

Data Challenge 2024: bilan au WLCG workshop

D. Bouvet
L. Duflot



Les objectifs

- ◆ Tester si les infrastructures actuelles peuvent atteindre 25 % des besoins du HL-LHC, soit *en moyenne* :
 - ◆ RAW data export $T0 \rightarrow T1$: 1000Gb/s \rightarrow 250Gb/s
 - ◆ Distribution des données autres $T0 \leftrightarrow T1 \leftrightarrow T2 \leftrightarrow T2$: 400 Gb/s \rightarrow 100 Gb/s
- ◆ Pour atteindre ces chiffres en moyenne, il faut prendre de la marge : ne pas saturer le réseau et prendre en compte que la charge n'est pas constante dans le temps
- ◆ **Objectif 1 (modèle minimal) : 1.2Tb/s**
- ◆ Pour les échanges hors RAW, ces chiffres correspondent au besoins vitaux, la situation actuelle est qu'ont transfert beaucoup plus de données pour accéder aux ressources de calcul
- ◆ **Objectif 2 (modèle flexible) : 2.4Tb/s**



Schedule

	Monday 12/02/2024	Tuesday 13/02/2024	Wednesday 14/02/2024	Thursday 15/02/2024	Friday 16/02/2024	Saturday 17/02/2024	Sunday 18/02/2024
ALICE	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1
ATLAS	T0 → T1	T0 → T1	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2
CMS	T0 → T1	T0 → T1	T0 → T1 → T2	T1 → T2	T1 ↔ T2	T1 ↔ T2	T1 ↔ T2
LHCb		T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1
DUNE	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2
Belle II	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1

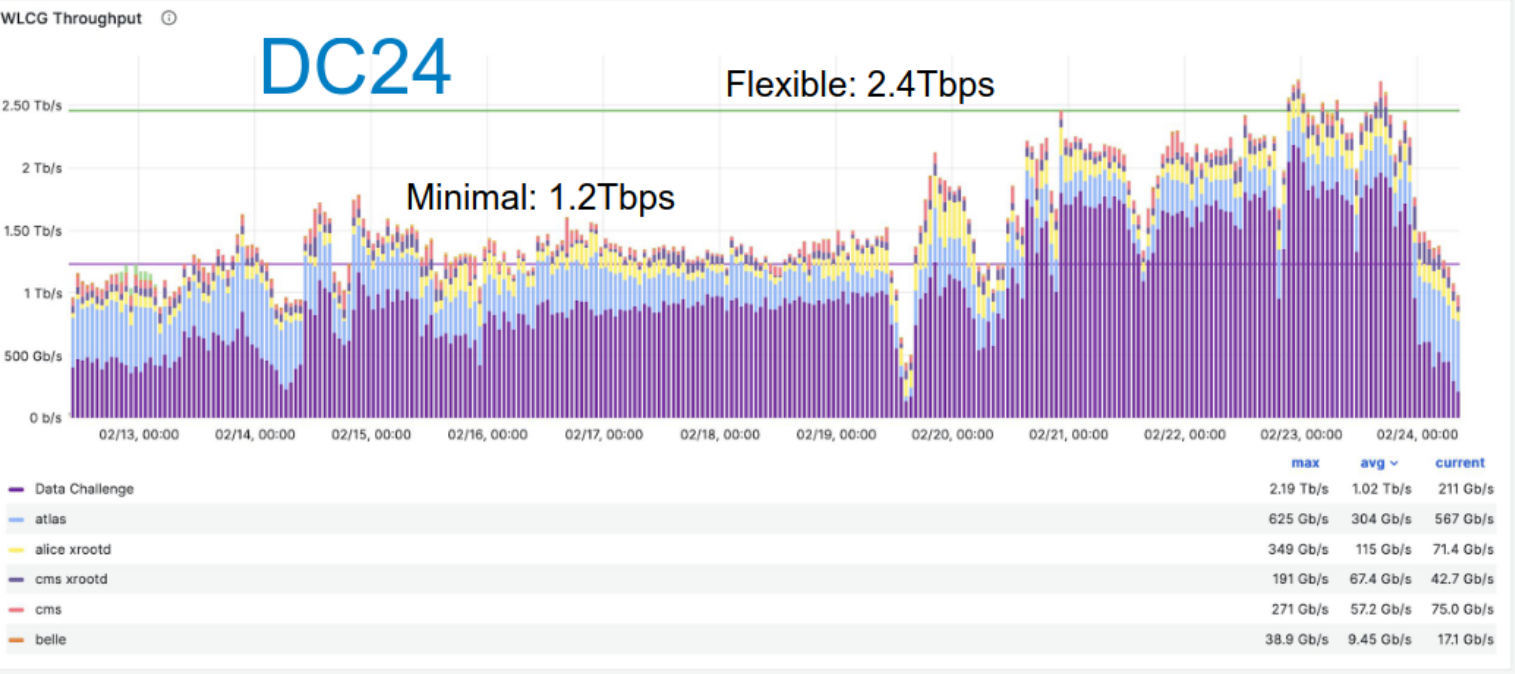
SUMMARY							
T0 exports minimal rates (ALICE+ATLAS+LHCb+CMS)	529.7 Gbps	650.3 Gbps	650.3 Gbps	650.3 Gbps	650.3 Gbps	650.3 Gbps	650.3 Gbps
T0 exports (DUNE + Belle II)	18.5 Gbps (belleII)	18.5 Gbps (belleII)	18.5 Gbps (belleII)	18.5 Gbps (belleII)	18.5 Gbps (belleII)	18.5 Gbps (belleII)	18.5 Gbps (belleII)

	Monday 19/02/2024	Tuesday 20/02/2024	Wednesday 21/02/2024	Thursday 22/02/2024	Friday 23/02/2024	
ALICE	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	yellow: "reduced minimal" (only T0 export)
ATLAS	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	blue: minimal scenario
CMS	AAA T1 → T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	red: flexible scenario
LHCb	T0 → T1	T1 Tape Recall	T1 Tape Recall	T1 Tape Recall	T1 Tape Recall	
DUNE	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	T0 → T1 → T2	
Belle II	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 == SURF , T1 == FNAL, T2 == Storage sites

SUMMARY					
T0 exports high rates (ALICE+ATLAS+LHCb+CMS)	449.56 Gbps	895.56 Gbps	895.56 Gbps	895.56 Gbps	895.56 Gbps



DC24 - main result



DC24 met the (main) goals:

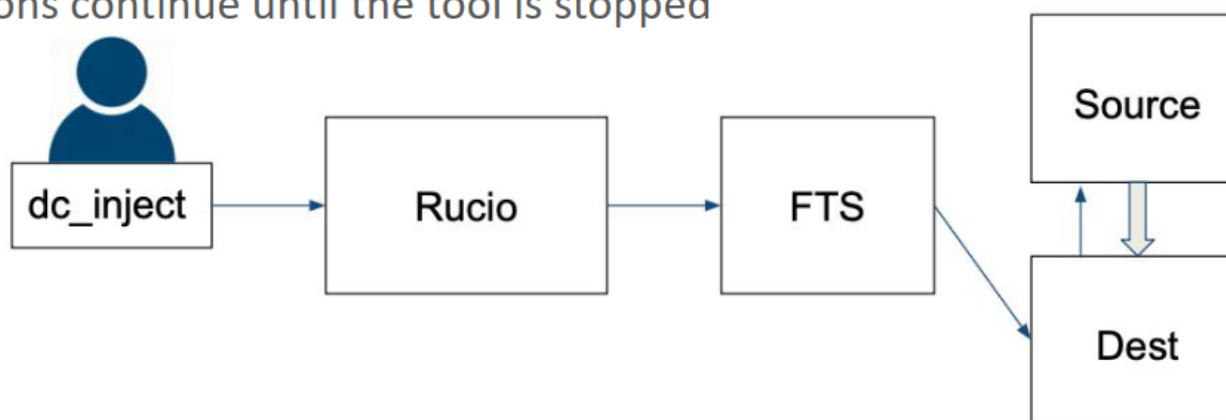
- Achieved full throughput of minimal model (1st week)
- Push for flexible target (2nd week)




Point de vue des expériences

Talk

- ALICE continued normal production work, transferring HI datasets CERN -> Tier 1s
 - Using XRootD (not FTS) with ALICE tokens and the jAliEn transfer system (34PB)
 - Tuning via the number of parallel streams
- LHCb queued up blocks of transfers from CERN to Tier 1s (disk then tape), then read the data back from tape to disk
 - Using FTS and Dirac
- ATLAS and CMS used a tool written by Mario to submit Rucio rules periodically
 - New Rules injected every 15 minutes
 - Data had lifetime of 1-3 hours (before eligible for deletion)
 - Choose to have Rucio select large datasets first (== large files first)
 - Injections continue until the tool is stopped




11



ALICE: Ongoing data export from 2023 HI run

Time evolution T1s




Centre	Target rate GB/s	Average achieved GB/s
CNAF	0.8	0.98 (+20%)
IN2P3	0.4	0.6 (+40%)
KISTI	0.2	0.25 (+22%)
GridKA	0.6	1.12 (+90%)
NDGF	0.3	0.35 (+15%)
NL-T1	0.1	0.25 (+150%)
RAL	0.1	0.58 (+500%)
CERN	10	14.2 (+40%)

DC24 was a great success for ALICE, achieving above target rates at every site, with minimal interference, and no effect on other activities




EOS -> Disk link



- ▶ Target throughput (14GiB/s) was achieved during the first day
- ▶ Lower throughput later
 - ▶ Some sites finished transferring their part during the first day so were no longer contributing to overall throughput
 - ▶ Submissions were slow and not optimal
 - ▶ Submission agent got stuck a few times, that was also a contributing factor

3

Disk -> Tape link




- ▶ Target threshold (14GiB/s) crossed several times
 - ▶ Max around 35GiB/s
 - ▶ Spikier throughput because of the nature of the link and submission agent problems

5

Objectifs atteints mais pas assez de volume de donnée pour le voir sur une longue période

Staging




- ▶ Target throughput (9.58 GiB/s) was achieved during the first two days of the test
- ▶ Lower throughput later
 - ▶ Some sites finished transferring their part and were no longer contributing

8

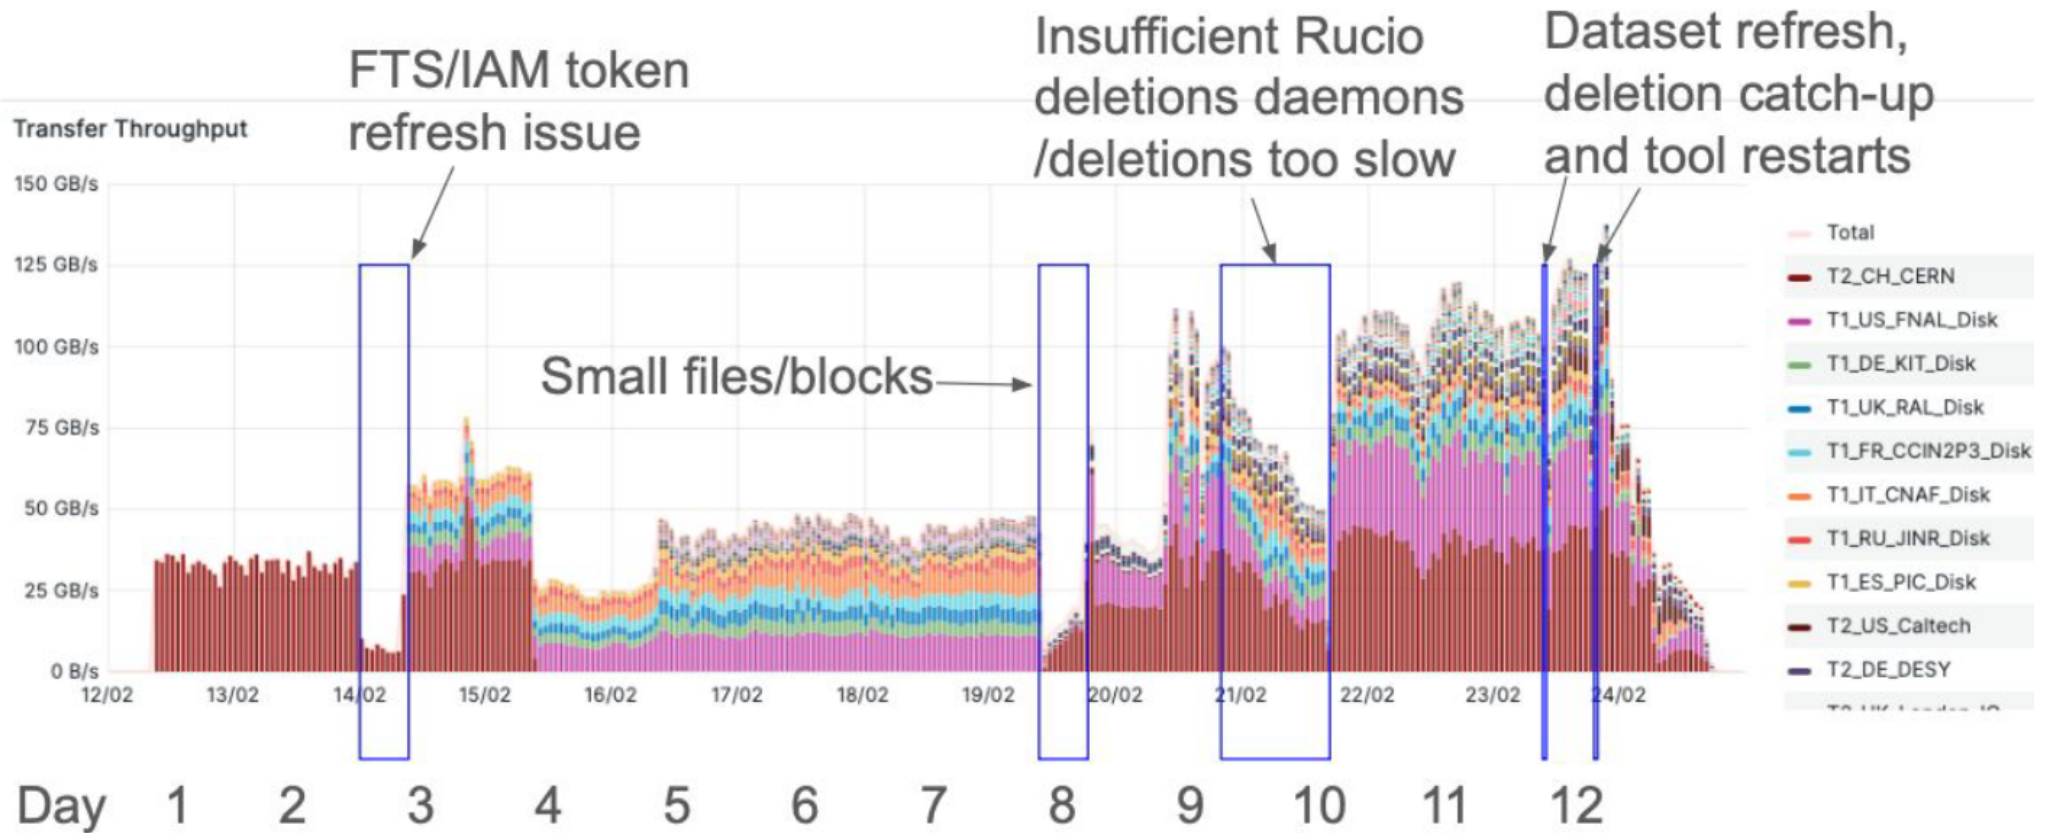


- Daily exercise menu with increasing complexity
- T0 export, T1s to T1s and T1s to T2s, AAA
- First week targets were mostly met easily
- Overall target of ~125GB/s was reached with significant effort
 - A few hundred links maximum (Prod + DC)
 - More data injected than the target required

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



CMS - notable issues




These marked periods were excluded as 'known problems' when analysing results



- Generally considered success for highlighting bottlenecks, though rates hampered by the really large number of links

- Injections on >1200 links every 15m
 - ~2000 links with production
- Short data sets lifetime 1h -> 2h -> 3h
- Helped highlighting problems that wouldn't have been seen otherwise


Transfers Throughput (all final states from enr_complete) ⓘ



- None of the bottlenecks were due to the network specifically
 - Some sites had the LHCOPN link down but had alternative paths
- Some sites struggled mostly due to storage limitations
 - 17 problems were reported on GGUS
- T0 export rates were not achieved
 - Re-run T0-T1 export tests post DC (see later)



ATLAS - notable issues



CMS - Site analysis

- Some sites better source than destination and vice versa (compared to expectation)
- RAL (UK) subsea cable to the LHCOPN was broken throughout the first 4 days
- JINR (RU) may have had issues they did not mention; possible network issues in general
- CNAF (IT) problem with number of connections and FTS not pulling back after failures
- IN2P3 (FR) config in FTS not allowing sufficient connections (CMS team at fault)

Day	Scenario	JINR		FNAL		IN2P3		RAL		PIC		KIT		CNAF	
		DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC
1	T0 Export	1.42	N/A	1.13	N/A	1.09	N/A	0.76	N/A	1.18	N/A	1.16	N/A	1.17	N/A
2	T0 Export	1.46	N/A	1.12	N/A	1.10	N/A	0.50	N/A	1.17	N/A	0.94	N/A	1.17	N/A
3	T0Export, T1Export	1.31	0.62	1.08	0.88	1.33	1.03	0.72	0.99	1.18	1.06	1.10	1.06	1.28	0.93
4	T1 Export	N/A	0.37	N/A	0.91	N/A	1.12	N/A	0.76	N/A	1.05	N/A	0.95	N/A	1.00
5	T1-Export, Prod-out	1.18	1.72	1.15	0.87	1.25	0.89	0.98	1.01	1.21	1.09	1.23	0.77	1.17	0.77
6	T1-Export, Prod-out	1.14	2.42	1.18	0.88	1.47	0.88	0.72	0.81	1.17	1.03	1.19	0.76	1.18	0.95
7	T1-Export, Prod-out	1.19	2.19	1.15	0.87	1.22	0.87	0.81	1.04	1.20	0.98	1.21	0.73	1.16	1.02
8	AAA	1.30	N/A	N/A	1.10	1.39	N/A	1.31	N/A	1.31	N/A	1.70	N/A	1.32	N/A
9	All	0.38	0.34	0.87	0.84	0.57	0.57	0.95	1.02	1.25	0.86	0.86	0.56	0.65	0.25
10	All	0.70	0.34	0.98	0.74	0.58	0.65	0.56	0.99	0.70	0.66	1.03	0.98	0.63	0.28
11	All	0.63	0.33	0.91	0.73	0.43	0.76	0.77	1.05	1.09	0.84	0.91	1.09	0.69	0.24
12	All	0.40	0.54	0.92	0.86	0.89	1.00	0.85	1.15	1.21	0.87	1.13	0.89	0.78	0.29



ATLAS - Site analysis

- NDGF had a bug in the storage heavily affecting writing rates
- RAL had internal network and gateways problems
- BNL digesting files too quickly for the injected unprioritized rates, there was comb like patterns in the rates
- Day 8 was affected by FTS DB defrag 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%	



ATLAS - T0 export repeated tests

- T0 export rates are the most important use case and were not achieved
- The rates weren't achieved because they were queued behind production
 - T2 traffic is non negligible in ATLAS (42% dst, 25% src)
- Tests were repeated injecting one site at the time
 - Rates improved for the majority of sites
- Some differences:
 - SARA was testing 800 Gb/s after DC24; was injected with much larger rates
 - RAL wanted to test writing directly to tape in the second test; other limitations were identified
 - NDGF resolved the dcache bug that was affecting them

Site	T0 Export	DC24 best rates on day 1,2	% of expected rates	T0-T1 one T1 at the time	% of expected rates
<u>BNL-ATLAS</u>	60	<u>31.5</u>	53%	<u>61.3</u>	102%
<u>FZK-LCG2</u>	32	<u>26.4</u>	83%	<u>42.2</u>	132%
<u>IN2P3-CC</u>	38	<u>43</u>	113%	<u>50.9</u>	134%
<u>INFN-T1</u>	23	<u>19.3</u>	84%	<u>33.5</u>	146%
<u>NDGF-T1</u>	15	<u>13.8</u>	92%	<u>28.2</u>	188%
<u>SARA-MATRIX</u>	15	<u>12.2</u>	81%	<u>274.1</u>	1827%
<u>pic</u>	11	<u>12.3</u>	112%	<u>18.1</u>	165%
<u>RAL-LCG2</u>	38	<u>15</u>	39%	<u>27.2</u>	72%
<u>TRIUMF-LCG2</u>	25	<u>23.9</u>	96%	<u>27.2</u>	109%
T1 summary	257	197.4	77%	562.7	219%
T1 summary -SARA	242	185.2	77%	288.6	119%



- ◆ **DESY HH+ZN** : saturation du réseau sinon OK
- ◆ **KIT** :
 - ◆ +100G ajouté pour le challenge
 - ◆ Tuning dCache
 - ◆ Aucun problème
- ◆ **NGDF** :
 - ◆ Minor dCache Bug & HAProxy config broke 3rd party copies under load
 - ◆ Some network bottlenecks observed between CERN and NDGF
- ◆ **PIC** : pas de problème, tests de NOTED
- ◆ **SARA** :
 - ◆ Minor dCache (tuning) related issues for LHCb
 - ◆ Saturated internal link



INFN-T1

- DC24 very useful to identify bottlenecks – very much in favour of running such exercises + preparatory tests and repetitions of the tests
- However, DC is a stress test greatly impacting sites and overloading storage endpoints
 - SAM tests should be aware of ongoing DC.
 - Significant production load during the challenge – sometimes heavy impact
- Wish for FTS to regulate injection based on success rate to avoid overloading storage endpoints
- Differences between FTS and site monitoring observed



- Planning to align the StoRM WebDAV instances dedicated to CMS (observer higher failure rates on servers with less cores)
- Re-think LHCb hardware configuration so to accommodate their workflow
- Recently upgraded StoRM WebDAV in all the endpoints to improve efficiency + planning to introduce performance markers in StoRM WebDAV



RAL

- DC24 was very successful
- By end of DC24, throughput was 2 – 3 times higher than ever achieved before
- LHCOPN link cut under the sea, fixed by February 16
 - Failover to LHCONE needed tuning
- Tape services during DC24 meet expectations. Post DC24 tape tests for ATLAS helped to expose areas for improvement.
- Hit 25Gbit/s limitations of gateway hardware
 - 100Gbit/s capable hardware has been purchased
- Lots of experience gained, trying to balance the load across the gateways



Point de vue de Rucio

- ◆ Globalement pas de problèmes importants
- ◆ Le DC a mis en lumière des problèmes de contention entre daemons (soumission FTS, effacer les règles expirées) ou de gestions à améliorer pour les sites surchargés (effacement de fichiers)
- ◆



General overview and impressions



- The data challenge was a success for FTS
 - The sum of all FTS instances: ATLAS, CMS, LHCb and Pilot
 - 33 million file-transfers and 249 PB over a two week period
 - Best transfer concurrency provided by a single instance:
 - FTS ATLAS broke the previous record (also held by FTS ATLAS)
 - 20K concurrent transfers for 17 hours (previous record peak was 9K)
 - Half of the FTS transfers used token authentication
- There were challenges
 - Too much fire fighting behind the scenes with respect to **fts3-atlas.cern.ch**
 - Many thanks to the database-on-demand team for quickly increasing the DB RAM
 - Defragmentation of the **fts3-atlas.cern.ch** DB was not completed
 - FTS maximises concurrent-transfers per link whereas data challenges maximise data throughput



FTS ATLAS database was overloaded



- It was believed queuing more transfers would increase data throughput
- DB slowed and practically stopped during weekly defrag (every Monday at 10:00am)
- FTS token-refreshing was the main cause of the overload
 - FTS continually polled DB for near-to-expire tokens
 - SQL queries had not been optimised
- Urgent fix required because FTS ATLAS also handled non-DC24 transfers
- Database-on-demand team quickly increased DB RAM from 80 to 120 GB
- FTS team migrated a lower-priority FTS instance out of a high-performance DB server
- FTS team migrated FTS Pilot into the newly vacated DB server
- ATLAS split DC24 load across FTS ATLAS and Pilot



Slow FTS ATLAS optimizer



- FTS optimizer increases or decreases the amount of transfers on a link based on the link's current performance
- Usually takes 6 to 12 minutes for the optimizer to run
- 3 hour optimizer runs were observed during DC24
- Slow down was a linear function of file-transfer queue-length
- Hourly restart of daemons meant a full optimizer run could not complete
- Slowdown and restarts “froze” the majority of concurrency decisions
- **Data manager changes were effectively ignored**



- ◆ L'unité de base de FTS est le transfert individuel
- ◆ FTS tente de maximiser le *nombre* de transferts sur un lien / une destination, ce n'est pas la même chose que d'optimiser la bande passante globale (si les divers liens vers une destination ont une bande passante très différente)
- ◆ **Post-DC24**
 - ◆ Optimisation du « token refresher »
 - ◆ Regarde si peut implémenter un algorithme de « back-pressure »
 - ◆ Travail sur l'optimizer
 - ◆ Priorité sur les transferts



- ◆ Les implémentations dans rucio et FTS ont été terminées juste avant le DC24
 - ◆ Peu de temps pour tester en charge
 - ◆ Choix faits rapidement et pas concertés
 - ◆ Failles de sécurité possibles (dont des tokens apparaissant dans des logfiles...)
- ◆ FTS
 - ◆ En particulier le renouvellement des tokens a été problématique : les tokens arrivent avec des durées de vie très variables, il faut donc les renouveler rapidement → surcharge de la DB. Un pb de renouvellement a pour conséquence que FTS se retrouve avec des transferts en queue qui vont échouer
- ◆ ATLAS a arrêté les transferts par tokens au cours de la 2eme semaine pour tenter d'attendre l'objectif de bande passante



- ◆ « success » mais en fait l'objectifs de bande passante soutenue en 2eme semaine non atteint.
- ◆ 2eme semaine chaotique entre FTS qui essaie de maintenir en vie fts-atlas et ATLAS/CMS qui tentaient de pousser au max
 - ◆ Peu de temps pour regarder les T2, peu de temps pour discuter e.g. des sites FR
- ◆ Le canal de communication FR via chat.in2p3 n'a pas bien marché, difficile de savoir ce qui se passe
- ◆ Aujourd'hui : retour du CC (talks) et des T2 (si possible, dans les discussions), leçons à tirer pour configurer nos sites pour le DC26

