WLCG Workshop 2022 - Highlights

7-9 November 2022 (Lancaster UK)

J. Flix

PIC
port d'informació
científica

Institut de Física d'Altes Energies
Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

WLCG + Rucio Workshop

The **WLCG Workshop 2022** took place in Lancaster University, UK
[7th-9th November]

The **5th Rucio Community Workshop** was co-located after the WLCG WS
[10th-11th November]

Thanks Roger and his team!

WLCG Workshop 2022 structure

The **WLCG Workshop 2022** focused on how we foresee the **evolution of our infrastructure towards HL-LHC**

- LHC experiments session
- Other experiments future plans session: Belle-II, SKAO, DUNE
- DOMA session
- HPC, cloud, and opportunistic resources session
- AAI and Tokens session
- Facilities session

+ **WLCG Collaboration Board** [closed session]

+ **Reception** (@University) & **social dinner** (@Hotel Morecambe, beach side)

WLCG Workshop 2022 structure

PIC
port d'informació
científica

Monday 7th Nov.

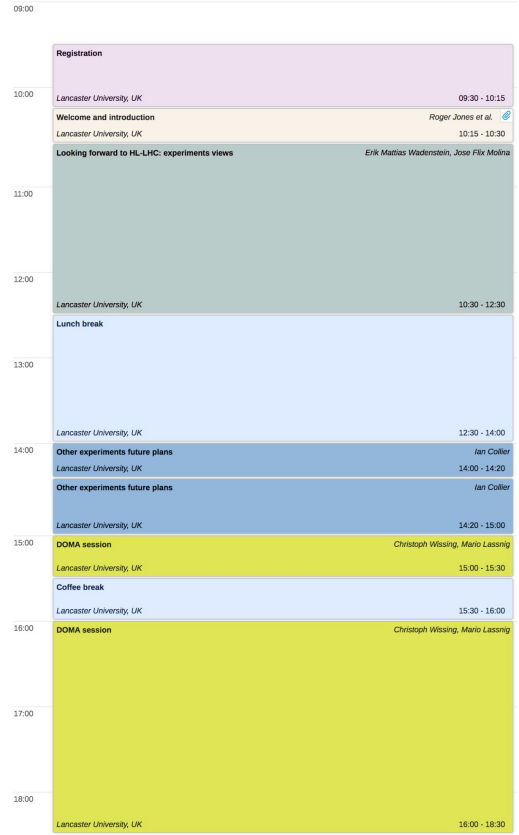

+Social dinner

Tuesday 8th Nov.

+Reception

Wednesday 9th Nov.

WLCG Workshop in Numbers

PIC
port d'informació
científica

- **80 Physical Attendees + 25 remote**
  - First in-person WLCG Workshop after the pandemic, quite well attended!
  - Happy to see new faces and old friends
- **40 Institutes from 17 Countries**

Local organising committee: M. Doidge, G. Edwards (administrative), G. Hand, R. Jones, P. Love, C. Noble (ex-officio), S. Simpson

Programme Committee: J. Andreeva, L. Betev, S. Campana, D. Cameron, D. Cohen, I. Collier, A. di Girolamo, M. Girone, J. Flix, M. Jouvin, M. Lassnig, B. Lin, M. Litmaath, A. McNab, D. South, G. Stewart, M. Wadenstein, C. Wissing

Workshop conveners: [Looking forward to HL-LHC: experiments views] M. Wadenstein, J. Flix; [Other experiments future plans] I. Collier; [DOMA session] C. Wissing, M.Lassnig; [HPC, cloud, and opportunistic resources] A. Filipcic, M. Girone; [AAI and Tokens] M. Litmaath, T. Dack; [Facilities] A. Dewhurst, O. Rind; [Workshop wrap-up] S. Campana

https://indico.cern.ch/event/1162261/

Presentations

- **38 Contributions + 1 panel (HPC)**
  - Scheduled total duration 16h30m
- **Live document** filled during the workshop:
  https://docs.google.com/document/d/1JzlmXERJJglaV7EDNC9EN4AgFVrKL3q5ePybKcWTC9g/edit
- **Lots of discussion** occurred in the sessions
  - Much better than in virtual events - more interaction
  - Even in the coffee and lunch breaks that happened in the same room

Great Venue!

Great social experiences!

Looking forward to HL-LHC: experiments views

[summary by J. Flix / M. Wadenstein]

PIC port d'informació científica

ALICE

- Operates according to the new model designed for Run3 and Run4
- The computing model favors local data access - further storage consolidation is welcome

ATLAS

- Several operational changes ongoing. Conservative & aggressive R&D activities towards HL-LHC. TDR by 2024
- successfully runs on a diverse set of resources. Engaging with industrial partners on R&D projects, enhancing opportunistic resource usages

CMS

- Run3 is in many ways a testbed for HL-LHC, maximizing the physics reach
- Aware of the use of network: Improving monitoring is a priority!
- CMS has been working with slim data formats for analysis for several years
- Efficient use of HPCs is expanding: goal to have this be transparent to central operations
- CDR by 2023

LHCb

- 30x increase in throughput from the upgraded detector: a big challenge in data handling & network & resources
- Progress in using HPCs. DIRAC developments to support HPCs.
- Analysis Facilities: bottom-up approach, collecting use cases towards a more structured activity

Looking forward to HL-LHC: experiments views

[summary by J. Flix / M. Wadenstein]

PIC port d'informació científica

DPM EoL affecting some sites: transition periods might affect availability of resources. Direction and supervision is needed

Data Challenges are useful for the experiments and sites. Important to re-evaluate the (next) targets soon

WLCG **tokens transition** timeline v1.0 agreed. Work in progress. No major problems expected here

Analysis Facilities are new elements in the system that can provide specialised hardware/software (GPUs, ML tools)

- Modern analysis approaches, e.g. declarative, Pythonic, HEP specific and from the data science community. Interactive
- Reduced data formats are a must

More **HPCs** could be integrated if minimum requirements are met, though HPCs differ considerably

- Can we rely on such resources for pledging?

OS evolution: plan from GDB, endorsed by the WLCG MB. It might need to be revised?

Rising **electricity costs**, increasing **HW costs** and **delivery issues** everywhere

- It clearly affects our extrapolation model for future resources
- Sites may need to reduce power consumption: preferences from the experiments - mitigation plans set up

Exploiting **non-x86** architectures coming: GPU, ARM, …

Broad consensus that we should accept **HEPscore** and use it for WLCG pledges

Other experiments future plans

[summary by I. Collier]

Great to hear from **Belle II**, **DUNE** & **SKA**

DUNE & Belle II are well integrated with WLCG

- Making effective use of resources and common tools at scale
  - Rucio, FTS, DIRAC, VOMS, CVMFS, etc
  - Belle II working on migration to tokens, https & webdav
- Compute and data requirements that fit well in WLCG infrastructure
- Clear roadmaps for both usage and capabilities – not going to duplicate here

SKA Regional Centre Network

- Expecting global data footprint comparable to WLCG around 2028-29
- At least some resources in facilities shared with WLCG
  - Opportunities to support one another
- Prototyping work began earlier this year
  - Teams forming, defining problem space, starting to test ideas.
    - One highlighted example was AAI – with a dedicated INDIGO IAM instance run at RAL
- Compute will have a very different mix – much big data sets – much bigger jobs
- Global data movement much more familiar
  - Tools and approach should be recognizable to us
  - For example, Rucio testbed moving data between (currently small number of) sites around the world

DOMA: Data Challenge Preparation

[summary by C. Wissing, M. Lassnig]

Time window appears to be fixed: **~March 2024** (before pp run)

- Start more formal process to organise
- Participation of DUNE and Belle II

Adjustment of **target rate(s)**

- Lowering to 20% or 25% seems a better match with present LHC schedule - discuss it - fix it - approval by MB

Definition of **technical content**

- Authentication will be token based for disk endpoints
- SDN functionalities
  - Opt-in for interested sites
  - Might be easier for some NRENs than for other - Approach relevant parties early

Ramp-up exercises

- More detailed scope and dates to be worked out
- Please step forward if there are particular tests you'd like to do

Preparation of **"execution environment"**

- Production endpoints vs Test endpoints
- "Challenge infrastructure and not the operations teams"

DOMA: Miscellaneous Items

[summary by C. Wissing, M. Lassnig]

PIC
port d'informació científica

Token based authentication for data transfers

- Decide about porting features of GsiFTP to Http/WebDAV (e.g. multi-stream) if necessary
- Coordinate timeline with WLCG AuthZ working group

Tape REST API

- Roll out plan for all T1s

WLCG data transfer monitoring

- Focus Xrootd monitoring deployment initially at CERN and FNAL (main sources for CMS pileup mixing)

Collaboration beyond LHC experiments

- A number of topics have been addressed in the context of ESCAPE
  - Joining efforts where same tools are being used (e.g. Rucio, FTS, DIRAC …)
  - Analysis facilities
  - Usage of shared sites & infrastructures, e.g. storage consolidation
  - Common AAI solutions
- Foster exchange with "close" projects: Belle-2, DUNE, SKA

HPC, cloud, and opportunistic resources

[summary by A. Filipcic, M. Girone]

PIC
port d'informació científica

Experiments overview

- HPCs used regularly by all 4 experiments, ~30 centers
- Most HPCs run MC, some everything
- CVMFSExec works for most without cvmfs mounted
- X86_64 mostly, with some GPU usage

MadGraph5

- Perfect example of using new cpu features (x8 speedup) and gpu offloading (60x) speed up
- Can significantly increase event generation rate, other generators might benefit as well
- Foreseen for production

Benchmarking:

- New arch (ARM) and GPU studies, new workloads to be defined
- I/O benchmarking essential for HPCs

HPC, cloud, and opportunistic resources

[summary by A. Filipcic, M. Girone]

PIC
port d'informació científica

Clouds:

- ATLAS fully automated on Google with k8s, similar on UVic cloud
- Accounting and authorization non trivial (IGTF)
- LHCONE integration - need to understand from traffic and security point of view

The vision of HPC in EU:

- 7B euro investment till 2027, Quantum Computing top priority
- Large investment in R&D, training and education, competence centers, centers of excellence
- Spread usage in industry

US HPC summary:

- (summary of summary is too hard…)
- Essential to work on WLCG HPC strategy document

Panel:

- For new architectures, essential to have a full chain development infrastructure at CERN
- Standardization of SW development, using cross-architecture sw platforms
- No need for significant ML infrastructure yet
- Find a proper/common way to communicate with HPC centers and organizations
- Cloud usage in the future depends on TCO

Token transition: discussion points

[summary by T. Dack, M. Litmaath]

PIC
port d'informació científica

Complexity of Token integration for Jobs vs Transfers

- Some scenarios and processes have multiple levels of detail
- The transition process will have systems become more complicated before things simplify
- There are many not fully answered questions here, but work to tackle and figure this out is underway

VOs with Multiple/Incompatible Tokens Issuers

- There are many reasons for different formats, and these will need to be considered and accounted for to not break for different “customers”

Refresh Tokens and lifetimes

- How to maintain a balance of security optimal and process optimal, and understand the risks of each option

Token transition: things to take forward

[summary by T. Dack, M. Litmaath]

PIC port d'informació científica

UIs – Not just for Users, but for Site Admins and Testing

- Make things available earlier than for the generic user

Can we connect CentOS7 EOL to the transition process?

- There may be areas where X509 systems can be retired alongside the OS
- Fingers crossed for some easy wins!

FTS must consider other communities which may not be orchestrated by Rucio or DIRAC

Token lifetimes will also need to be rediscussed

- Short lifetimes may lead to unreasonable implications for services that support our workflows

Questions to be answered around running IAM in HA

Facilities: HEPScore

[summary by O. Rind, A. Dewhurst]

PIC
port d'informació científica

Key agreements:

- The **first** production version of **HEPScore** will cover **x86 and possibly ARM**
  - If not ready for ARM, VOs will continue to work for future inclusion of ARM workloads
- In **April 2023** we will **switch from HS06 to HEPScore** with a 1:1 conversion
- Sites will be expected to only **benchmark new hardware** with HEPScore
  - Old hardware does not need to be re-benchmarked
  - However, sites wishing to re-benchmark old hardware may do it
- While it is not needed initially, in future the **APEL accounting** system should evolve to handle multiple benchmarks
  - Need to track the mix of benchmarks being used

Facilities: Energy crisis

[summary by O. Rind, A. Dewhurst]

It is likely that this winter at least some sites will be asked to **limit their power usage**

Easiest option for both Sites and Experiments is to **turn off some old hardware** however this may not be sufficient

Rod's work to **limit clock speeds** looks the most promising way forward that doesn't involve switching on / off hardware

- Urgently seeking **other sites to further investigation**

Facilities: Hardware costs

[summary by O. Rind, A. Dewhurst]

PIC
port d'informació científica

Presentations from UK and US perspectives

- Costs in recent years have been flat or going up!
- Not obvious what new normal will be

Energy costs have always made up a **significant fraction of the total cost of ownership** and this fractions looks to be increasing

- Running hardware for longer isn't necessarily an option

Possible solutions involve making **use of ARM** or **HPC resources**

Flat budget model study group kick-off meeting after the workshop:

- Agenda: https://indico.cern.ch/event/1223854/

Facilities: Other topics

[summary by O. Rind, A. Dewhurst]

PIC
port d'informació científica

Analysis Facilities

Lots of ongoing development - use the technology as building blocks, without locking ourselves into a predefined architecture, and engage users to adapt the technology to their needs

Networking Topics

Lots of work to do and decisions to be taken in advance of DC24. We do not want sites purchasing equipment earlier than necessary to meet the defined goals.

Tier-2 as Cloud

New tools and continuing development have been making the idea of a “cloud-like” k8s-based Tier 2 increasingly viable

Summary

- **First in-person WLCG workshop since 2019** (Adelaide!)
- Very **productive and pleasant event**
- **Good feedback** from attendees
- **Live notes**, **full recordings** (available at some point soon)
- **Good coverage of topics** and identified list of **actions** from the workshop
- Will try to continue the series of workshops as **in-person-meetings**
- **Thanks** to all that contributed to make this happen, in particular huge thanks to **R. Jones and his team** at Lancaster University