DC24 Planning and Near Term Activities

Shawn McKee / University of Michigan

WLCG DOMA General Meeting

https://indico.cern.ch/event/1299111/

July 5, 2023

WLCG Data Challenges

The **WLCG Data Challenges** are a ~biennial series of four increasingly-complex exercises which started in 2021 and are aimed at demonstrating readiness at the HL-LHC scale.

Next data challenge is targeting **25%** of HL-LHC scale and includes T1/T2

Such exercises are only meaningful if there is worldwide, multi-experiment participation; they are planned and managed by this group, WLCG DOMA (Data Organization, Management and Access, see DC24 update from Feb MB mtg).

These data challenges provide many benefits, allowing **sites**, **networks** and **experiments** to evaluate their progress, motivate and validate their developments in hardware and software and show readiness of technologies at suitable scale.

Context and Overview

- There is a significant amount of possible work foreseen to fully take advantage of the upcoming WLCG Network Data Challenge 2024 (DC24)
  - The data challenge should not be viewed as just reaching (or not) a specific metric, but rather as an opportunity to demonstrate and evaluate new capabilities at scale.
  - **Beneficial** capabilities then get folded into our production infrastructure, improving things for the next data challenge…
- WLCG DOMA needs to provide the coordination, vision and visibility
  - What is the big picture for DC24: current status, planning and ideas?
  - What are the critical things that need action?

Planned DC24 Work Areas

As part of USATLAS facility planning they have identified some areas of work for Data Challenge 2024 (DC24)

- perfSONAR site network debugging prior to DC24
- Enabling WLCG Site Network load monitoring
- Deployment of traffic marking for data transfers
- Demonstration of packet pacing at two or more sites
- Network load testing to find bottlenecks

These are relevant for WLCG as a whole and each of these requires effort and planning to ensure we can benefit from the DC24 process

WLCG/LHCONE Identified Activities

In addition to what has been planned within USATLAS there are a number of areas intended to demonstrate feasibility and benefit from WLCG and LHCONE

- GEANT flow monitoring (enabling user visible R&E flow monitoring)
- NOTED (traffic steering across multiple available links)
- P4 packet analysis (use of programmable switches to track/account-for traffic)
- Routing adaptations for commercial collaborators (cost optimization via R&E networks)
- SENSE Rucio/FTS/XRootD interoperation testing (UCSD, Caltech, FNAL, others)
- Use of Tokens for (ALL?) Data Transfers (ATLAS, others?)
- Integration of new analysis capabilities in parallel with DC24

What Should We Do ASAP?

The WLCG DOMA group is organizing DC24 activities by providing visibility for activities proposed by groups, sites or experiments and roughly defining targets and timelines.

- We are NOT imposing tasks, goals or activities from above

This means that those interested in demonstrating useful capabilities need to be self-organizing.

I propose we solicit interested parties to define: **Objective for DC24, Who & What/When** for their areas of interest and provide a clear location and mechanism for sharing the results.

Template for DC24 Projects

Description: provide a brief paragraph of what the project will do during DC24 including planned benefits, scope and how to evaluate success.

Timeline: Provide one or more **dates** where a specific goal, mini-challenge or deliverable will be completed AND **who** will be tasked with ensuring the work is completed.

Metrics: Describe one or more metrics that will be gathered to enable evaluation of the work and its benefits and the source of the metrics (existing or to be created, etc.)

Example: Flow/Packet Marking

Description of objective for DC24: We intend to mark a significant amount of WLCG network traffic via flow marking (UDP Fireflies) and/or packet marking (for IPv6 traffic). Significant means at least 5% of traffic during DC24 for two or more participating VOs. The goal is providing visibility of all R&E traffic anywhere in the network, giving experiments feedback to optimize net-use and networks info/control

Timeline: We will need to have mini-challenges to develop the capabilities and metrics required

- July 2023: We will demonstrate marked flows from two or more production sites using dCache and Xrootd
  - **Who**: Shawn McKee, Marian Babik
- July 2023: We will show initial accounting of flows via the ESnet Stardust system.
  - **Who**: Andy Lake, Shawn McKee
- July 2023: Flowd available in EPEL for EL8/EL9
  - **Who:** Steve Traylen, Marian Babik
- September 2023: We will demonstrate marked IPv6 packets from two or more production sites using dCache and Xrootd.
  - **Who**: Garhan Attebury, Marian Babik
- September 2023: We will demonstrate the use of packet accounting of marked IPv6 packets.
  - **Who:** Carmen Moreira, Edoardo Martelli, Yatish Kumar(? Or other ESnet), Jeronimo Bezerra(?)
- December 2023: Identify and demonstrate traffic volume accounting by VO and activity globally and at one or more net locations.
  - **Who**: Marian Babik, Andy Lake, GEANT

Metrics: During DC24 we intend to track

- The amount of traffic that is marked by packets (MarkedPacketTotal [TB]) and by flow (MarkedFlowTotal[TB])
- The number of UDP fireflies sent (estimated) and received (#Sent, #Recieved)
- The total size of traffic, broken out by owner and activity (by flow or by packet) at one or more network locations (ESnet, GEANT?, Internet?) and globally via the ESnet Stardust receiver. Estimate December 2023

Items for DOMA

I think DOMA should provide a well defined **location** and **set of tools** to organize and make visible the activities proposed by various groups.

- It is very important that we all understand the proposed activities, goals and objectives to allow collaboration

Can we clarify what happens in DC24 to better understand how activities plug in to the process?

- Is there a global framework to run the DC24 activities and what is the foreseen mechanism to coordinate activities and experiments?

We need the Tier-2 sites and target contributions defined **ASAP**

- Naively, scale T2s to T1s by storage+compute values?

Summary

We (WLCG) need to **clarify** and **document** existing plans, mini-challenges and goals.

We are only ¾ of a year away from DC24 and many of the current activities have **significant** work still to do (on top of our day jobs).

We have an **opportunity** to leverage DC24 to drive technology deployment, to show value and to demonstrate capabilities at scale.

Question, Comments, Discussion?

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

perfSONAR Site Network Debugging

This activity is intended to be completed BEFORE DC24.

Involves: WLCG Network Throughput WG, IRIS-HEP fellows and participating sites.

- Do we have volunteers to help run this for ATLAS sites?

Rough timeline:

- **Summer 2023**: Upgrade to newest perfSONAR, upgrade hardware for toolkits, verify site perfSONARs work
- **Fall 2023**: Utilize infrastructure (perfSONAR, alerting, analytics) to find and fix network issues (most egregious first)
- **Winter 2024**: Final debugging pass on largest sites

Enabling WLCG Site Monitoring

Objective: Provide site network monitoring (total In/Out traffic) every minute, gathered by CERN Monit.

WLCG Monitoring TF has begun a campaign to get this enabled at the largest sites by Fall 2023

Involves: WLCG MonTF and larger WLCG sites can benefit from understanding site network loads. We need to verify proper deployment, configuration and continued operations going forward.

Packet Pacing “Demonstrator”

The Research Networking Technical WG (RNTWG) has shifted its focus from packet/flow marking to packet pacing

Objective: To pace data transfers to better utilize our networks, preventing microbursts and increasing overall throughput.

Goal: Have one or more sites pacing their traffic during DC24

Involves: RNTWG, Participating sites, ESnet, LHCONE. Looking for additional interested participants.

Network Load Testing / Bottleneck ID

Hiro will provide the details in the next talk but this effort is designed to identify the current and evolving data transfer capacities of our sites and their infrastructure.

Involving: Hiro, ATLAS sites outside the US? CMS?

Important for us to benchmark where we are for DC24 and critical for identifying end-to-end bottlenecks.

Additional Areas?

ATLAS as a whole is trying to enable storage tokens for production data transfers for DC24.

- Should our sites take a role in this?

What about using our Analysis Facilities (AFs) as a component of DC24?

- Can we identify and enable new methodologies and applications that can improve user experience?
- Can we show them at scale during DC24? Who?

Other areas of interest (RUCIO/Sense, Comm Cloud access)?