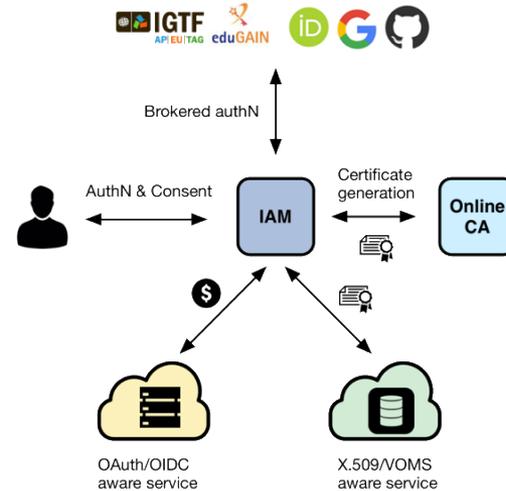
TPC and AAI

WLCG is planning to evolve AAI toward token based Auth/AuthZ and Federated Identities

The WLCG task force is finalizing the token profile as last item

While this has a much broader scope than DOMA, TPC offers a well confined use case to start with

Rucio is integrating tokens. Storage is preparing to manage them.



From Simone Campana @ LHCC 10/09/19

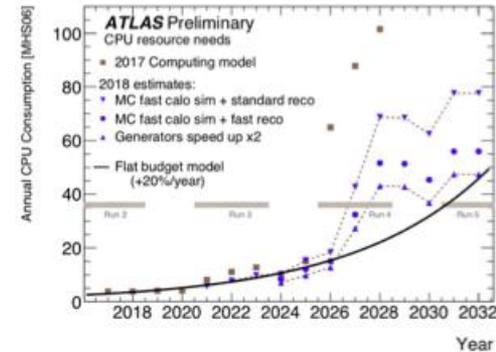
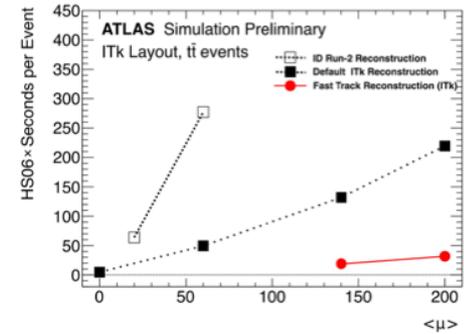
Software topics

□ Several active HSF working groups

- Event generators
 - Several workshops and meetings
- Reconstruction and software triggers
 - Common topics: GPUs, real time analysis, links to other communities
- Data Analysis working group
 - From DOMA to final analysis
 - Future analysis models, role of ML, etc.
- Software frameworks
 - Just set up, conveners nominated

□ Lots of work in experiments on software portability and performance

- Use of HPC
- Lots of work on tuning simulation; fast simulation (and where it is appropriate)
- Performance and portability:
 - Adaptation of frameworks to accommodate heterogenous code (CPU+accelerators)
 - Portability libraries: Kokkos, Alpaka, SYCL, etc
 - Can there be one codebase for all architectures?



Simulation

- ❑ Is a major cost driver (~50% total computing cost)
- ❑ Long term supportability/portability/performance is essential
 - Lot of effort in the world on portability to new architectures
 - Not as one-off to e.g. GPU-version-x
 - but we need restructuring of GEANT4 *framework & data structures* to enable efficient and ongoing portability and tuning/optimization & offloading key parts to accelerators/co-processors
 - Need a major effort on simulation for the future to tie in all of these R&D efforts
 - This is going to be a many-year effort
 - Is it GEANT4 or is it a new framework?
- ❑ This is where we really need to invest effort in the future

WLCG Risk Register

| ID | Risk | Risk | | | Owner | Effect | Mitigation |
|-------------------|--|------------|--------|----------|------------------|--|---|
| | | Likelihood | Impact | Severity | | | |
| Funding | | | | | | | |
| 1 | CERN Tier 0 budget insufficient | 3 | 4 | 12 | CERN | Cannot fulfill CERN commitments; knock-on effect to other Funding Agencies | Reduce scope of Tier 0 responsibility |
| 2 | Flexibility of CERN project budget lost | 2 | 4 | 8 | CERN | Cannot adapt spending to often changing LHC parameters and requirements | Work with RPC to retain project flexibility |
| 3 | No PCC for Run 4 | 3 | 4 | 12 | CERN | Unable to deliver Tier 0 commitment; Lack of flexibility between online/offline at CERN | Significantly reduce scope of Tier 0 responsibility |
| 4 | WLCG Funding Agencies lack funding | 4 | 3 | 12 | Funding Agencies | Cannot provide required resources to experiments; | Prioritise physics programme of the experiments, via the LHCC and experiment managements |
| 5 | Physics request for resources not affordable | 3 | 3 | 9 | Experiments | Not all physics programmes can be done, or will be delayed | |
| Technology | | | | | | | |
| 6 | Evolution of current technology insufficient | 4 | 4 | 16 | WLCG | Unpredictability of costs; Unable to provide required resources with constrained budgets | Ensure software as efficient as possible; Prioritise physics programme of the experiments Ensure flexibility in software and hardware - adaptable software, optimize between CPU, storage, networks |
| 7 | Offered resources not matched to requirements (e.g. HPC, GPS, etc) | 3 | 3 | 9 | WLCG | Inefficient usage | Ensure software as efficient & portable as possible; Prioritise physics programme of the experiments |
| 8 | Resources fragmented by types (HPC, GPU etc | 3 | 3 | 9 | WLCG | Inefficient usage; some resources may only support certain workflows; loss of general flexibility | Prioritise physics programme of the experiments |
| 9 | Tape market shrinks/fails | 2 | 4 | 8 | WLCG | Major rethink of value of data vs re-run of experiment, Need new data models; Drastic reduction in overall storage capacity | Content delivery models with relatively few data repositories; Rethink of value of multiple data copies; Assess alternative archiving solutions |
| 10 | HDD technology changes - expensive, hard to use | 2 | 4 | 8 | WLCG | Storage costs become too high | |
| 11 | Inadequate WAN; or competition with other sciences, or complex landscape | 3 | 3 | 9 | WLCG/GEANT/NRENs | Risk of sufficient bandwidth to support data delivery model | Work with NRENs, GEANT, etc.; Concepts like LHCCONE to prioritise and shape network traffic |
| Software | | | | | | | |
| 12 | Lack of investment in software skills | 3 | 3 | 9 | WLCG/HSF | Inability to sufficiently improve performance of core software | Training, education, hackathon programmes with HSF; projects such as IRIS-HEP, & other initiatives |
| 13 | Geant4 does not invest in long term strategy for performance improvement (x5-10) | 2 | 4 | 8 | EP/SFT | Cannot afford sufficient simulation; physics analysis suffers | Start major R&D program to devise strategies for fast MC, as well as better use of compute architectures |
| 14 | ROOT does not invest & prioritise I/O performance and data formats | 2 | 4 | 8 | EP/SFT | Significant bottleneck in data performance, streaming, access - starves CPU resources; this is also key for content delivery models required to address storage costs and concerns | Start R&D program on I/O and data management in ROOT, prioritise above other functions; has major cost impact |
| 15 | Event generators for higher order not optimised | 2 | 4 | 8 | HSF/Generators | Event generation too slow, becomes significant additional cost | HSF working group with generator community. Ensure progress is made |
| 16 | Experiments do not improve computing model performance | 2 | 4 | 8 | Experiments | Lack of investment of effort to reduce reconstruction time, decrease storage needs, etc. | Experiment working groups on all aspects of computing and storage: nanaAOD, reduce reconstruction time, fast simulation, etc. Must include work on data models. |
| 17 | Collaboration between experiments diverges | 2 | 3 | 6 | WLCG | Multiple solutions & duplication of effort too expensive to support in constrained budgets | Ensure common solutions deliver and adapt to majority of use cases |
| 18 | 3rd party software no longer supported | 3 | 1 | 3 | WLCG | gridftp, OSS tools; would need to replace or support in the community | Remain flexible to ongoing changes at all levels |
| Other | | | | | | | |
| 19 | Licensing limitations | 3 | 1 | 3 | CERN/Experiments | Costs rise too high | Only limited Oracle use today? |
| 20 | EC (& other e.g. EOSC) diverges from WLCG needs | 2 | 2 | 4 | WLCG | Contradictory requirements on collaborating sites | Engage early enough and steer, guide, advise; e.g. ESCAPE etc. |
| 21 | Long term reliance on short term funding (e.g. EGI, OSG) | 2 | 2 | 4 | WLCG | Risk that services lose support | Ensure key services have independent support commitments within WLCG |
| 22 | Security environment changes and breaks federation/trust model | 2 | 4 | 8 | WLCG | Disruptive to operation; would require significant investment in changing fundamental sharing model | Ensure our trust networks, policies, awareness are proactively adapting to the changing landscape |
| 23 | Lack of operational support at sites | 3 | 3 | 9 | WLCG/DOMA | Unreliability of services and resources | Focus storage and data services at large sites; serve data to smaller sites |
| 24 | GDPR not fully complied with | 3 | 2 | 6 | WLCG | Impact on services and collaboration in particular if some sites cannot run services. | Due diligence and clear policy frameworks. |

Likelihood:
 1: never expected to happen
 2: could happen but very unlikely
 3: could well happen
 4: will probably happen

Impact:
 1: we can deal with it, no problem
 2: a bit of a hassle but not too bad
 3: can be managed, but with significant effort
 4: crisis

<https://wlcg-docs.web.cern.ch/wlcg-docs/boards/MB/Risks/RiskRegister.html>



Conclusions

- ❑ Overall Run 3 will probably fit within cost expectations
- ❑ Preparations for HL-LHC ongoing
 - Significant efforts underway, bearing fruit
 - Infrastructure, software
 - Still some way to go
 - Long term software performance and adaptability is going to be a key area to address
- ❑ Skills retention for sw & computing remains a concern