WLCG Data Challenge 24: LHC experiment experiences

WLCG workshop, DESY
2024-05-14

Katy Ellis (RAL), Alessandra Forti (U of Manchester), Alexander Rogovskiy (RAL), Latchezar Betev (CERN)
Christoph Wissing (DESY), Mario Lassnig (CERN)
On behalf of the DC24 community

Introduction: LHC and its High Luminosity upgrade

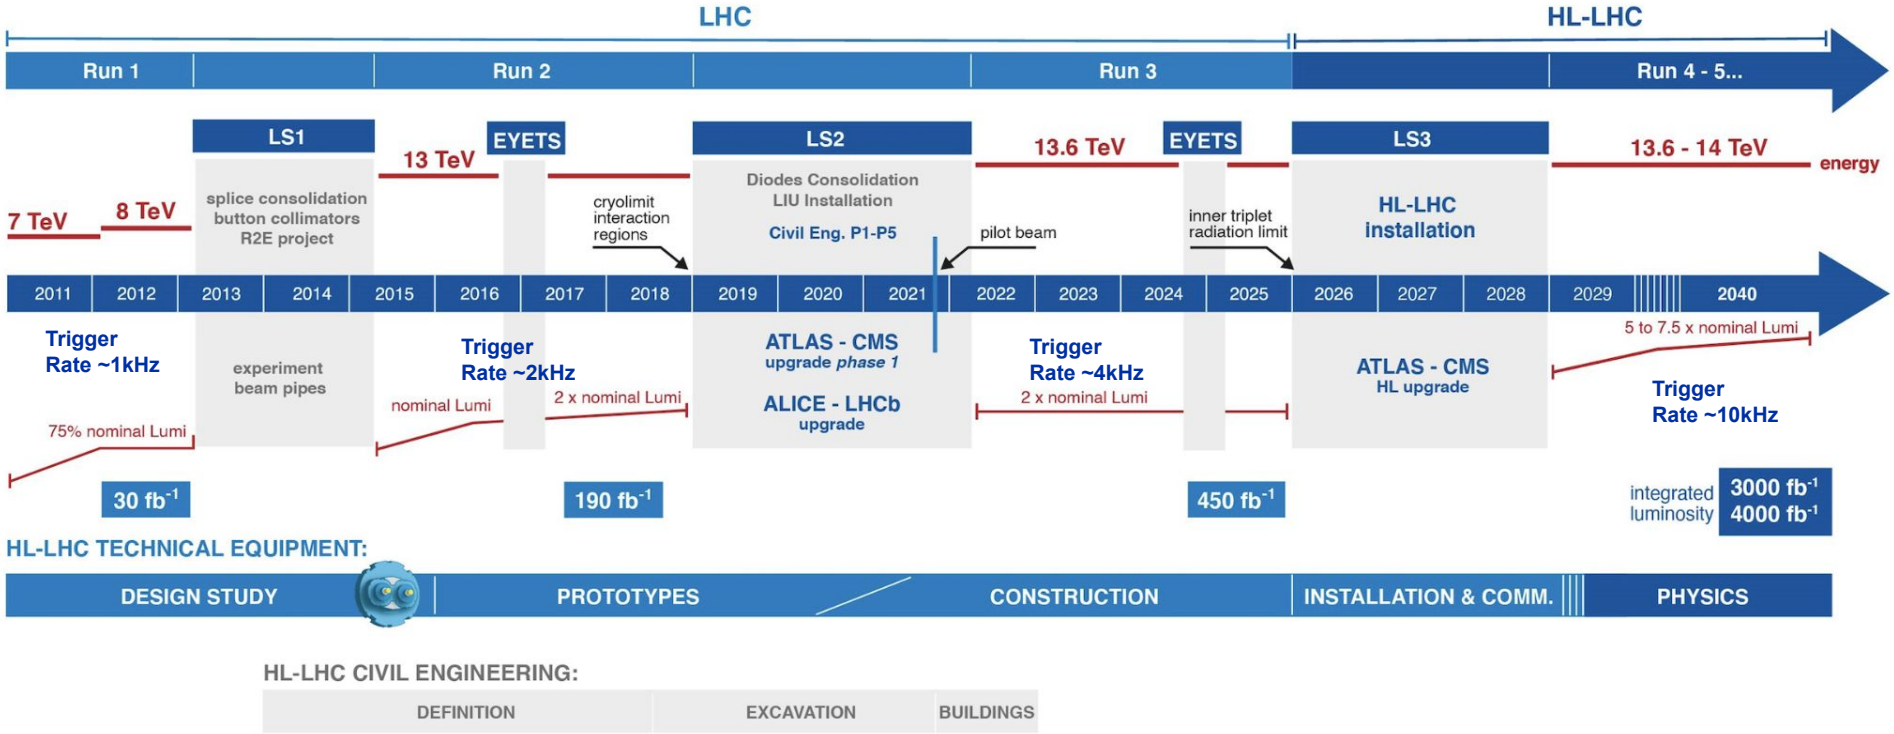

Figure adapted from:
Zerlauth, Markus & Bruning, Oliver. (2024). Status and prospects of the HL-LHC project. DOI; 615. 10.22323/1.449.0615.

Data Challenges for HL-LHC

- WLCG has mandated to execute data challenges (DC) for HL-LHC
  - Demonstrate readiness for expected HL-LHC data rates by a series of challenges
  - Increasing volume/rates
  - Increasing complexity (e.g. additional technology)
  - A data challenge roughly every two years

- DOMA is the coordination and execution platform
  - Data Organization Management & Access
    - Forum across all LHC experiments to address **technical** challenges
  - DC coordination across the LHC experiments and beyond
    - Suited dates
    - Reasonable targets
    - Functionalities
  - Help in orchestration
- No pressure on sites to increase their capacity
  - But can we improve the existing infrastructure?

Year	% of HL-LHC
2021	10
2024	25
2026?	50
2028?	100

Modelling the rates for HL-LHC (T0 export)

- ATLAS & CMS T0 export (T0 to T1s)
  - 350PB RAW per experiment, per year, taken and distributed during typical LHC uptime of 7M seconds
  - => 50GB/s or 400Gbps
  - Plus 100Gbps estimated for prompt, derived data
  - 1Tbps for CMS and ATLAS combined
- ALICE & LHCb T0 Export
  - 100 Gbps per experiment estimated from Run-3 rates

WLCG data challenges for HL-LHC - 2021 planning
https://zenodo.org/records/5532452

- Network needs to be bigger than the average, estimated rates:
  - Factor of 2 for *bursts*
  - Another factor of 2 for *overprovisioning*
- But the challenges only need to fill 50% of the network requirement
  - (So remove one factor of two when calculating rates for data challenges)

- T0 export HL-LHC target for data challenges:
  - 2*(1000Gbps (ATLAS and CMS) + 400Gbps (LHCb and ALICE)) = 2.8Tbps

Modelling the rates for HL-LHC (additional)

- ATLAS & CMS T1 export (T1s to T2s)
  - The same data exported to T1s will later be exported to T2s for reprocessing over the same period
  - Same cumulative rate as T0 export

- **Minimal Model**
  - T0 -> T1s -> T2s : **4.8Tbps for the expected HL-LHC requirement**

- **Flexible Model**
  - Adds traffic relating to Monte Carlo outputs, etc.
    Means doubling the Minimal Model: **9.6Tbps for the expected HL-LHC bandwidth needs**
    Data flows over many links (source/destination pairs)

- DC24 target: 25% of the above = 1.2Tbps (minimal), 2.4Tbps (flexible)

Planning DC24

- Preparation started one year in advance
  - Monthly DOMA General checkpoints
  - Dedicated workshop in Nov 2023
- Agreement on dates
  - 2 weeks before beam operation in 2024
- Some agreement on schedule
  - See next slide
- Full transfers from disk to disk
  - Not just artificial network traffic
  - VOs used real data
- Experiments had room to define their goals
  - ALICE and LHCb involved tapes
  - ATLAS and CMS did not
- Preparation of monitoring
- Rehearsals and pre-tests

DC24 workshop slides

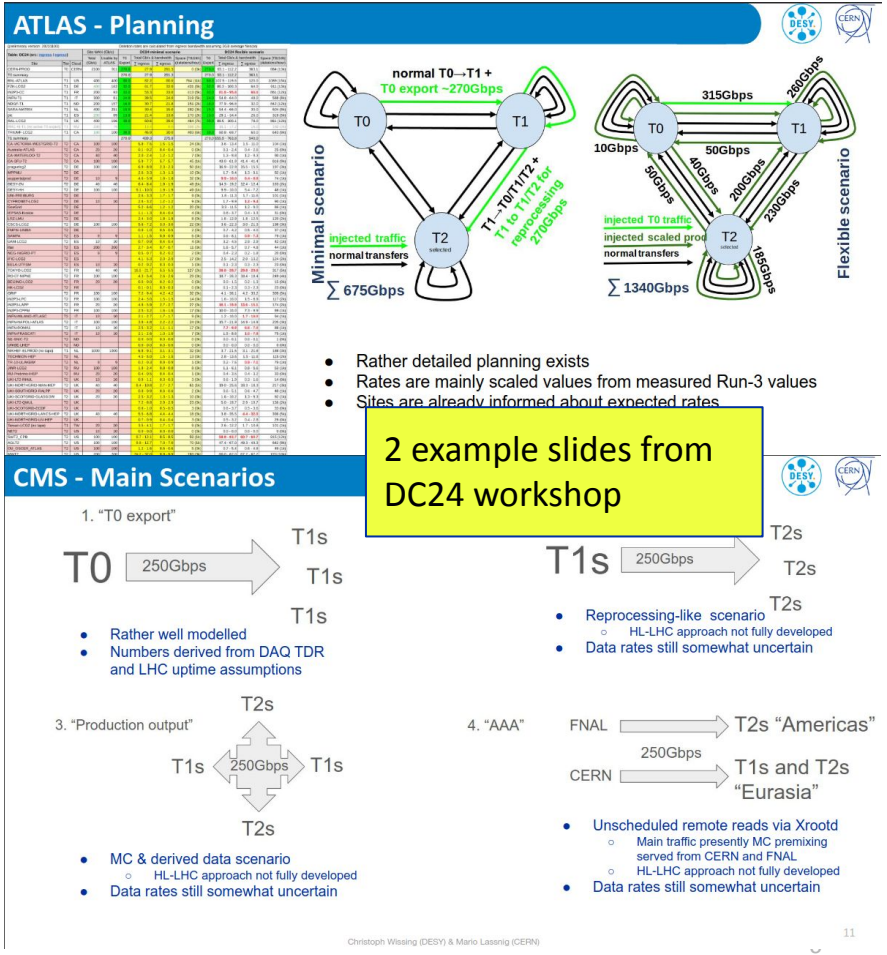

Schedule

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
	12/02/2024	13/02/2024	14/02/2024	15/02/2024	16/02/2024	17/02/2024	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	Tuesday	Wednesday	Thursday	Friday	
	19/02/2024	20/02/2024	21/02/2024	22/02/2024	23/02/2024	yellow: "reduced minimal" (only T0 export)
ALICE	T0 → T1	T0 → T1	T0 → T1	T0 → T1	T0 → T1	blue: minimal scenario
ATLAS	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	T0 ↔ T1 ↔ T2	red: flexible scenario
CMS	AAA T1 → T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	T0 → T1 ↔ T2	
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	

Pre-tests

No time to talk about these!

DC24 - main result

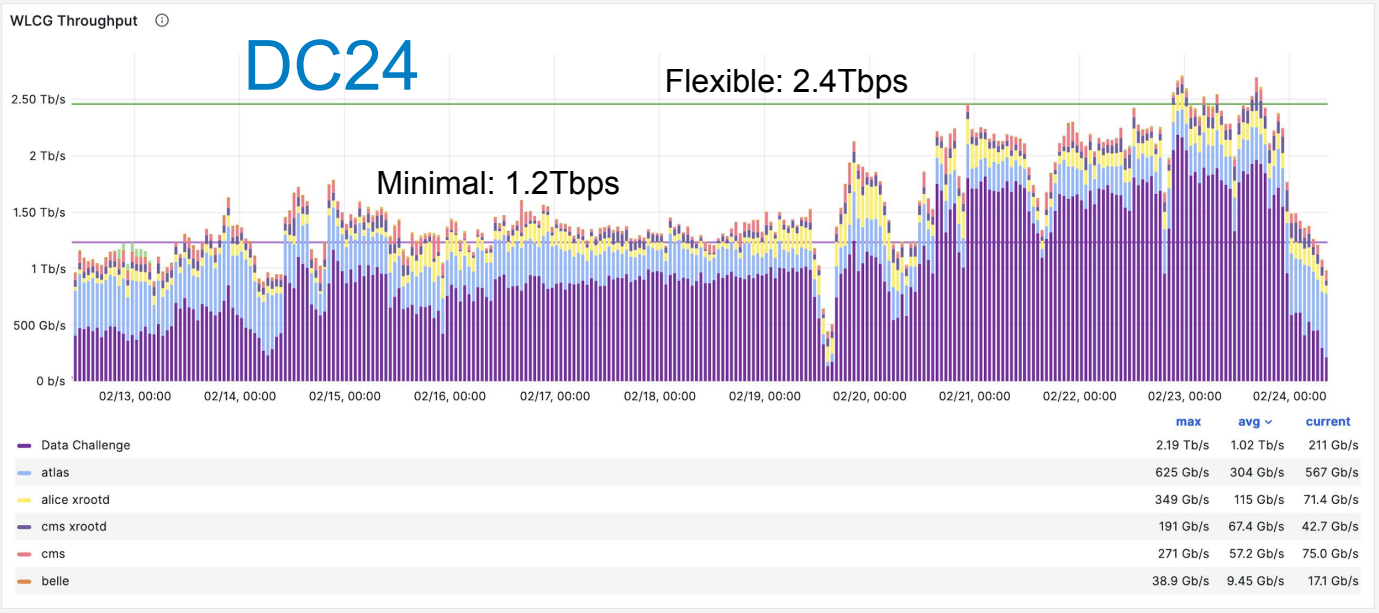

DC24 met the (main) goals:

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

DC24 vs DC21

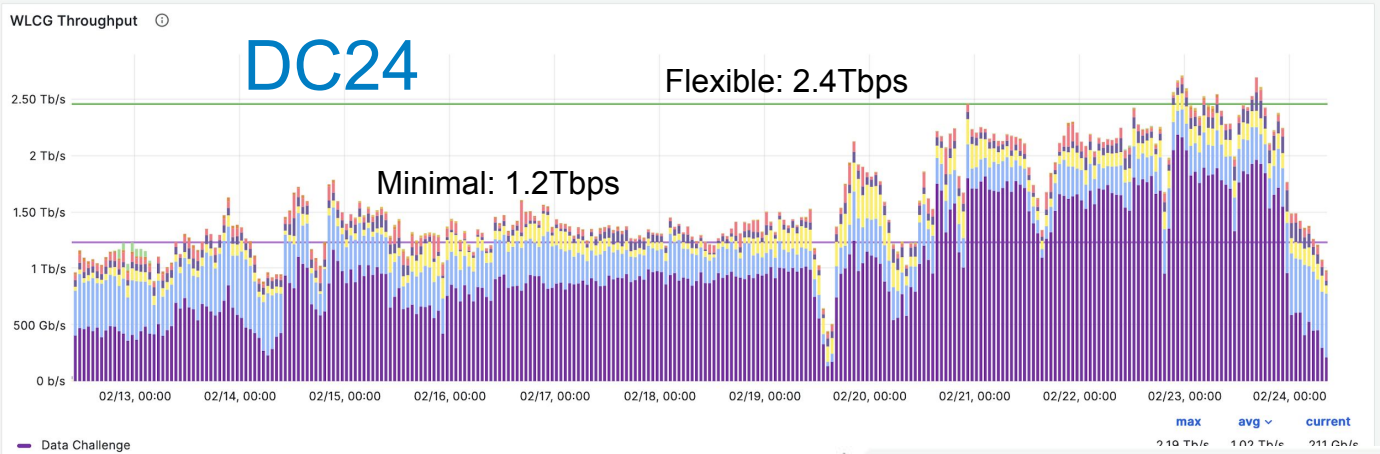

DC24 met the (main) goals:

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

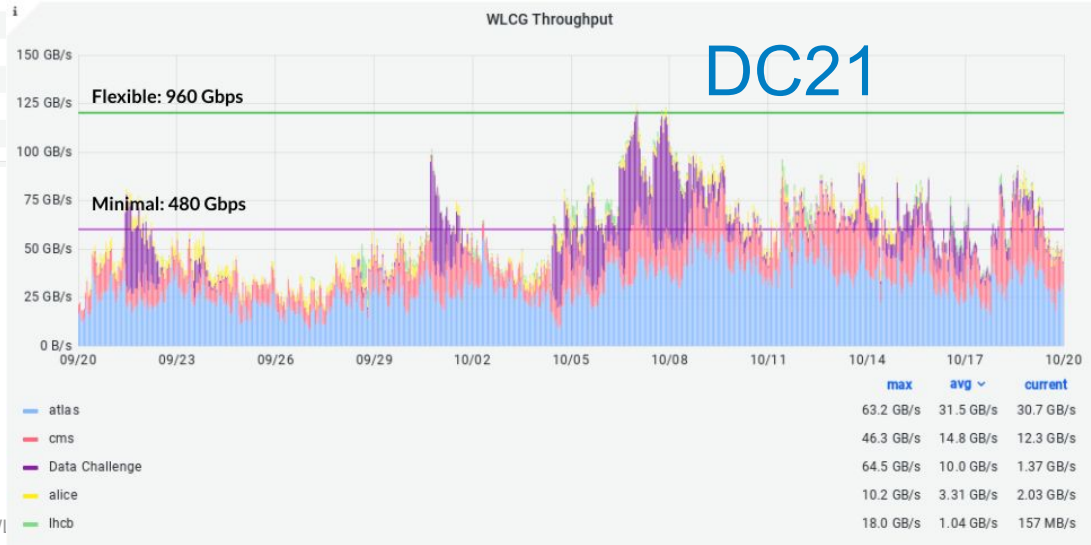

Network Data Challenges 2021 wrap-up and recommendations
https://zenodo.org/records/5767913

Methods and tools

- 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

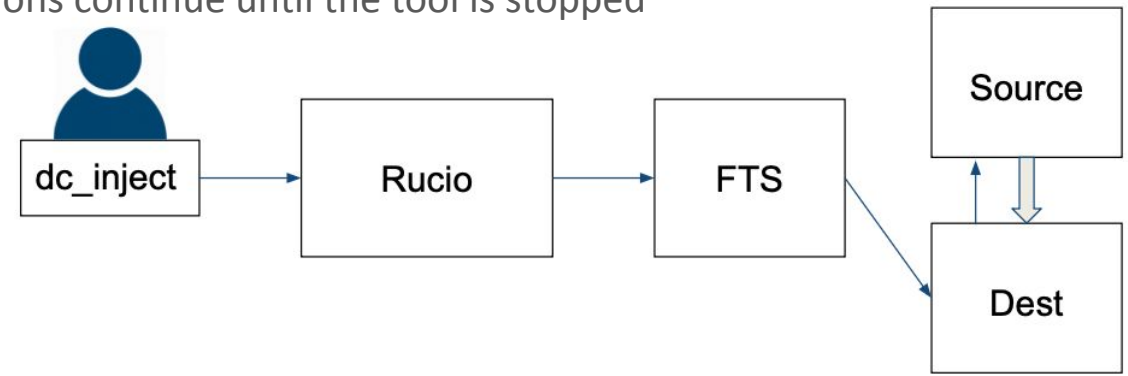

VO targets and results

ALICE: Ongoing data export from 2023 HI run

Time evolution T1s

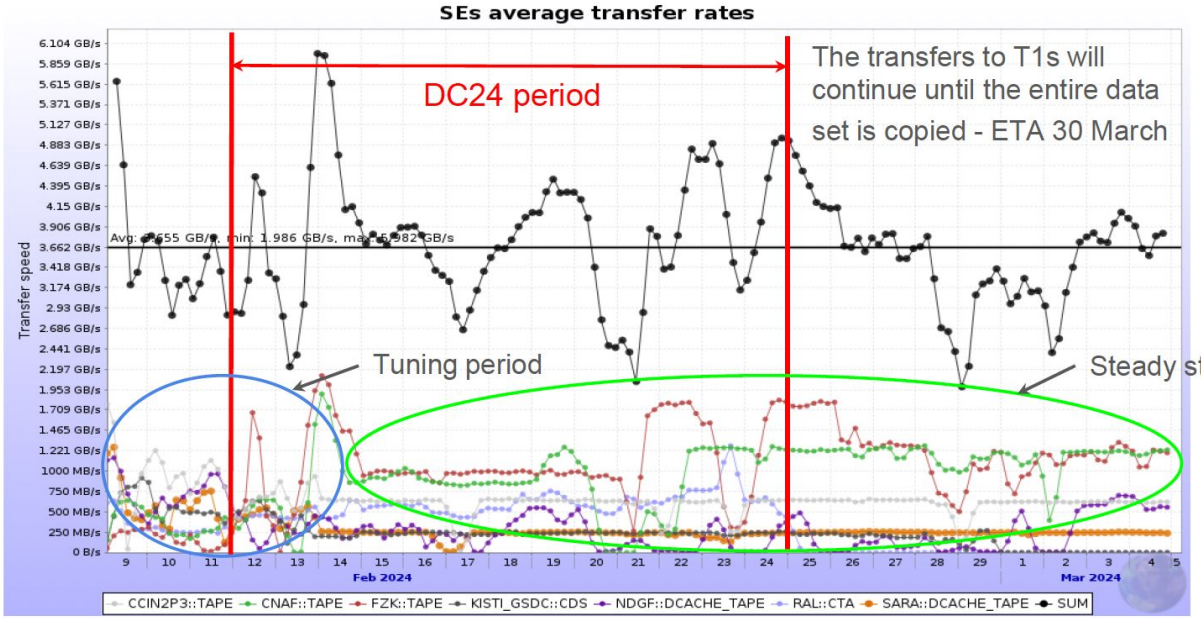

Centre	Target rate GB/s	Average achieved GB/s
CNAF	0.8	0.98 (+20%)
IN2P3	0.4	0.6 (+40%)
KISTI	0.2	0.25 (+22%)
GridKA	0.6	1.12 (+90%)
NDGF	0.3	0.35 (+15%)
NL-T1	0.1	0.25 (+150%)
RAL	0.1	0.58 (+500%)
CERN	*10*	*14.2 (+40%)*

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

ALICE: Transfers to Tier 0 tape

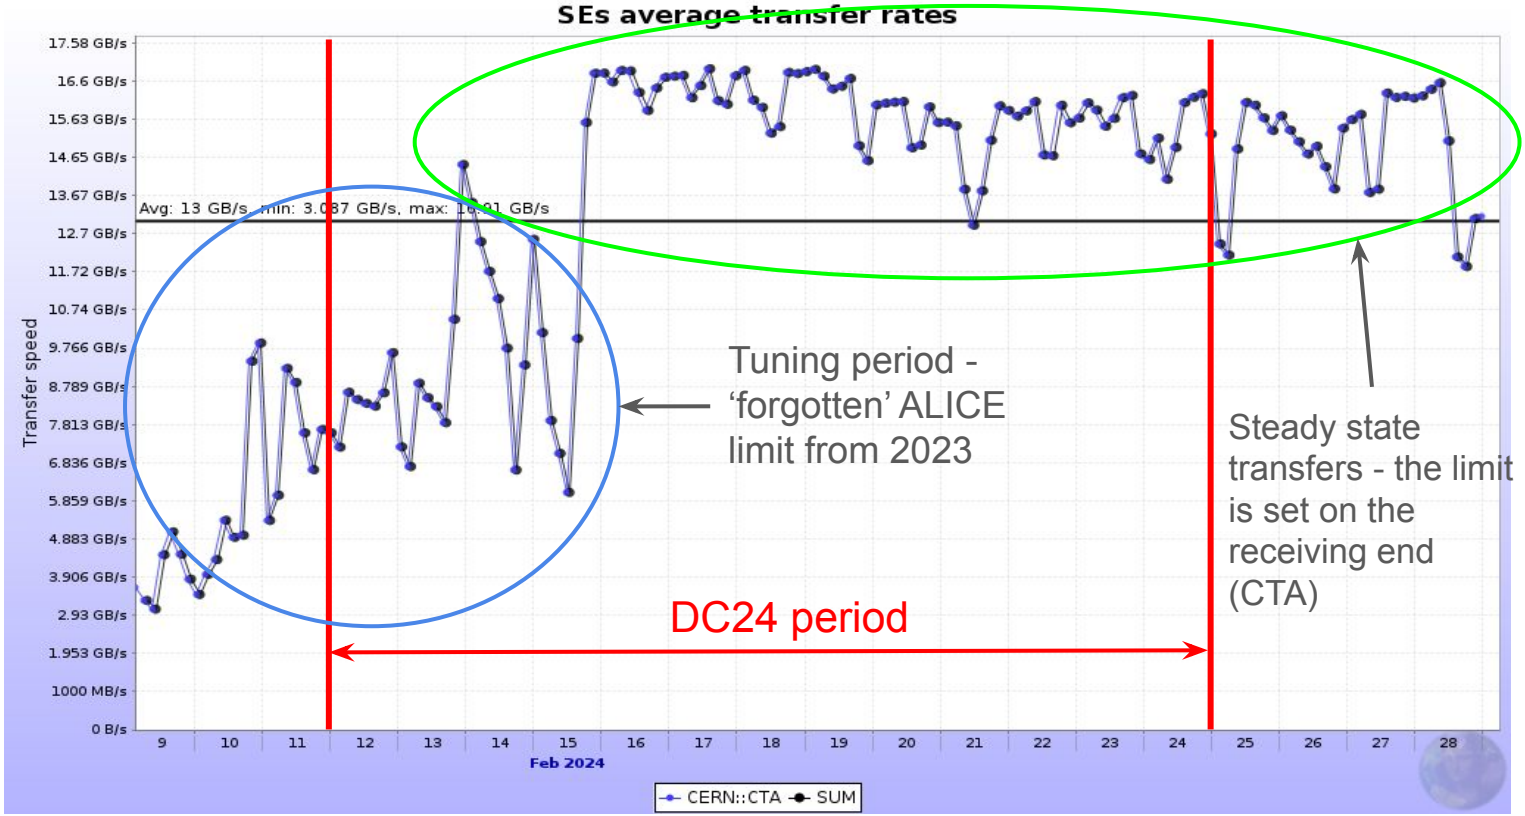

LHCb: Export from CERN to Tier-1s

EOS -> Disk link

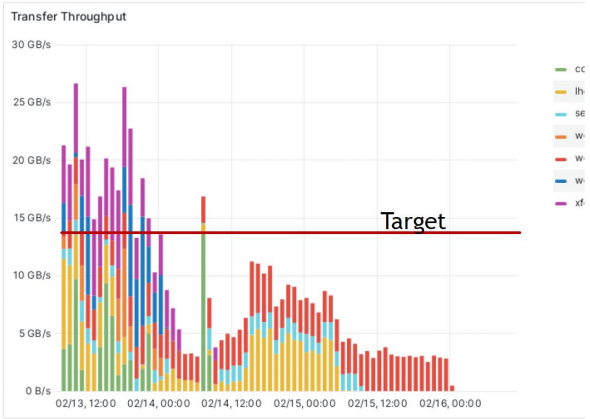

- Target throughput (14GiB/s) was achieved during the first day
- Lower throughput later
  - Some sites finished transferring their part during the first day so were no longer contributing to overall throughput
  - Submissions were slow and not optimal
  - Submission agent got stuck a few times, that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

Disk -> Tape link

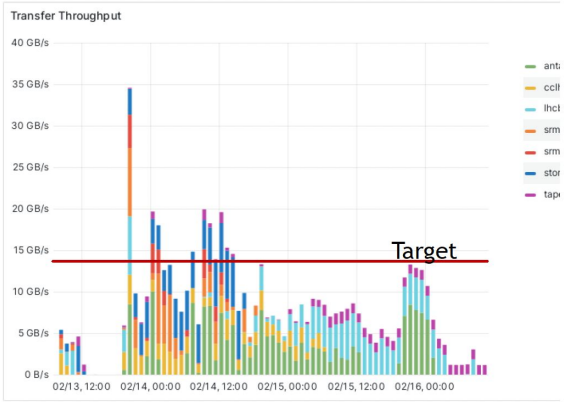

- Target threshold (14GiB/s) crossed several times
  - Max around 35GiB/s
  - Spikier throughput because of the nature of the link and submission agent problems

LHCb: Staging exercise

Staging

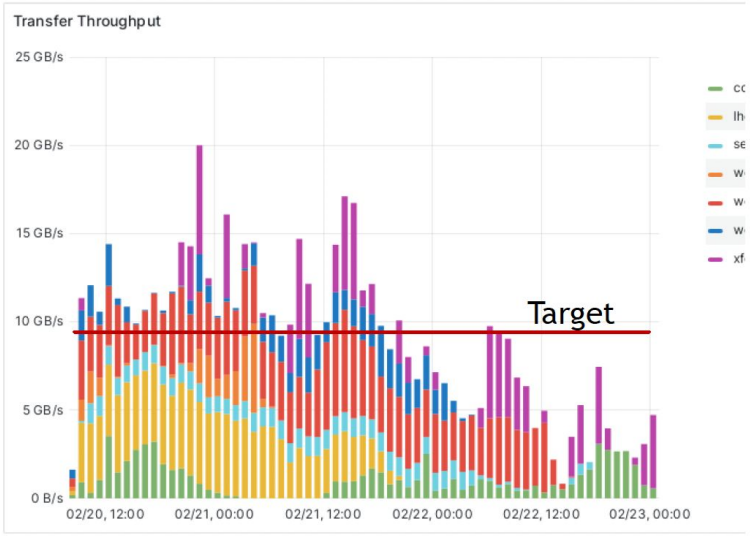

- Target throughput (9.58 GiB/s) was achieved during the first two days of the test
- Lower throughput later
  - Some sites finished transferring their part and were no longer contributing

8

CMS

- Daily exercise menu with increasing complexity
- T0 export, T1s to T1s and T1s to T2s, AAA

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

- 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

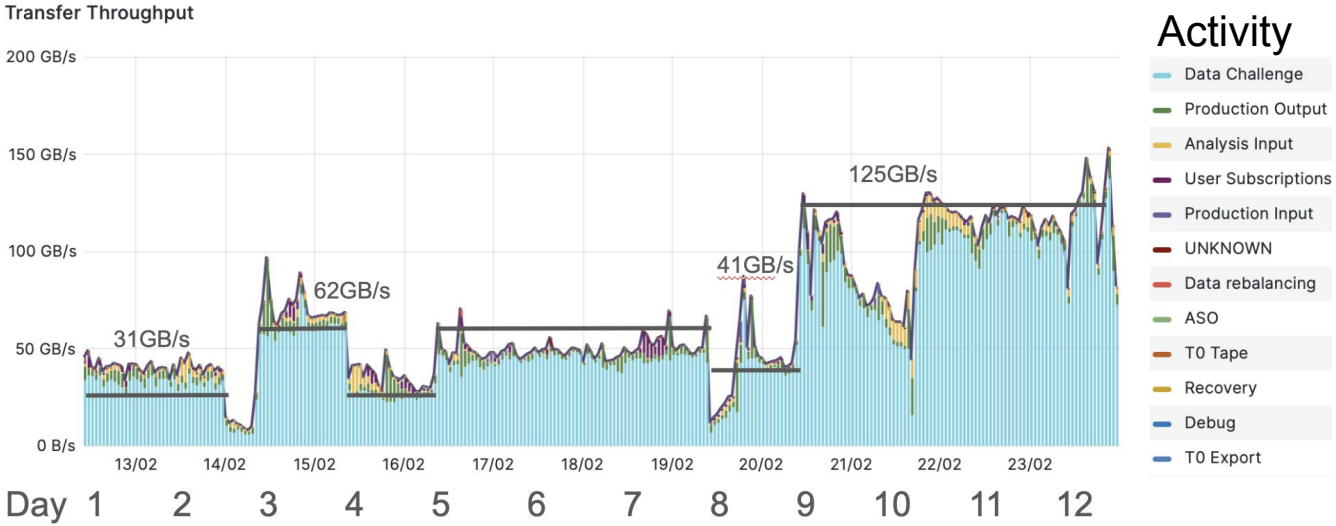

ATLAS

- 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
- 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)

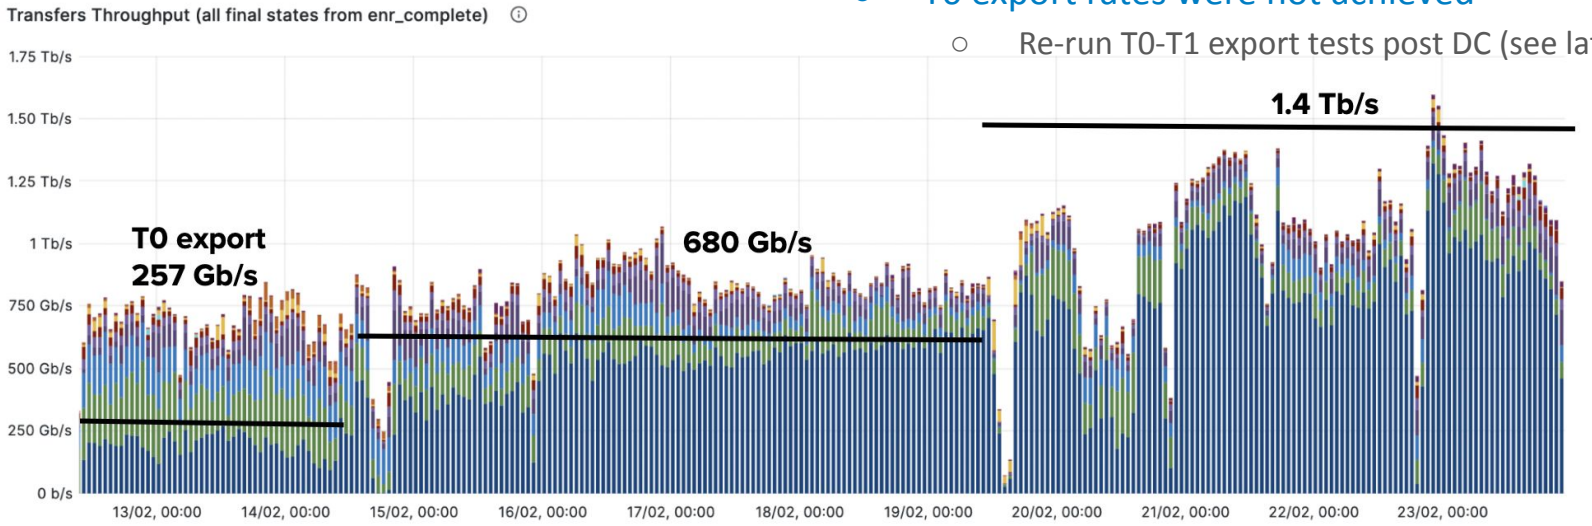

Observations and issues

ALICE - notable issues

- T0 rate - could not achieve 10GB/s in the first 3½ days
  - Reason - a limit of max active transfer threads was forgotten in the ALICE transfer system (self inflicted)
- 3h interruption of transfers to GridKA
  - Dead xrootd service on a disk buffer - repaired by restarting the service
- 24h interruption at CNAF
  - Too high rate observed on disk buffer, pinpointed to reads for md5sum calculation
  - Solved by adding SSD to the buffer
- All of the above are relatively trivial running issues, no structural or software problems identified

CMS - notable issues

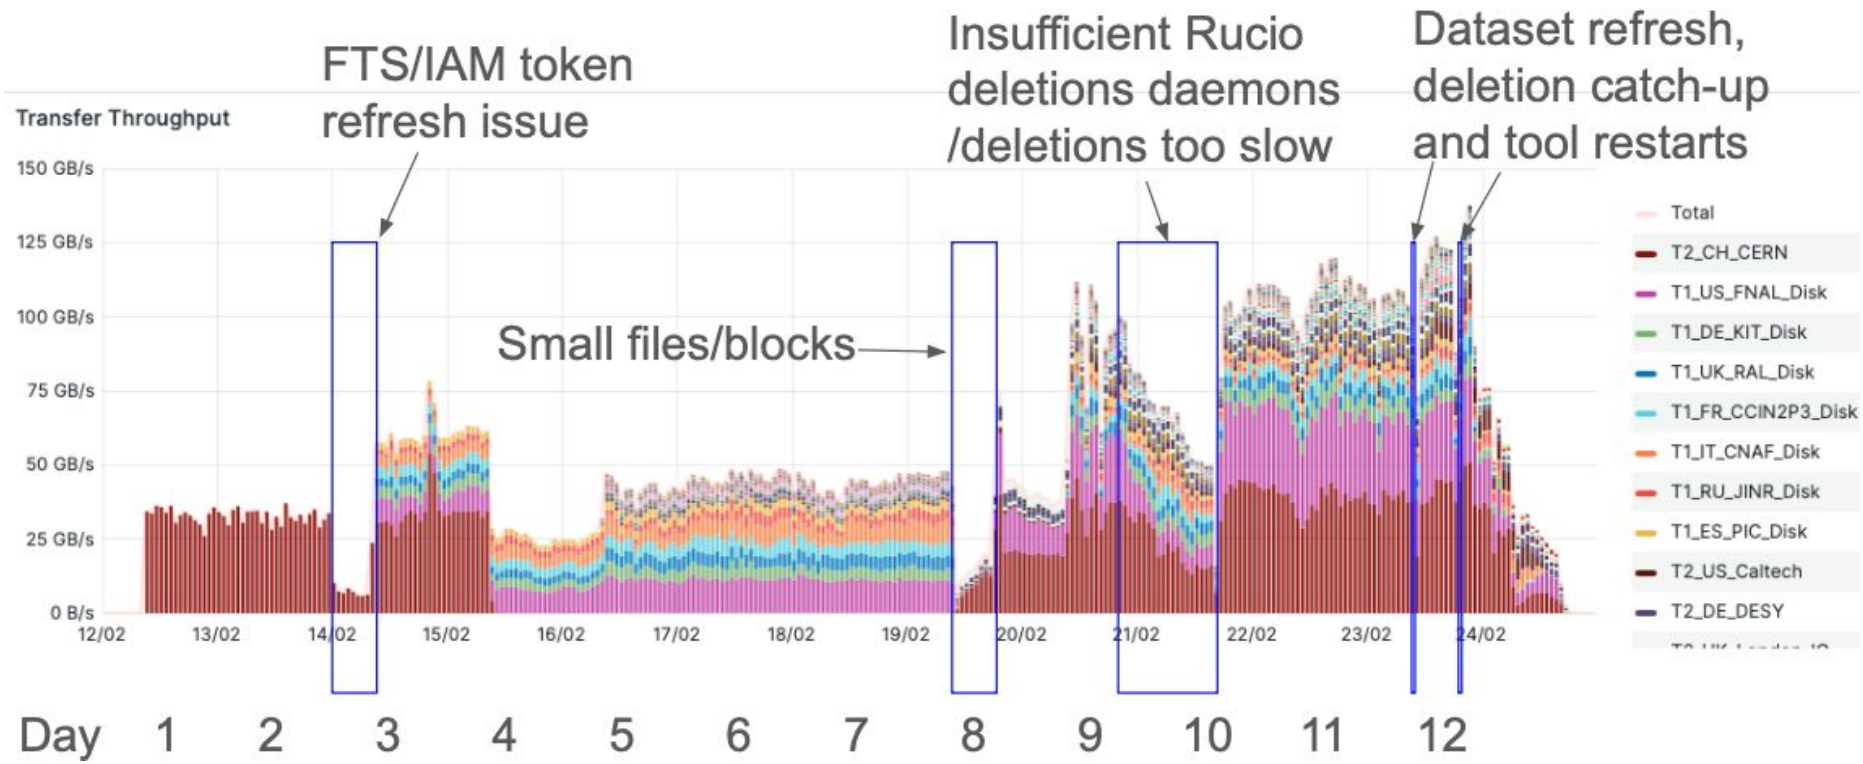

These marked periods were excluded as ‘known problems’ when analysing results

ATLAS - notable issues

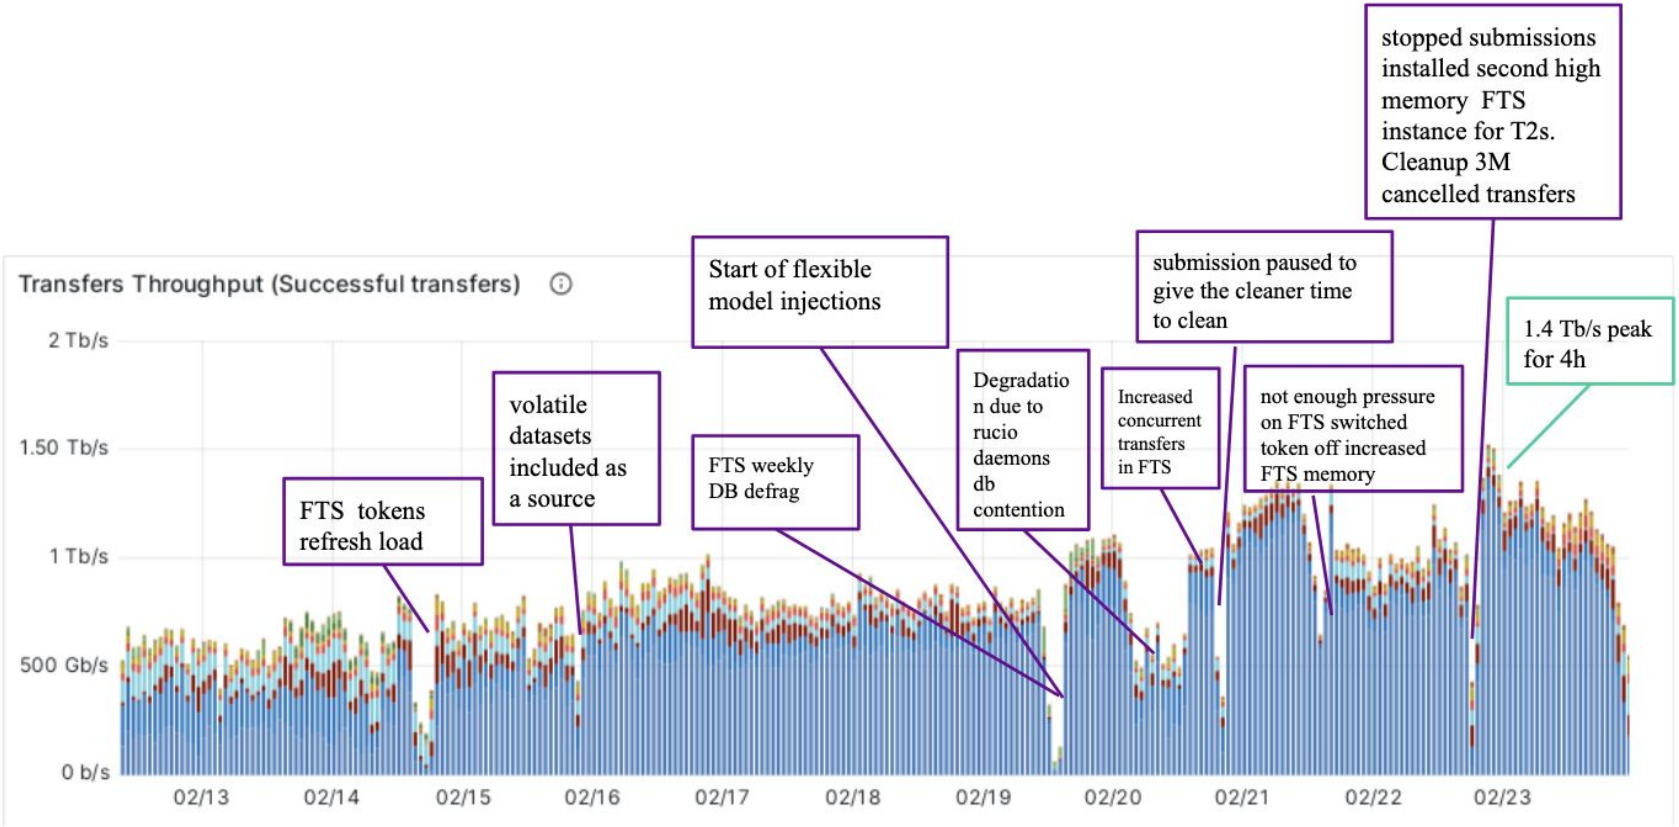

Token based Authentication

- Distributed infrastructure became ready just in time for DC24
  - FTS pre-release with token support
  - Rucio with base set of features for ATLAS and CMS
  - Deployment campaign to prepare storage elements
- About half of the transferred DC injected traffic via token authentication
  - Very high load on IAM by LHCb caused problems for LHCb token transfers
    - Used typically 3 tokens per transfer
  - ATLAS switched tokens off at the end of 2nd week
    - Refresh very expensive for FTS
  - CMS managed to maintain token auth throughout
    - But with a limited implementation
  - Valuable experiences gained with token usage at production scale

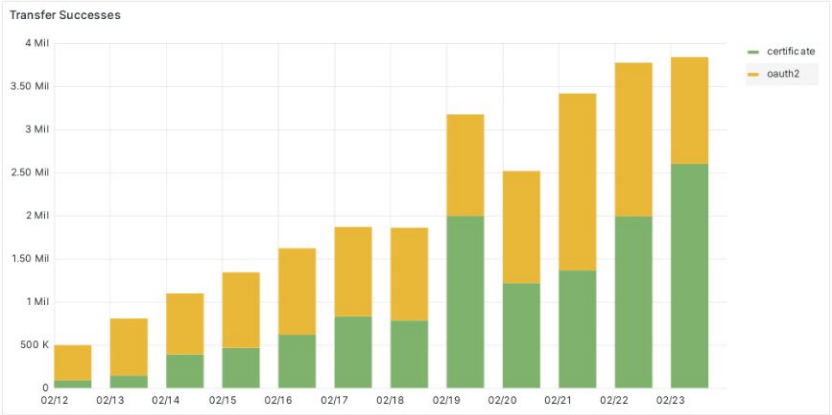

FTS operating at unprecedented scales

- Particularly FTS ATLAS instance survived thanks to permanent baby-sitting by FTS team
  - Database surgery in production
  - Increase of hardware resources
- Improved understanding of current FTS scaling
  - Optimizer cycle needed several hours
  - FTS has no concept of storage back pressure
  - FTS treats all links with the same activity with equal priority
- FTS team started to iterate developments items and related priorities with stakeholders of the community
- First official FTS release with token support this spring

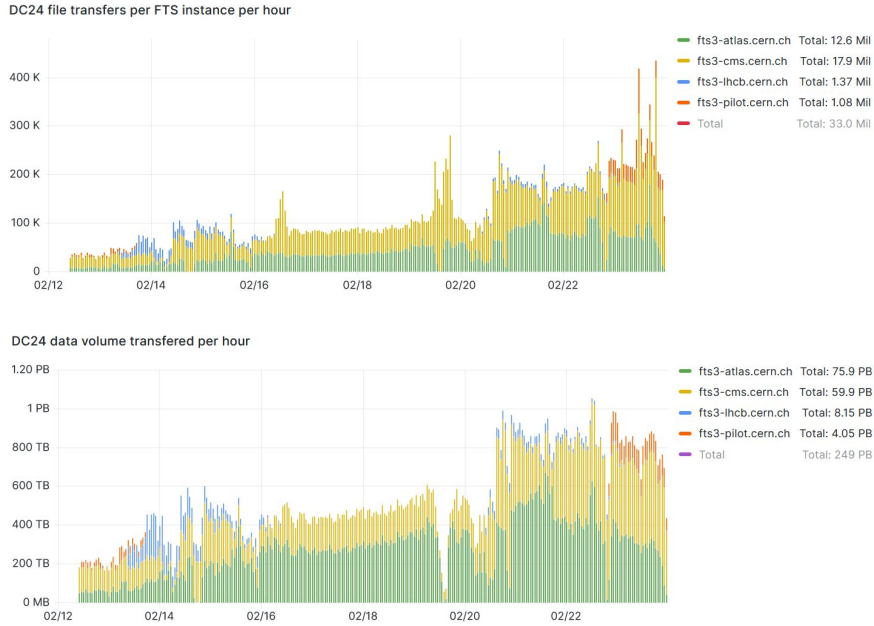

Plots show DC injected 'activity' only
Parallel ongoing production not included!

In-depth analysis

ATLAS and CMS - Site studies

- Each site had a daily target, and an observed rate.
- However, if the site did not hit the target, there could be several reasons for this:
  - FTS
  - Deletion issue
  - Other end of the transfer
  - Network - we see little evidence for this being the bottleneck
- Not every link has been studied - there are too many!

Next 2 slides - tables showing performance of Tier 1s per day compared to target

- Expected = Our target rate for that day, including all injected DC24 traffic
- Observed = Monitored average rate according to monit-grafana (certain periods excluded)
- Ratio = Observed / Expected (if value is 1 or above, site has met the target)
- Ratio colour scheme:
  - Green - ratio is >0.9; yellow - ratio is 0.7-0.9; orange - ratio is 0.5-0.7; red - ratio is <0.5

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)

		JINR		FNAL		IN2P3		RAL		PIC		KIT		CNAF	
Day	Scenario	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

- T0 export rates were not achieved - expected rates were 257 Gb/s
- Best rates achieved in the first 2 days with 192 Gb/s
  - T1 problems explain the low rates also in the the first 2 days
- Progressive degradation of the rates with more links being injected
  - T0-T1 treated the same as T2-T2 by FTS with no prioritization within an activity

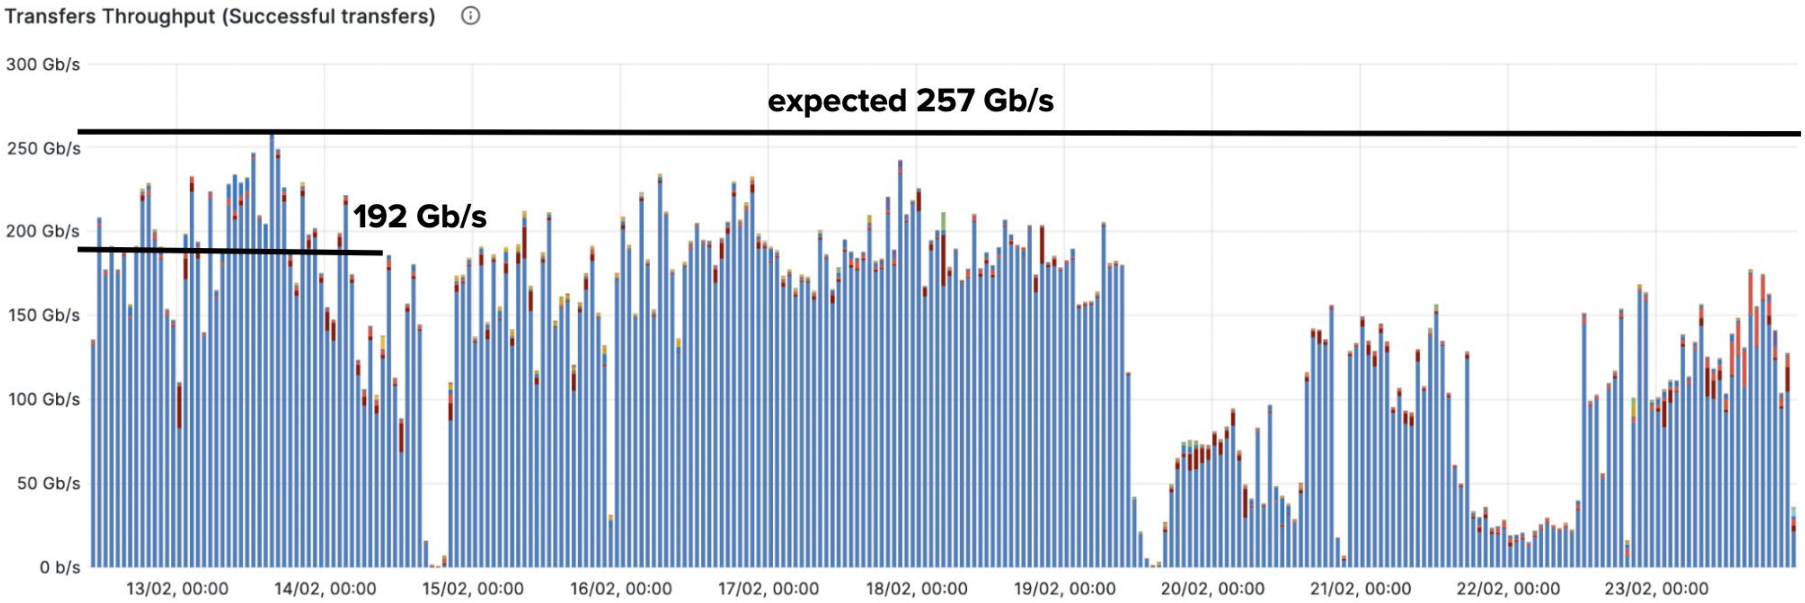

ATLAS - T0 export repeated tests

- 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 T1at the time	% of expected rates
BNL-ATLAS	60	31.5	53%	61.3	102%
FZK-LCG2	32	26.4	83%	42.2	132%
IN2P3-CC	38	43	113%	50.9	134%
INFN-T1	23	19.3	84%	33.5	146%
NDGF-T1	15	13.8	92%	28.2	188%
SARA-MATRIX	15	12.2	81%	274.1	1827%
pic	11	12.3	112%	18.1	165%
RAL-LCG2	38	15	39%	27.2	72%
TRIUMF-LCG2	25	23.9	96%	27.2	109%
T1 summary	257	197.4	77%	562.7	219%
T1 summary -SARA	242	185.2	77%	288.6	119%

CERN

- CMS: On the whole, CERN performed very well. No evidence to suggest otherwise
- ATLAS: CERN rates on the previous table were a reflection of the problems ATLAS had elsewhere in the infrastructure. **Not** a CERN problem.

CMS - scenario study

Day	T0 Export	T1 Export	T1s<->T1s	T2s<->T2s	AAA CERN to T2s	AAA FNAL to T2s	T2s->T1s	Special	Σ scenarios
1	1.11								1.11
2	1.05								1.05
3	1.11	0.99							1.05
4		0.83							0.83
5		0.79	1.09	0.59					0.79
6		0.86	1.10	0.56					0.81
7		0.83	1.11	0.59					0.81
8	1.29			0.92	1.18	0.98			1.08
9	0.61	0.54	0.77		0.96	0.74	0.73	0.90	0.70
10	0.83	0.62	0.67		1.05	0.67	0.70	0.83	0.75
11	0.71	0.64	0.80		0.92	0.60	0.85	0.84	0.73
12	0.82	0.70	0.92		0.86	0.67	0.89	0.22	0.71

Conclusions and outlook

Observations & remarks

- There are other bottlenecks than network bandwidth
  - Maintenance of DC injections was challenging
    - FTS instances got pushed to their limits, particular the ATLAS one
    - Keeping up with deletions is not trivial, systems not designed for best scaling here
    - Already ideas how to integrate data injector natively into Rucio
  - It needs time before a complex system reacts to parameter changes
    - The parameter space is huge
    - Not many attempts to re-adjust (very few per day)
  - A number of CMS sites asked for more (than planned) traffic to exercise their WAN connectivity

Final report

- To be delivered in time for the DESY Workshop **(NO EXTENSION!)**
  - Pre-structured document is here
  - Everyone involved in DC24 was asked to fill in their sections
  - CW & ML are editing the final document
    - A difficult task!

What to do before the next Challenge

- Aftermath of DC24
  - Derive 'lessons learned'
    - What went well, where were bottlenecks, organizational improvements …
  - Set priorities for ongoing developments
    - VO & community specific tools, e.g. Rucio, FTS, Storage middleware, Network equipment
- Planning of next DC
  - So far nothing is set except the global target of **50%** of expected HL-LHC throughput
  - Timeframe
    - Likely in 2026 or even later
    - Almost certainly in LS3, which makes scheduling much easier for LHC experiments
  - Experience shows that planning needs to start early (1 year before at least)
  - Recommended mini challenges throughout the period to test new functionalities and sites
  - Participating experiments
    - LHC experiments, hopefully again Belle-2 and DUNE
    - Interest expressed by JUNO, SKA, Neutrino experiments in Japan
  - Full implementation of tokens, mandated usage of tape, network tools, higher sophistication