UK ATLAS mini DC

A. Forti

LHCOPN/LHCONE

19 March 2025

Why miniDC?

- miniDCs should test smaller bits and be organised in a way that is more easily repeatable than a full DC
    - Either on demand or organised once or twice a year
    - New technology/hardware/network becomes available
        - Capacity & Capability

UK miniDC Goals

- **T1 tests:**
  - Read/Writes to Echo (disk)
    - Capacity tests
      - DC24 T0→RAL
      - Fill 200 Gb/s T1↔RAL
    - Cutting LHCOPN to/from RAL
      - Test failover to LHCONE
    - Test T1 JANET link to UK T2s
    - Stress test new checksum code
      - xrootd gateways
  - Read/Write to Antares (tape)
    - Newly attached LHCOPN link
    - New EOS buffer
  - Stress test 100Gb gateways at RAL
    - One writing to Ceph
    - One writing to CephFS
  - .....

- **T2 tests:**
  - Capacity tests T2s
    - Replicate DC24 rates using only UK sites
      - Sustained rates for at least 6h
    - Increase rates to max bandwidth at 100Gb/s sites
  - Capability testing under discussion
    - Jumbo frames
    - Token AAI
  - ....

– done
– postponed
– to be done

Dates tests

- Need to be careful to not lose the **“mini”** along the way doing too many things in one test
  - Originally all in one week ➜ not so mini anymore
- Eventually one ➜ many miniDC
  - **T2 tests:** Week starting 9th of December 2024
    - ATLAS extended to 2 weeks to test a site ingress/egress one at the time
  - **T1 tests:** 3 March 2025
    - Disk
  - T1 tests: TBD
    - Tape
  - T1 tests: TBD
    - 100 Gb/s boxes
    - Stress test link to T2s

Setup&Rates

DC24 rates (Mb/s)									
ingress	QMUL	RALPP	LANCS	MAN	GLA	RAL	DC24 rates	ingress Space TB	egress space	Full mesh space
QMUL		1421	7105	2842	1421	14211	27000	194	224	419
RALPP	978		978	391	196	1957	4500	32	53	85
LANCS	7526	1505		3011	1505	15053	28600	206	221	427
MAN	7136	1427	7136		1427	14273	31400	226	90	316
GLA	2935	587	2935	1174		5870	13500	97	50	147
RAL	9286	1857	9286	3714	1857		26000	187	333	520
Total	27861	17237	27440	11132	6406	51362	131000			
egress	QMUL	RALPP	LANCS	MAN	GLA	RAL	DC24 rates	egress space TB	QMUL	RALPP
QMUL		1274	6368	2547	1274	12737	24200		0	9
RALPP	891		891	357	178	1783	4100		6	0
LANCS	6789	1358		2716	1358	13579	25800		49	10
MAN	6409	1282	6409		1282	12818	28200		46	9
GLA	2652	530	2652	1061		5304	12200		19	4
RAL	14393	2879	14393	5757	2879		40300		104	21
Total	31135	7322	30714	12438	6970	46221	134800		224	53

- Setup:
  - dc_inject used in DC24
    - script to create rucio rules according to rates per link
  - T2: DATADISK (prod space) ➔ SCRATCHDISK (user space, no particular permission needed)
  - T1: DATADISK ➔ DATADISK (quota setup, but scratchdisk not big enough for T1 rates)
- Replicate the DC24 rates
  - Rates per link re-calculated to take into account reduced number of links
    - This info is needed by dc_inject
  - For T1 simply split evenly between T0 and T1

Tier2 results

T2 summary

- Full summary and detailed notes for each site
- Similarly to DC24 upstream network not the bottleneck
- Most problems stem from
  - ToR switches
  - Gateways
  - Storage backend
  - Ability to digest or
  - source the data
    - Space availability at the destination
    - Data availability at the source
  - Concurrent production traffic
  - Checksum calculations in ingress at xrootd sites

Site	Test	Bw	Expected	Achieved	ingress%	Bw	Expected	Achieved	egress %
Manchester	DC24	40	28	16	57	40	31	38	122
Lancaster	DC24	100	29	35.3	120	100	26	26	100
QMUL	DC24	100	27	56	200	100	24	49	200
RALPP	DC24	100	4.5	35.5	777	100	-	-	-
Glasgow	DC24	20	13.5	5.3	40	20	12	10	83
Lancaster	100 Gb/s e/i	100	80	35	44	100	80	85	106
QMUL	100 Gb/s e+i	100	80	27	34	100	80	24	30

T2 tests: Manchester

- Uplink 40 Gb/s 4 xrootd gateways + cephfs
- Most problems stemming from completely different internal network usage and reduced number of servers accepting data transfers
  - from 58 independent servers 10 Gb/s ➜ 4 gateways 25 Gb/s + 40 ceph interconnected data servers
- egress improved to the bandwidth limit by
  - Spreading gateways in racks with better connection
  - Tuning xrootd redirectors parameters to avoid accumulation of transfers on a single gateway
- ingress couldn’t go above 16 Gb/s we think internal network is the culprit
  - Next steps
    - Upgrade old 10 Gb/s ToR switches
    - Review networking with University network team

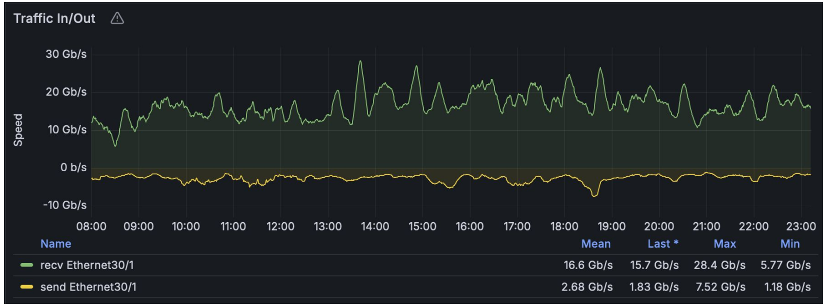

T2 tests: Lancaster

- Uplink: 100 Gb/s 6 xrootd gateways + cephfs
- Notable problems in ingress with checksum calculations doubling the internal traffic
    - read-back checksumming is a problem at **all xrootd sites** (see later)

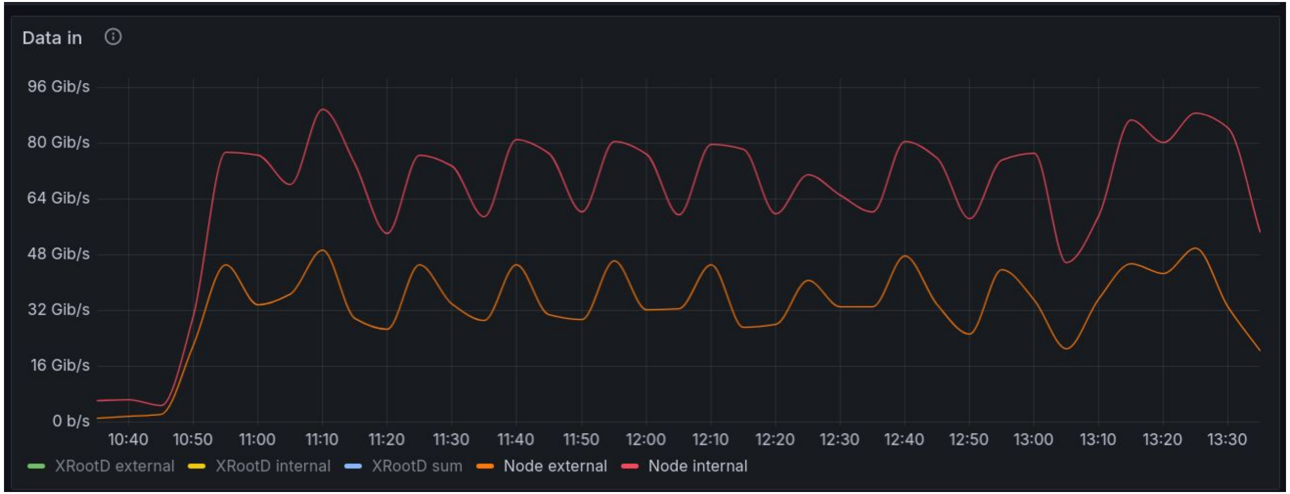

- DC24 rates achieved with ease
    - 100Gb/s in ingress hampered by lack of space

T2 tests: QMUL

- Uplink 100 Gb/s storm + lustre
- QMUL reached all the targets but noticed that the storage at high rates particularly with r/w was running “red hot!”
  - Note that production often saturates the available bandwidth already particularly in egress
- To improve the rates further aside of BIOS tweaks the only way is to upgrade the storage but this poses the usual questions of how much it is necessary and when/

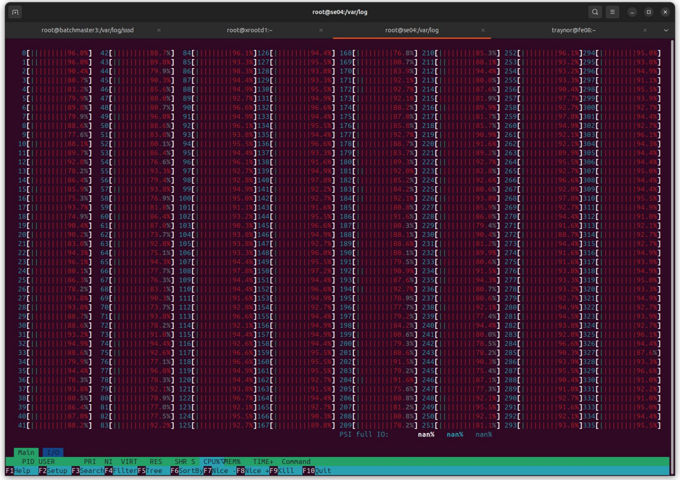

T2 Tests: Glasgow

- Uplink: 20 Gb/s single gateway + ceph
- Problem: increase of deletion time
  - This is really a problem during a DC where fast cycles of injections and deletions is necessary ➔ this slows down rates since FTS cannot copy until the file is deleted
- Problem: checksum calculation
  - Need to add more gateways or beef up the current one
  - Software components not as reliable to do that yet

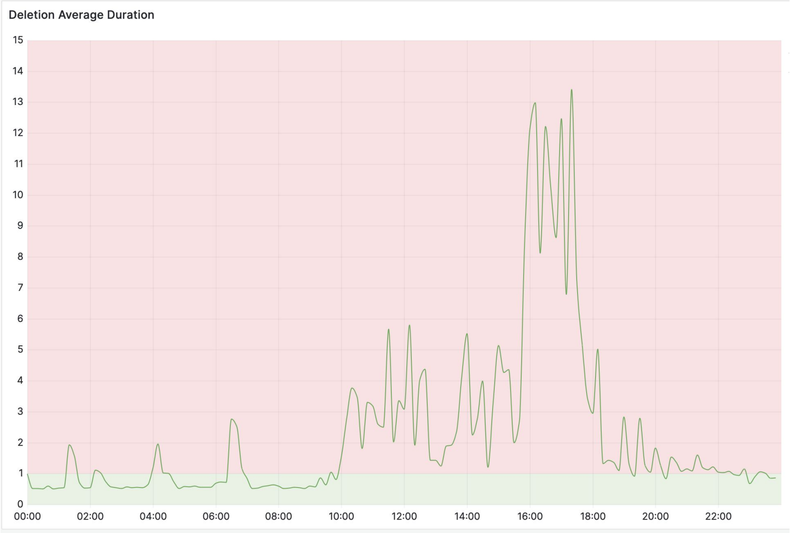

T2 tests: RALPP

- Uplink: 100 Gb/s dcache
- Only T2 shared with CMS but **ATLAS T3**
  - testing the network at 100 Gb/s was limited by the storage space
- ingress achieved 50 Gb/s
  - Increase to 600 concurrent transfers caused a degradation
  - OOM errors started to appear
    - Caused by limits in dcache later corrected by the site
- egress not enough data at the site to really try
  - Perhaps could have worked with file override (?)

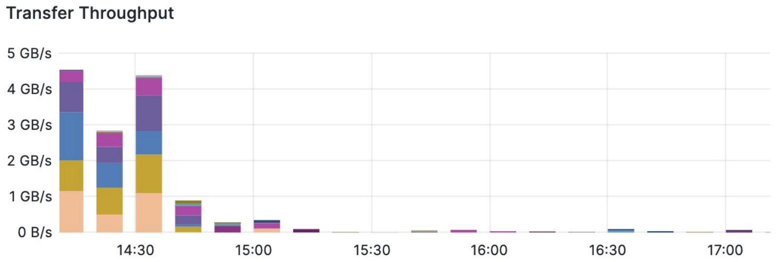

	min	max	avg	current
RAL-LCG2	0 B/s	1.15 GB/s	151 MB/s	0 B/s
UKI-NORTHGRID-LANCS-HEP	0 B/s	1.08 GB/s	149 MB/s	0 B/s
UKI-LT2-QMUL	0 B/s	1.34 GB/s	144 MB/s	0 B/s

Tier1 results

T1 tests

- Main goal was to test LHCOPN and echo
- ATLAS used also the Europeans T1
  - Spread the load
  - Avoid space issues in egress
  - Check how things change adding new sites

Date	Source	Destination	Rates Mb/s	Test
3/3/25	CERN-PROD_DATADISK	RAL-LCG2-ECHO_DATADISK	44000	T0 export
4/3/25	CERN-PROD_DATADISK	RAL-LCG2-ECHO_DATADISK	44000	T0 export interrupting LHCOPN failover to LHCONE
5/3/25	CERN-PROD_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	IN2P3-CC_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	INFN-T1_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	NDGF-T1_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	NIKHEF_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	PIC_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
5/3/25	SARA-MATRIX_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	T1-T1 ingress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	CERN-PROD_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	IN2P3-CC_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	INFN-T1_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	NDGF-T1_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	NIKHEF_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	PIC_DATADISK	10000	T1-T1 egress over LHCOPN
6/3/25	RAL-LCG2-ECHO_DATADISK	SARA-MATRIX_DATADISK	10000	T1-T1 egress over LHCOPN
7/3/25	CERN-PROD_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	IN2P3-CC_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	INFN-T1_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	NDGF-T1_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	NIKHEF_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	PIC_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code
7/3/25	SARA-MATRIX_DATADISK	RAL-LCG2-ECHO_DATADISK	10000	checksum code

General results

- UK T1 disk miniDC was a success
  - Most of the problems we had during DC24 have disappeared
    - T0 export rates
    - General rates
    - Tokens
  - LHCOPN cut didn't affect rates
  - Checksum new code doesn't have obvious side effects

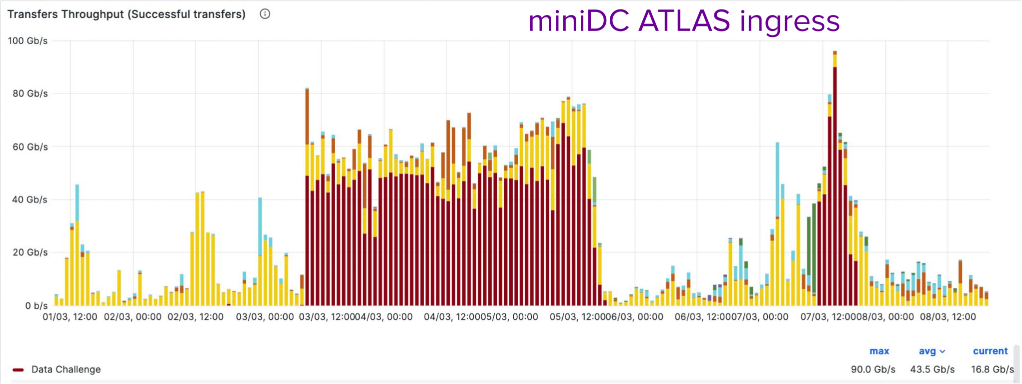

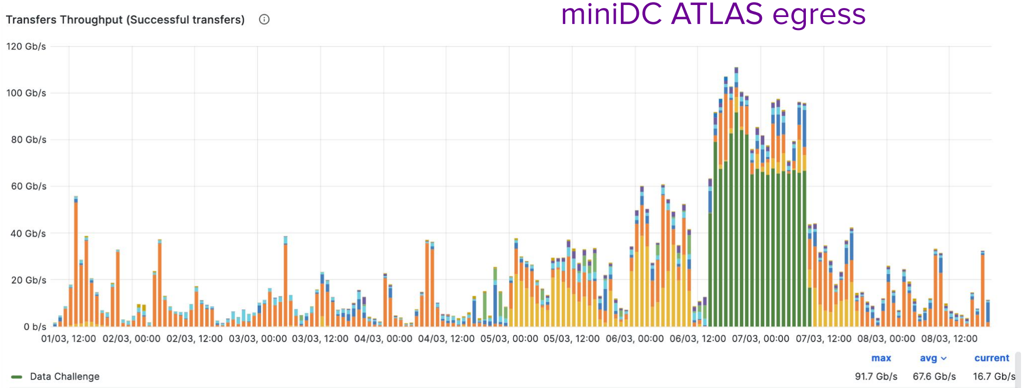

T0 export rates

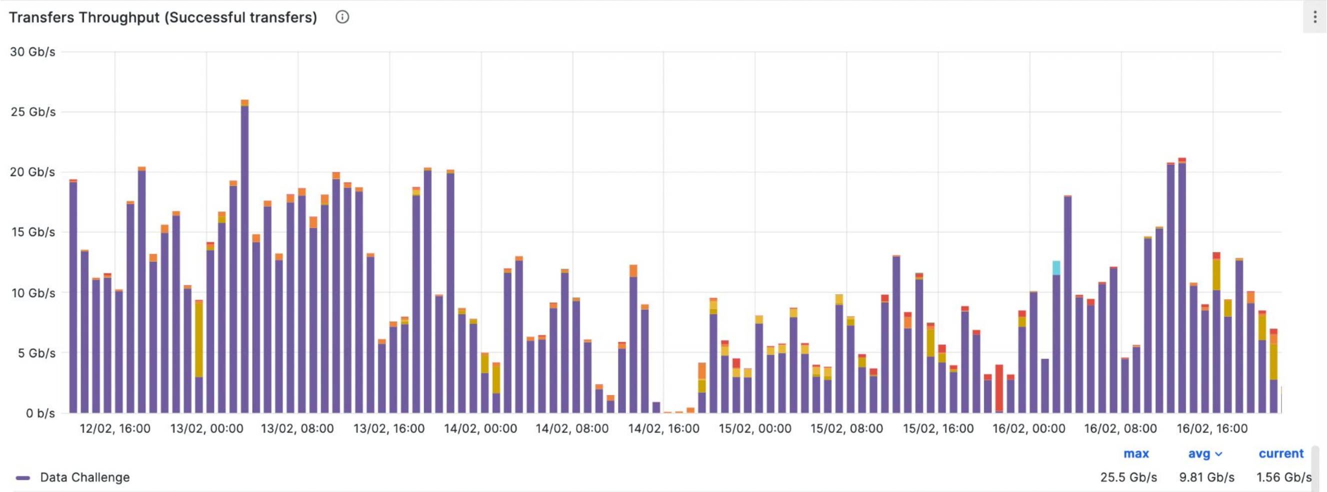

DC24 T0 export expected rates 44 Gb/s achieved ~ 10 Gb/s wobbling due to several concurrent problems

miniDC T0 export achieved ~51 Gb/s sustained. Expected still 44 Gb/s

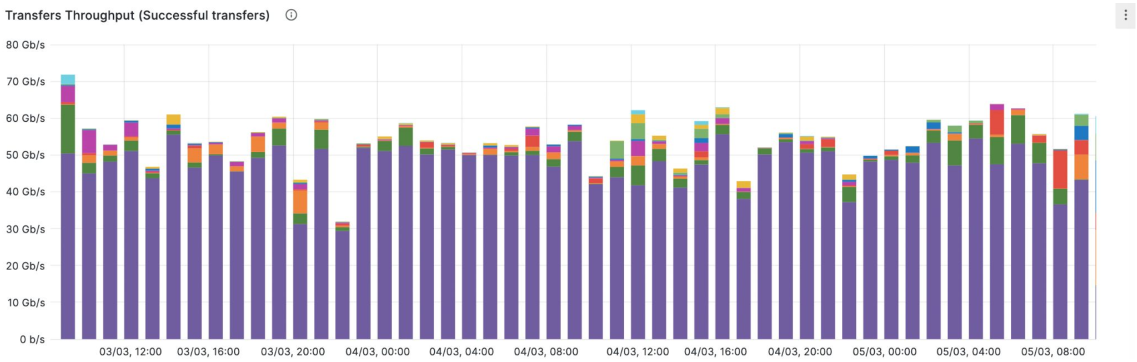

T0 export: tokens v x509

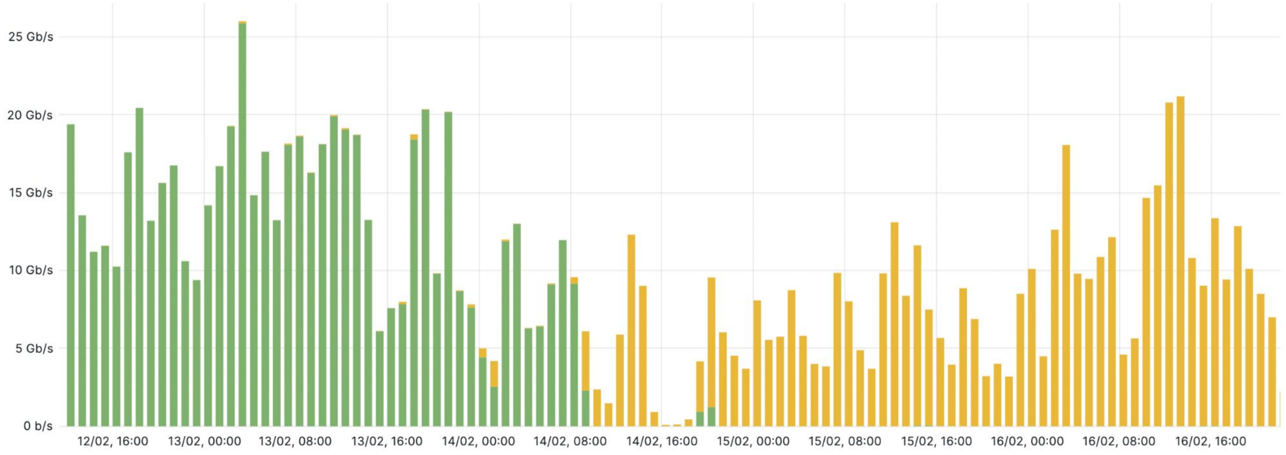

DC24 tokens had to be switched off after 2 days

miniDC T0 export almost exclusively done with tokens

LHCOPN cut


LHCOPN UK-RAL Total - 200Gbps
100 Gb
50 Gb
0 b
-50 Gb
-100 Gb
-150 Gb
-200 Gb
-250 Gb
03/03, 12:00
03/03, 16:00
03/03, 20:00
04/03, 00:00
04/03, 04:00
04/03, 08:00
04/03, 12:00
04/03, 16:00
04/03, 20:00
05/03, 00:00
05/03, 04:00
05/03, 08:00
min max avg current total
in total - RAL to CERN 0 b 75.5 Gb 11.1 Gb 14.1 Gb 6.34 Tb
out total - CERN to RAL 0 b 230 Gb 132 Gb 72.1 Gb 75.0 Tb

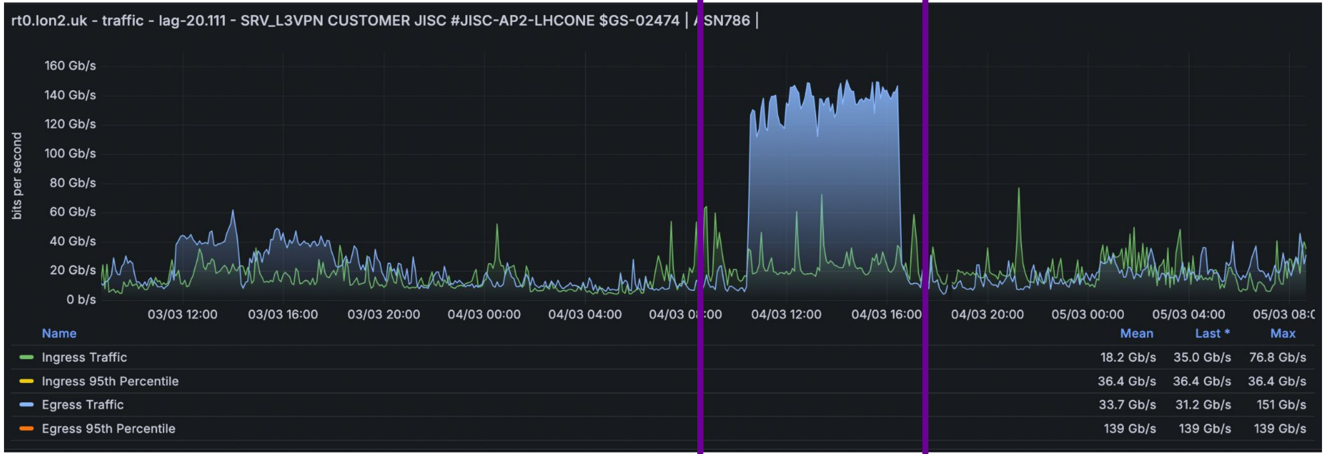
rt0.lon2.uk - traffic - lag-20.111 - SRV_L3VPN CUSTOMER JISC #JISC-AP2-LHCONE $GS-02474 | ASN786 |
bits per second
160 Gb/s
140 Gb/s
120 Gb/s
100 Gb/s
80 Gb/s
60 Gb/s
40 Gb/s
20 Gb/s
0 b/s
03/03 12:00
03/03 16:00
03/03 20:00
04/03 00:00
04/03 04:00
04/03 08:00
04/03 12:00
04/03 16:00
04/03 20:00
05/03 00:00
05/03 04:00
Name Mean Last * Max
Ingress Traffic 18.2 Gb/s 35.0 Gb/s 76.8 Gb/s
Ingress 95th Percentile 36.4 Gb/s 36.4 Gb/s 36.4 Gb/s
Egress Traffic 33.7 Gb/s 31.2 Gb/s 151 Gb/s
Egress 95th Percentile 139 Gb/s 139 Gb/s 139 Gb/s

LHCOPN cut

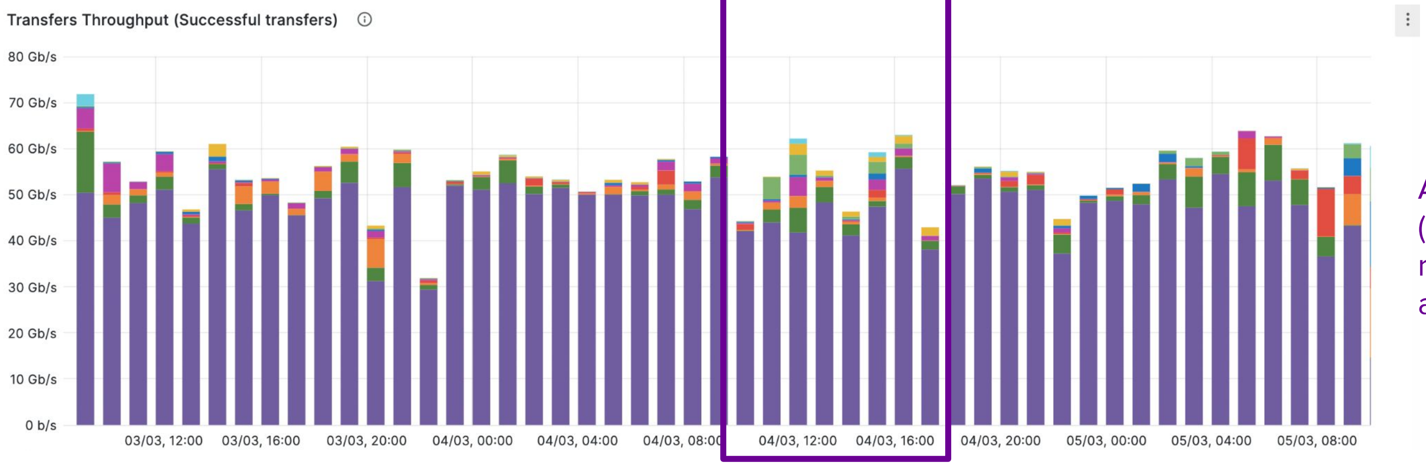

ATLAS LHCOPN (T0+T1) rates remained in average

Traffic from/to T1

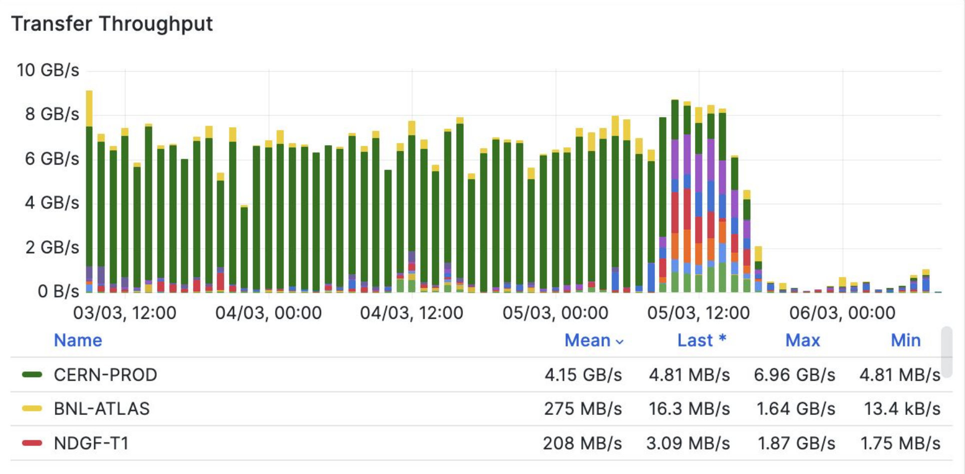

Ingress smooth switch from CERN only miniDC traffic to mixture of T1s but added load on gateways added a tail of time outs. 100 Gb/s nics should solve this

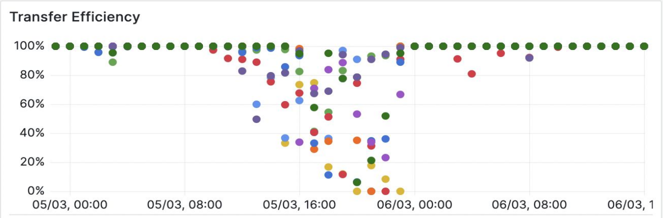

egress tooth comb pattern due to T1s absorbing files too fast. Positive but might be also due to the system injecting smaller files.

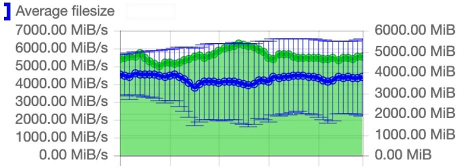

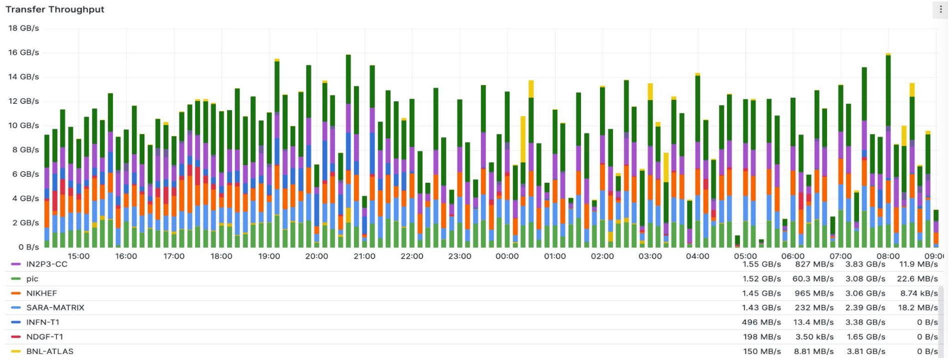

Checksum code

- RAL (Jyothish) wrote code for XrdCeph, the echo plugin, to calculate checksum on-the-fly, i.e. no readback
    - It is hoped that this code can be deployed eventually also at standard xrootd sites (Lancaster, Manchester, Durham, Bristol, Brunel...)
- This miniDC was the first stress test of the code.
    - The code didn't add load to the gateways where it was deployed and didn't affected their throughput
        - Deployment for now is concurrent to read-back checksumming to validate the checksums for the next few months

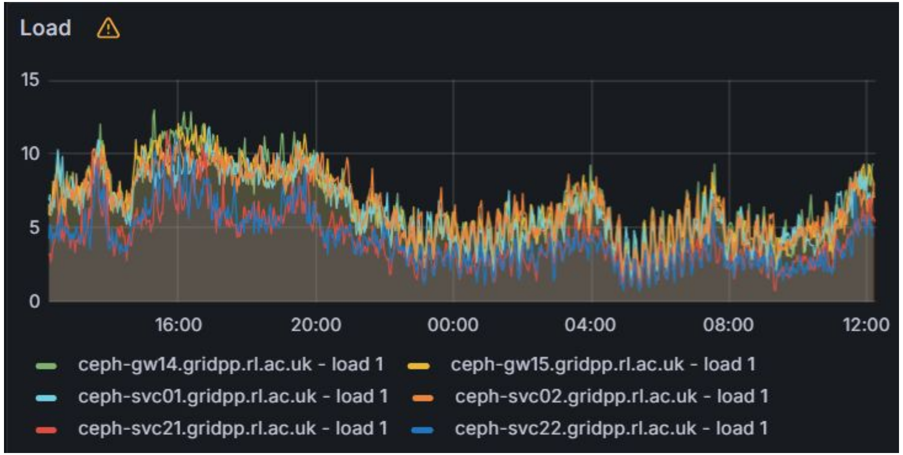

Some gateways have new code some don't, load is similar

Tools

dc_inject

- Used dc_inject for all the tests
- Pros:
  - **Extremely** easy to use
  - Tests the whole chain of services
    - Bottlenecks can be anywhere and this highlights it
- Caveats:
  - Tests the whole chain of services
    - If there is a bottleneck upstream it isn't testing downstream
    - Needs knowledge of FTS/rucio to interpret
  - Stochastic behaviour really good for global rates
    - More unpredictable on targeted tests
- There is effort to make it more robust and incorporate it in rucio
- But we should bundle also other tools, useful to tests and debug
  - **Ideal is to have a collection of tools at least documented to make it easier for people to do consistent tests.**

Conclusions

- System is complicated every little nut and bolt can affect the rates and has to be studied.
- Not only network ;)

Backup