Object Storage: How chunky would you like your data?

Matt Brown*, Amulya Chevuturi, Matt Fry

*Contact: matbro@ceh.ac.uk

UK Centre for Ecology & Hydrology

What is Object Storage?

- A new type of data storage on the cloud
- Unlike traditional file-systems: not hierarchical and accessed over HTTP
- A ‘tenancy’ (disk) contains ‘buckets’ (folders) which contain ‘objects’ (files/chunks of files)

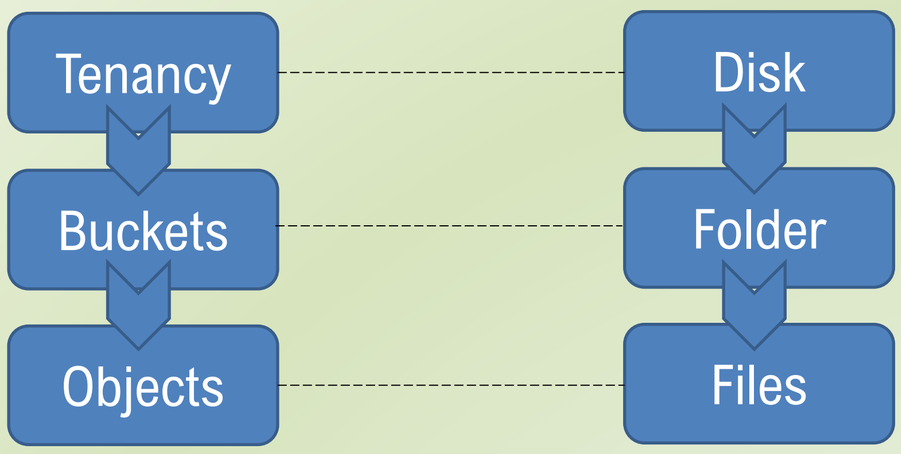

Why is it needed?

- Far cheaper per byte than disk
- Faster and better support of parallel processing
- Important for the next generation of multi-petabyte datasets, models and digital twins
- Easier for sharing externally

What did we do and why?

- With object storage you can only download one whole object at a time
- This is very inefficient for large datasets
- Instead you can divide the ‘objects’ into ‘chunks’ of each file using the **zarr** file format

But how best to do this?
What dimensions to chunk?
What size to make the chunks?

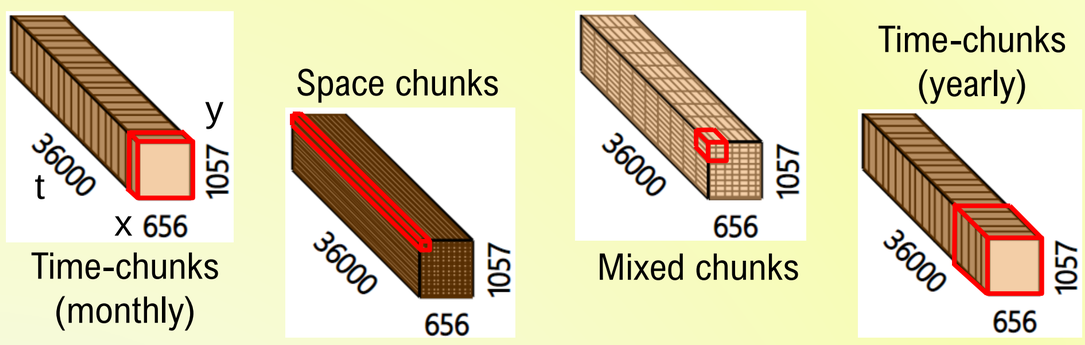

- How you chunk and the type of data analysis you are doing strongly affects the performance of the object store
- We attempted to quantify this and provide demonstration notebooks for accessing and analysing datasets on object storage

Key message: At least a 10× slow-down in performance with ‘bad’ chunking

- Chunking can be ‘bad’ with an analysis:
  - At every timestep (left) on data **heavily chunked in time** **OR**
  - At every grid-point (right) with data **heavily chunked in space**
- Because reading the whole matrix for just one grid-point of interest is inefficient!
- Chunk byte-size also impacts: Many small chunks mean large overheads; **few large chunks** mean unnecessary data load. Network bandwidth also plays a role.
- **Mixed chunking** reduces the impacts: All analyses have ~equal but average speed

Grid-box analysis at each timestep

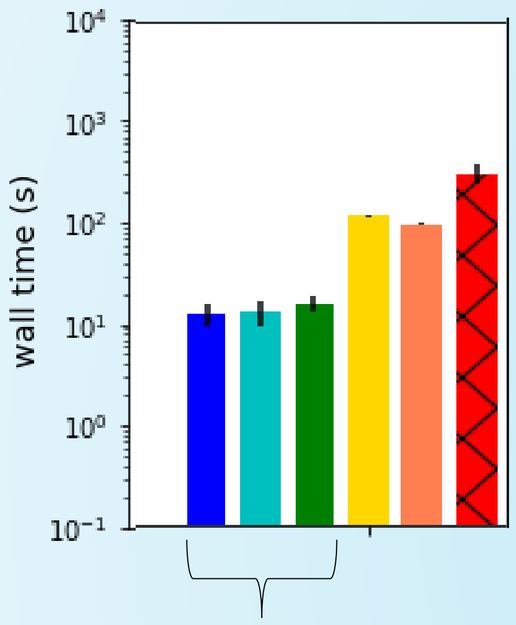

Spatial & mixed chunking

10km spatial chunking (14MB)
100km spatial & yearly (14MB)
100km spatial & 10yr (140MB)

Time-period analysis at each grid-point

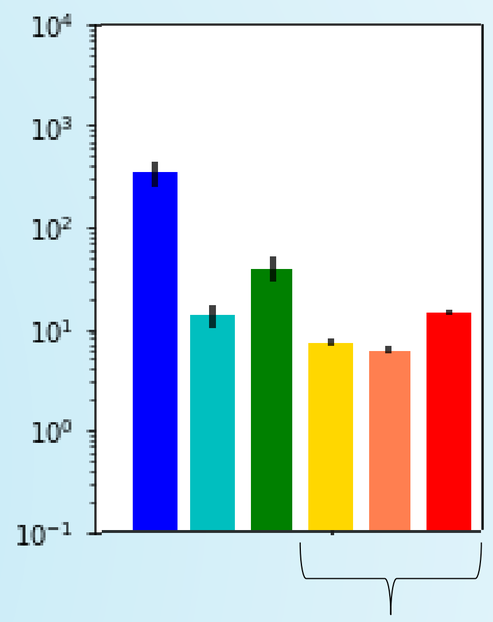

Time chunking

Seasonal chunking (240MB)
Yearly chunking (930MB)
Monthly chunking (79MB)

Hydro-JULES

DataLabs

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

UKRI Natural Environment Research Council