ATLAS Network Data Challenge 2024 (DC24): Results and Plans for DC26

Shawn McKee / University of Michigan

5th Global Research Platform Meeting
(https://grpworkshop2024.theglobalresearchplatform.net/)
September 16, 2024

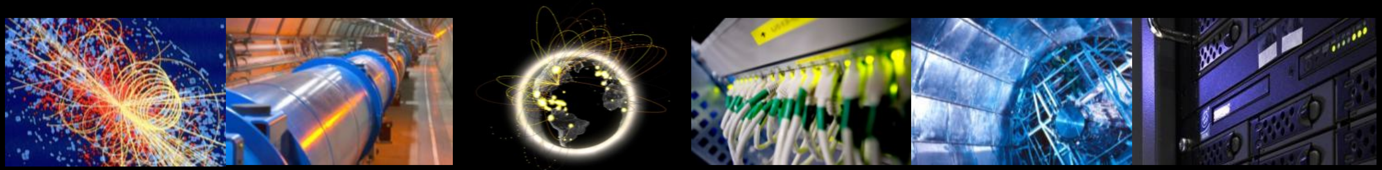

WLCG
Worldwide LHC Computing Grid

Open Science Grid

Introduction

I am a Research Scientist at the University of Michigan Physics Department working on the ATLAS project.

My roles within ATLAS include being the Distributed Data Management co-coordinator as well as the USATLAS Facilities and Distributed Computing co-manager.

This presentation represents **my** ATLAS perspective on the WLCG Data Challenges but may not represent official ATLAS perspective (I am speaking from my own experience working on and preparing for data challenges for ATLAS).

Please feel free to ask questions as I go or save them up for the end of the talk.

ATLAS, LHC and the WLCG

For those not familiar with High-Energy Physics, I want to quickly introduce ATLAS, the LHC and the WLCG:

ATLAS (http://atlas.cern/) is the largest of four particle detectors that measure and record the particle collisions at the LHC. The primary scientific goal is to quantitatively measure and discover properties of the Standard Model (SM) of particle physics.

The **L**arge **H**adron **C**ollider (**LHC**) at the European Laboratory for Particle Physics (CERN https://home.cern/) is the most powerful particle accelerator in the world. Highly energetic protons, traveling almost at the speed of light around a 27 kilometer long ring in both directions, are steered to collide head-on, creating new particles and new interactions to probe fundamental natural laws.

The **W**orldwide **L**HC **C**omputing **G**rid (**WLCG**) (https://wlcg.web.cern.ch/) collects resources worldwide and enables their usage by the LHC experiments as a distributed computing facility. WLCG is co-ordinated by CERN. WLCG is managed and operated by a worldwide collaboration between the experiments (ALICE, ATLAS, CMS and LHCb) and partners with EGI (European Grid Infrastructure), OSG (Open Science Grid), and NeIC (Nordic e-Infrastructure Collaboration).

ATLAS and LHC Running Organization

The LHC schedule since its startup is organized into long periods of particle collisions (called **Runs**) and periods of repair/maintenance (called Shutdowns). Between Runs are **Long Shutdowns** for upgrades and replacements

A new schedule is under discussion which may extend the current **Run-3** and move **HL-LHC** to 2030

We will know more later this fall…

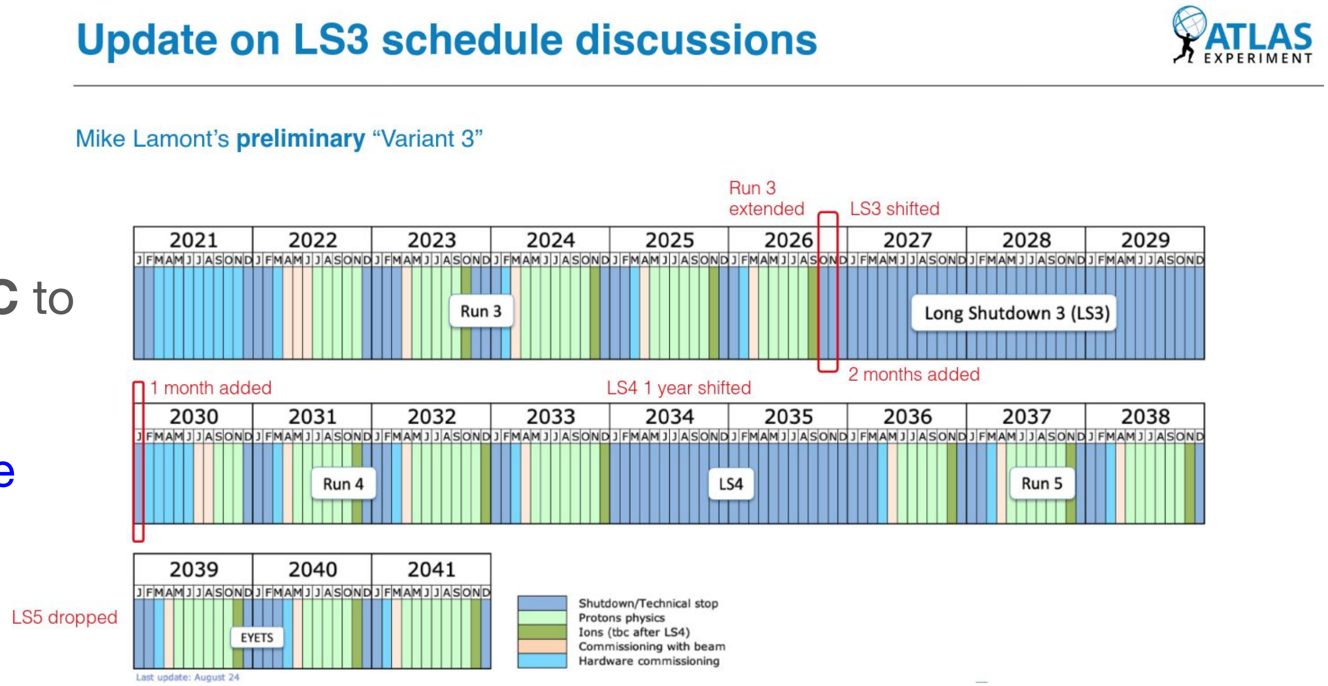

High-Energy Physics Resource Challenges

The WLCG experiments generate **~200 Petabytes** of raw data yearly from their detectors and additionally generate a similar volume of simulated and transformed data, which must be tracked, managed and made globally available.

- Since WLCG resources (computing, storage, services, etc) are also globally distributed, networking becomes a critical component for being able to pursue the scientific goals of the experiments.
- Additionally the resource requirements, in terms of the amount and type of computing accessible and the volume of storage, continue to increase while budgets remain (so far) flat.

This means WLCG experiments, like ATLAS, need to be innovative to do more with the resources we can afford and need continual development and improvements to meet our future needs.

WLCG has planned a series of **data challenges** to evaluate how our infrastructure is developing to meet future needs and focus effort on identified bottlenecks.

Open Science Grid
WLCG Worldwide LHC Computing Grid

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 (**DC26** or possibly **DC27**) targets **50%** of HL-LHC scale and includes T1/T2 and any improvements we can integrate into our infrastructure. (**DC24** from February this year was **25%** of HL-LHC)

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.

I believe it is critical to fully participate in future challenges, both by preparing and testing before each and analyzing the results after each.

DC24 Total Throughput

WLCG Throughput

Flexible target (green line) was to be met for 48 hours

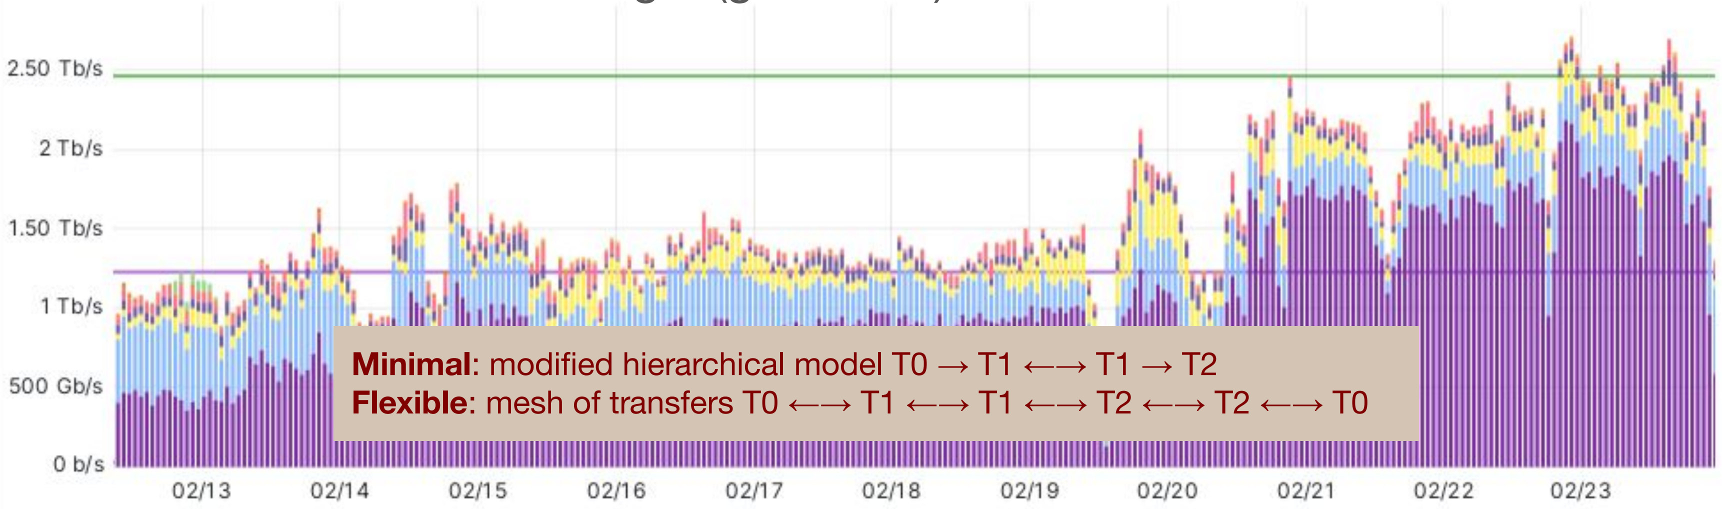

Minimal: modified hierarchical model T0 → T1 ←→ T1 → T2
Flexible: mesh of transfers T0 ←→ T1 ←→ T1 ←→ T2 ←→ T2 ←→ T0

	max	avg	current
Data Challenge	2.19 Tb/s	1.03 Tb/s	587 Gb/s
atlas	608 Gb/s	298 Gb/s	547 Gb/s
alice xrootd	349 Gb/s	114 Gb/s	43.8 Gb/s
cms xrootd	191 Gb/s	66.1 Gb/s	40.2 Gb/s
cms	271 Gb/s	57.0 Gb/s	73.8 Gb/s

Annotated ATLAS DC24 Throughput

DC24 met the baseline target, but not the flexible one

Wide-area Networking was NOT the bottleneck

FTS was one of the primary limitations

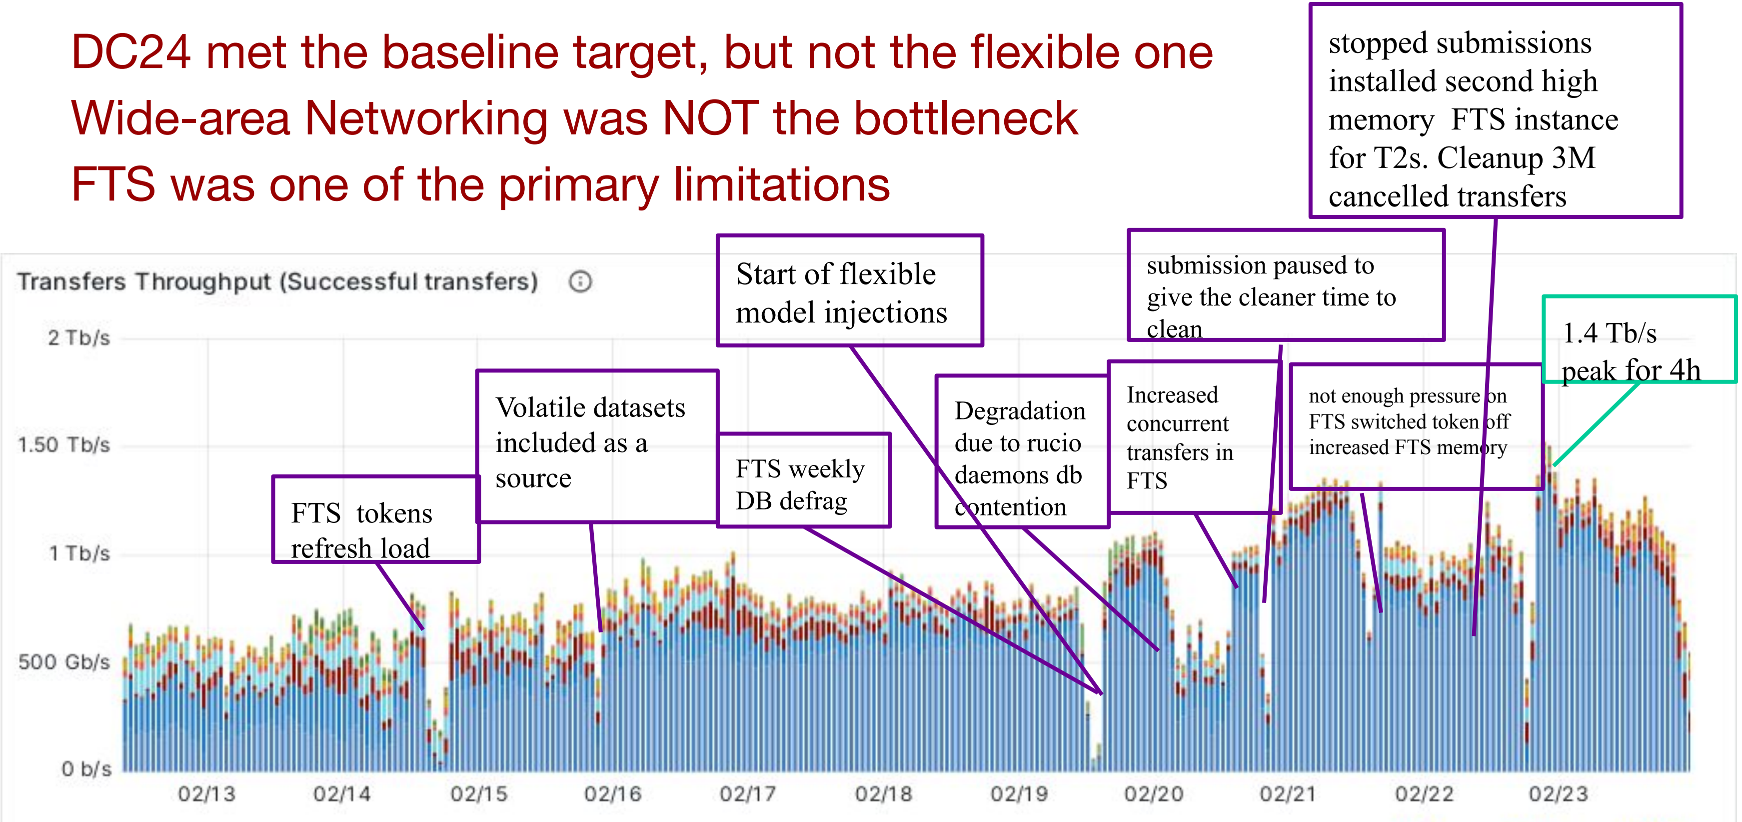

What Have We Learned from Data Challenges?

The two WLCG Network Data Challenges we have run so far have been very helpful for the experiments, sites, R&E networks and technology innovators.

- **Sites** were able to identify bottlenecks that were not obvious before
- **R&E networks** were able to gain understanding about how the various participants data flows might interact with each other across their topology.
- The **WLCG experiments** and partners were also able to identify where bottlenecks in software, services and architecture exist
- **Technology proponents** were able to do “at-scale” testing to inform software and service evolution

The following slides show some of the results from DC24 in more detail

DC24 Dashboard By Experiment

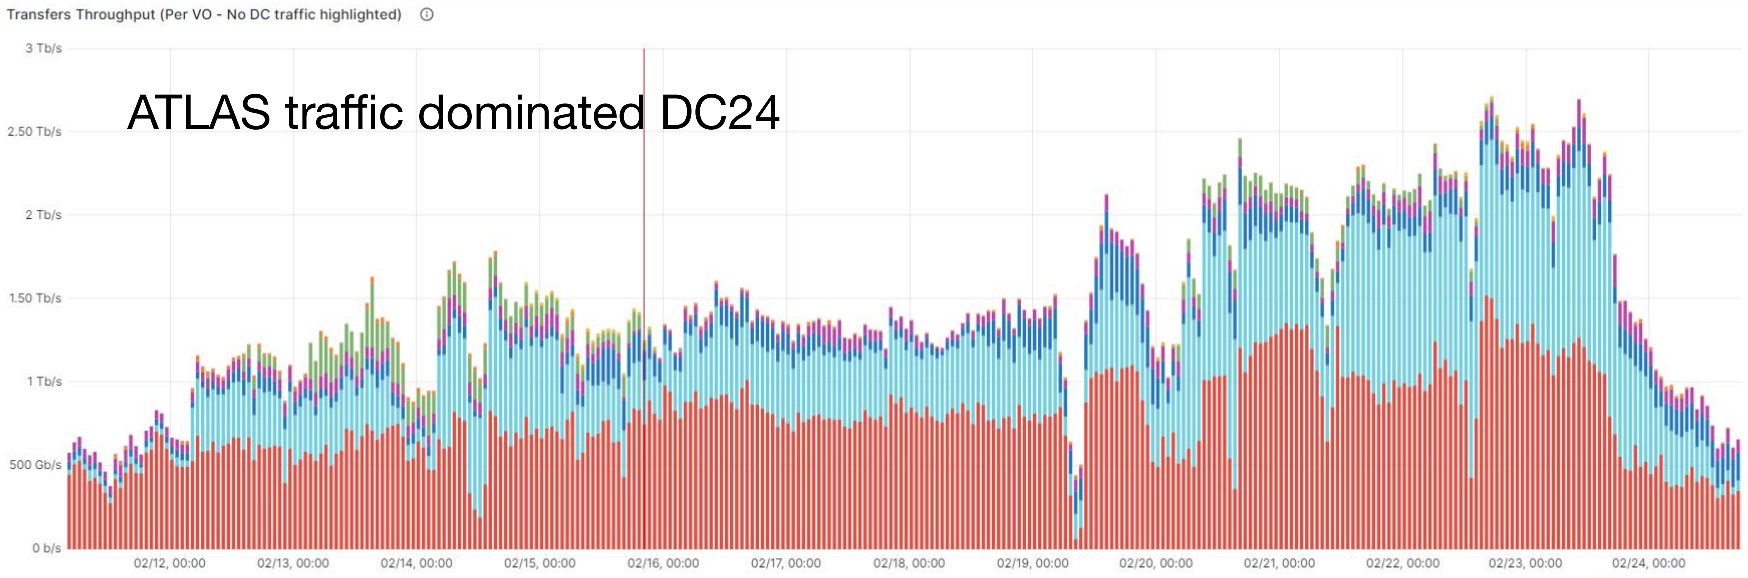

	max	avg	current
atlas	1.52 Tb/s	780 Gb/s	370 Gb/s
cms	1.16 Tb/s	455 Gb/s	39.8 Gb/s
alice xrootd	349 Gb/s	111 Gb/s	22.7 Mb/s
cms xrootd	191 Gb/s	67.3 Gb/s	0 b/s
lhcb	382 Gb/s	43.2 Gb/s	300 Mb/s
belle	38.9 Gb/s	9.04 Gb/s	1.20 Gb/s
dune	28.6 Gb/s	5.36 Gb/s	273 Mb/s

Open Science Grid

WLCG Worldwide LHC Computing Grid

LCG

DC24 by Country or Tier

Source Countries

	total	percentage
USA	22 Mil	23%
Switzerland	17 Mil	17%
UK	8 Mil	9%
Germany	8 Mil	9%
Italy	6 Mil	7%
France	6 Mil	7%
UNKNOWN	4 Mil	4%

Destination Countries

	total	percentage
USA	19 Mil	20%
Switzerland	11 Mil	12%
Italy	11 Mil	12%
UK	8 Mil	9%
Germany	8 Mil	9%
France	6 Mil	6%
Sweden	5 Mil	5%

Source Tiers

	total	percentage
1	36 Mil	37%
2	35 Mil	37%
0	16 Mil	16%
-1	4 Mil	4%
-1	3 Mil	3%

Destination Tiers

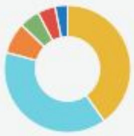

	total	percentage
2	38 Mil	40%
1	37 Mil	39%
0	8 Mil	8%
3	5 Mil	5%
-1	4 Mil	5%

US delivered our % as Source, not quite as Destination

Tier-2s matched Tier-1s for number of transfers as source or destination

WLCG Site Network Monitoring

For DC24, a site network monitoring campaign was undertaken to provide better visibility into each site's capabilities (see CERN Gitlab)

- This was a result of DC21 noting a deficiency in our monitoring
- For DC24 we just wanted total IN/OUT for each site
- For DC26, we may want to improve the level of detail (traffic by experiment)

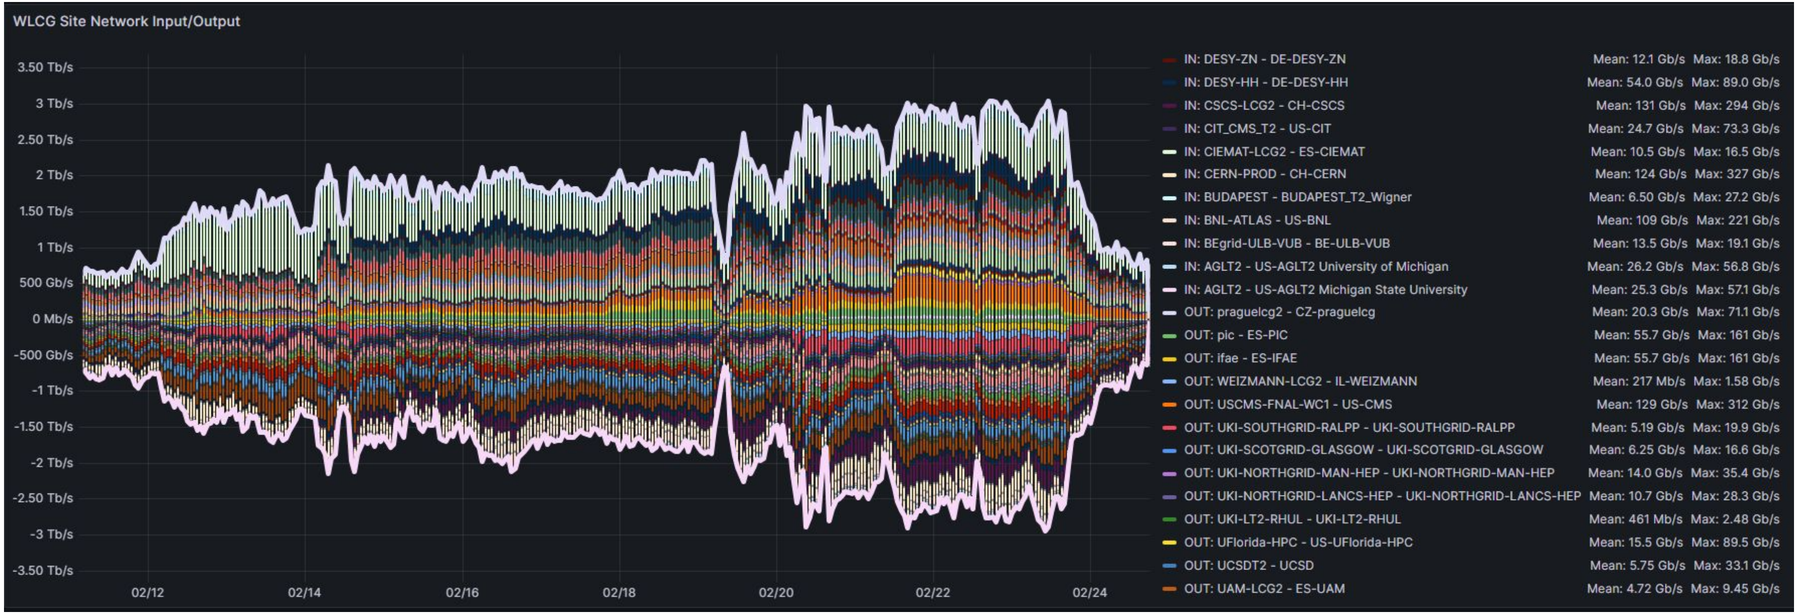

Transatlantic Links (ESnet Stardust Monitoring)

These performed well despite pre-DC24 concerns, with peaks up to 0.5 Tbps

Total capacity during DC24 on ESnet provided links was 0.9 Tbps

[ESnet Monitoring]

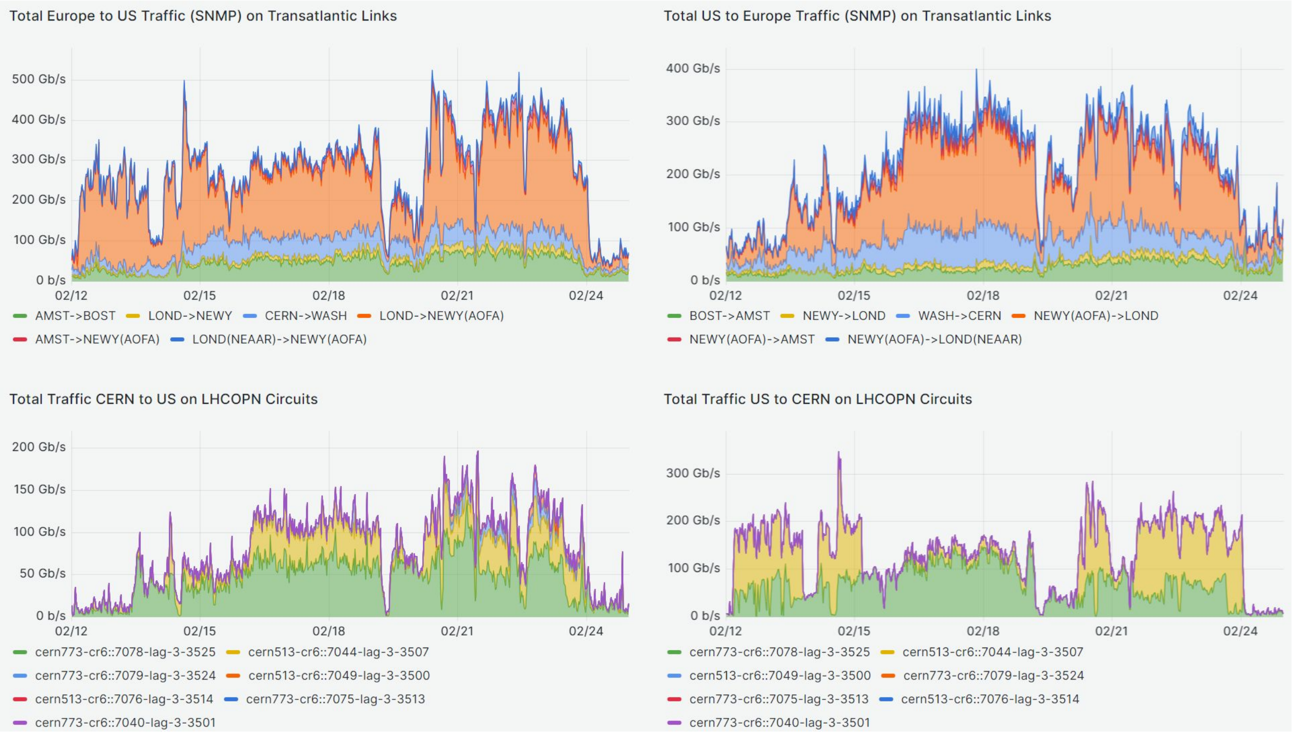

Network Load Testing / Bottleneck ID

We had a very successful test within the US **prior** to DC24

Joint US ATLAS and CMS network challenge

- BNL was able to push data out to MWT2 and AGLT2 at up to ~400 Gbps.
- When the reverse direction data flow (into BNL) was added, there was a drop in the original direction (out of BNL) after some delay.
    - Disk IO issue at source and/or destination?

USATLAS DC24 Overview

How did DC24 go for USATLAS?

- In general our sites performed well BUT we did identify issues
- **AGLT2**: Not stressed but prior testing showed 80Gbps bottleneck at UM
- **MWT2**: Some issues with dCache performance being investigated
- **NET2**: Some storage/transfer issues encountered. Only had 10 Gbps
- **SWT2**: Problems with network capacity, issues with OU deletion and storage performance, missing site network monitoring

We are continuing to analyze how our sites and software performed, especially while monitoring systems still have the data at finest granularity.

We also tested new technologies but not uniformly across the whole facility (more in 2 slides)

Overall ATLAS DC24 Issues Noted

- Low T0 Export to T1 throughput needs to be understood (DDM Monitoring, DC24 Monitoring, LHCOPN Network)
- FTS (our File Transfer Service) is NOT throughput aware and slow transfers end up dominating the capacity
- MANY site issues
  - Slow deletion at RAL - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165358
  - Worsened NDGF-T1 - https://ggus.eu/index.php?mode=ticket_info&ticket_id=164846
  - Timeouts to Milano - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165356
  - SSL errors to CNAF - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165355
  - Timeouts to FZK - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165393
  - **Timeouts at OU_OSCER_ATLAS - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165362, https://ggus.eu/index.php?mode=ticket_info&ticket_id=165379**
  - Expired tokens in the FTS causing problems at DESY-HH - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165397
  - Timeouts at UKI-SCOTGRID-GLASGOW - https://ggus.eu/index.php?mode=ticket_info&ticket_id=163552
  - "Unexpected server error" to NIKHEF during pre-DC24 test - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165263
  - "Unexpected server error" to UKI-NORTHGRID-LANCS-HEP - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165394
  - Dark data caused by heavy load at TRIUMF-LCG2 - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165343
  - IFIC tickets is not blaming DC24 but errors stopping with the end of it - https://ggus.eu/index.php?mode=ticket_info&ticket_id=165395
  - IN2P3-CC being overloaded and HC putting the site into test for lack of storage free connections
    - Cured by reducing the number of connections in FTS
  - **SWT2 large wave of jobs in transferring state (concurrent with a wave of evgen jobs)**
  - FZK QMUL slowed down transfers
  - INFN-T1 - (one) stuck doors while retrieving JWK from IAM https://ggus.eu/index.php?mode=ticket_info&ticket_id=165355
  - Some sites applied storage limits tuning, e.g. FZK (GGUS:165393), TRIUMF (GGUS:165364)…, there were few more and not all communicated with GGUS => for final report we should also ask sites what they observed (e.g. SARA internal? throughput saturated their links (GGUS:165359), INFN also observed huge traffic (GGUS:165355), we don't fully understand much higher throughput on some links,

Technologies Tested for DC24

We had a number of "capability" tests that were run as part of DC24

- We tested the **Scitags** (see https://www.scitags.org site) firefly capability and ESnet provide a global receiver and associate monitoring.
- ATLAS and CMS tested "**storage tokens**" for some periods during the data challenge and found limitations in the implementation (too many tokens were needed for the tested granularity; a token issuer limitation)
- CMS tested **SENSE/Rucio** (SDN coupled to data management) components involving UCSD, Caltech, Nebraska and FNAL
- CERN tried some **NOTED** (splitting large flows to other paths dynamically) tests during DC24
- We planned to do tests related to traffic optimization via Jumbo Frames, new protocols and packet pacing but didn't manage to get it in time.

What is a “mini-challenge”

For ATLAS, we found great benefit in the pre-DC24 testing we undertook and realized that having easy to use tools to run “mini-challenges” on demand would be very powerful.

What do we mean by “mini-challenge”. Here is a possible definition:

- A mini-challenge is a **lightweight** way to test capacity or capability with one or more sites

The goal is to make it easy to test and track both our capabilities and capacities, finding and fixing bottlenecks, identifying bad architectures and hardware and improving our visibility into how our sites perform as part of a globally distributed infrastructure.

- What is critical is that this should NOT require any expert involvement which currently prevents on-demand testing to be easy to do.

Goals for DC26

For DC26 (or DC27 if it moves later) we are targeting:

- All sites should be moving the majority of their data via **IPv6**
- We should have a few **IPv6-only** sites for each experiment
- At least 80% of the traffic should be **identified via SciTags**
- At least 50% of the traffic should be using **jumbo frames**
- **Rucio/SENSE** to be used by few Production sites
- Sites should be able to easily utilize **90% of their declared WAN bandwidth** for an extended period (many hours to days)
- **Network traffic monitoring** should be able to track throughput by network type (LHCOPN, LHCONE, Research & Education, Commercial/Commodity)

All of these areas could benefit from continual testing via mini-challenges

Transforming our Sites

The data challenges provide us with an opportunity to evaluate our existing hardware, software and architecture to identify bottlenecks, limitations and misconfigurations.

Given that HL-LHC is ~6 years away, now is the perfect time to re-evaluate our site's hardware configuration and architecture so that we can have a suitable baseline ready for HL-LHC requirements.

- Six years of hardware purchases can fully replace our current hardware
- Incrementally transforming sites should allow a smooth transition in capability

It is **critical** that sites understand how they fit into our globally distributed infrastructure so they can meet the HL-LHC requirements and use-cases.

- Mini-challenges are a great opportunity to understand our current capabilities, identify bottlenecks and prototype new technologies.

Transforming our Sites (cont'd)

An example of the above is the case of Caltech (CMS site) that went from performing at 80Gbps during DC24 to ~250Gbps by analyzing the results from DC24 and upgrading and tuning their system

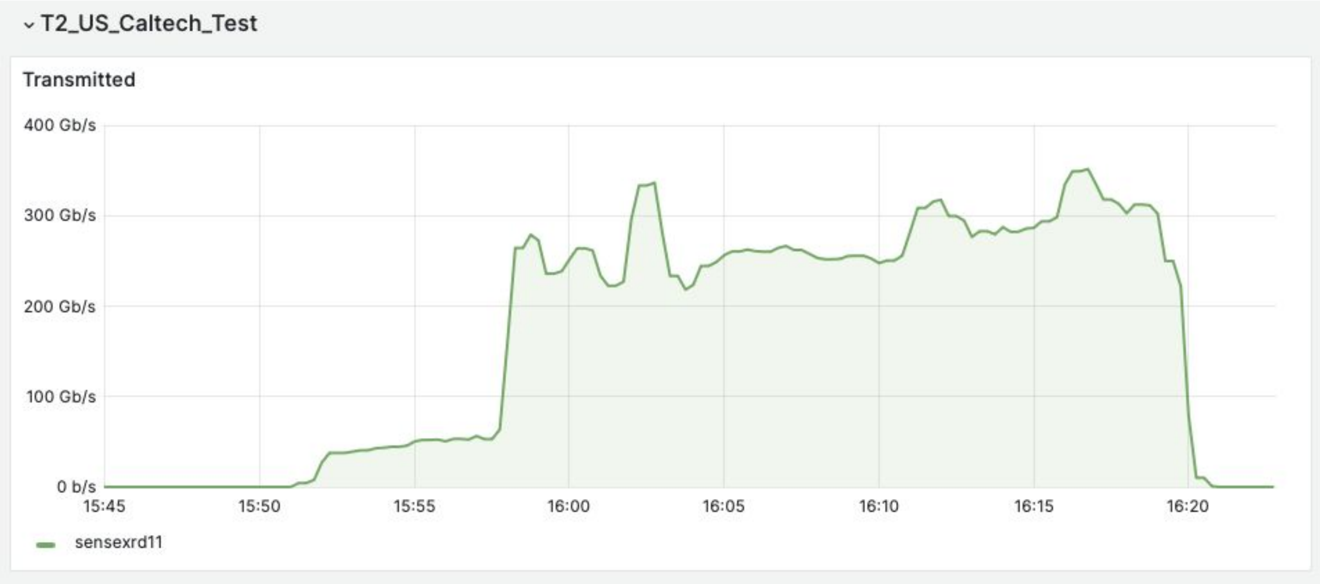

Caltech (Prod) => UCSD (test) HTTP-TPC transfer throughput

Status & Plans for Scitags

The Scitags Initiative has been underway for a while with a goal of allowing easy traffic identification anywhere in the network.

- We have an IETF draft to define what we are doing

For DC24 we had a few sites running **xrootd** and **EOS** storage infrastructure sending out “fireflies” which identified their traffic.

- Fireflies are UDP JSON-formatted packets targeting destination port 10514 which contain the experiment (owner) and activity as well as the associated flow information (src-ip,src-port,dst-ip,dst-port, protocol)
- We also have targeted IPv6 packets to contain the same information in the “flow-label” field and next steps will be getting that production ready
- There will be associated demos at SC24 coming in November
- **dCache** is our next targeted storage infrastructure to enable

Preparing Technologies and Capabilities

HL-LHC will require more resources than we can currently afford.

- To address this, the experiments are working hard to optimize workflows
- New technologies and capabilities will play a critical role in bridging the gap

The WLCG data challenges are designed to regularly test where we are relative to where we need to be for HL-LHC.

Possible technologies to test and, if beneficial, integrate

- **New / improved storage servers** (Gen5 PCIe, NVMe, new NICs, etc)
  - Define/document LHC server best practice for hardware and configuration
- **SciTags** (traffic identification anywhere in the network)
- **Traffic optimization** (via Jumbo Frames, pacing, new protocols)
- **Network Orchestration** (SENSE/Rucio, NOTED, GNA-g efforts, etc)
- Improvements (alternatives) to **WebDAV and Xrootd protocols**
- Improvements to **storage elements** (dCache, Xrootd, STORM, EOS, etc)
- Evolution of **Distributed Data Management** (Rucio, FTS, etc)

Mini-Challenges and Ongoing Testing

As noted before, the mini-challenges prior to DC24 turned out to be very beneficial for finding problems in our infrastructure and we should plan to have regular mini-challenges going forward

- The injection tool, used for DC24, is relatively easy to use. We could/should use it to run periodic tests with individual sites e.g. CERN to US sites help them understand possible limitations
- Hiro Ito / BNL has developed a load tester usable without “experts”
- Regional sets of sites (up to “North America”) should be tested simultaneously by USATLAS/USCMS to verify we don’t conflict at PoPs
- Mini-challenges should also include tests of technologies and new capabilities

Ongoing mini-challenges a few times a year provide important guidance and validation for site changes in hardware, software and tunings.

Open Science Grid

WLCG Worldwide LHC Computing Grid

LCG

Testing, Benchmarking and Documenting Tunings

Many of our sites have tuned and documented site software and services but these may have last been tested 5, 10 or more years ago.

Give the drastic changes in operating systems and software over even a few years, it is important to revisit this area.

We are seeking **sites** and **experiment** and **application teams** to benchmark various tunings for many different components in use: widely used applications, operating system, storage, computing and networking.

The **goal** is to create a new, **current** reference set of tunings available (perhaps added to the ESnet Fasterdata pages; maybe augment TuneD).

Near To Mid-Term Plans

Both ATLAS and CMS have plans to begin new testing this Fall.

- USATLAS wants to find the new “maximum” between its Tier-1 at BNL and the various Tier-2s
- CMS and ATLAS in the UK are planning a regional mini-challenge at the end of November beginning of December.
- ATLAS would like to undertake some capacity based mini-challenges early in the new year and capability based challenges in Summer 2025)

One important area of work is better mini-challenge infrastructure which should allow any of the sites or technology proponents to run tests on demand.

Summary

We need to **clarify** and **document** existing plans, mini-challenges and goals for the next year and for DC26/DC27

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

This work I have described can benefit from and contribute to the larger community, e.g., for a "Global Research Platform"

Question, Comments, Discussion?

Acknowledgements

Thanks to Justas Balcas, Diego Davila and Asif Shaw for their contributions to the slides

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:

- IRIS-HEP: NSF OAC-1836650 and PHY-2323298

Backup Slides

Background Material

Here are some resources we know about:

Presentations

- WLCG Data Challenge 2024 (DC24) Status and Plans Related to ATLAS DDM (Jun 2023)
- DC24 Planning and Near Term Activities (Jul 2023)
- USATLAS Data Challenge 2024 Take-aways (Feb 2024)
- Medium to Long Term Network Plans for ATLAS and CMS (Mar 2024)
- DC24 Network Activities & Results (May 2024)

Some Google Docs

- WLCG/DOMA Data Challenge 2024: Final Report
- USATLAS Milestones/MiniChallenges for Next WLCG Data Challenge in 2024
- Planning Mini-Challenges for US ATLAS Facilities and Distributed Computing
- NOTES: USATLAS Facility Status and Evolution Discussion

DC24 Links

Official DC24 report
https://zenodo.org/records/11402618

DC24 Network Activities and Results:
https://docs.google.com/presentation/d/1s0VvbXEpj1PN9umFT8wgsHsHmG9EYucymbalKNrvuKQ/edit#slide=id.p1

Katy Ellis LHCONE/LHCOPN DC24 presentation:
https://docs.google.com/presentation/d/1Tm3pCMkfHj5KHTW3PXbgS7mdHf72Ir27qr1JgMbrnRg/edit#slide=id.g1ea89411ecb_0_4

Next Steps Towards DC26:
https://docs.google.com/presentation/d/1mMx6QaihWJWpbVEQgxNjZXRT5_s4SkBTXu0SpELtuvI/edit#slide=id.gd170caf633_1_0

DC24 ATLAS Retrospective:
https://docs.google.com/presentation/d/1Lh_D57BvWn13AFClhhuczm-j-tKV-yMez_oD4yYUtBo/edit#slide=id.gd170caf633_1_0