

WLCG Workshop 2022 - live notes

Monday morning

Looking forward to HL-LHC: experiments views

ALICE (M. Litmaath)

ALICE has 95% of data processing done in the LAN and a write/read ratio 1/30. Would be good to understand how much this is for other experiments. Those numbers will likely drive the decisions about the computing model.

In ALICE an AF is simply a Grid-enabled site with a particular choice of the hardware to favour high I/O workflows (TBC). Different from what the HSF AF working group is looking at.

Alice is also moving to multicore jobs, possibly to whole node allocation in future. How to allocate a whole node? This is challenging for a site, and not everyone would be able to support this.

Storage consolidation: storage best located at sites which can operate storage at scale. (~1PB today). How does ALICE plans to support storageless sites? (is that the 5% WAN traffic in job processing mentioned above?) ALICE will run Monte Carlo at those sites.

ATLAS (D. South)

ATLAS highlights how important the IAM support is for the future planning of tokens migration

Could WLCG better support sites in deploying middleware? May be using some community based structure like WLCG discourse,

For ATLAS, turning off CPUs within a few hours cycles is not optimal.

CMS (K. Ellis)

LHCb (B. Couturier)

Monday afternoon

Other experiments' future plans

Belle-II (S. Pardi)

SKA (R. Bolton)

DUNE (M. Kirby)

DOMA

Storage & transfers with tokens(P. Vokac)

Basic token support in Rucio and FTS.

Upcoming Dirac 8 brings support for tokens.

Tokens should be obtained transparently - details to be figured out in the WLCG AuthZ group.

WLCG JWT storage compliance testbed is developing.

No DPM support for tokens planned - timely phase out required.

Requirements from the experiments have been collected

Storage systems must work with both X509 and tokens for a transition period.

Target: DC24 with tokens only. Requires deployment campaign to finish well in 2023.

No firm plans yet for tape/archival storage with tokens yet.

Some clarification needed regarding features that GridFTP no has but might not be ported to WebDAV.

Improving of error messages from WebDAV/Tokens is a cross experiment effort,

Development requirements need to be clearly made and communicated. High level plan exists in Rucio, some details need to be specified and worked out.

If X509 is to be decommissioned a clear end date needs to be defined. (Otherwise support needs to be provided by storage providers forever).

Capabilities should be used in favour of groups, that should make site configuration easier.

Handling of user data is not discussed in depth yet.

HTTP TAPE REST API: Status & Plans (M. Patrascoiu)

Specification V1.0 exists.

Existing implementations need some integration with experiment frameworks
Review and a new minor release, v1.1 will happen in 2023.
CERN CTA wants to deploy Tape rest API by the end of 2022, RAL wants to follow by Feb 2023.

Timeline for all T1s needs to be defined

dCache HTTP API already included in 8.2, in long-term-support release
BNL upgrading to 8.2 in December, ready for API testing (DUNE already on 8.2)
Priority in API? Staging metadata (incl. priority, activity) is not in the spec document, storage is free to decide(?)
Bulk-submission vs spoon-feeding requests, will stay the same, not part of the API
dCache will ignore priority
Might be possible to develop "something" at the glue layer between storage system (dCache) and tape system (hpss), but this is a long-term development prospect that does not yet have a design sketch

IMHO Every time a VO requests the ability to control and reorder prioritisation normally is only achieved by limiting efficiency and overall throughput (be it tape staging rates, FTS throughput rates, Single file access rate etc.

Dataflow Monitoring (B. Garrida Bear)

Most important:

- Improve Xrootd monitoring
- Harmonization of transfer monitoring
- Site provided monitoring

Can you do site monitoring without packet marking? answer is yes, it's two distinct things
typical site has multiple links, it adds up counters from the different network devices
packetmarking/flowlabels is to investigate the activity/purpose/owner

what can we learn from the the total site in/out rates: zeroth-order understanding,
understand bottlenecks, capacity vs actual used/utilisation

can we focus on a couple of sites for xrootd monitoring (mostly cms focused?)
dcache + xrootd/eos - primarily CERN and FNAL being the main sources for pileup-libraries.

network topology: how deep we want to monitor / e.g. multi-science labs / segregation
human-readable document: mandatory vs optional (don't provide e.g. security-important data)

Network Management Enhancements for the HL-LHC Era (Diega Davila Foyo)

lifetime of circuits: at which scale do you open them?

it takes time to create the service, not good for one file, should be on dataset level

very tricky to diagnose network problems, esp. when crossing network organisation boundaries: single pipe should make things "easier" (or at least not more complicated)

Does it also work for xrootd? yes, also works for http certificates: how does it work with IP address mapping? statically assigned IP SANs

What to focus on testing for 2024? If CERN is participating then that's a major focus. then also tier-2s on ESnet.

challenge is deployment, maybe demonstrate for "tricky" links?

vo conciliation: lhcb/dirac integration? to not eat away lhcb traffic

tier-1s in europe? please speak to NRENs NOW, typically not a lot of alternate paths

To control rates between sites and VOs when worried about multiVO multisite orchestration, is control via FTS link settings a good enough measure (ie use SENSE to control routes and rates between sites/xrootd server pairs, use FTS link control for intra-VO rates?)

Data Management and Data Access: lessons learnt from ESCAPE (X. Espinal)

Not only is there benefit to getting different communities to use the same technological solutions, there can also be benefits in support effort and hardware procurement by sharing single set of resources.

The path to the next Data Challenge (M. Lassnig)

"Clear description of expected bandwidth (average and peak) and access pattern" not only useful for network providers but also for storage providers (how much local access i/o rate / connections are expected at the same time.)

Making sure that we exercise the "right challenge", we want to challenge the network infrastructure and not the Ops teams. Needs certainly some preparation within the experiments.

Tuesday morning

HPC, cloud, and opportunistic resources

Experiments and HPC (D. Cameron)

ATLAS HPC usage in the last years is dominated by Vega.
Main workflow is MC sim (not all of the workflows are suitable) - done similarly by all of the experiments

In slide 8: using `cvmfs_preload` is missing (used by CMS in Barcelona Supercomputing Center)

Before pledging the resources, we need to be sure we account well for the usage of HPCs.

MadGraph porting to GPUs (S. Roiser)

How does the performance picture change when using a more realistic / more diverse process, like $p > t \sim j j$?

Since $g g > t \sim g g$ has almost exclusively diagrams that interfere, isn't this a worst-case scenario for the CPU?

Were the SYCL/KOKKOS ports actually full ports to an alternative language, or just minimal changes to the initial CUDA version? Wondering if the picture would change if this was done as a full re-write by someone who was focused on the portability language.

Benchmarking HEP workflows on HPC (D. Southwick, D. Giordano)

Challenges in integrating cloud infrastructures in production (F. Barreiro)

Question

From Rohini Joshi: Thanks Fernando! You mentioned PanDA running on multiple clusters and I was wondering if there is work done towards Infra as Code and/or how you manage running and porting these services across Cloud providers and clusters?

Answer: Just to avoid misunderstanding, on the cloud the only component running is the "batch queue" in the form of a k8s cluster. Panda/Harvester run centrally at CERN. I do manage my Google infrastructure through Terraform (<https://github.com/gcp4hep/terraform-scripts>). On Amazon I create resources on the command line/manually, because the project is less mature/active. Terraform scripts are also cloud dependent, i.e. my Google script is no good for Amazon. Despite the initial learning curve, managing the clusters through Terraform makes things easier for me and if the Amazon project would continue, I would eventually also script it.

The vision of HPC and Quantum Computing in Europe (Daniel Opalka)

US HPC blueprint workshop summary (F. Barreiro)

It's time to review the documents we have wrt HPC usage and centers. Things have evolved, and it might be time to write a coherent document on how to approach HPCs (Simone C. comment)

HPC/Cloud-as-a-service panel

Antonio - CMS

Zach - ATLAS
Latchezar - ALICE
Federico, Ben - LHCb

Experiments asked to answer some questions posted in the agenda.

- "New technologies (non x86_64, GPU, FPGA, APUs, QC) need large SW development investment. How will you handle the effort (manpower, common projects)? Does the effort suffice? Are there high risks of failure to port SW to new platforms?"
- Use of ML techniques is growing in many science domains, including HEP. How do you see the evolution of resource needs for specialized AI architectures? Would access to HPC sites resources help?
- There are many cross-domain forums, advisory groups, bodies, consortia (at least in EU) promoting and shaping HPC future (competence centers, centers of excellence, ...). Would WLCG benefit to join them? US, EU, Japan, other countries, is there an organized effort to support HEP on HPCs?
- Clouds, what are the expectations for public and private cloud usage in the future?"

It could be good that answers are posted from the conveners (summary) or asking the experiments to add some notes in indico

ATLAS summary (Zach)

On non-x86 resources:

We support CMS's mention of build resources for new platforms. We are still in a period of early adoption for most non-x86 resources. We in ATLAS have mostly focused on infrastructure work and work that is generically useful independent of the accelerator or language solution that we settle on. There are a few successful applications on GPU, like FastCaloSim and calorimeter clustering. There is a risk in this period that we adopt a language that is dead in a few years, as it seems that we are in a period with a large number of different solutions and it seems unlikely all of them survive an attempt at standardization. There is a lot of interest and excitement about accelerators, but little expertise, and to that end training programs are very useful.

On AI-ML:

The use is growing in analysis, but in the medium (5-year) term, it is unlikely to rise above the level of a few percent of the experiment's computing. Nevertheless, having some capability is very important. There are some disruptive R&D projects that could significantly enhance our use of and need for AI/ML (and related accelerators) like ML-based fast simulation (beyond our current fast simulation which uses ML but not as extensively as some proposals) or GNN-based tracking. Generally, these projects use common tools like Keras, Tensorflow, and Pytorch, so the adoption of accelerators is relatively straightforward.

On EuroPHC engagement:

We have great colleagues involved in EuroHPC and the corresponding groups worldwide. As discussed earlier in the session, it would be useful to update the WLCG document on HPCs, perhaps into two documents: one "administrator-facing" document for

those who are running an HPC, to explain what technical support and solutions are needed, and one “funding-agency-facing” document for groups considering building an HPC, to explain what general features are important to the LHC experiments. Having some ability to shape HPC evolution is very useful, keeping in mind of course that we don't need to use every HPC worldwide, but a few large machines are sufficient to serve the community.

On Clouds:

We support much of what CMS said. As Fernando's presentation showed, things are working quite well, and there are a few clearly-defined items that still need a bit of work. Because these are resources that one must pay for, we have left it up to the sites to decide whether they would like to deploy resources via a cloud. Some have chosen to do so, as CMS mentioned; this also includes Canada, which are operating a quite successful academic cloud. We are working on a Total Cost of Ownership for cloud resources, as mentioned this week, which we hope will be helpful to groups trying to decide whether cloud is cost-effective; our plan is to provide guidance about things to look for when setting up an agreement, rather than to make a strong recommendation for or against the use of cloud resources.

Maarten: network problems to interconnect Clouds (not a trivial thing).

Tuesday afternoon

AAI and Tokens

Token transition state of affairs (M. Litmaath)

Discussion

Rcauth: Service that needs certificates, not just use tokens but allows IAM to AuthN/Z for a service which cannot cope with tokens

Not discussed nor overly considered to WLCG - part of the architecture but does not currently play a role. Not about users - about endpoints that support X509. May be useful in the future.

When talking about job sub/data transfers, different scenarios? Presumably it's the same mechanism? Separation of data transfer from job submission. Tokens orders of magnitude more complex - tokens only used for one communication due to exchange and renewals. How far are we from the full scenario?

Some work to do. Need to discover whether we want to do 1/1 token/job, etc. have some flexibility. Will need to be steered by operational exercise. Even one single token would be an improvement if it could not delete - compared to proxies. Minimal use can bring big advantages - then to understand how far to push things

From a site point of view, may need to deal with three different issuers - ensure that all work, as don't have separate storage space

Had an opportunity to make things the same, but now have multiple solutions

Will need to ensure to have the tools needed - as it is possible for it to break for one "customer"

No quick way out - various reasons for different profiles. On WLCG started with AARC to come up with an EU standard for tokens. The further this went, the less it was liked - also needed to work with OSG.

Modify SciTokens, but now EGI sites need to support multiple

Can be done, but requires attention to detail

Security risk with Worker node - doesn't change between X509 and tokens. Could be more strict with tokens due to fine tuning, however the hard part is storage of these. Feels these are complimentary, security relies on capability of storage system

Always be implementation details for SE flavour

Maarten's opinion it's better - there are no all powerful credentials at the WNs. Instead give the WN just enough to do its job and no more. Details can be tuned, and different experiments may reach different conclusions - some things may be easier/difficult in different places.

For data management - there will be teething problems, and the advent of tokens will not make configuration easier to begin with and will need people to support each other.

Cannot assume the whole grid will be token ready simultaneously - not presenting a complete blueprint, but rather where we are, what we have, and the plans for moving forward

Petr: confident that they can configure storage for grid token usage, but not for personal home areas but this will not use tokens

The focus has been on the main workflows for a VO.

Pain-points to be highlighted and documented, to share and support.

Slide 12 - 70% progress. CMS progress is not included here. In CMS, already 100% move to new Condor CEs and now exchanging tokens with all CEs. All ARC CE using REST interface as well. Check the linked page for further information and to understand where this is missing.

Token lifetimes - the amount of time FTS needs to complete a job may be longer than the token lifetime, what happens here?

Refresh tokens - can refresh access tokens if they are still needed at expiry, ie the job lasting longer than the token

For Site admins, the UI is useful for testing components? Is there a plan for tokenising the UI.

Will note this down - already token support in some areas, but will need to understand.

As this is more in user area, it has not been focussed on as much

Possibly a natural point where X509 is ended with Centos7. Then use tokens with the new OS in mid 2024

It is a nice point, has not been discussed

Were seen as orthogonal, would be good to take advantage of this in some areas.

FTS will need a clear part in the workflow definition to avoid producing a solution which is not WLCG specific and alienate others

Good point - can be actioned from the CL.

Will need to allow FTS to control these things, and not just rely on Rucio and Dirac to sort these things

But who will define these outside of WLCG?

Not known, but knows who will get the tickets and questions when they come up

Do they know what they want to do?

Some do, others less. One community submits via WebFTS. Want WLCG technology to reach outside of the WLCG

Will need to discuss to understand pros & cons

DC24 may not be the same as Run4 - things may change, but the aim is to move away from X509 and voms

If we design for short lived and then find no 24/7 support - and then things get extended

Will need to find a solution - cannot allow things to be too boxed in by security concerns with shorter lived tokens

Can you have multiple IAM servers providing access?

In theory, yes. Voms today are all at CERN for the HR-DB interaction.

Wednesday morning

HEPScore

HEPScore – a new CPU benchmark for the WLCG (Randall Sobie)

<https://indico.cern.ch/event/1162261/contributions/5092745/attachments/2543843/4380269/Sobie-WLCG-HEPScore.pdf>

Discussion

Questions were asked about if it is better to normalise HEPScore == HS06? Domenico believes that the advantages of this outweigh the drawbacks (as stated in the slides).

Doug asked if there was a benefit of using HEPScore compared to HS06 in terms of the amount of actual work done. Randy said yes.

Pepe asked if there was a benefit of switching to HEPScore given how close the correlations were. Randy answered that we could probably get away with it now but in future the discrepancies will get bigger.

Dan Traynor - Do we see a benefit for turning on Hyperthreading. Randy answered that there was a slight improvement.

Simone asked about the ARM question. When should this be ready? Alastair has said that sites would only need a working benchmark for procurement by June next year, if the benchmarks could be ready by then. Maarten pointed out there is a difference between a benchmark and the VO being able to keep a batch farm full.

Simone summarised by saying he believes it would be better to go with the x86 HEPScore and allow the VOs to do solid work on ARM.

HEPScore - transition plan to the new benchmark (accounting) (Julia Andreeva)

<https://indico.cern.ch/event/1162261/contributions/5117866/attachments/2544039/4380614/NewBenchmarkWS.pdf>

Discussion

Doug asked how the finance sector makes the transition between accounting systems. They don't have a step function.

Helge thinks that having a obviously different HEPScore vs HS06 value will allow problems in the accounting to be quickly picked up.

Helge believes it is essential that the accounting system has the capability to use multiple benchmarks. Adrian notes that we have to trust sites with the data they input.

Alessandra asks if a site wants to move benchmark, why wouldn't they re-benchmark the whole site. Julia said it would be possible to maintain two different systems. Alessandra believes it will make things very complicated.

Domenico believes we will need more records in the accounting system. One of the main reasons to not re-benchmark is that some sites could lose capacity, which we wouldn't want to happen.

Simone believes that we should agree for HEPSCore == HS06 and this means that the migration can just "happen by construct". He also said that we need to be able to monitor this.

Oxana notes that there may be completely different benchmarks for other communities and that the accounting portal should be able to deal with several.

Adrian notes that the contractors have 2 developers and a project manager at the moment to improve the accounting portal. Alessandro diGi notes that while Apel has been providing a very reliable service for 10 - 15 years it hasn't demonstrated that it can rapidly develop things and he wouldn't mix it with the transition.

Analysis Facilities

Analysis Facilities (Alessandra Forti)

https://docs.google.com/presentation/d/1Rph1UtlcRvsjTBsZBrXdYLSWuH10cPAh0RF5dIOZRTA/edit#slide=id.gd170caf633_1_0

Discussion

Doug says that we don't have a common definition and those that pushing the technology are defining the requirements. Doug feels we should reach out more to find out what the end users actually need. We can't just discuss this within HSF.

Alessandra agrees but in the HSF AF meetings there are often experiment experts to report. For example the NAF UK and the SWAN&Coffea meetings were solid users reports. Also the people pushing for things are those working on them.

Domenico likes the monitoring idea and thinks it should be done from the beginning with central monitoring in mind. Alessandra knows by experience you need first a schema and a collection method before going to MONIT. A WG is necessary to make a solid case.

Energy Issues

Evolution of the WLCG Energy Needs (Simone Campana)

<https://indico.cern.ch/event/1162261/contributions/5124364/attachments/2543687/4381025/WLCGEnergyNeeds.pdf>

Discussion

Andrej: How do you achieve PUE of 1.1? Liquid cooling? - No, air cooling

Tony: Uses increased air flow from outside. 1.1 is average over the year.

Dave B.: PUE is not going to get much better. Hardware will not improve much, other than switching to new architectures.

What are "services" in the pie chart? A: Other CPUs providing services for lab, build services, etc.

Mattias: What about the office space needed for the scientists? Should this be part of the calculation?

Tony: CERN has an environmental report including CO2 emissions, but not clear how to answer this particular question

Risk of loss of computing resources: CMS experiment perspective (Katy Ellis)

<https://indico.cern.ch/event/1162261/contributions/5124354/attachments/2544121/4380774/21109%20-%20CMS%20Electrical%20Power%20WLCG%20Workshop.pdf>

Discussion

Energy conservation considerations at DESY (Thomas Hartmann)

https://indico.cern.ch/event/1162261/contributions/5124395/attachments/2544224/4380976/DESY_Cluster_Energy.pdf

Discussion

Voluntary Load-shedding During Peaks (Rodney Walker)

[https://indico.cern.ch/event/1162261/contributions/5117873/attachments/2544149/4380828/PeakLoadShedding_WLCG\(1\).pdf](https://indico.cern.ch/event/1162261/contributions/5117873/attachments/2544149/4380828/PeakLoadShedding_WLCG(1).pdf)

Discussion

Andrej: We have a rack by rack system for suspending jobs, can be adapted to do CPU throttling as well.

What is the power saving? Approx a factor of 4

Brian: Consider geographical location as well as cooling issues in the Summer?

Ian: Impressive savings with clocking down, but how much does it slow down the jobs?

Was not measured, assumed to be proportional but might depend on job efficiency

Alessandra: OS total control on depends on the CPU vendor. AMD doesn't work at the moment. For me killing the WNs for 4 hours is not doable. I mean it is but with some heartache particularly for users

Will this shed enough load? Is throttling a more efficient approach than suspension?

What fraction of the total node power is reduced by throttling CPU? (See Simone's talk)

Budgets and Hardware Costs

UK hardware cost trends and on how to reduce energy consumption (David Britton)

https://indico.cern.ch/event/1162261/contributions/5134740/attachments/2544372/4381256/21108_lancaster.pdf

Discussion

Doug: Disk prices look like they're flattening out faster than indicated in the plot

Partly due to exchange rate

Doug: How well do you know the energy source mix in the UK? Have you talked to the power companies about the effect of your reduction

We are tiny for them. However, we should be in the frame of mind where we are trying to be more green

Simone: We keep coming back to the idea of providing different amount of resources in different times of the year. We can do this, e.g. can pledge by quarter in CRIC. We should reconsider potential models with varying pledges throughout the year.

Might need to coordinate scheduling across sites and within experiments to do this
Planning can be done, but usage influenced by things like conferences.

Might need new metrics to monitor this.

Consider a working group to consider things like carbon footprint and optimising hardware purchases and usage? Need to think seriously about accounting definitions as well.

Dwayne: Having pledges by quarter is quite good, but fossil fuel use is most common in the winter, and most of the resources that pledge to ATLAS share the same winter, so this could lead to less compute available overall in the winter. This can work if ATLAS is able to exploit a move from a flat pledge to a pledging scheme where more is available in the summer than we do in the winter but...

David S: ...ATLAS does the reverse, it sends a lot of its simulation in the winter in preparation for conferences in the early spring

Flat budget and HW costs: USA reflections (Chris Hollowell)

https://indico.cern.ch/event/1162261/contributions/5124388/attachments/2544315/4381138/hw_costs.pdf

Discussion

Pepe: Is cost model study public?

No, it's internal, but if there's interest, may be able to make a redacted version available

Dwayne: If it's a question of being green, does offloading compute to HPC centres really help?

Alastair: The idea that costs won't come back down may be overly pessimistic. We've seen purchasing delays coming down.

Chris: Typically, we have seen when costs increase they don't come back down, but it's not impossible. Not clear they will come down when supply chain issues abate.

Tony: Supply chain issues and cost may be due to different sources and not necessarily correlated.

Capability improvements may bring cost/HEPSpec down

Tony: The question will be "what flat cash gets you"

Networking Topics

WLCG Networking Topics from the LHCONE/LHCOPN Meeting (Shawn McKee)

https://indico.cern.ch/event/1162261/contributions/5124358/attachments/2544019/4380566/WLCG%20Networking%20Topics%20from%20the%20LHCONE_LHCOPN%20Meeting.pdf

Discussion

Alastair asked if we should make the PerfSonar plots more visible to "shame" sites into keeping the services up to date. Shawn said that the plots are used all the time in some circles but agreed they could be regularly shown more widely.

Data Challenge 24. What throughput should we be aiming for? Simone commented that with the previous data challenge he thought that some sites could/should have done a lot better.

Randy commented that SKA is likely to come online at roughly the same time as HL-LHC and that the Canadian SRC will share the same site as their current Tier-2.

Doug commented that he is grateful that DUNE is being included as there is an ~80% overlap with sites. While DUNE's data rates are not large, if they have a supernova they would need high capacity for a short period of time.

Matt asked when the DC after DC24 would be. Shawn said that was still TBD.

Shawn wants to know what the mechanism should be for data challenges. Previously we have used Rucio, however are additional features needed? Doug is wondering if more people effort is required to organise these events.

Tier 2 Operations

K8s @Victoria (Ryan Taylor)

<https://indico.cern.ch/event/1162261/contributions/5124359/attachments/2543972/4380486/20221109%20WLCG%20workshop%20-%20UVic%20T2%20k8s%20.pdf>

Discussion

Ofer asked about security concerns in multi-tenant k8s environments. Ryan said there are tools for enhancing and managing pod security.

Doug asked how Kubernetes can ensure that fairshares are being respected. Ryan mentioned a variety of tools available (e.g. Volcano) but noted that currently they support a single VO.