


# To Exascale via Exabytes: An Analysis Problem (as well as an HPC problem)

## Workflow:

- Develop Models (**years**)
  - Integrate Models (**months**)  
(using multiple HPC platforms)
  - Move data (avoid  $n \times n$  data movement)
  - Evaluate Integrations  
(involves comparison, **years**)
  - Long data lifetimes!
- Future:
- More output! Bigger evaluation Teams!

## Joint Analysis System: JASMIN

- Private Cloud with Archive
  - Batch Compute Service
- For
- Atmospheric Science
  - Earth Observation
  - Environmental Genomics  
and more.



## One Recent Example:

- N512L180 Model (25 km global with 180 levels)
- “Understanding” not “Projection”
- 2 year integration: 1.8M core hours (Archer)
- 100 TB output (high temporal resolution too)
- Output cf Obs, other models?
- Moved it from ARCHER to JASMIN at  
> 50 TB/week (while others moving data too).

## Analysis Environment

- £10m investment
- 12 PB Fast Disk
- Curated Archive
- Petascale Tape
- 4000 Cores
- 2 TB/s I/O
- Lightpaths elsewhere

