

DC24 ATLAS retrospect


A. Forti,

GridPP51

27 March 2024

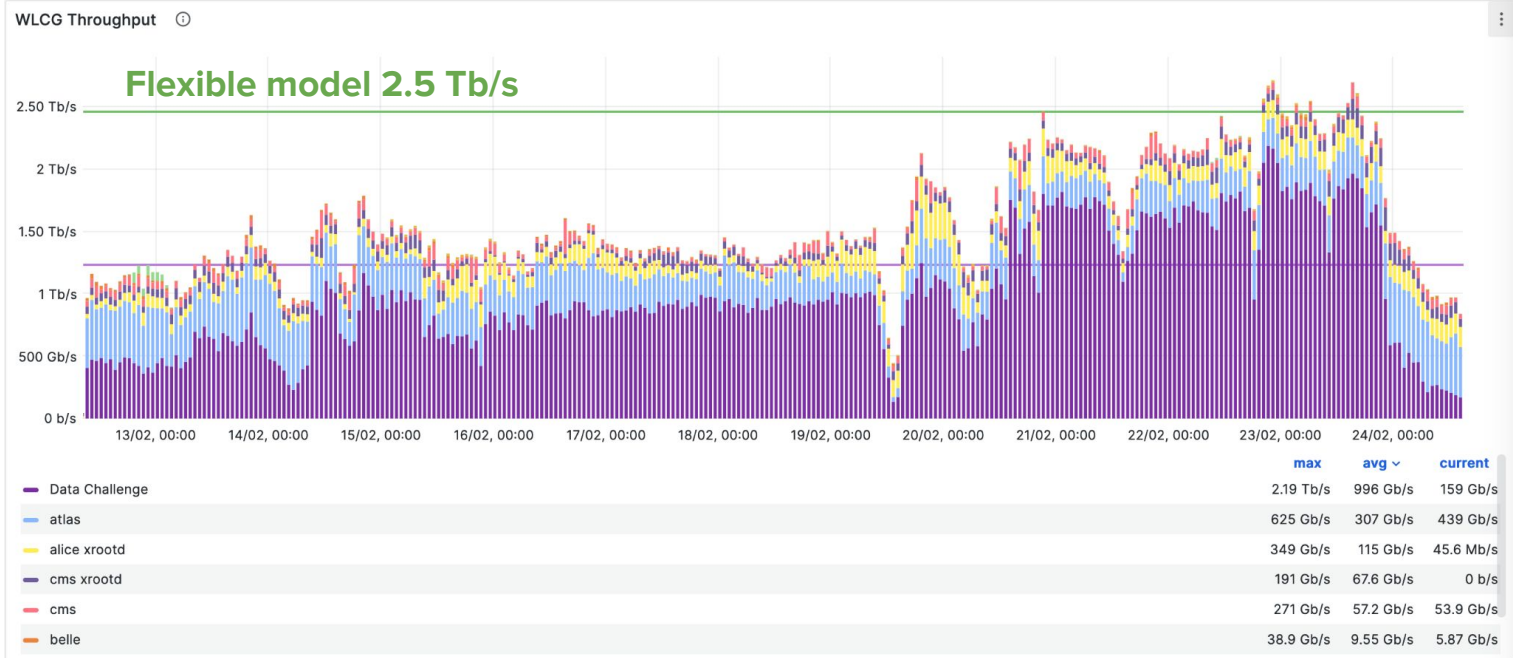


DC21 → DC24



DC21: Production traffic dominant

DC24: DC traffic dominant




ATLAS

- Challenge pushed the whole system
 - Injections every 15 minutes on ~1200 links
 - ~2000 links if we include production
 - 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
- This helped highlighting problems that wouldn't have been seen otherwise in the whole infrastructure




Some explanations



FTS snowball

- 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
 - fts3-atlas 10k -> 25k concurrent transfers
- 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
 - Doesn't prioritise faster transfers
- Tokens put further load on the system




Sites

- Some sites struggled mostly due to storage limitations.
 - Either it wasn't possible to open enough parallel connections (IN2P3-CC)
 - they had a problematic bug (NDGF),
 - or a bottleneck on the gateways due to hardware limitations (RAL).
 - Rates exceeding the expected values and storage not coping (INFN-CNAF)
- Some Tier2s also reported having problems
 - Lancaster had to double the number of gateways from 4 to 8,
 - SWT2 and other sites had a long wave of jobs in transferring state
 - MILANO and other sites saw a large amount of timeouts.
- Overall the number of problems reported, considering the amount of data pushed through, is reasonable
 - 17 problems were reported or GGUS tickets open (list in backup slides)



Deletions

Deletion Average Duration



- Increase in deletion time particularly during the second week.
 - Problem was general but some sites had much higher times than others
 - It will need further investigation to see how it maps with storage types


T0 - T1

- T0 export rates were not achieved
 - Lack of prioritization in FTS
 - Increased concurrent transfers to all T1 to 5k and from T0 to 1k
- T0 exports have been rerun one T1 at the time
 - For at least 6h

Site	T0 Export	DC24 first 2 days	% of expected rates	T0-T1 single test	% of expected rates
CERN-PROD	310.7	197	-	-	-
BNL-ATLAS	68.4	31.5	46%	61.3	90%
FZK-LCG2	39	26.4	68%	42.2	108%
IN2P3-CC	44.2	43	97%	50.9	115%
INFN-T1	28.3	19.3	68%	33.5	118%
NDGF-T1	24.4	13.8	57%	28.2	116%
SARA-MATRIX	19.3	12.2	63%	27.4	1420%
pic	13.3	12.3	92%	18.1	136%
RAL-LCG2	44.4	15	34%	27.2	61%
TRIUMF-LCG2	29.3	23.9	82%	27.2	93%
T1 summary	310.6	197.4	64%	562.6	181%
T1 summary -SARA	291.3	185.2	64%	288.6	99%

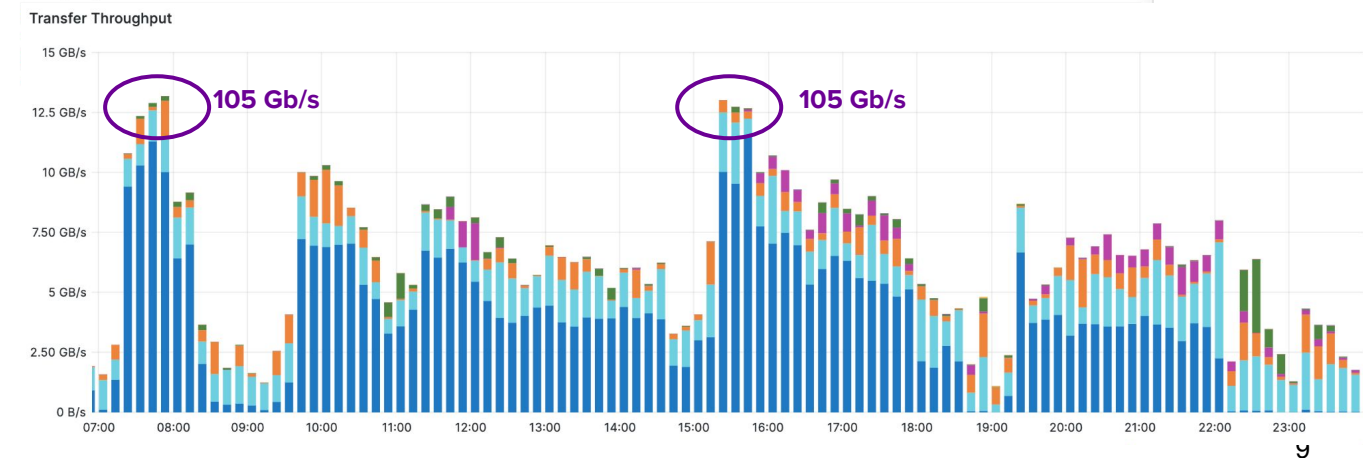

- Testing T1s one by one shows in most cases the original problem was a congested FTS
- TRIUMF was slightly lower but very stable rates
- BNL tends to absorb data too quickly
- RAL will give more details for the WLCG report
 - Internal tape - disk traffic in both direction is always present at RAL
 - 40-50 Gb/s internal pushes any other traffic out
 - We tried both echo and antares for different reasons the rates were below
 - 27.2 Gb/s single test was the first 3h of the test
 - after that antares was stable at 10 Gb/s when we removed deletions
 - antares is not connected to LHCOPN yet

UK: RAL general




- Interval is that of the last day when ATLAS could push the rates higher and Tier2s were involved.

- Destination traffic dominated by the data challenge
 - 70 Gb/s vs 97 Gb/s
- Source more mixed picture
 - 61.6 Gb/s vs 76 Gb/s
- Combined production traffic in other intervals after the challenge has reached Dc24 levels with peaks of 105 Gb/s
 - No doubt the DC24 tuning helped




UK: T2



Green injected

Site	Ingress	Egress
UKI-NORTHGRID-MAN-HEP	31.4	28.2
UKI-SOUTHGRID-RALPP	4.5	4.1
UKI-SCOTGRID-GLASGOW	13.5	12.2
UKI-LT2-QMUL	5.9	3.2
UKI-NORTHGRID-LANCS-HEP	28.6	25.8
UKI-NORTHGRID-LIV-HEP	0.5	0.3

- QMUL traffic was production but ingress was at the rates of DC24
- RALPP concurrent CMS traffic
- Man/Lancs hit their bandwidth limits
- Glasgow should have been much higher



Conclusions

- Positive: system was definitely stressed and it cracked in places
 - Aim of a challenge is finding bottlenecks not only achieve rates
 - Limitations of certain setups were highlighted and, where possible, corrected on the fly.
 - In other places it will require more thinking
- UK
 - RAL needs to do tuning
 - Tier2s generally well but in some cases already hitting the bandwidth
- My recommendation for WLCG is that these DC should be made more often and be less overloaded
 - UK could also agree to do internal challenges

