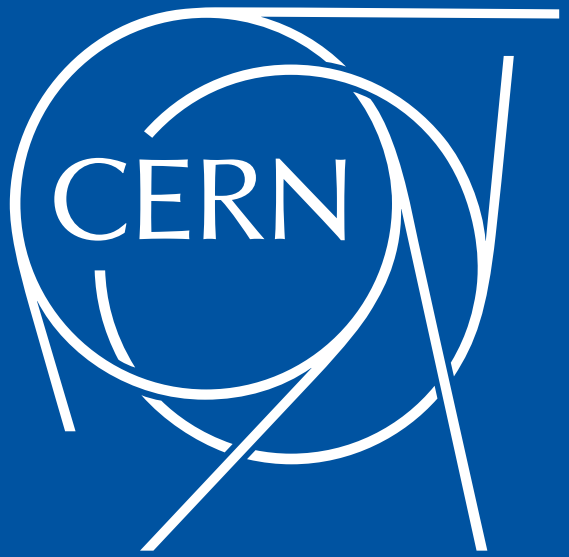
CERN

Understanding network traffic via the analysis of data from File Transfer Service

NOTED

Joanna Waczyńska Edoardo Martelli

March 24, 2021

1. Introduction & Technical side
2. Effects
3. Tool developing = Traffic forecasting
4. Conclusion

Section 1

Introduction & Technical side

The goal of the NOTED project

We would like to optimized transfers of LHC data focusing by network problem (like saturation).

How?

- What is generating traffic?
  For WLCG, most network traffic is generated by FTS. We therefore analyse the data transfers via FTS in order to estimate when any network optimisation should be applied.
- Do the transfers contain information about the network topology/configuration?
  No, but the NOTED uses information from the CRIC database to identify the site network prefixes (IPv4/IPv6) of the storage elements involved and then groups transfers.

Figure: Identification and grouping of relevant transfers.

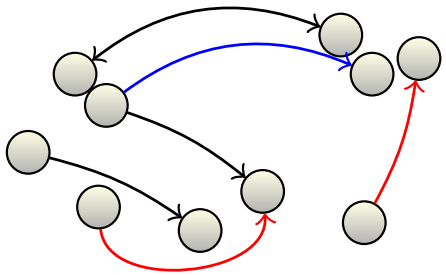

(a) Transfers from FTS between endpoints. Red represents inactive transfers, black active, blue transfer which is now inactive, but it was active.

Figure: Identification and grouping of relevant transfers.

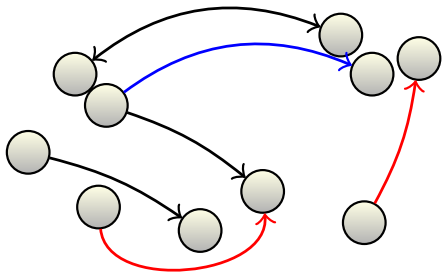

(a) Transfers from FTS between endpoints. Red represents inactive transfers, black active, blue transfer which is now inactive, but it was active.

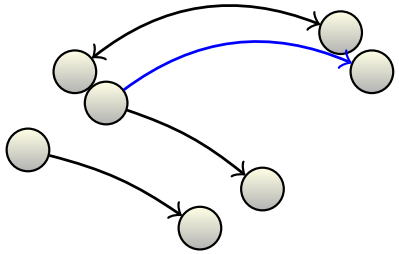

(b) Filter transfers: inactive transfers are discarded. This operation could be defined by input settings.

Figure: Identification and grouping of relevant transfers.

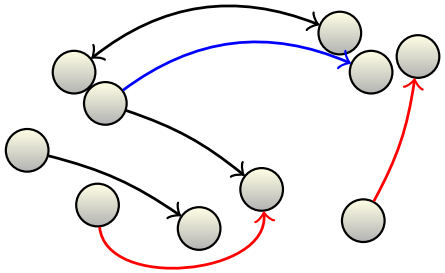

(a) Transfers from FTS between endpoints. Red represents inactive transfers, black active, blue transfer which is now inactive, but it was active.

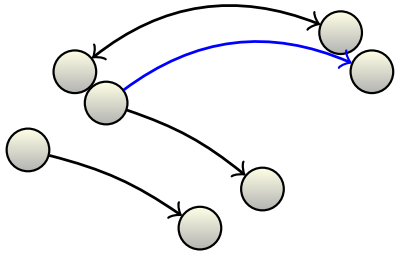

(b) Filter transfers: inactive transfers are discarded. This operation could be defined by input settings.

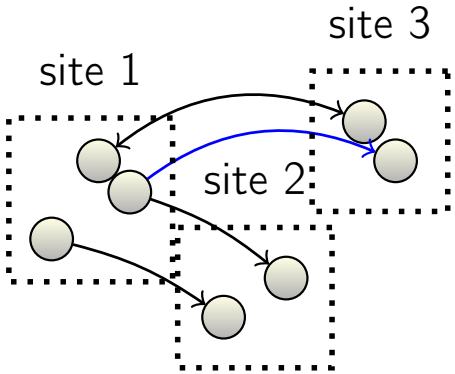

(c) Translate FTS endpoints to unique sites using CRIC, then transfers are grouped using {source,destination} key.

How?

- Do we have enough knowledge about transfers to know which links will be used?

How?

- Do we have enough knowledge about transfers to know which links will be used?
  No. Why?

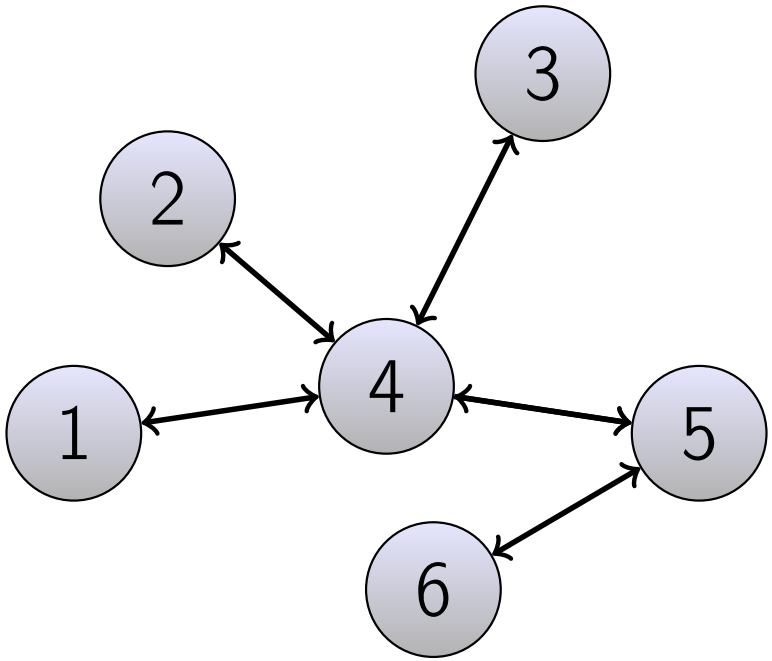

Figure: Example of network topology.

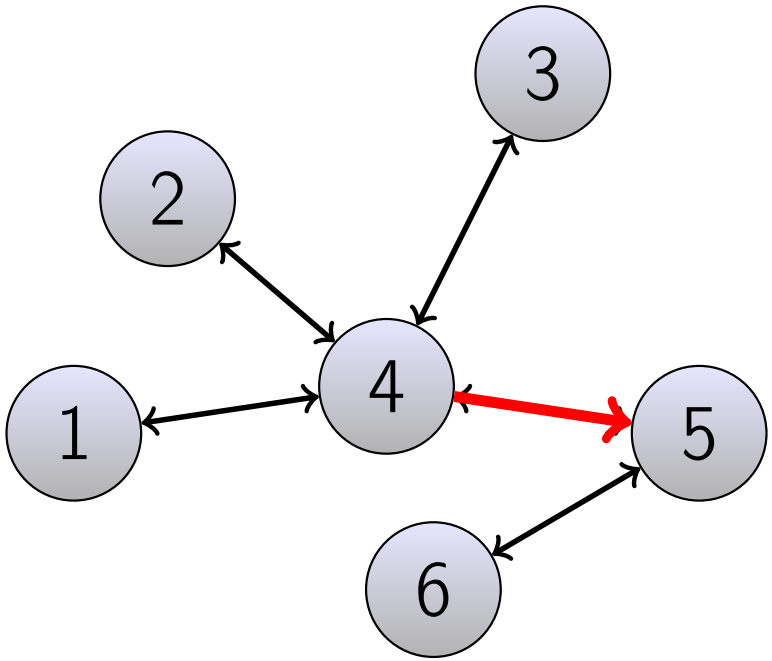

Figure: Example of network topology.

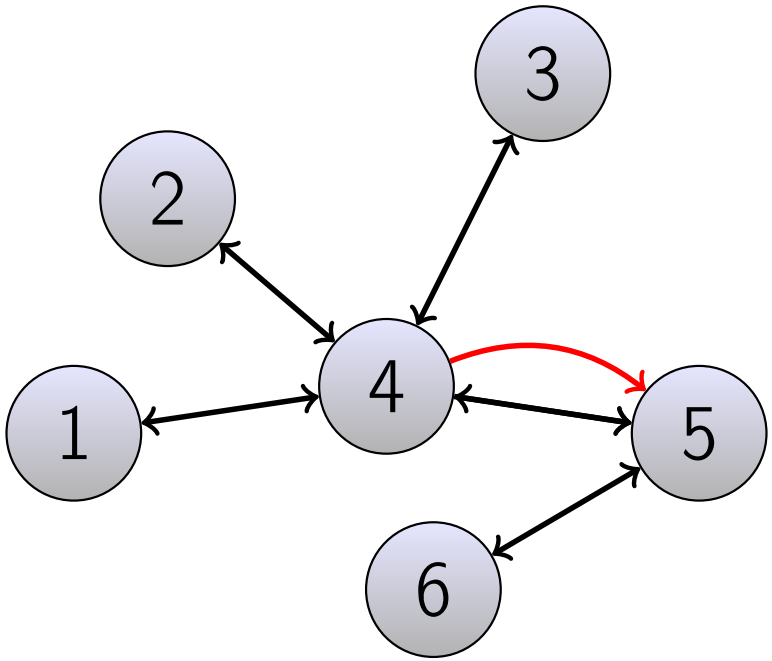

Figure: Example of network topology.

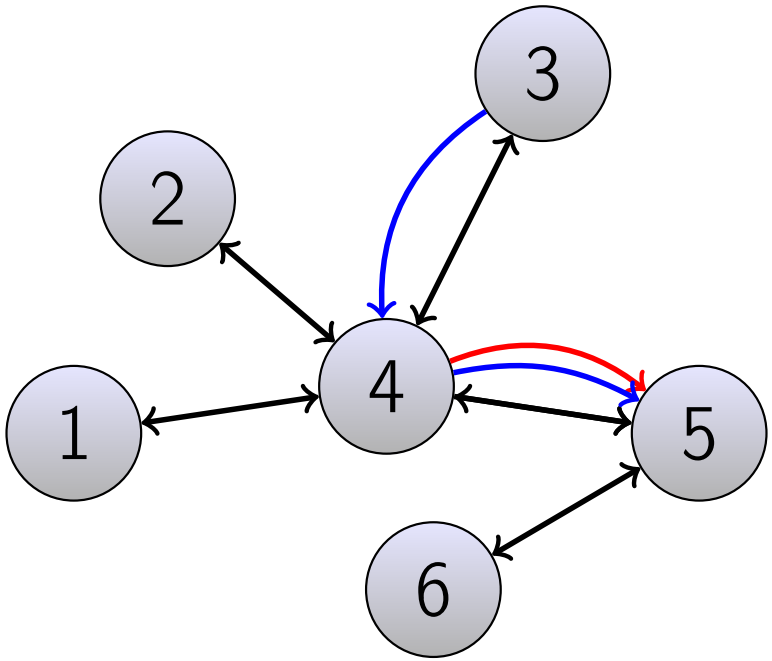

Figure: Example of network topology.

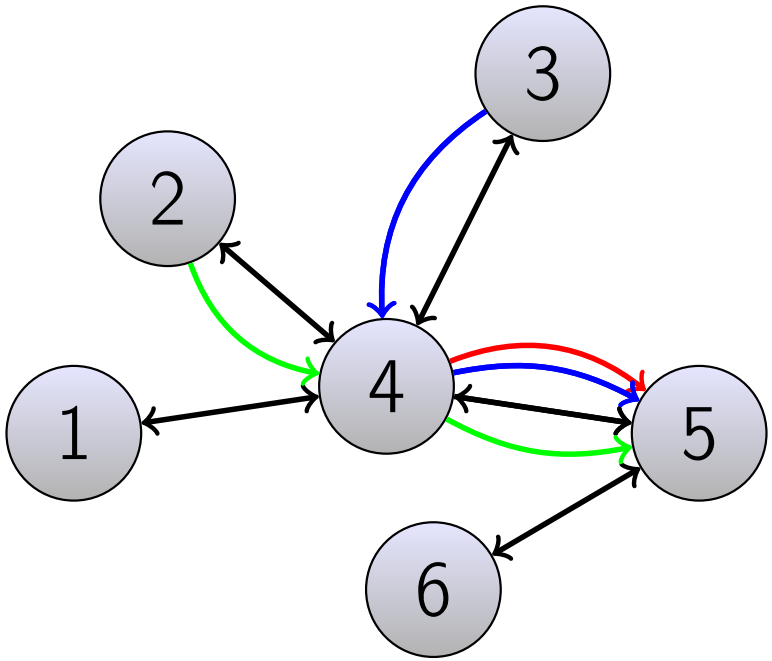

Figure: Example of network topology.

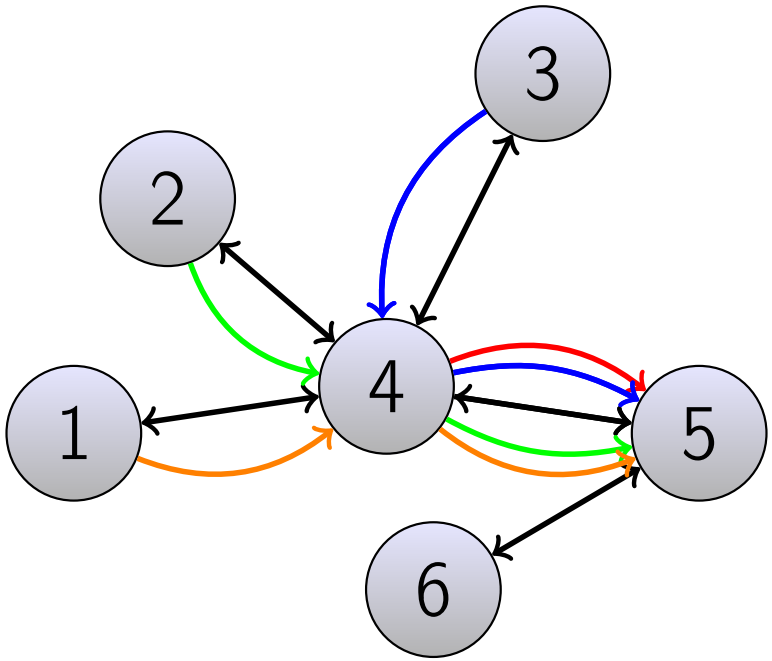

Figure: Example of network topology.

How?

The aggregation stage is a key element of the making decision process as it enables us to combine the impact of potentially independent FTS decisions to give information about the impact on one or more network path segments. We can define our own controller.

Controller configuration

We can define our own controller:

- obligatory
  - what transfers we should consider? For example: CERN -> PIC, or {CERN & TRIUMF-SFU } ->PIC
  - (future optional) define limits: for example focus only for transfers more then 10TB.
- optional:
  - contact information – send information email about decisions
  - postpone decision about removing path when main transfers are finished
  - postpone decision about removing path when FTS reduced throughput, and transfers do not generate enough traffic to saturate link, but transfers have not finished yet
  - corresponding controllers

Section 2

Effects

Detection

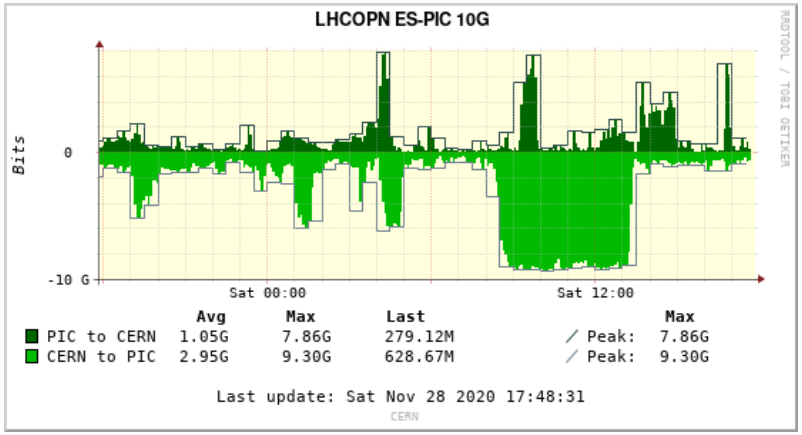

(a) Utilization of the LHCOPN 10Gbps link CERN-PIC on 27/28-11-2020. FTS was the main traffic source.

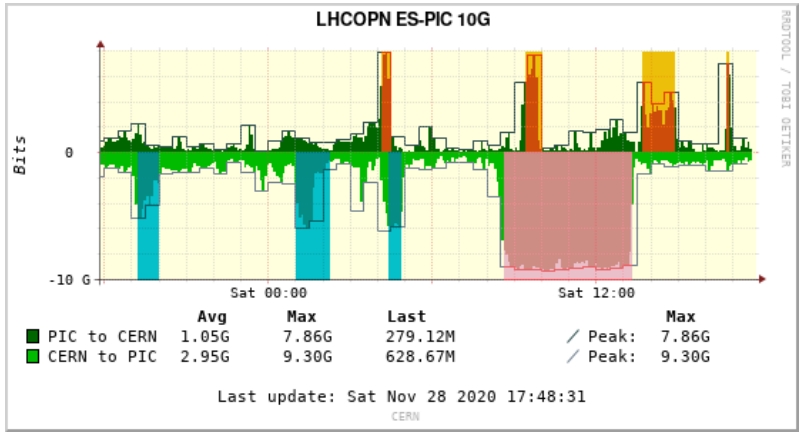

(b) Large transfers detected by the NOTED.

Figure: Juxtaposition of two graphs: from the network monitoring website showing the traffic observed between PIC and CERN and the period when NOTED requested and removed network re-configurations.

LHCOPN

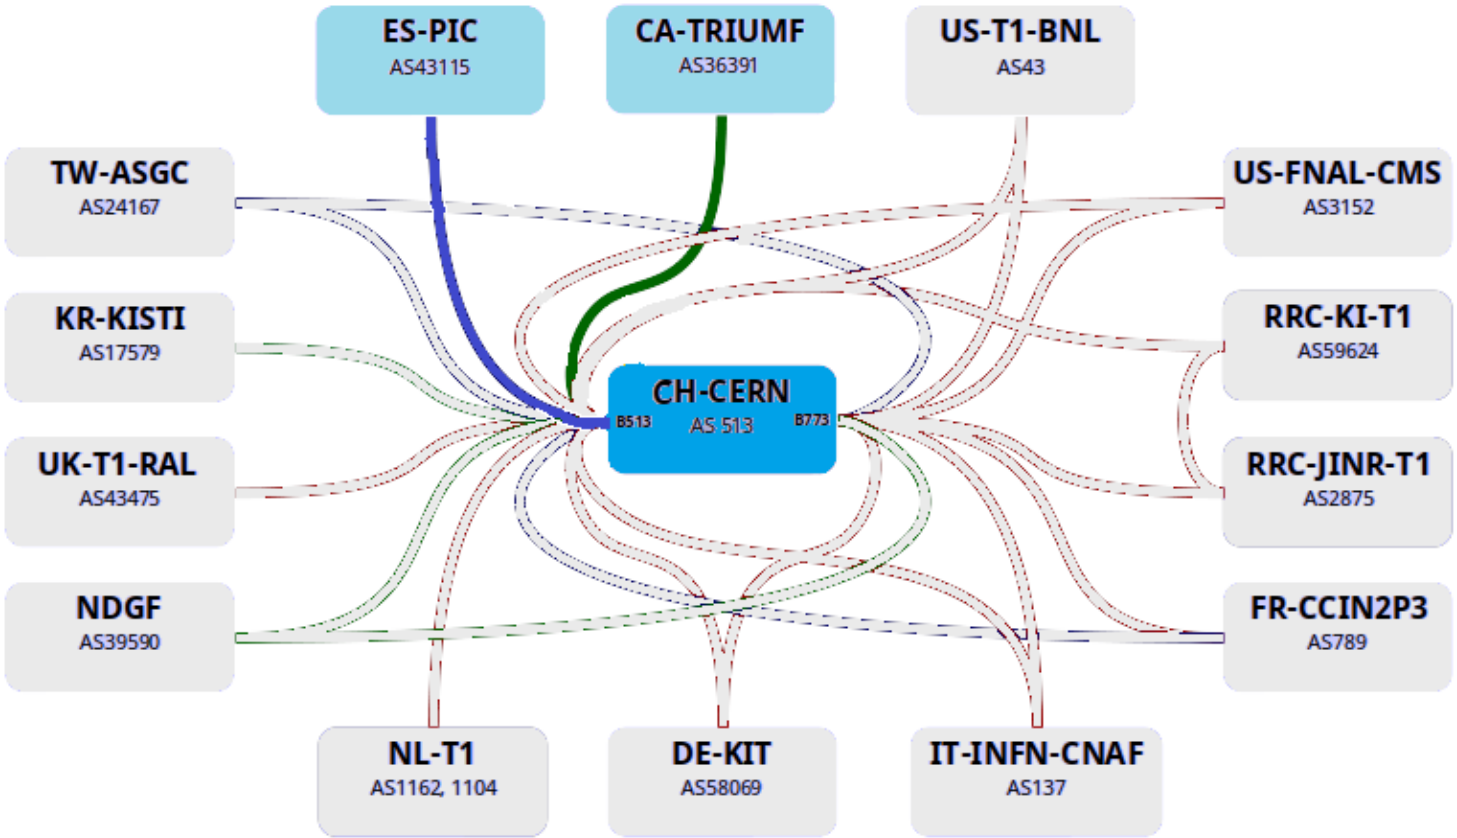
ES-PIC
AS43115
CA-TRIUMF
AS36391
US-T1-BNL
AS43
TW-ASGC
AS24167
US-FNAL-CMS
AS3152
KR-KISTI
AS17579
RRC-KI-T1
AS59624
CH-CERN
B513
AS 513
B773
UK-T1-RAL
AS43475
RRC-JINR-T1
AS2875
NDGF
AS39590
FR-CCIN2P3
AS789
NL-T1
AS1162, 1104
DE-KIT
AS58069
IT-INFN-CNAF
AS137

LHCOPN

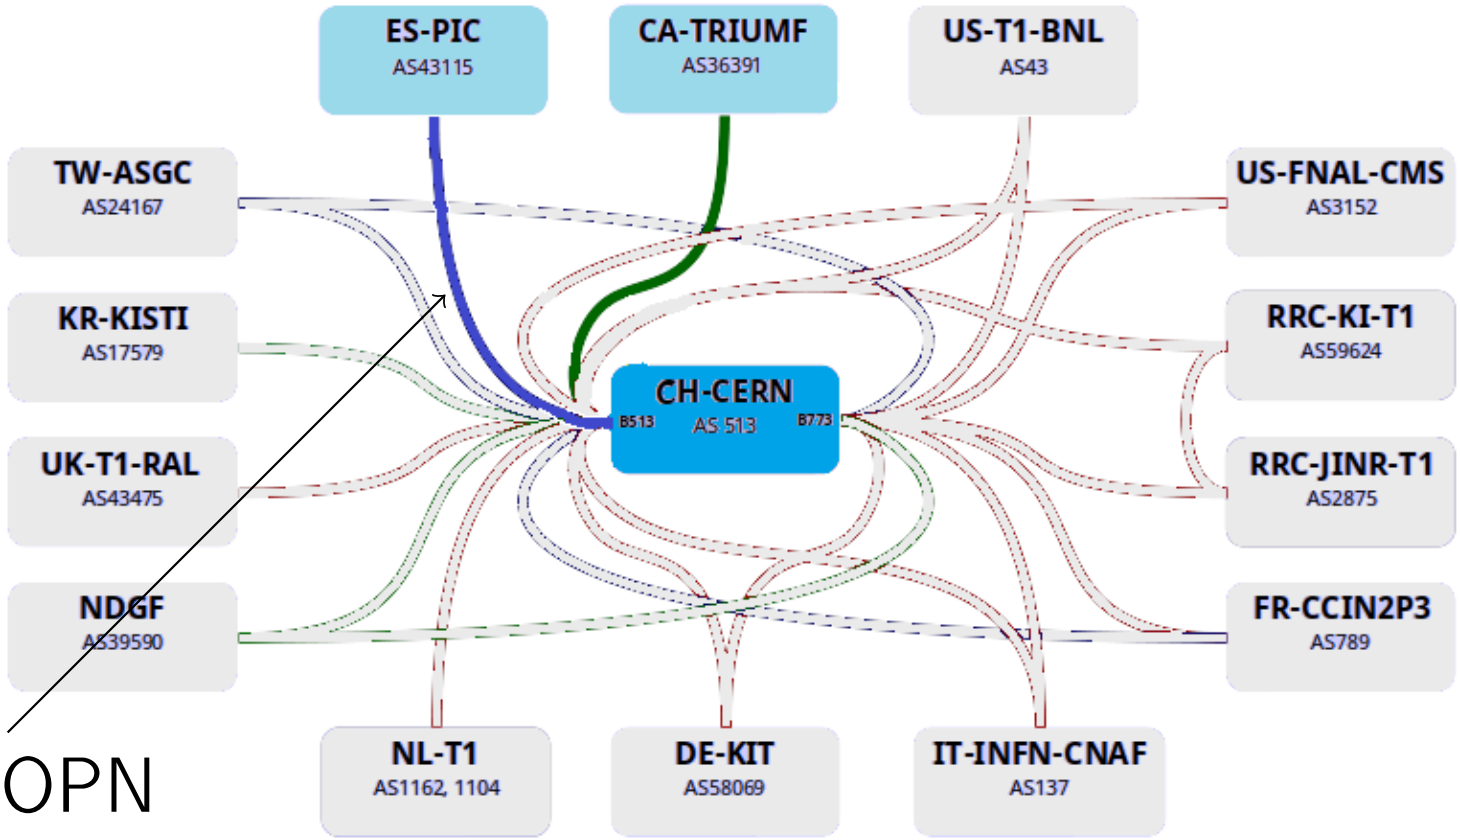

LHCOPN
ES-PIC 10Gbps

LHCOPN

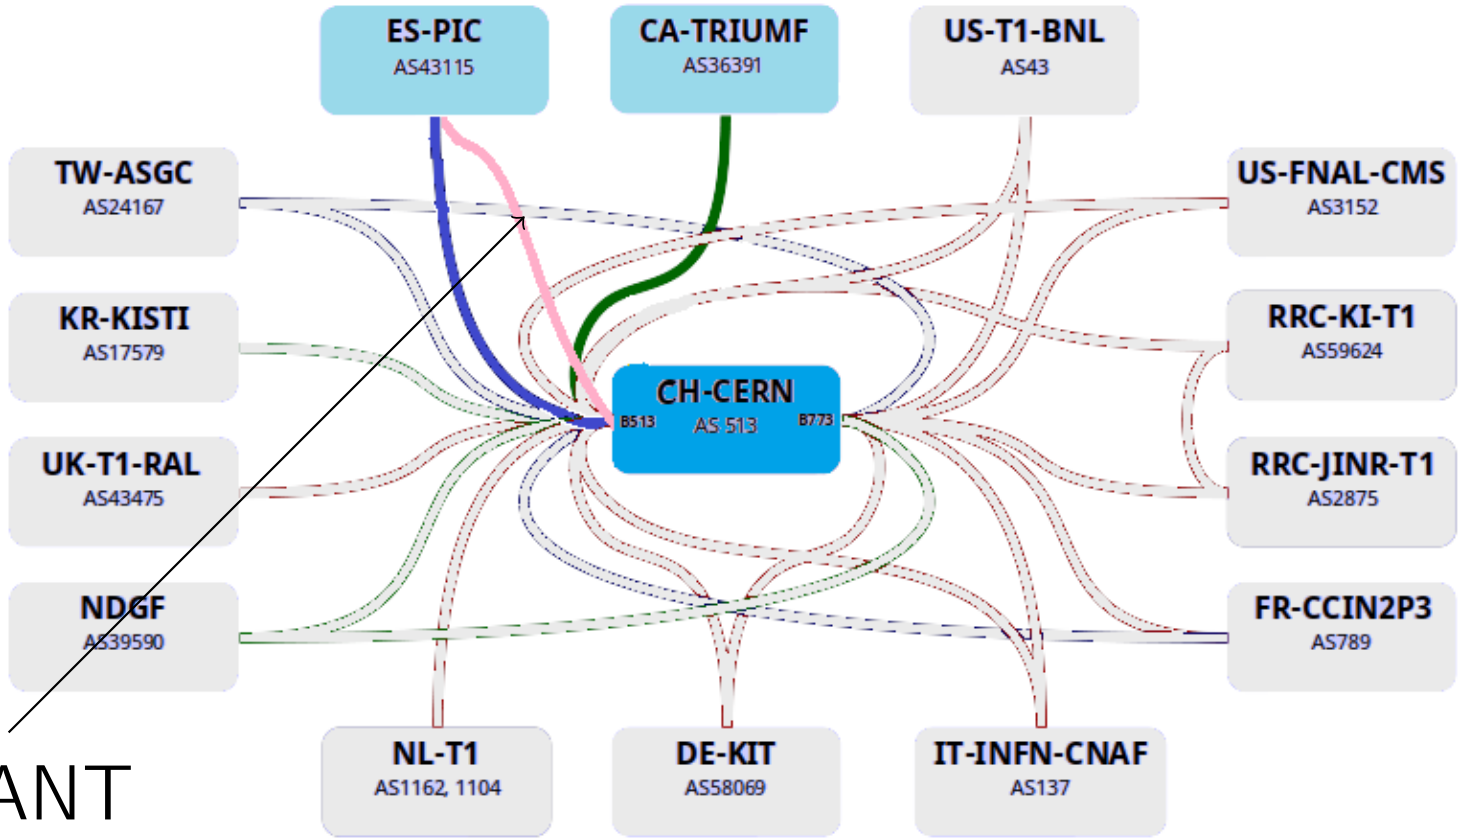

GEANT
LHCONE

LHCOPN

SFU v2127
10Gbps

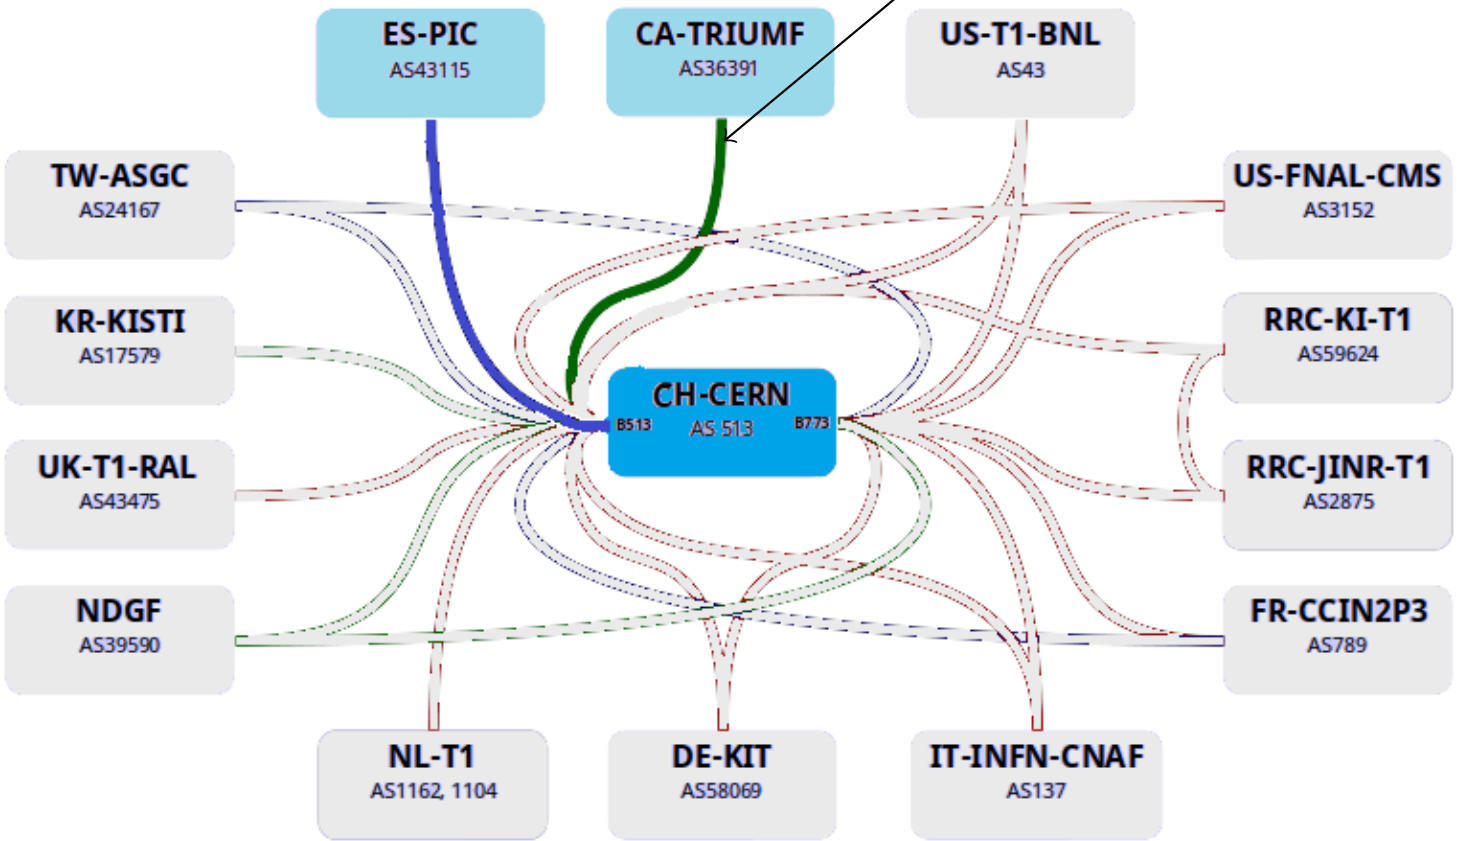

LHCOPN

SFU v2128
50Gbps

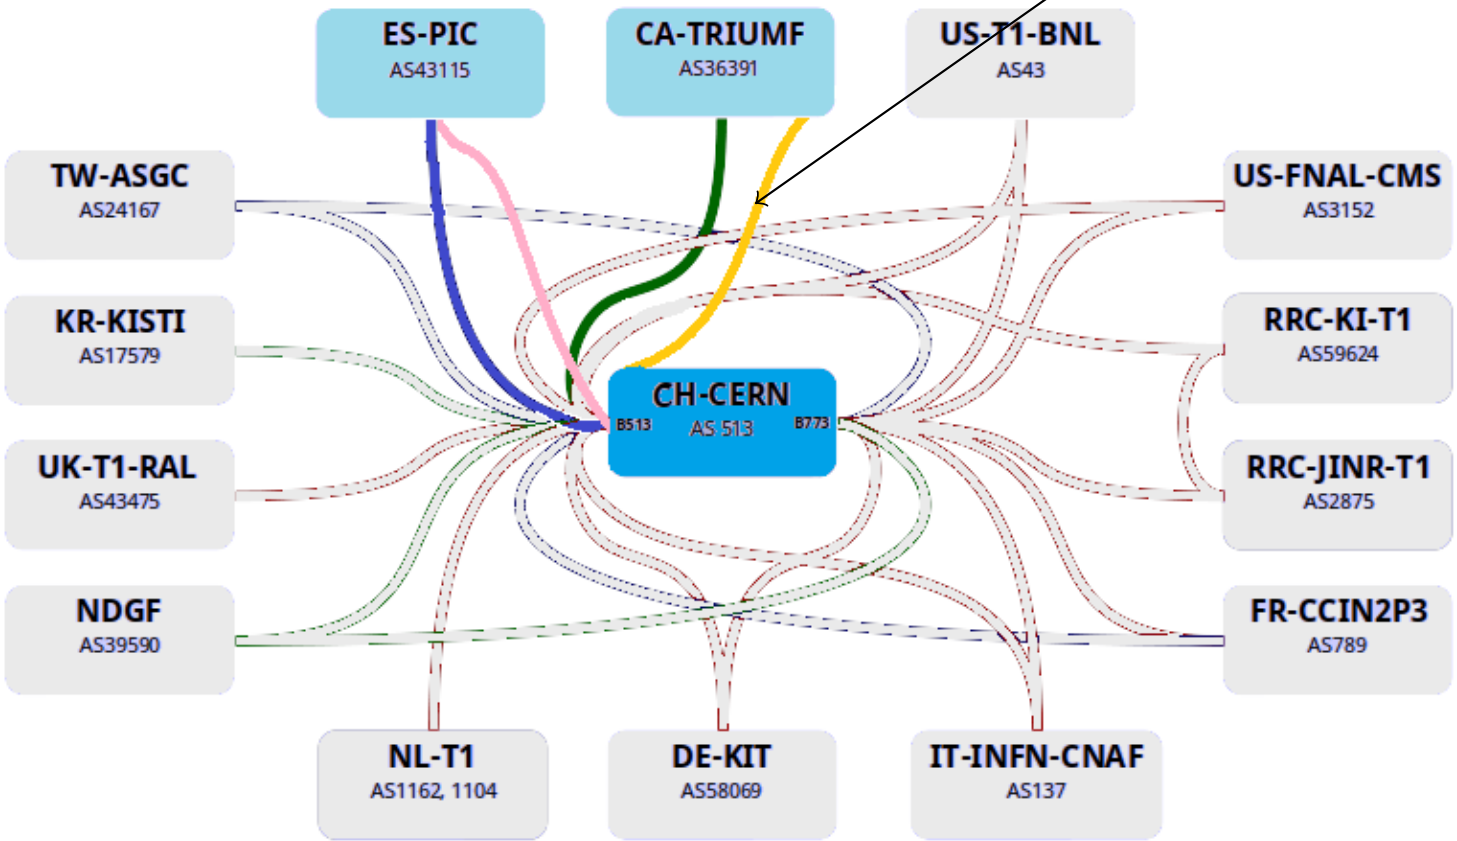

Test with PIC (balancing traffic)

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed

Figure: Report between two endpoints. New bulk (our test transfer) started at 9:15 (UTC). Test transfer contained 3060 files (avg file size >6GB); totalling around 145 TB.

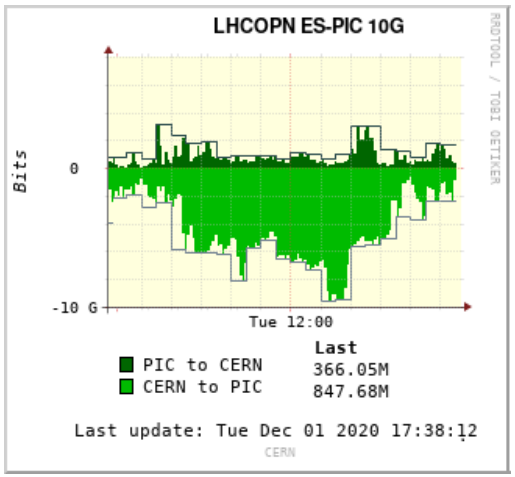

(a) Observation of traffic passing through the 10Gbps path from 01-12-2020. Multiple overlapping transfers can be observed.

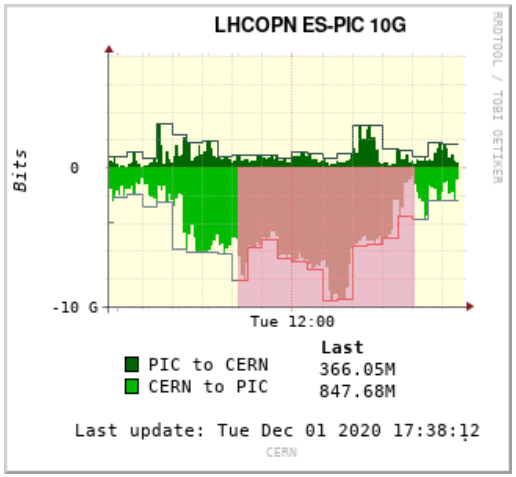

(b) This NI controller focused on transfers from CERN. It detected our test transfer and recommended network optimiisation only for that.

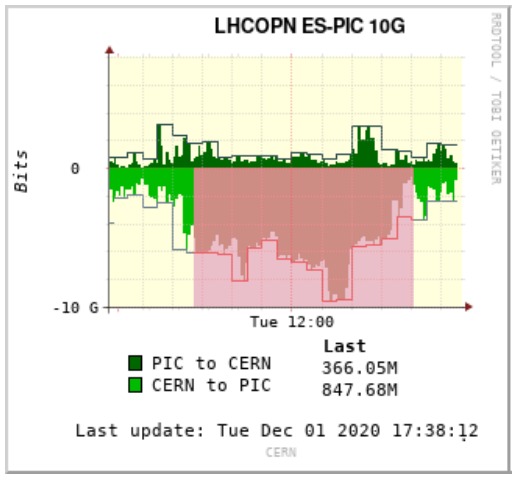

(c) The second NI controller detected and aggregated, all FTS transfers transiting CERN to reach PIC, including those generated by other Tier1s. Network optimisation was thus recommended for a longer period.

Figure: Operation of two NI controllers. The red colouring shows increased traffic due to transfers from FTS.

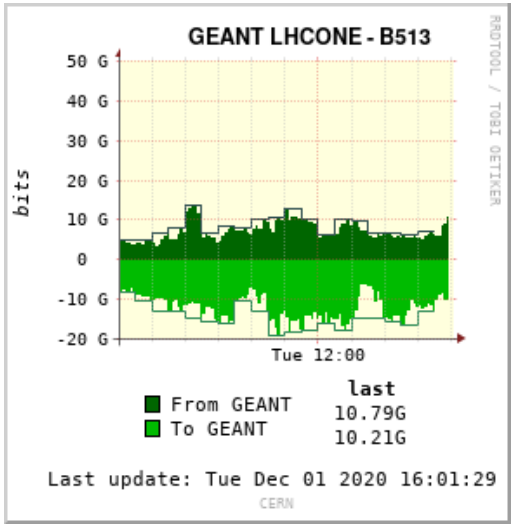

(a) Observation of traffic passing through the added LHCONE_GEANT path from 01-12-2020. Link GEANT was added around 10:15 and in result traffic was balanced.

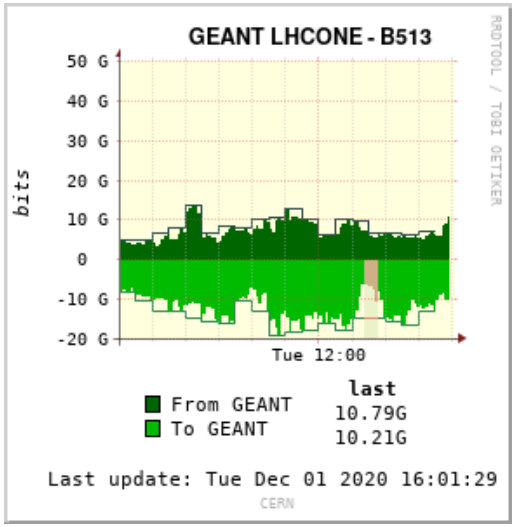

(b) The yellow area presents the period when LHCONE_GEANT path wasn't used to load-balance the traffic between CERN and PIC. The Network traffic decrease in the LHCONE_GEANT path is visible.

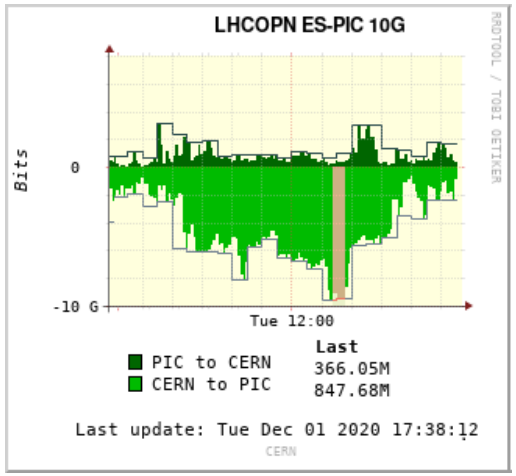

(c) The yellow area presents the period when the LHCONE_GEANT path wasn't used to split traffic between CERN and PIC. Network traffic increased in the LHCONE_ES_PIC path.

Figure: Impact of removing the added path on the network traffic of the observed links.

Test with TRIUMF-SFU Tier1

Scenario: site connected with a direct low speed link (VLAN 2127, 10Gbps) with possibility of requesting a larger bandwidth dynamic circuit
(VLAN 2128, >50Gbps).

Tried to use a dynamic circuit provisioned by AutoGOLE, but a firewall issue prevented the reservation to work. We will be very happy to continue the cooperation.

Test with TRIUMF-SFU 12-03-2021 (forwarding traffic)

3 test transfers:

- CERN -> TRIUMF-SFU
- TRIUMF-SFU -> CERN
- CERN -> TRIUMF-SFU

LHCOPN v2127 10Gbps was default link since 8:30 to 18:30.

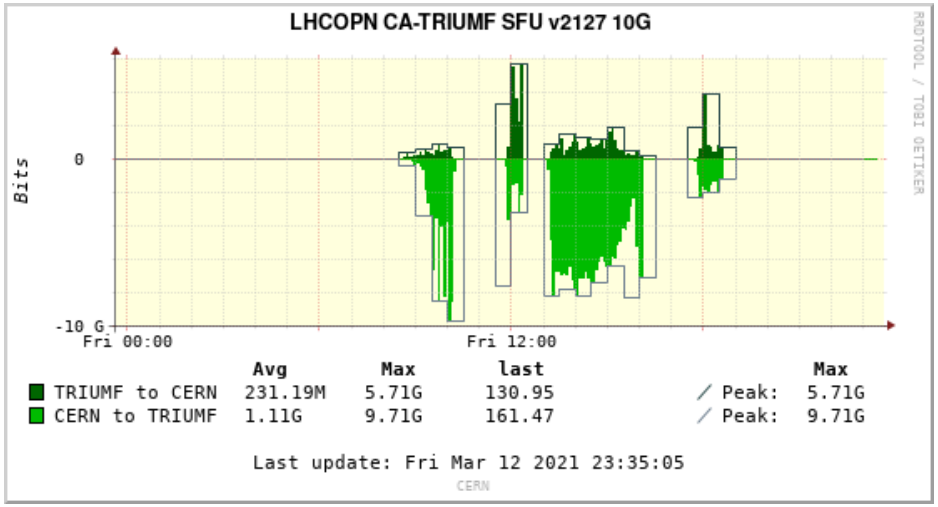
LHCOPN CA-TRIUMF SFU v2127 10G
Bits
0
-10 G
Fri 00:00
Fri 12:00
Avg Max last Max
TRIUMF to CERN 231.19M 5.71G 130.95 / Peak: 5.71G
CERN to TRIUMF 1.11G 9.71G 161.47 / Peak: 9.71G
Last update: Fri Mar 12 2021 23:35:05
CERN
RRDTOOL / TOBI OETIKER

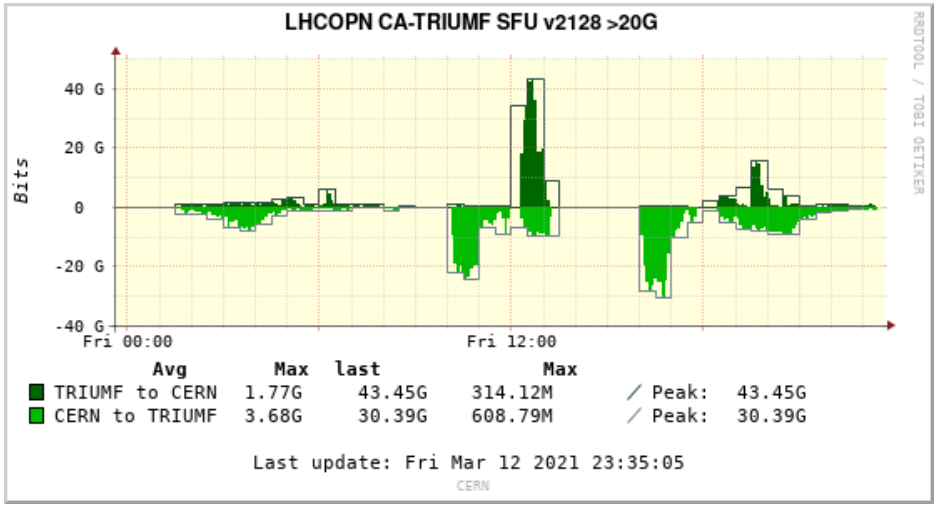
LHCOPN CA-TRIUMF SFU v2128 >20G
Bits
40 G
20 G
0
-20 G
-40 G
Fri 00:00
Fri 12:00
Avg Max last Max
TRIUMF to CERN 1.77G 43.45G 314.12M / Peak: 43.45G
CERN to TRIUMF 3.68G 30.39G 608.79M / Peak: 30.39G
Last update: Fri Mar 12 2021 23:35:05
CERN
RRDTOOL / TOBI OETIKER

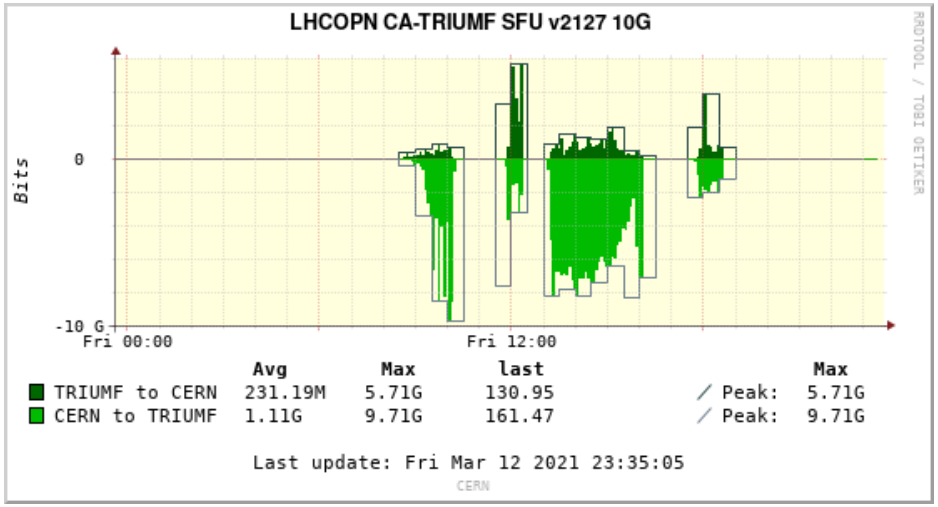
LHCOPN CA-TRIUMF SFU v2127 10G
Bits
0
-10 G
Fri 00:00
Fri 12:00
Avg Max last Max
TRIUMF to CERN 231.19M 5.71G 130.95 / Peak: 5.71G
CERN to TRIUMF 1.11G 9.71G 161.47 / Peak: 9.71G
Last update: Fri Mar 12 2021 23:35:05
CERN
RRDTOOL / TOBI OETIKER

+

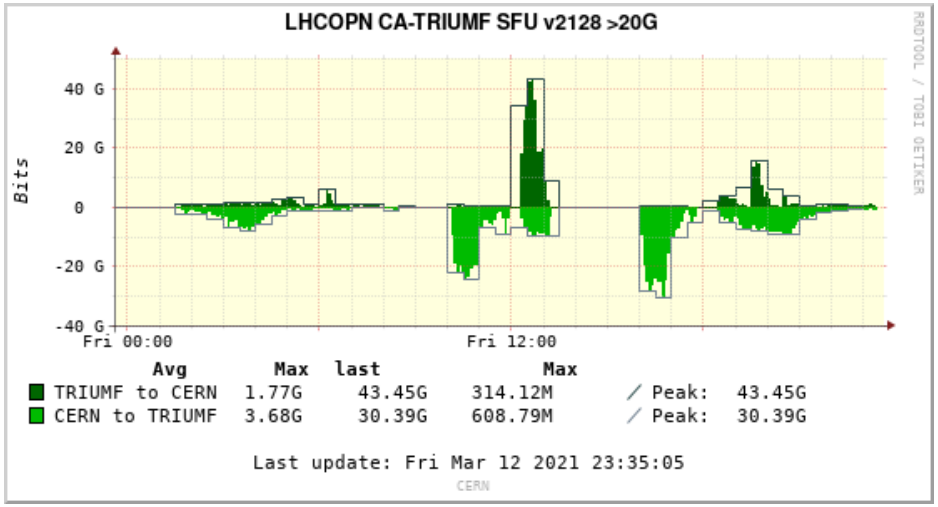
LHCOPN CA-TRIUMF SFU v2128 >20G
Bits
40 G
20 G
0
-20 G
-40 G
Fri 00:00
Fri 12:00
Avg Max last Max
TRIUMF to CERN 1.77G 43.45G 314.12M / Peak: 43.45G
CERN to TRIUMF 3.68G 30.39G 608.79M / Peak: 30.39G
Last update: Fri Mar 12 2021 23:35:05
CERN
RRDTOOL / TOBI OETIKER

=

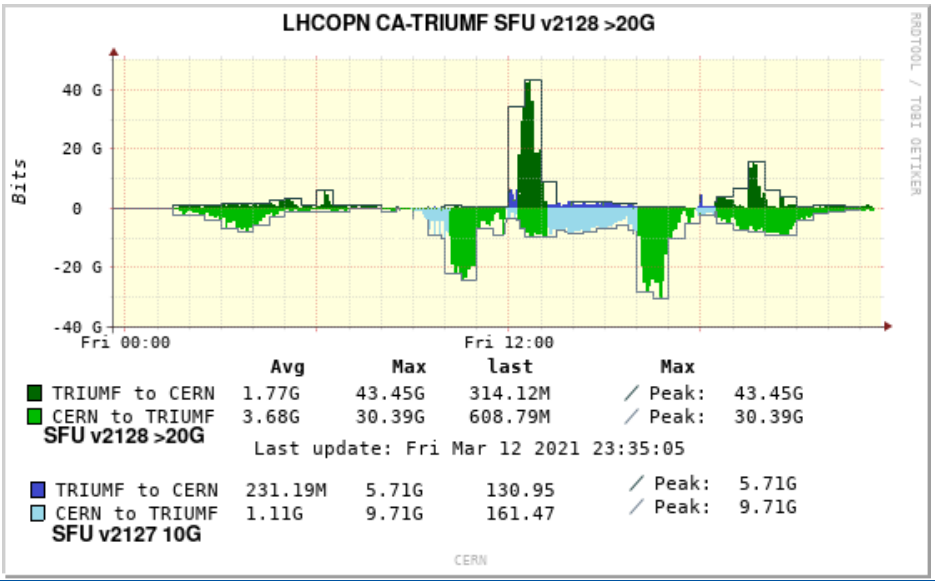
LHCOPN CA-TRIUMF SFU v2128 >20G
Bits
40 G
20 G
0
-20 G
-40 G
Fri 00:00
Fri 12:00
Avg Max last Max
TRIUMF to CERN 1.77G 43.45G 314.12M / Peak: 43.45G
CERN to TRIUMF 3.68G 30.39G 608.79M / Peak: 30.39G
SFU v2128 >20G
Last update: Fri Mar 12 2021 23:35:05
TRIUMF to CERN 231.19M 5.71G 130.95 / Peak: 5.71G
CERN to TRIUMF 1.11G 9.71G 161.47 / Peak: 9.71G
SFU v2127 10G
CERN
RRDTOOL / TOBI OETIKER

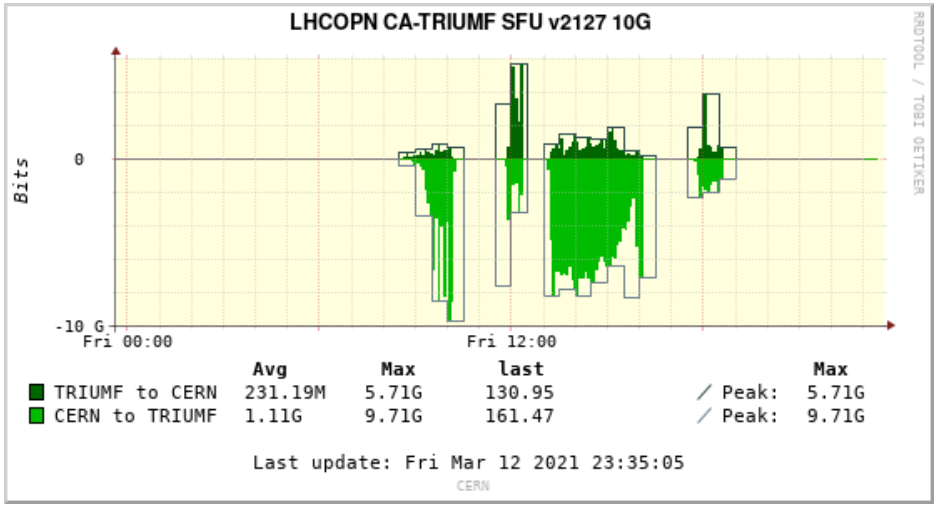

+

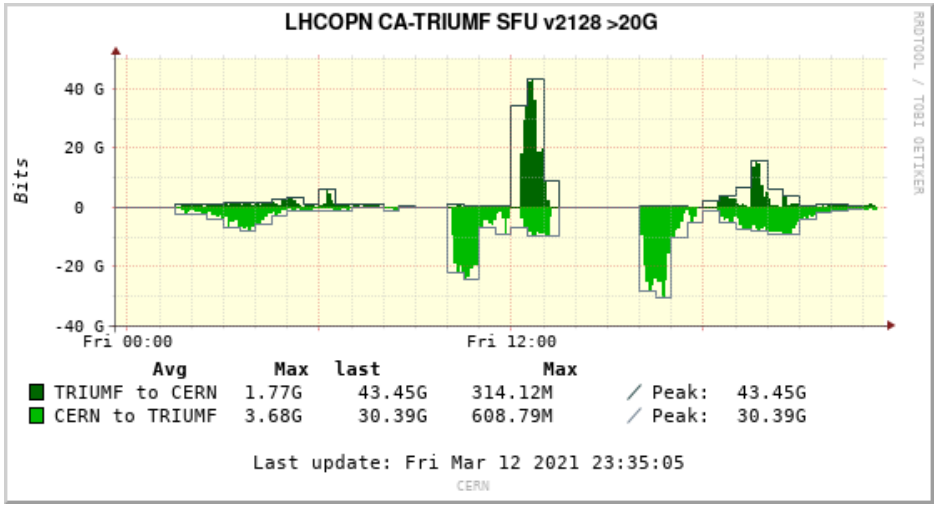

=

start test day

stop test day

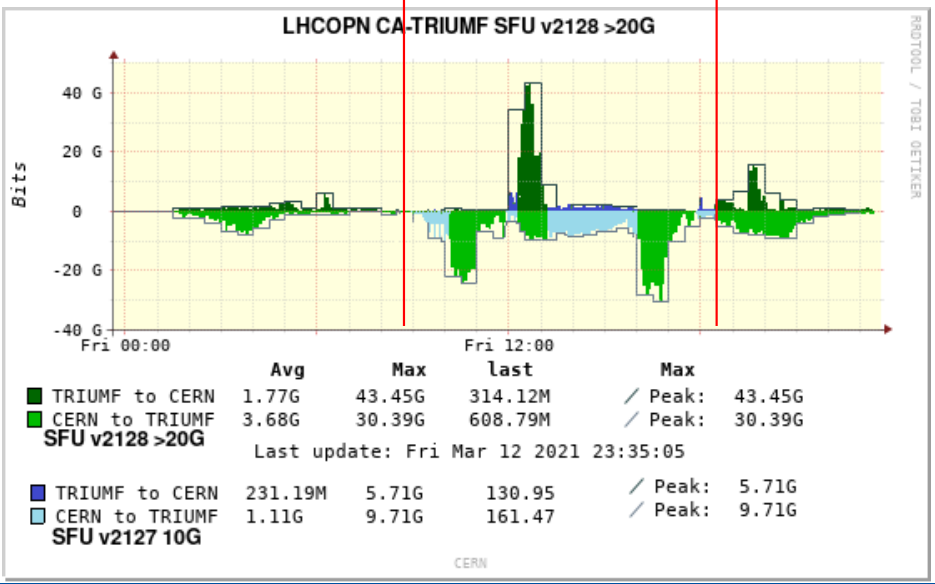

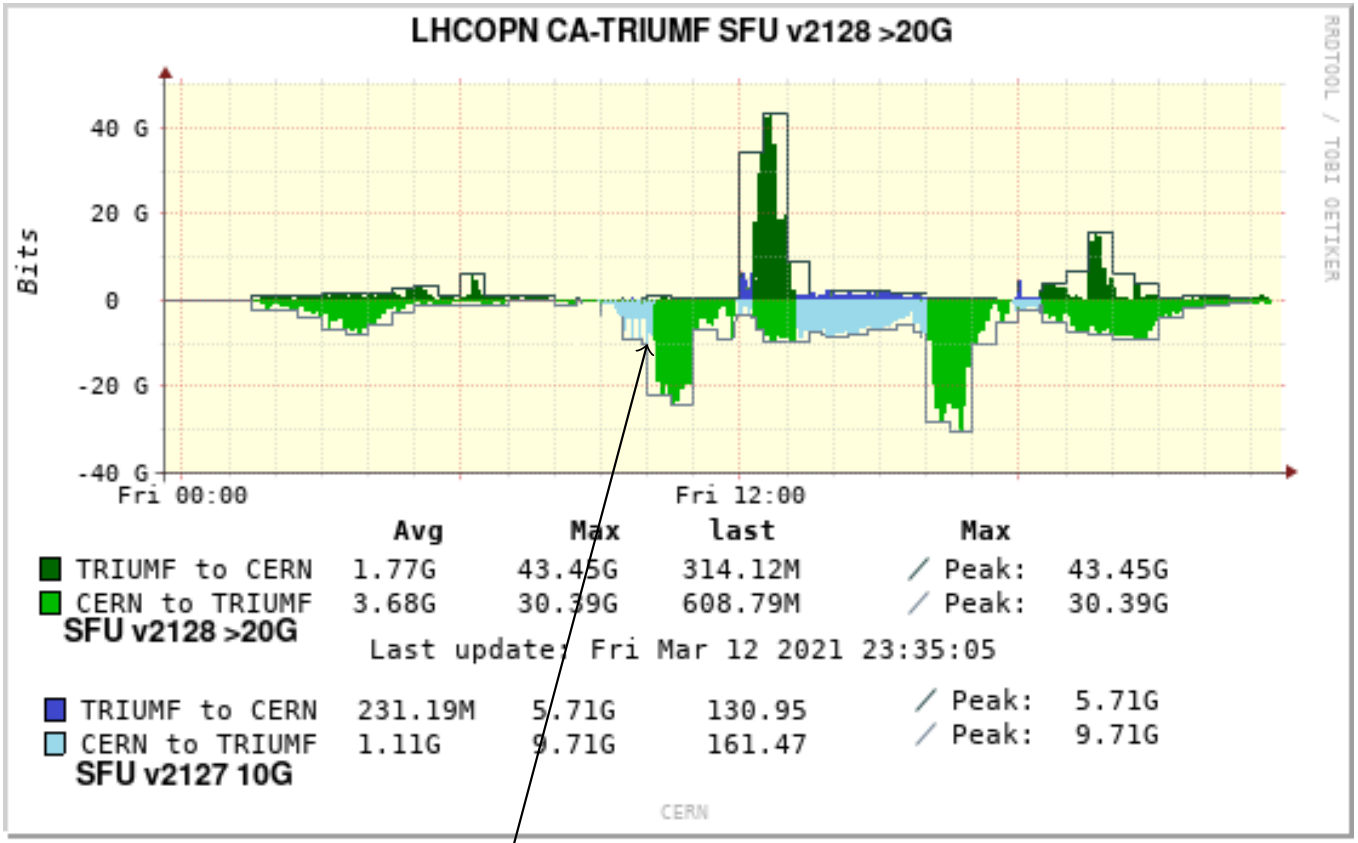

1st test transfer from CERN -> TRIUMF-SFU was detected; tool decided reconfigured network, and changed link SFU v2127 (10Gbps) to SFU v2128 (>50Gbps)

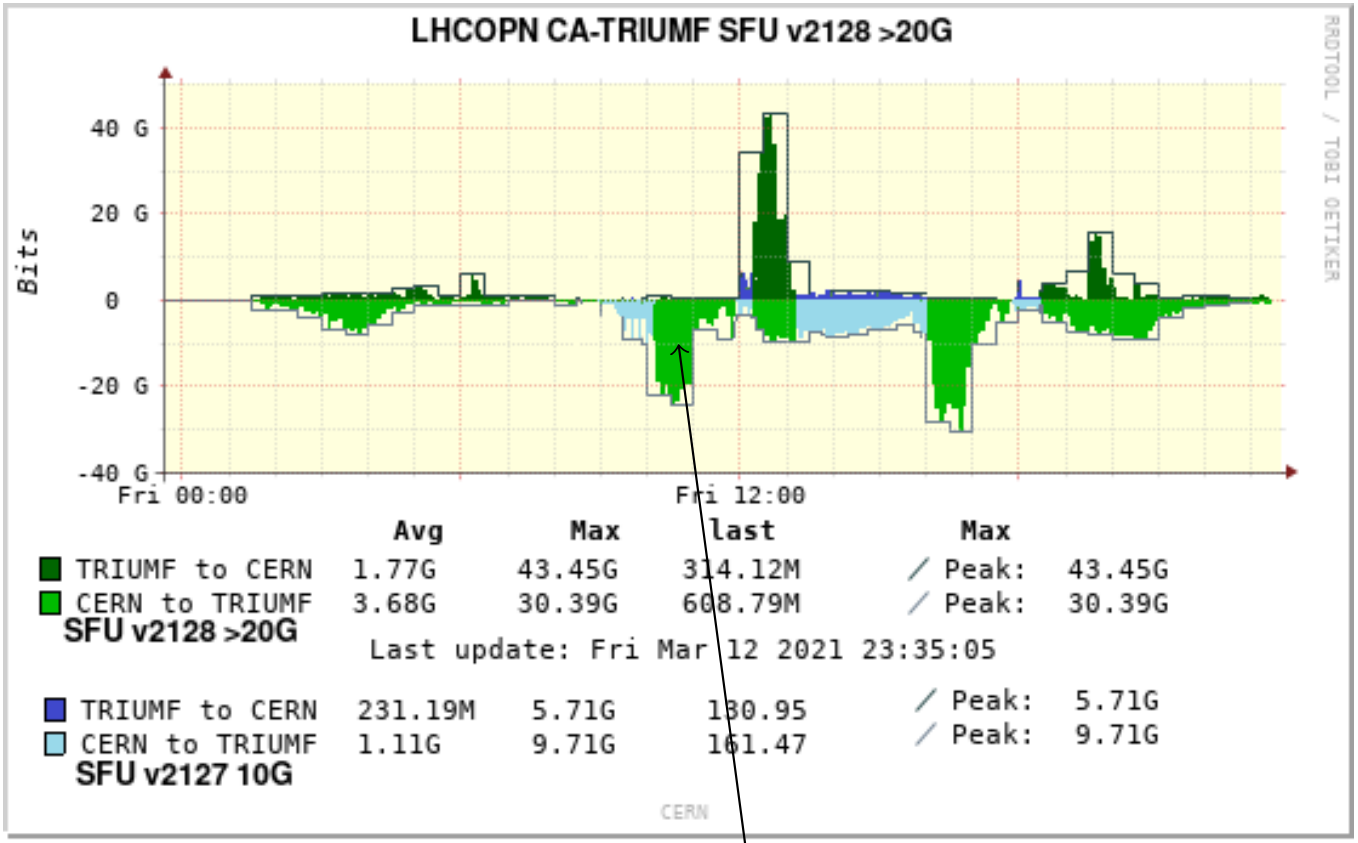

The FTS optimiser allowed to increase throughput and in the result we could generate traffic around 20Gbps

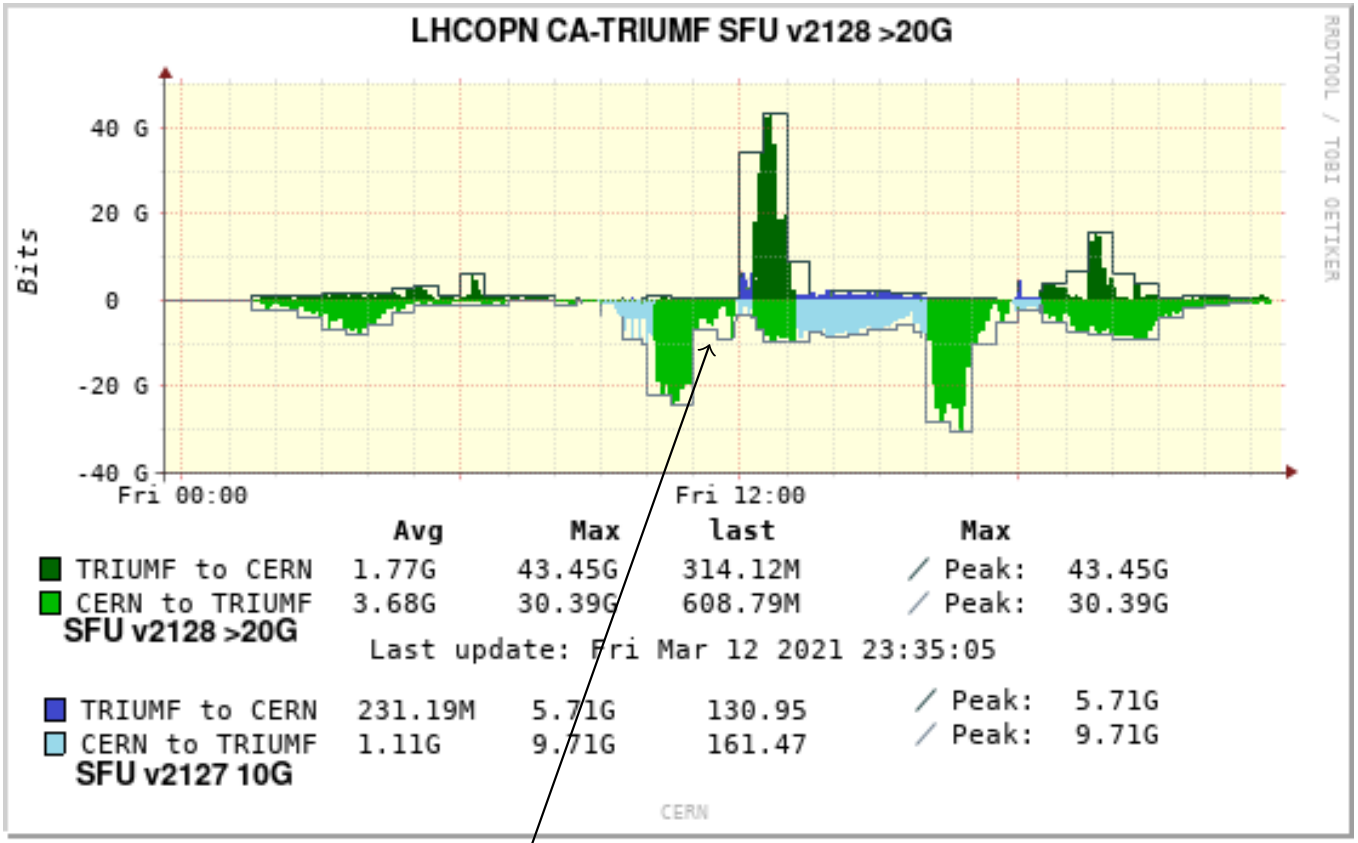

At the end of our test transfer, the tool detected a second huge transfer, so the reconfiguration decision was postponed for a certain amount of time. After couple of minutes the v2128 link was changed to the v2127

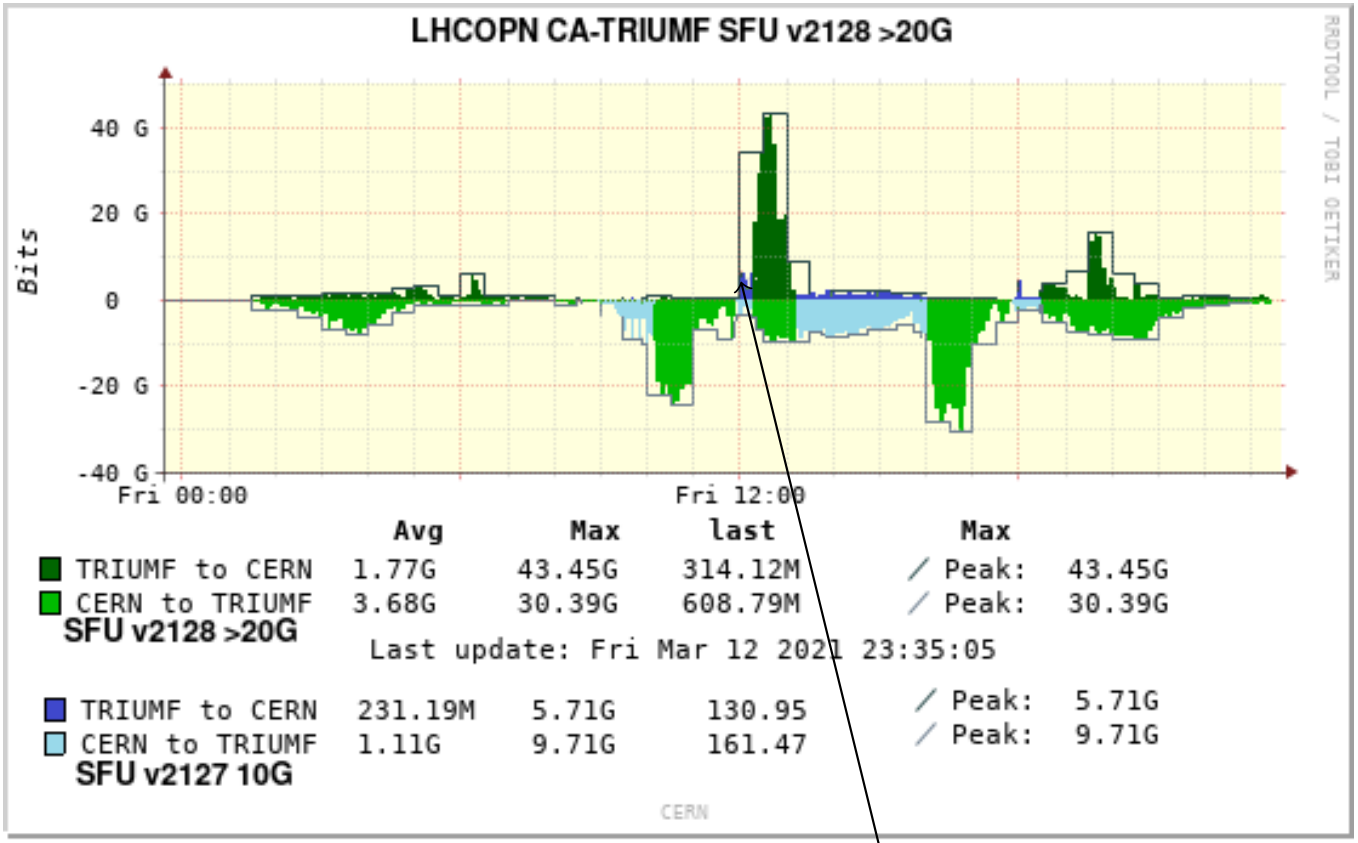

The 2st test transfer from TRIUMF-SFU -> CERN was detected;

tool decided reconfigured network, and changed link SFU v2127

(10Gbps) to SFU v2128 (>50Gbps)

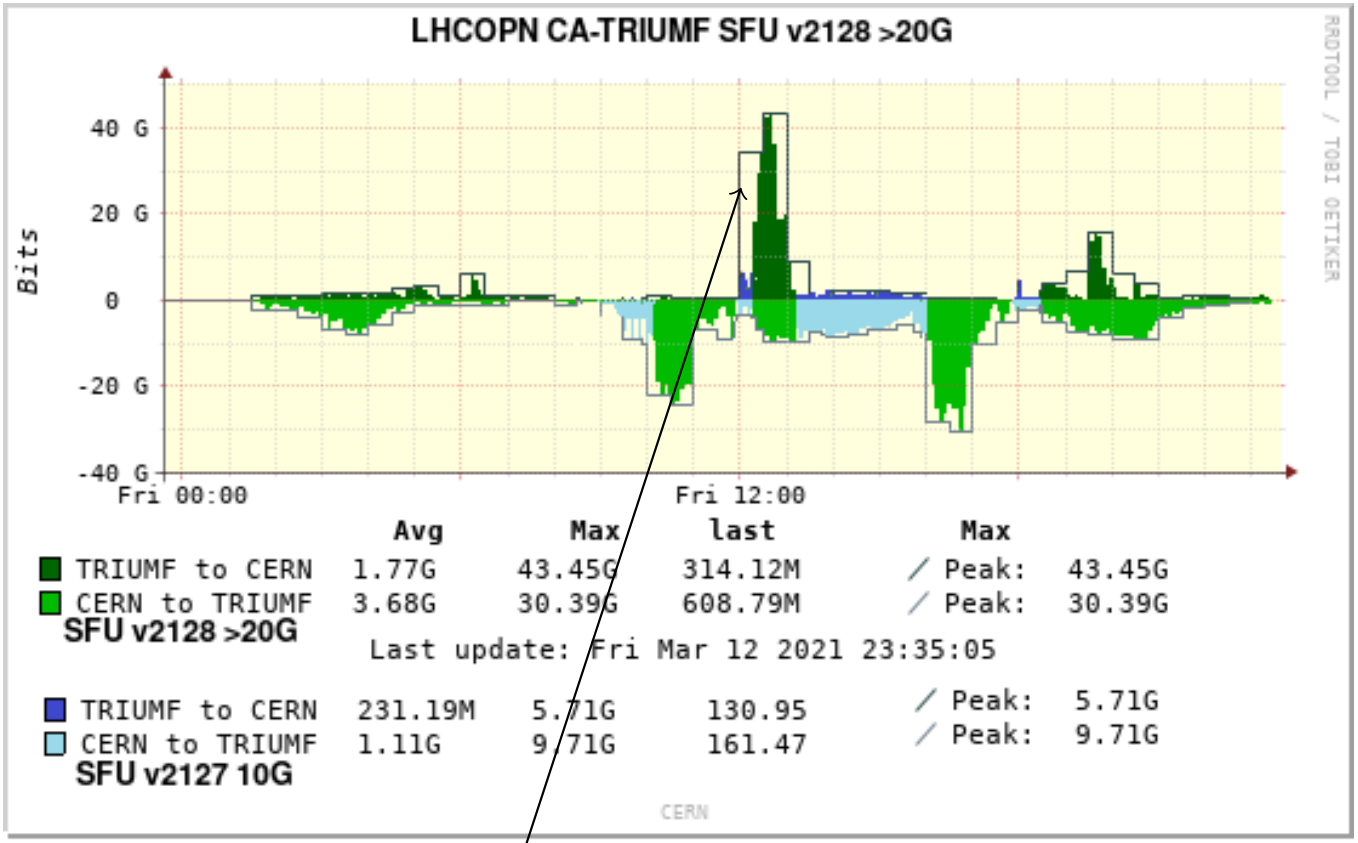

The FTS optimiser allowed increased throughput and in the result we could generate traffic around 40Gbps!

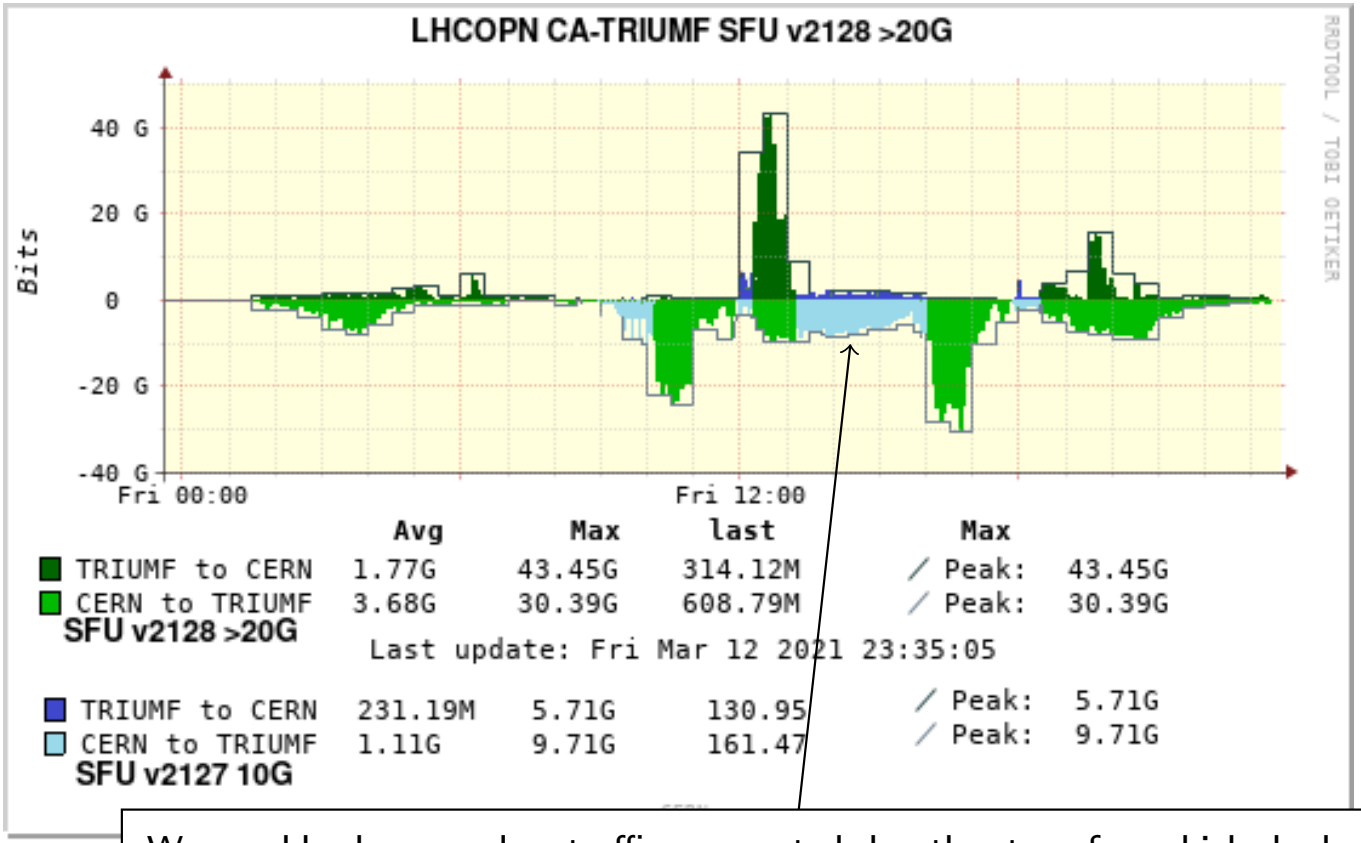

We could observe also traffic generated by the transfer which had been detected during 1st test transfer for few hours. BUT! FTS didn't generated traffic more then 10Gbps, so when our test transfer (TRIUMF-SFU -> CERN) finished – tool reconfigured links and we used v2127 again.

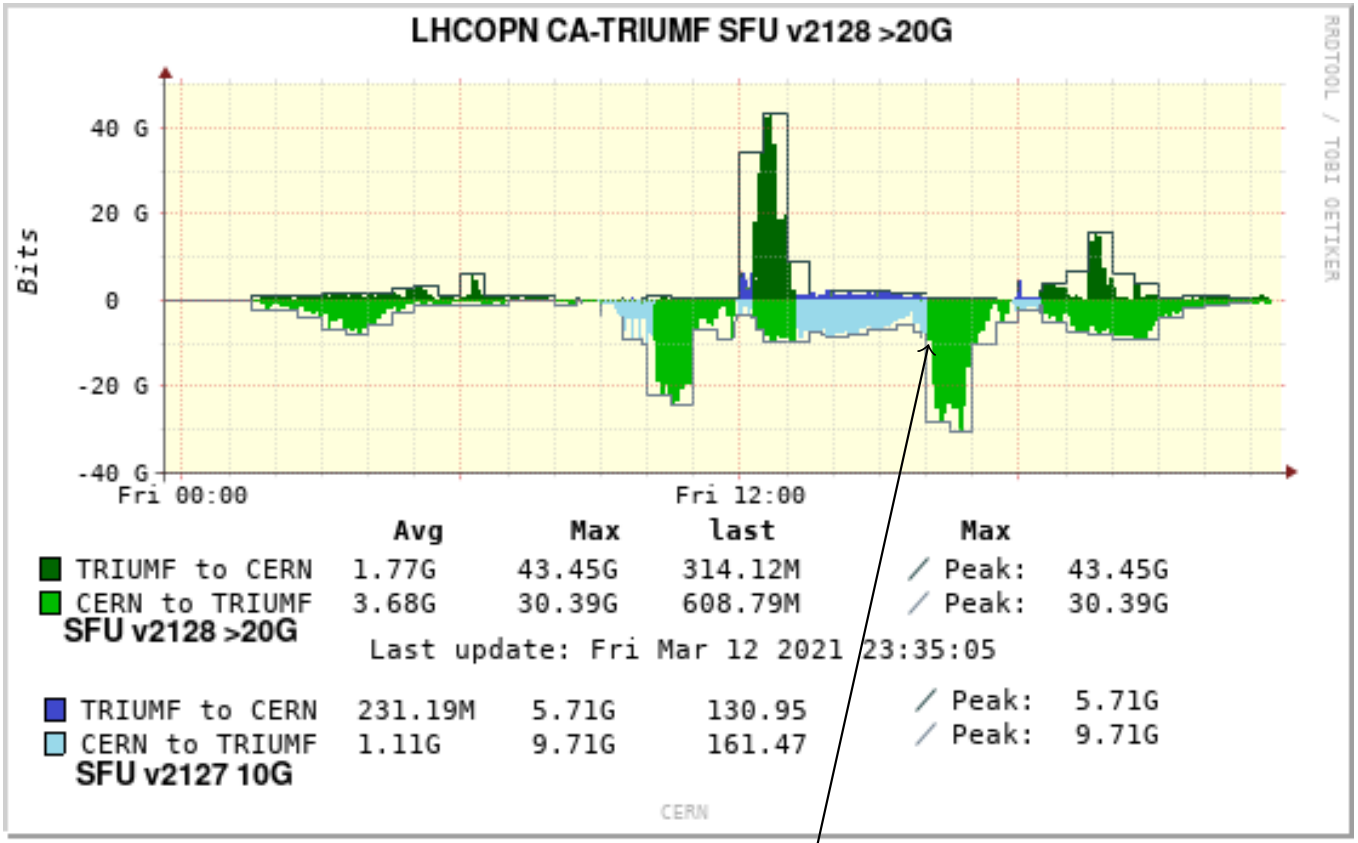

Last 3rd test transfer from CERN -> TRIUMF-SFU was detected; link was changed into v2128; FTS optimizer allowed send transfer faster and we could observed traffic around 30Gbps and after transfer finished tool returned v2127.

Section 3

Tool developing = Traffic forecasting

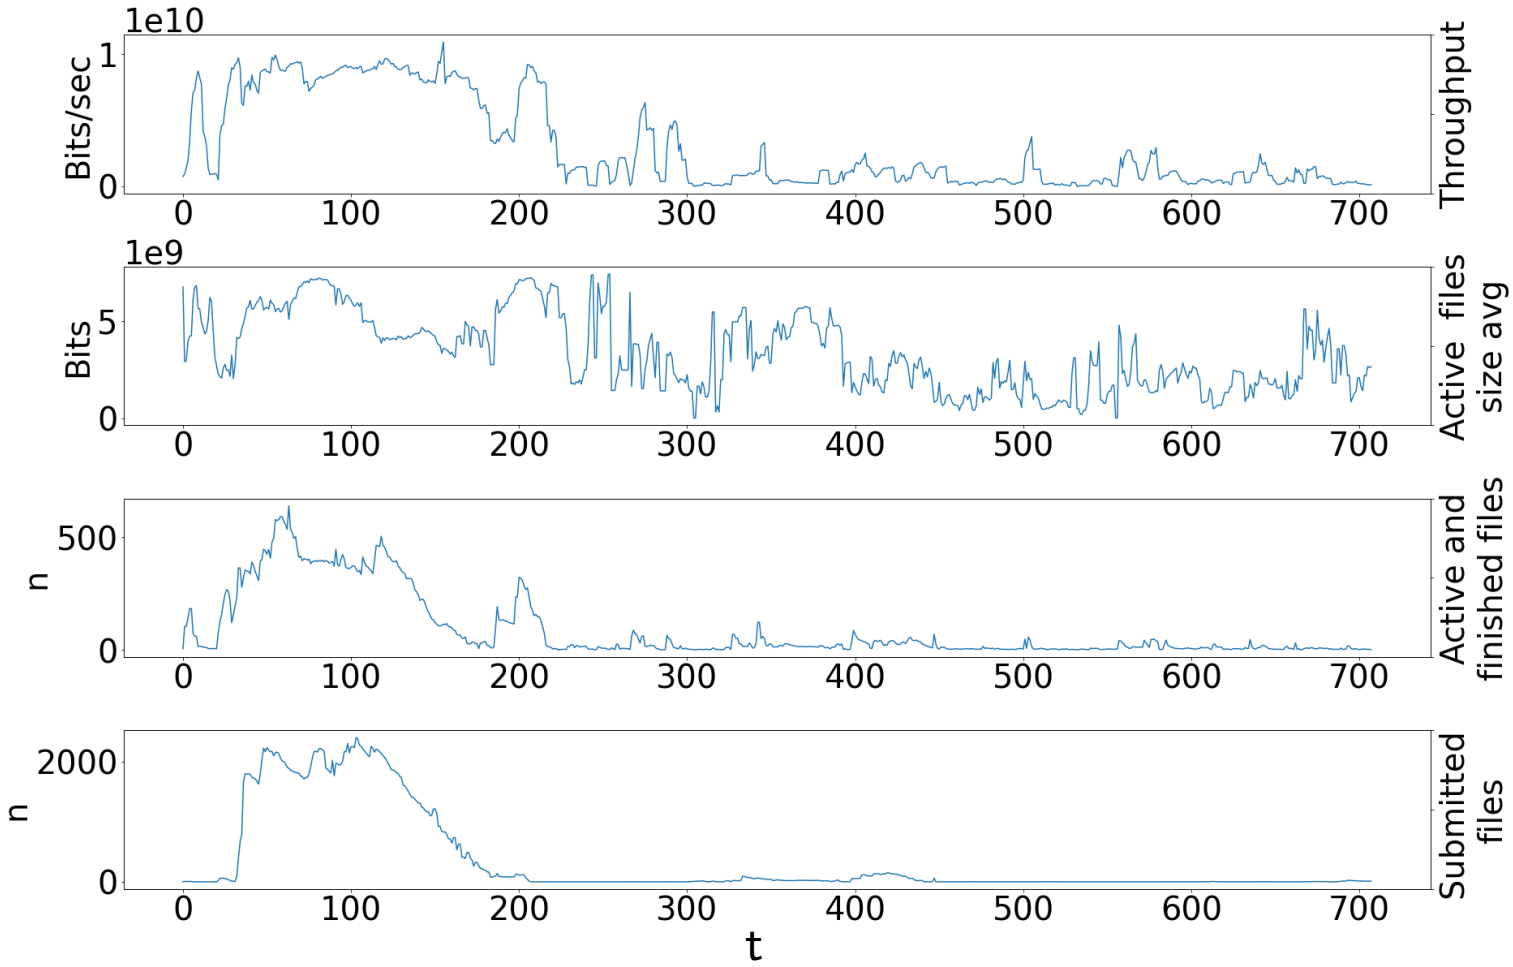

*Figure presents data used to as test dataset during modeling

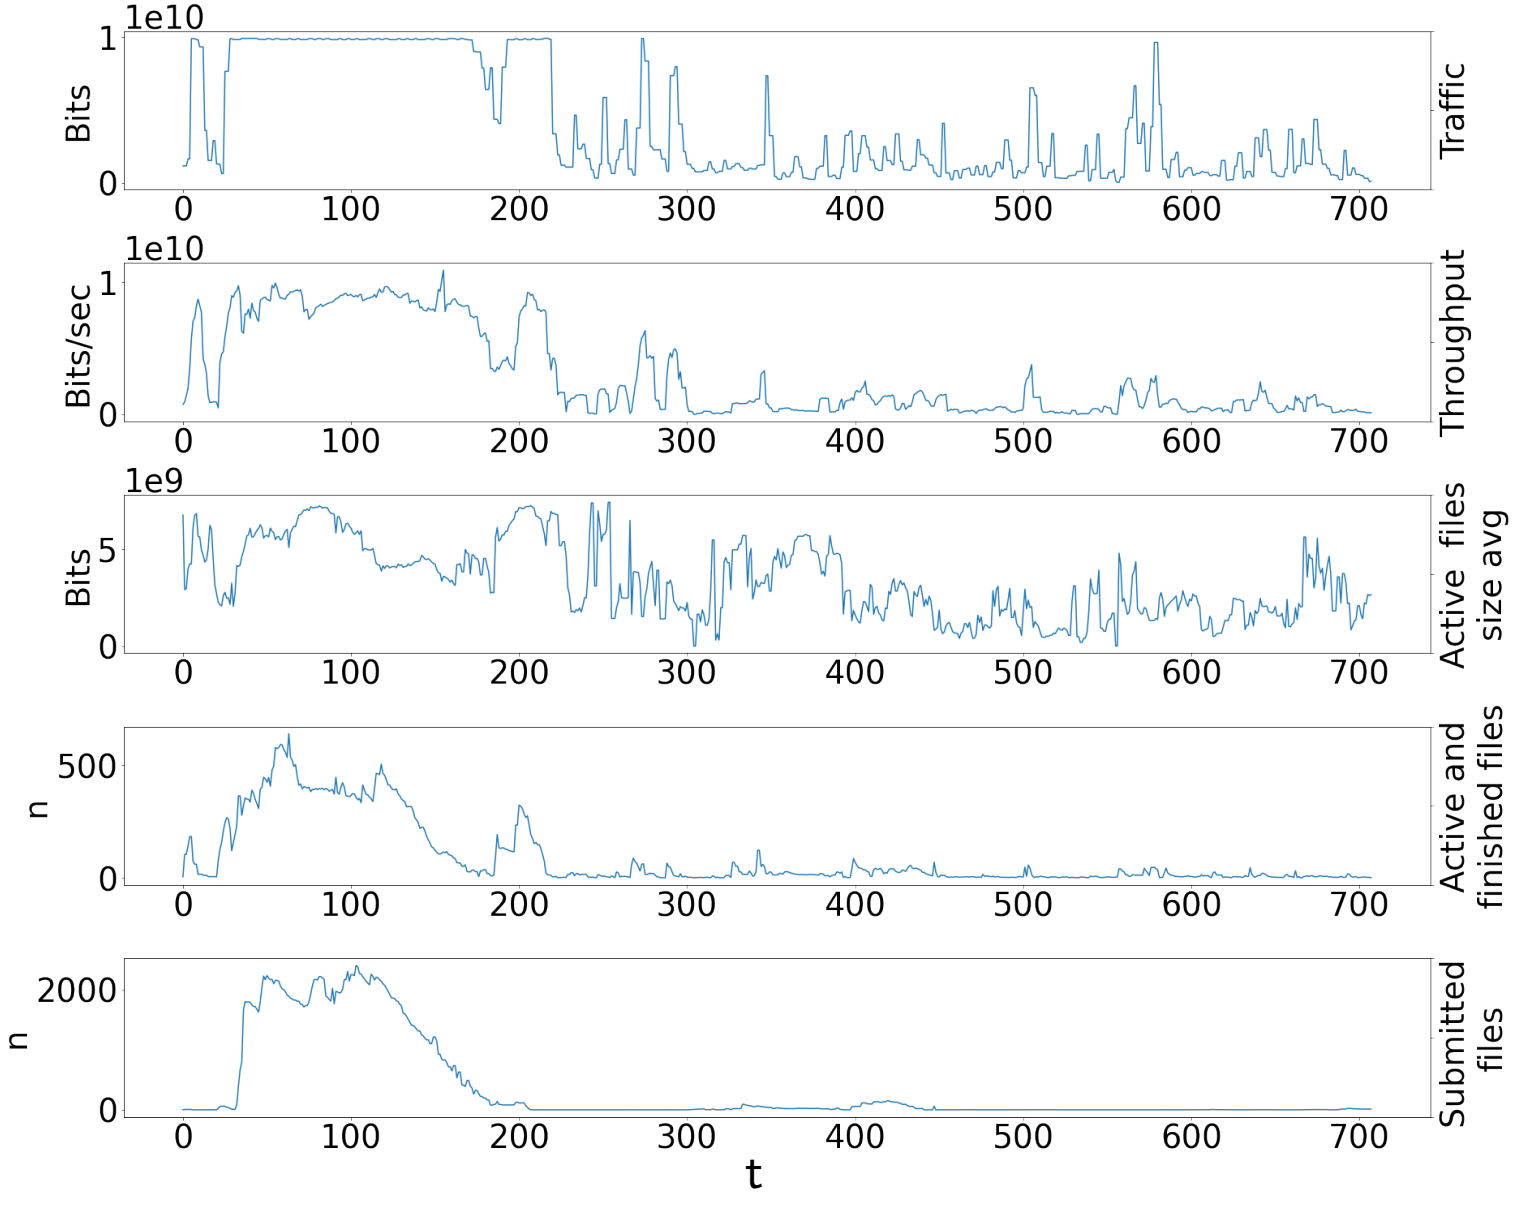
1e10
Bits
Traffic
1e10
Bits/sec
Throughput
1e9
Bits
Active files size avg
n
Active and finished files
n
Submitted files
t
NetStat
FTS

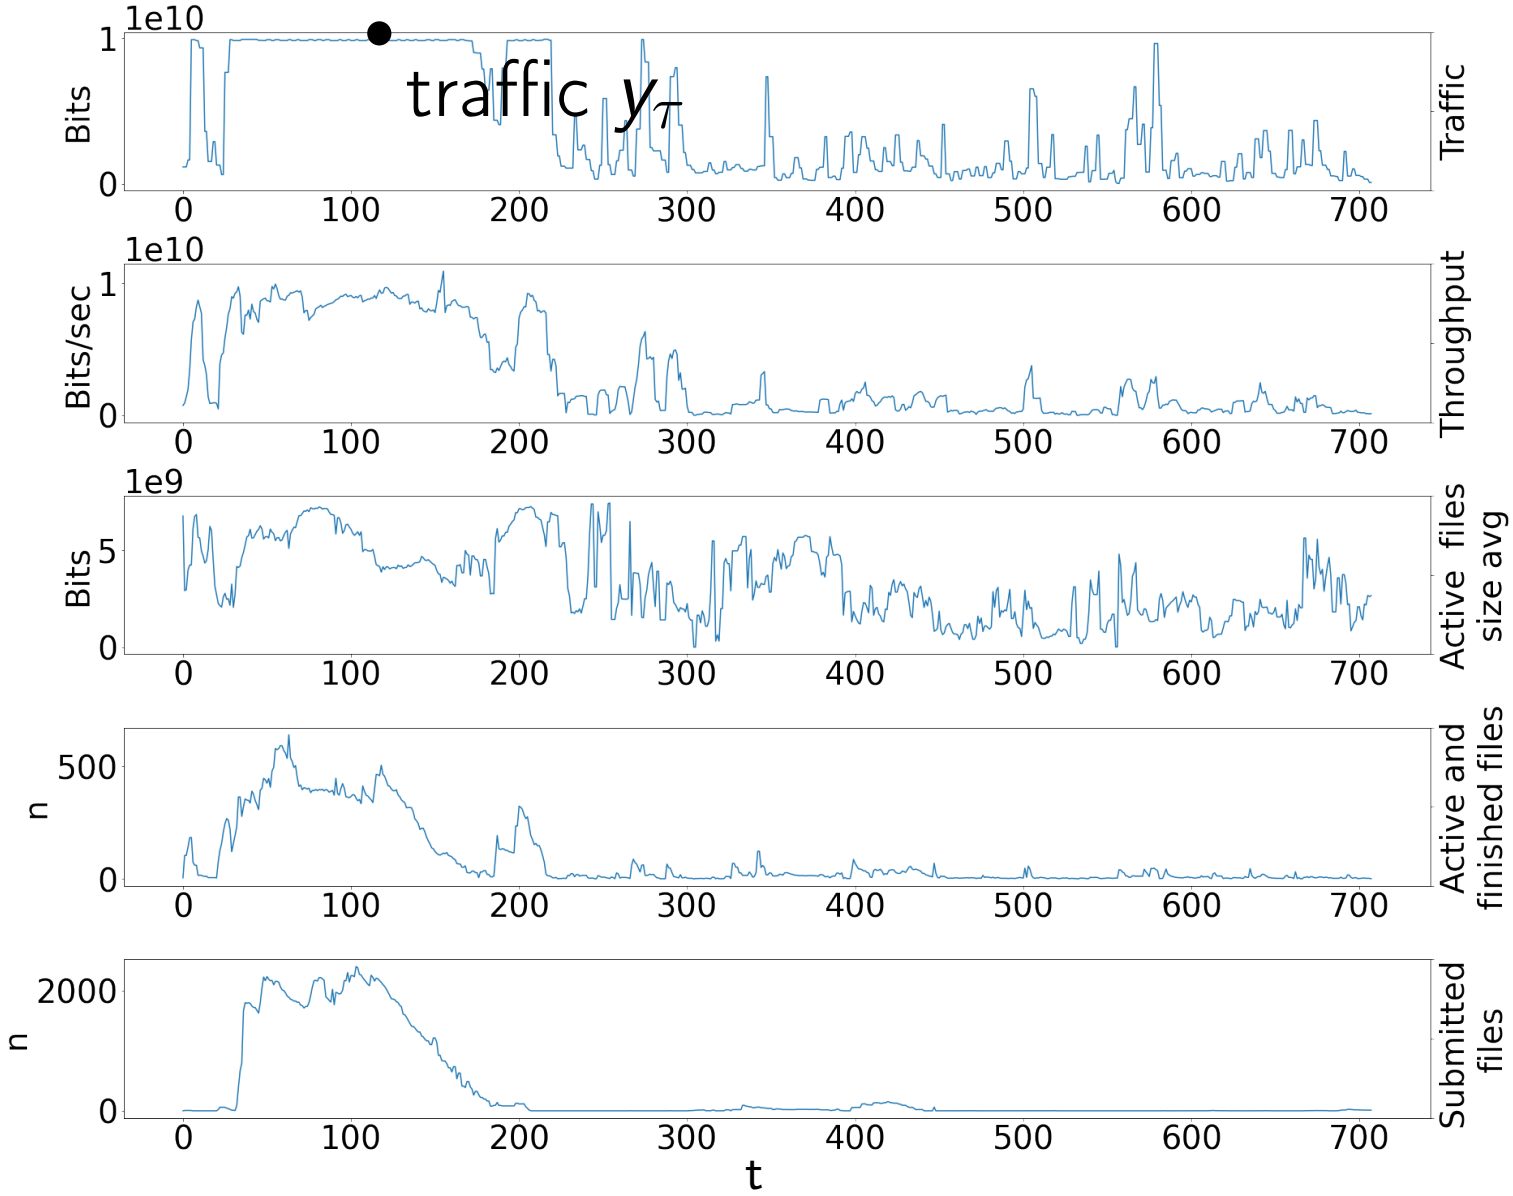
traffic y_τ
1e10
Bits
Traffic
1e10
Bits/sec
Throughput
1e9
Bits
Active files size avg
n
Active and finished files
n
Submitted files
t
0
100
200
300
400
500
600
700

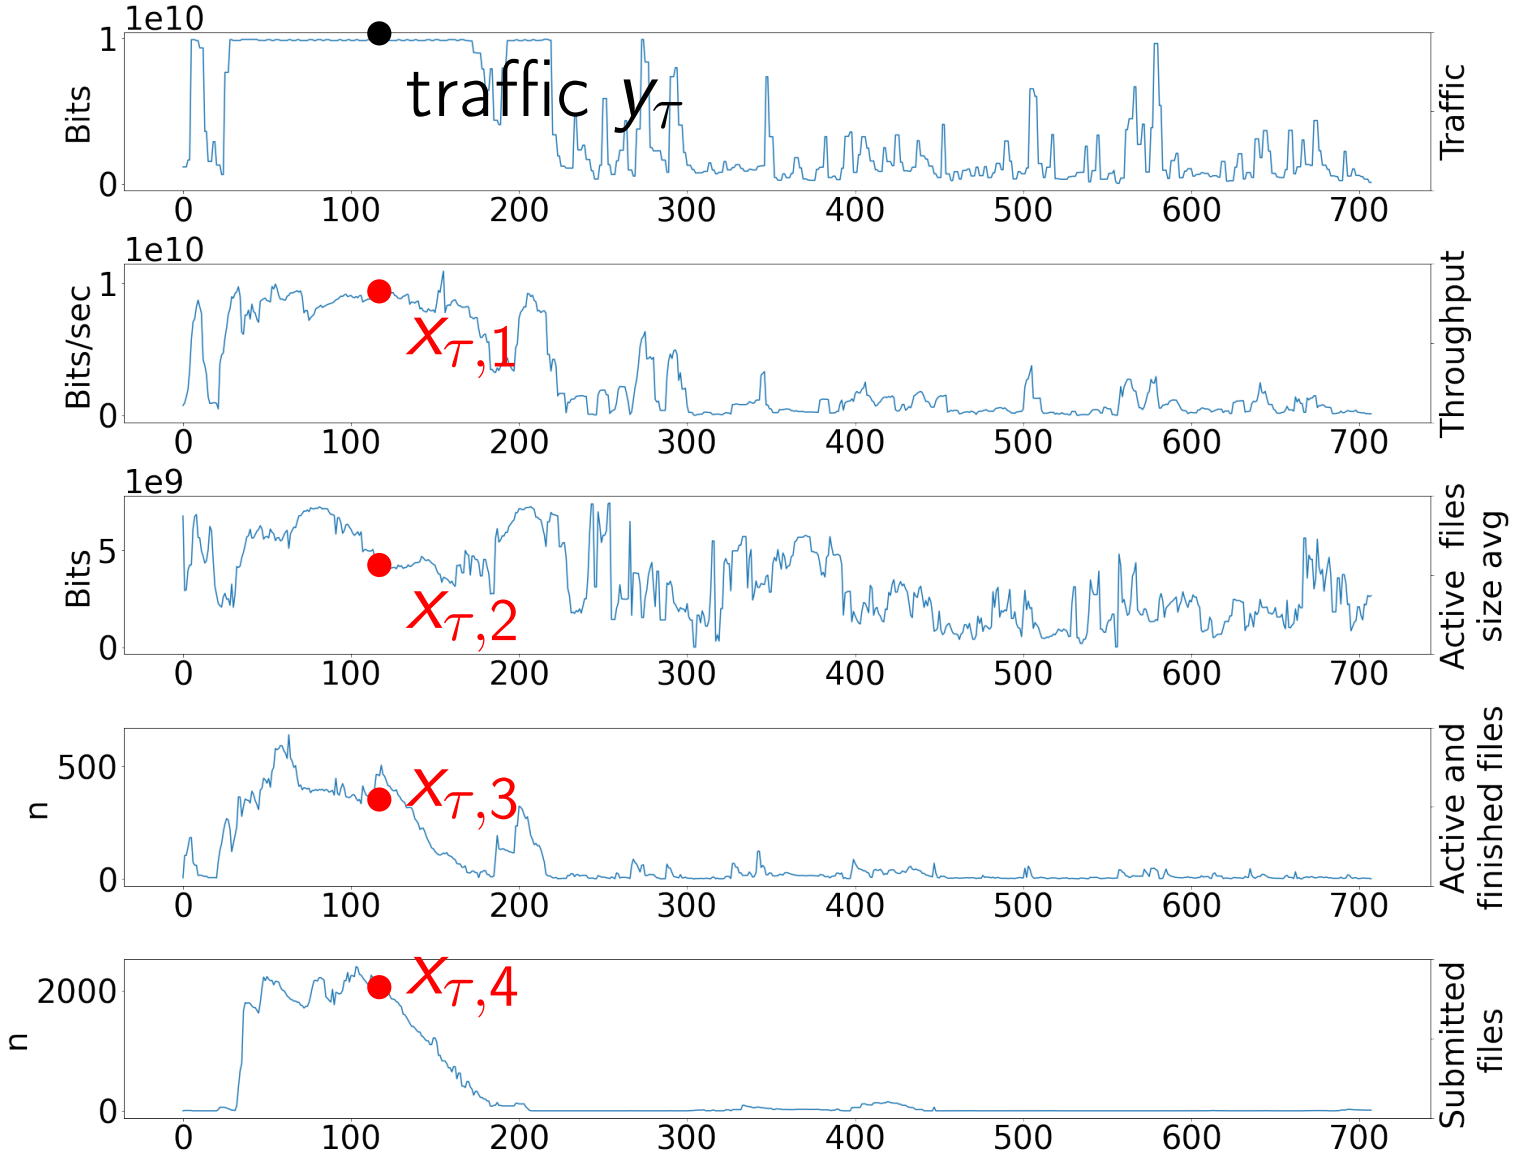

+ link capacity (as limit)

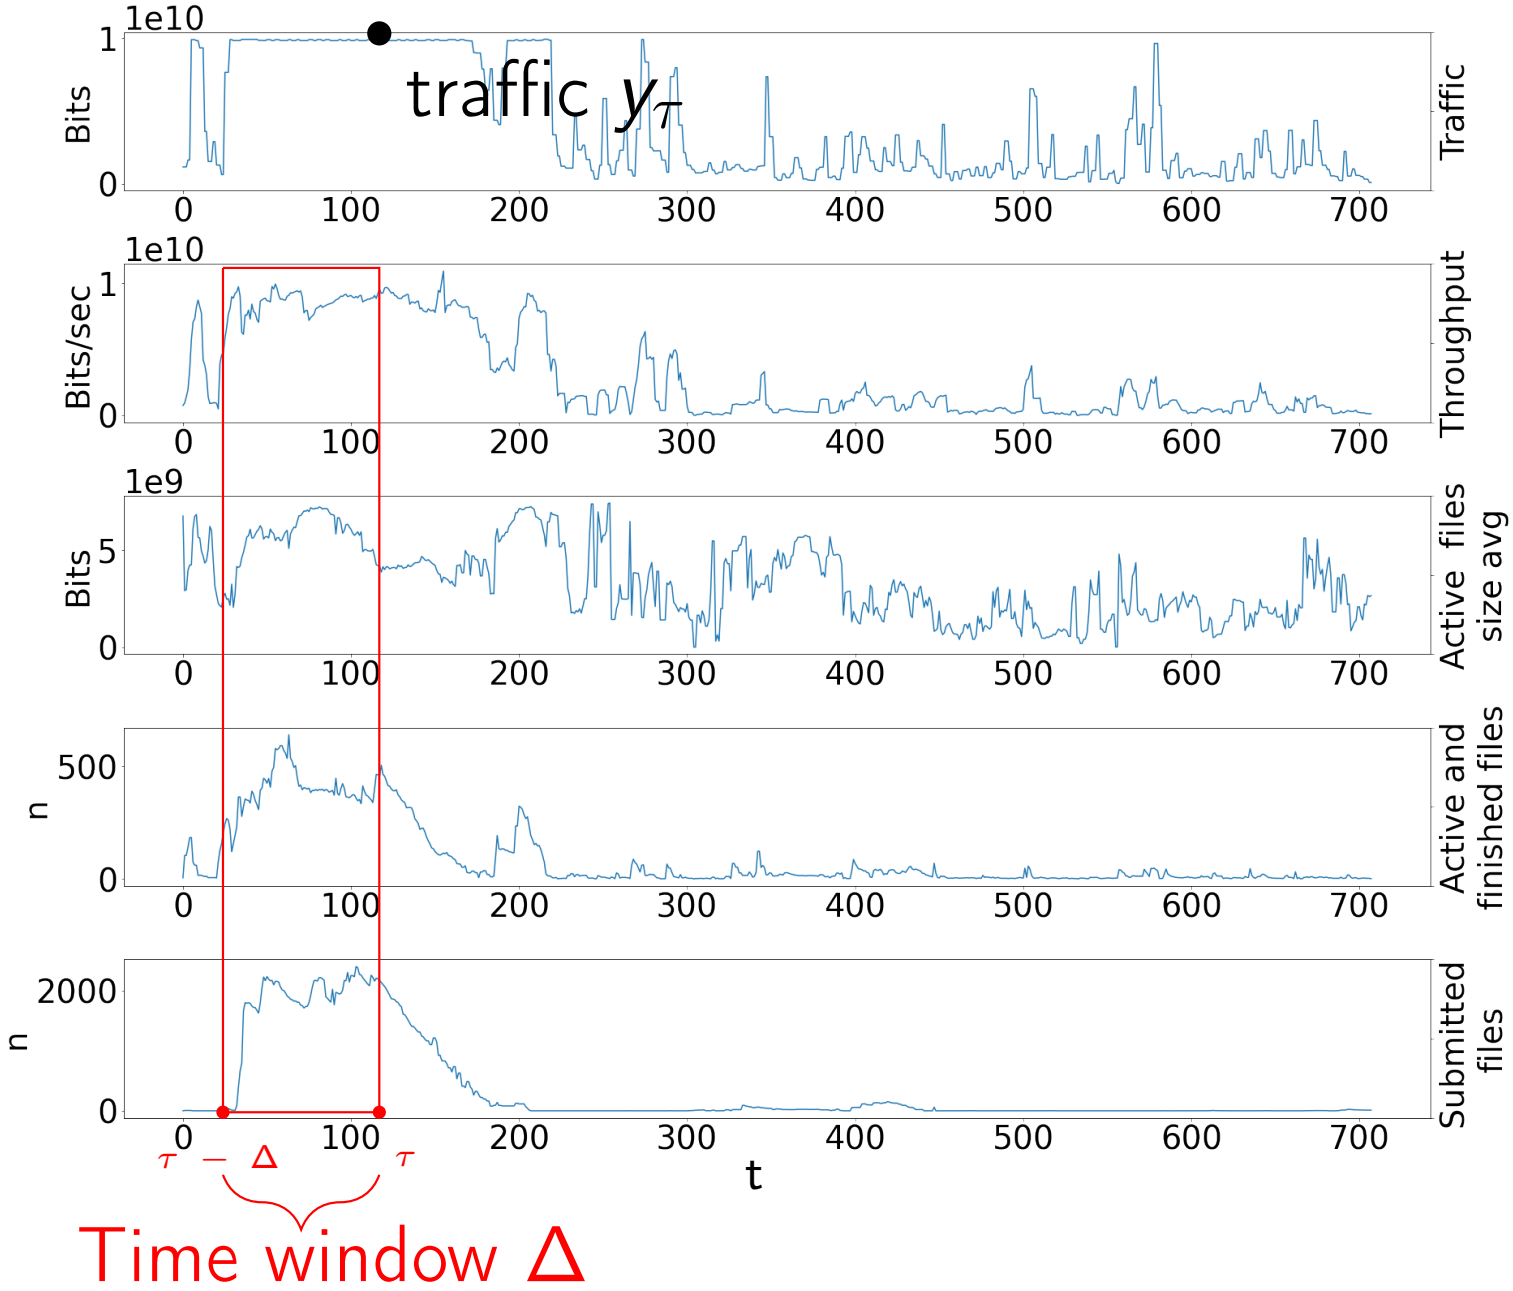
traffic y_τ
1e10
Bits
Traffic
Bits/sec
Throughput
1e9
Active files size avg
n
Active and finished files
Submitted files
t
$\tau - \Delta$
τ
Time window Δ

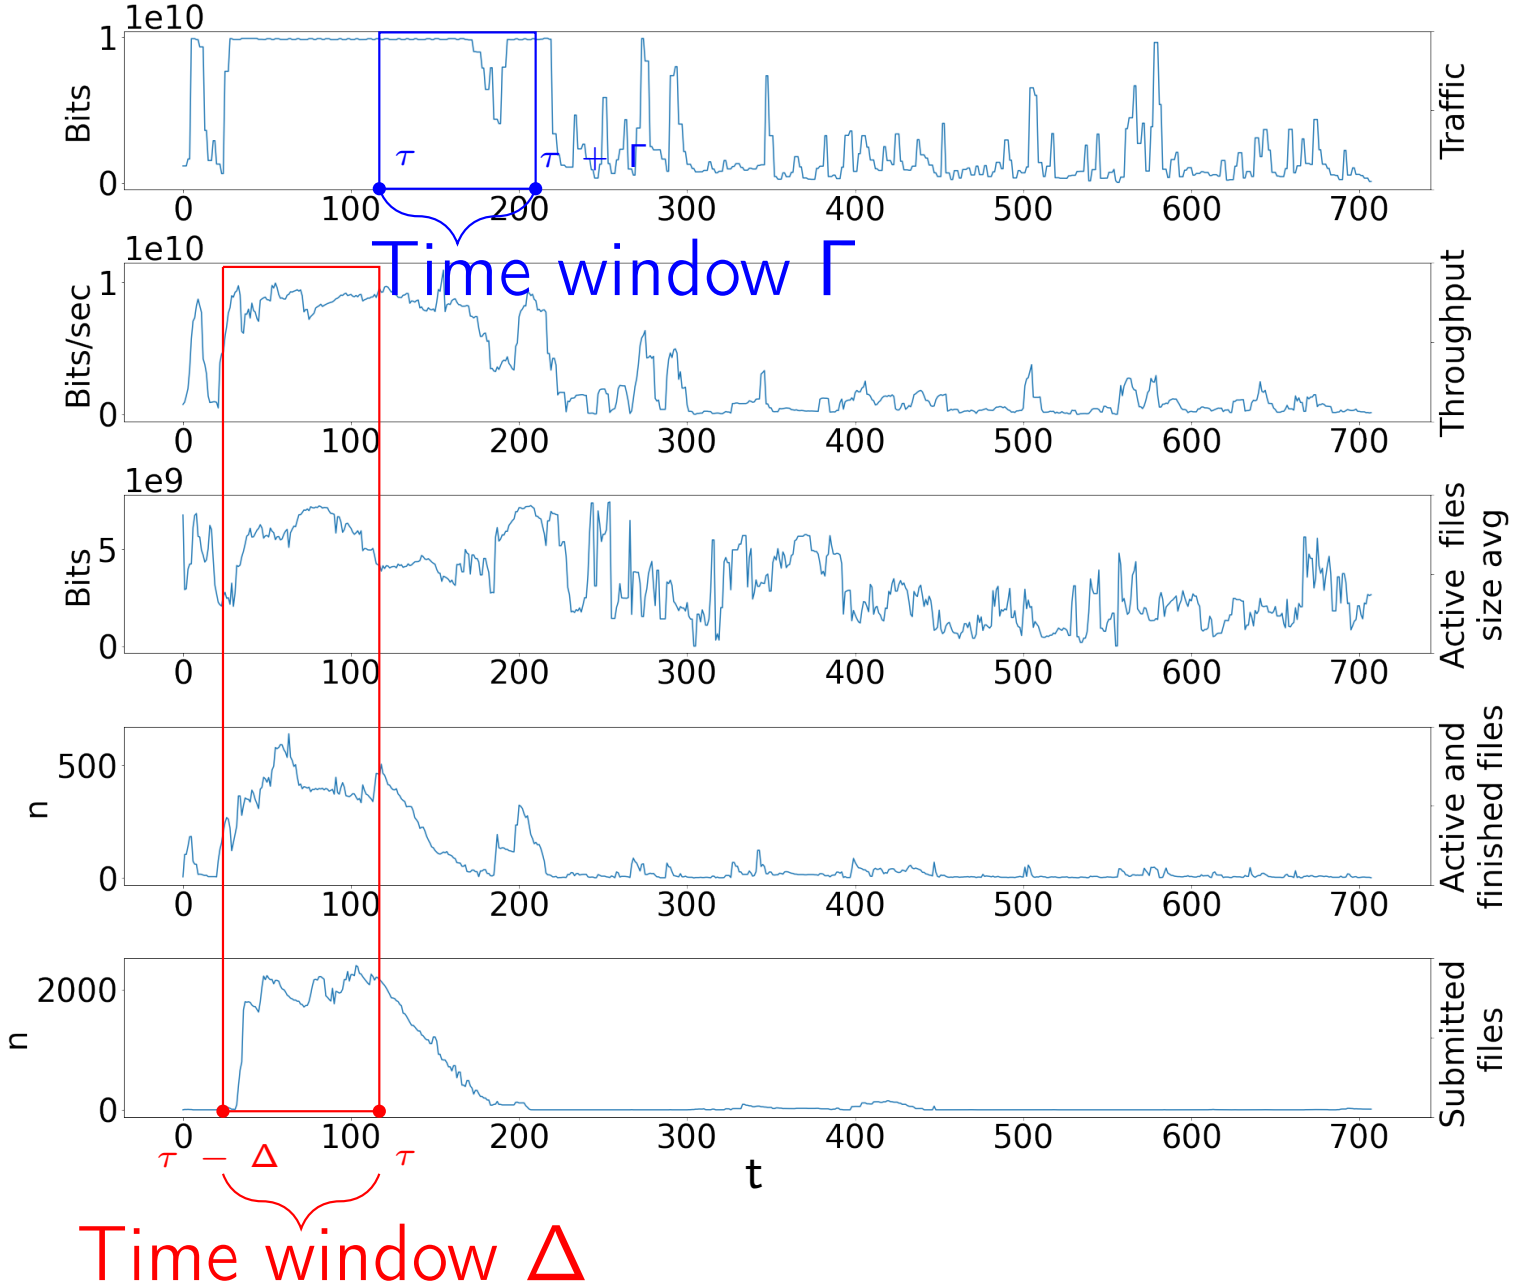
1e10
Bits
Traffic
τ
$\tau + \Gamma$
Time window Γ
1e10
Bits/sec
Throughput
1e9
Bits
Active files size avg
n
Active and finished files
n
Submitted files
t
$\tau - \Delta$
τ
Time window Δ

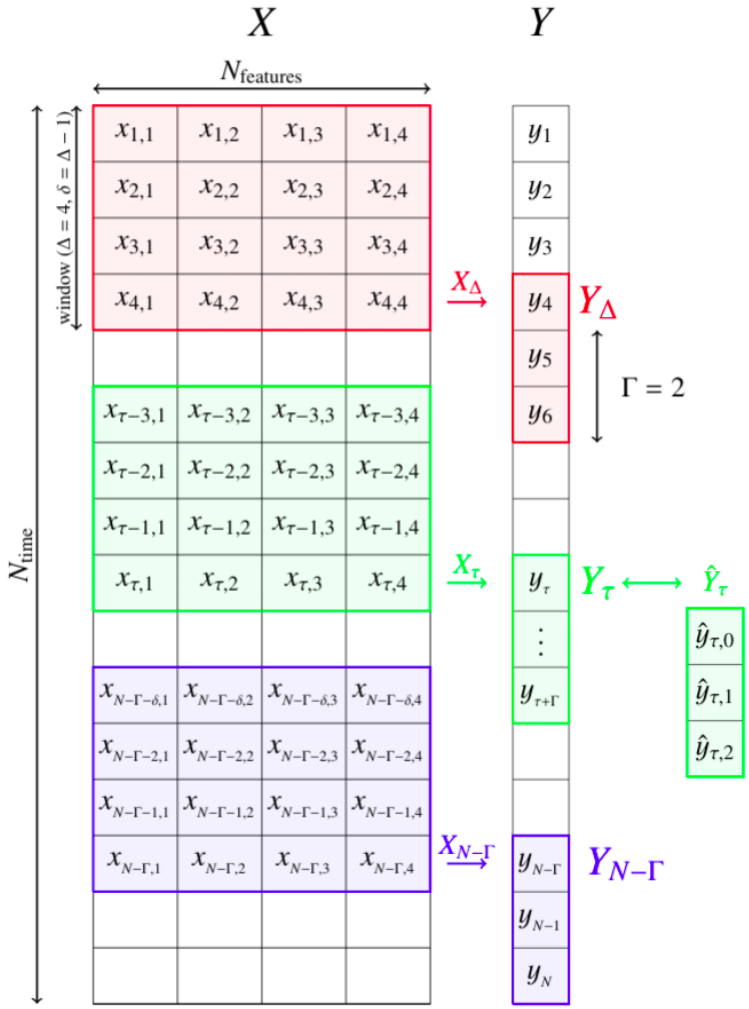

Figure: Schema input data set X, and dependent variables Y

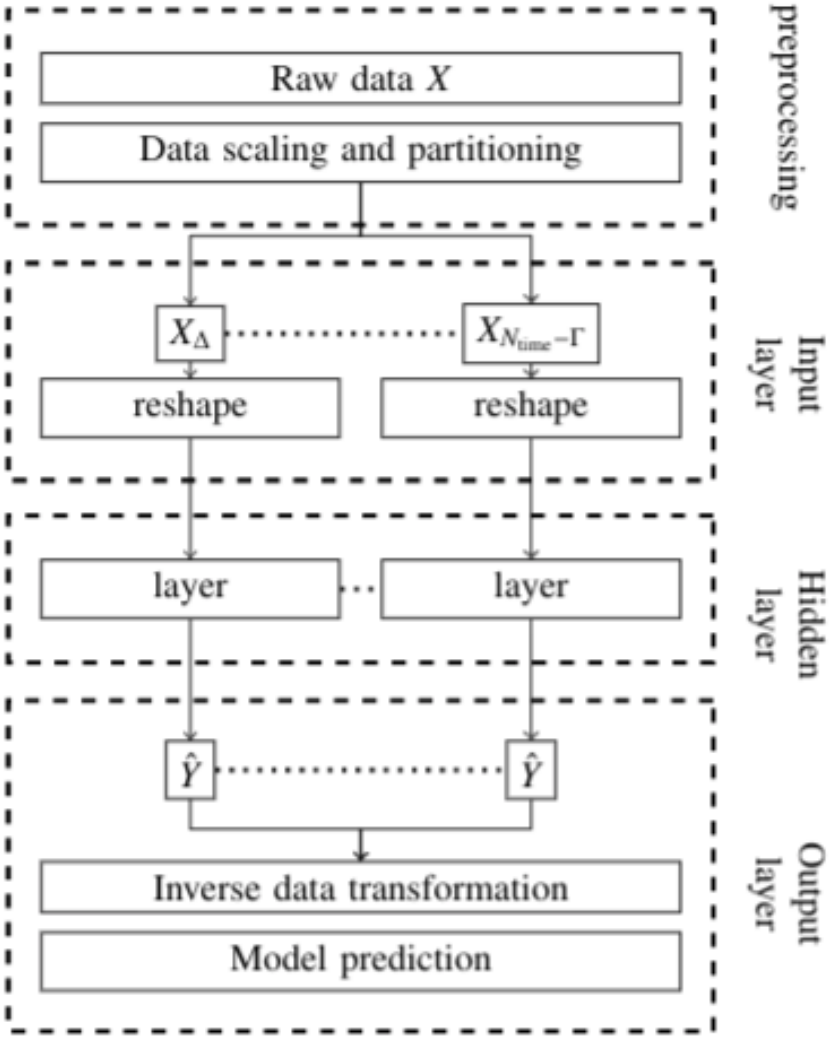

For the Conv-LSTM model we consider additional pre-processing steps and calculate the throughput exponential moving average over the last 15 minutes, and estimated size values

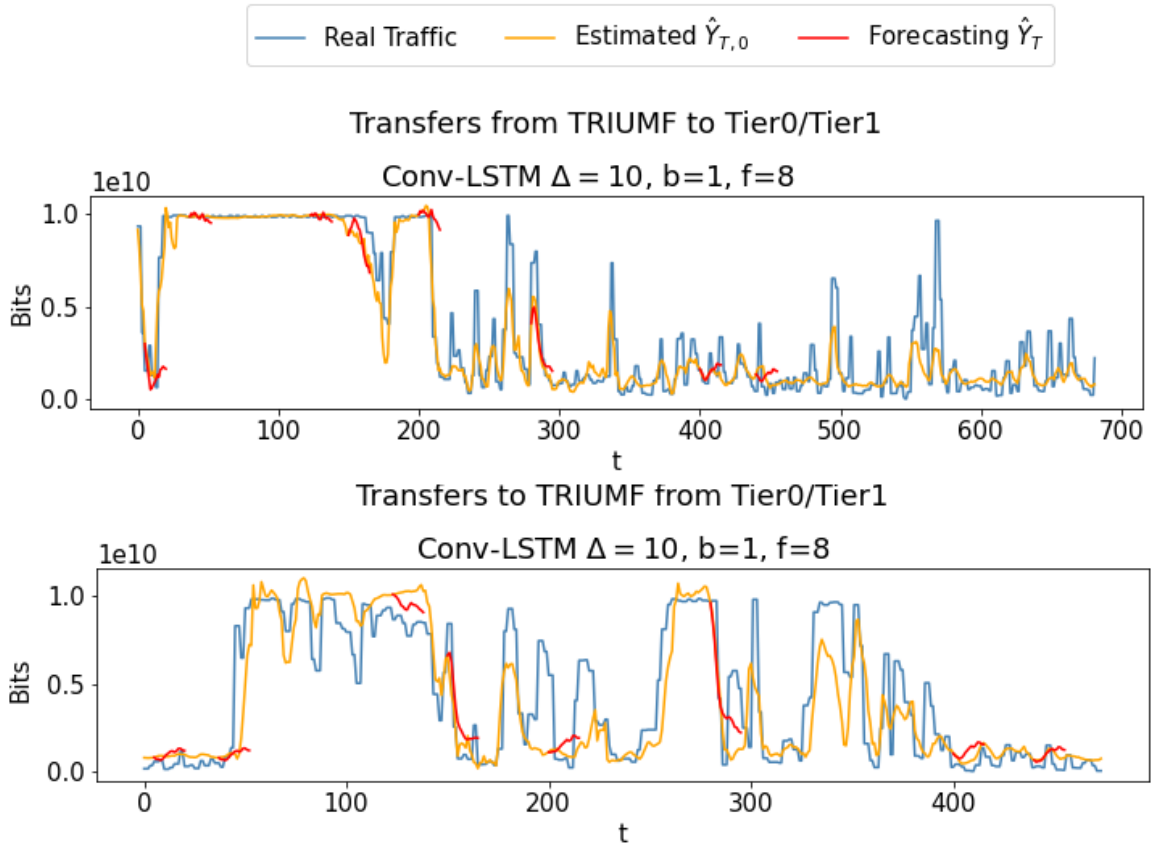

Figure: 1. We trained model to predict traffic based on information about transfers from FTS (from TRIUMF-SFU to Tier0/Tier1). 2. Pictures present results on two data sets: from TRIUMF-SFU to Tier0/Tier1 and from Tier0/Tier1 to TRIUMF-SFU. Forecasting is based on aggregated information about transfers from last 20 minutes ($\Delta = 10$); b (batches); f(filters). Forecasting $\hat{Y}_\Gamma$ is predicted Γ future values for chosen t. Here $\Gamma = 15$ (30 min).

Section 4

Conclusion

Conclusion

- We can understand network traffic via the analysis of data from FTS (!)

Conclusion

- We can understand network traffic via the analysis of data from FTS (!)
- We can classify transfers and display those that will generate /are generating large network traffic

Conclusion

- We can understand network traffic via the analysis of data from FTS (!)
- We can classify transfers and display those that will generate /are generating large network traffic
- grouping endpoints by sites (prefixes), aggregating sites according to the controller configurations allows for fully conscious automation of decision-making regarding network reconfiguration.

Conclusion

- We can understand network traffic via the analysis of data from FTS (!)
- We can classify transfers and display those that will generate /are generating large network traffic
- grouping endpoints by sites (prefixes), aggregating sites according to the controller configurations allows for fully conscious automation of decision-making regarding network reconfiguration.
- using Conv-LSTM model we are able to forecasting traffic regardless of the transfer behavior based on short history (time window).

Conclusion

- We can understand network traffic via the analysis of data from FTS (!)
- We can classify transfers and display those that will generate /are generating large network traffic
- grouping endpoints by sites (prefixes), aggregating sites according to the controller configurations allows for fully conscious automation of decision-making regarding network reconfiguration.
- using Conv-LSTM model we are able to forecasting traffic regardless of the transfer behavior based on short history (time window).

Thank you for your attention!

Contact:

Joanna Waczyńska
joanna.waczynska@cern.ch
joanna.waczynska@gmail.com
IT-CS-CE

Edoardo Martelli
edoardo.martelli@cern.ch
IT-CS-CE

Sofia Vallecorsa
Sofia.Vallecorsa@cern.ch
IT-DI-OPL

Edward Karavakis
edward.Karavakis@cern.ch
IT-ST-PDS


CERN

supplement

Why NOTED?

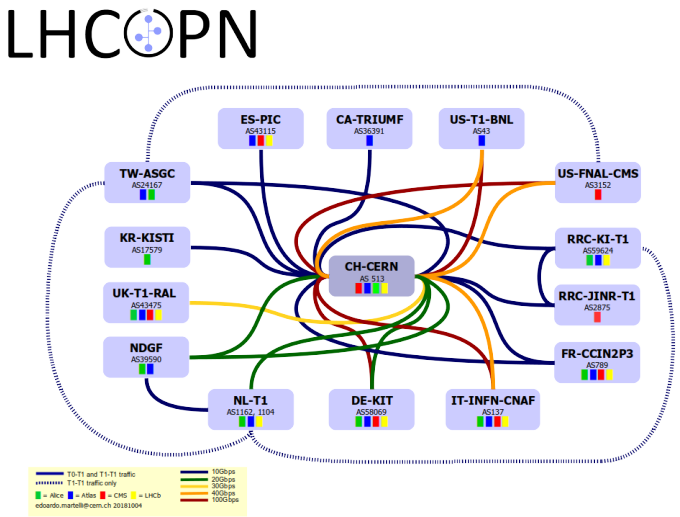

Figure: LHCOPN (Large Hadron Collider Optical Private Network) topology

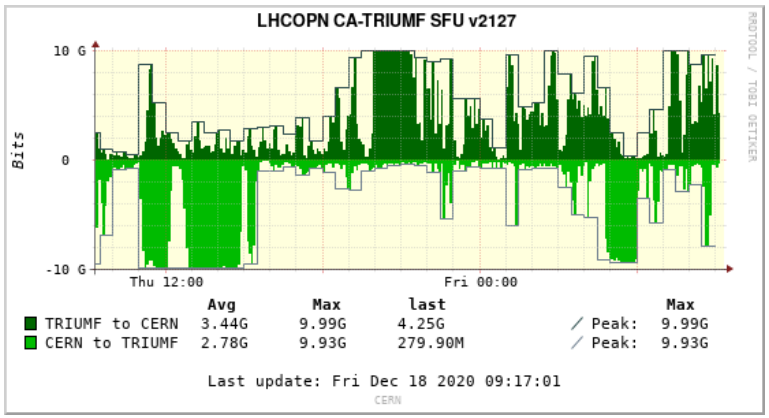

Figure: Network traffic observed on the LHCOPN path between CERN and TRIUMF. Link saturation occurs in both directions.

Solution - how avoid saturation

- Add extra path/link to balancing traffic.
- Reconfigure network, and in result move all traffic from one link to 2nd link.

How?

- Automatically decides when the link will be saturated for a long period of time. Manually modify the configuration of network devices

- Automatically decides when the link will be saturated for a long period of time, and automatically modify the configuration of network devices (SDNC).

FTS details - how transfer report look like

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

DESTINATION

SOURCE

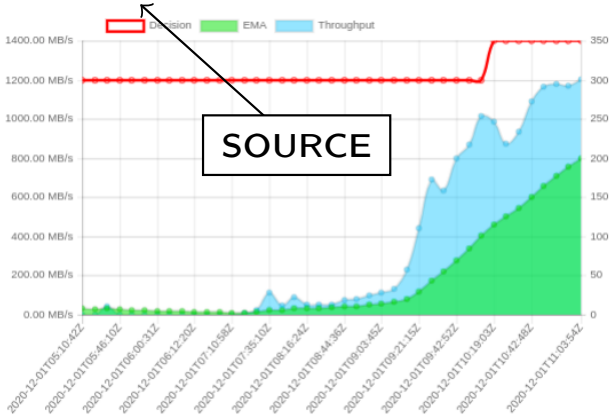

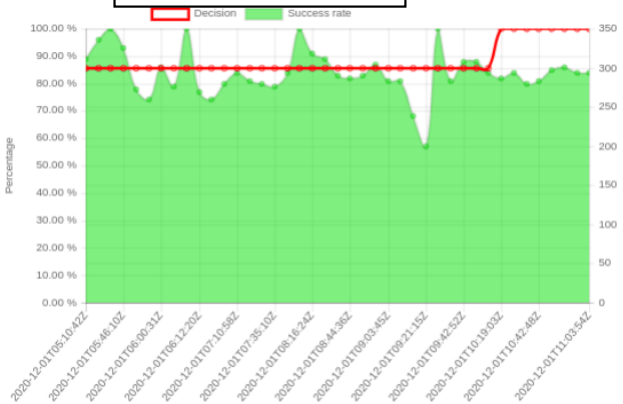

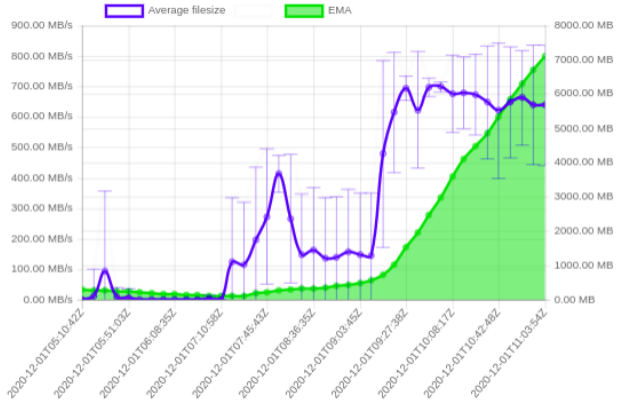

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

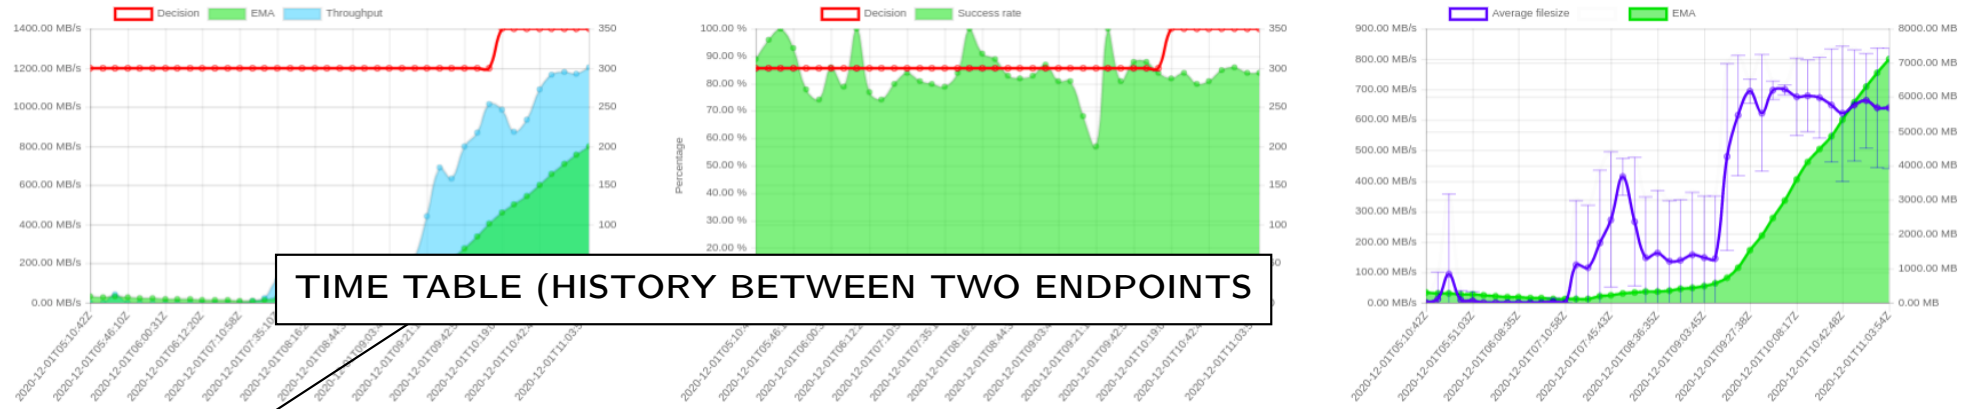

TIME TABLE (HISTORY BETWEEN TWO ENDPOINTS

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

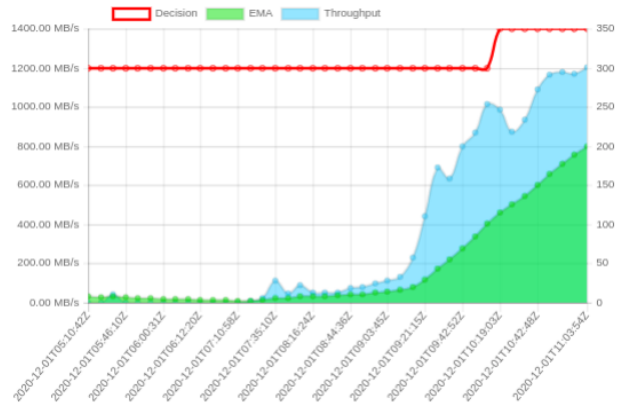

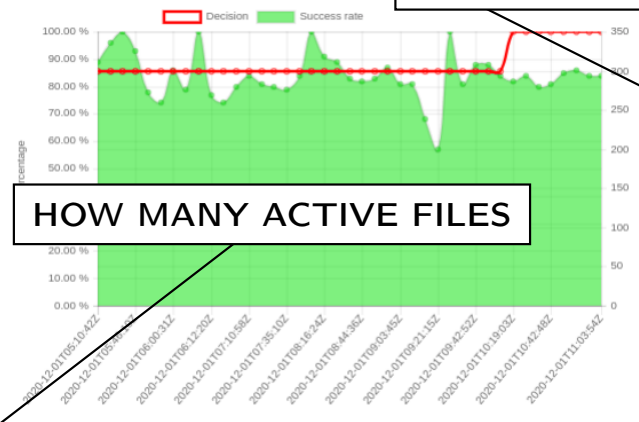

AVERAGE FILESIZE

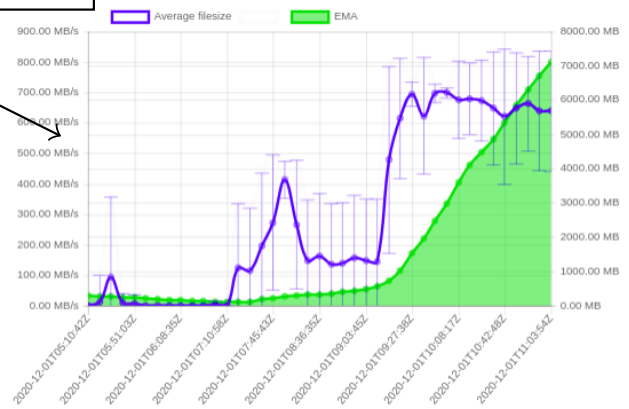

HOW MANY ACTIVE FILES

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

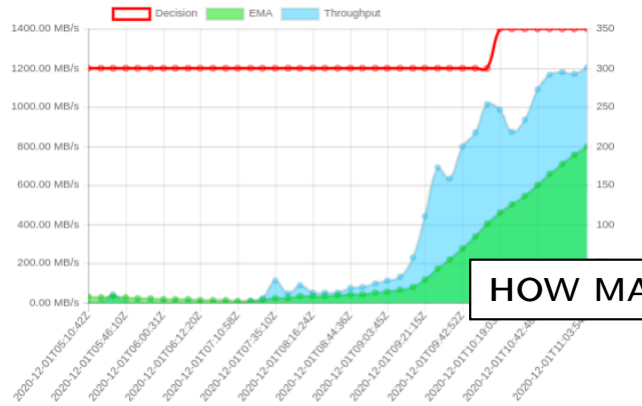

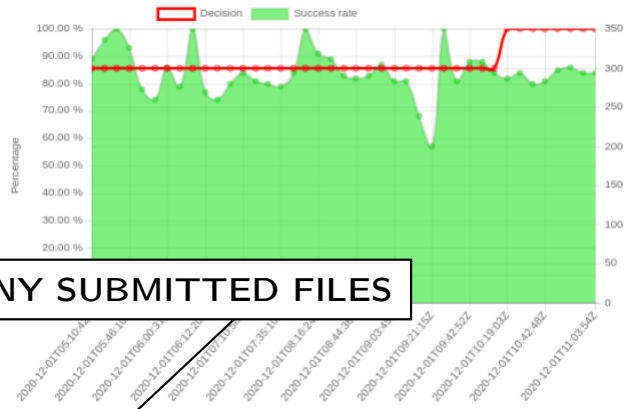

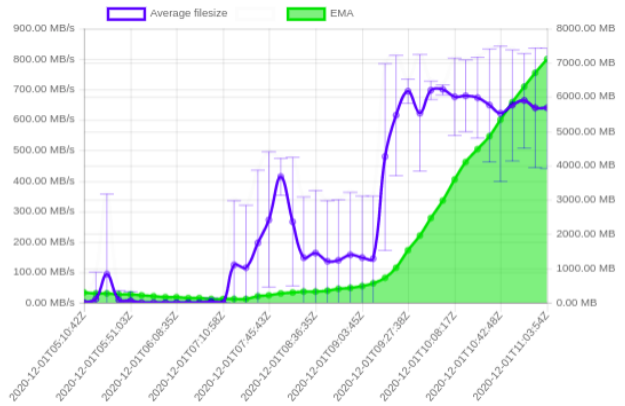

HOW MANY SUBMITTED FILES

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

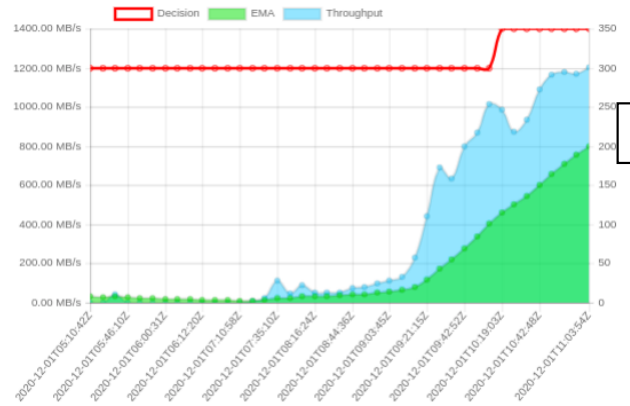

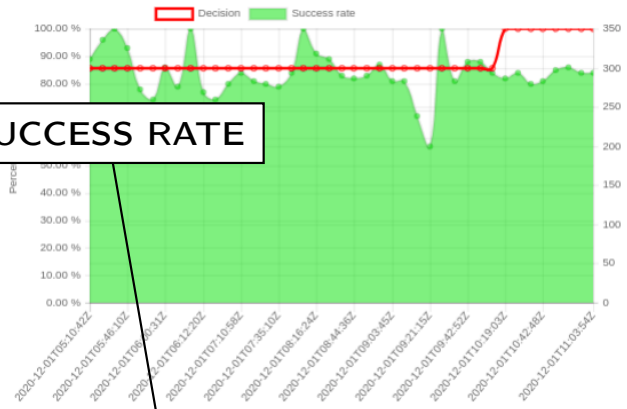

SUCCESS RATE

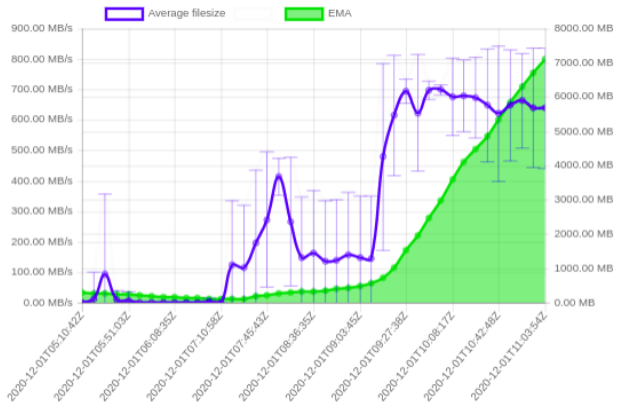

First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

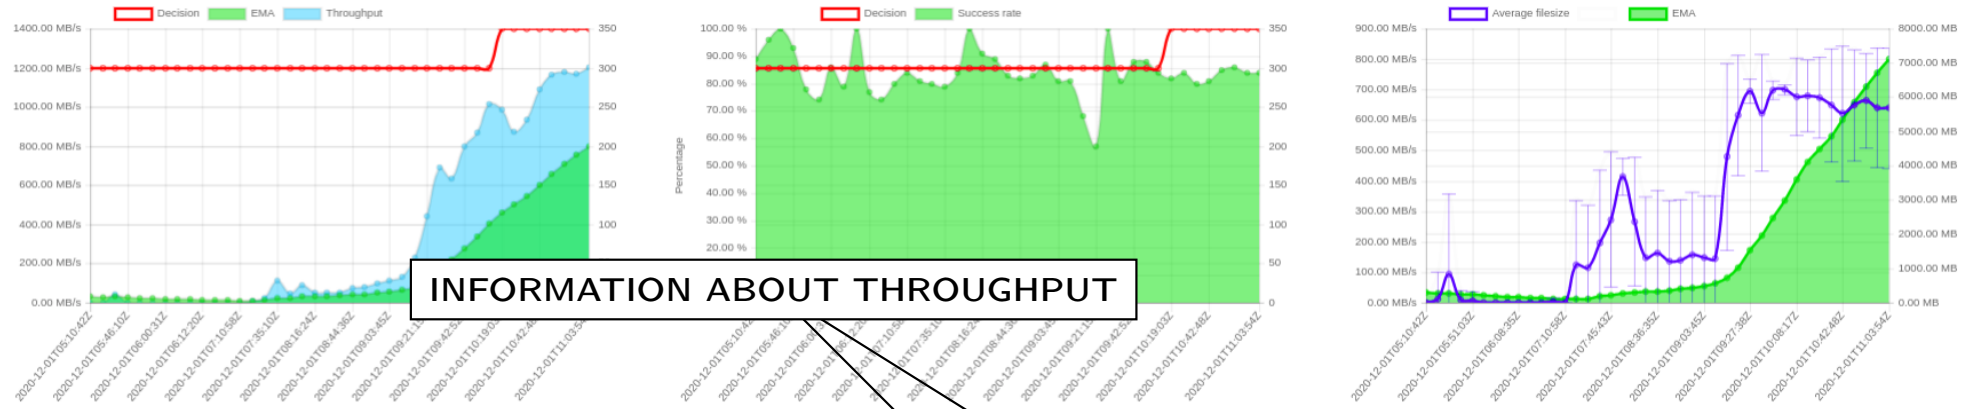

INFORMATION ABOUT THROUGHPUT

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

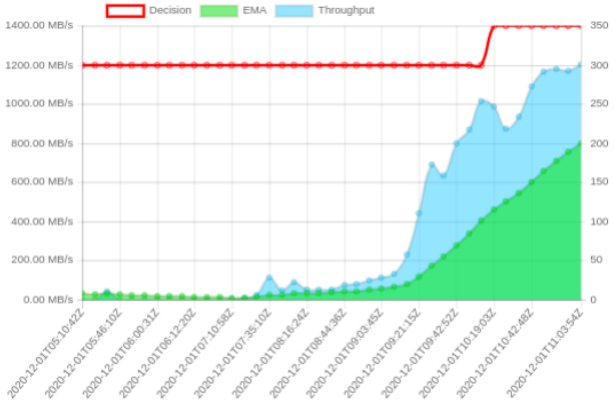

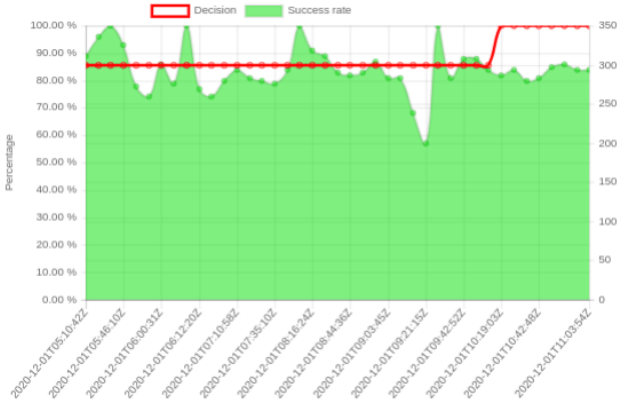

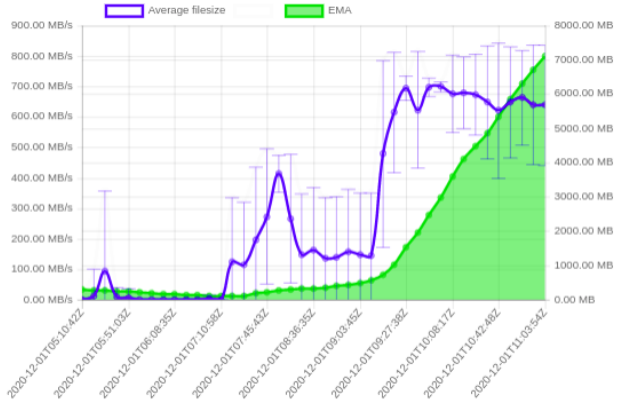

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	133	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

NEW BULK

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

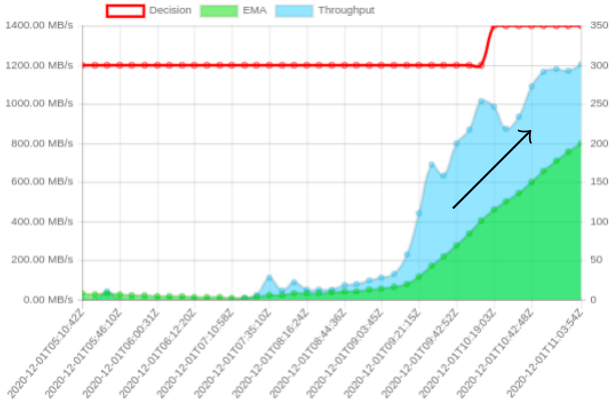

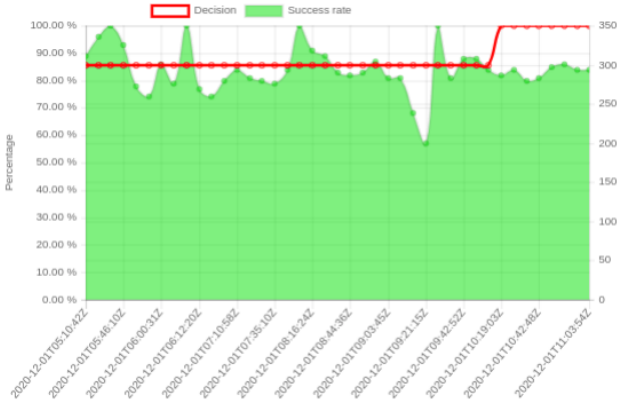

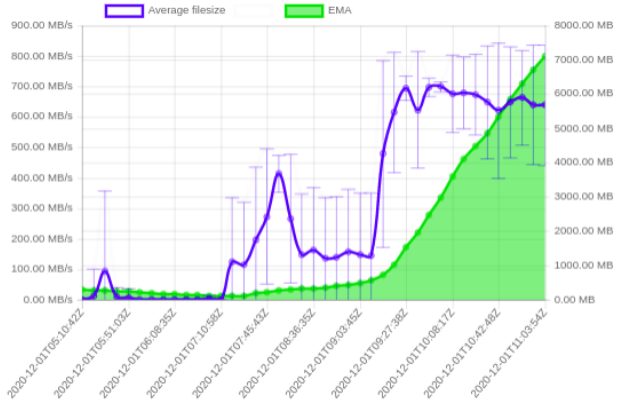

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	192	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

NEW BULK

Details for davs://eosatlas.cern.ch → davs://webdav-at1.pic.es

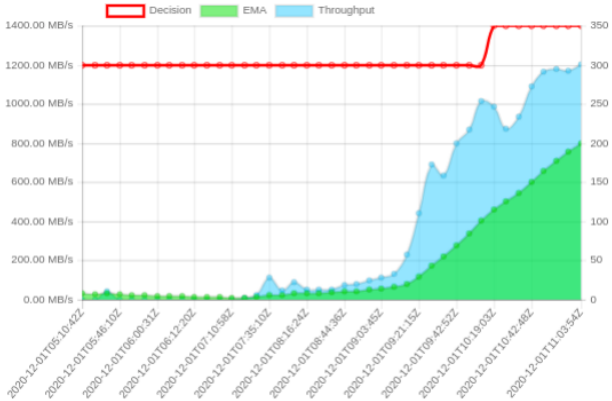

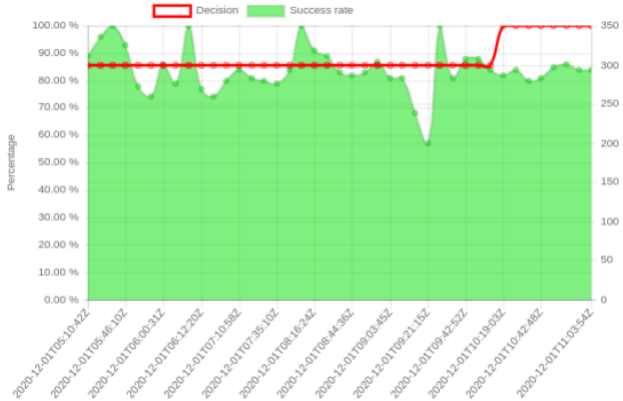

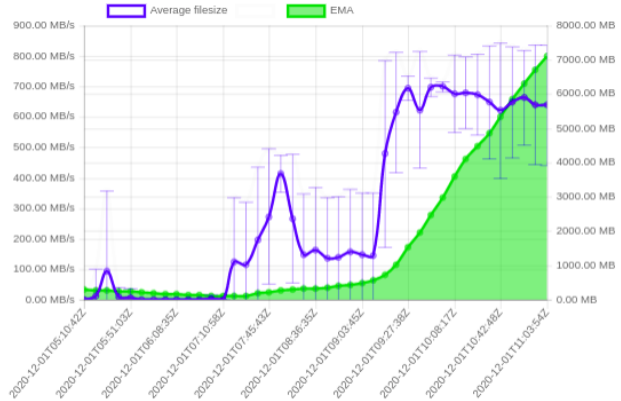

First | Previous | 1 | Next | Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	800.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	184	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	600.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

LINK WASN'T EMPTY

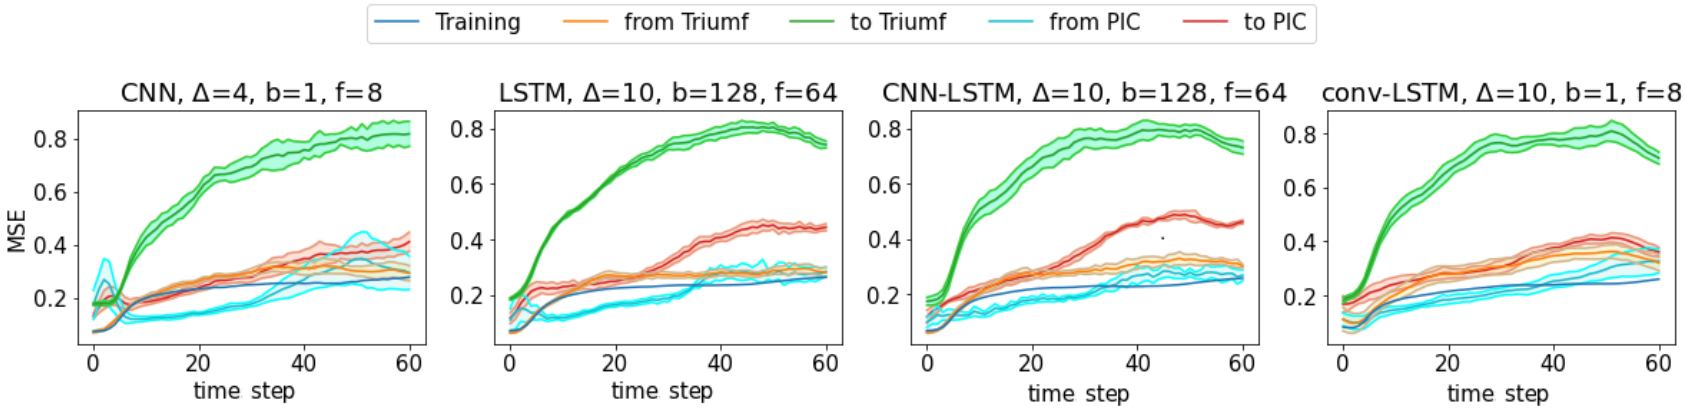

Figure: Average MSE and its variance with respect to the forecasting steps (here: $\Gamma = 60$).

Table: Comparison of model parameters on the test data set representing transfers from TRIUMF to Tier0/Tier1. $MSE_{\Psi,0}$ and MSE_{Ψ} means respectively MSE_0 and MSE average during Ψ period. Ψ period means saturation and short drop between two saturation periods (slide 41). $\Gamma = 15$ (30 minutes). S is the standard deviation over 10 training repetitions.

Δ	Model	Batch - Filters\ Units	$MSE_{\Psi}(\Gamma)$	$S(MSE_{\Psi}(\Gamma))$	$MSE_{\Psi,0}$	$S(MSE_{\Psi,0})$
4	CNN	1 - 8	0.206	0.007	0.206	0.009
	LSTM	128 - 64	0.224	0.008	0.042	0.005
	CNN-LSTM	128 - 64	0.233	0.015	0.060	0.007
	CONV-LSTM	1 - 8	0.159	0.012	0.048	0.007
10	CNN	1 - 8	0.223	0.095	0.223	0.010
	LSTM	128 - 64	0.185	0.012	0.025	0.006
	CNN-LSTM	128 - 64	0.188	0.011	0.021	0.006
	CONV-LSTM	1 - 8	0.125	0.008	0.036	0.008


CERN