Pre-DC24 tests: update

DOMA general, 6th Dec 2023
Katy, Lisa, Panos

T0 export test on Thursday 30th Nov

- As planned, tested CERN-EOS to each CMS Tier 1 site
- Started at half of the DC24 planned T0 export rate; then ramped it up to the full DC24 planned T0 export rate
- Test lasted 24 hours from 10am CET, coinciding with ATLAS
- Used the injector tool and “Data challenge” activity

Combined VO result (all transfers, all source/dest)

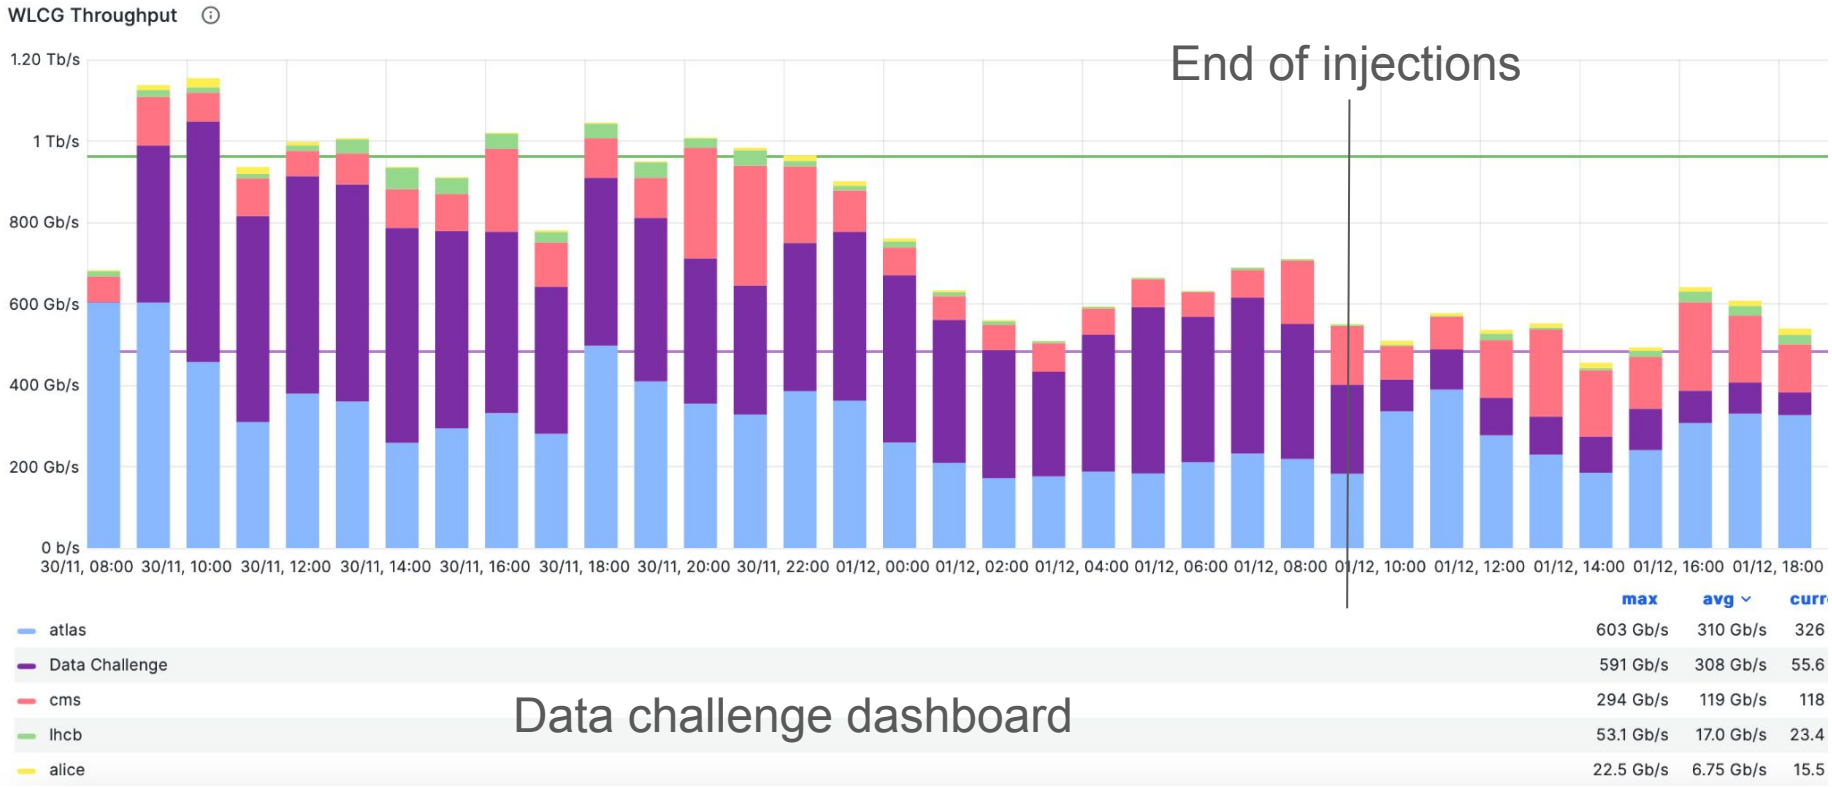

Data challenge dashboard

Timeline (CET)

10:00 - First submissions with half of the DC24 projected rate

10.53 - Increased outbound transfers from CERN from 1000 to 2000 (CERN FTS only)

14:00 - Realised that deletions were not happening so we were submitting rules that created no transfer, reducing the overall rate. Reaper pods restarted and deletions commenced

16.37 - Submissions with full DC24 projected rate

10:03 - Stopped submitting any transfers (Friday morning)

Results - CMS only, by activity

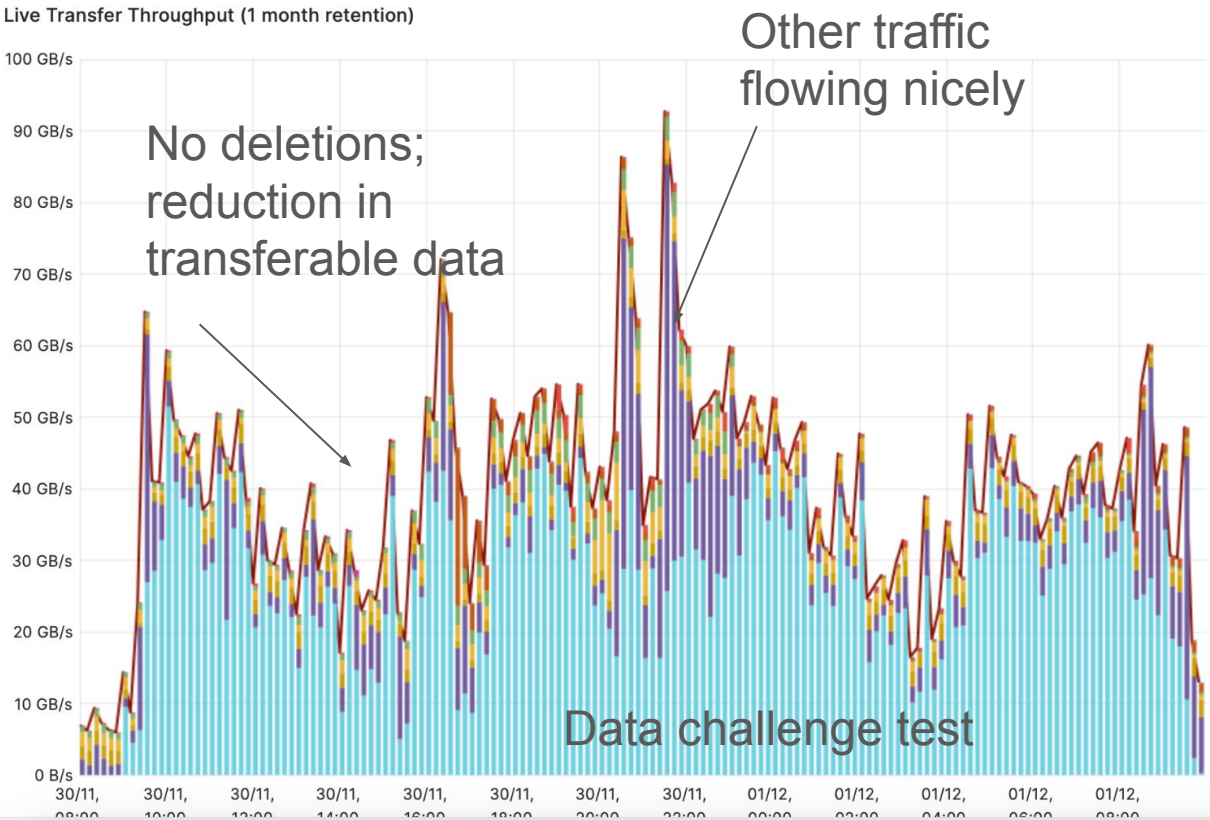

	max	avg
Total	92.7 GB/s	39.7 GB/s
Data Challenge	51.5 GB/s	27.0 GB/s
Production Output	59.6 GB/s	6.90 GB/s
T0 Tape	2.67 GB/s	1.83 GB/s
Analysis Input	9.76 GB/s	1.71 GB/s
ASO	3.37 GB/s	1.10 GB/s
User Subscriptions	21.9 GB/s	793 MB/s
Data rebalancing	3.38 GB/s	348 MB/s
Functional Test	74.3 MB/s	42.4 MB/s
Production Input	639 MB/s	6.16 MB/s
T0 Export	61.9 MB/s	3.49 MB/s
Data Consolidation	21.3 MB/s	293 kB/s
UNKNOWN	3.58 MB/s	118 kB/s

CMS only, DC test traffic only - by dest site

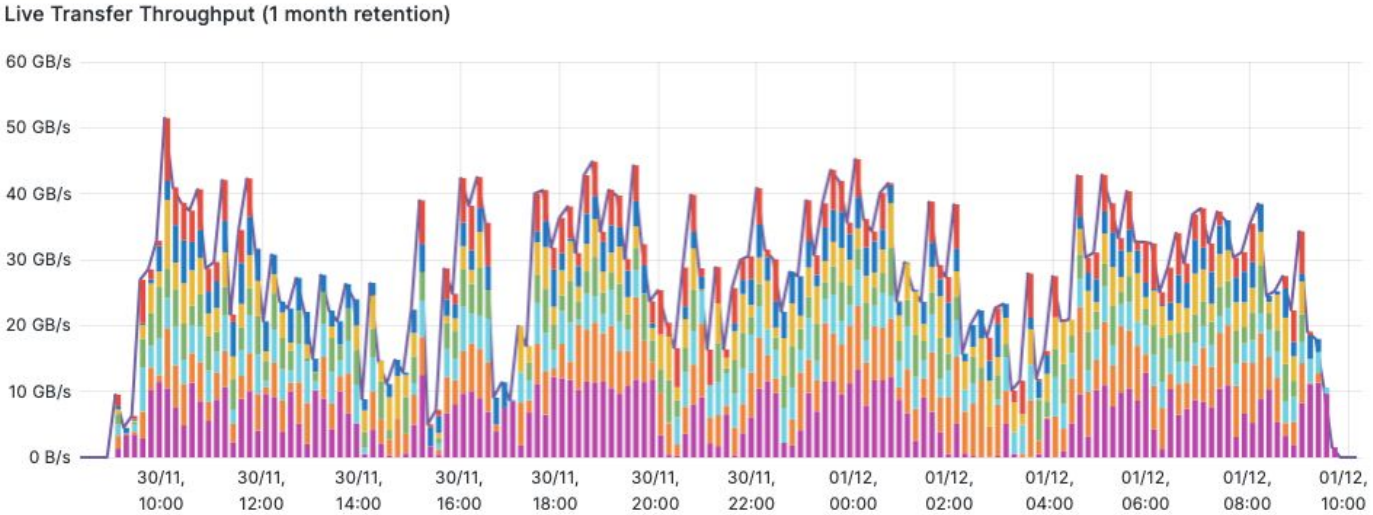

	max	avg
Total	51.5 GB/s	26.8 GB/s
T1_US_FNAL_Disk	13.4 GB/s	6.32 GB/s
T1_IT_CNAF_Disk	13.2 GB/s	4.94 GB/s
T1_FR_CCIN2P3_Disk	6.80 GB/s	3.41 GB/s
T1_DE_KIT_Disk	7.82 GB/s	3.37 GB/s
T1_ES_PIC_Disk	10.4 GB/s	3.33 GB/s
T1_UK_RAL_Disk	6.51 GB/s	2.75 GB/s
T1_RU_JINR_Disk	9.55 GB/s	2.71 GB/s

Summary of results by site

(1-hour binning)

(Obtained by eye)

Reminder: pressure dropped

SIte / T0 export GB/s	Half target	Full target	Peak rate	Sustained rate	Avg (24 hours)
T1_DE_KIT_Disk	1.626	3.252	4.79	~4	3.4
T1_ES_PIC_Disk	0.650	1.301	5.83	~4	3.35
T1_FR_IN2P3_Disk	1.675	3.349	5.18	~4	3.44
T1_IT_CNAF_Disk	2.114	4.227	8.05	~7	4.97
T1_RU_JINR_Disk	1.801	3.602	5.90	~4	2.73
T1_UK_RAL_Disk	1.257	2.513	5.12	~4	2.77
T1_US_FNAL_Disk	6.503	13.007	11.4	~10	6.37
Total	15.625	31.25			

Observations

- According to FTS, the transfer success rate was very high
- Outbound connections from CERN was set relatively low (1000, then 2000)
- FNAL was the only site not using the CERN FTS (has own FTS)
- All European sites (low latency) hit their full rate targets for T0 export
  - During peaks AND for the average across 24 hours!
- JINR hit the target during peaks and could easily sustain the target rate, but suffered particularly from the deletion problem (more on this later)
- FNAL was close to the target - may respond better to larger datasets and longer rule lifetimes?

Operational experience using the tool (Panos)

- Overall great experience
- Configuration:
  - Injection interval 15 minutes
  - Rule lifetime 2 hours
  - Expiration delay 30 minutes
- Logs of the various executions can be found here https://cernbox.cern.ch/s/ftt1RLQAkuRGKMX
- 3 different errors were observed:
  - Rule already exists - Mainly when re-starting the tool
  - Unique constraint (CMS_RUCIO_PROD.REPLICAS_PK) violated, which we suspect is a problem of deletions. We observed this error when we were using big datasets which probably didn't have enough time to be deleted
  - DID not found, which means that the dataset was not in the source anymore.

⇨ Following up these inconveniences in a Jira ticket

Summary

- A very successful and useful test
- CERN performed very well
- Many targets hit at T1s
- A lot of experience gained using the injector tool
- Ideas on how to further improve our experience for a smooth and relaxing DC24!
-

The next test…

- Not so much focussed on big rates
- To test what happens when we submit a large number of rules in Ruio (with the tool)
- How does Rucio cope? How does FTS cope?
- We get experience with a larger ‘menu’ - i.e. testing many more links with the injector tool than ever before.
- Would be good to do this before Christmas, because if there is a problem then we can repeat in January.
- No particular need to run jointly with ATLAS unless they would like to!