

Mini-DC: UK edition

CMS & LHCb perspective

Katy Ellis, LHCONE/OPN workshop, 19th March 2025



Science and
Technology
Facilities Council



Introduction to UK mini-DCs

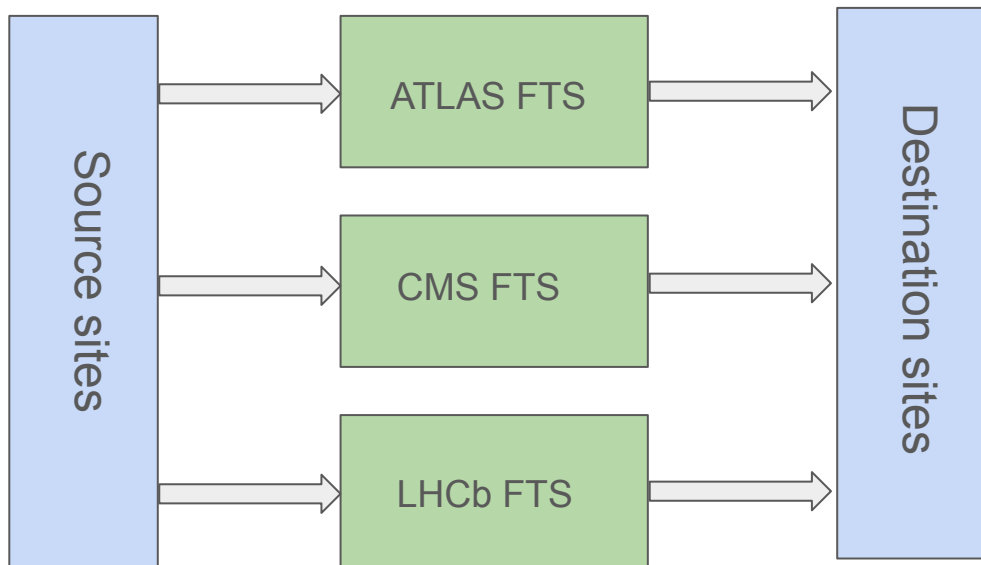
- In DC24 Tier 1s were put under a lot of pressure
 - UK T1 RAL struggled to hit the targets set
 - Partly for reasons out of site control, e.g. LHCOPN was cut
 - But big improvements were made in other areas
- In DC24 Tier 2s were not given much attention
 - VO teams were focussed on the major links and other problems
 - Since DC24 many changes at UK sites
 - New storage technologies adopted
 - New data centre (QMUL)
 - Sites wanted to know their capabilities

Test overview

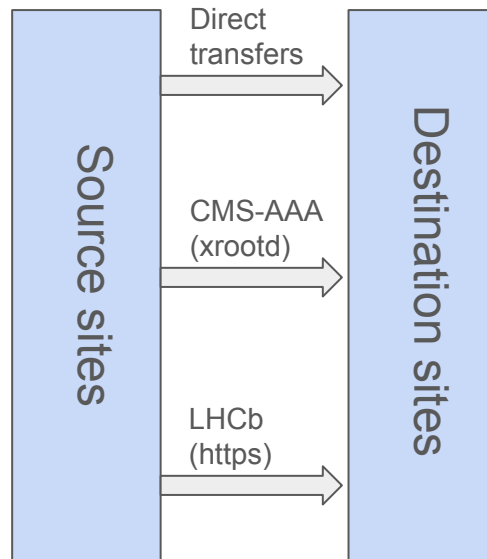
- Tier 2 tests - week of 9th December 2024
 - CMS sites tested: IC, RALPP, Brunel
 - RAL T1 used as a source and sink for transfers
 - CERN also used as a source
- Tier 1 tests - week of 3rd March 2025
 - Disk endpoint (Echo) rigorously tested
 - (Tape tests will come later, following hardware and network changes)
 - CERN and CMS European T1s used as source and sink for transfers
- Each VO used similar file size ~5GB
- ATLAS and CMS ran the `dc_inject` tool which creates and deletes rules in Rucio periodically, and continuously until stopped
- LHCb created a new tool to inject data continuously (unlike previous methods)
 - Could be used in subsequent DCs

Controlling the flow?

Throttling by number of concurrent connections on link or storage

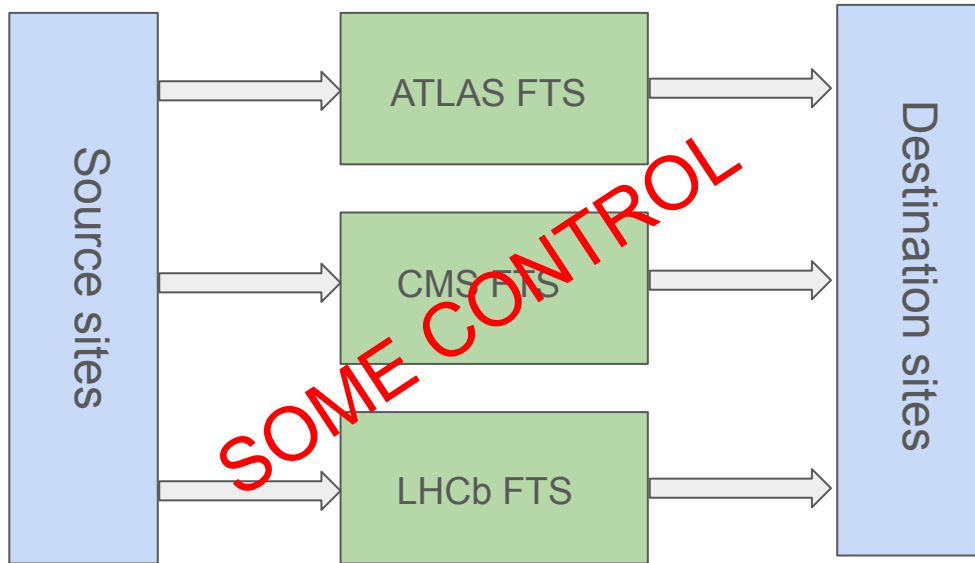


N.B. usually data does not actually pass through the FTS

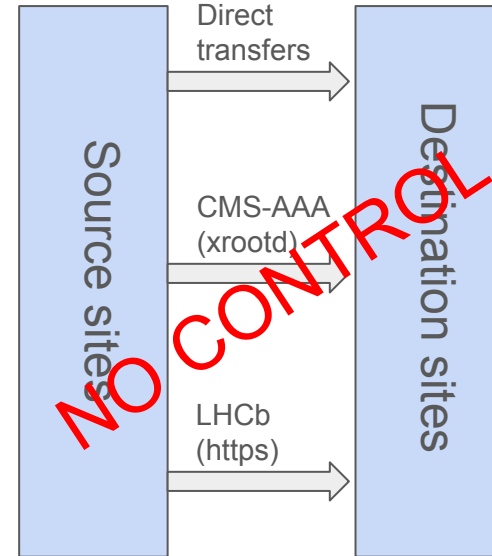


Controlling the flow?

Throttling by number of concurrent connections on link or storage







N.B. usually data does not actually pass through the FTS



Tier 1 tests

Schedule and high-level results

Day	Plan	Executed
Monday	Half-fill the OPN: writes to Echo from CERN	More than half-filled 
Tuesday	Simulate 'cutting' of OPN; fallback to LHCONE	Observed slight drop 
Wednesday	Revert to OPN and saturate; data from CERN/T1s	Some disturbance to gws 
Thursday	Reads from Echo; data to CERN/T1s	Easily achieved 
Friday	Tests with UK Tier 2 sites; postponed	Tests with checksum-on-the-fly

Injector tool issues (timestamp is CET)

2025-03-03 14:56:07 :: T2_CH_CERN->T1_UK_RAL_Disk :: Injection messed up: The creation of the replication rule failed at this time. Please try again later.;;Details: (cx_Oracle.IntegrityError) ORA-00001: unique constraint (CMS_RUCIO_PROD.REPLICAS_PK) violated -- SKIPPING

Oracle DB error (no Rucio rules created for 4 cycles of 15 minutes = 1 hour)

Skipping bad dataset cms:/ExpressPhysics/Run2024F-Express-v1/FEVT#efae4f11-3ffd-417a-bf63-6f69c48fc723

Skips data that is either removed from source or already at destination (rules are still created with alternative datasets)

Appears a lot for smaller T1s with less unique data?

2025-03-04 02:12:23 :: T2_CH_CERN->T1_UK_RAL_Disk :: Adding rule for cms:/ExpressPhysics/Run2024F-Express-v1/FEVT#995507dc-b565-48a8-837a-937e32cf4547

2025-03-04 02:13:748 **WARNING** Waiting 0.25s due to reason: server returned 504

2025-03-04 02:14:24,003 **WARNING** Waiting 0.5s due to reason: server returned 504

2025-03-04 02:15:24,507 **WARNING** Waiting 1.0s due to reason: server returned 504

2025-03-04 02:15:25 :: T2_CH_CERN->T1_UK_RAL_Disk :: Injection messed up: An unknown exception occurred.;;Details: no error information passed (http status code: 504) -- SKIPPING

504 error (???) (no Rucio rules created)

2025-03-04 19:58:18 :: T2_CH_CERN->T1_UK_RAL_Disk :: Adding rule for cms:/ExpressPhysics/Run2024E-Express-v1/FEVT#0af43c1c-a273-47e0-b5f2-91f04537ca3a

2025-03-04 19:58:19 :: T2_CH_CERN->T1_UK_RAL_Disk :: Injection messed up: An unknown exception occurred.;;Details: no error information passed (http status code: 500) -- SKIPPING

500 error (no Rucio rule created)

Many errors appear multiple times

Injector tool crash (timestamp is CET)

2025-03-04 20:13:21 :: T2_CH_CERN->T1_UK_RAL_Disk :: we have 1923 fresh datasets out of 1932

Skipping bad dataset cms:/ExpressPhysics/Run2024G-Express-v1/FEVT#560c4e1b-968d-432f-8e26-453d7eea2744

Skipping bad dataset cms:/ExpressPhysics/Run2024G-Express-v1/FEVT#87202a56-7b07-4f81-b707-e93649ea21cc

Exception in thread Thread-1:

Traceback (most recent call last):

```
File "/usr/local/lib/python3.8/threading.py", line 932, in _bootstrap_inner
    self.run()
```

```
File "/usr/local/lib/python3.8/threading.py", line 870, in run
    self._target(*self._args, **self._kwargs)
```

```
File "dc_inject.py", line 136, in run_transfer_thread
    if is_dataset_bad(src, dst, dataset['did'], rucio_client):
```

```
File "dc_inject.py", line 81, in is_dataset_bad
    rse_list = list(rucio_client.list_dataset_replicas(scope, name))
```

```
File "/cvmfs/cms.cern.ch/rucio/x86_64/rhel8/py3/current/lib/python3.6/site-packages/rucio/client/replicaclient.py", line 348, in list_dataset_replicas
    raise exc_cls(exc_msg)
```

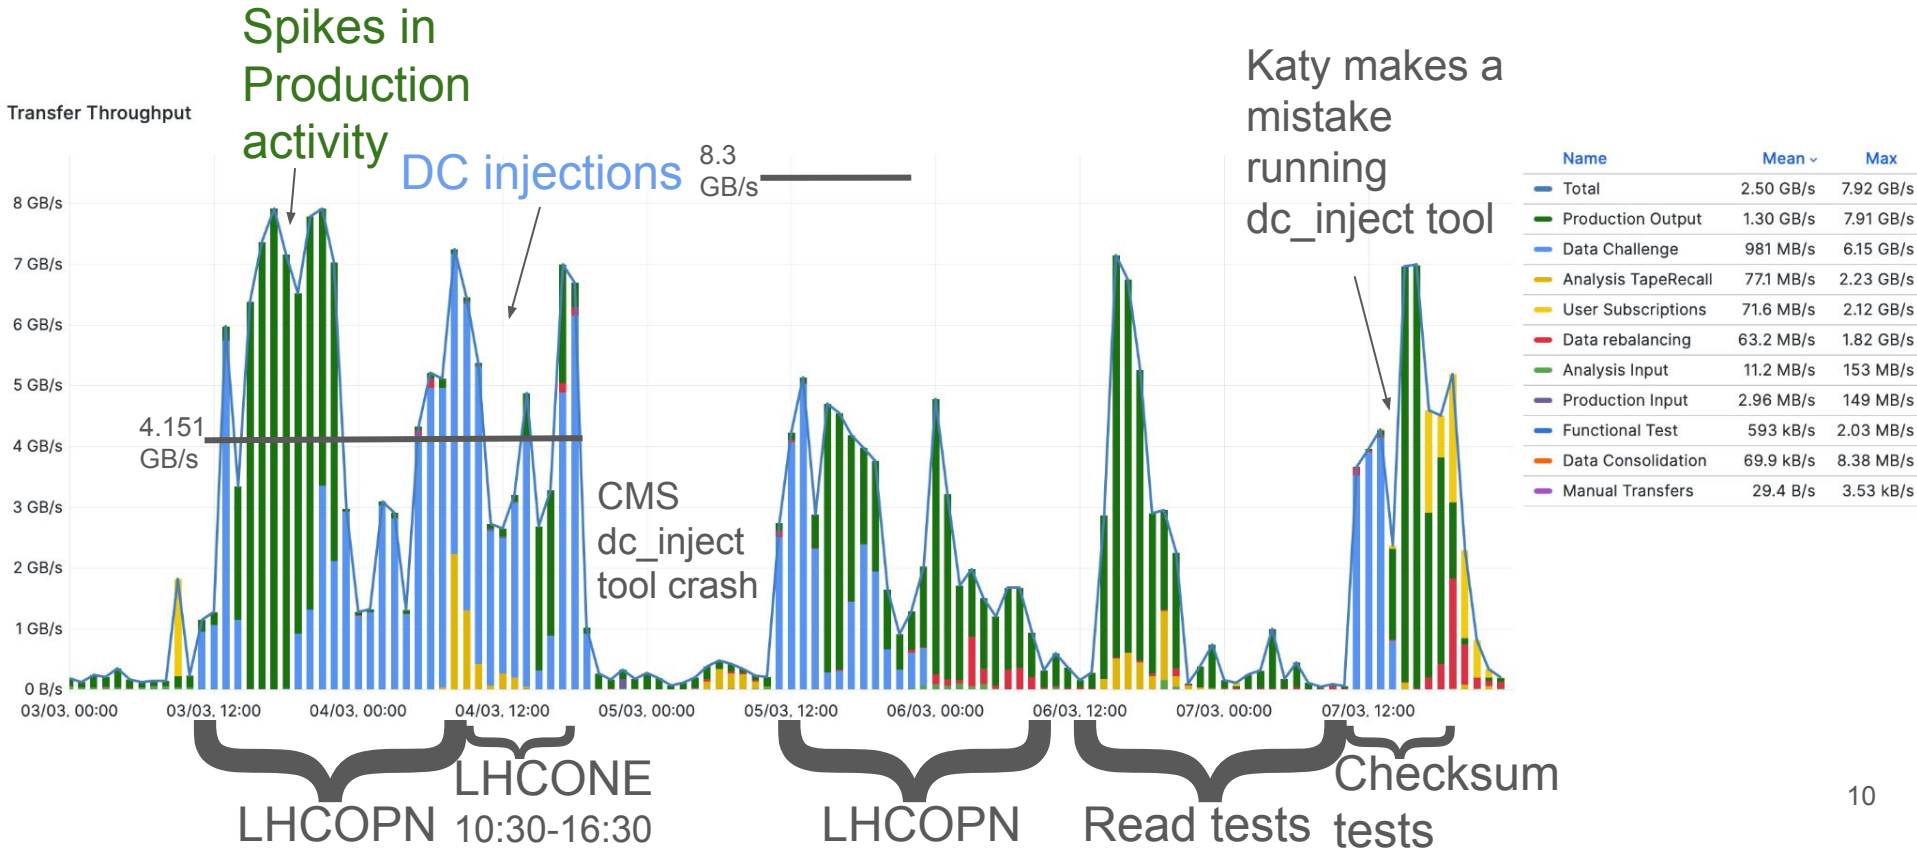
rucio.common.exception.DatabaseException: Database exception.

Details: (cx_Oracle.DatabaseError) ORA-02391: exceeded simultaneous SESSIONS_PER_USER limit

(Background on this error at: <https://sqlalche.me/e/20/4xp6>)

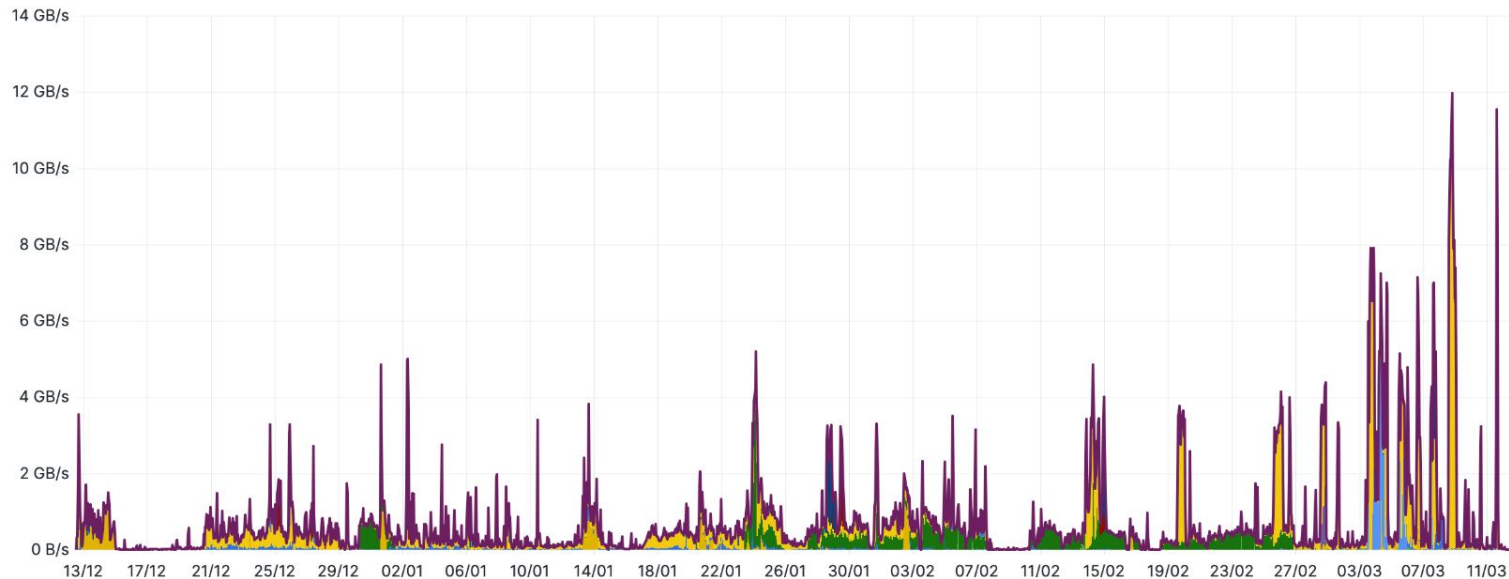
After the crash, probably no 'used datasets' lists are dumped

FTS Transfers to Echo: CMS



CMS FTS transfers to Echo from last 3 months

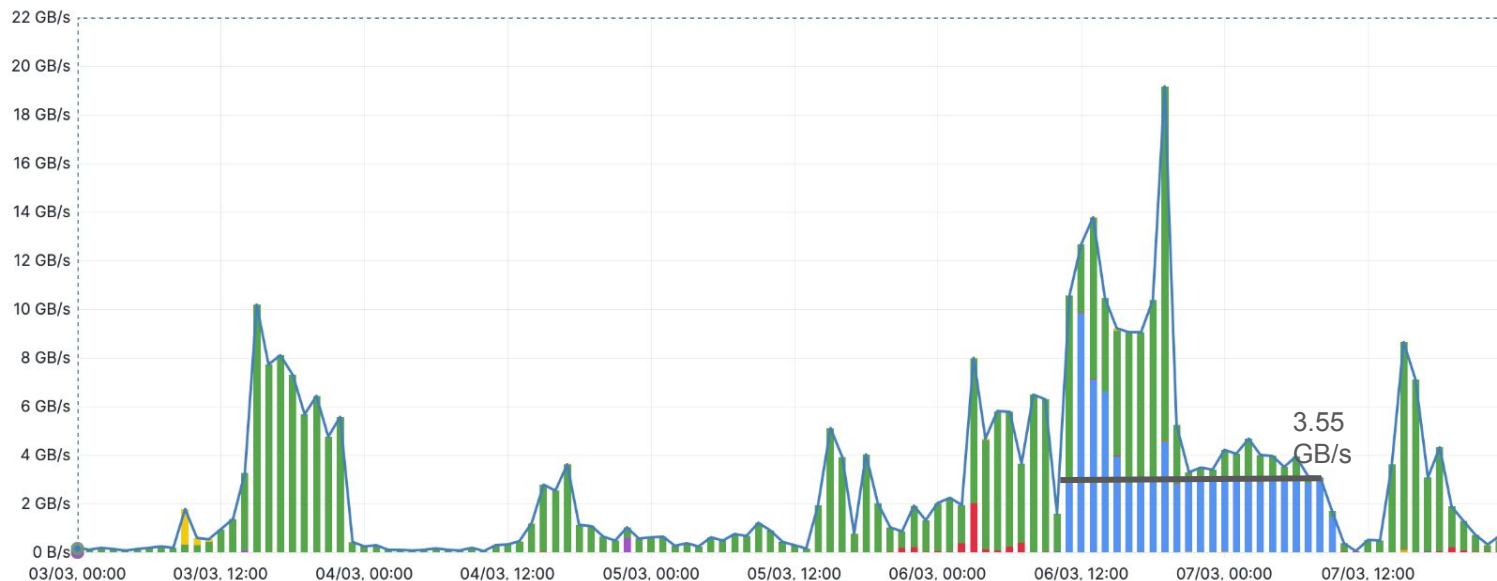
Transfer Throughput



Name	Mean	Max
Total	725 MB/s	12.0 GB/s
Production Output	371 MB/s	11.9 GB/s
Production Input	136 MB/s	4.75 GB/s
Data Challenge	54.5 MB/s	6.15 GB/s
Analysis TapeRecall	51.1 MB/s	4.58 GB/s
User Subscriptions	38.2 MB/s	4.45 GB/s
Data rebalancing	36.5 MB/s	1.82 GB/s
User AutoApprove	14.7 MB/s	2.58 GB/s
Recovery	9.70 MB/s	2.76 GB/s
Analysis Input	6.44 MB/s	586 MB/s
Manual Transfers	5.82 MB/s	2.24 GB/s
Functional Test	518 kB/s	4.20 MB/s
Data Rebalancing	108 kB/s	196 MB/s
T0 Export	92.5 kB/s	27.4 MB/s
Data Consolidation	3.88 kB/s	8.38 MB/s

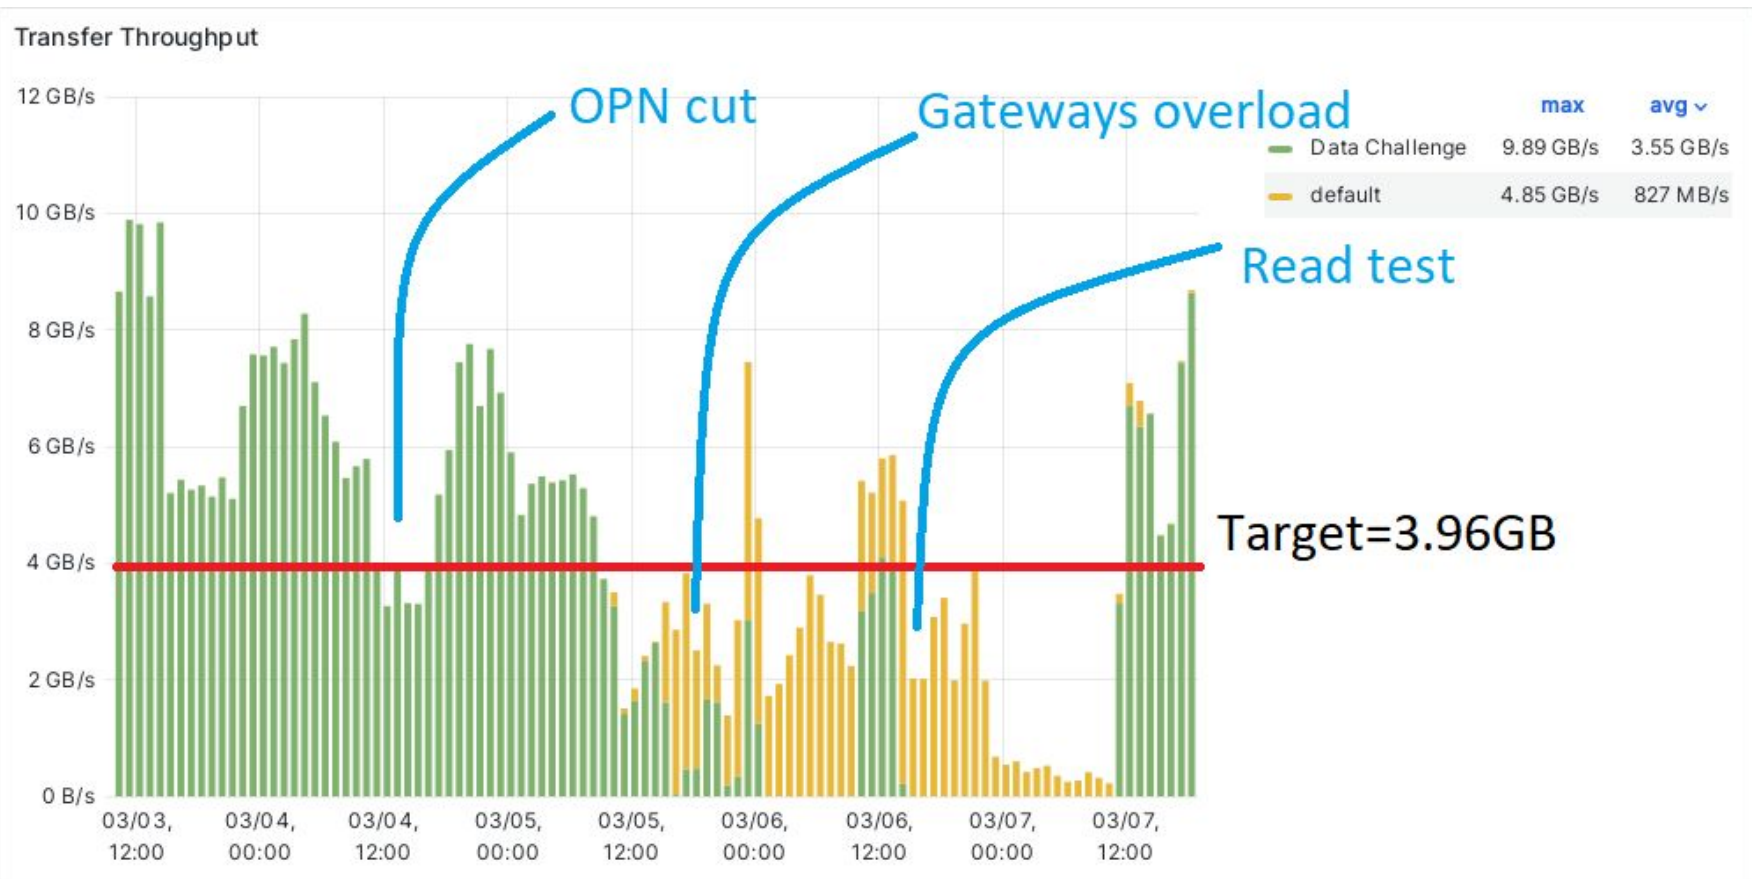
FTS Transfers from Echo: CMS reads

Transfer Throughput

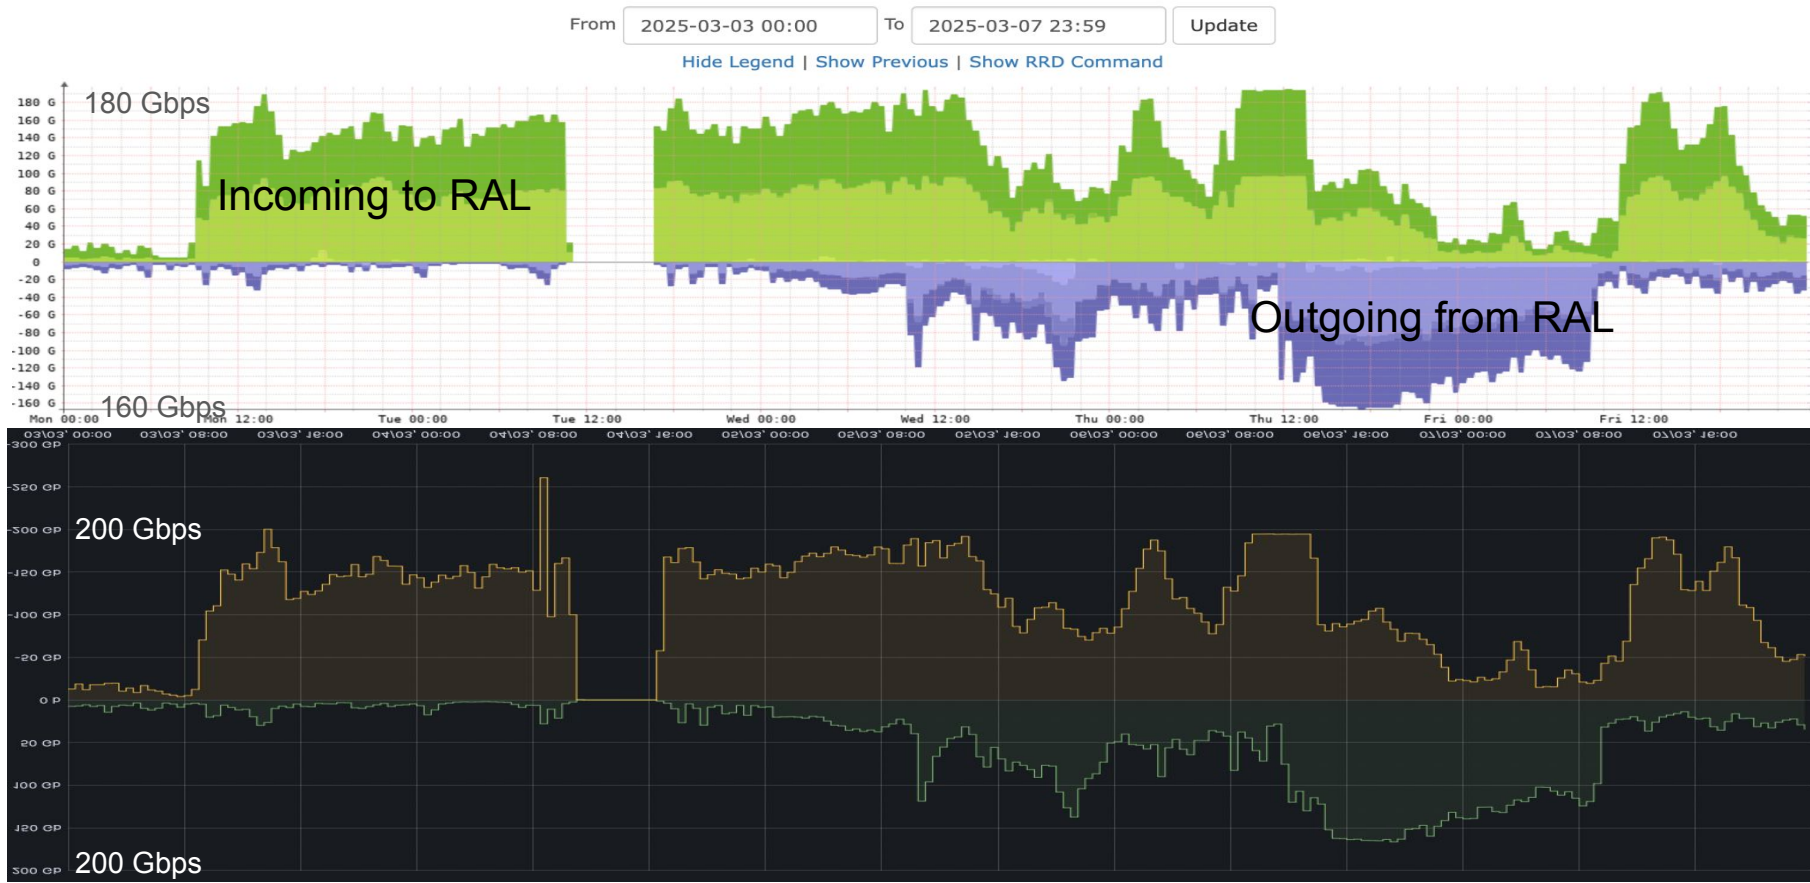


Name	Mean ~	Max
Total	2.94 GB/s	19.1 GB/s
Production Output	2.17 GB/s	14.6 GB/s
Data Challenge	697 MB/s	9.85 GB/s
Data rebalancing	41.3 MB/s	2.04 GB/s
User Subscriptions	21.6 MB/s	1.45 GB/s
Production Input	5.58 MB/s	590 MB/s
Analysis TapeRecall	1.43 MB/s	109 MB/s
Functional Test	802 kB/s	2.40 MB/s
User AutoApprove	485 kB/s	36.3 MB/s
Analysis Input	179 kB/s	8.56 MB/s
Data Consolidation	69.9 kB/s	8.38 MB/s

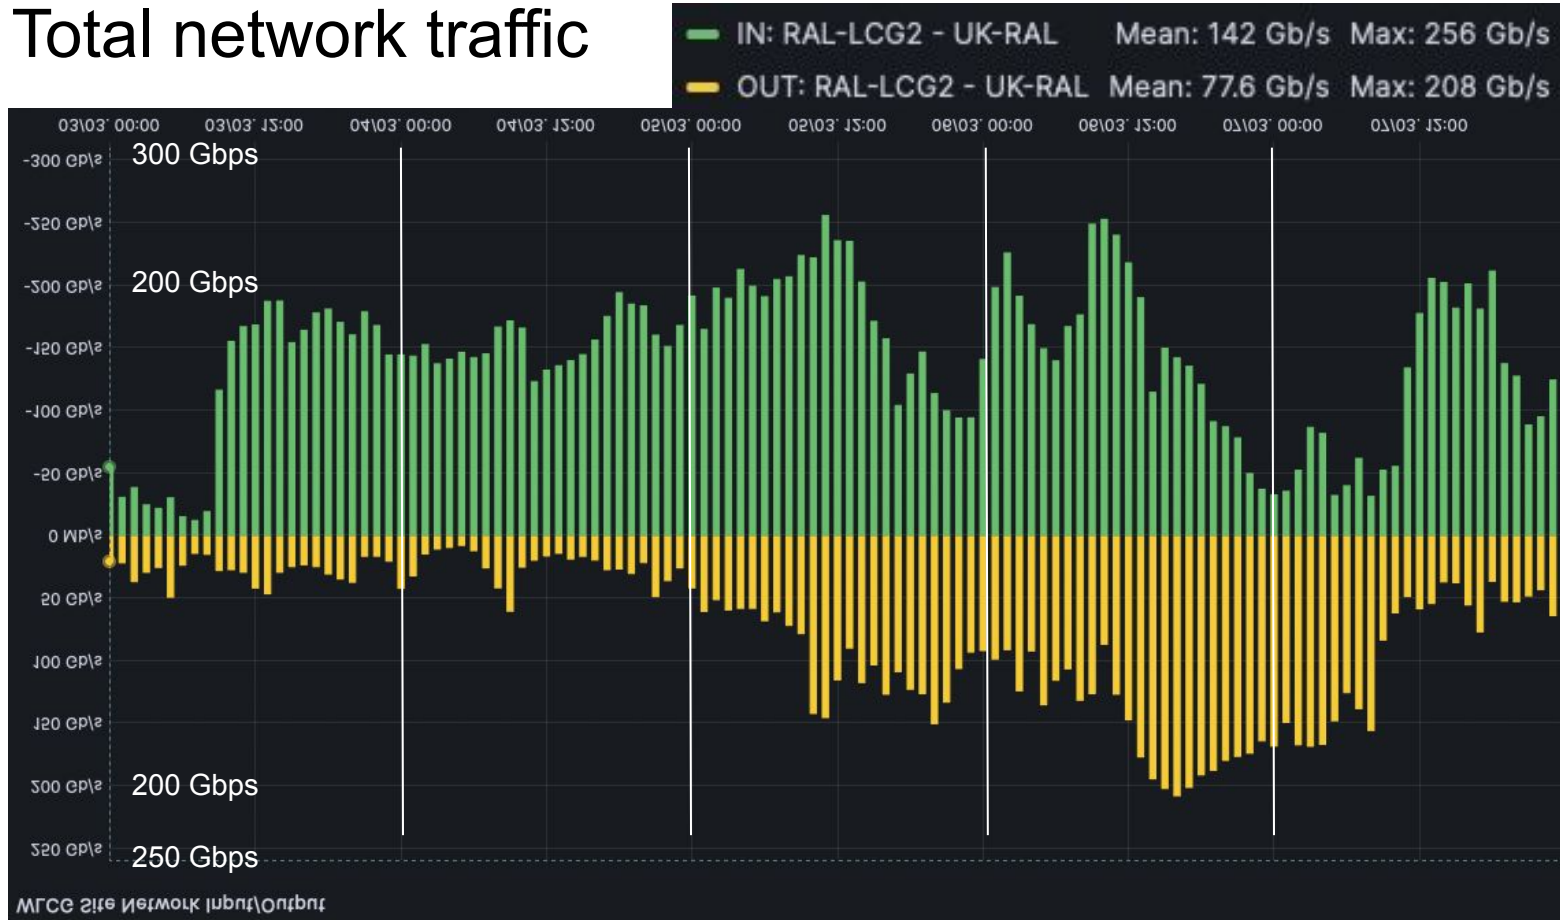
FTS Transfers to Echo: LHCb



Comparing network monitoring (OPN)



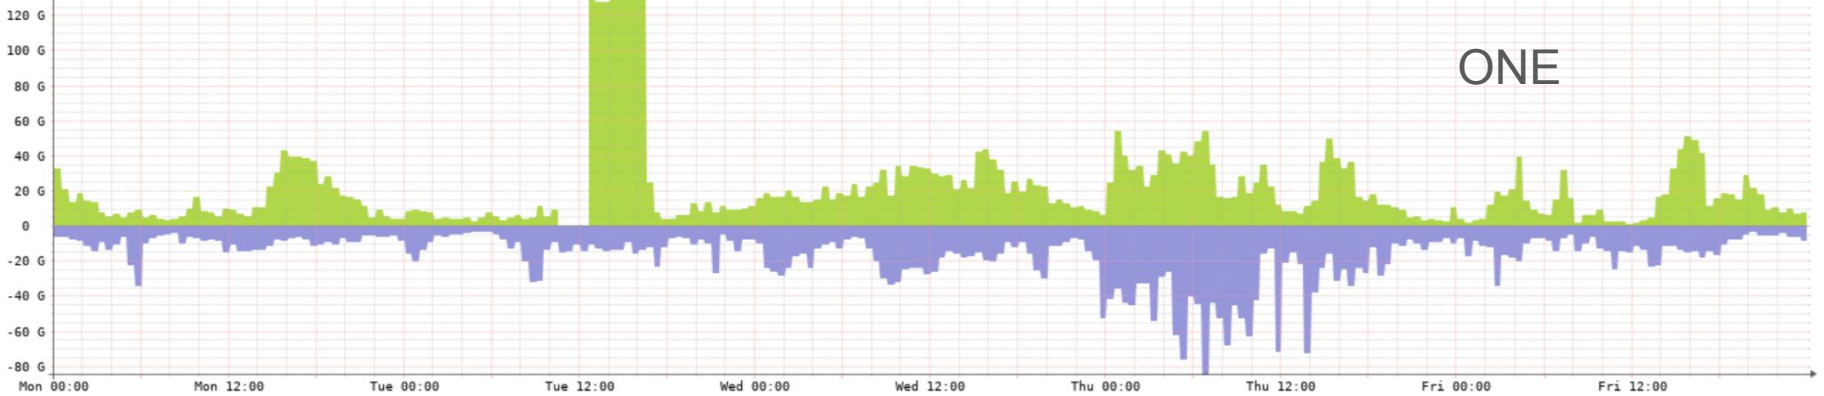
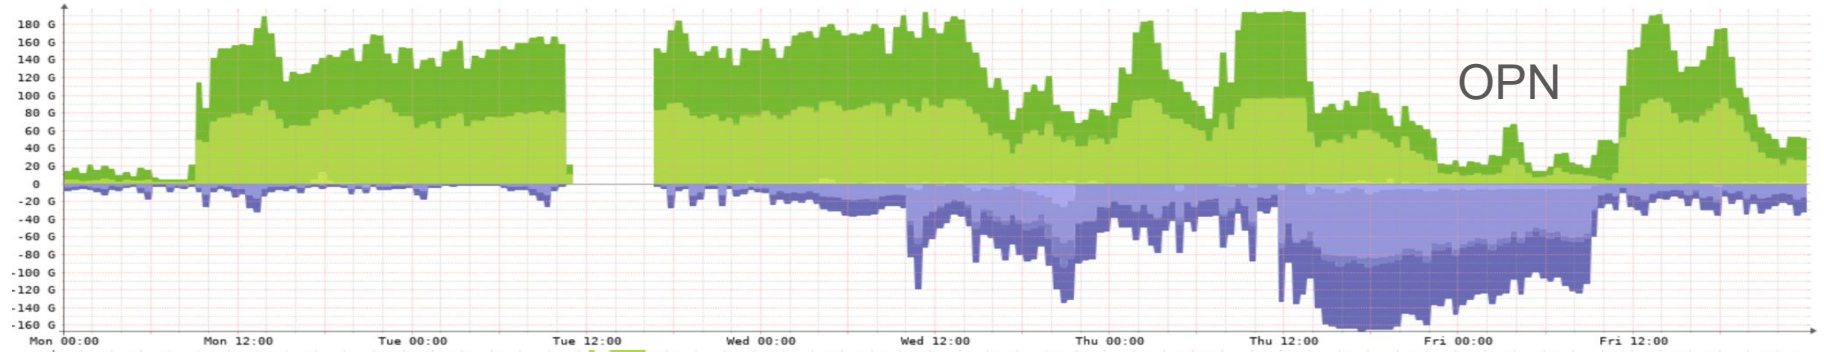
Total network traffic



LHCOPN and LHCONE

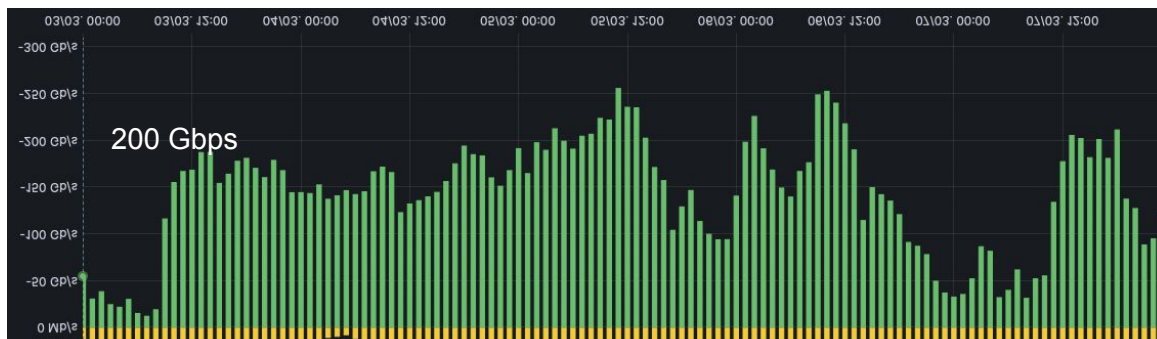
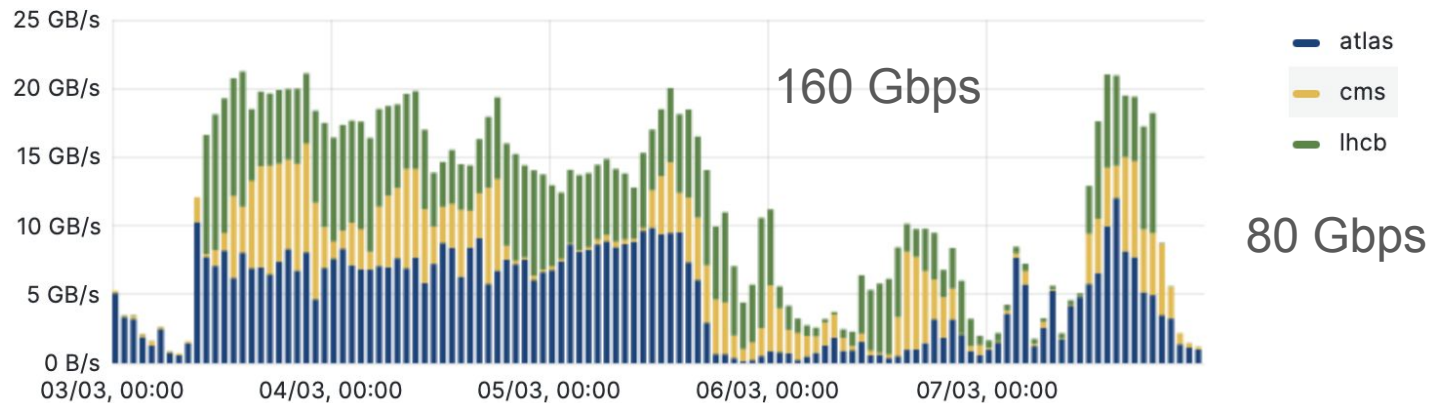
From To

[Hide Legend](#) | [Show Previous](#) | [Show RRD Command](#)



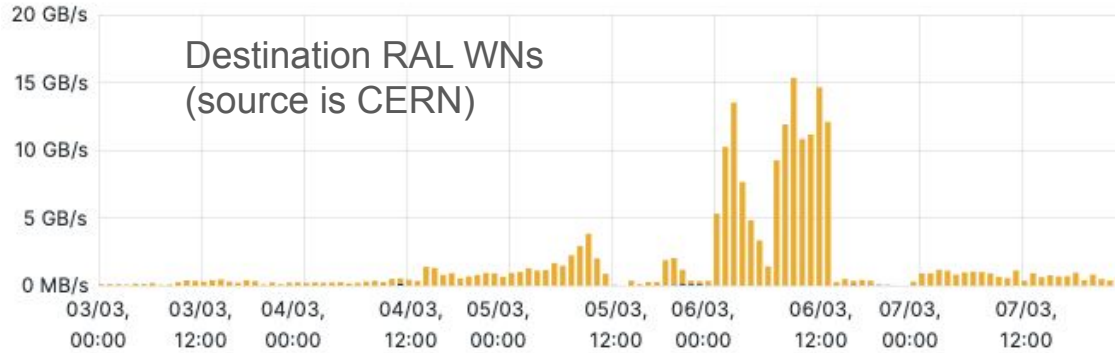
FTS vs total (IN to RAL)

Transfer Throughput



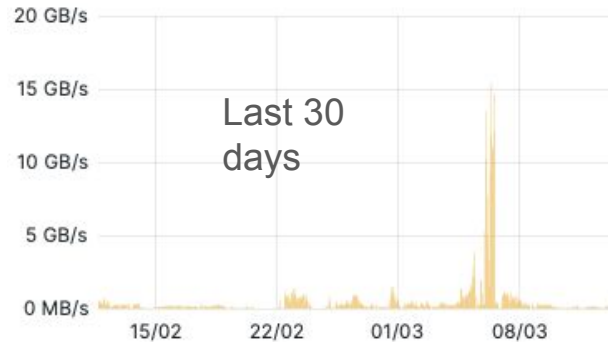
Direct reads from batch farm to remote site

Transfer Throughput

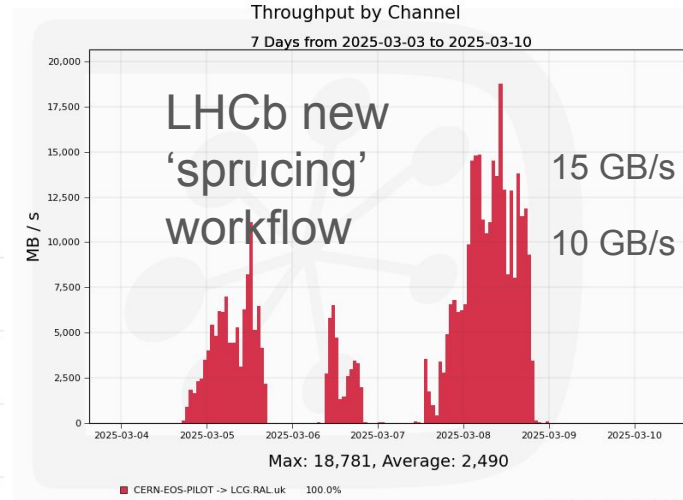


Shoveler (CMS-AAA)

Transfer Throughput



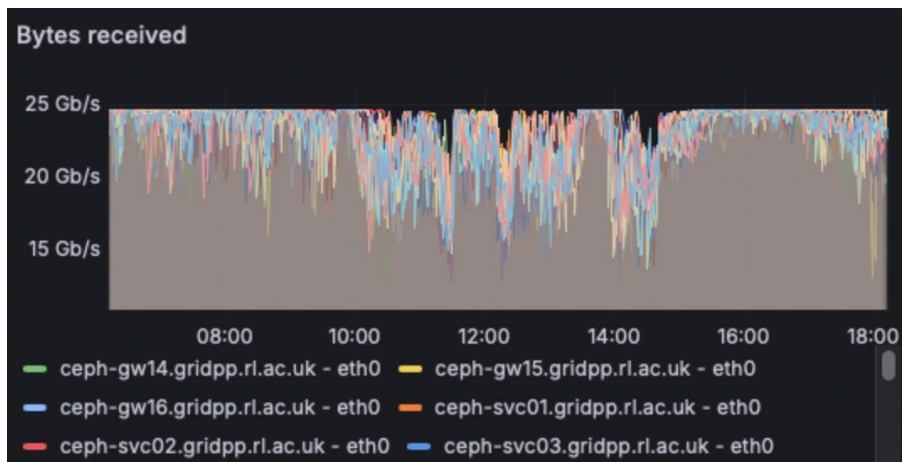
- atlas
- cms
- lhcb



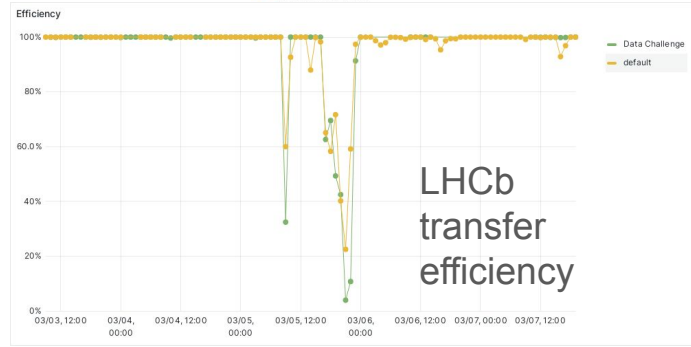
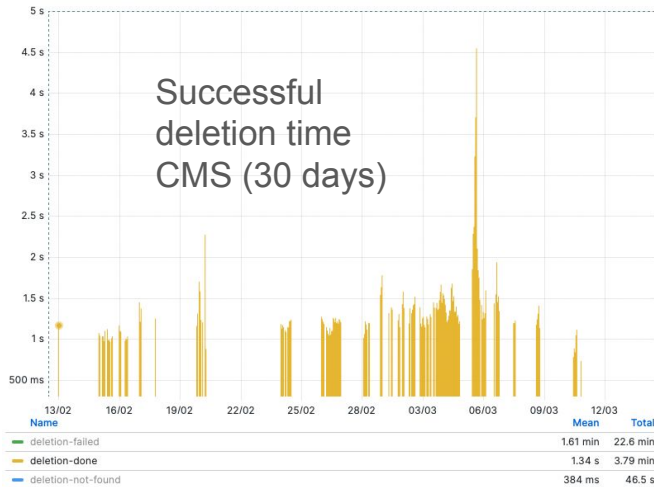
Deletions

- RAL disk storage technology is intrinsically slower at performing deletions (as it actually needs to delete data rather than just unlink the files)
 - Around 1 second per file instead of ~hundreds of ms
 - Interesting to observe ATLAS deletions are quicker for similar file size
- This caused problems during DC24 when the gateways were under maximum pressure
- (In DCs, the deletion rate REALLY has to keep up with the transfer rate, which is only somewhat artificial)
- Normal operation in Rucio has deletion commands issued in a site-by-site cycle
 - To mitigate slower deletes at RAL disk, we create a bespoke deletion daemon to only target RAL

Gateways maximally stressed (Wed)



Deletion Average Duration



T1_UK_RAL rdr.echo.stfc.ac.uk org.cms.SE-WebDAV-7crt-write: UNKNOWN

T1_UK_RAL rdr.echo.stfc.ac.uk org.cms.SE-WebDAV-8crt-directory: UNKNOWN

T1_UK_RAL rdr.echo.stfc.ac.uk org.cms.SE-WebDAV-99summary: UNKNOWN

CRITICAL 2025-03-05T16:55:00Z to 2025-03-05T17:15:00Z 20 minutes

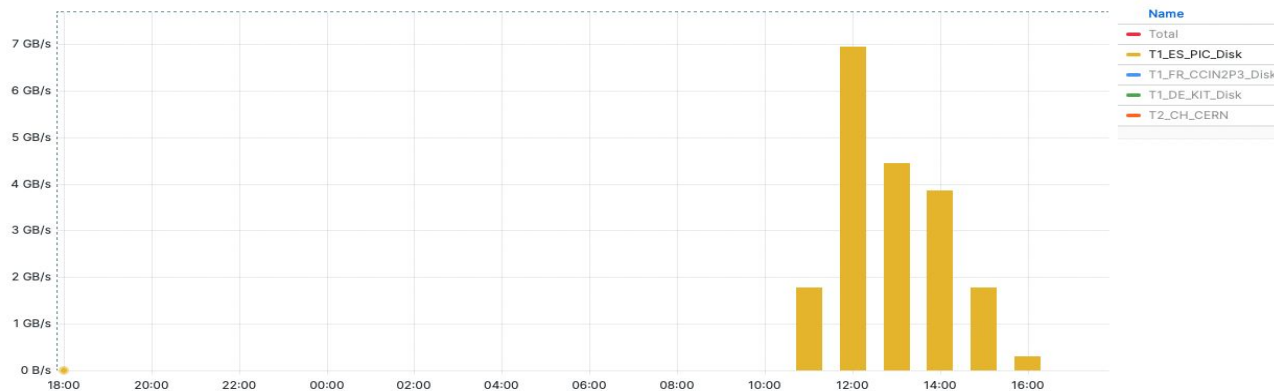
(Accidentally) stress testing PIC (Thur)

Reads from RAL, with data transferred to a range of CMS T1s and CERN

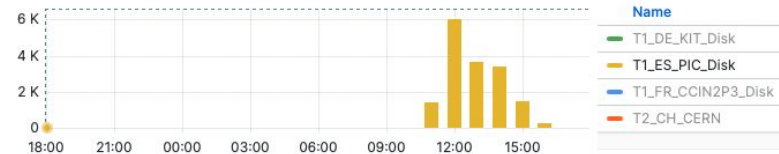
T1_UK_RAL_Disk, T2_CH_CERN, 7803
T1_UK_RAL_Disk, T1_DE_KIT_Disk, 7803
T1_UK_RAL_Disk, T1_FR_CCIN2P3_Disk, 7803
T1_UK_RAL_Disk, T1_ES_PIC_Disk, 50000

~

Transfer Throughput



Transfer Successes



Transfer Failures

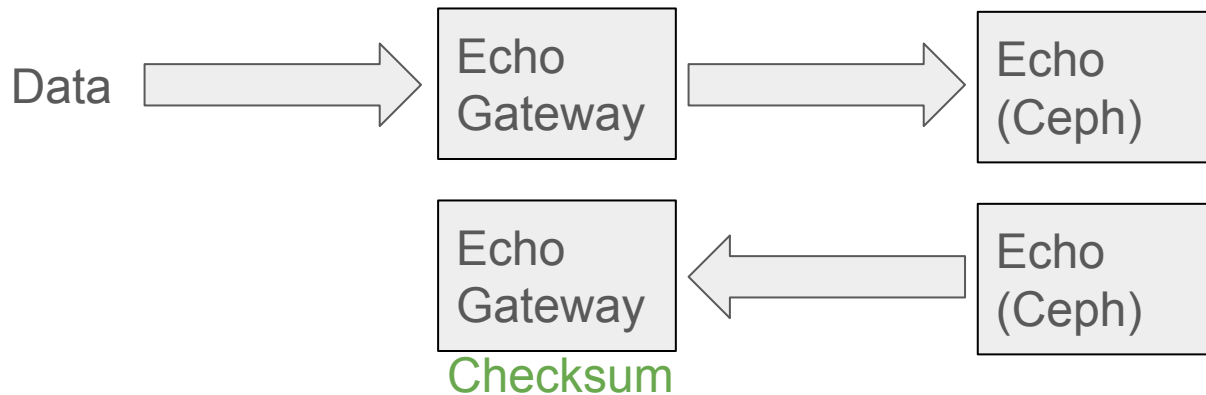


No data

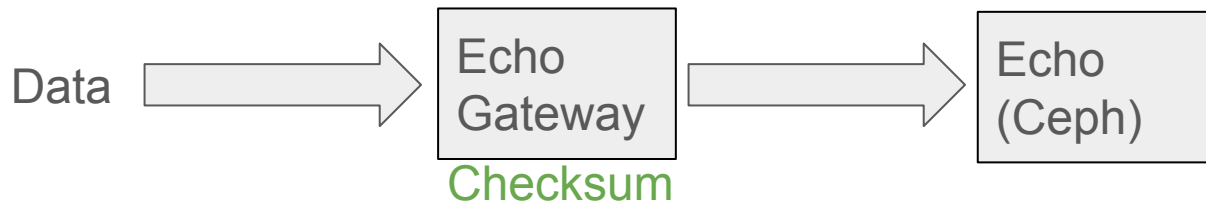
Antonio was
very pleased!

Checksum on-the-fly

- The old checksum method:



- The new checksum method (on-the-fly):



Checksum on-the-fly: what happened?

- Both new and old methods run simultaneously and checksums compared
- Performance of the new algorithm was fine
- Some bugs to clean up
- A couple of possible mismatches which are being investigated
- Once this is fully tested we will remove the old method checksum, and should see some gains in gateway throughput

https://monit-grafana.cern.ch/d/dd6c5e41-6946-47b9-bb1d-50ac89bc7499/user-tunikode-fts-queue-metrics?orgId=11&var-source=davs:%2F%2FAtlasdtn01.usatlas.bnl.gov&var-destination=davs:%2F%2Fantares.stfc.ac.uk&var-source_se=davs:%2F%2Fcmsdcadisk.fnal.gov&var-dest_se=davs:%2F%2Frdr.echo.stfc.ac.uk&tab=query&from=now-30d&to=now&var-bin=1d

Per-file transfer speed

Per-file = throughput / active_count

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2025-03-18T00:26:18Z	1000	278	10305	100.00%	1.88 GiB/s	1.88 GiB/s	0	Good link efficiency. Increment. Hit upper range limit
2025-03-18T00:25:39Z	1000	268	10352	100.00%	1.77 GiB/s	1.77 GiB/s	0	Good link efficiency. Increment. Hit upper range limit
2025-03-18T00:25:01Z	1000	277	10359	100.00%	1.77 GiB/s	1.77 GiB/s	0	Good link efficiency. Increment. Hit upper range limit

Time	data.source_se	data.dest_se	data.active_count	data.throughput
Mar 4, 2025 @ 15:29:30.359	davs://eoscms.cern.ch	davs://rdr.echo.stfc.ac.uk	766	4,736,329,630.439
Mar 4, 2025 @ 15:28:54.315	davs://eoscms.cern.ch	davs://rdr.echo.stfc.ac.uk	805	4,768,777,080.548
Mar 4, 2025 @ 15:28:17.267	davs://eoscms.cern.ch	davs://rdr.echo.stfc.ac.uk	841	4,741,144,937.96
Mar 4, 2025 @ 15:27:39.228	davs://eoscms.cern.ch	davs://rdr.echo.stfc.ac.uk	867	4,752,035,292.219



Source	Avg per-file rate
CNAF T1	4 MB/s
IN2P3 T1	8 MB/s
FNAL T1	6 MB/s
Caltech	3 MB/s
Nebraska	3 MB/s

Tier 2 tests (CMS only)

CMS Tier 2 sites

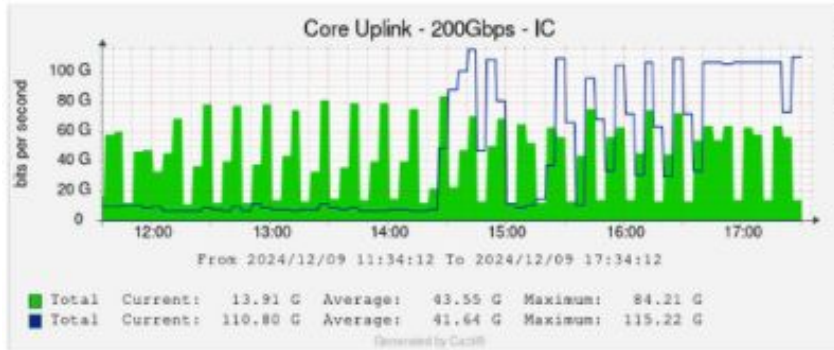
- T2_UK_London_IC (Imperial College London)
 - 100Gbps link, CMS Tier 2
- T2_UK_SGrid_RALPP (RAL T2 site)
 - 100Gbps link, Tier 2 for ATLAS, CMS (+many more)
- T2_UK_London_Brunel
 - Requested to be write-tested up to 30Gbps and 40Gbps if possible
 - New storage technology adopted (Ceph replaced DPM)
- (T2_UK_SGrid_Bristol)
 - Small site, not tested as unavailable during test week

Talk from DOMA-General: <https://indico.cern.ch/event/1495675/>

Tier 2 results - IC

Each site was easily able to sustain DC24 write rates from RAL Tier 1, and was then pushed harder...

IC went over their theoretical limit...



Simultaneous writes and reads

Image from Daniela. Green is writes and blue is reads

Tier 2 results - RALPP

A small adjustment to the number of possible connections allowed RALPP to fill their 100 Gb/s bandwidth

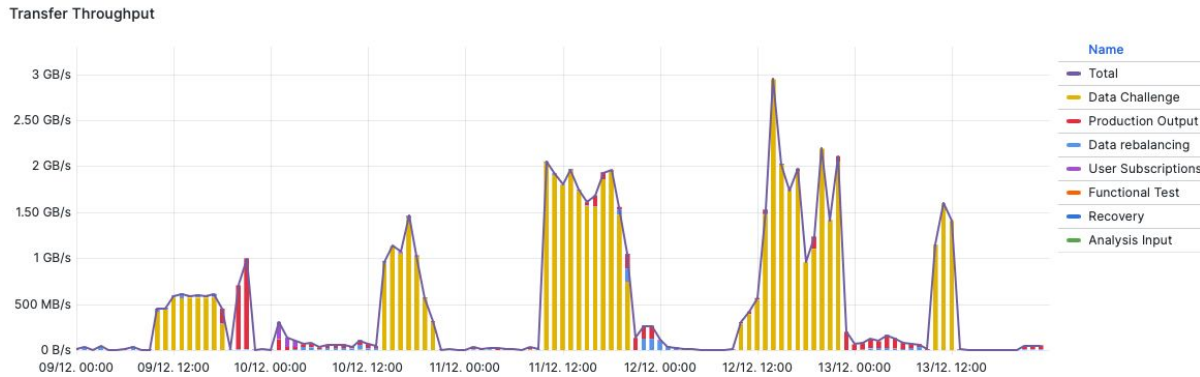
We can also view RALPP activity in the nice new network monitoring:



Traffic of RALPP, mostly ingress. The earlier spiky behaviour is due to the site absorbing data faster than the tool is injecting. The final hour is where the data injected is high enough to create a sustained plateau.

Tier 2 results - Brunel

- Attempt to write at 30 Gbps; a sustained rate around 15 Gbps was achieved
- However, significant failures were seen
 - FTS Optimiser started see-sawing up and down on the 'decision' over how many concurrent transfers to do
- From the site point of view the bandwidth was filled...but this included a large number of failing transfers
- Thursday, the number of concurrent transfers allowed in FTS was manually limited to 300



Site is interested in further testing

Future plans UK

- Test the Antares Tape endpoint
 - Will soon have new front end
 - Will be connected to the LHCOPN for the first time
 - Alternative methods needed on the VO side
- Contribute to testing as needed, e.g.
 - New 100 Gbps gateway machines at RAL (will start to replace 25Gbps gateways)
 - Other improvements or changes at any of the UK sites

Summary

- Testing the UK infrastructure was a successful and worthwhile undertaking
- The RAL disk endpoint performed well - big improvements since DC24
- Saturation of the LHCOPN; successful fallback to LHCONE
- Imperial College and RAL T2 filled their network
- Improvements made at Brunel - further testing required
- Slow per-file transfers?
- Lots of things still to improve!