

Table of Contents

Notes for the preparation of the WLCG workshop of February 1-3, 2016.....	1
Topics.....	1
Raw input.....	1
OSG.....	1
ATLAS.....	1
CERN IT.....	1
GridPP.....	2
DESY.....	2
CMS.....	3
Belle II.....	3
ALICE.....	4
HTCondor.....	4
LHCb.....	4
NDGF.....	4
Tier-1.....	5
Security.....	6
USATLAS.....	6
OSG.....	6
Summary of input.....	7

Notes for the preparation of the WLCG workshop of February 1-3, 2016

Topics

Raw input

OSG

In addition to Technical Forum topic list from Ian [\[?\]](#), I would add:

- Can we revise our approaches to the pledges / MoU to be more friendly to non-traditional-WLCG resources? I.e.,
- Make the pledging compatible with allocations at supercomputing centers (which are done in HS06-hours, not HS06-cores as we do in WLCG).
- Understand how resources (such as HPC centers) that are not going to sign a MoU can fit into the security and accounting framework.

Topics I'd be interested in contributing to the discussion:

- Evolutions of data federations.
- Deployment of IPv6.
- Need for / evolution of the traditional CE / batch system setup.
- Mechanisms for small computing sites (volunteer computing in Ian's slides).
- Info services: SAM-like functionality and service discovery.

ATLAS

- WLCG long term evolution (the new forum initiated by the MB). Work plan, strategies, etc ..
Includes:
 - ◆ Rationalisation and simplification of the infrastructure: consolidate and federate storages, diskless centers, network/storage/CPU interplay, volunteer computing for small centres e.g. T3s, a model for the future
 - ◆ Which middleware do we need for the above? Do we need more than what we have? Less?
 - ◆ Data access and data transfer protocols. What is the way forward?
 - ◆ CPU benchmarking: some conclusion on the various discussions. What should we do?
 - ◆ WLCG accounting (was very broken till very recently). Validation, improvements. Agree on a limited set of functionalities we need (and we need them working).
 - ◆ Network, network operations and network monitoring
 - ◆ Information System (GOC/OIM, BDII, Rebus, you name it). Converging on the various discussions on the working group and rationalisation of the system.

CERN IT

- Data Management:
 - ◆ x years of federation - what are experience / plans
 - ◇ is the approach used/useful
 - ◇ is the term federation well defined (I would claim it is not - federation of namespaces vs clustering of distributed storage (eg HU/CH) - all with very different semantics)
 - ◇ how do experiments / wlcg see the role of federation in the medium term ? [basically the proposals by SDC/DSS which are being discussed at the montenegro meeting]

◇ how do xroot / http federation/redirection integrate [also part of the above discussion]

- Storage Management
 - ◆ impact of cloud storage approaches (eg distributed key value pair services (eg ceph/s3) and devices (object disks))
 - ◇ integration plans by the different storage sw providers (sw providers)
 - ◇ any visible gains for sites/experiments so far? (active sites)
 - ◇ do experiment plan to integrate new technologies on the media level (behind existing storage systems) or as direct end-user service? (discussion)
- Infrastructure Analysis
 - ◆ progress in automating metrics aggregation / processing / and missing correlation data between storage - batch - experiment workflow
 - ◆ any real success with statistical or ML approaches to understand operations/resource usage at sites / experiments / global wlcg level?
 - ◇ popularity prediction CMS (Valentin/Tony)
 - ◇ scrutiny input on popularity
 - ◆ hadoop use cases
 - ◇ for analytics (MR, spark, python, root, jupyter, zeppelin)
 - ◇ for RDBMS (Oracle) replacement (eg dashboard, ATLAS event index)
 - ◇ status of service and remaining issues (upgrade time, technical support)

GridPP

- Sites simplification at the last workshop we talked about simplifying the sites setup going from eliminating services at grid sites to new batch-less and storage-less setups. Work is currently on going to simplify the info system, there should be work on going to look at the elimination of APEL boxes and reducing the number of storage protocols. Though until we move to a different type of storage I don't see that happening which brings me to the next point
- Storage evolution and introduction of object stores. ATLAS is working on the Event Service and this relies on Object Store type of storage. The most common being ceph.
- Review of the requirements for batch systems and elimination of at least obsolete batch systems.
- Progress on the use of containers both to install software and run jobs
- Monitoring this requires a review of the SAM tests.
 - ◆ Data Analytics and automation we should explore the effort going on within the community to exploit data analytics frameworks to monitor the system and if this information can be harvested and integrated.
- Work on parallelization and use of new architectures. GPUs/ARM/..... this is more on the experiments but as a site I'm interested to know direction and timelines for things to happen.

DESY

So as a high level session or even just as information for the audience, it would be good to learn what WLCG feels to be responsible for and if this matches the ideas of the experiments, non-CERN technology providers and sites. What is the split in duties between WLCG and EGI (for example). For some years, it turned out that direct contacts to the experiments is more useful than with WLCG, at least for non-CERN technology providers.

There is of course the problem of funding for non-CERN technical teams. In order to get funding, we need to involve ourselves in EU projects. However, currently there are no projects, which would allow us to work on storage for WLCG. What are the future plans of WLCG on going for larger EU funding, allowing non CERN technology providers to join?

On the technical 'Ian' list:

We are certainly very interested in what WLCG is envisioning in the domain of data management and data access. e.g. : Are there planned significant changes in the storage model, e.g. on federated resources?

Secondly, but not less important, are ideas on "Long term distributed computing models", e.g. what is the impact of the availability of local cloud infrastructures (OpenStack and ONE) or the upcoming 'container' technology of the computing model of the experiments. And maybe what is the role of HPC in that area?

CMS

The list of topics from the future Technical Forum spans over many disciplines and CMS is certainly interested in them. Covering them at the workshop would be great. On Ian's Slide 6, the first 3 topics under Compute and the Networking part are priorities for CMS. Topics that we would like to discuss are also on working with industry, openlab, other communities and initiatives over EU and oversea (including OSG initiatives, HSF and many others). How will we do this, when, on what. CMS would like to know if there are plans.

Belle II

Since early August, we have started the large-scale MC production campaign (of course, relatively smaller than LHC experiments, though). Through this experience, I would like to propose the following items as input from us. But please forgive me if these items were already discussed in the past workshops. In addition, I am not familiar with political relations within the WLCG (i.e. what was developed by WLCG, and what was not), so I may propose the items which should not be discussed in WLCG. Anyway, let me propose a couple of things as an observer.

1) non-VO specific FTS dashboard monitor

Since the WLCG+alpha share the infrastructure, it is good that we have the FTS as a common tool and the FTS dashboard as the common monitoring of the file transfer.

In the Belle II computing model, access to remote data is a key concept, it would be nice to have a common tool together with a common monitoring. The current "tools" are VO specific and the WLCG transfer dashboard is dedicated to those specific tools.

2) GGUS ticket

The Belle II distributed computing utilizes both grid and non-grid sites. I also believe that LHC experiments have utilized much more non-grid sites than Belle II, including HPC, vacuum, etc. Actually, DIRAC allows us to utilize the sites without grid middleware (CE), but issue tracking (with GGUS) is limited to the sites registered in GOCDB. There might be a room to consider how to treat non-grid sites within WLCG.

3) perfSONAR mesh.

It is really nice to see the network condition with perfSONAR mesh. Now, Thanks to Shawn McKee, we, the Belle II started using the perfSONAR mesh. Is the perfSONAR maintained in WLCG ? We would like to keep using it.

4) network-SE correlation monitoring tool

Is there any monitoring tool combining the perfSONAR mesh, FTS dashboard, GOCDB downtime information, etc. ? In order to transfer the job output data from the source SE to the destination SE, there are many components which affect the data transfer performance, for example, stability of the network path, the

status of the SE, heavy network traffic, etc. It can be check one by one with some web portals. I just thought it may be good to have a common monitor to check the correlation of the network condition, status of the SE, etc.

For the moment, my proposal puts a bit weight on the monitoring issue. If I come up with more topics, I will let you know. But, please keep in mind. This is just our thoughts.

ALICE

An initial selection of topics that ALICE would like to see discussed in Lisbon:

- Data: A global scalable name space based on EOS and/or CVMFS. Allowing to have a "grid of regional clouds".
- Networking: IPv6 deployment plan.
- Benchmarking and performance monitoring: A common, free benchmark, whose results are made available through MJF files on the hosts.

HTCondor

I can see value in thinking/discussing two "types" of topics - topics that address the "how" and topics that discuss the "what".

Given the scope and challenges of the WLCG I believe that how we work together will hold the key to our success. I wonder how much interest/motivation is out there to address organization, allocation of resources, decision making,

When it comes to the "what", it will be useful to address the challenges we face in defining, implementing, monitoring and verifying resource provisioning policies.

LHCb

Apart from ongoing efforts e.g. in the areas of file access & protocols, cloud computing (storage?) we thought it would be good also to give some report about experiences in Run 2, especially as the LHCb Turbo stream has caught some attention in Barcelona. Also LHCb is discussing and brainstorming for Run3 computing this November in a dedicated workshop and possible can present a digest of this meeting. Would this be ok? what are the thoughts of other VOs?

NDGF

indeed, there is one peculiarity in our region, namely, that we all rely on heavily shared resources which position themselves as HPC, ad some even appear now and then in Top500. We do it already for 10+ years, and we even developed middleware for it (yes, the ARC-CE), but it seems that we hid it so well, nobody realized that we routinely run an LHC Tier1 on HPC facilities 😊 And now some other countries start exploiting such possibilities, so it might be that we'll get a company. So, perhaps one of the topics should be indeed around usage of non-dedicated resources, and particularly various national HPC facilities.

From Ian's slides, I'd say that most topics in the "Compute" bullet are very relevant: yes, HPC resources have to be (are, actually) included into WLCG; yes, CEs are needed because there is hardly any other way to include HPC machines; yes, batch systems are needed because there is no HPC machine without such; and we even do volunteer computing using essentially the "CE-model".

One other aspect which is quite acute for us is federated storage vs networking: we are the only Tier1 which has internally federated storage, so we know very well that it is very taxing on networks. And if it is taxing on networks, it is taxing on taxpayers' pockets who pay for research networks. So this subject is also of interest -

how to balance the insatiable VO requirements with finite bandwidth 😊

Tier-1

1. virtualisation and cloud. What does it really mean for wLCG? The ideas on this are all over the map. At one end of the spectrum we keep wLCG like it is, it's just all the WNs are now virtual instead of bare metal; nothing changes at all. The other end of the spectrum is that we completely drop the idea of submitting jobs, we submit virtual machines, or perhaps even submit requests for instantiation of virtual clusters at the sites, and the experiments have their own software for managing what happens on the virtual clusters.
Moving more towards the latter model brings new challenges for shared sites (and hence for all of wLCG). The original intent of the grid was to have a shared infrastructure, the concept is built in and hence in practice is not a problem. Sharing is also not a problem for a cloud when it is only 30% full, however sharing is not so obvious when the system is 100% full. wLCG-dedicated sites will have no problem in such a world, however shared sites like many of the european Tier-1s will have to find a way of dealing with the new situation. Without some good model and associated developments, the amount of opportunistic resources available to wLCG will decrease.
2. wLCG support model. Partially linked to the above, as support for the "give me a virtual cluster" is going to be much different than support for the current model. Aside from this, the support model used by the experiments and by wLCG has several issues:
 1. basic support by sites. Not all sites answer GGUS tickets, upgrade middleware releases, reply to security challenges, fix their broken storage elements, etc. My impression is that this costs wLCG and the experiments quite a bit of effort, it's not clear to me whether the amount of computing we get from such sites justifies the effort.
 2. support demands by experiments to sites. Experiments have linked support manpower with computing resources in many ways that only (in some cases) makes political sense (ie funding). Most of what a so-called "site person" or "Tier-1 liaison" does is debug experiment code; the only link to the specific site is that the experiment code happened to have failed at that site. In the worst cases, one has many site support people debugging the exact same failures at many separate sites. If we move to a more cloudy world this becomes even more so; virtual clusters are identical from site to site, the only things that remain site-specific are experiment-independent issues like VM performance, networking bandwidth, allocation management etc. People are increasingly harder to fund, work is needed, mostly within the experiments, to separate true site support (things that really inherently have to do with a site) from experiment distributed computing operations; and also to continue to improve experiment frameworks to facilitate debugging.
3. convergence of our AA infrastructure with the rest of the world. AARC is supposed to "do this" as far as I can tell; I think wLCG should engage with these folk to help them keep their eyes on the ball. Security folk sometimes drift off into an abstract version of the world "can we talk to them about a few concrete use cases, and how / when they see these being supported? The use cases need to be very concrete: something like "a professor at the university of utrecht wants to submit an ALICE analysis job without having to deal with grid certificates; using only his utrecht university credentials (username / password), how can this be achieved? The AARC people would illustrate how to do this, by having actual product names on the boxes in the diagrams instead of abstract concepts.
4. Meeting structure. Not a big issue but important. We have operations meetings, ops coordination, middleware readiness, grid deployment board, management board, overview board, collaboration board, HEPiX, WLCG workshops, CHEP "why do we have all these meetings? Can we reduce the number and frequency? It gives us time to do some more work, and have more interesting things to talk about (and less status reports ;-)) when we actually do meet in person. btw my opinion is that this many-meetings phenomenon has become possible due to ubiquitous wireless & laptops. 20 years ago I used to get really angry if a meeting was unproductive, as it was keeping me from doing work. Now if a meeting is unproductive, I have the option to tune out and continue working on something else.

Security

I would simply propose "The future of academic computing security", covering our trust model, roles and responsibility, security organisation, and interactions with the rest of the world.

USATLAS

- Efficient use of processing in all its forms, in particular opportunistic. In ATLAS this focus has led to developing the Event Service with work underway to deploy it pretty much everywhere.
- Efficient data management: minimal use of storage in favor of maximal use of networking. Leveraging web approaches/technologies like CDN. Intelligent dynamic data delivery, favoring caching and streaming over placement. Uniformity/consolidation in data handling protocols & tools (xrootd + http). Object stores, supporting fine grained data flows and ease of data access.
- In both of the above, emphasis on the 'flow' in workflow and dataflow. Fine grained, agile, highly automated, informed by rich dynamic information systems
- smooth (transparent, elastic) integration of cloud resources including particularly commercial clouds. (Commercial clouds have arrived; Amazon spot market resources are now slightly cheaper than BNL Tier 1 procured resources, including (the small) associated storage costs, by a recent metric from M. Ernst). Bringing commercial resources in as a full player in the mix, evaluating them against in-house facilities with their associated costs. Not only bulk processing but also services like DBs, analytic platforms, messaging etc.
- bottom-up rather than top-down model for developing common distributed software, and leveraging open source/web tech as much as possible in doing so. How can WLCG/dist sw make use of HEP Software Foundation, or alter it to make it useful.

OSG

1. The future of the information system, or what is the minimum information we need to make scheduling decisions given that all submissions are pilots?
2. Delegated security, or can we do without enduser level certificates?
3. relevance of WLCG beyond the LHC
4. The LHC last mile problem towards publication, and the relevance of computing beyond T0,T1,T2.

The first two items are reasonably obvious. So let me comment on the 2nd two.

I am detecting a shift in opposite directions in EU and US, and think it is worth pointing this out to people. In the US, what we do is increasingly gaining relevance outside the LHC community, to a point where I wouldn't be surprised if in 5 years the LHC is less than 50% of what we do. Already today we are at ~60% or so.

This influences where we are focusing, and how things get done in that the distributed infrastructure is becoming increasingly more distributed, and more diverse. E.g. 20% of OSG is physics other than LHC. 10% is life sciences. 10% everything else in science.

I don't know the situation in the EU and Asia, and would like to hear about it. Because if its anything like the US then it begs the question of what is WLCG and what should it be in the future? Take Mu2e, Nova, XENON1T, LZ, IceCube, sPHENIX, Polarbear, glueX as examples for physics experiments that all are experiments OSG serves today, and will do so increasingly in the future, but are not part of the WLCG. Some of these have European stakeholders (Nova, XENON1T), most probably don't. Are there similar issues in Europe and Asia? Should we rethink by expanding what we think of as WLCG's scope, like we do now with Belle? And what does this mean for the future technology discussions? What does it mean for federation? Central management? Central control?

In the US, we increasingly talk to CIOs at Universities. Are there similar discussions happening in EU or

Asia? Does WLCG have a mandate or interest to look at University computing broadly worldwide? Should we? Or is it better for WLCG to be a club of national labs that serve the LHC? I.e. stick to a smaller scope, and focus on doing well there.

At the same time, in the US there is a strong push towards being able to use supercomputers, new architectures, etc. This leads to tension and challenges in dynamic range from giga to exascale that OSG needs to deal with. Does anybody in EU or Asia have similar issues of dynamic range?

Summary of input

Topic	Popularity	Examples
WLCG in general		
WLCG scope and responsibilities	+++	Organization, allocation of resources, decision making, including non-LHC experiments
Funding	++	Plans to get EU funding, how to provide funding to software projects essential to WLCG
Collaboration with other communities	+	Industry, Openlap, other EU/US projects
Pledges and MoU for non-traditional WLCG resources	+	Make pledging compatible with allocations at supercomputing centres (HS06*hours rather than HS06*cores)
Commercial clouds	+	Evaluate commercial cloud resources as full players and their costs compared to in-house facilities (also for databases, analytics platforms, messaging)
WLCG support model	++	Clear separation between basic site support and experiment support, liaisons, extension of GGUS to non-Grid sites
Security		
Security and identity federations	++	Send jobs without an X509 certificate
Security for non-pledged resources	++	Fitting HPC centres in the security framework
Access to computing resources and provisioning		
Resource provisioning policies	+	Defining, implementing, monitoring and verifying resource provisioning policies
Running on HPC facilities and other opportunistic resources	++	Learn from NorduGrid (which uses a lot of HPC resources), Event Service
Clouds and resource sharing	+	Understand how to run on clouds on shared, finite resources
Site simplification and small sites	+++	Understand how to reduce the services to be run and how to integrate small clusters offering volunteer resources, diskless centres
Evolution of the CE+batch system approach	++	Current and future alternatives to this model, new middleware
Virtualisation, clouds and containers	+++	Virtual WNs managed by the site vs. virtual clusters instantiated by the experiments vs. container technologies, transparent and elastic integration of cloud resources
Data management and federated storage		
Status and future of storage federations and federated storage	++++	Usefulness, medium term evolution (see Montenegro workshop), integration of xrootd/http federations
New technologies for storage management	++++++	Impact of cloud storage, storage-less setups, event/object stores, changes in storage model, implications on network, intelligent data delivery, CDN, caching and streaming vs. placement
Storage protocols	+++	How to reduce the amount of required storage protocols
Networking		

IPv6 deployment	+	Discussion on timescale and impact
Monitoring, accounting and information systems		
Combined network-storage-transfer metrics	++	Aggregation of data from perfSONAR, FTS monitoring, SE status information
Data analytics on infrastructure	++	Application of Machine Learning techniques to understand operations/resource usage at sites/experiments/global WLCG level
Monitoring tools and non-LHC VOs	+	Extending FTS and perfSONAR monitoring to other VOs (e.g. Belle II)
Review of SAM tests	+	To make sure that they still fulfil the needs of the experiments and evolve SAM as needed
Information System	++++	Understand its future and what is really needed from it to make scheduling decisions, how it could be simplified, how to use it for service discovery and SAM tests, use it for high automation of workflows and dataflows
Accounting for non-pledged resources	+	Fitting HPC centres in the accounting framework
Accounting data quality	+	Better validation, limited but clear set of agreed functionalities
Running conditions and resource models for Run3 and beyond		
Experience in Run2	+	Reports on valuable experiences, e.g. the LHCb Turbo stream
Run3-4 discussions	+	Updates on Run3-4 requirements
New architectures	+	Push towards supercomputers and new architectures, causing tension and challenges in dynamic range from giga to exascale
HEP software		
Development of common software	+	Bottom-up approach, leveraging open source/web technology, how WLCG can make use of the HSF
Benhmarming		
CPU benchmarking	+	Draw conclusions and decide how to move forward

-- AndreaSciaba - 2015-09-23

This topic: LCG > WLCGWorkshop2016

Topic revision: r21 - 2015-11-10 - AndreaSciaba



Copyright © 2008-2015 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback