

# Tier-2 Sites Networking status

Duncan Rand and sites

# Site WAN bandwidth and throughput

1. Is the WAN connection bandwidth sufficient?
    - Based on site role, size and network forward look
  2. How well is the site able to make use of what they currently have?
    - Assessed from aggregate FTS throughput
  3. FTS individual file transfer rates
- perfSONAR is useful for monitoring network quality
    - Not going to discuss perfSONAR

# Site WAN connectivity

- The Network Forward Look suggests connectivity sites are expected to need, e.g. 100 Gb/s for 'large' sites
- Some sites have that connectivity and some await it
- Arguably we also desire that sites have the ability to transfer data at sufficiently high rates
  - Ideally we would like them to be able to fill their existing link
- For FTS, throughput is the product of individual file transfer throughput rate and number of active transfers
- Many things can hinder flows, e.g. poor filesystem performance or institutional security devices such as firewalls

# Aggregate FTS throughput

- Aggregate throughput is best measured with actual transfers
  - The Data Challenge useful in stressing sites
- Attempt to assess download ability by examining peak FTS download rate over last 6 months **with hourly sized bins** (typically in GB/s)
- A proxy for maximum achievable throughput and likely to be less than the peak network flow seen on network monitoring (typically 5 min resolution in Gb/s)
- A possible future measure could be the average rate calculated from the time taken to download a large (50-100TB) dataset from a reliable site e.g. CERN

# Information requested from sites

- *Current capacity of your Tier-2 network connection out to JANET (i.e. is it shared, is it dedicated, what is the limiting bandwidth)*
- *Campus JANET network connection capacity*
- *Upgrade planning / expectation if any*
- *Blockers if any*
- *Traffic plots for Tier-2 I/O rates over say last few months if you can get them*
- *Site IPv6 status (separate talk)*

# Reminder: Network Forward Look (September 2021)

- In the next 2 years, the “large” Tier-2 sites (Glasgow, Imperial, Lancaster, Manchester, QMUL, RAL) are likely to need up to 100 Gbit/s for all LHC traffic. In fact, several are already connected at 100 Gbit/s which demonstrates excellent forward planning with Jisc.
- We expect the “other” Tier-2 sites to require 10-20 Gbit/s for LHC traffic, though there are some “medium” sites that currently have a high network I/O rate and a connection bandwidth similar to the larger sites.
- The summary of the scale of connection required is:
  - **Tier 1: 200 Gbit/s**
  - **Tier-2 (large): 100 Gbit/s**
  - **Tier-2 (other): 10-20 Gbit/s and (and in some cases possibly 40 Gbit/s)**

# Tier-2 site connectivity to Janet

- Assume large need: 100 Gbit/s, med: 20-40 Gbit/s, small: 10 Gbit/s
- Notation: site (**existing connection/required connection** in Gbit/s)
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (10/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# Tier-2 site connectivity to Janet


- Assume large need: 100 Gbit/s, med: 20-40 Gbit/s, small: 10 Gbit/s
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)



# Imperial Tier2 Network Report

- Site currently has 100G to the outside world
- Also on LHCONE
- No upgrades currently planned
- (Fully IPv6 enabled)

- Inbound transfer in March 2022 from CERN to Imperial
- Volume: 150 TB
- Peak rate : 95 Gbps for ~ 2 hours



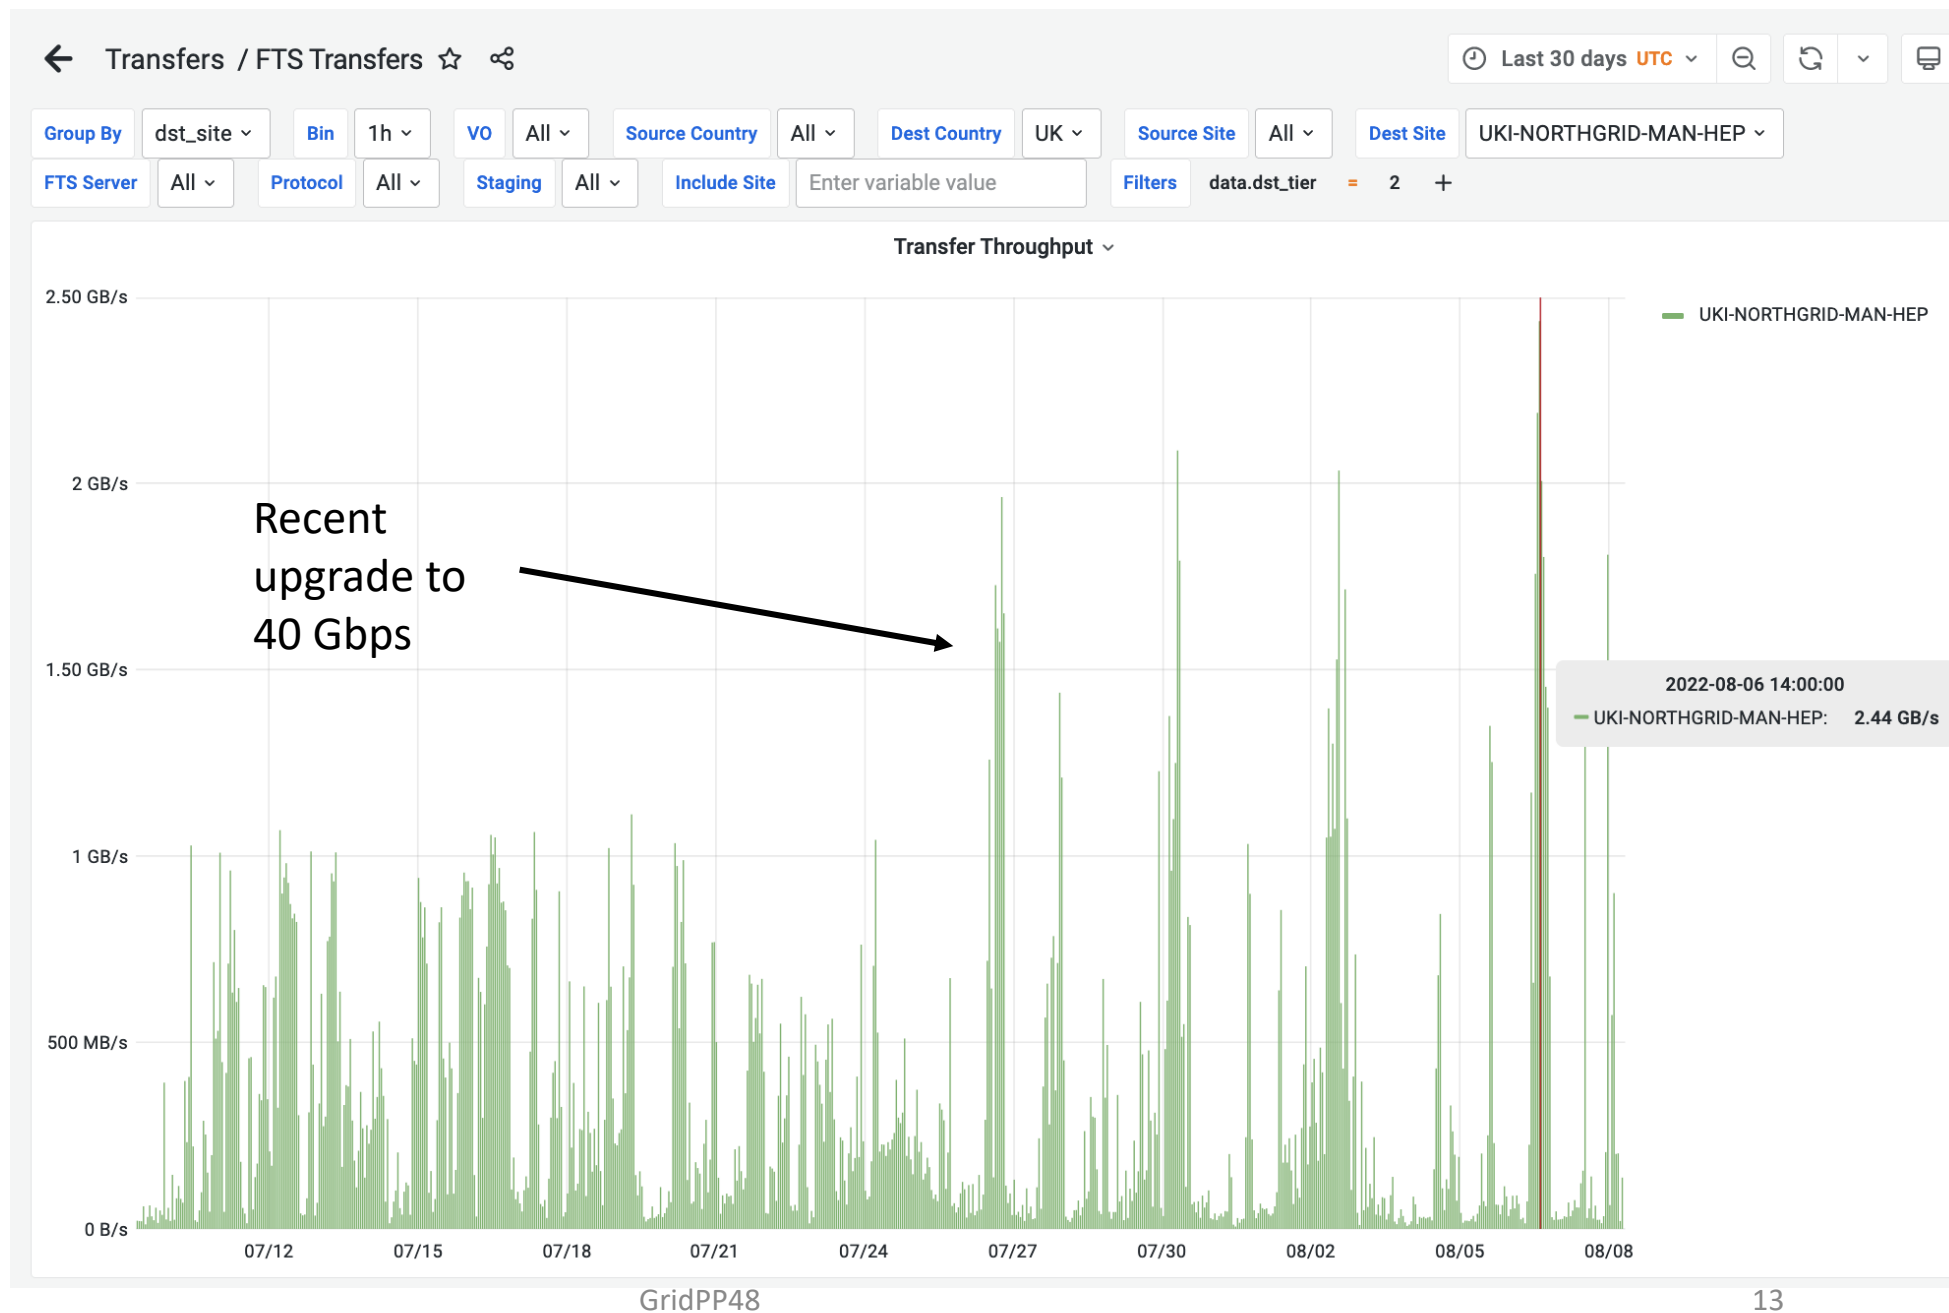
# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# Manchester

- University has 100G connection to Janet
- 40G allocated to UKI-NORTHGRID-MAN-HEP

Peak over  
last 6  
months:  
  
17 Gb/s (2.44  
GB/s)



# Oxford Tier-2 Network Status

## GRID Cluster

- The GRID cluster is currently connected to the backbone at 10Gbit.
- 2\*10Gbit connection upgrade is paid for and it will be enabled in August 2022.

## Connection to Janet

- The University is currently connected to Janet at 2\*20Gbit.
- The University Janet connection upgrade to 2\*100Gbit is in place and will be configured into production in July 2022.

## IPV6

- The Central IT Services are in the early stages of a project to roll out a production IPV6 service across the University. They have applied for funding from the University.

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)


# Lancaster

- Bandwidth: 40Gb/s (4\*10)
- Occupying the "new" backup link that now runs parallel to the regular link, southwards.
- No plans to upgrade to 100Gb in the near future, but groundwork for such an upgrade is part of scheduled updates to the hardware over the next few years.
- As of the 18th of August we effectively doubled the bandwidth of our new CEPH/xroot SE, but not fully exercising it or our external link - will need to regain nucleus status (soon).
- We need to exercise this link before we can justify moving towards a bandwidth upgrade.
- Had IPv6 for years (configured using DHCPv6), and the site is outside the institution firewall.



Group By: dst\_site ▾ Bin: 1d ▾ VO: All ▾ Source Country: All ▾ Dest Country: UK ▾ Source Site: All ▾ Dest Site: UKI-NORTHGRID-LANCS-HEP ▾ FTS Server: All ▾ Protocols: All ▾  
 Staging: All ▾ Include Site: Enter variable value Filters: data.dst\_tier = 2 +

Transfer Throughput ▾



Upgrade to 40G



Peak  
inbound  
rate over  
last 6  
months:

20 Gb/s  
(2.5 GB/s)


01/09/2022

Brunel: site has 40G (4x10G)

Peak over last  
6 months:

11.2 Gb/s  
(1.4 GB/s)

01/09/2022



# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)



# Required Questions

- Current status

We currently have a 2x10Gb/s link direct to the JANET router. At present, this is not available to the site as I have bonded the link from the gateway to the switch. Internally, it's almost all 10Gb/s ethernet

- Upgrade planning if any

None planned at present as it doesn't seem necessary given traffic rates

- Any existing agreements between JISC and site authorities

The grid site is outside the Bham firewall and everyone seems happy with this arrangement

- Expected timeline/blockers...

N/A

- IPv6 status

Stalled. The gateway can route packets but after sometime trying to get DNS entries out of the university, I got told they aren't using it yet and don't have plans to in the near future (!). I will try to resurrect this in the coming months

# UKI-SOUTHGRID-BRIS-HEP

- Current status:
  - Up to 20 Gbit/s, IPv4 + IPv6
  - Some issues observed with perfsonar → looking into it (takes time, low on FTEs)
- Upgrade planning
  - Thinking about 25 Gbit/s on NIC level
  - Some rumours about University of Bristol going to 100 Gbit/s
- Any existing agreements between JISC and UoB
  - Query sent to IT Networks, those within GridPP with connections to JISC probably also know more than site admins
- Expected timeline/blockers: staff oversubscribed

# Sussex and Sheffield

	Janet Connection	Tier-2 Connection	Forward look	Comment from Site Admin
UKI-SOUTHGRID-SUSSEX	3 x 10Gbs	1 x 10Gbs shared	Network modernisation project about to start	
UKI-NORTHGRID-SHEF-HEP	None shared 1 x 10Gbps link for LHC / Grid / HPC	1 x 10Gbps uplink	No bandwidth upgrade is planned for LHC / Grid / HPC activities.	All routes through Institution firewall.

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), , RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# QMUL

- Current capacity of your Tier-2 network connection out to Janet
  - In principle we should have 20Gb/s using the backup link. However due to a number of issues we are sharing the main link and have been limited to 10Gb/s
- Campus Janet network connection capacity
  - QMUL has at least one 100Gb/s connection.
- Upgrade planning / expectation if any
  - New 100Gb/s capable network is in the final stages of being commissioned. New firewall has been tested in anger (first test found the rules had been written in backwards!). Will go in to production on the 1st September. New IPS functionality will then be commissioned, once that is done the old IPS will be removed (which limits us to 20Gb/s) and we should then get access to 100Gb/s link ~ November.
- Blockers if any
- Traffic plots for Tier-2 I/O rates over say last few months if you can get them
  - we are able to utilise 10Gb/s for sustained periods both in and out of the site.
- Site IPv6 status
  - Storage and worker nodes now have working IPv6



# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# ScotGrid Glasgow

- **Current WAN status:**  $2 \times 10$  Gbps (of which we can use c. 10 Gbps before causing issues for central IT networking team)
  - Site link to campus is 40 Gbps with 40 Gbps (fully manual) failover
- **Upgrade planning:**
- **Existing JISC / site agreements:**
- **Expected timeline / blockers:**
  - Campus-wide network renewal programme has just moved into phase 4 (“Detailed design and implementation of the Core, Distribution and Initial Edge solutions”)
  - Campus development has also resulted in multiple large new buildings coming online (JMcCS LTH, M-S ARC, Clarice Pears) while others (James Watt North) are redeveloped
  - Together, this has been occupying much of our excellent but extremely overworked networking team’s time for several years, and will continue to do so for some time
- **IPv6 status:**
  - Enabled on perfSONAR
  - Campus IPv6 performance is poor due to constraints of existing infrastructure
    - It works, but please don’t use it
  - Hopefully, this situation will improve as part of network renewal programme (see above)

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# RALPPD

	Janet Connection	Tier-2 Connection	Forward look	Comment from Site Admin
UKI-SOUTHGRID-RALPP	2x100 Gbit/s Active + 2x100 Gbit/s Passive Connections	2x40 Gbit/s from T2 router to the Border routers but 2x10 Gbit/s bottleneck to the T2 Router	Hardware procured at the end of last FY to upgrade the T2 router to 2x100 Gbit/s to the Boarder Routers and 2x100Gb downlink. Discussing connections to the Tier 1 as their network design finalises	Tier 2 Network is a DTZ and is connected to LHCONE

- RAL site connected to Janet at 200G
- RALPPD on LHCONE
- Some further internal campus work needed to provide good connectivity to Janet?

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# RHUL Tier-2 Network Report

- Current status
  - 10 Gb/s active and failover links dedicated to Tier2 - does not seem currently to be a limiting factor for our site
- Upgrade planning
  - Given size of site expect to need 20-40 Gb/s in the future, discussed upgrade path with JISC
  - They seemed to agree in principle to upgrade to 100 Gb/s for whole college, but would no longer offer two separate links; interim option of additional 10Gb/s links also discussed
- Expected timeline/blockers...
  - Change from separate to shared link requires major re-work of border network, but this is planned anyway in ~1 year
  - Would need to buy new equipment even to support the interim option of additional 10Gb/s links, so maybe not worth it
  - We have a good relationship with the local network team, but they are under-resourced and find it hard to get around to our requests

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# Durham

	Janet Connection	Tier-2 Connection	Forward look	Comment from Site Admin
UKI-SCOTGRID-DURHAM	10Gbps (Technically there's multiple 10Gbps, we just have one dedicated to us)	10Gbps	Moving to 2x 10Gbps in 2021. Pushing for minimum of 40Gbps by 2022/2023. Wanting to aim for 100Gbps by 2025.	Dedicated Science DMZ



# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# GridPP Edinburgh Networking

- 2x10Gbps uplink from GridPP Tier2, all servers on 10Gbps.
- 100Gbps to ACF. ACF running efficiently close to physical capacity.
- University of Edinburgh started major network upgrade in 2020. Almost complete.
- Main bottlenecks other than 20Gbps comes from using shared compute cluster with other researchers.
- IPv6 support at UoE is good, but slow uptake means maintenance issues. E.g. dropped IPv6 for ~1.5mo during vlan migration.

# Tier-2 site connectivity to Janet

- Assume large need: 100, med: 20-40, small: 10
- Sites with  $\geq 100$  Gbit/s to Tier-2 already or imminently
  - Imperial (100/100)
- Sites with  $\geq 100$  Gbits/s to campus and less to Tier 2
  - Manchester (40/100), Oxford (20/20)
- Other sites with no problem
  - Lancaster (40/100), Brunel (40/20), Birmingham (10/10), Bristol (20/10), Sussex (10/10), Sheffield (10/10)
- All other sites with actual or looming problem
  - QMUL (10/100), Glasgow (10/100), RALPPD (20/100)
  - RHUL (10/20), Durham (10/20)
  - ECDF (10/20), Liverpool (10/20)

# Liverpool Tier2 Network Status and Plots

- Current Capacity
  - University has 2x20Gbps links out to Janet run as active-active
  - Border routers and firewall operate at 40Gbps
  - Two routers serving Physics datacentre, each with 2x10Gbps links to core
  - But outbound is restricted to 3Gbps as “this aging router has a tendency to drop all routing when busy with high loads, which especially happens with IPv6 traffic.”
- Upgrade planning
  - University to migrate to the new JISC NW network infrastructure between now and May '23 (delayed by pandemic).
  - 2x20Gbps links to Janet will be upgraded to 2x40Gbps links
  - Physics network area to be upgraded in financial year 22/23
  - Existing connection to core will be upgraded to 2x40Gbps

# Liverpool Tier2 Network Status and Plots

- Timeline/blockers
  - Upgrade of core Tier2 cluster network in the next year, to be discussed with local group and University
  - Some indication central IT wants to replace and control data centre switches due to 'Cyber Essentials and new compliance rules'.
  - Specifications, management, and funding may be blockers.
- IPV6 status
  - Subnet allocated for years pending on the upstream upgrades to allow production use
  - Tested on Perfsonar

# The Goal: Network Throughput

- The goal is to transfer a given dataset in as short a time as possible
- A high aggregate transfer rate requires a sufficient number of sufficiently high rate transfers, preferably enough to fill the WAN connection and with a high success rate
- The faster the individual transfers the better

# Monitoring

- How can we monitor these parameters such as throughput and number of active transfers?
- Chain of events:
  - Rucio/DIRAC submit set of transfer jobs to FTS3 service
  - FTS3 optimiser attempts to maximise throughput by increasing the number of transfers on efficient routes and reducing the number transfers on inefficient routes
  - Each individual file transfer throughput recorded
- Examine statistics of transfers

Transfer '1f5328f4-190e-11ed-b89e-fa163ec3b00c' FINISHED

VO: atlas

- Delegation ID: 647795c349808332

Submitted time: 2022-08-11T00:39:56Z

Job finished: 2022-08-11T08:32:23Z

Priority: 3

Bring online: -1

Archive timeout: -1
- Received by fts-atlas-006.cern.ch

Job expires: 2022-08-18T00:39:56Z

Overwrite flag: Y

Job type: N

Cancel flag:


Pin lifetime: -1

Target QoS:

Metadata:

```
{"issuer": "rucio", "multi_sources": false, "auth_method": "certificate"}
```

Total size	Done	Submission time	Start time	Running time	Avg. file throughput	Current job throughput
20.24 GiB	20.24 GiB	2022-08-11T00:39:56Z	2022-08-11T00:49:13Z (+557s)	23524 s	3.70 MiB/s	-



Showing 1 to 4 out of 4




- SUBMITTED
- DELETE
- READY
- STAGING
- ARCHIVING
- ACTIVE
- STARTED
- CANCELED
- FAILED
- 4 FINISHED
- NOT\_USED

File ID	File State	File Size	Throughput	Remaining	Start Time	Finish Time	Staging Start	Staging End	Archiving Start	Archiving End	
+ 4276118072	FINISHED	5.05 GiB	6.98 MiB/s	-	2022-08-11T07:08:33	2022-08-11T07:21:17	-	-	-	-	Log
davs://gridftp.atlas-swt2.org:1094/xrd/datadisk/rucio/mc16_13TeV/8a/39/A0D.17289443._000700.pool.root.1?copy_mode=pull											
davs://fal-pygrid-30.lancs.ac.uk:443/dpm/lancs.ac.uk/home/atlas/atlasdatadisk/rucio/mc16_13TeV/8a/39/A0D.17289443._000700.pool.root.1											
+ 4276118073	FINISHED	5.05 GiB	1.77 MiB/s	-	2022-08-11T07:14:14	2022-08-11T08:03:25	-	-	-	-	Log
davs://gridftp.atlas-swt2.org:1094/xrd/datadisk/rucio/mc16_13TeV/3d/4b/A0D.17289443._001094.pool.root.1?copy_mode=pull											
davs://fal-pygrid-30.lancs.ac.uk:443/dpm/lancs.ac.uk/home/atlas/atlasdatadisk/rucio/mc16_13TeV/3d/4b/A0D.17289443._001094.pool.root.1											



Details for [davs://gridftp.atlas-swt2.org](#) → [davs://fal-pygrid-30.lancs.ac.uk](#)

Storage	Current active transfers	Max. active transfers
davs://gridftp.atlas-swt2.org	304	300
davs://fal-pygrid-30.lancs.ac.uk	9	300



First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (last 1min)	Throughput	EMA	Diff	Explanation
2022-08-11T08:46:52Z	100	9	72	100.00%	18.38 MiB/s	17.26 MiB/s	0	Queue emptying. Hold on.
2022-08-11T08:44:24Z	100	7	76	100.00%	25.73 MiB/s	17.14 MiB/s	0	Queue emptying. Hold on.
2022-08-11T08:41:47Z	100	8	76	100.00%	22.70 MiB/s	16.18 MiB/s	0	Queue emptying. Hold on.
2022-08-11T08:39:02Z	100	9	76	100.00%	20.57 MiB/s	15.46 MiB/s	0	Queue emptying. Hold on.
2022-08-11T08:36:21Z	100	9	76	100.00%	17.19 MiB/s	14.89 MiB/s	0	Queue emptying. Hold on.

Discover

14 hits

New Save Open Share Inspect

Filters 3 Search

KQL

Jun 10, 2022 @ 11:04:00.0 → Jun 10, 2022 @ 11:04:30.0

Refresh

data.t\_final\_transfer\_state: OK × data.throughput exists × data.dst\_site is one of UKI-LT2-IC-HEP, UKI-LT2-QMUL, UKI-NORTHGRID-LANCS-HEP, UKI-NORTHGRID-MAN-HEP, UKI-SCOTGRID-GLASGOW, RAL-LCG2 ×

Add filter

monit\_prod\_fts\_enr\_compl... ▾

Jun 10, 2022 @ 11:04:00.000 - Jun 10, 2022 @ 11:04:30.000 — Auto ▾

- Selected fields
- t data.dst\_site

# data.f\_size

t data.src\_site

# data.throughput
- Available fields
- t \_id

t \_index

# \_score

t \_type

t data.activity

# data.block\_size

# data.buf\_size

t data.channel\_type


# data.chk\_timeout

t data.dest\_srm\_v

t data.dst\_country

t data.dst\_experiment\_site

\* data.dst\_federation



Time	data.src_site	data.f_size	data.throughput	data.dst_site
> Jun 10, 2022 @ 11:04:07.742	CERN-PROD	2,353,626,136	130,757,007.556	UKI-LT2-IC-HEP
> Jun 10, 2022 @ 11:04:01.522	CERN-PROD	2,205,448,515	122,524,917.5	UKI-LT2-IC-HEP
> Jun 10, 2022 @ 11:04:14.909	CERN-PROD	2,186,989,220	80,999,600.741	UKI-LT2-IC-HEP
> Jun 10, 2022 @ 11:04:04.636	RAL-LCG2	3,983,110,477	28,050,073.782	UKI-LT2-QMUL
> Jun 10, 2022 @ 11:04:21.983	MWT2	2,722,560,578	17,341,150.178	UKI-NORTHGRID-LANCS-HEP
> Jun 10, 2022 @ 11:04:27.181	TRIUMF-LCG2	10,640,778,457	16,943,914.74	UKI-NORTHGRID-LANCS-HEP
> Jun 10, 2022 @ 11:04:06.959	TRIUMF-LCG2	10,674,853,634	16,837,308.571	UKI-NORTHGRID-LANCS-HEP
> Jun 10, 2022 @ 11:04:22.162	TRIUMF-LCG2	10,727,982,979	16,632,531.75	UKI-NORTHGRID-LANCS-HEP
> Jun 10, 2022 @ 11:04:26.946	TRIUMF-LCG2	10,472,068,680	16,388,213.897	UKI-NORTHGRID-LANCS-HEP
> Jun 10, 2022 @ 11:04:19.444	TRIUMF-LCG2	10,701,412,186	16,363,015.575	UKI-NORTHGRID-LANCS-HEP

[Save](#) [Share](#) [Inspect](#) [Refresh](#)Filters **4**

Search

KQL



Last 1 month

[Show dates](#)[Refresh](#)

data.throughput exists ×

data.t\_final\_transfer\_state: Ok ×

data.dst\_country: UK ×

data.src\_site: CERN-PROD ×

**data.dst\_site: Descending** ⌵**Average data.throughput** ▼**Max data.throughput** ⌵

UKI-LT2-IC-HEP

120,885,758.057

183,278,544

UKI-LT2-Brunel

99,941,166.737

351,756,096

UKI-SOUTHGRID-BRIS-HEP

74,816,256.016

105,970,360

UKI-NORTHGRID-MAN-HEP

74,159,915.838

119,148,616

UKI-LT2-RHUL

38,352,586.041

110,484,104

UKI-NORTHGRID-LANCS-HEP

35,696,278.273

114,968,312

UKI-SCOTGRID-ECDF

25,999,478.11

111,554,296

UKI-SOUTHGRID-RALPP

25,184,360.95

137,857,952

RAL-LCG2

21,871,003.729

143,209,904

UKI-LT2-QMUL

19,955,966.351

109,235,968

UKI-NORTHGRID-LIV-HEP

10,500,957.637

79,440,360

UKI-SCOTGRID-GLASGOW

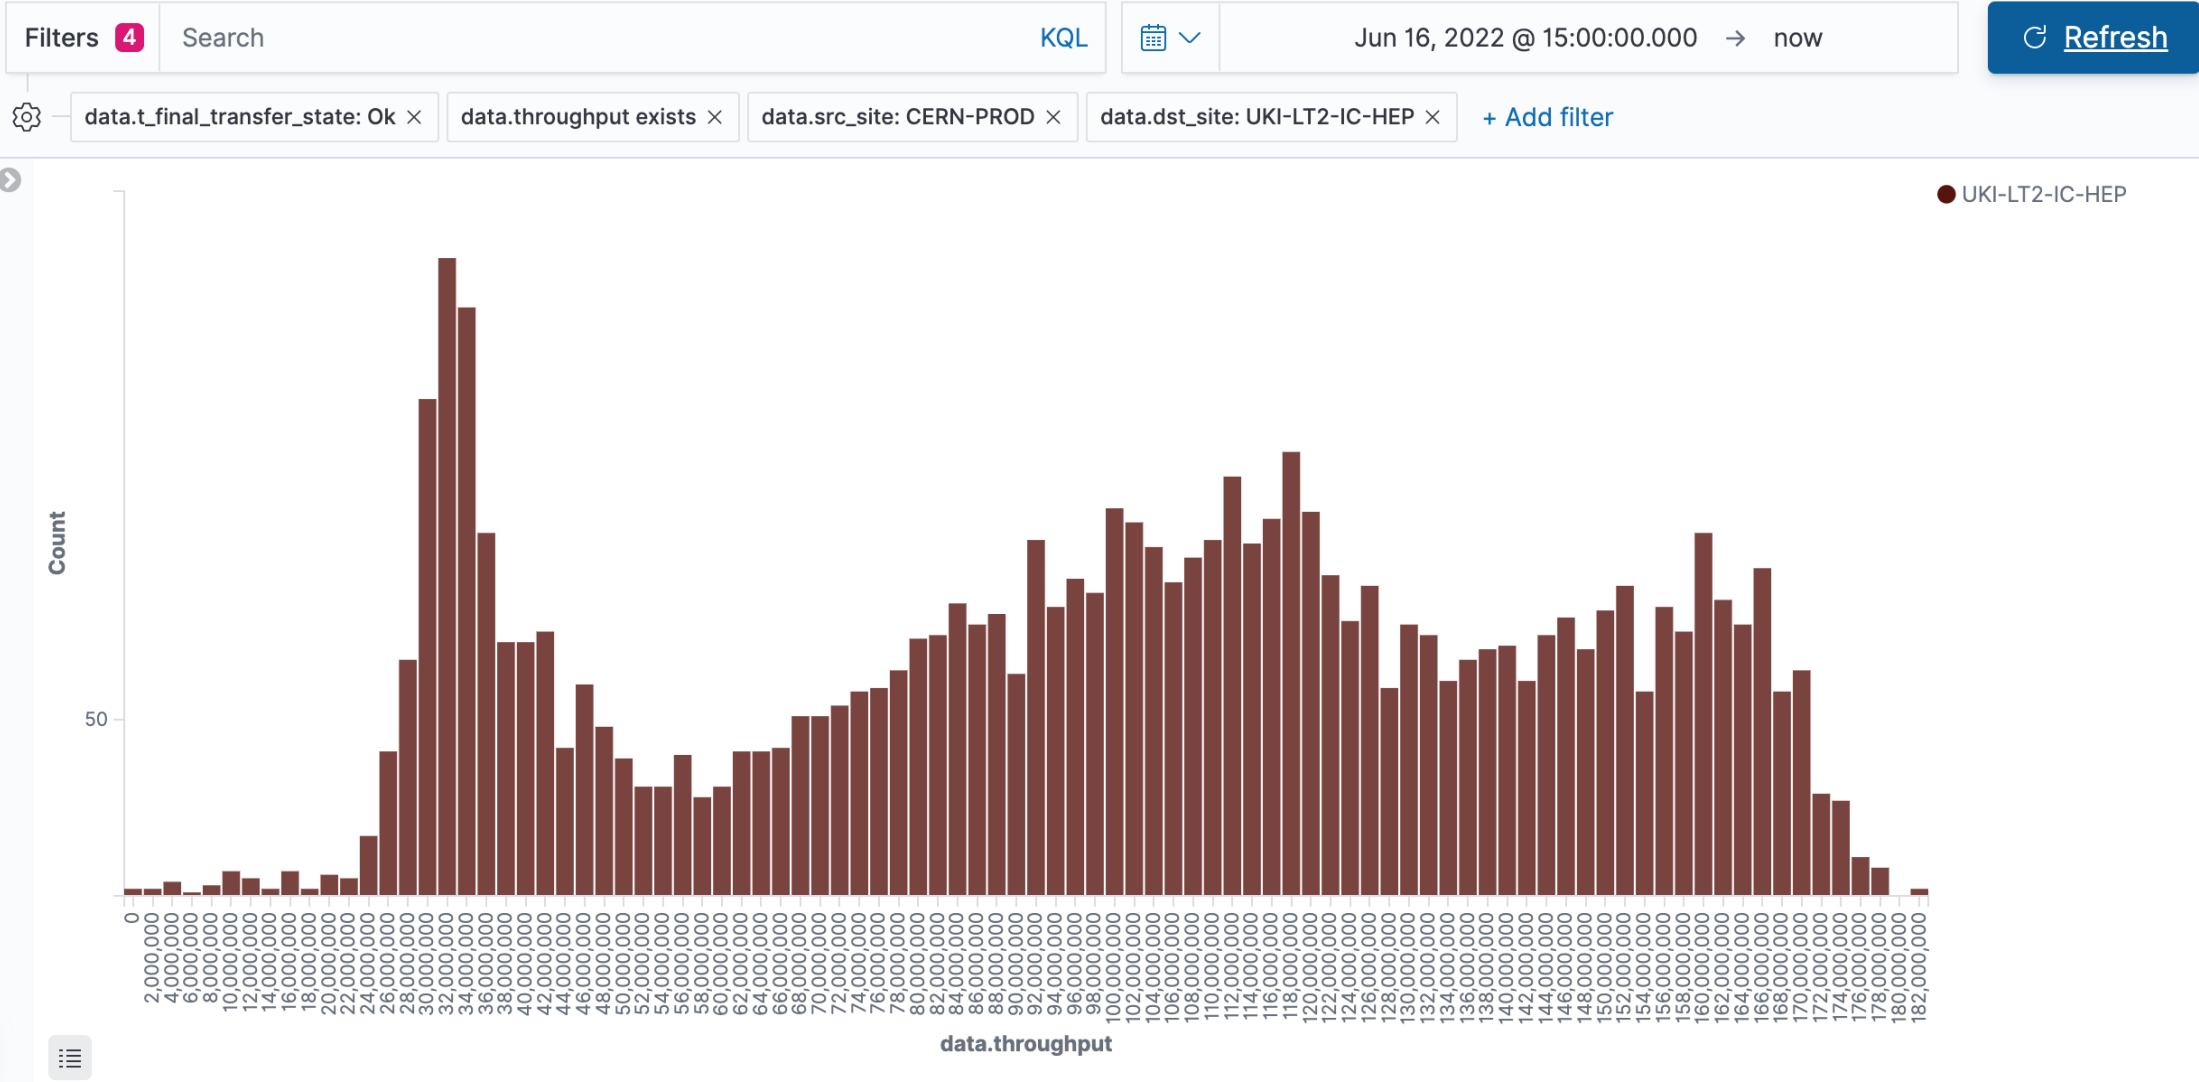
8,223,674.529  
GridPP48

43,664,140

# Larger sites transfer rates

- Imperial – good throughput
- Lancaster has good connection (40G), but new storage element
- Contention at several sites
  - Glasgow – 10G,
  - Manchester – historically 10G
  - QMUL – 10G recently
- Can we see these issues in the individual file transfer rate data?


Individual file transfer throughput rate from CERN to Imperial



Peaks at 32 MB/s, transfers up to ~180 MB/s

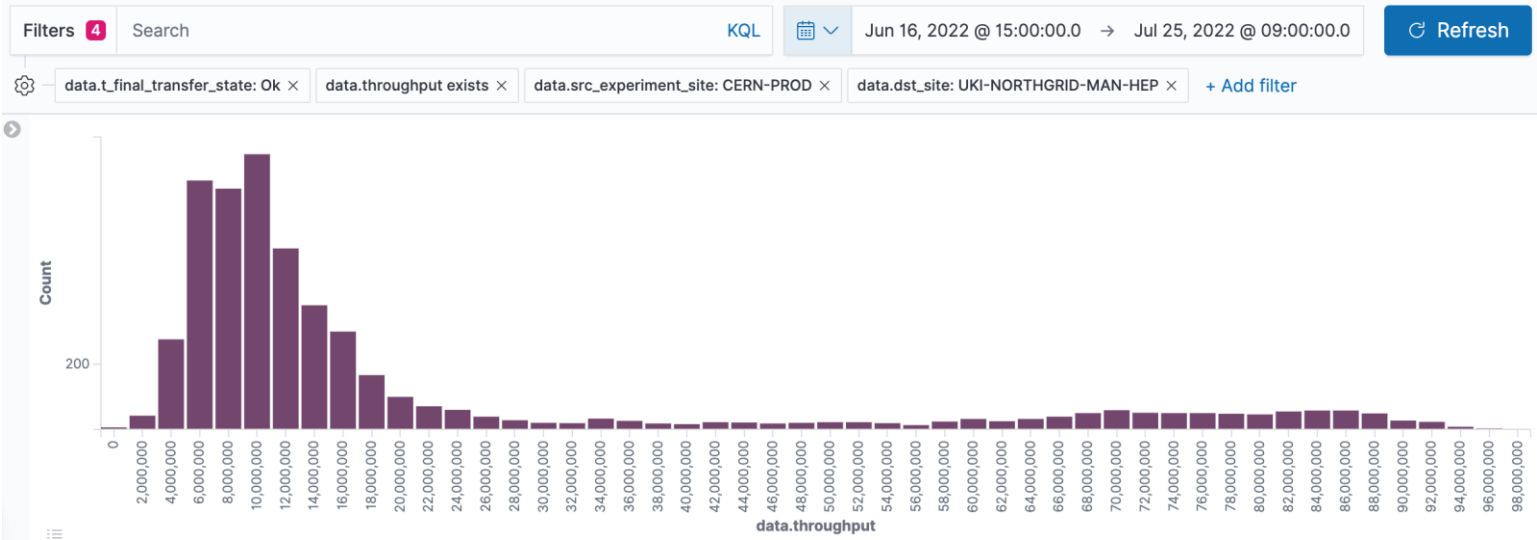
# Individual file transfer rates from CERN to Manchester

Peak at ~8 MB/s, and another smaller one at around 86 MB/s




Individual file transfer rates from CERN to Manchester


10 Gb/s link (peak at 10 MB/s) - contention?



40 Gb/s link (peak at 86 MB/s)



# Individual file transfer rates from CERN to Lancaster




Peak at ~6 MB/s, and another smaller one at around 84 MB/s

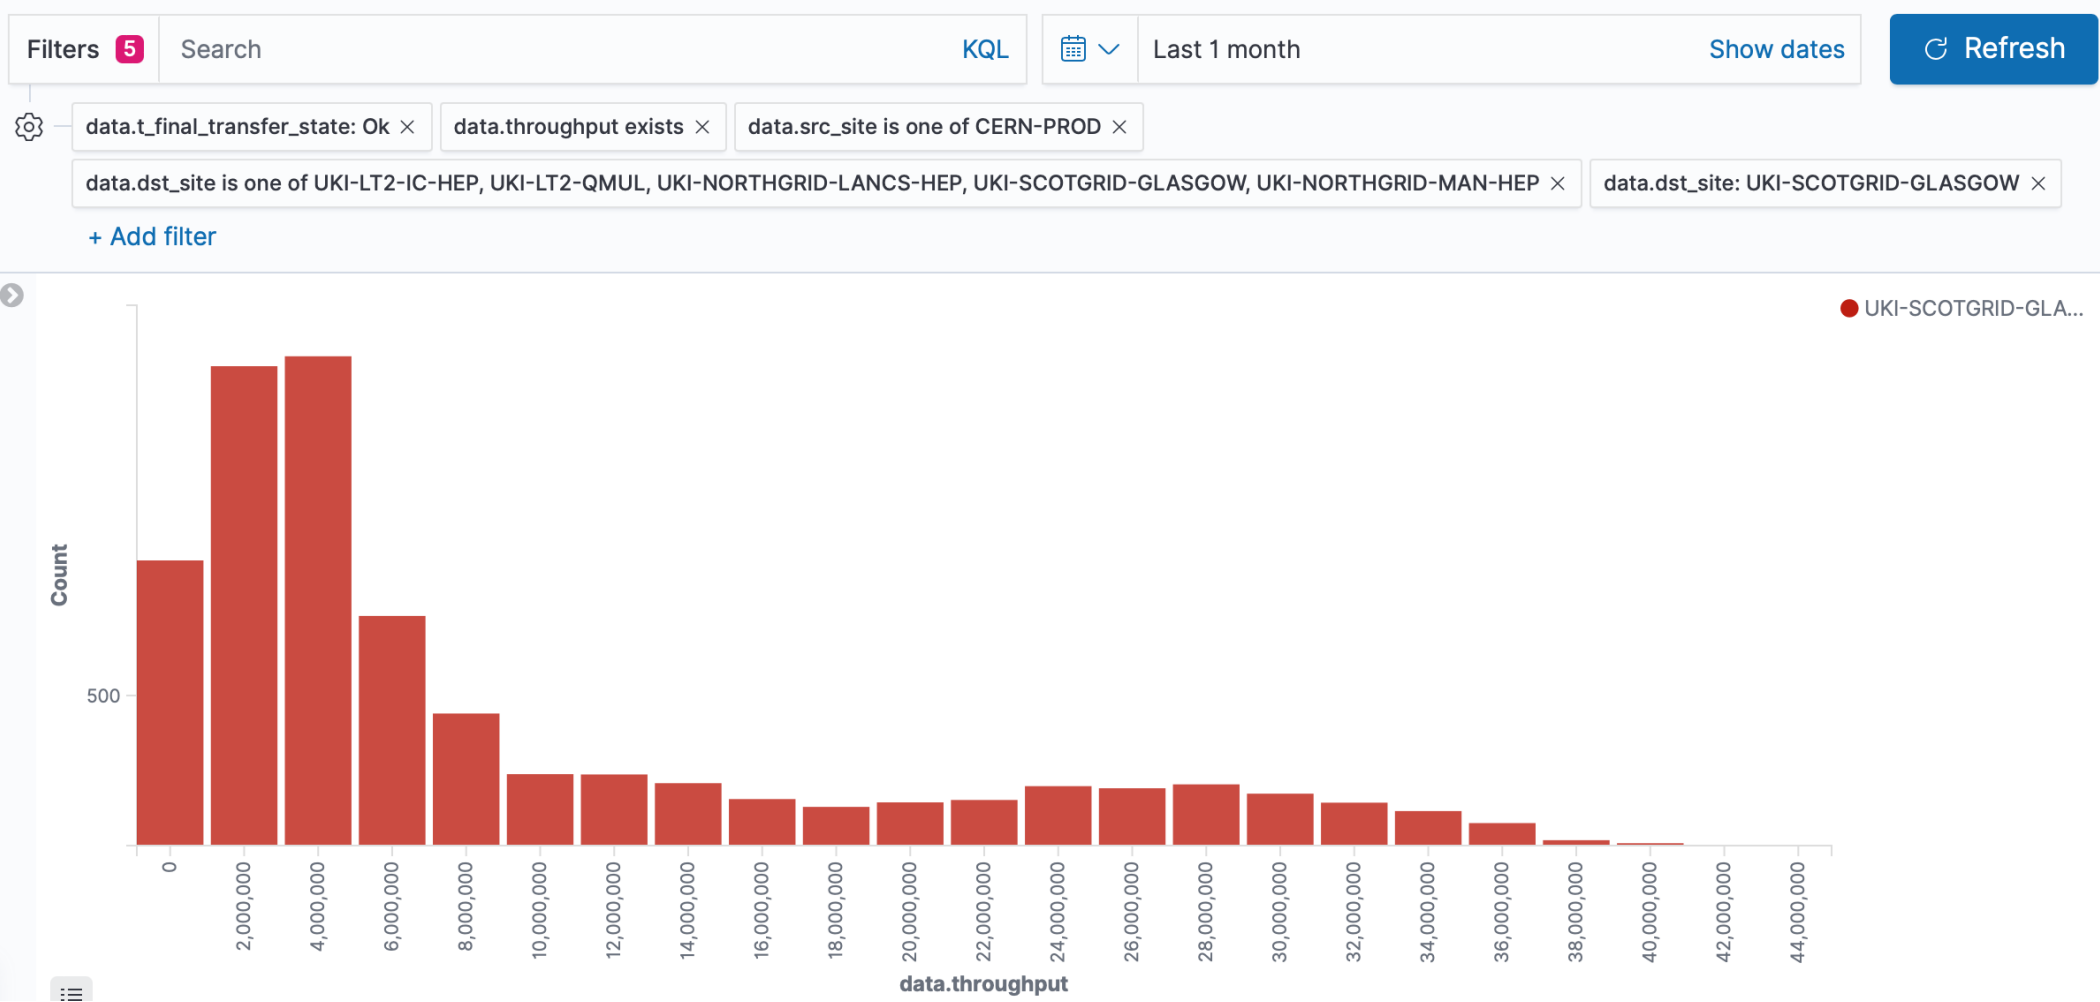


# Individual file transfer rates from CERN to QMUL

Peak at ~2 MB/s




# Individual file transfer rates from CERN to Glasgow




Peak at ~4 MB/s

# Individual file transfer rates from CERN to RALPPD

Peak at ~4 MB/s



Try to get sites up to  
40 Gb/s or 100 Gb/s



# Summary

- Need to ensure the GridPP sites are able to transfer data at sufficiently high rates
- Several GridPP sites have already had their connection to the Janet network upgraded
- Some are waiting for the campus network to be modified to provide GridPP sites with higher bandwidth
- But many campus network teams are overloaded thereby delaying campus network improvements
- Other sites are not achieving full use of their existing connection
  - Look at site network details and estimates of FTS flows
  - Individual file transfer rates can help understand issues
- Jisc Network Performance group are trying to help by strengthening the community of “Research Network Infrastructure Engineers”

# Site contributions

# Oxford GridPP Tier-2 Network Status

Pete Gronbech

May 2022

# Oxford Tier-2 Network Status

## GRID Cluster

- The GRID cluster is currently connected to the backbone at 10Gbit.
- 2\*10Gbit connection upgrade is paid for and it will be enabled in August 2022.

## Connection to Janet

- The University is currently connected to Janet at 2\*20Gbit.
- The University Janet connection upgrade to 2\*100Gbit is in place and will be configured into production in July 2022.

## IPV6

- The Central IT Services are in the early stages of a project to roll out a production IPV6 service across the University. They have applied for funding from the University.




# Snapshot of Grid Cluster traffic

## 9<sup>th</sup> May 2022


### Ten-GigabitEthernet1/0/1 (PHYSICS) on dcdist-beg

[Back to frodo](#)

#### The last hour




#### The last day




01/09/2022

#### The last week



#### The last month

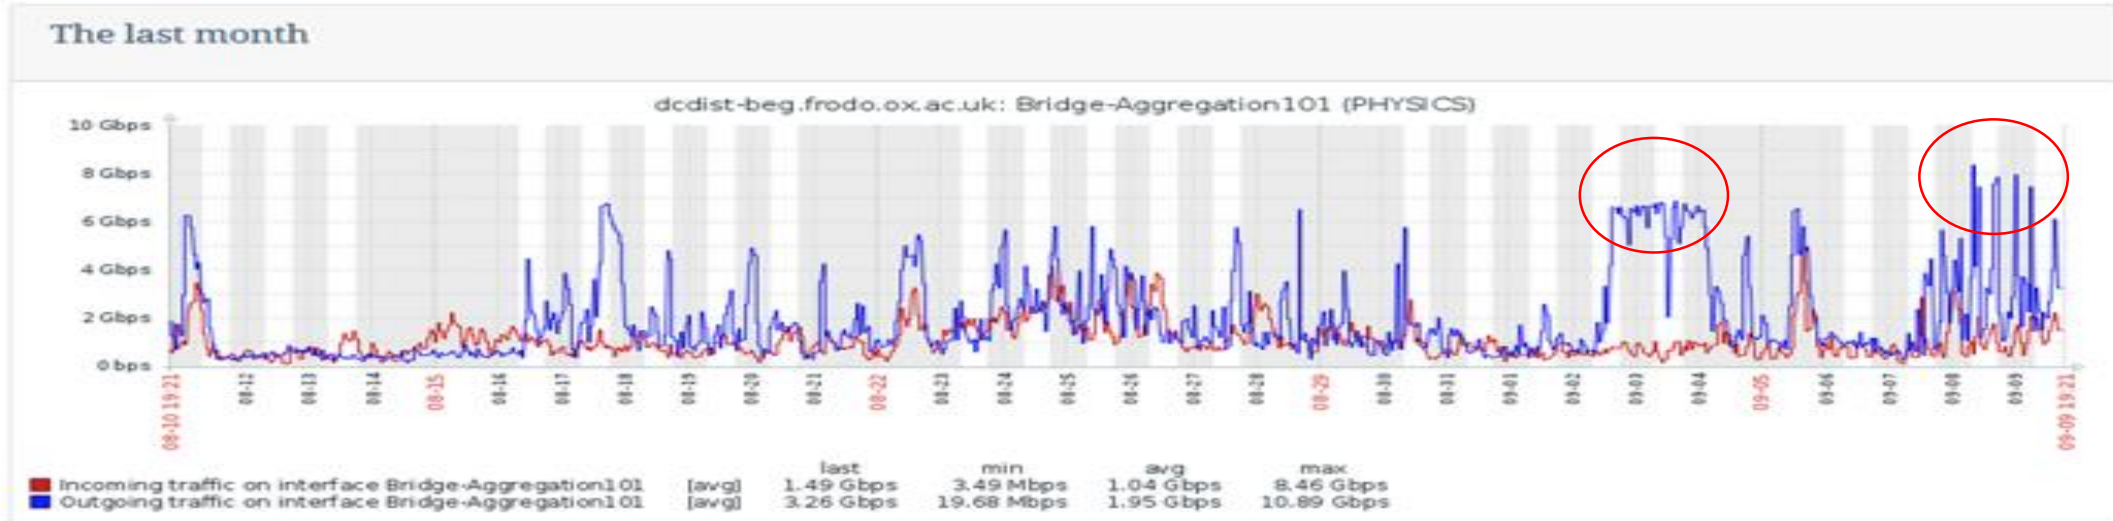


GridPP48

57


# Previously tested dual 10Gbit from the cluster in September 2021

- The Grid traffic bursts shown below were seen contributing to load peaks on the 3<sup>rd</sup> and 9<sup>th</sup> September on the overall university traffic graph.
- The average University total traffic peaked at ~10Gbit during that time.




# Grid Cluster connected at 2\*10Gbit ~16:55 on 9<sup>th</sup> September 2021 (for 1 month test)

- Can see spikes above 10Gbit from the Grid cluster on the 10<sup>th</sup> September.



- More spikes up to 18.4Gb/s on 15<sup>th</sup> September. Unfortunately, the Xcache server was not being utilised between 7<sup>th</sup> and 18<sup>th</sup> September.



# Summary

- Current average Grid traffic is perhaps lower than back in September 2021.
- Expect traffic to increase when a loan Xrootd xcache server from RAL is installed.
- Expect newer network switches and worker node later in the year.
- Upgrades to University JANET connection to 2\*100Gbit and 2\*10Gbit connection to cluster in few months.





- Current status

We currently have a 2x10Gb/s link direct to the JANET router. At present, this is not available to the site as I have bonded the link from the gateway to the switch. Internally, it's almost all 10Gb/s ethernet

- Upgrade planning if any

None planned at present as it doesn't seem necessary given traffic rates

- Any existing agreements between JISC and site authorities

The grid site is outside the Bham firewall and everyone seems happy with this arrangement

- Expected timeline/blockers...


N/A

- IPv6 status

Stalled. The gateway can route packets but after sometime trying to get DNS entries out of the university, I got told they aren't using it yet and don't have plans to in the near future (!). I will try to resurrect this in the coming months




# Traffic Plots (12h)






# Traffic Plots (8w)



node\_network\_rate\_in(instance="epggw1.ph.bham.ac.uk:9100")




node\_network\_rate\_out(instance="epggw1.ph.bham.ac.uk:9100")

# UKI-SOUTHGRID-BRIS-HEP

- Current status:
  - Up to 20 Gbit/s, IPv4 + IPv6
  - Some issues observed with perfsonar → looking into it (takes time, low on FTEs)
- Upgrade planning
  - Thinking about 25 Gbit/s on NIC level
  - Some rumours about University of Bristol going to 100 Gbit/s
- Any existing agreements between JISC and UoB
  - Query sent to IT Networks, those within GridPP with connections to JISC probably also know more than site admins
- Expected timeline/blockers: staff oversubscribed



[Hide Legend](#) | [Show Previous](#) | [Show RRD Command](#) | To show trend, set to future date



#### Source

dice-io-37-00.acrc.bris.ac.uk  
2001:630:e4:2810:137:222:79:1  
[Host info](#) ▾

#### Destination

t2ps-bandwidth2.physics.ox.ac.uk  
2001:630:441:905::b  
[Host info](#) ▾

#### Report range

← 1 year →

2021-05-19 11:45

2022-05-19 11:45


Submit

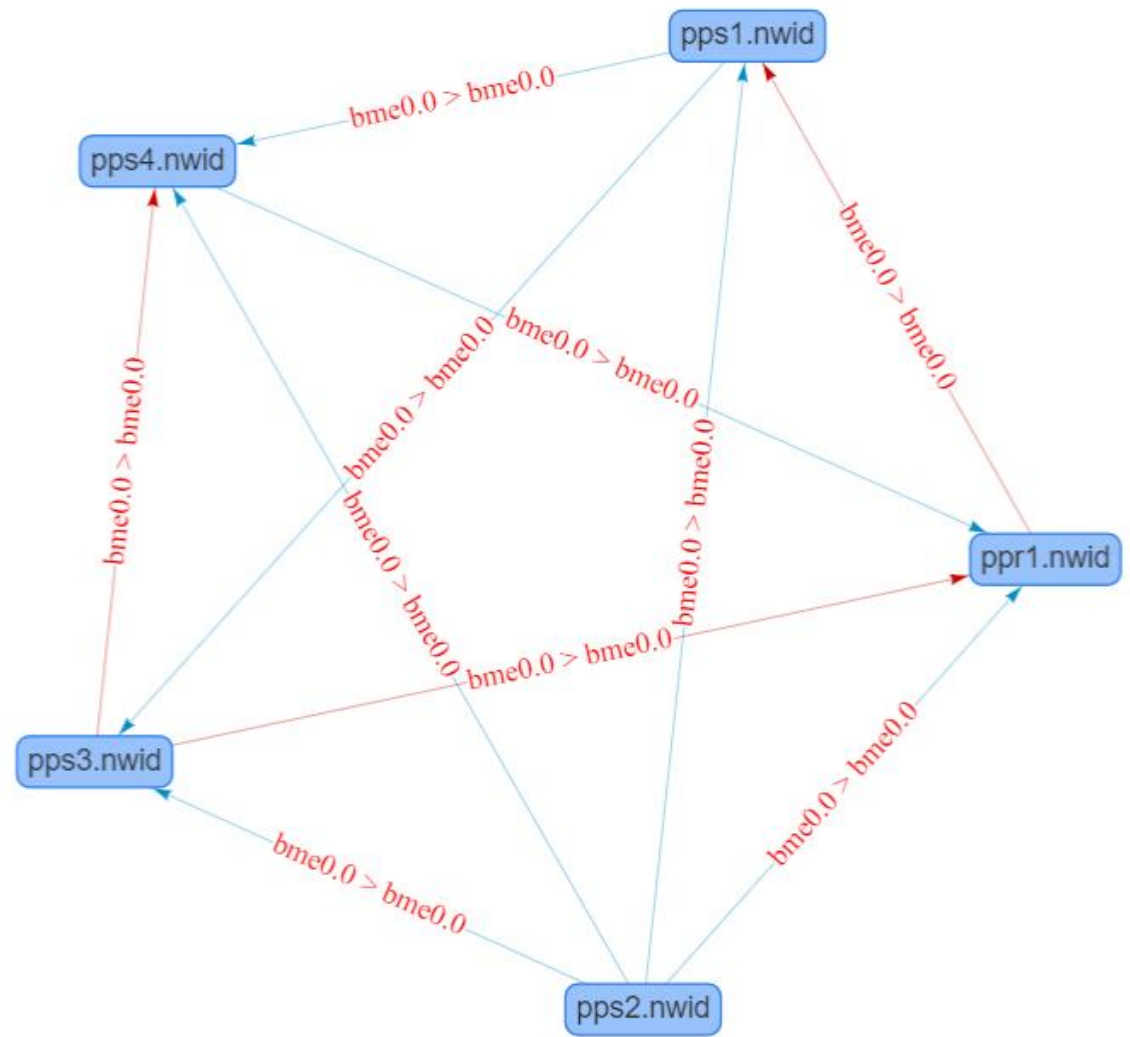
Wed, 19 May 2021 10:45:17 GMT to Thu, 19 May 2022 10:45:17 GMT

Show/hide chart rows ☒ Throughput ☒ Packet Loss ☒ Latency ☒ Application Response Time

Tput (TCP) Tput (UDP) Loss (UDP) Loss (one way) Loss (rtt) Retrans Latency (one way) Latency (rtt) DNS HTTP

Settings Forward Reverse Failures





# Liverpool Tier2 Network Status and Plots

- Current Capacity
  - University has 2x20Gbps links out to Janet run as active-active
  - Border routers and firewall operate at 40Gbps
  - Two routers serving Physics datacentre, each with 2x10Gbps links to core
  - But outbound is restricted to 3Gbps as “this aging router has a tendency to drop all routing when busy with high loads, which especially happens with IPv6 traffic.”
- Upgrade planning
  - University to migrate to the new JISC NW network infrastructure between now and May '23 (delayed by pandemic).
  - 2x20Gbps links to Janet will be upgraded to 2x40Gbps links
  - Physics network area to be upgraded in financial year 22/23
  - Existing connection to core will be upgraded to 2x40Gbps

# Liverpool Tier2 Network Status and Plots

- Timeline/blockers
  - Upgrade of core Tier2 cluster network in the next year, to be discussed with local group and University
  - Some indication central IT wants to replace and control data centre switches due to 'Cyber Essentials and new compliance rules'.
  - Specifications, management, and funding may be blockers.
- IPV6 status
  - Subnet allocated for years pending on the upstream upgrades to allow production use
  - Tested on Perfsonar

WAN throughput  
to University Core  
Network (20G  
uplink)

