


Science and
Technology
Facilities Council

Natural
Environment
Research Council

```
diff.plot()
```

```
[79]: <matplotlib.collections.QuadMesh at 0x7fa27b744a58>
```







Cloud futures for CMIP data – evaluating object storage models and re-evaluating federation for data distribution

AGU, Session IN032-02, 14 December 2020

Ag Stephens, Philip Kershaw, Alan Iwi, Matthew Jones, Bryan Lawrence, Neil Massey, Ruth Petrie, Matt Pryor, Eleanor Smith

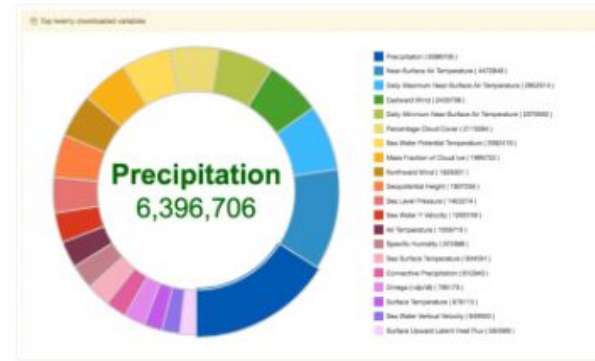
Earth System Grid Federation: 10+ years of operations



ESGF Federation



Data usage

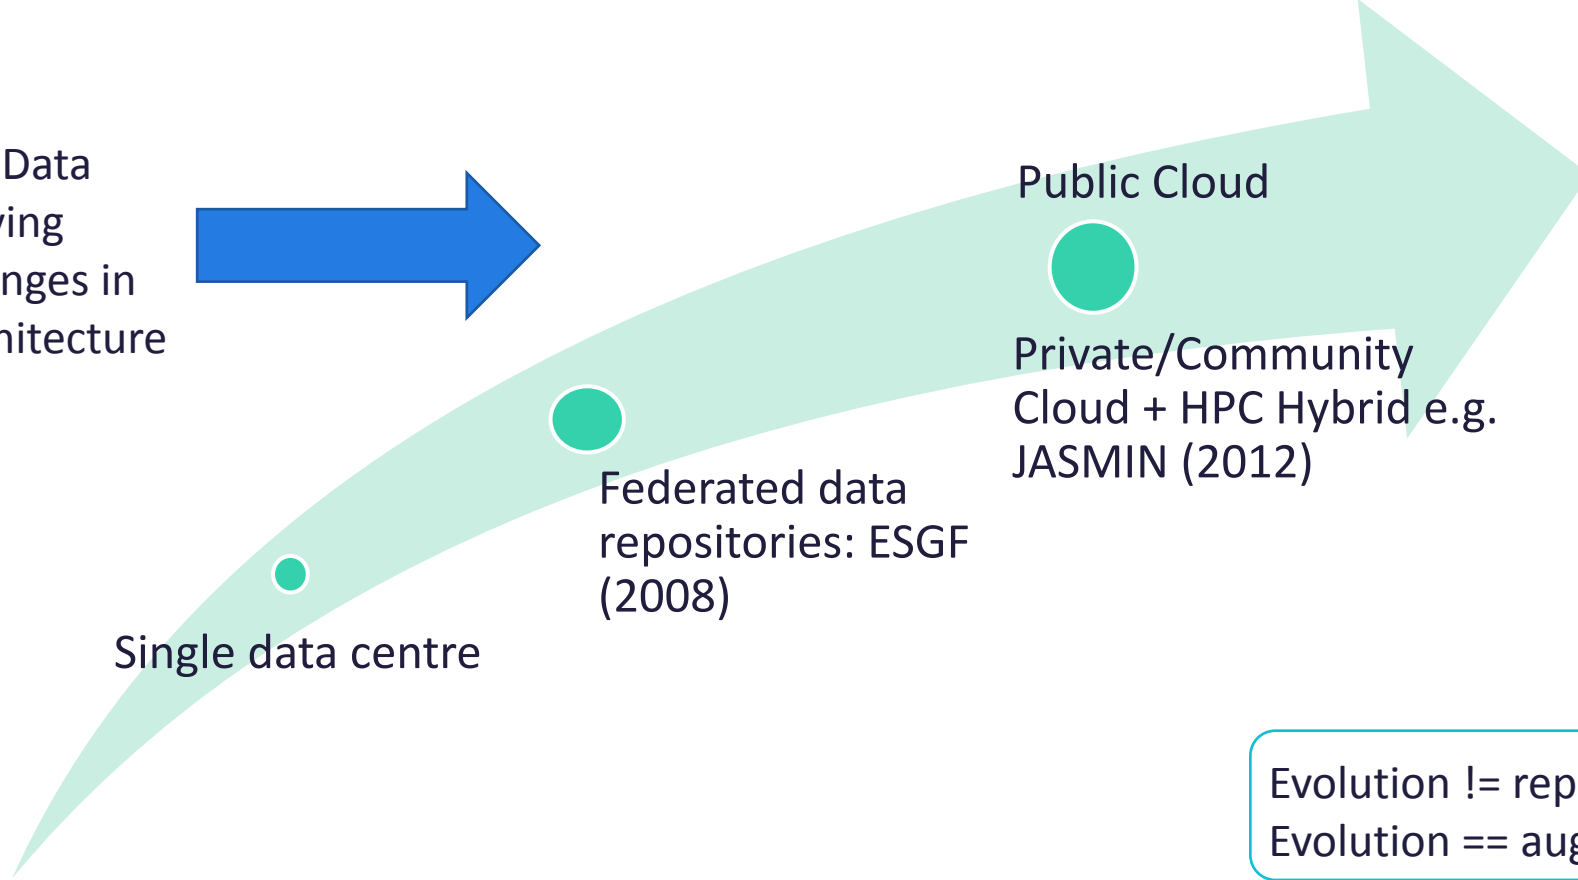


Data publication



Evolution to regional clusters aggregating data + co-locating it alongside processing capability

Big Data driving changes in architecture

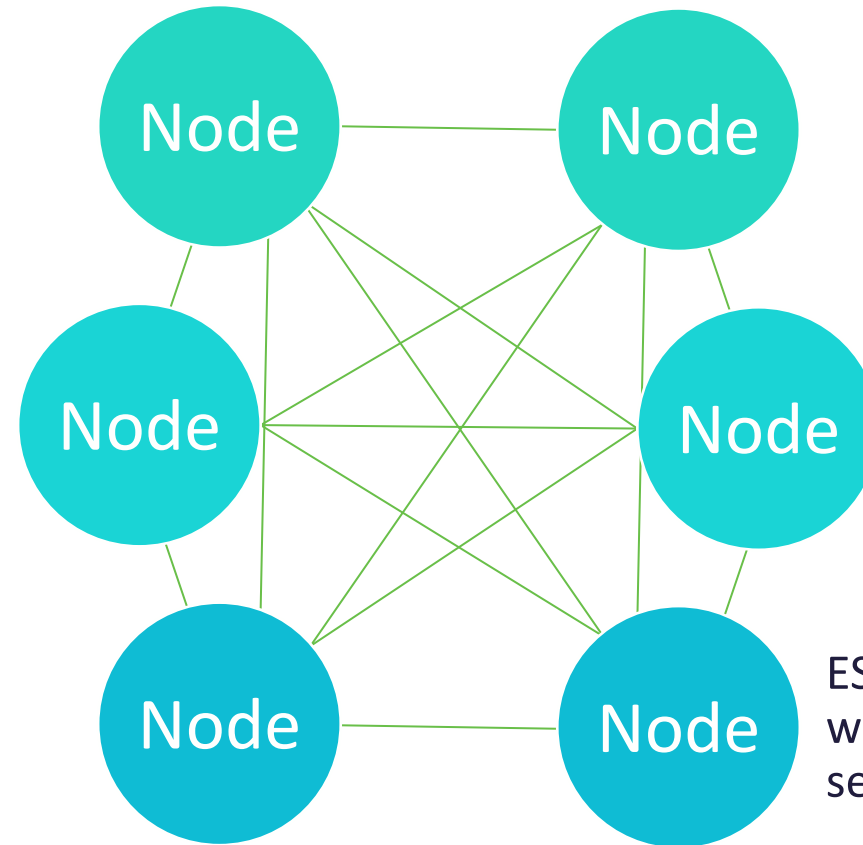


Data Analysis Platforms

- Co-located compute and data
- Analysis Ready Data (ARD)

Evolution != replacement
Evolution == augmentation + replacement

Federation enables scaling and resilience



ESGF Nodes at sites around the world deliver data and associated services

Public Cloud has in-built capabilities for scaling and resilience




Aggregate data and services into the cloud

There are technical, policy and financial challenges and opportunities for wider adoption of public cloud



Object Storage

Public cloud popularised object storage as a convenient alternative to POSIX storage.




Large-Volume Scientific Data

Overall costs for hosting large volumes of scientific data on public cloud remains high when compared to on-premise hosting.



Science and
Technology
Facilities Council

Natural
Environment
Research Council



Centre for Environmental
Data Analysis

SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL



National Centre for
Atmospheric Science

NATURAL ENVIRONMENT RESEARCH COUNCIL

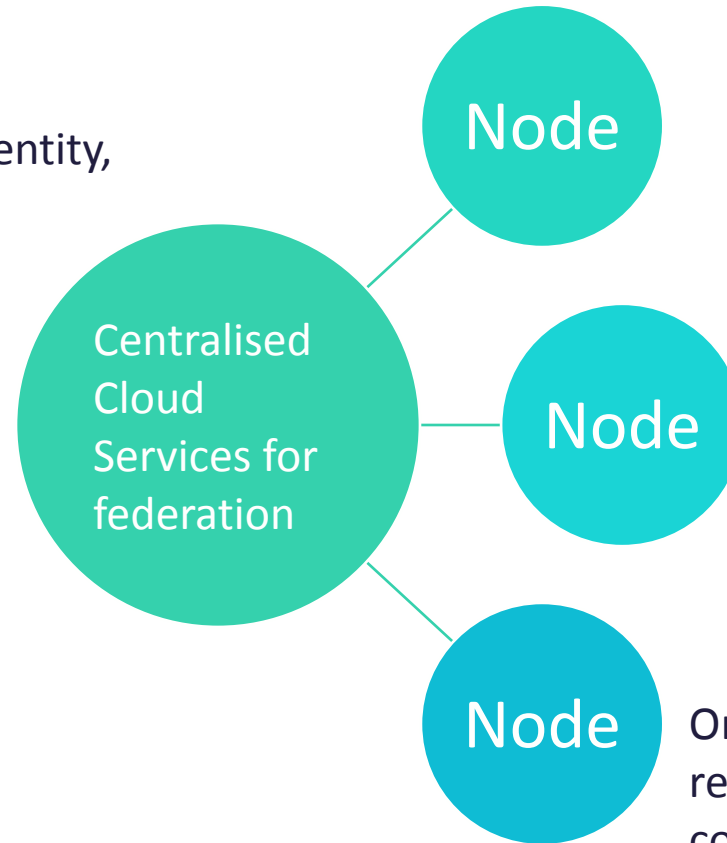


National Centre for
Earth Observation

NATURAL ENVIRONMENT RESEARCH COUNCIL

Federation and Cloud can augment and complement one another


Federation glue: search, identity, metrics, monitoring public cloud-hosted



On-premise nodes for big data serving – research community labs, private and community clouds, computing centres



Science and Technology Facilities Council
Natural Environment Research Council



Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL



National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL




National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Build a cloud-compatible s/w stack to maximise portability between public cloud and on-premise

Public (commercial) cloud



technical agility for migration



On-premise:
private/community cloud



JASMIN

[Example of National Facility
for the research community]



Kubernetes




Object Store

DevOps, Infrastructure-as-Code approach



Science and
Technology
Facilities Council

Natural
Environment
Research Council



Centre for Environmental
Data Analysis

SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL



National Centre for
Atmospheric Science

NATURAL ENVIRONMENT RESEARCH COUNCIL



National Centre for
Earth Observation

NATURAL ENVIRONMENT RESEARCH COUNCIL


Develop a Strategy for CMIP Data with Object Store for JASMIN and ESGF

Performance-based: evaluate different technologies

- S3netcdf4 (<https://github.com/cedadev/S3-netcdf-python>)
- NetCDF4-python with HTTP range GET extension
- Xarray / zarr

Functional and user-directed: populate a subset of CMIP Data using xarray / zarr

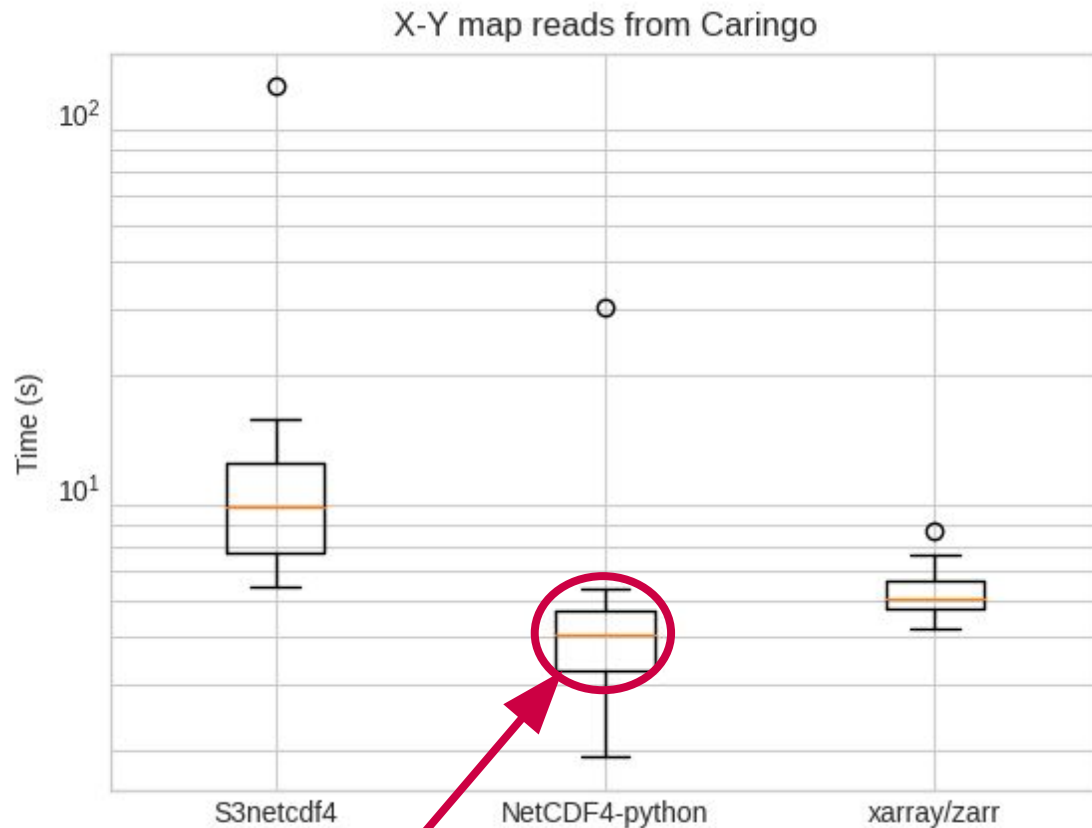
Object Store performance test design



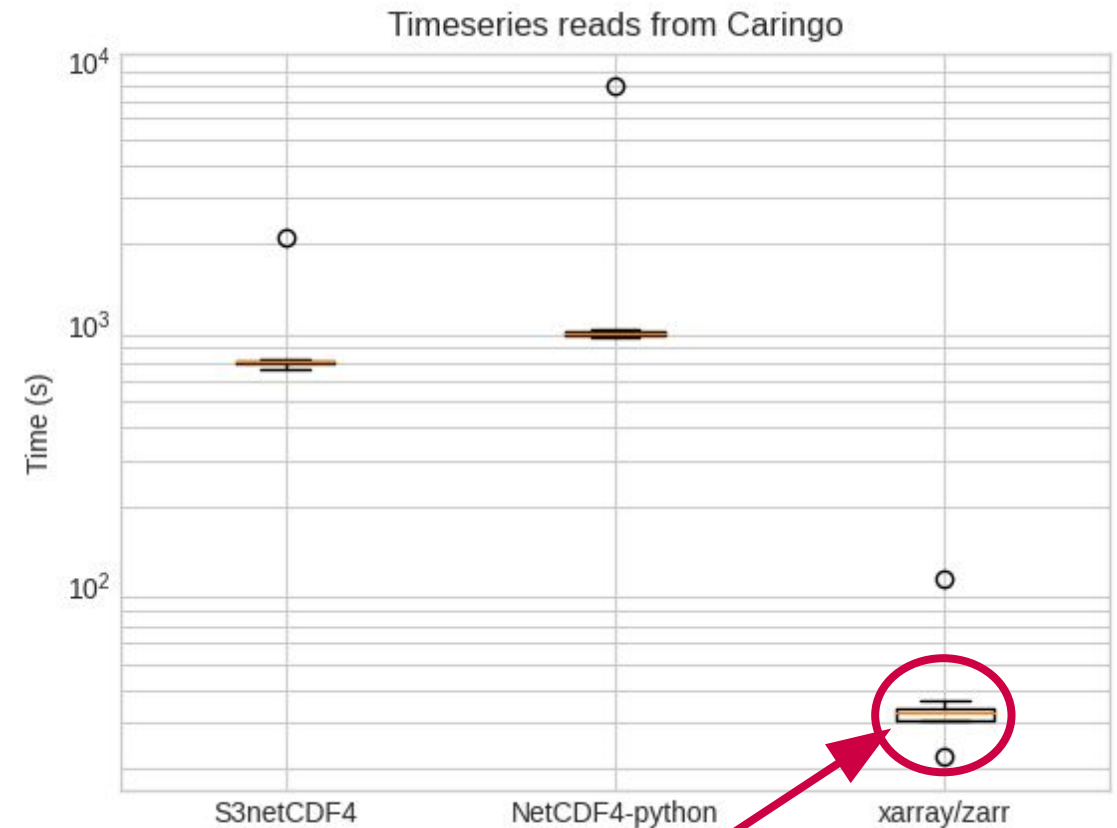
- Use JASMIN *Caringo* object store (S3 compliant interface)
- Use serial client queries – x/y and time series slicing
- Used BCC test dataset, relative humidity
- Object / chunking (□ ~250MB objects)

Serialisation	Time steps	Height	latitude	longitude
S3netcdf4	60	19	160	320
NetCDF4-python	7200	19	160	320
xarray/zarr	68	19	160	320

Object testing *preliminary* results reflect the chunking strategies; more analysis needed



Complete slices in lat/lon allow for efficient retrieval



xarray/zarr performed well on time series queries

CMIP6 Subset on JASMIN Object Store: functional evaluation



Ongoing work and next steps



Collaboration on CMIP data on object store: Pangeo community, AWS ESGF node, DKRZ, ...




Search:

- object store and traditional HTTP file serving
- ESM profile for STAC?



Long-term preservation and archiving – to address:

- Checksums
- Versioning
- Reconstruction from data corruption



Interface for multiple storage media

- S3 as interface for tape and object store

Acknowledgements + Contact Details



- This work has been carried out through IS-ENES3, a project funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Philip.Kershaw@stfc.ac.uk

@PhilipJKershaw

