



## **Benchmarking OPeNDAP services for modern ESM data workloads**

S. Pascoe (1), R. Wilkinson (2), and P. Kershaw (1)

(1) Centre of Environmental Data Archival, STFC, UK, (2) Tessella plc., Abingdon, UK

Experience from the 5th Coupled Model Intercomparison Project (CMIP5) suggests that end-user download of today's Earth System Model outputs is constrained by the effective bandwidth of our global network links. One solution is to provide higher-level data access services which allow users to select only the data they require. OPeNDAP is a well established RESTful protocol for sub-setting data in the NetCDF data model that has the potential to meet this need. However, for OPeNDAP to be a viable alternative to basic file transfer, OPeNDAP services must scale to the data volumes required to serve today's ESM data whilst keeping server load at manageable levels.

We present the dapbench framework for testing the performance and security constraints of OPeNDAP servers under serial and parallel workloads. We compare the performance of commonly used OPeNDAP server implementations for request patterns typical of high-level NetCDF tools and analyse the affect of ESG Federation security filters on the performance of serial and parallel workloads.

As a possible optimisation of OPeNDAP services we suggest constraining the allowed requests to a cachable subset and evaluate the effect of HTTP proxy caching performance. The results from these experiments will be used to identify the challenges to effective use of OPeNDAP in an analysis context.