

Tape and Network Data Challenges

A. Forti, X. Zhao, M. Lassnig

LHCOPN

12 October 2021



DC2021

- Network Challenges planned every two years with increasing bandwidth and different tests
- In 2021 the data challenge was split in **two** challenges run back to back
 - Network Data Challenge 4-8/10/2021
 - Tape Challenge 11-15/10/2021
- Network Data Challenge 2021 goals
 - Commission HTTP-TPC
 - Demonstrate we can fill 10% the bandwidth that is requested at HL-LHC scale
- Tape Challenge goal
 - Goal is to validate tape bandwidth on WLCG tape sites (Tier0, Tier1s and even Tier2 if available), at the Run3 scale, for all experiments



Network Challenge targets

	%ATLAS	%CMS	% Alice	% LHCb	ATLAS+CMS Network Needs (Gbps) Minimal Scenario in 2027	Alice Network Needs (Gbps) Minimal Scenario in 2027	LHCb Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Flexible Scenario in 2027
T1									
CA-TRIUMF	10	0	0	0	196	0	0	196	393
DE-KIT	12	11	21	17	473	85	66	624	1247
ES-PIC	4	5	0	4	169	0	16	185	370
FR-CCIN2P3	13	10	14	15	448	56	58	562	1124
IT-INFN-CNAF	9	15	26	24	472	105	95	673	1345
KR-KISTI-GSDC	0	0	12	0	0	50	0	50	99
NDGF	6	0	8	0	111	31	0	142	285
NL-T1	7	0	3	8	143	12	34	188	376
NRC-KI-T1	3	0	13	5	51	51	21	123	247
UK-T1-RAL	15	9	3	27	472	10	109	591	1183
RU-JINR-T1	0	5	0	0	103	0	0	103	207
US-T1-BNL	23	0	0	0	453	0	0	453	906
US-FNAL-CMS	0	45	0	0	909	0	0	909	1817
(atlantic link)					1362	0	0	1362	2723
Sum	100	100	100	100	4000	400	400	4800	9600

Network numbers from the [HL-LHC document](#)

Above global HL-LHC numbers for minimal and flexible models split by T1 and VO

Below the rates split by T1 for the minimal model for DC transfers this year the next years.

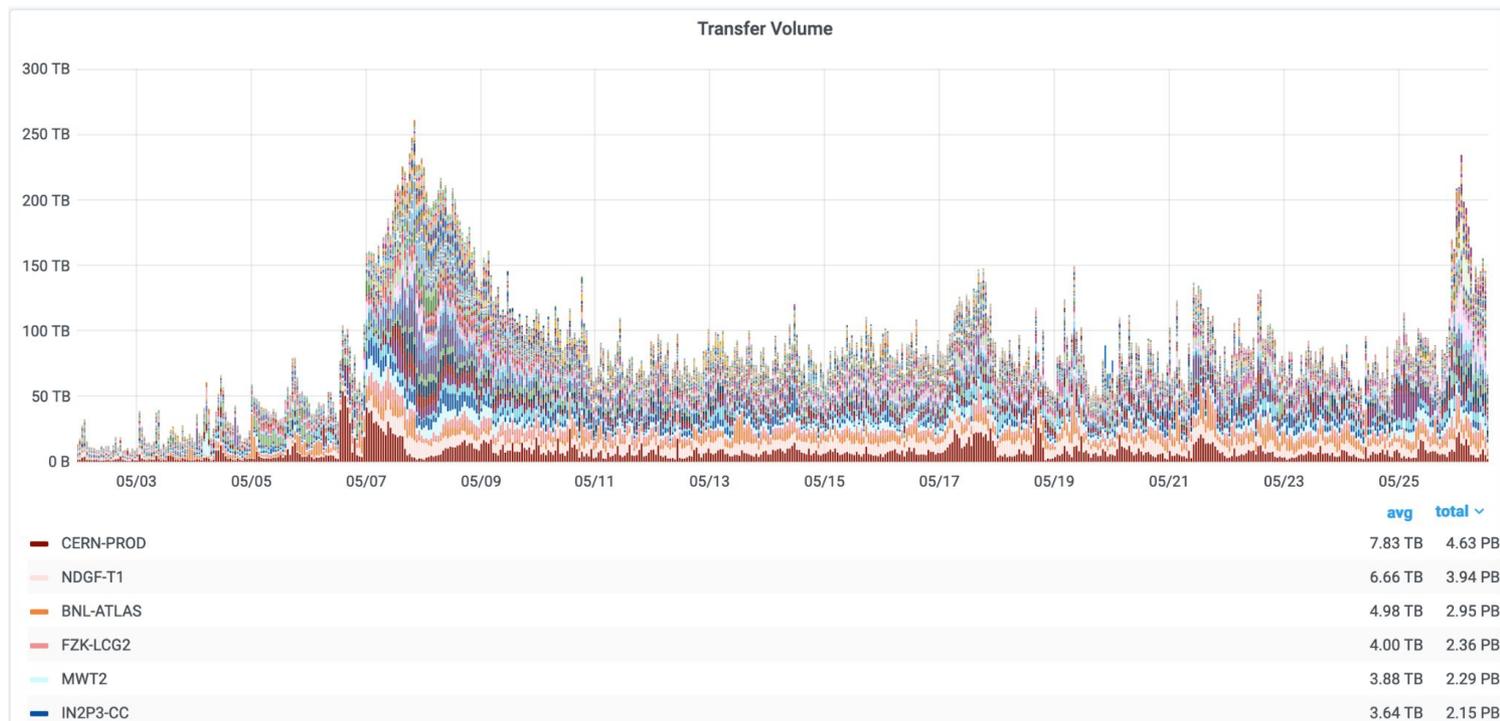
- Numbers in this table are 1/2 of the target.
- ingress/egress 50/50

	Data Challenge target 2027 (Gbps)	Data Challenge target 2025 (Gbps)	Data Challenge target 2023 (Gbps)	Data Challenge target 2021 (Gbps)
T1				
CA-TRIUMF	98	59	29	10
DE-KIT	312	187	94	31
ES-PIC	93	56	28	9
FR-CCIN2P3	281	169	84	28
IT-INFN-CNAF	336	202	101	34
KR-KISTI-GSDC	25	15	7	2
NDGF	71	43	21	7
NL-T1	94	56	28	9
NRC-KI-T1	62	37	19	6
UK-T1-RAL	296	177	89	30
RU-JINR-T1	52	31	15	5
US-T1-BNL	227	136	68	23
US-FNAL-CMS	454	273	136	45
(atlantic link)	681	408	204	68
Sum	2400	1440	720	240



Network DC Models

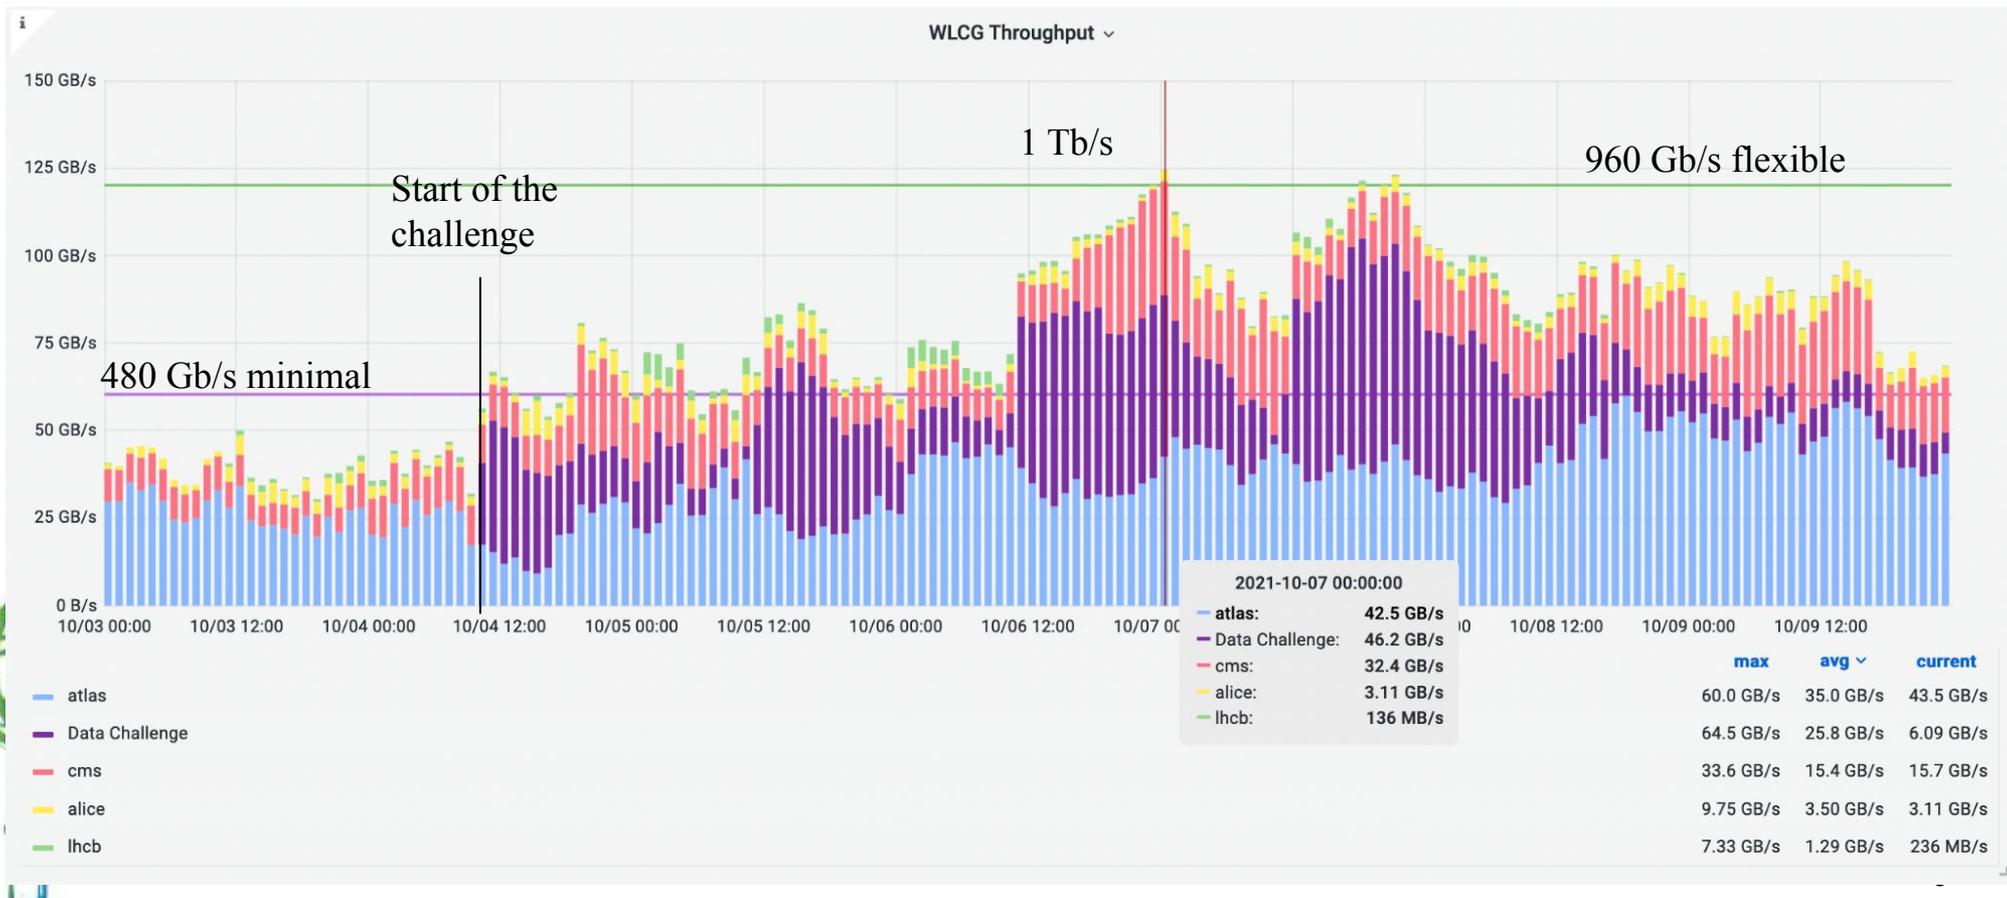
- Two models very different
 - Minimal: hierarchical model T0-T1-T2 traffic
 - Flexible: chaotic model currently more realistic
- T0-T1 Data Taking rates only thing in common
- Challenge planned using chaotic realistic model



Example of standard atlas traffic per destination site

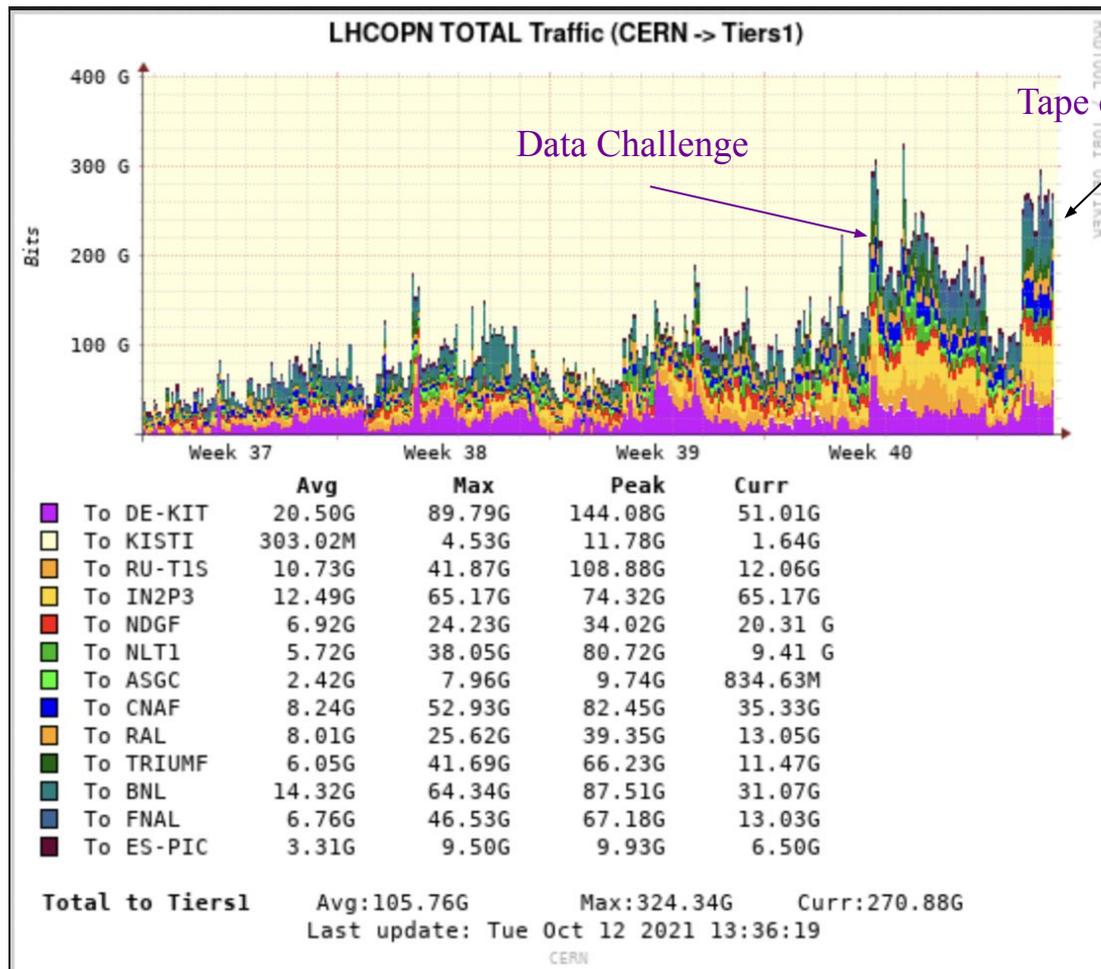
Global Traffic

- Global traffic FTS + xrootd
 - xrootd Alice and part of CMS production traffic
- But DC injected using FTS
 - Most numbers for FTS only
- 20% of minimal model, 10% of flexible



Matching with network traffic

- Matching FTS traffic with network traffic not straight forward
 - For example TC peak higher than DC but less data injected

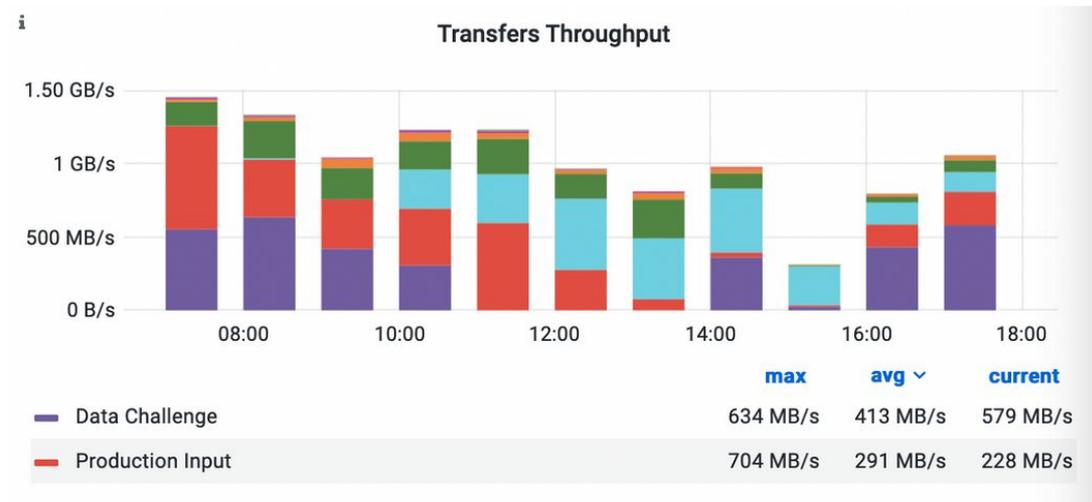


- 3 Cases

- T2
- T1
- T0

Manchester T2

- Very simple case
 - single 10Gb/s pipe
 - Mostly incoming traffic from ATLAS Data Challenge
 - Same monitoring (grafana)
- Almost overlapping
 - Particularly the drop on the right

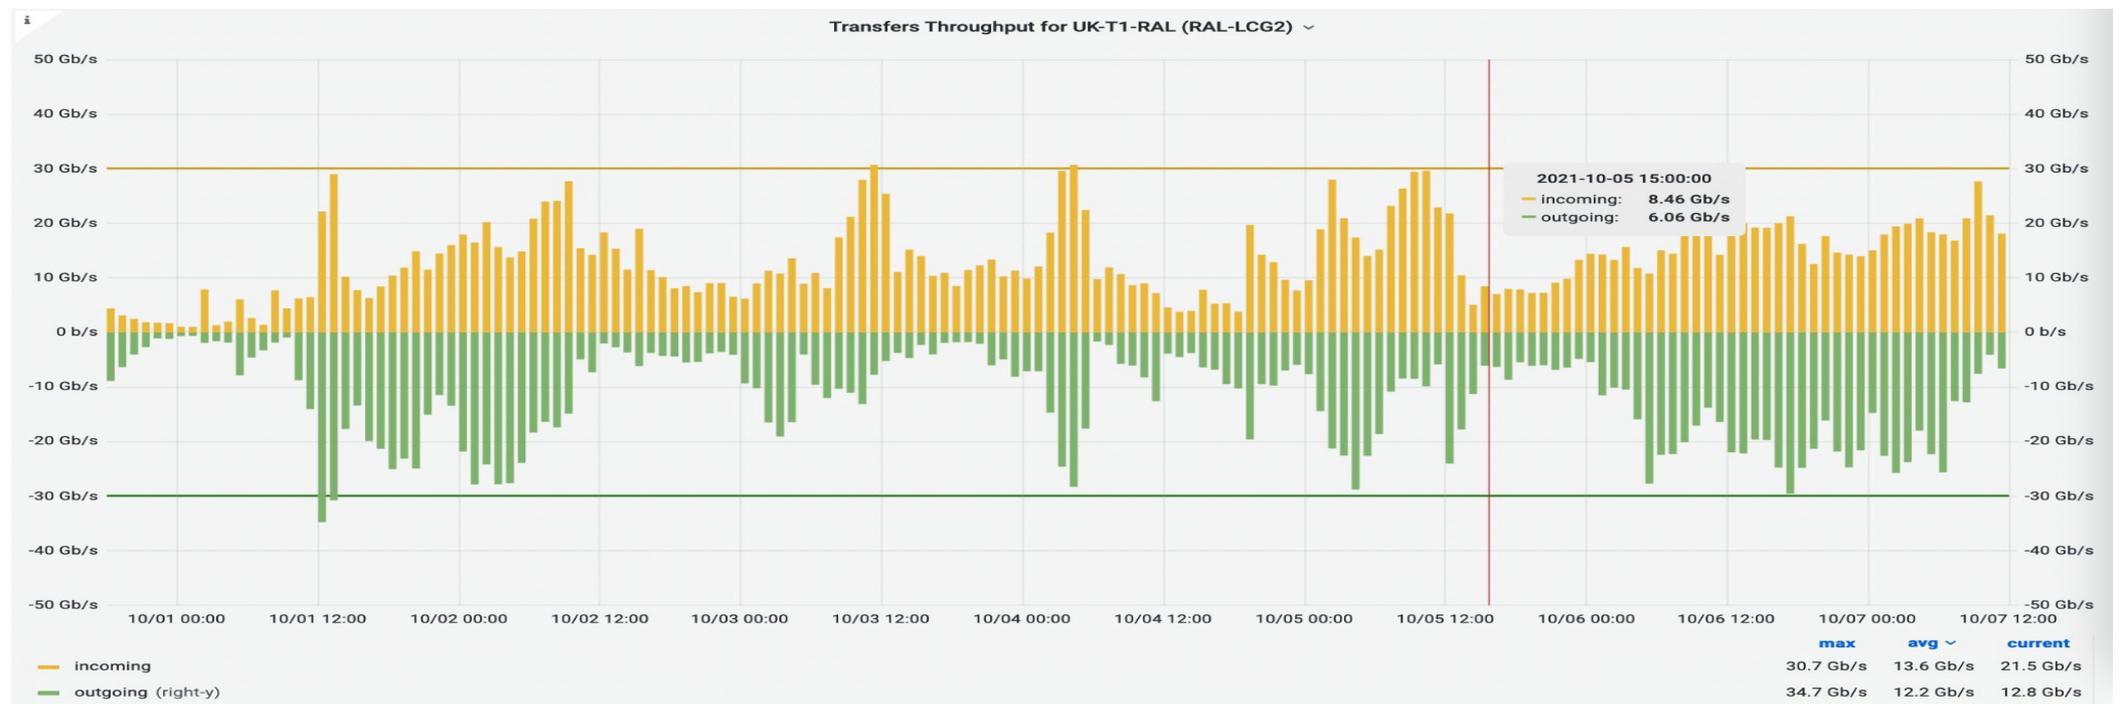
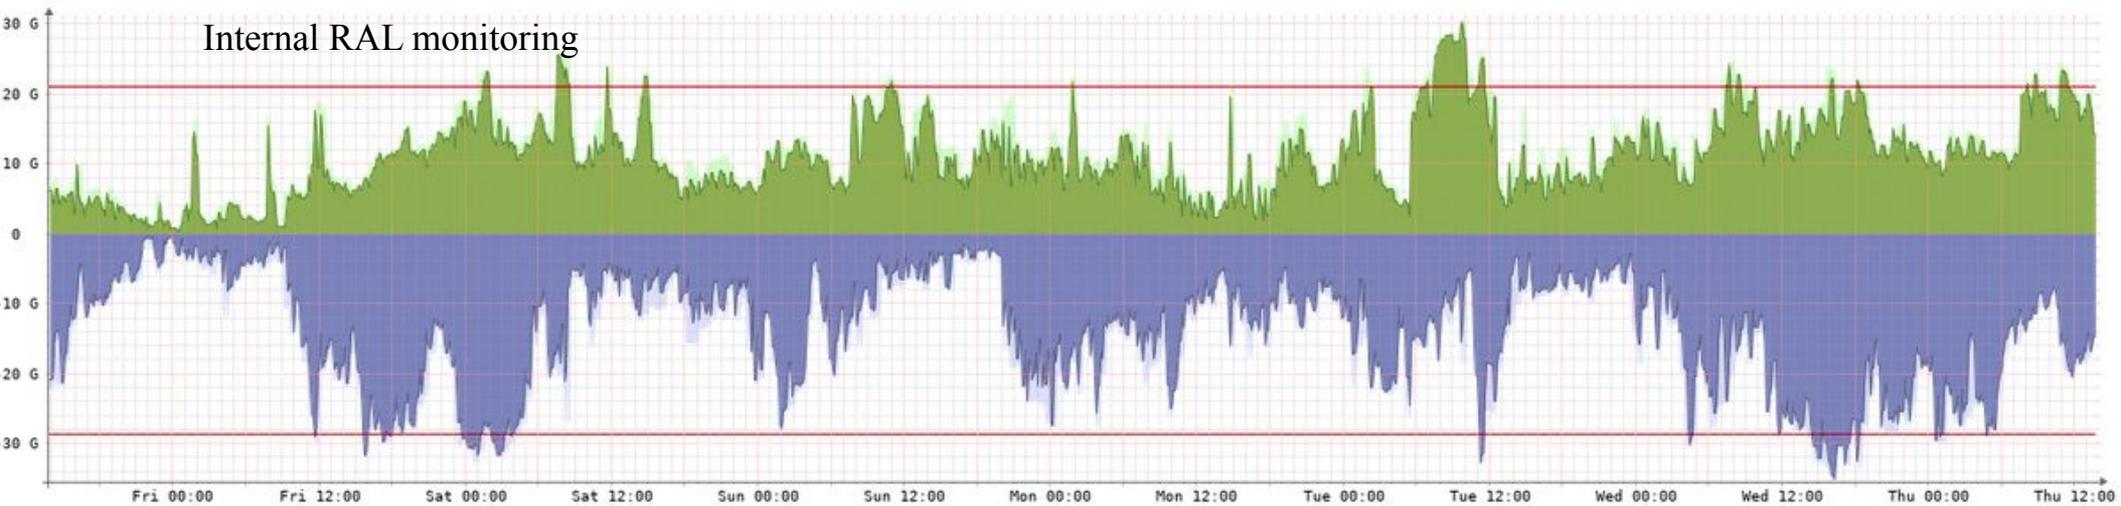


RAL T1

- More complicated network
- Different monitoring
 - Internal
 - CERN LHCOPN
- IPv4/IPv6 traffic

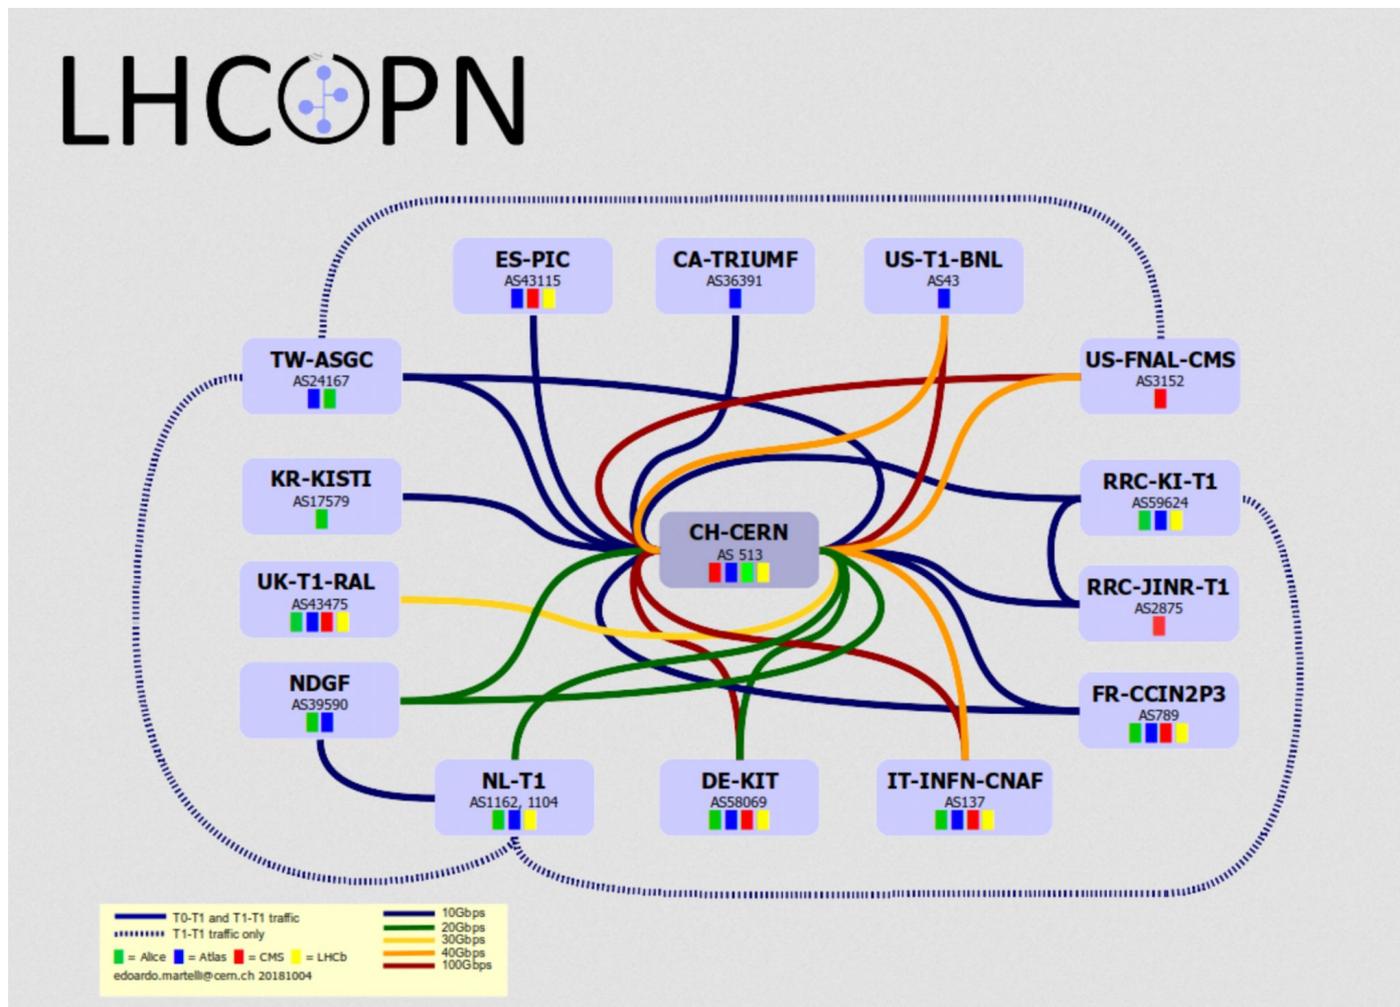


RAL T1



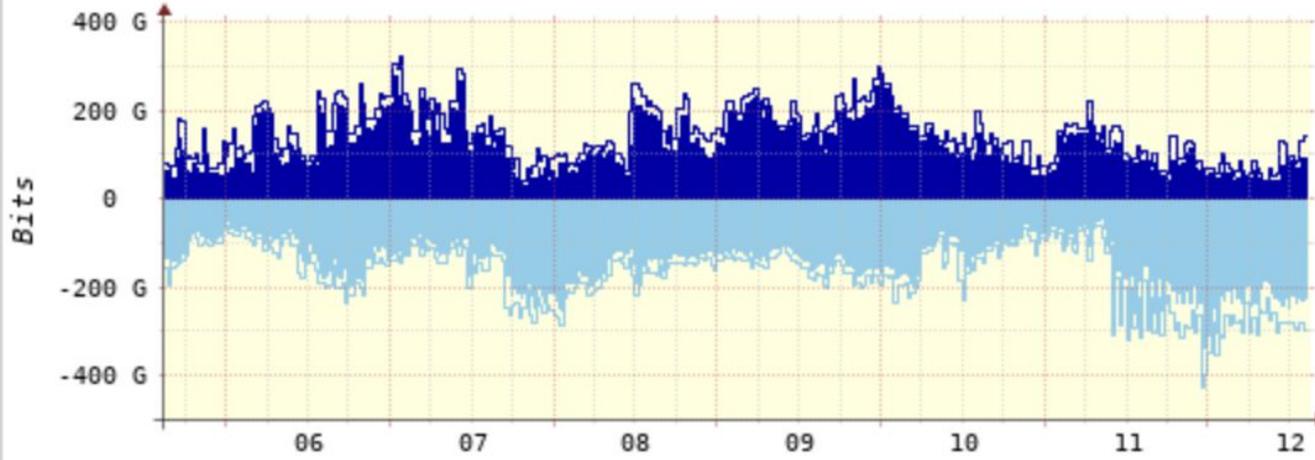
CERN T0

- What T0-T1 traffic?
 - T0-T1 FTS traffic \neq LHCOPN T1-T0-T1



CERN T0

CERN WLCG Traffic to LHCOPN and LHCONE



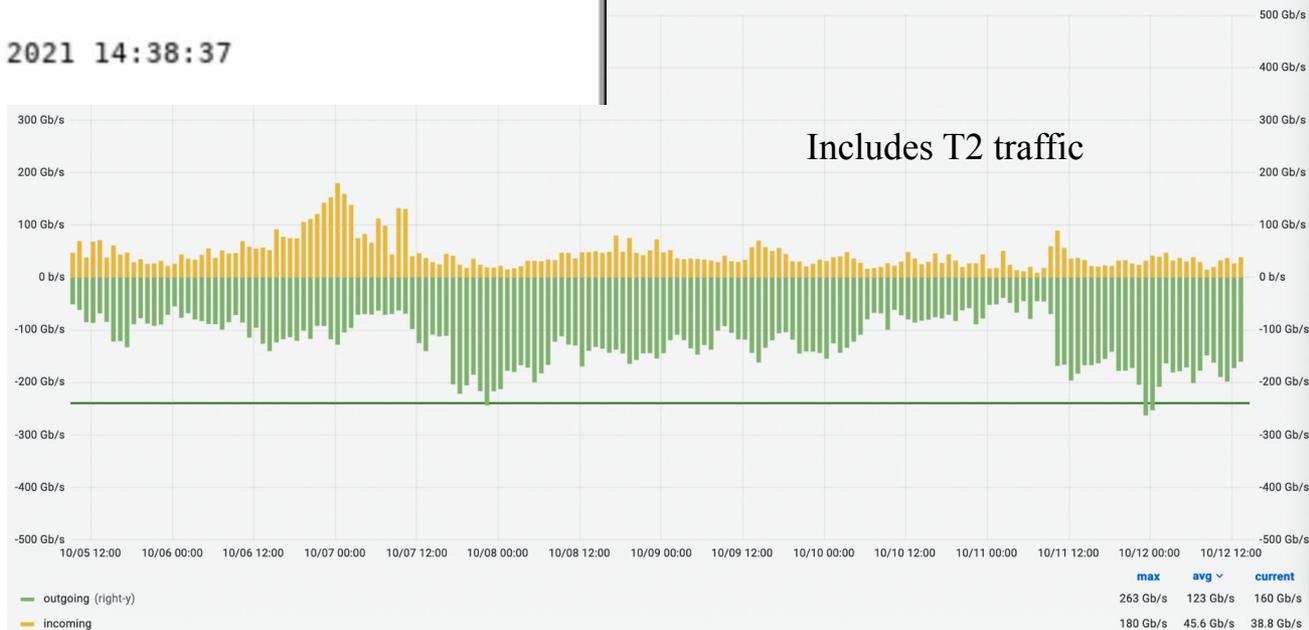
	Avg	Max	Peak	Curr
■ CERN to TierXs	147.21G	341.15G	424.95G	227.72 G
■ TierXs to CERN	124.34G	288.74G	322.29G	92.57 G
／ from CERN peaks				
／ to CERN peaks				

Last update: Tue Oct 12 2021 14:38:37

CERN

Shape for outgoing traffic matches not clear the rates do and incoming traffic different missing xrootd traffic might be a possibility...

Throughput for CERN-T0 (CERN-PROD)



Includes T2 traffic

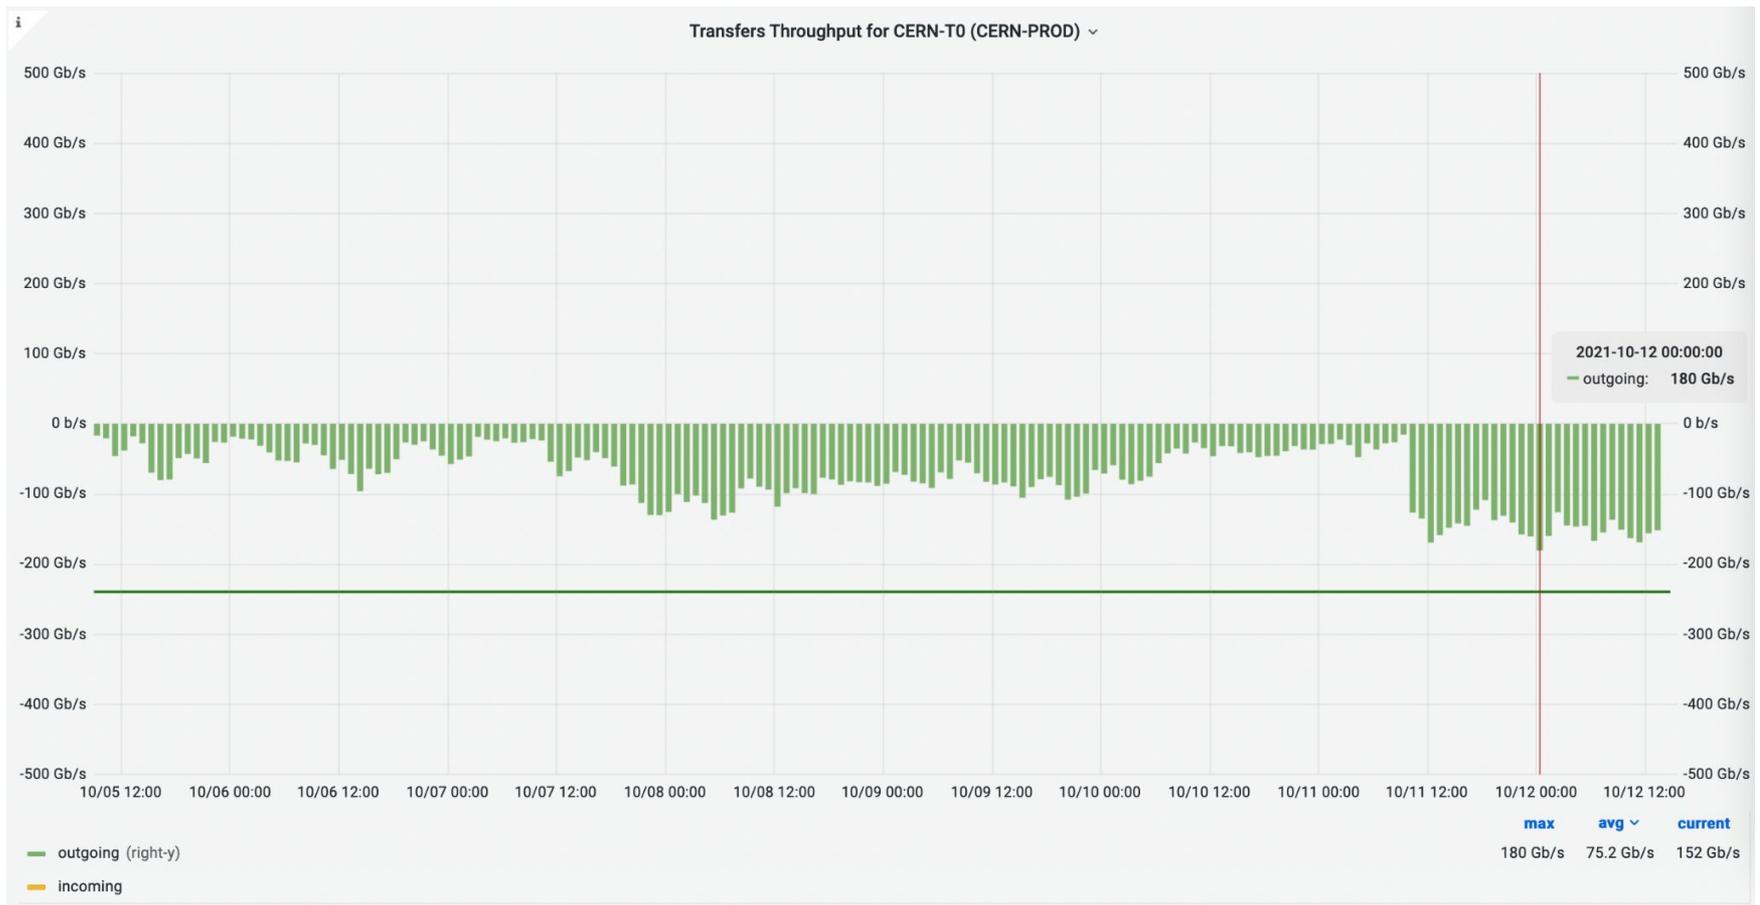
— outgoing (right-y)
— incoming

max	avg	current
263 Gb/s	123 Gb/s	160 Gb/s
180 Gb/s	45.6 Gb/s	38.8 Gb/s



FTS T0-T1

- Couldn't really fill the bandwidth during the DC
- Tape Challenge doing better
- Maybe dedicated tests needed



Conclusions

- Network DC was a success
 - Reached 10% of flexible model with chaotic rather than hierarchical traffic
- Consistency between DC monitoring and network monitoring
- Other sites should do similar matching work?
- Need to better understand T0-T1 case with another round of tests
 - Tape challenge with run3 rates filling bandwidth more consistently
 - There is no data deletion between injections and the traffic is more sustained 180 Gb/s over an expected 240 Gb/s
 - Could be done again with larger amount of data



Backup



Links

- [HL-LHC network needs and data transfer challenges](#)
- [Tape Challenge Planning doc](#)
- <https://its.cern.ch/jira/projects/DOMATPC>
- <https://gitlab.cern.ch/wlwg-doma>
- doma-data-challenges-development@cern.ch
- <https://mattermost.web.cern.ch/wlwg-gdb/channels/wlwg-data-challenges>
- [Tape Challenge Slack channel](#)
- [DOMA indico category for meetings](#)

