

DC 2024 report from CMS

Lisa Paspalaki¹, Katy Ellis², Hasan Ozturk³, Panos Paparrigopoulos³, Rahul Chauhan³, Andres Manrique Ardila⁴,

*with inputs from many teams,(site admins, FTS team, wlcg team)

1. Purdue University , 2, Science and Technology Facilities Council STFC (GB) 3. CERN, 4. Fermilab

DOMA General Meeting

06/03/24

Preface

- CMS took part in the Data Challenge 2024 (12-23 of February)
- Previous challenge 10% of HL-LHC
- **Target goal: 25% of HL-LHC**
- In communication with other VOs
- Pushed to the max our current infrastructure-- make sure we can achieve our target goals
- Rucio-FTS was used -- tested in an new scale
- Used the rucio injector tool
- Test new infrastructure tools -- such as token authorization and monitoring
- Several pre-tests was performed to assure success of the challenge


CMS time Plan

Date	12 Feb	13 Feb	14 Feb	15 Feb	16 Feb	17 Feb	18 Feb	19 Feb	20 Feb	21 Feb	22Feb	23 Feb
	T0 export	T0 export	T0 export	T1 export	T1 export	T1 export	T1 export	AAA	T0 export	T0 export	T0 export	T0 export
			T1 export		Prod. output	Prod. output	Prod. output		T1 export	T1 export	T1 export	T1 export
									Prod. output	Prod. output	Prod. output	Prod. output
									AAA	AAA	AAA	AAA
Scenario(s)	1	1	1,2	2	2,3	2,3	2,3	4	1,2,3,4	1,2,3,4	1,2,3,4	1,2,3,4
Rate (GB/s)	31	31	62	31	62	62	62	31	125	125	125	125
Rate (Gb/s)	250	250	500	250	500	500	500	250	1000	1000	1000	1000

- Each day(s) had a different menu, different configuration on the injector tool
- Variety of type of transfers and combinations, (T0 export, T1s to T1s and T1s to T2s, AAA)
- Plan was in synchronization with ATLAS, no AAA test for ATLAS
- constant monitoring of the challenge by Data Challenge Team


Data Challenge Dashboard

WLCG Throughput ⓘ




	max	avg	current
Data Challenge	2.19 Tb/s	990 Gb/s	19.9 Gb/s
atlas	625 Gb/s	308 Gb/s	385 Gb/s
alice xrootd	349 Gb/s	116 Gb/s	123 Gb/s
cms xrootd	191 Gb/s	68.0 Gb/s	73.6 Gb/s
cms	271 Gb/s	57.4 Gb/s	86.0 Gb/s
belle	38.9 Gb/s	9.50 Gb/s	1.08 Gb/s

CMS Overview



- [CMS-FTS Dashboard](#)
- Grouped by activity
- Activity “Data Challenge”


CMS Overview



- [CMS-FTS Dashboard](#)
- Group by site destination

Example: Day 1 and 2 T0 export

Overall target 31.25 GB/s




Dest	Target Rate (GB/s)	Avg rate (GB/s)
T1_DE_KIT_Disk	3.252	3.69
T1_ES_PIC_Disk	1.301	1.59
T1_FR_CCIN2P3_Disk	3.349	3.63
T1_IT_CNAF_Disk	4.227	3.90
T1_RU_JINR_Disk	3.602	4.22
T1_UK_RAL_Disk	2.513	1.69
T1_US_FNAL_Disk	13.007	14.5

We exceed the expected rates for most of the sites on the first day for T0 export

Day 3- T0 and T1 exports

Overall target 62.5 GB/s




[Dashboard](#)

Day 5-7 - T1 and production output

FTS was limiting the number of parallel transfers


Below target rate

Changing FTS configuration achieved higher rates



Transfer Throughput

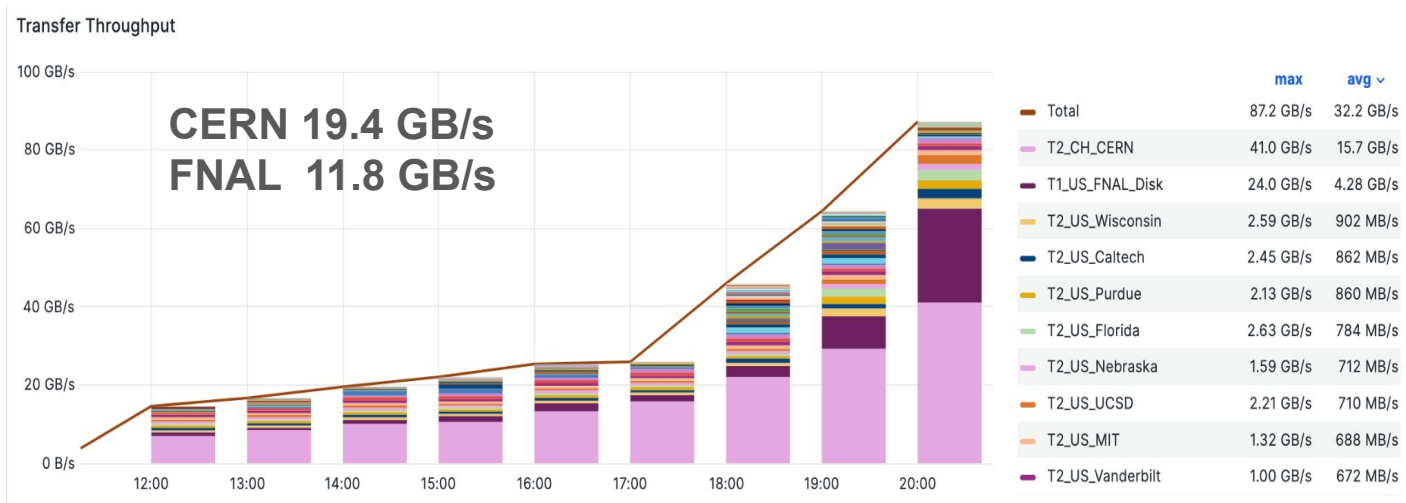
Overall target 62.5 GB/s



	max	avg
Total	70.5 GB/s	48.1 GB/s
T2_CH_CERN	10.2 GB/s	7.94 GB/s
T1_US_FNAL_Disk	5.95 GB/s	3.54 GB/s
T2_UK_London_IC	3.62 GB/s	2.51 GB/s
T2_DE_DESY	2.76 GB/s	2.09 GB/s
T1_RU_JINR_Disk	2.50 GB/s	2.02 GB/s
T1_IT_CNAF_Disk	3.01 GB/s	1.95 GB/s
T1_DE_KIT_Disk	2.80 GB/s	1.77 GB/s
T1_FR_CCN2P3_Disk	2.98 GB/s	1.72 GB/s
T2_US_Vanderbilt	2.54 GB/s	1.35 GB/s
T2_CH_CSCS	1.64 GB/s	1.21 GB/s

[dashboard](#)

Day 8- AAA (CERN/FNAL to T1s, T2s)




Grouped by source rse

Rate was not achieved immediately because the injector tool used small files as input

When fixed, we achieved the rates

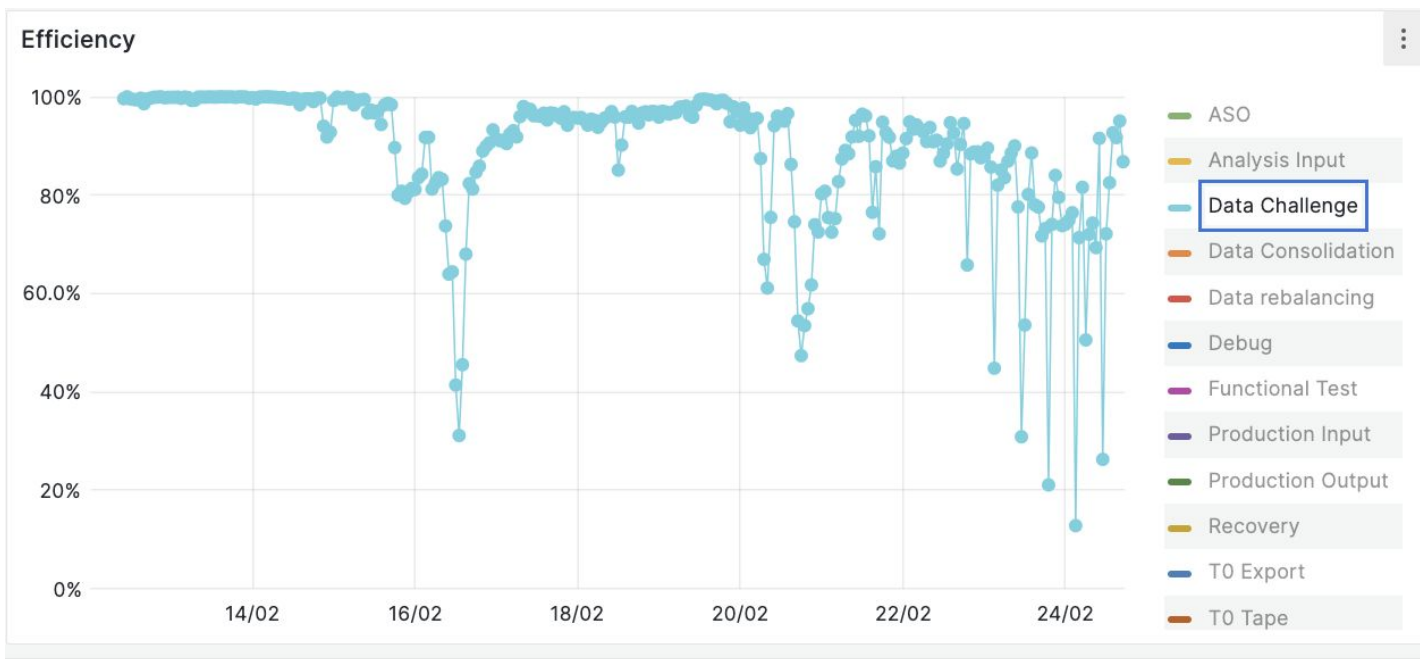
[dashboard](#)

Day 9-10 maximum rates



February 21st: Rucio could not handle deletions due to large backlog had to act to achieve target rates the next days

Efficiency of Activity Data Challenge




High efficiency through most of the challenge

Efficiency dropped when we start pushing, but reached high values soon after

Data Challenge Deletions

Deletion Throughput (1 month retention)



During the challenge we had to monitor closely the deletions and make sure that they can keep up

Issues-Bottlenecks during the challenge

- The challenge run without major issues for all the days
- Constant monitoring and fine tuning from DC team
- Issues discussed in [Rucio meeting](#)

Rucio - related issues

- Existing Reapers (8) could not keep up with deletions- had to increase them, sometimes have an specific reaper for RSE that were short of space

Future: grouping RSEs with specific pods

- Size of files affects the number of FTS requests and thus the achieved rate
- Had to kill and resubmit with larger files to achieve the target rates

Future: find ways to better monitor the submitter backlog

General Remarks

- some sites were out of production-- some of the extra traffic was to compensate for this
- some sites wanted to test their network at higher rate -- hence pushing up the injection
- We plan to do a more detailed analysis focusing at individual sites and whether they achieved individual targets

Summary - Milestones Achievements

- **Successful test for CMS and other VOs**
- Tests our system to the maximum
- First challenge with token authentication for 25 sites!
- Had to monitor and fine tune parameters- we understand our system better
- Started to discuss the paper structure
- A lot post mortem analyses will follow, GDC meeting and DC workshop