ATL-SOFT-PROC-2025-028
11 February 2025

ATLAS WLCG Data Challenge 2024 planning and implementation

Alessandra Forti[1,*], *Petr* Vokac[2], *Mario* Lassnig[3], *Dimitrios* Christidis[3], *Shawn* McKee[4], and *Vincent* Garonne[5] on behalf of the ATLAS Computing Activity

[1]School of Physics and Astronomy, University of Manchester, Oxford Rd, Manchester, M139PL, UK
[2]Czech Technical University in Prague, Faculty of Nuclear Sciences and Physical Engineering, Břehová 7, Prague, 115 19, Czech Republic
[3]European Organization for Nuclear Research (CERN), 1211 Genève 23, Switzerland
[4]University of Michigan: Ann Arbor, MI, US
[5]Brookhaven National Laboratory: Upton, NY, US

Abstract. ATLAS is participating in the WLCG Data Challenges, a bi-yearly program established in 2021 to prepare for the data rates of the High Luminosity HL-LHC. In each challenge, the transfer rates are increased to ensure readiness for the full rates by 2029. The goal of the 2024 Data Challenge (DC24) was to reach 25% of the HL-LHC expected transfer rates, with each experiment deciding how to execute the challenge based on agreed general guidelines and common dates. The ATLAS challenge was designed to test the ATLAS distributed infrastructure across 66 sites and was carried out over 12 days, with increasing rates and more complex transfer topologies, putting significant strain on the system. It was also the first time the new OAuth 2.0 authorization system was tested at such a large scale. This paper will discuss the planning of the challenge, the tools used to execute it, the agreed-upon transfer rates for the connections, and finally, the achieved results and any unrealized goals, along with an analysis of the bottlenecks. We will then describe how the challenge itself was executed, the results obtained, and the lessons learned. Finally, we will look ahead to the next challenge, currently scheduled for the end of 2026 or early in 2027, with 50% of HL-LHC rates.

1 Introduction

The High Luminosity LHC (HL-LHC) increase in data output poses interesting questions for the WLCG [1] infrastructure and its capacity to absorb it. In preparation for this, the infrastructure has to prepare both storage and network to absorb the data, as well as the central data management, transfer and authorization services to sustain the throughput. For this reason, WLCG in 2021 started a series of bi-yearly data challenges with progressively increasing rates.

In our planning documents, the DC rates are expressed in two models: one called minimal based on the original WLCG hierarchical topology with data flowing from Tier0 to Tier1s and from Tier1s to Tier2s; and one called flexible which allows a full mesh of data transfers,

*e-mail: Alessandra.Forti@cern.ch

flowing from every site to every other site independently from the tier level, resulting in twice the rates as the minimal model. Tier0 export of RAW data from Tier0 to Tier1s is at the core of both models rates calculations and it is the most important transfer activity.

The minimal model represents the backbones of the transfers, but ATLAS has been using a flexible topology for years since the Tier2s connectivity and reliability became stable enough to make them run data intensive applications and not only simulation. For this reason, the flexible model rates are now taken as the targets for each data challenge.

- T0 export: T0 → T1
- Minimal model: T0 → T1 → T2
- Flexible model: T0 ←→ T1 ←→ T1 ←→ T2 ←→ T0

The 2024 Data Challenge (DC24) is the second one listed in the planning document [2]. It builds on the experience of the first data challenge (DC21) [3] which ran at 10% of the HL-LHC rates, where DC24 ran at 25% (originally at 30%). The most important difference between between the DC21 and DC24 are the rates but most of all the contribution of the production traffic to the challenge and the amount of traffic ATLAS had to inject to reach the targets. In DC21 the production traffic contributed 60% while in DC24 it contributed 36%. This, of course, affected the way ATLAS ran the challenge and put a real strain on the central services and some of the sites as we will discuss later in this paper.

2 Goals and metrics

ATLAS divided the challenge into three parts, injecting progressively more data into an increasingly complex mesh of network links as described in section 5. The rates were calculated for every link taking into consideration the following aspects

- HL-LHC Data Challenges rates in the DC21 planning document
- Each site bandwidth usable by ATLAS
- Active (injected rates) and passive (only production traffic) contributions.
- Number of TB of disk space necessary at sites to sustain the rates
- Number of required deletions to keep the space free for following injections
- Number of concurrent transfers needed to sustain the rates

In Figure 1 we report the numbers for the Tier1s. A similar larger table was produced for the Tier2s. These numbers were used to calculate the traffic per link which the tools that generate transfers needed in input.

Table: DC24 (src)	Site WAN (Gb/s)	Common to all scenarios	DC24 minimal scenario			DC24 flexible scenario			FTS active inbound / outbound
	Usable by ATLAS	T0 Export	Total Gb/s & bandwidth		Space [TB/24h] (deletions/hour)	Total Gb/s & bandwidth		Space [TB/24h] (deletions/hour)	
Site			∑ ingress	∑ egress		∑ ingress	∑ egress		
CERN-PROD	891	257.0	23.4	282.5	246 (3505)	88.9	392.8	937 (13330)	454 / 2037
BNL-ATLAS	400	60.0	84.5	67.1	892 (12681)	119.8	124.9	1263 (17964)	719 / 851
FZK-LCG2	144	32.0	55.9	35.5	590 (8386)	92.9	65.5	980 (13939)	473 / 410
IN2P3-CC	177	38.0	59.8	43.0	631 (8976)	93.5	77.7	987 (14032)	543 / 429
INFN-T1	62	23.0	36.3	26.0	383 (5447)	61.2	46.1	645 (9177)	230 / 209
NDGF-T1	149	15.0	44.6	23.3	471 (6692)	95.6	33.7	1009 (14345)	593 / 106
SARA-MATRIX	238	15.0	31.0	16.4	327 (4650)	60.1	30.2	634 (9020)	164 / 139
pic	85	11.0	17.1	12.5	181 (2570)	29.0	20.9	306 (4355)	141 / 150
RAL-LCG2	177	38.0	64.7	40.3	683 (9709)	92.8	81.0	978 (13915)	1595 / 663
TRIUMF-LCG2	100	25.0	38.2	27.8	402 (5723)	60.0	50.9	632 (8996)	322 / 434

Figure 1: ATLAS DC24 estimated rates

At global level the ultimate goal was to sustain 1.4 Tb/s for several hours, ambitiously for 48h. In addition, ATLAS tested the new authorization Oauth2.0 infrastructure based on Tokens [6].

3 Method to generate the traffic

DC24 traffic was the sum of production traffic and traffic injected using the same infrastructure used by production and using real datasets. It is important to note that we wanted to test the whole chain of data transfer and not only the network.

We used dc_inject, a script that creates Rucio [4] rules to replicate data with a certain lifetime according to the rates per link calculated above. The tool was very stable and for the next challenge, an improved version will be incorporated into Rucio directly to make it extremely easy for sites to adopt it for smaller tests ("mini-challenges") between the planned WLCG challenges.

To generate the rates per link using real data, it is important to select unique data so that transfers from different sources do not interfere with each other copying the same file. It is also important to be able to select the average file size for the transfers because the number of transfers may change by several orders of magnitude depending on the average file size. Much of the origin data in the selected threshold (avg 5-7 GB size, 1000 files) was from temporary datasets used by production, which were automatically removed as production computational tasks progressed. This meant eventually the input datasets list had to be refreshed every 12 hours to keep the number of available fresh datasets big enough for our injection rates.

A big constraint in running a challenge of this size using real data is the limited amount of space available at sites. This forced DC24 to be executed with a really fast cycle of injections and deletions which, combined with the large number of links, has pushed the whole infrastructure quite hard.

ATLAS had 66 sites participating to various degrees: the Tier0, 9 Tier1s and 56 Tier2s. Because of this, the max number of links to inject with extra traffic was 1200 (2000 if we add only the smaller sites participating only with production traffic). These 1200 links were injected continuously at intervals of 15 minutes to ensure sustained traffic. During DC21 ATLAS tried much longer intervals but this produced waves of traffic rather than sustained.

Datasets lifetime tested 1h, 2h and 3h. Short lifetimes determine how long the data can stay on the storage. After the end of the lifetime, the data is deleted if the space is needed. When the lifetime of the datasets was set at 3h, the space at some sites ran out because not enough data was deleted; at 1h too many transfers were canceled while still in the queue, so eventually 2h was used for most of the challenge.

ATLAS ran the DC24 transfers in backfill mode to ensure the production traffic was prioritized with respect to the data challenge one and minimize the disruption. However this had some consequences on the DC24 rates at Tier1s: the internal traffic from disk to tape pushed the DC24 one out and didn't count in the final results because it was internal.

4 DC24 operation, achieved rates

The challenge, as agreed in WLCG / DOMA, ran for 12 days, and ATLAS split in three parts increasing the number of links to inject and using a more complicated model (described in section 1) at each stage.

- 2 days only T0 export (9 links)
- 5 days minimal model (350 links)

- 5 days flexible model (>1200 links)

The results per day for each T0 and T1 and for the sum of T2 are reported in 2 where it is possible to note 3 things:

- Some Tier1s were not performing, particularly in ingress. This will be addressed in section 6
- The increase in the number of injected links at high rates in the second week degraded the rates of both Tier1s and Tier2s. The reasons for this will be explained in section 5.
- T0 looks also very red and orange but it is a reflection of the first two points which we will also explain in section 5.

Day	Scenario	BNL-ATLAS		FZK-LCG2		IN2P3-CC		INFN-T1		NDGF-T1		pic	
		dst	src	dst	src	dst	src	dst	src	dst	src	dst	src
1	T0 → T1	25.68	N/A	29.76	N/A	35.6	N/A	21.84	N/A	12.56	N/A	10.48	N/A
2	T0 → T1	35.1	N/A	13	N/A	41	N/A	23.52	N/A	9.79	N/A	14.5	N/A
3	T0 → T1 ↔ T1 → T2	61.6	67.1	47.4	42.2	43.8	39.3	32.1	28	7.72	26.5	18.4	10.8
4	T0 → T1 ↔ T1 → T2	65.3	79.7	61.8	58.5	64.6	47.2	31.8	50.1	4.92	22.7	30.3	15.2
5	T0 → T1 ↔ T1 → T2	63	116	81.3	78.4	75.6	56.6	37.8	52.3	7.59	18.1	32.7	13.1
6	T0 → T1 ↔ T1 → T2	73.7	98.9	85	77.9	71.1	51	39.1	60	4.8	20.2	29.5	21.8
7	T0 → T1 ↔ T1 → T2	65.7	94	79.6	102	63.6	44.8	33.7	69.5	2.2	11.2	33.6	43.8
8	T0 ↔ T1 ↔ T1 ↔ T2	52.8	77.3	59.5	56.5	38.9	50.8	33.7	20	2.99	33.1	24.5	19.1
9	T0 ↔ T1 ↔ T1 ↔ T2	87.9	80.7	51.6	63.6	40.1	34.8	46.1	48.6	2.41	33	39.3	28.8
10	T0 ↔ T1 ↔ T1 ↔ T2	90	95.9	43.7	97.5	39.6	36.8	47.6	50.5	21.9	32.4	54	43.4
11	T0 ↔ T1 ↔ T1 ↔ T2	110	96.8	58.8	82.1	42.1	44.6	55.9	53.4	16.3	44.8	50.7	38.3
12	T0 ↔ T1 ↔ T1 ↔ T2	89.8	84.2	52.4	51.8	34	38.7	64.6	56.4	27.2	67.2	48	38.3

Day	Scenario	RAL-LCG2		SARA-MATRIX		TRIUMF-LCG2		T2 summary		T0 summary			
		dst	src	dst	src	dst	src	dst	src	dst	src		
1	T0 → T1	12.16	N/A	12.64	N/A	19.92	N/A	N/A	N/A	N/A	188		
2	T0 → T1	12.5	N/A	18.9	N/A	24.2	N/A	N/A	N/A	N/A	201		
3	T0 → T1 ↔ T1 → T2	16.7	40.2	34.3	65.3	33.3	27.6	299	141	19.8	141	>90%	
4	T0 → T1 ↔ T1 → T2	25.2	44.7	35.8	92.2	35.5	28.3	346	124	19.6	173	70-90%	
5	T0 → T1 ↔ T1 → T2	23.1	52.2	36.3	89.2	49.2	46.3	387	134	25.9	197	50-70%	
6	T0 → T1 ↔ T1 → T2	27.4	23.6	30.6	95.5	40.9	41.1	337	104	20.3	201	<50%	
7	T0 → T1 ↔ T1 → T2	27.6	20.4	47.2	86.5	53.7	43.4	341	91.7	17.1	190		
8	T0 ↔ T1 ↔ T1 ↔ T2	29.4	47.1	37.7	29.1	37.3	19.9	400	311	54	100		
9	T0 ↔ T1 ↔ T1 ↔ T2	32.3	39.1	59.4	84	51.7	42.7	447	330	89.8	139		
10	T0 ↔ T1 ↔ T1 ↔ T2	43.9	43	92.9	72.3	62.8	52.5	435	337	94.4	97		
11	T0 ↔ T1 ↔ T1 ↔ T2	51.9	56	111	73.8	66.8	42.1	445	406	127	138		
12	T0 ↔ T1 ↔ T1 ↔ T2	72.7	58.8	115	70.8	72.9	31.5	418	407	158	174		

Figure 2: Daily results for each T1, T0 and the sum of T2

5 Central Services

The central services involved in DC24 were Rucio (data management service), FTS [5] (file transfer service), and IAM [7] (token provider service). In this section, we will examine how their behavior affected the rates. Figure 3 shows the several drops in rates caused by one problem or another and the amount of effort required to correct them and keep the rates steady during the challenge.

5.1 rucio

Rucio could handle the extra load quite well until the second week when DB contention between the submitter and the cleaner was identified as causing slow submission rates and a hot patch was applied. A further drop was caused by pausing submission to allow the cleaner to catch up and then submissions were resumed. The amount of Submitters (3×) and Cleaners (4×) was increased to exclude them as a bottleneck, after that everything flowed as expected.

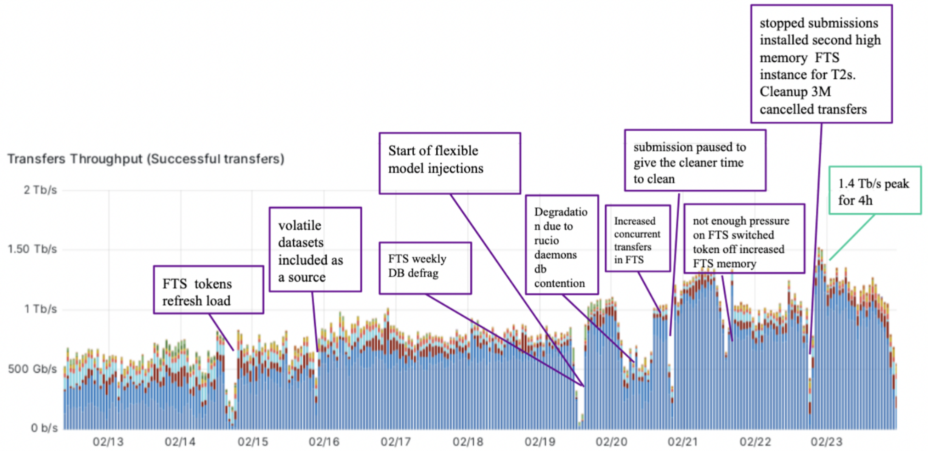

Figure 3: Effects of central services drops

5.2 FTS

The FTS team did an admirable job to keep the ATLAS FTS instances going, tuning and expanding on the fly, and orchestrating a number of transfers that they never reached before. However, the DC highlighted some weak points that need to be addressed. At the time of writing, the FTS team has already done a good amount of development to correct what we have observed and the first release should be in the second quarter of 2025. During DC24 the normal FTS production servers hardware wasn't enough and the FTS team had to upgrade the memory of the production instance and eventually we had to add a second instance to get the Tier2s going. But the major problems were due to FTS orchestrating links, not throughput, and how it prioritizes the transfers.

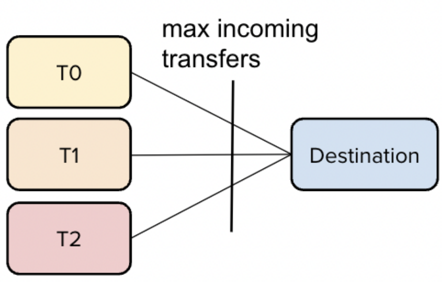

Figure 4: Max incoming transfers per destination

The first problem we encountered was that FTS is a transfer orchestrator, and it optimizes transfers based on the number of transfers per link over many links, it doesn't optimize throughput. The only way to increase throughput is to increase the number of parallel transfers, but, to avoid flooding sites with transfers they cannot absorb, FTS limits the number of transfers per destination and per link as represented in the diagram in Figure 4. These are tunable parameters and ATLAS spent quite some time over the 12 days tweaking them trying to find a balance and maximize throughput. In the future though both the number of transfers and the throughput will need to be managed. The number of transfers is important because ATLAS needs to move many files between many locations, and transfer rates are important because we have to build resilient networks and be able to exploit them.

The second aspect that affected the rates is that FTS has a list of activities with weights and prioritizes transfers according to those weights, but, within an activity, all the transfers are equal. This affected ATLAS in two ways. The first was that since DC24 traffic ran in backfill mode, and FTS manages also internal tape transfers at Tier1s, the DC24 traffic was queued behind waves of tape traffic which wasn't accounted for in the challenge. The second was that within the DC24 activity, the transfers were ordered on a first-in first-out basis. And during the second week, there were thousands of Tier2s slower transfers taking precedence over Tier1 to Tier1, and, even worse, over T0 export to Tier1s. In Figure 5 it is possible to clearly see the degradation of the T0 export rates which had eventually to be rerun in another exercise post DC24. The rates were already not great due to some Tier1s not being able to achieve their rates but the degradation with the increase of links is evident. We could have split the traffic into multiple DC24 activities but the monitoring would have been much more complicated and prioritization according to fast links or according to a manual configuration within an activity is an improvement ATLAS wants to have.

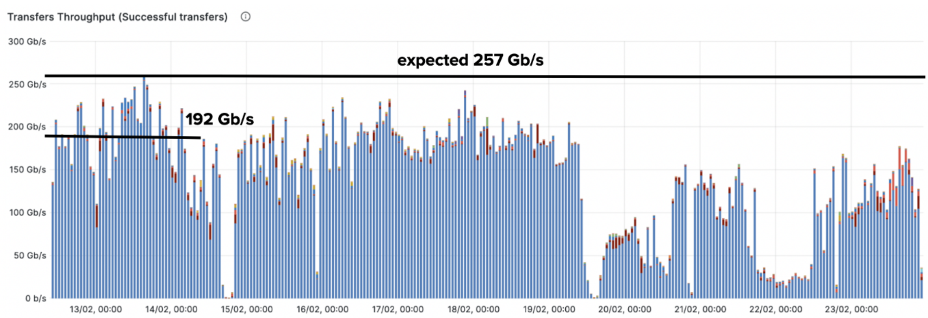

Figure 5: T0 export rates degradation

5.3 tokens

WLCG experiments should move from an x509 AAI infrastructure to a token based Oauth2.0 one. This requires much development on all grid services and for what concerns Data Challenges, Rucio and FTS have released their token based code just before DC24. DC24 was the first time ATLAS did transfers using Tokens for authorization on a really large scale. The result was a partial success. 26% of the transfers were done with Tokens which was really positive, however the use of Tokens slowed down the throughput. The main reason was that Tokens for security reasons must be short lived and need to be continuously refreshed and stored during a transfer and this put an extra load on the FTS DB and the IAM service providing the Tokens which put an extra overhead on the transfers time. In the second week, we switched off token refreshing hoping that 6h would be enough for a transfer to complete but that was optimistic and after we started to have authorization failures we switched completely off the token authorization and progressed only with x509. ATLAS, and the FTS and Rucio teams will have to agree on policies to reduce the number of necessary Tokens.

6 Sites

Part of the bottlenecks ATLAS found were at the sites. Overall 17 tickets were opened about problems caused by the DC24 traffic. Considering the amount of traffic pushed through that was considered reasonable. As we have seen in Figure 6, some Tier1s didn't perform because of storage limitations that affected heavily the rates at these sites. Examples of such problems that had to be corrected and were retested with much improved rates are below:

- One site had to tune its gateways network cards. The tuning they applied is already having positive effects on their current production traffic which occasionally now has sustained higher rates than the data challenge while during the challenge they couldn't go above 74% of the required rates. Their LHCOPN [8] network link was cut during the first week. That was a good test of the reliability of the failover network which kicked in immediately.
- Another site had a bug in dCache which heavily affected the incoming traffic. The bug created also background failures at Run 3 rates and could be ignored, but during the DC24 it crippled the site. This has been fixed after the challenge and the T0 export rates were twice what was expected when we retested.
- A third site couldn't sustain more than 600 connections in ingress and 600 in egress. During the second week this was not enough to sustain all the connections and their rates dropped dramatically. It has already beefed up their storage to correct this problem.
- A fourth site performance was lower than expected despite it could reach really higher rates within the US and in previous tests from CERN. In this case, there were two problems in different stages of the challenge. First a monitor granularity problem. The storage was absorbing data much faster than it was injected. For several minutes between each injection the rates were 0, but the monitoring had 1 hour bin, which averaged down instead of showing a comb pattern. The second problem was a problem of higher latency from CERN which affected the rates and we should have injected more than we calculated to compensate.

7 Mini Data Challenges

ATLAS found DC24 a positive experience. However, running a full DC is a really big effort and the aim is to test multiple things, accordingly it is better to arrive at the full challenge with most problems ironed out and use the DC to "put it all together". The concept of mini Data Challenges (mini-DC) to test single sites when they deploy new technologies or hardware; and to test new features in central services or at any rate different capabilities becomes important. Mini-DCs should be easy to run on demand whenever it is needed. to help with this ATLAS is putting the effort into streamlining dc_inject and incorporating it in Rucio. At the time of writing UK and US clouds already started to run their mini-DC and France is planning theirs for 2025.

8 Conclusions

Overall the challenge was a useful exercise designed to stress all the systems and help us find bottlenecks in the infrastructure not only in the network but also in the storage and in the services, which in the end were the most problematic compared to the network. Steps to solve these bottlenecks were taken both during the challenge and afterward. ATLAS has reached global rates it never achieved before running at 1.4 Tb/s for about 4 hours with a mixture of injected and production traffic as we did during the 2021 Data Challenge (DC21). Figure 6 shows the achieved rates over the two weeks. In DC24, the injected rates were the dominant ones. During the challenge 12 days ATLAS transferred 108 PB, 68 of which were injected. Several participating sites have started more frequent tests to make sure new technologies and hardware they adopt in the meantime don't become a bottleneck at the next challenge.

We want to explicitly acknowledge the support of the National Science Foundation, which supported this work through OSG: NSF MPS-1148698 and IRIS-HEP: NSF OAC-1836650 grants. Furthermore, we acknowledge our collaborations with the CERN IT, WLCG, WLCG DOMA, LHCONE/LHCOPN communities, and all the sites who also participated in this effort.

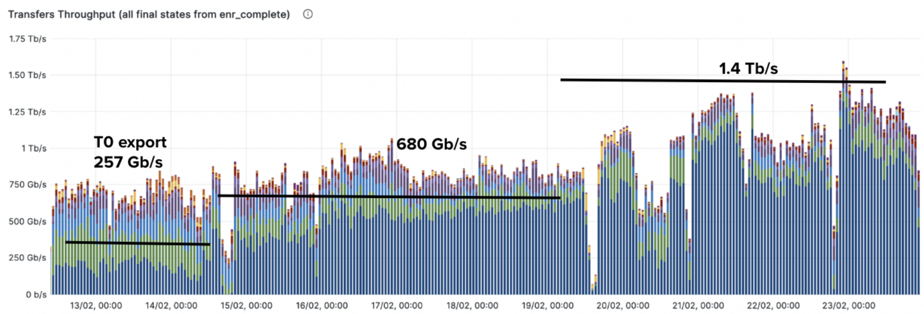

Figure 6: ATLAS achieved rates during DC24

References

[1] J. Shiers, The Worldwide LHC Computing Grid (Worldwide LCG), Comput. Phys. Commun. **177** (2007) 219–223, doi:10.1016/j.cpc.2007.02.021

[2] S. Camapana WLCG data challenges for HL-LHC - 2021 planning https://doi.org/10.5281/zenodo.5532452

[3] A. Forti *et al.*, Network Data Challenges 2021: wrap-up and recommendations https://doi.org/10.5281/zenodo.5767913

[4] M. Barisits *et al.*, Rucio: Scientific data management. Comput. Softw. Big Sci. **3** (2019) 11, doi:10.1007/s41781-019-0026-3

[5] S. Murray *at al.*, FTS Service Evolution and LHC Run-3 Operations. EPJ Web of Conferences **295**, 01031 (2024)

[6] T. Dack *at al.*, WLCG Transition from X.509 to Tokens. Status, Plans, and Timeline. EPJ Web of Conferences **295**, 04054 (2024) https://doi.org/10.1051/epjconf/202429504054

[7] *INDIGO IAM*, https://indigo-iam.github.io/v/current/ (2024), accessed: 2025-02-07

[8] E. Martelli, Evolving the LHCOPN and LHCONE networks to support HL-LHC computing requirements. EPJ Web of Conferences **295**, 07016 (2024), https://doi.org/10.1051/epjconf/202429507016