

Data transfer on COSMA

Alastair Basden, George Beckett, James Perry,
Jonathan Hollocombe, Kieran Leach, Mark
Lovell, Shaun de-Witt, Paul Walker, Billy Carlyle

Durham NFCS+ data sharing workshop

COSMA

- Includes:
 - The DiRAC Memory Intensive Service
 - 8th generation: COSMA8
 - HPC Hardware Lab
 - Including the DINE/DINE2 clusters
 - Many GPU/CPU/network/storage/tech testbeds
- Est 2001

Data on COSMA

- COSMA is a data generator
 - Principally large cosmology simulations
- This data needs sharing:
 - Transferring to other sites
 - Available for analysis and download
- COSMA is a data consumer
 - Hosting a tape archive for DiRAC
 - Enabling researchers to transfer data onto it

Data transfer setup

- We follow a flexible approach
 - If someone asks for something, we try to provide
 - Databases
 - Web hosting/downloads
 - Data streaming capabilities
 - Data transfer tools
 - Globus
 - Rucio end point
 - S3 storage
 - StorJ
 - Batch queue access for post processing
 - Scratch and long-term staging locations

The cosmological dataset

- Cosmology datasets are large
 - O(PB)
 - Not easy to transfer or host
- The VirgoDatabase:
 - An online portal for extracting subsets of data
 - Highly effective, but old

EAGLE Database

Documentation

Welcome Database user.
Streaming queries return unlimited number of rows in CSV format and are cancelled after 1800 seconds.
Browser queries return maximum of 1000 rows in HTML format and are cancelled after 90 seconds.

NEW: EAGLE particle data can be downloaded via [this page](#) (log in with your EAGLE database account)
See the documentation paper for details.

Query (stream)
Query (browser)
Help

Maximum number of rows to return to the query form: 10

Demo queries: click a button and the query will show in the query window.
Holding the mouse over the button will give a short explanation of the goal of the query. These queries are also available on this page.

Magnitudes:

Images:

Subhalo:

Dust:

Metadata queries: The SQL statements under these buttons provide examples for querying and managing the state of a private database.
Holding the mouse over the button will give a short explanation of the goal of the statement.

Public Databases
- DIMONLY
- Pfiducial_models
- Physics_vars

Private (MyDB) Databases
- nqc550_db (rw)

ICC
VIRGO
GAVO
GERMAN ASTRONOMICAL
VIRTUAL OBSERVATORY

HDF5 Streaming Service

- Remote access to HDF5 data
 - Recursive downloads of full directories
 - Individual files
 - HDF5 groups, datasets, attributes within a file
 - Streaming via the messagepack format
 - Easy to use Python library interface
 - With a similar interface to h5py
 - But for remote files
 - Dataset slices as numpy arrays

Use cases

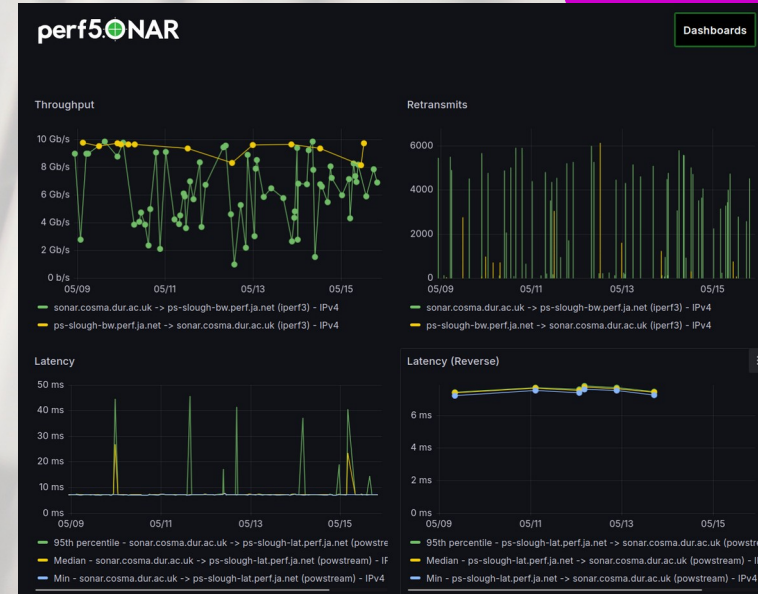
- Users wanting to analyse large datasets on facilities they have access to
 - Without needing a full data dump

Fast data transfer


- **bbcp:**
 - Now deprecated in many places
 - More complex to use than some tools
 - Fast and optimisable
 - Command line
 - Requires accounts at both ends

Monitoring


- Part of the JISC PerfSonar
- And login node/DTN monitoring



Average Bytes Received Per Second



Average Bytes Transmitted Per Second



Globus

- Free version
- Single node
- Coupled with system LDAP
- MTU 9000 on data transfer node

Ad-hoc transfers

- Usually researchers needing to put data on COSMA
 - Due to closures of other DiRAC services
 - Or lack of quota at other sites
 - Globus or “Rsync at night approach”

StorJ on-premise cloud storage

- A self-hosted system
 - Edinburgh, Leicester, Durham
 - Looking to expand
- Offers S3 interface
- And StorJ native protocol
- Makes use of spare/old disk
 - High erasure coding to cope with failure
- Web interface, rclone, aws, etc
 - Native protocol can make use of aggregate data bandwidth from all sites
 - S3 protocol goes through a gateway at Durham

The COSMA logo is displayed on a server rack. It features the word "COSMA" in a bold, white, sans-serif font, with a stylized white arc above the letters "O" and "S" that resembles a star or a signal. The background of the server rack is dark, and the logo is illuminated.

View your project statistics, check daily usage, and set project limits.

Team Info

All team members should use the same passphrase to access the same data.

Objects

Project total

720

Buckets

In this project

2

Access Keys

Total keys

3

Team

Project members

2

CO₂ Estimated

For this project

0 kg CO₂e

CO₂ Avoided

By using Storj

0 kg CO₂e

Storage

311.25MB Used

Limit: 25TB

25TB Available

[Set Storage Limit](#)

431.56MB Used

Limit: 9%

1000TB Available

[Set Download Limit](#)

Segments

736 Used

Limit: 10,000,000

9,999,264 Available

[Learn more](#)

Buckets

2 Used

Limit: 100

98 Available

[Learn more](#)

Storage Buckets

Buckets are where you upload and organize your data.

[New Bucket](#)

Search

Bucket	Objects	Segments	Storage	Download	Versioning	Lock	Date Created	Created By
backup-production	360	376	194.53MB	0.23GB	No	on	Jun 20, 2025	a.g.basden@durham.ac.uk
backup-studio	280	280	116.79MB	0.14GB	No	off	Jun 20, 2025	a.g.basden@durham.ac.uk



- Users
- Invites
- Projects
- Buckets
- Nodes
- ADMIN CONSOLE
- Admin Users
- Access Grant Inspector
- Audit Logs

Listing Users

Home Users

Paid (3) | Nil (0) | Free (7) | Frozen (0) | Legal Hold (0) | Pending Deletion (0) | Deleted (0) | Inactive (2)

Normalized email	Projects	Project bandwidth limit	Project storage limit	Project segment limit	Created at	Kind	Project limit	Full name	User as
OPS+DIRAC@STORJ.IO	1	25 GB	25 GB	10,000	11th Nov 2024 10:58 PM	Paid	1	Storj Dirac	None
ETHAN@STORJ.IO	1	25 GB	25 GB	10,000	11th Nov 2024 5:35 PM	Paid	1	Ethan Adams	None
A.G.BASDEN 97@CANTAB.NET	1	25 GB	25 GB	10,000	12th May 2026 2:58 PM	Free	10	Alastair Basden	None
BOHDAN@INTEGRATED-DIGITAL.COM	1	200 GB	200 GB	10,000	18th Feb 2025 2:30 PM	Free	3	bohdan	None
RYAN07BIGNELL+STORJ@GMAIL.COM	0	25 GB	25 GB	10,000	23rd Feb 2026 4:30 AM	Free	1		None
AKOS@INTEGRATED-DIGITAL.COM	1	25 GB	25 GB	10,000	18th Feb 2025 2:43 PM	Free	1	Akos Gyimesi	None
STUART.MCALPINE@FYSIK.SU.SE	1	25 GB	25 GB	10,000	8th Jul 2025 7:53 AM	Free	2	Stuart McAlpine	None
A.G.BASDEN@DURHAM.AC.UK	18	1,000,000 GB	25,000 GB	10,000,000	7th May 2025 9:06 AM	Paid	24	Alastair Basden	None
ETHANJADAMS@GMAIL.COM	0	25 GB	25 GB	10,000	18th Feb 2025 2:26 PM	Free	1		None
TEST@STORJ.IO	1	150 GB	150 GB	0	11th Nov 2024 4:16 PM	Free	5	Hiro Protagonist	None

Displaying Users 1 - 10 of 14

Browse

Buckets > backup-production > backups

[Upload](#) [New Folder](#) [Refresh](#) [Filter](#)

Search

Name	Type	Size	Date
default-20260416000029	Folder		
default-20260417000030	Folder		
default-20260418000031	Folder		
default-20260419000032	Folder		
default-20260420000033	Folder		
default-20260421000034	Folder		
default-20260422000035	Folder		
default-20260423000036	Folder		
default-20260424000037	Folder		
default-20260425000038	Folder		
default-20260426000039	Folder		

Items per page: 10 1-10 of 29

Federation of data transfer

- A single Office365 login?
- Multiple accounts across multiple sites?
- Globus seems to be highly regarded
- Bespoke solutions are appropriate in places

RUCIO end point

- Set up using nginx and webdav
- And using the LSST RUCIO server
 - Tests ongoing

Zero-copy data movement


- How can data be made accessible without needing to transfer the whole dataset?
 - The HDF5 streaming library is a good example
 - S3 export of Lustre is another possibility
 - Easy shared access to large datasets
 - Easier for read than for write

NFCS+ outputs

- <https://durham.readthedocs.io/en/latest/datafed>
-

Summary

- Nothing particularly clever
 - Mostly user-driven
 - And set up to match user requirements
 - We take a flexible approach



Alastair Basden
DiRAC / Durham University

A photograph of a server rack with a purple geometric overlay in the top right corner. The server rack is composed of multiple vertical units, each with a circular ventilation grille and a small rectangular display or indicator light. The image is slightly blurred, and the purple overlay is a solid, vibrant color.

Alastair Basden
DiRAC / Durham University