WLCG Data Challenge 24

LHCONE/LHCOPN meeting, Catania

2024-04-10

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

Introduction: LHC and its High Luminosity upgrade

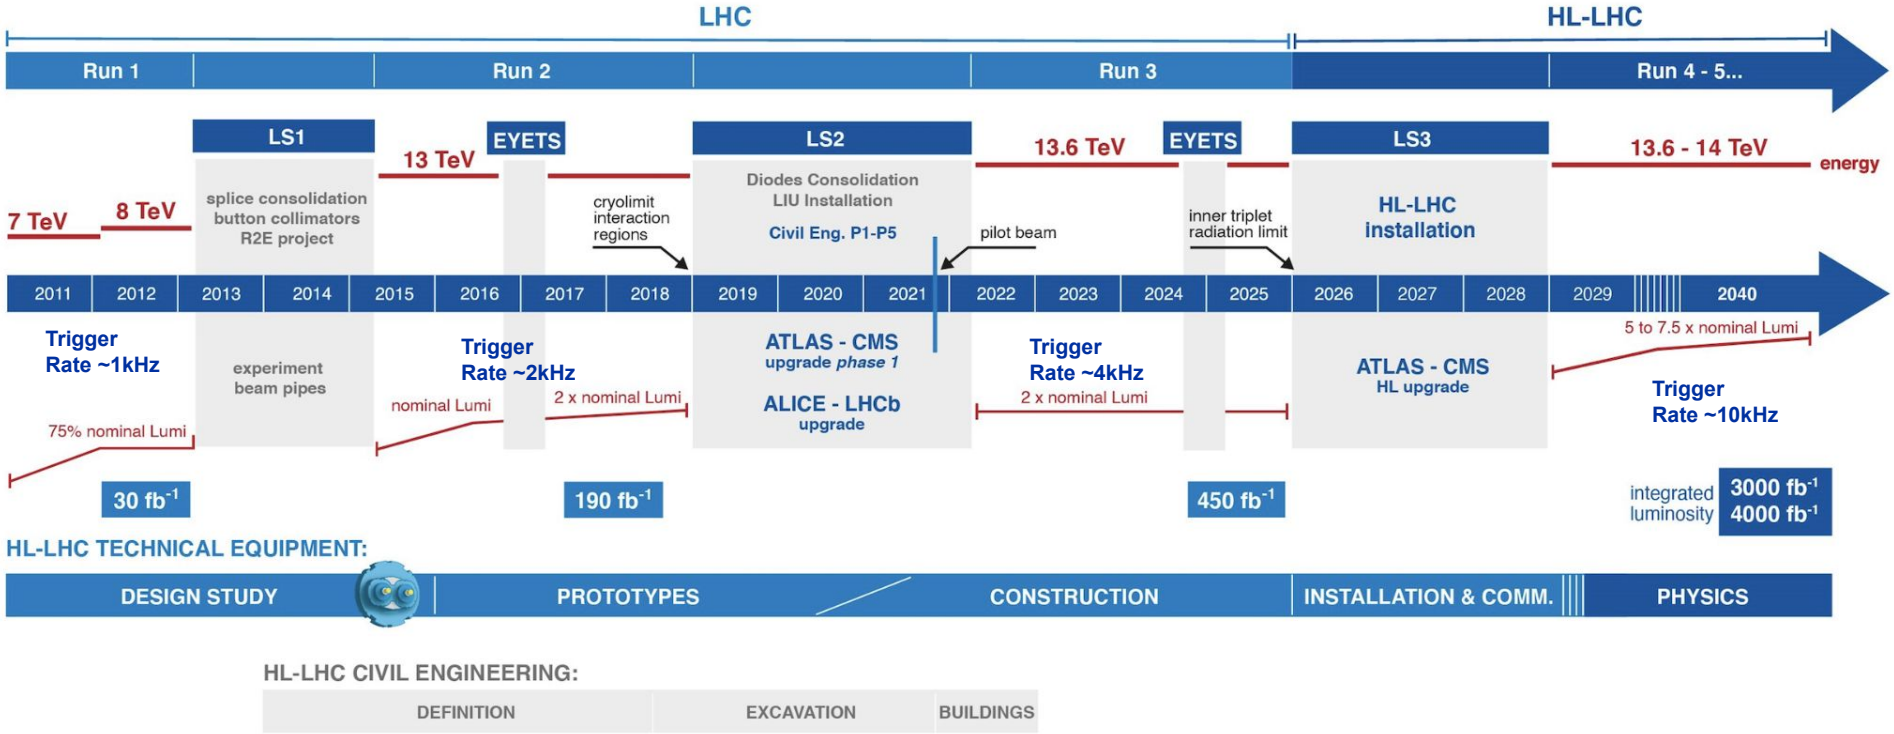

Figure adapted from:
Zerlauth, Markus & Bruning, Oliver. (2024). Status and prospects of the HL-LHC project. DOI; 615. 10.22323/1.449.0615.

Data Challenges for HL-LHC

WLCG has been 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
- Increase 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** needs and challenges
- For the DCs find agreements across the LHC experiments and beyond
  - Suited dates
  - Reasonable targets
  - Functionalities
- Help in orchestration

Dates and high level goals always approved by WLCG Management Board

LHC site structure

"Tier 0" == CERN

"Tier 1" => ~15 sites in total, receiving RAW data directly from CERN

"Tier 2" => ~100 sites, typically receiving data from Tier 1s

"Full mesh"

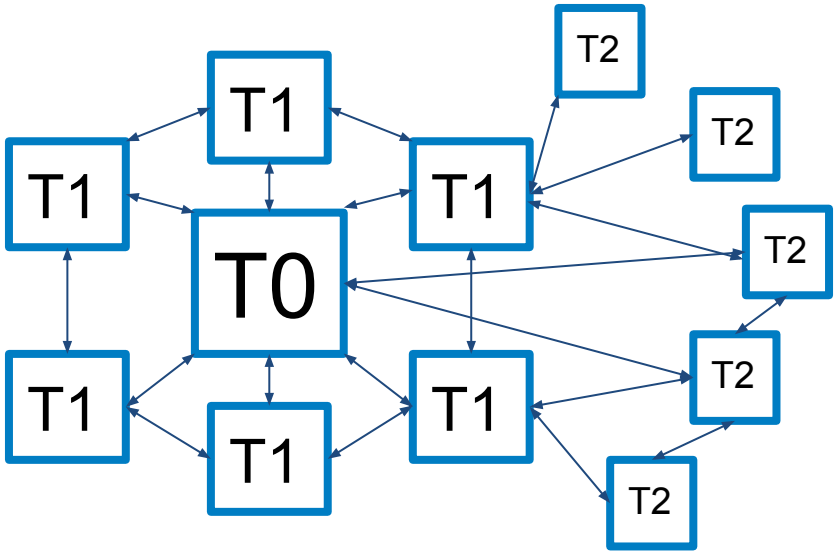

Recap of (initial) modelling & resulting rates for HL-LHC

ATLAS & CMS T0 to T1 per experiment

350PB RAW per year, taken and distributed during typical LHC uptime of 7M seconds

50GB/s or 400Gbps

Another 100Gbps estimated for prompt reconstruction data tiers (AOD, other derived output)

1Tbps for CMS and ATLAS summed

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

Minimal Model

Sum (ATLAS,ALICE,CMS,LHCb)*2(for bursts)*2(*overprovisioning*) = **4.8Tbps for the expected HL-LHC bandwidth needs**

Flexible Model

Assumes reading of data from above for reprocessing/reconstruction in 3 months (about 7M seconds)

Means doubling the Minimal Model: **9.6Tbps for the expected HL-LHC bandwidth needs**

However data flows primarily from the T1s to T2s and T1s!

Data Challenges target: **50% filling of expected** HL-LHC bandwidth needs

Planning of DC24

- Overall target: **25%** of HL-LHC throughput
  - Slightly lowered from originally 30% due to delayed start of HL-LHC
- Planning the DC24 program started well in advance
  - Agreement on dates
    - 2 weeks before beam operation in 2024
    - Full transfers from disk to disk
    - DC and production traffic counts
  - Experiments had room to define their goals
    - ALICE and LHCb involved tape
    - ATLAS and CMS did not
  - Preparation of monitoring
  - Regular preparation started one year before
    - Monthly checkpoint meetings
    - Dedicated workshop in Nov 2023

ATLAS - Planning

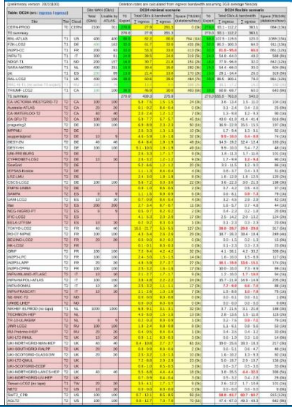

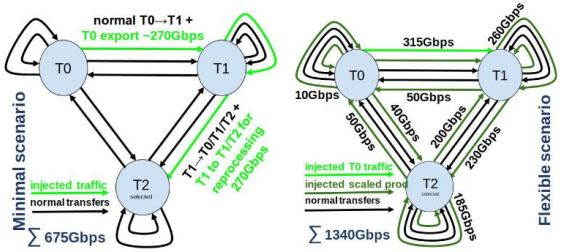

- Rather detailed planning exists
- Rates are mainly scaled values from measured Run-3 values
- Sites are already informed about expected rates

CMS - Main Scenarios

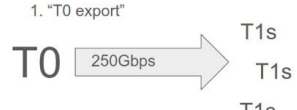

- Rather well modelled
- Numbers derived from DAQ TDR and LHC uptime assumptions

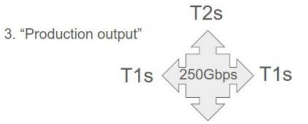

- MC & derived data scenario
  - HL-LHC approach not fully developed
- Data rates still somewhat uncertain

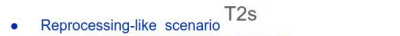

- Reprocessing-like scenario
  - HL-LHC approach not fully developed
- Data rates still somewhat uncertain

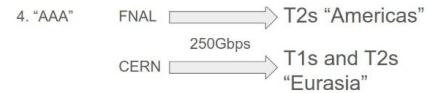

- Unscheduled remote reads via Xrootd
  - Main traffic presently MC premixing served from CERN and FNAL
  - HL-LHC approach not fully developed
- Data rates still somewhat uncertain

2 example slides from DC24 workshop

Christoph Wissing (DESY) & Mario Lassnig (CERN) 11

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-DC24 tests…

- Significant testing done well in advance
- No time to show results!
- Focussed on CERN to Tier 1s
  - Testing each T1 in turn, then all together
  - Also tested our infrastructure

DC21 - 10% of HL-LHC Throughput

However, we managed to reach 100% of the (minimal) DC21 target!

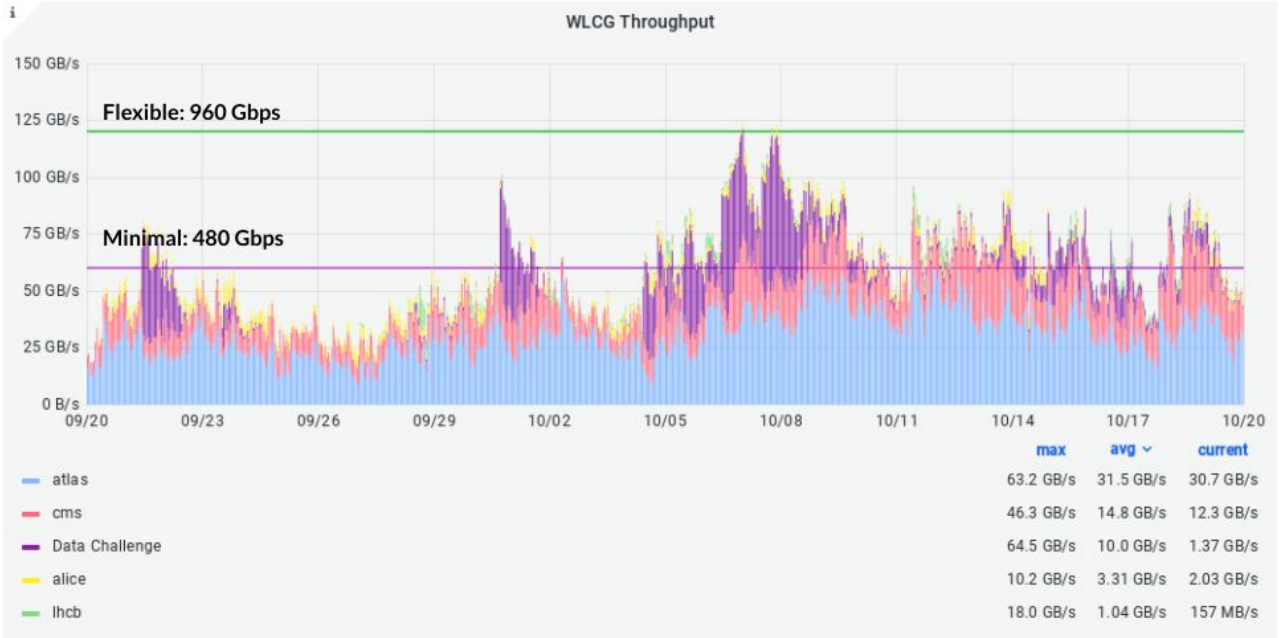

Network Data Challenges 2021 wrap-up and recommendations

https://zenodo.org/records/5767913

DC24 in one plot

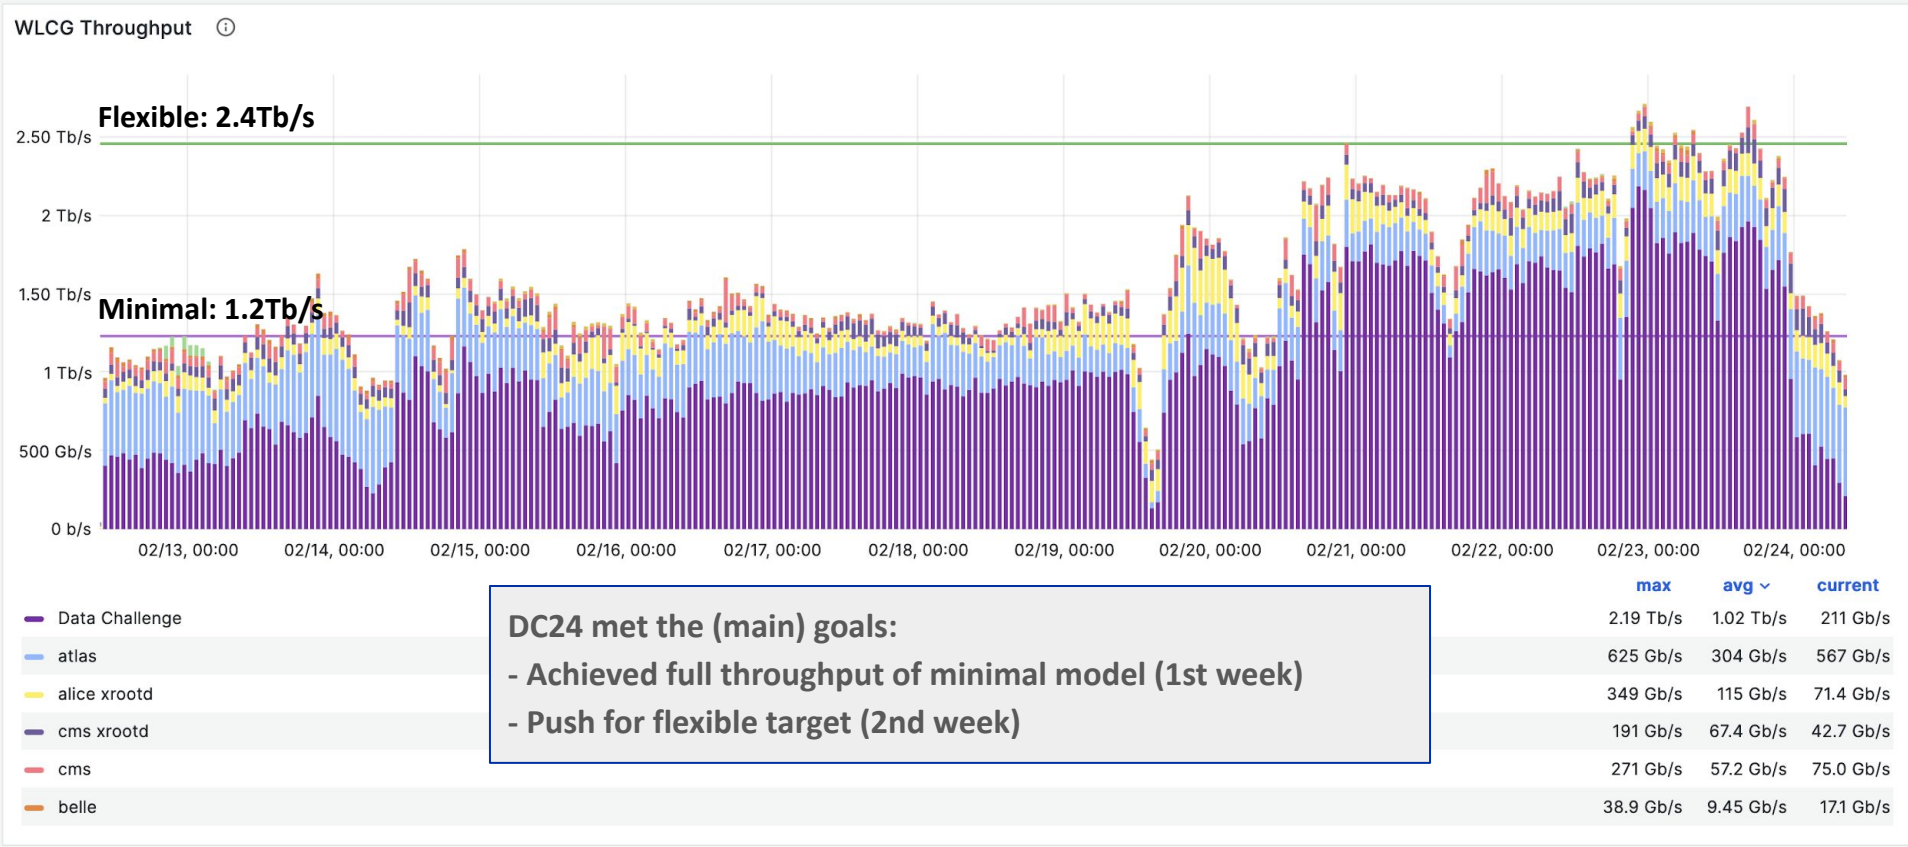

DC24 met the (main) goals:

- Achieved full throughput of minimal model (1st week)

- Push for flexible target (2nd week)

ATLAS and CMS workflow

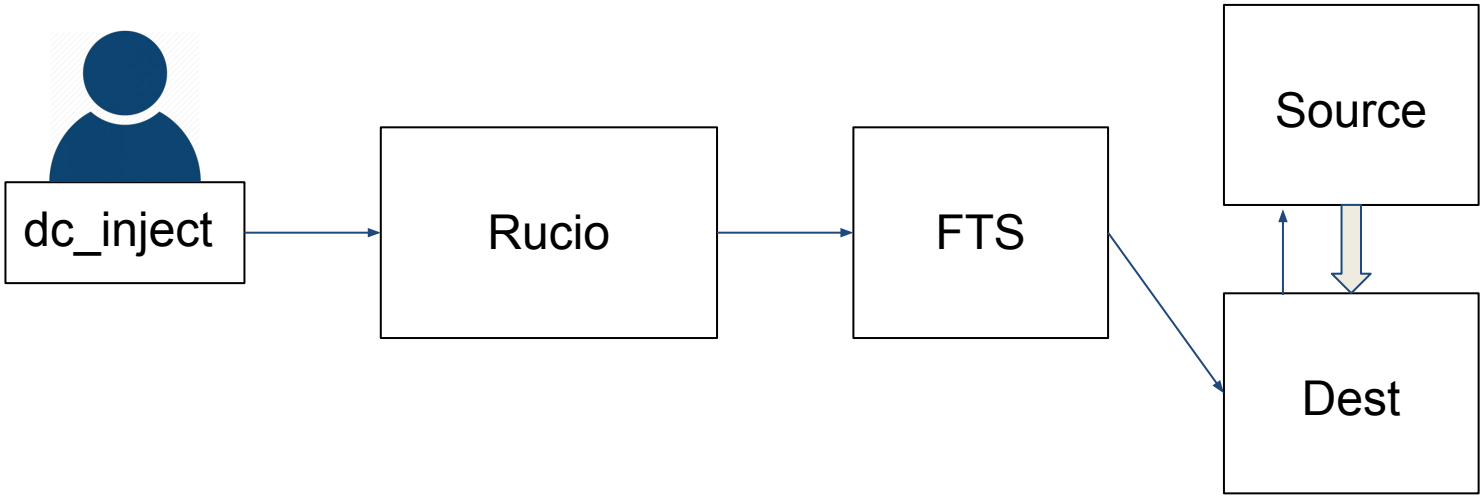

Notable issues

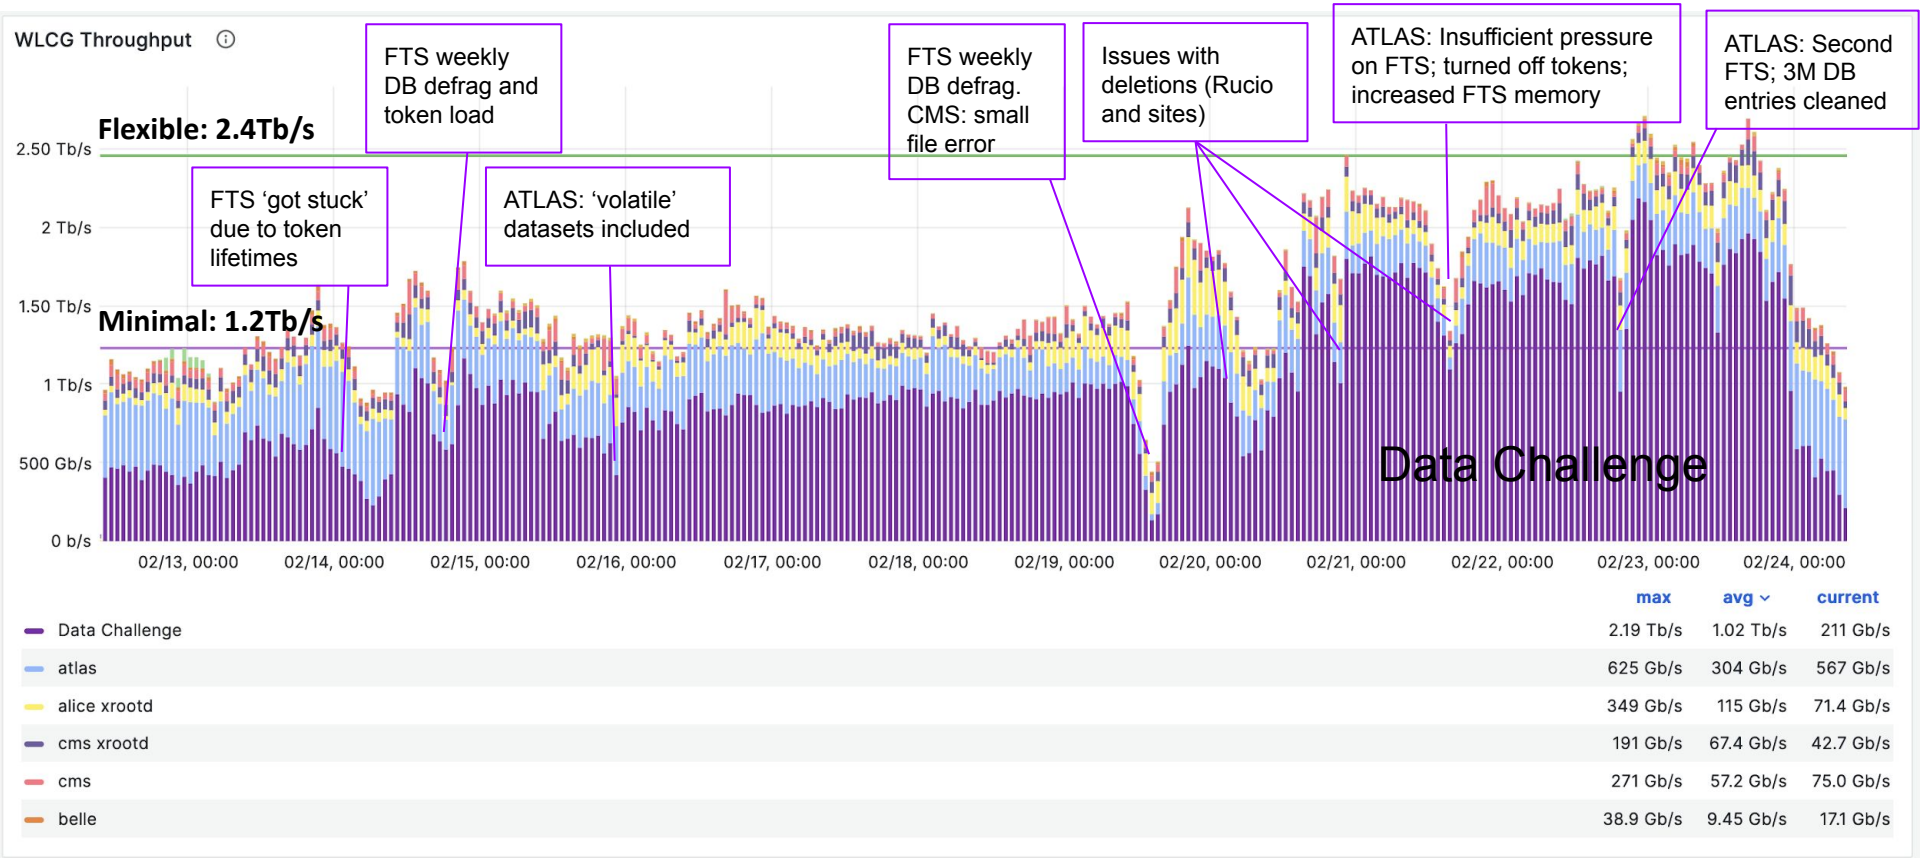

ATLAS

- Generally considered a success, though with some “homemade” issues
  - Injections on >1200 links every 15m
    - ~2000 links with production
  - Helped highlight problems that wouldn’t have been seen otherwise
- None of the bottlenecks were due to the network specifically
- FTS and Rucio central services affected the transfers more
  - T0-T1 had to be re-run after DC24
- Some sites struggled mostly due to storage limitations
  - 17 problems were reported on GGUS

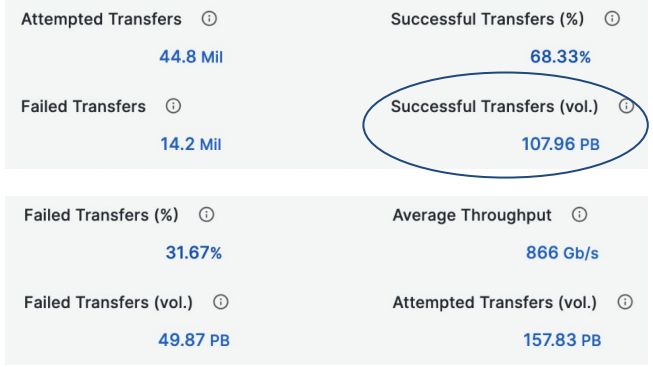

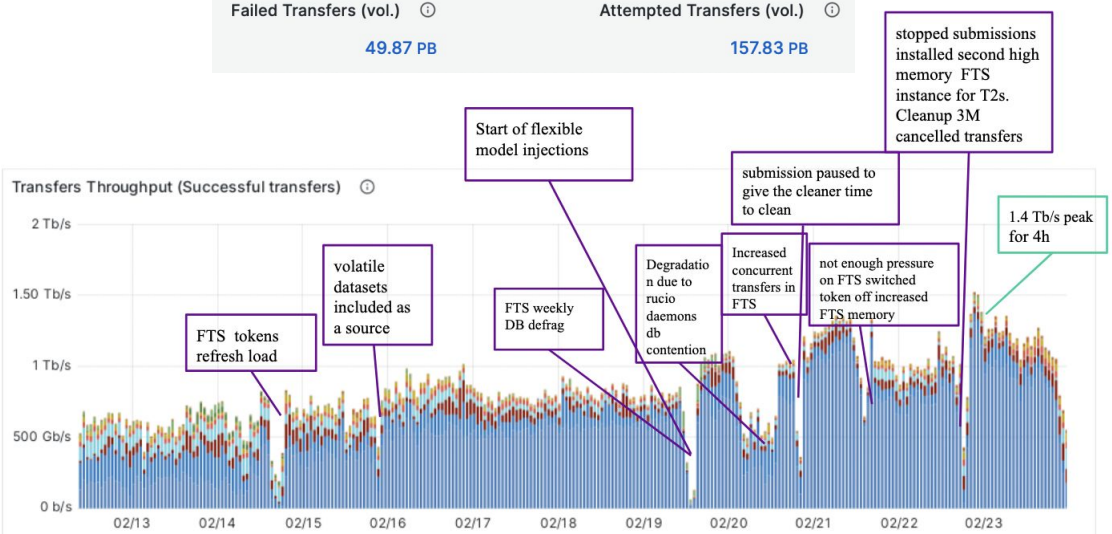

CMS

- Daily exercise menu with increasing complexity
- T0 export, T1s to T1s and T1s to T2s, AAA
- Overall target of ~125GB/s could be met
  - A few hundred links in total (Prod + DC)
- Some limitation in 'deletion performance'
  - Tuning of Rucio deletion pods

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

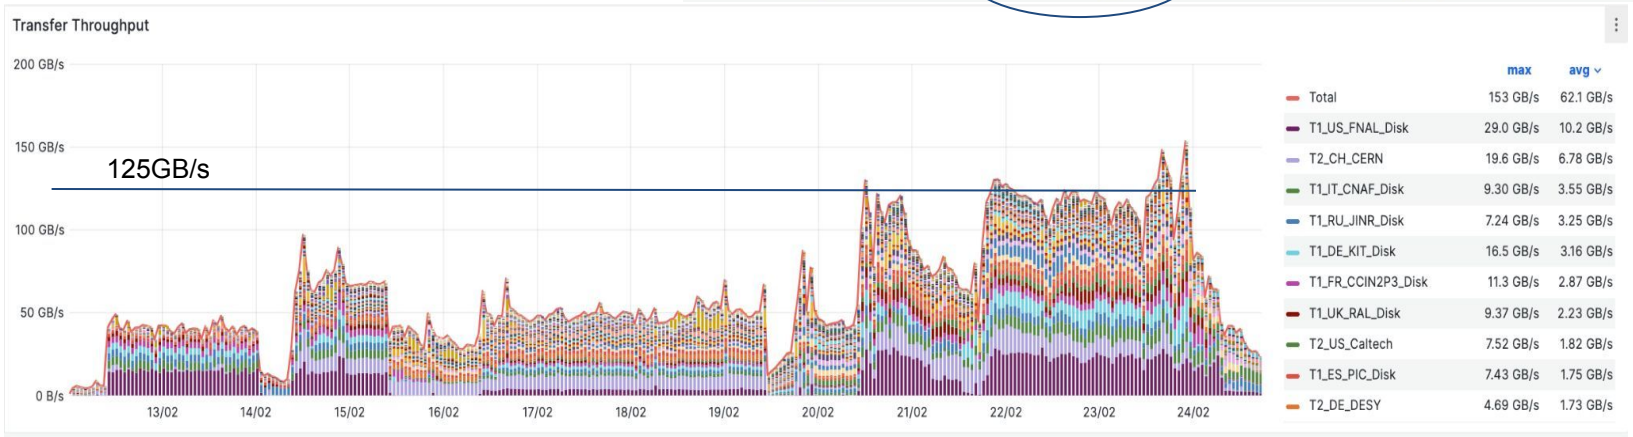

LHCb: Export from CERN to Tier-1s

EOS -> Disk link

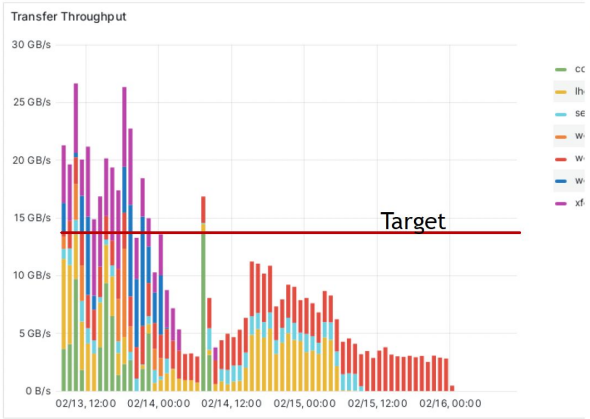

- Target throughput (14GiB/s) was achieved during the first day
- Lower throughput later
  - Some sites finished transferring their part during the first day so were no longer contributing to overall throughput
  - Submissions were slow and not optimal
  - Submission agent got stuck a few times, that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

Disk -> Tape link

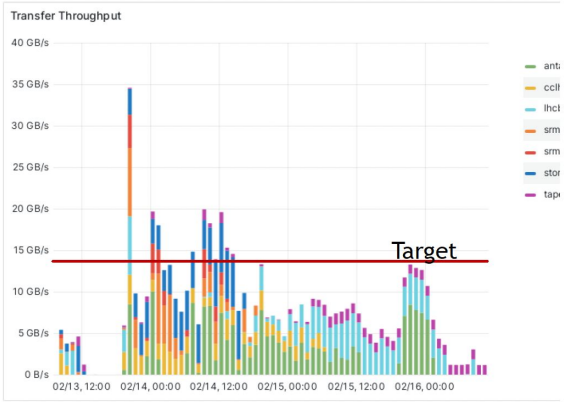

- Target threshold (14GiB/s) crossed several times
  - Max around 35GiB/s
  - Spikier throughput because of the nature of the link and submission agent problems

LHCb Credit: Alex Rogovskiy

LHCb: Staging exercise

Staging

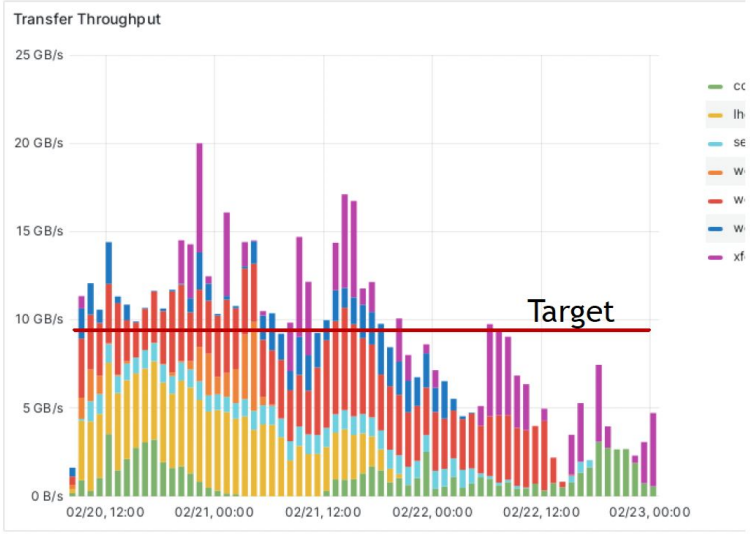

- Target throughput (9.58 GiB/s) was achieved during the first two days of the test
- Lower throughput later
  - Some sites finished transferring their part and were no longer contributing

8

ALICE: Ongoing data export from 2023 HI run

Time evolution T1s

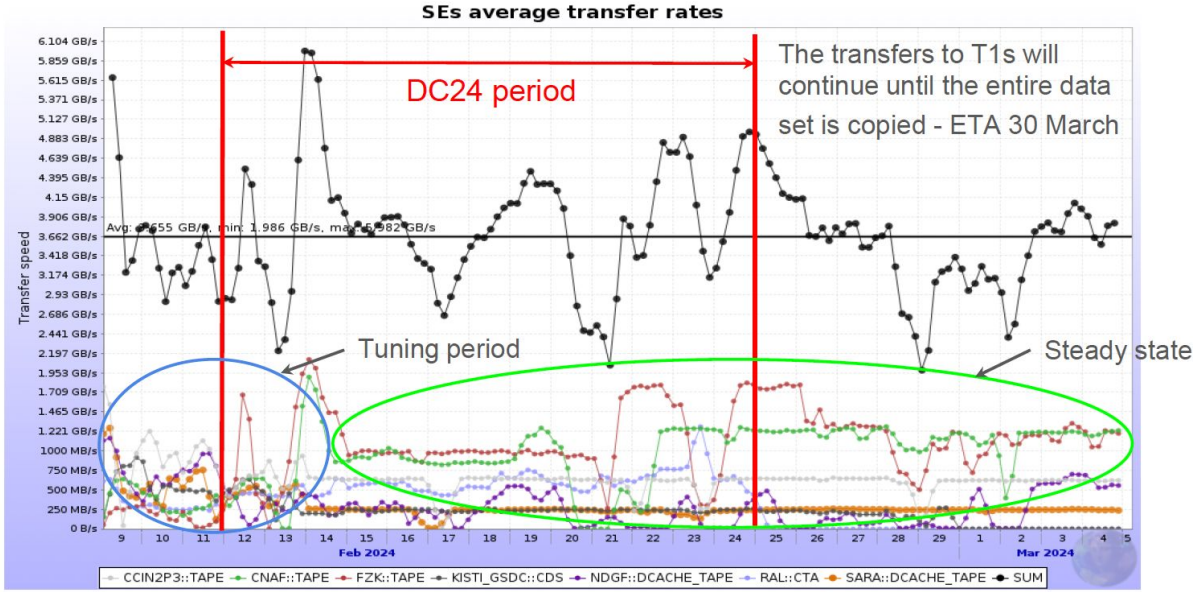

Centre	Target rate GB/s	Average achieved GB/s
CNAF	0.8	0.98 (+20%)
IN2P3	0.4	0.6 (+40%)
KISTI	0.2	0.25 (+22%)
GridKA	0.6	1.12 (+90%)
NDGF	0.3	0.35 (+15%)
NL-T1	0.1	0.25 (+150%)
RAL	0.1	0.58 (+500%)
CERN	*10*	*14.2 (+40%)*

ALICE Credit: Latchezar Betev and Maarten Litmaath

Belle II & DUNE

- Participation of non-LHC experiments in WLCG challenge for the first time
  - Belle II and DUNE fully included in planning process
  - Rates order of magnitude slower compared to LHC, flows often in opposite direction
- Belle II
  - Focus on traffic from KEK to RAW data centers and between RAW data centers
  - Targets were met
  - No obvious interference with LHC experiments
- DUNE
  - Focus on RAW data archiving & processing
  - Identified and improved some bottlenecks
  - Participation considered extremely useful

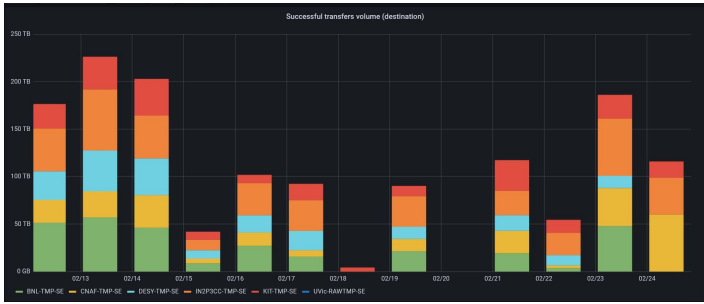

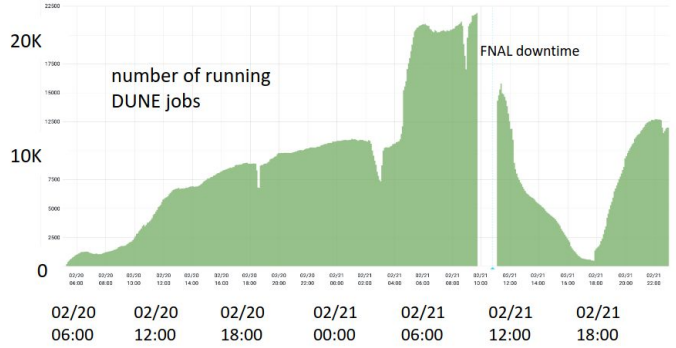

Credits: Silvio Pardi and Doug Benjamin

Token based Authentication

- Distributed infrastructure just became ready 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 for LHCb
    - Used 1 token per transfer
  - ATLAS switched tokens off at the end of 2nd week
    - Refresh very expensive for FTS
  - Valuable experiences gained with token usage at production scale
- Follow-up discussions in relevant forums to come

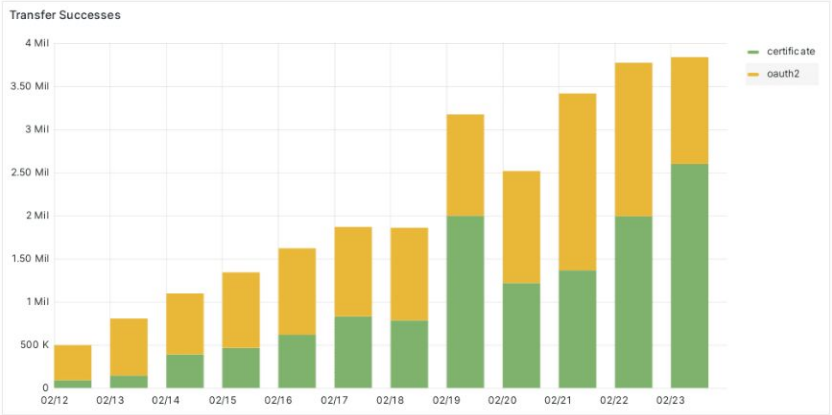

FTS operating at unprecedented scales

- 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 development items and related priorities with stakeholders of the community
- First official FTS release with token support this spring

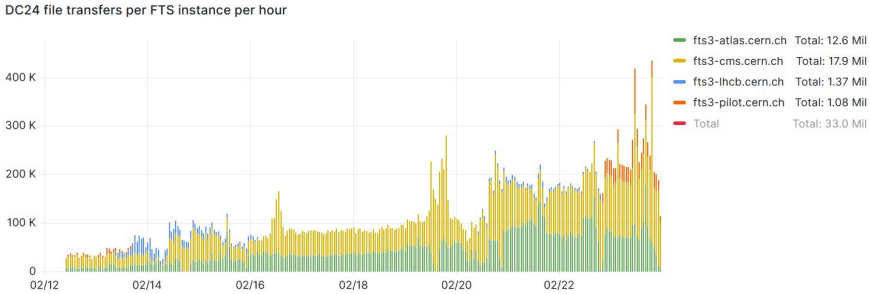

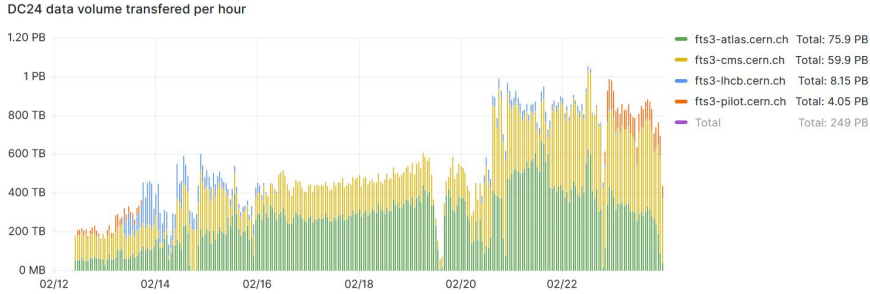

Plots show DC injected 'activity' only
Parallel ongoing production not included!

More pain - which should be a gain

- The main reason for not being able to sustain the DC24 target for 48 hours was…
  - FTS manages concurrent data transfers per link and **NOT** throughput
  - FTS treats all links with the same activity with equal priority
- FTS saturated all of its configured destination endpoints
- FTS CANNOT reach maximum throughput for the following configuration:

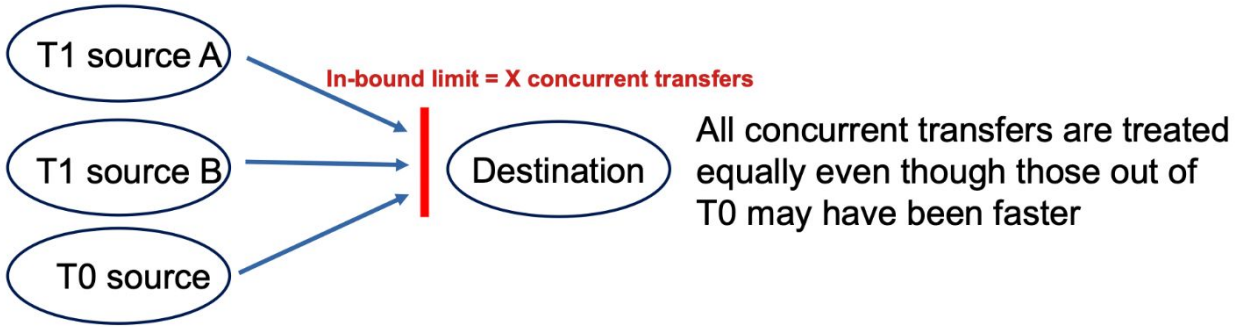

FTS Credit: Steven Murray

Site study

- 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
  - Etc.
  - (Network - we see little evidence for this being the bottleneck)
- Not every link has been studied - there are too many!

Tier 1 study (CMS)

- 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

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	1.42	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	1.46	1.10	N/A	0.50	N/A	1.17	N/A	0.94	N/A	1.17	N/A
3	T0-Export, T1-Export	1.31	0.62	1.08	1.31	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	N/A	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	1.18	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	1.14	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	1.19	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	1.30	1.39	N/A	1.31	N/A	1.31	N/A	1.70	N/A	1.32	N/A
9	T0-Export, T1-Export, Prod-out, AAA	0.38	0.34	0.87	0.38	0.57	0.57	0.95	1.02	1.25	0.86	0.86	0.56	0.65	0.25
10	T0-Export, T1-Export, Prod-out, AAA	0.70	0.34	0.98	0.70	0.58	0.65	0.56	0.99	0.70	0.66	1.03	0.98	0.63	0.28
11	T0-Export, T1-Export, Prod-out, AAA	0.63	0.33	0.91	0.63	0.43	0.76	0.77	1.05	1.09	0.84	0.91	1.09	0.69	0.24
12	T0-Export, T1-Export, Prod-out, AAA	0.40	0.54	0.92	0.40	0.89	1.00	0.85	1.15	1.21	0.87	1.13	0.89	0.78	0.29

Tier 0,1,2 study (ATLAS)

- 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 repeat

- 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%

UK study

- Would the 300G link between Geant and Janet be saturated?
- Due to the OPN being down to the UK, traffic fell back to LHCONE and Janet
- Saturation observed due to config issue
- General IP traffic moved to the backup link on 14th Feb

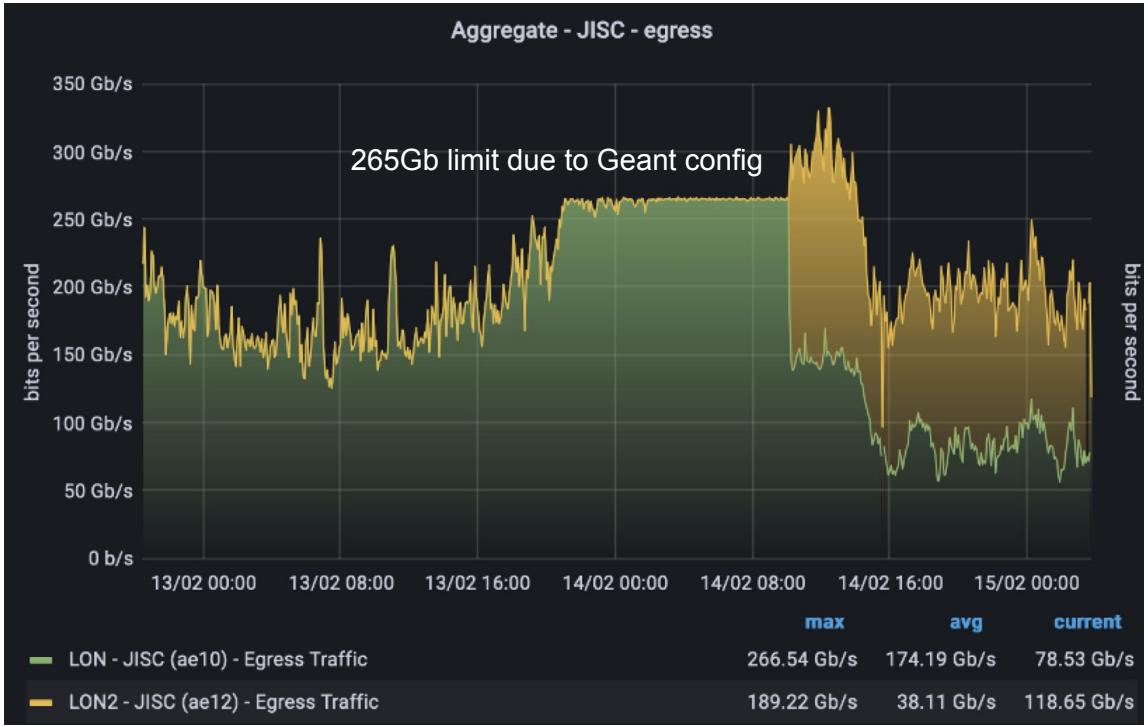

UK Credit: Duncan Rand

Monitoring and validation

- The data challenges have their own monitoring (FTS)
- New ‘network’ monitoring shows us the total in/out for each major site (51 sites?)
  - It is useful not having to ask individual sites for internal network plots!
  - However, ideally this should be validated by sites
  - VOs will struggle to validate as typically they only have FTS monitoring

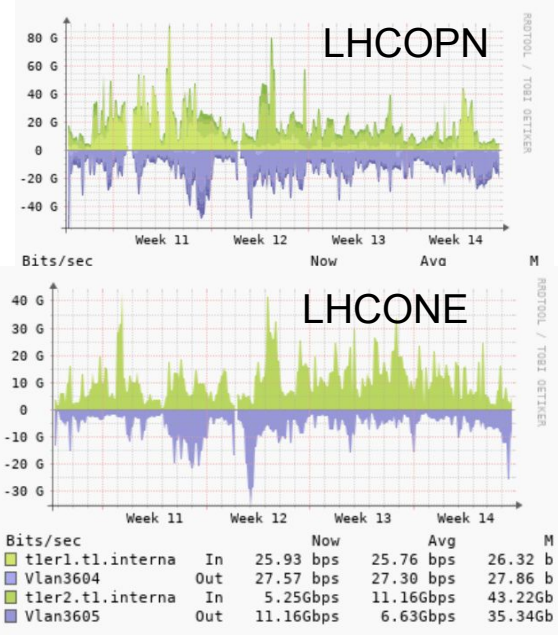

RAL internal (libreNMS) and new grafana network monitoring for last 4 weeks

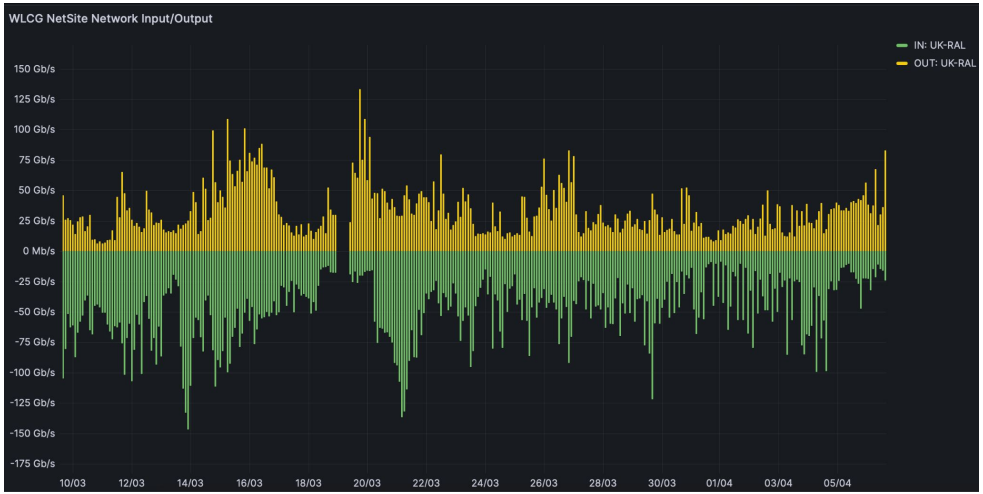

Monitoring and validation

- The data challenges have their own monitoring (FTS)
- New ‘network’ monitoring shows us the total in/out for each major site (51 sites?)
  - It is useful not having to ask individual sites for internal network plots!
  - However, ideally this should be validated by sites
  - VOs will struggle to validate as typically they only have FTS monitoring

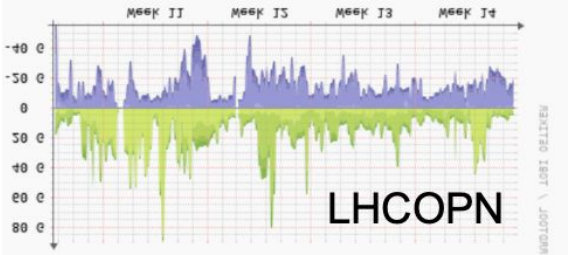

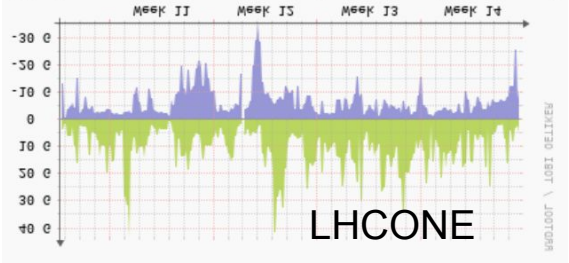

RAL internal (libreNMS) and new grafana network monitoring for last 4 weeks

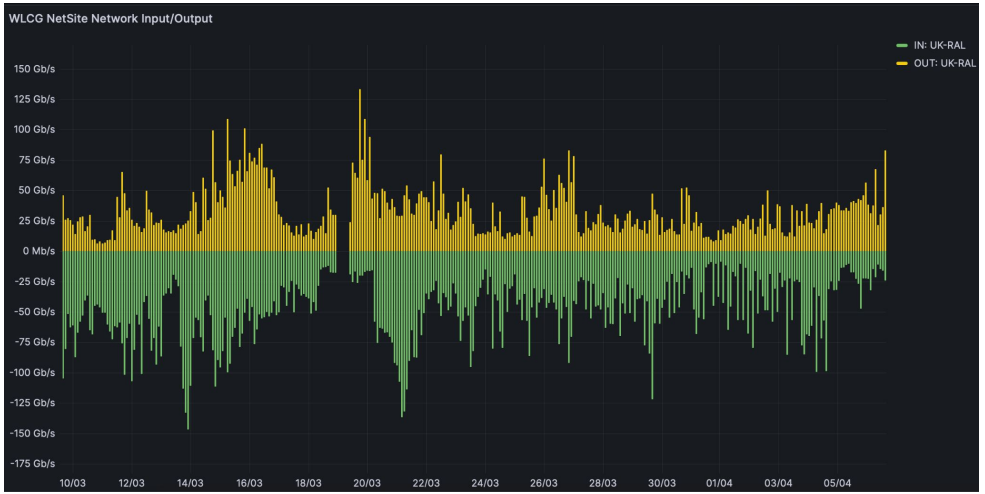

Monitoring and validation

- RAL validation?

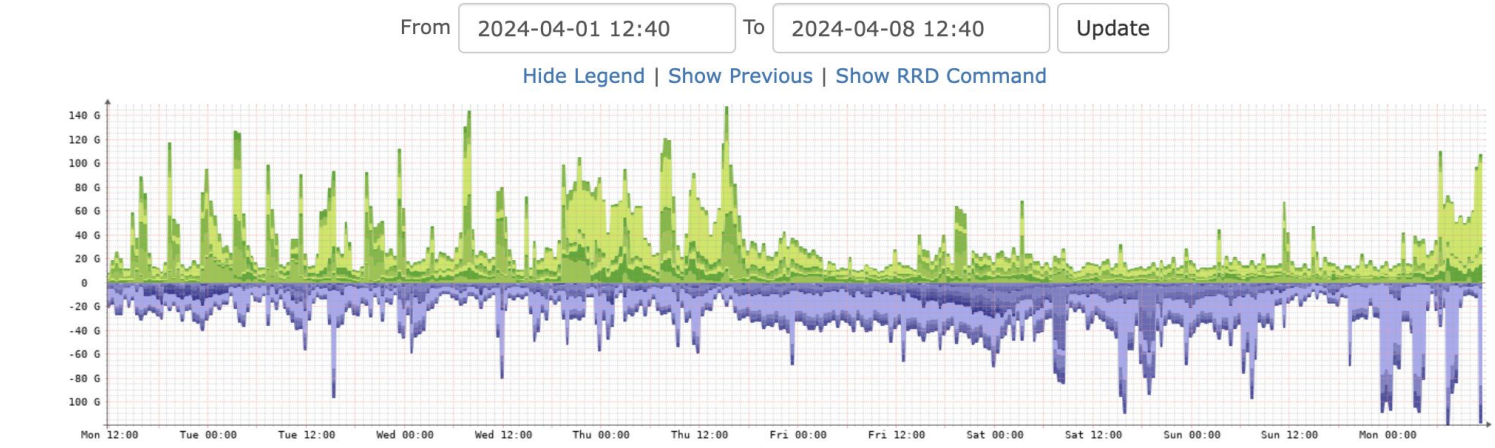

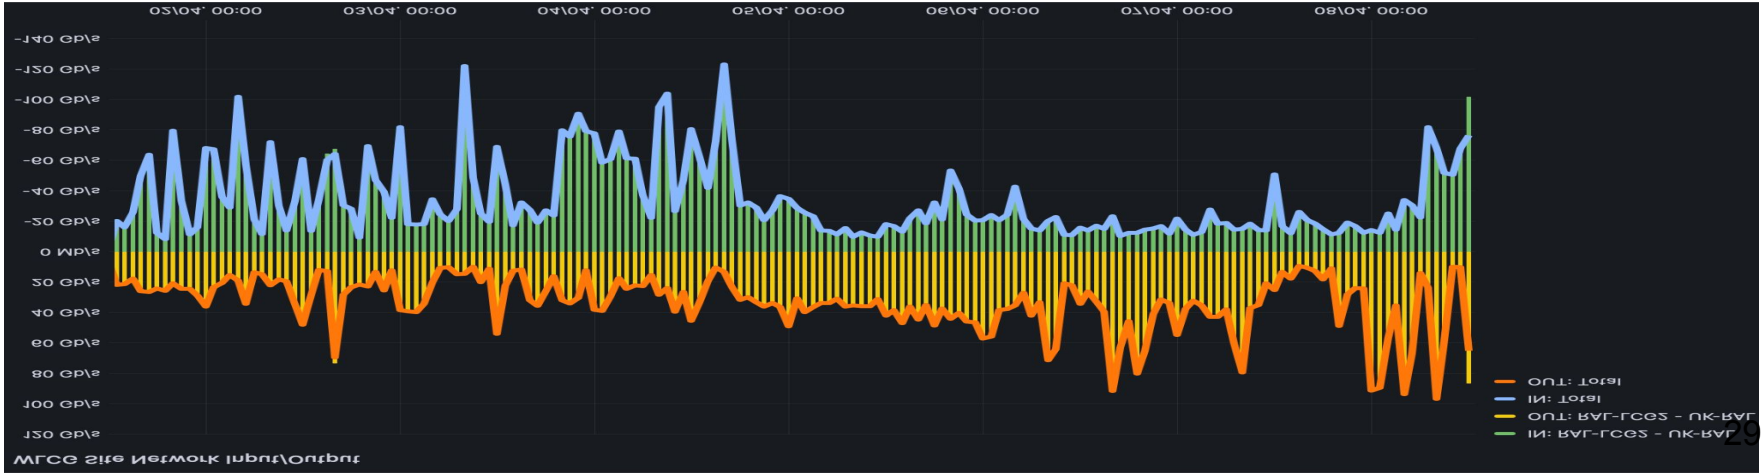

LHCOPN Routing from CERN to FNAL (CMS' biggest sites)

CERN to FNAL pre-tests. Nov 2023

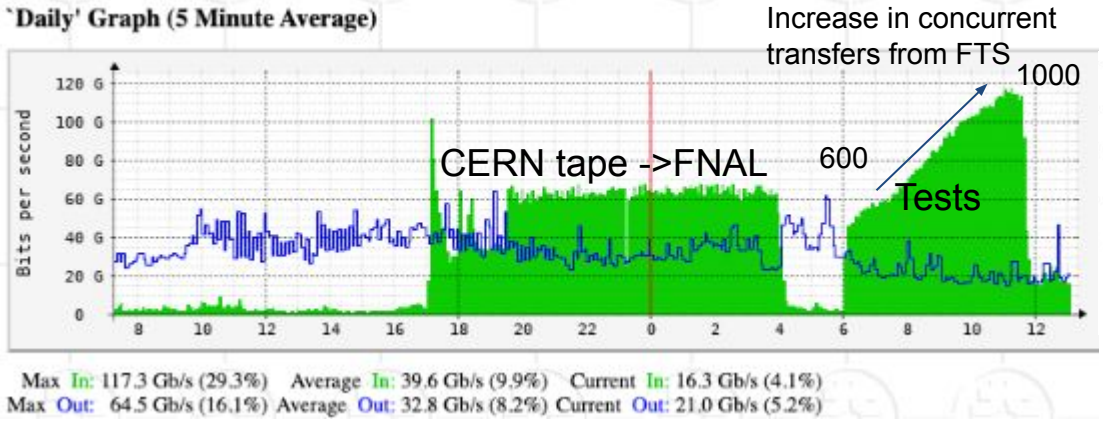

Equivalent FTS monitoring

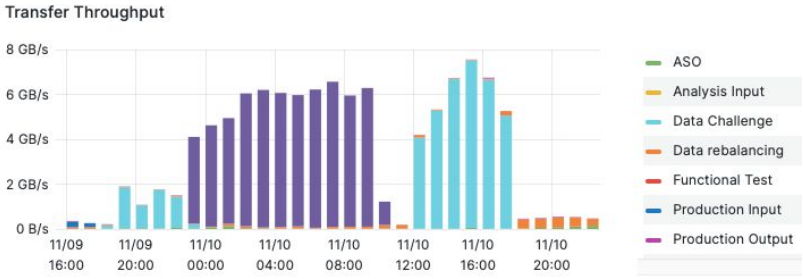

Various undersea cables cut…

Beyond throughput

- WLCG DCs should also (scale) test new technologies
  - Deployment can vary depending on level of matureness
- Some technical topics addressed in the context of DC24
  - Measures to improve monitoring
    - Site based network monitoring
      - Captures all traffic
    - Network flow marking
      - with *SciTags* and UDP *Fireflies*
  - Software Defined Networking (SDN)
    - NOTED
    - SENSE/Rucio
  - Low level network stack
    - Jumbo frames
    - BBRv2, BBRv3 TCP stacks

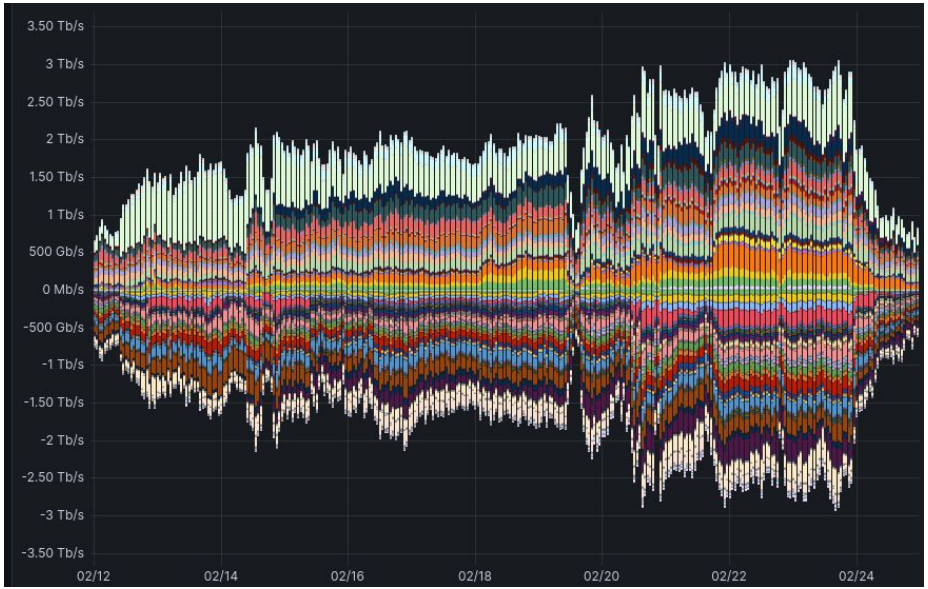

After the Challenge is before the next Challenge

Aftermath of DC24

- Derive 'lessons learned'
  - What went well, identify bottlenecks, organizational improvements …
- Set priorities of 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 **about 50%** of expected HL-LHC throughput
- Dates
  - Likely in 2026 or even later
  - Almost for sure in LS3, which makes scheduling much easier for LHC experiments
- Participating experiments
  - LHC experiments, hopefully again Belle-2 and DUNE
  - Interest (already expressed during DC24) by JUNO, SKA, Neutrino experiments in Japan
- Experience shows that planning needs to start early (1 year before at least)

Some random preliminary observations & remarks

- DC24 was a success - we learnt a lot
  - Testing in advance was very helpful
  - Post-DC24 analysis goes on…
- There are other bottlenecks than network bandwidth
  - Maintenance of DC injections was challenging
    - FTS instances got pushed to their limits, particular the ATLAS one
      - Overloading FTS was a concern
    - Keeping up with deletions is not trivial, systems not best designed for scaling here
    - Other issues with certain sites such as overloading or inability to sustain enough connections
  - It needs time before a complex system reacts to parameter changes
    - The parameter space is huge
    - Not much opportunity to make changes

Joint WLCG/HSF workshop at DESY

- May 13-17th in Hamburg
- Topics include data challenges, analysis facilitates, software tools and training
- Lots of details and information is available now on the workshop indico
- Book your accommodation now, DESY Hostel rooms are still available
- Dinner at the Altes Mädchen Craft Beer Brewery and Restaurant
- Registration is now open: 250€ until April 8th, rising to 275€, registration closes on April 26th

What AI thinks we are doing …

Bing Image Creator: "Worldwide LHC Computing Grid, Data Challenge Workshop, Happy Mood"

Bing Image Creator: "Worldwide LHC Computing Grid, Data Challenge Workshop, Serious Mood"