

Facility Data Challenge Planning

Shawn McKee / U Michigan

December 2, 2022

At the WBS 2.3 Face-to-face, SLAC

<https://indico.cern.ch/event/1201515/contributions/5141071/>



Introduction

We are in the midst of preparing for the next WLCG “Network” Data Challenge, having had numerous meetings and discussions (LHCONE/LHCOPN side meeting, WLCG meeting in Lancaster, WLCG BDT meetings, WLCG DOMA meetings, etc)

The question we as a facility need to consider is where we will focus our participation.

While the date is not yet fixed, it will likely be February/March 2024.

The consensus is that we need **interim** milestones and mini-challenges to ensure we are ready...



Topics from the WLCG Network DC Side Meeting

The side meeting (after the LHCONe/LHCOPN) was targeted towards identifying **milestones** and **mini-challenges** to prepare for the next data challenge.

- Notes are available at

<https://docs.google.com/document/d/1zWZMR5U6-nhX1Zo8u-TR0Nt9L3sUiTyQAo9c9Qsjts/edit>

However the discussion evolved in interesting ways and we didn't end up finishing many concrete steps or timelines.

The notes are useful to set the stage for our discussions today.

We should have some time this morning to discuss any of the topics and determine what action items and next steps are needed.

Discussion about DC24

During the side meeting after the LHCONE/LHCOPN, we discussed DC24 topics:

- Clarifying the Target DC24 Details
- Monitoring
- Transfer tools and applications
- Storage configuration and use-cases
- Network Capabilities, SDN and Orchestration
- Scope, experimental coordination, involvement of VOs outside of WLCG, HEP?

Goal was to create a table highlighting milestones or mini-challenges needed:

Milestone / Mini-challenge	Target Dates	Description
Milestone	March 2023	Get perfSONARs upgraded (to 5.x?) and operational across most important WLCG sites.
Capability	March to October 2023	Define and then instantiate performance metrics at all participating organizations for data challenge 2.
SDN milestones		Steps to get Rucio/SENSE SDN stuff working at sites willing/able to incorporate in DCs to do this

Network Topics from the DC24 Discussions

Some of the important network related areas we need to work on for DC24:

- Update and utilize **perfSONAR** to clean up links before DC24.
- Enable **packet marking** and **flow labeling** for a significant portion of DC24.
- Instrument and document **site networks**, for at least our largest sites.
- Plan for network orchestration using **Rucio/SENSE** for a set of production sites
 - We will need a set of sites who are interested in testing things out using their production systems.
- **Network planning**: we need to make sure our sites and their local and regional networks are aware of our requirements and timeline and are planning appropriately
 - May need to extend what ESnet has been doing in the US to other regions
- **IPv6** should be enabled everywhere not just because of packet marking, but because it will allow us to get back to a single stack sooner!

My Takeaways from the Discussion

For DC24, sites should **NOT** prematurely purchase equipment just to meet the target and the target (30%) is likely to be revised.

The “Flexible” target involving our Tier-2s is the relevant target (x2 from base)

Planning and executing on a set of quarterly milestones or mini-challenges in specific areas will be **crucial** for our success.

Since data challenges are meant to evaluate our capabilities when all WLCG experiments are operating at the same time, it may also be important to involve other VOs that will have a large global footprint to also participate.

Part of the value of DC24 will be to engage with sites, regional, national and international networks to ensure the **planning** and **timelines** for upgrades are well matched to our needs and as cost effective as possible.



Site Expectations

Rob had asked about any documentation listing Tier-2 requirements or expectations for the data challenges.

We can start to develop some informations by answering some questions:

- Networking: What is needed for the next data challenge?
- Networking: What are the logical times to plan for capacity / capability upgrades?
- Networking: Are sites in touch with their administration, campus network teams and regional network providers?
- Computing: What capacity will sites have by February 2024?
 - What network connectivity will compute nodes have at a site?
- Storage: What capacity will sites have by February 2024?
 - What fraction of the capacity will be usable for (mini)data challenges?
 - What is the typical streaming rate for storage servers at a site? (What is the range?)

Possible USATLAS Mini-Challenges / Activities

Some ideas for US focused mini-challenges / activities (and needed milestones)

- Run VP as the main queue at one or more US Tier-2s to explore its behavior and limitations.
 - This could represent a new mode of operation and one we may want to try at as large a scale as we can configure
- Develop and test Hiro's "load testing" framework to push the US sites and networks to identify bottlenecks and weak points
 - Need an easy way to select sites and data, introduce controllable loads, monitor the additional activity and be able to turn off or rollback the added load if the production infrastructure starts to fail.
 - Work on tuning for network, storage and ATLAS applications used by the framework.
- Get site network monitoring (from the WLCG Monitoring Task Force) in place at all our Tier-2s and the Tier-1 (currently only at CERN & AGLT2)
- Deploy packet and flow monitoring for our Tier-1 and Tier-2 sites

USATLAS Table Capturing Our Interests


I have copied the table from the data challenge discussion at CERN to a USATLAS specific document:

<https://docs.google.com/document/d/1wip-zf4uFmpEdWRTKliogUKBXTdLAp57V2Im9eKK6Sg/edit?usp=sharing>

The facility R&D document we have all been adding to should be an excellent source of topics we could integrate into the data (mini)challenge and milestones.


<https://docs.google.com/document/d/1uqHqygcY6Dcg1amPBZclH0d1x-5kBNvaxjc99UxviXU/edit#>

Please add your ideas to the table so we can capture possible activities for USATLAS related to the planning, mini-challenges and milestones for WBS 2.3



Questions, Comments, Discussion...

Let's discuss our plans....



Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- [OSG: NSF MPS-1148698](#)
- [IRIS-HEP: NSF OAC-1836650](#)



Relevant URLs

- OSG/WLCG Networking Documentation
 - <https://opensciencegrid.github.io/networking/>
- perfSONAR Infrastructure Dashboard
 - <https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7>
- perfSONAR Dashboard and Monitoring
 - <http://maddash.opensciencegrid.org/maddash-webui>
 - https://psetf.opensciencegrid.org/etf/check_mk
- perfSONAR Central Configuration
 - <https://psconfig.opensciencegrid.org/>
- Toolkit information page
 - <https://toolkitinfo.opensciencegrid.org/>
- Grafana dashboards
 - <http://monit-grafana-open.cern.ch/>
- ATLAS Alerting and Alarming Service: <https://aaas.atlas-ml.org/>
- The perfSONAR Dashboard application: <https://ps-dash.uc.ssl-hep.org/>
- ESnet WLCG Stardust Dashboard:
<https://public.stardust.es.net/d/XkxDL5H7z/esnet-public-dashboards?orgId=1>

Network Topic: Network Capacity Planning

Network capacity planning, both for **sites** and **network providers**

- Are sites aware of HL-LHC networking requirements?
 - Need to incorporate the requirements with an appropriate timeline: too early is too expensive while waiting too long risks crippling the usability of the site; also supply chain, implementation delays also need to be considered.
 - We need to make sure all our sites ARE aware and planning appropriately.
- Are campus network teams & administrators incorporating LHC needs into their planning?
 - It is very important that everyone involved is aware of the requirements and timescale so they can appropriately plan and time their upgrades.
- Are regional networks aware of both HL-LHC and site requirements?
 - If sites upgrade to new higher capacities, it has implications for their regional networks (as well as the R&E backbone providers).
 - Does the regional network plans mesh with those of its sites?

