ATLAS WLCG Data Challenge 2024

planning and implementation

A. Forti, P. Vokac, M. Lassnig,
S. Mckee, V. Garonne, D Christidis
on behalf of the ATLAS Computing Activity
CHEP24
October 2024

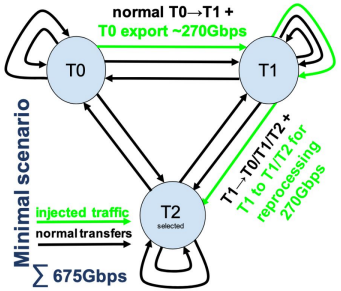

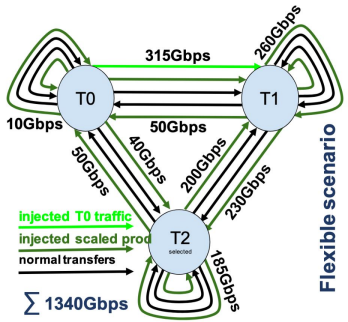

Introduction

- WLCG DC24 as a cooperative effort to optimize WLCG data transfers will be covered in the plenary
- This talk is only about ATLAS results in DC24
  - Primary goal:
    - 1.4 Tb/s aggregate for 48h
  - Secondary goal:
    - Test tokens AAI

DC rates and methods

ATLAS fully flexible

- ATLAS transfers topology is a fully connected mesh
- Large range of file sizes due to different ATLAS activities
  - O(10) kB - O(10)GB
    - Number of transfers is as important as transfer rates
- Two major levels of storage
  - Tape and disk
- 3 independent FTS instances

Large combination of different types of transfer

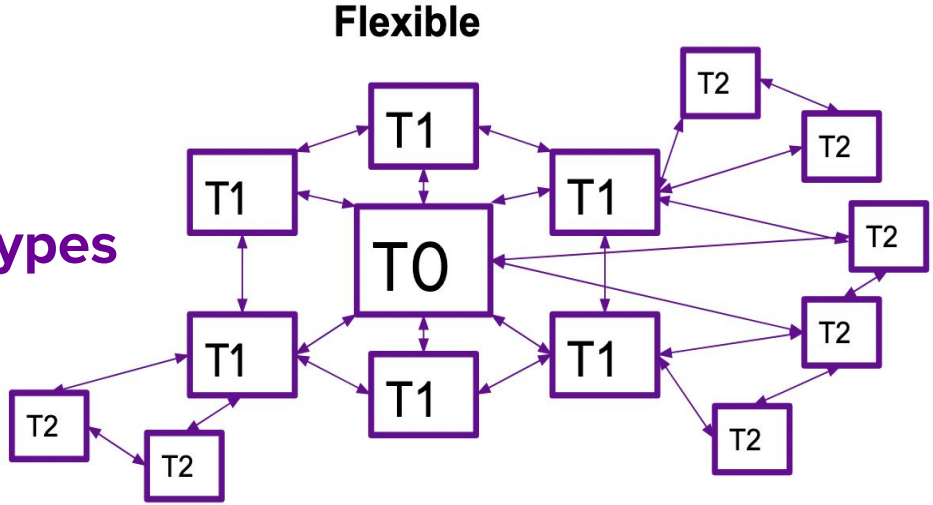

Timeline

- 2 weeks with increasing number of injections and complexity
  - 2 days: T0 ➔ T1 (9 links)
  - 5 days: T0 ➔ T1 ⟷ T1 ➔ T2 (~350 links)
  - 5 days: T0 ⟷ T1 ⟷ T2 ⟷ T2 ⟷ T0 (~1200 links)

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
12/02/2024	13/02/2024	14/02/2024	15/02/2024	16/02/2024	17/02/2024	18/02/2024
T0 → T1	T0 → T1	T0 → T1 ↔ T1 → T2	T0 → T1 ↔ T1 → T2	T0 → T1 ↔ T1 → T2	T0 → T1 ↔ T1 → T2	T0 → T1 ↔ T1 → T2

Monday	Tuesday	Wednesday	Thursday	Friday
19/02/2024	20/02/2024	21/02/2024	22/02/2024	23/02/2024
T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0

- 25% of HL-LHC expected rates

Method

- Challenge design to push the whole system
  - Used production infrastructure:
    - rucio (data management) + FTS (file transfers service)
- Number of sites
  - 66 → 1 T0, 9 T1s & 56 T2s
- Injections every **15 minutes on ~1200 links**
  - ~2000 links including production transfers
  - Pushed FTS really hard to orchestrate
- Short datasets lifetime **1h -> 2h -> 3h** to keep the space free
  - Pushed the deletions rates up
  - Pushed rucio to maintain a balance between submissions and deletions
  - 3h space was running out in some places
- **Data Challenge traffic backfilling**
  - DC just another FTS activity

How we calculate the rates

- For each of the links, we had to calculate:
  - The ingress and egress target rates
  - Taking into account available bandwidth at sites
  - The number of transfers necessary to achieve those rates
  - The number of deletions necessary per hour
  - The average over a period of time is used to compare to the targets.

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

Results

DC24 in a Nutshell

- 107 PB moved in ~12 days 😊
  - avg 0.82 Tb/s
- Only touched target rates 🤔
  - max 1.4 Tb/s for ~4h
- None of the bottlenecks were due to the network specifically
  - FTS and Rucio central services affected the transfers more
    - Almost daily FTS tuning
  - Storage at sites also affected the rates either due to hardware limitations or m/w bugs or tuning

Attempted Transfers	Successful Transfers (%)
44.8 Mil	68.33%
Failed Transfers	**Successful Transfers (vol.)**
14.2 Mil	107.96 PB

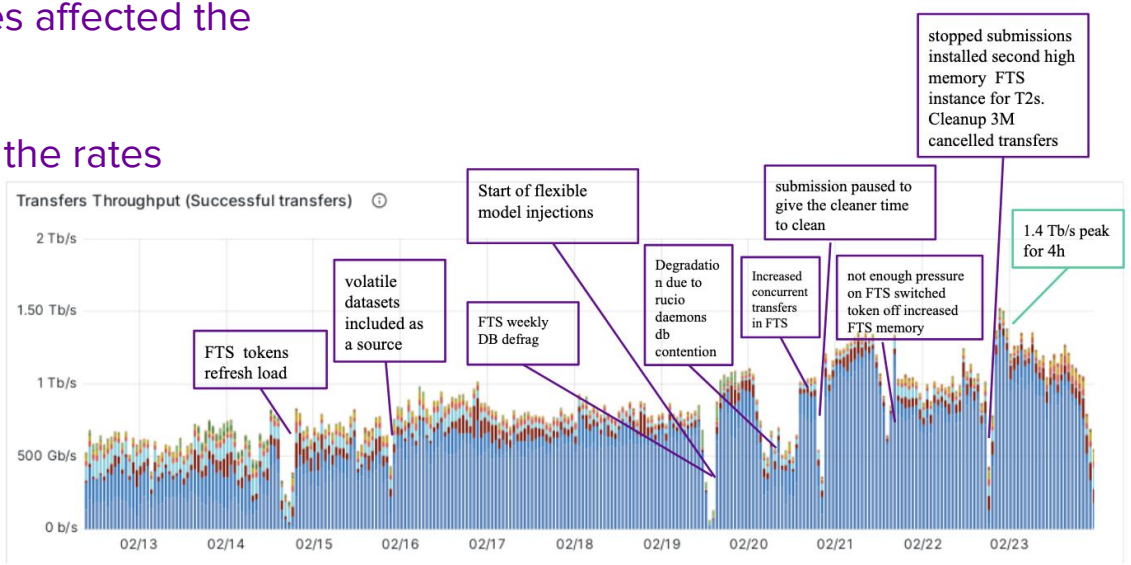

General results

- 3 large T1s had either hardware, network or MW problems
  - These problems became apparent with extra rates
  - T0 rates affected by this
- Day 8 was affected by FTS operations
- Second week was affected by the really large number of transfers

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 ↔ T2 ↔ T0	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 ↔ T2 ↔ T0	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 ↔ T2 ↔ T0	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 ↔ T2 ↔ T0	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 ↔ T2 ↔ T0	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
4	T0 → T1 ↔ T1 → T2	25.2	44.7	35.8	92.2	35.5	28.3	346	124	19.6	173
5	T0 → T1 ↔ T1 → T2	23.1	52.2	36.3	89.2	49.2	46.3	387	134	25.9	197
6	T0 → T1 ↔ T1 → T2	27.4	23.6	30.6	95.5	40.9	41.1	337	104	20.3	201
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 ↔ T2 ↔ T0	29.4	47.1	37.7	29.1	37.3	19.9	400	311	54	100
9	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	32.3	39.1	59.4	84	51.7	42.7	447	330	89.8	139
10	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	43.9	43	92.9	72.3	62.8	52.5	435	337	94.4	97
11	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	51.9	56	111	73.8	66.8	42.1	445	406	127	138
12	T0 ↔ T1 ↔ T1 ↔ T2 ↔ T2 ↔ T0	72.7	58.8	115	70.8	72.9	31.5	418	407	158	174

>90%	
70-90%	
50-70%	
<50%	

T0 export

- Another view on Tier-0 outgoing rates
- Complete degradation in the second week
  - Mostly because of the number of unprioritized transfers within the DC activity
  - On top of the Data Challenge activity being in backfill mode

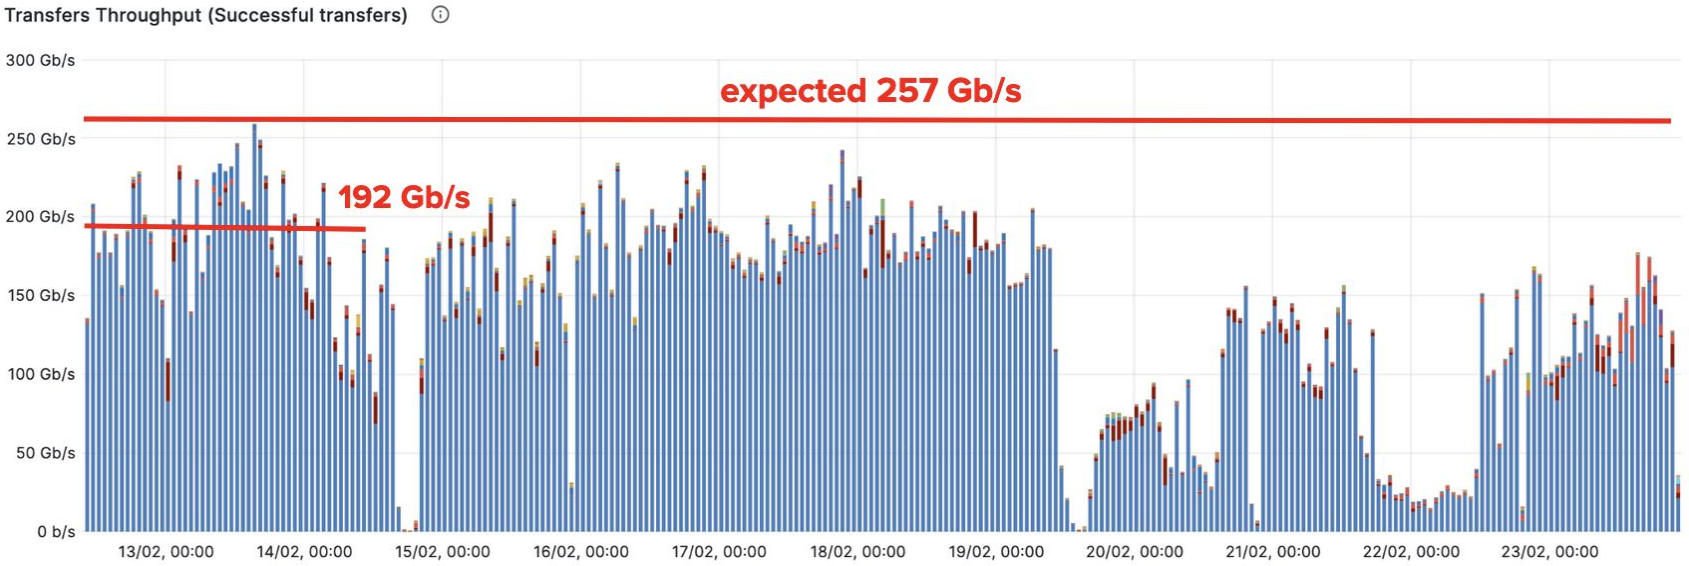

Results explained

- FTS orchestrates **transfers per link over many links**
  - **Doesn't orchestrate throughput**
  - To increase throughput we had to increase the number of allowed parallel transfers by an over an order of magnitude
- Has a concept of fair share per activity
  - Doesn't have a concept of links priorities within an activity, i.e. all links are equally treated T0-T1 same level as T2-T2
    - Could prioritise faster transfers or more important channels
- Testing also new authz system with tokens put further load on the system

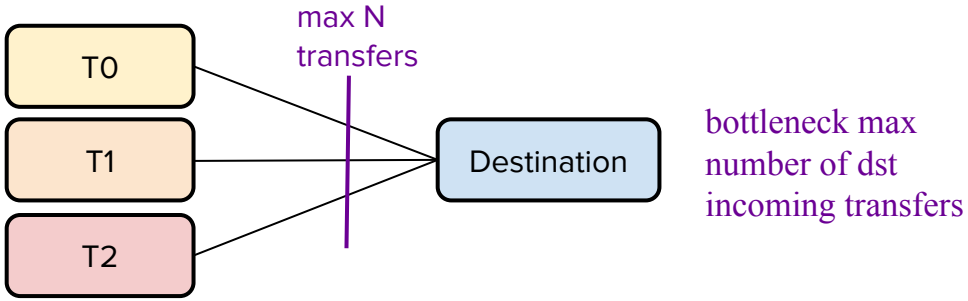

- **Agreed with FTS** problems to solve first for next challange

Tokens

- During DC24 ATLAS tested also the new token based authorization
  - **Plus:** 26% of transfers with tokens ➜ success
  - **Minus:** Load on the FTS/IAM (token providing service) was really high
    - Worked well up to second week
    - Switched off completely to achieve rates in the last day

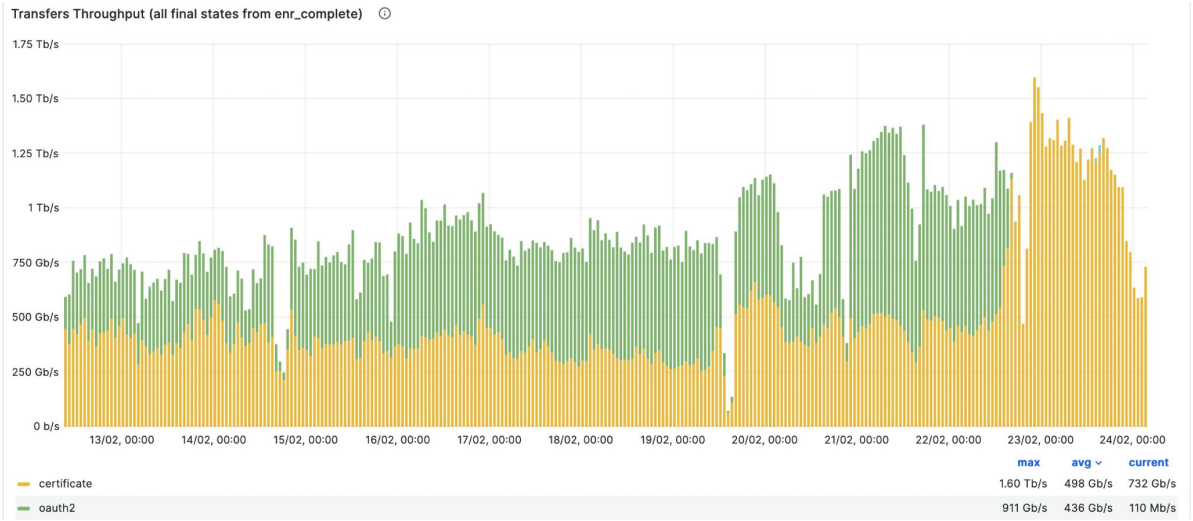

Positive outcome

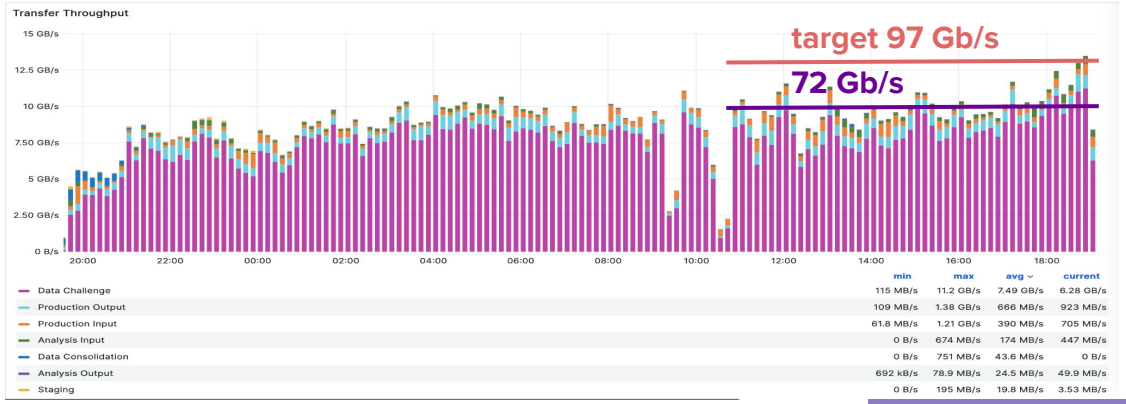

- DC24 ATLAS results were below average
  - Best rates 72 Gb/s
  - Expected 97 Gb/s
- Bottlenecks on the gateways
  - Incorrect Explicit Congestion Notifications (ECN) configuration on gateways.
  - Incorrect network tunings on gateways leading to packet loss.

Benefits of the DC24: improved production rates! RAL ingress

- Recent production traffic after tuning and increasing the number of gateways
  - Rates shown 110 Gb/s sustained over 12h
  - **Concurrent** with CMS AAA traffic at >100Gb/s

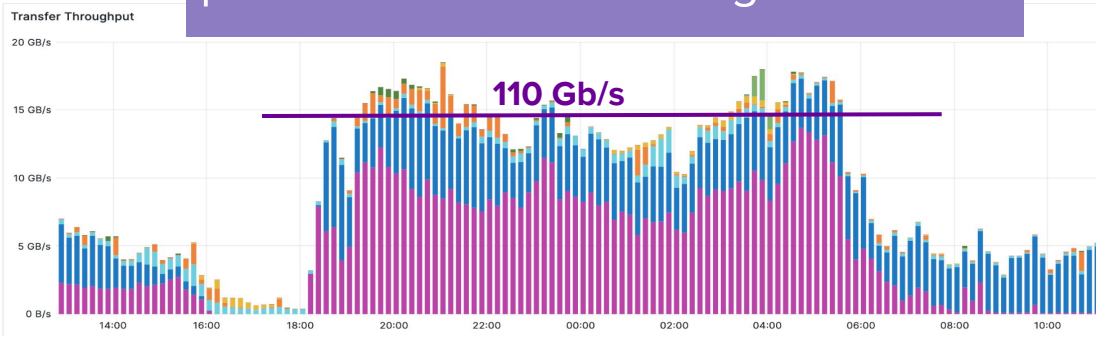

Future

Going forward

- Network
  - DC26 2xDC24 rates (50% of the HL-LHC traffic)
  - Orchestration (SDN, NOTED, SENSE/rucio)
  - Optimization (jumbo frames, traffic pacing, new protocols)
  - Visualization (scitags to label experiments traffic)
- Tokens
  - In DC24 not a priority, in DC26 will have to be battle tested
  - Currently agreeing on tokens policies to test in a miniDC
- Tape testing
  - Not in DC24 but 27% of traffic is to and from tape
  - Tape intrinsically more complicated
  - Each site tape system behaves differently
- Network Monitoring
  - Scitags
  - Tape monitoring
  - Reduce discrepancies between network based and FTS based monitoring
  - Better granularity in the FTS/xrootd dashboards.

Intermediate mini challenges

- Finding all the bottlenecks during the challenge can still be a strategy
  - But with increased traffic it will be more difficult to reach the targets
- **Need mini-challenges** between one challenge and the other
  - Future mini-challenges might be focused on particular sites, regions, technologies or applications.
- ATLAS is identifying when, where and how to execute mini-challenges to test the infrastructure and applications at suitable scale.
  - Streamlining tooling to manage such mini-challenges without requiring expert support
  - Agreeing with sites what to test and when
- Need timeline and priorities also to coordinate with external contributors
  - For example in DC24 tokens development was tested during the challenge not before!

Conclusions

Conclusions

- DC24 demonstrated to be a really **useful exercise to find bottlenecks at every level**
  - The system is complex and slow to change when parameters are tuned
- More **consistent and frequent cooperation** between stake-holders is necessary
  - Particularly for what concerns services development and advanced testing
- Mini intermediate DC to test new technologies and state of scalability fundamental to prepare for DC26 at 50% of the rates
- Method will need to be revised
  - Both to solve DC24 shortcomings and to add new tests