

# CLIP-C WP5 Workshop on tape archive integration with ESGF

September 22-23rd, 2014. ECMWF, Reading, UK.  
(<http://www.ecmwf.int/en/about/contact-us/location>)

The Earth System Grid Federation provides a unified mechanism for discovery and download for on-disk datasets in the earth system sciences and is underpinning the design of the CLIP-C portal for climate impact indicators. So far delivery of data from tape archive has not been part of the ESGF system, although some ESGF sites have implemented mechanisms specific to their institutions. For CLIP-C there is a requirement to deliver EURO4M data from the MARS tape archive at LIU into the ESGF system and there is a wider need to develop a standard approach to exposing data in deep storage for ESGF.

This workshop aims to specify a system architecture for implementing tape archive integration with ESGF through addressing the specific use cases of the CLIP-C project.

## Attendees

STFC:

- Stephen Pascoe ([Stephen.Pascoe@stfc.ac.uk](mailto:Stephen.Pascoe@stfc.ac.uk))
- Ag Stephens ([Ag.Stephens@stfc.ac.uk](mailto:Ag.Stephens@stfc.ac.uk))

LIU/NSC:

- Per Lundqvist ([perl@nsc.liu.se](mailto:perl@nsc.liu.se))
- Hamish Struthers ([struthers@nsc.liu.se](mailto:struthers@nsc.liu.se))
- Prashanth Dwarakanath ([pchengi@nsc.liu.se](mailto:pchengi@nsc.liu.se))

SMHI:

- Michael Kolax ([michael.kolax@smhi.se](mailto:michael.kolax@smhi.se))
- Sébastien Villaume ([sebastien.villaume@smhi.se](mailto:sebastien.villaume@smhi.se))

UKMO:

- Mark Elkington ([mark.elkington@metoffice.gov.uk](mailto:mark.elkington@metoffice.gov.uk))

ECMWF:

- Manuel Fuentes
- Baudouin Raoult
- Kevin Marsh

## Agenda

### Monday September 22nd

Session 1 : Introduce the problem (10:30 - 12:00)

- The CLIP-C WP5 aims [Stephen Pascoe]

- Specific solution for EURO4M
- Recommend a standard for offline datasets to ESGF
- Introduce WP5 Tape use cases
  - Serve EURO4M from tape [Sébastien Villaume]
  - Allow users to order products from UKMO [Mark Elkington]
- Prior work on ESGF tape integration [Stephen Pascoe]
  - esgf-stager (<https://github.com/ESGF/esgf-stager>)
  - ESGF and long term storage at DKRZ

### Session 2 : Current technology and work to date on solutions (12:00 - 13:00)

- MARS deployment at LIU [Per Lundqvist]
- LIU proposal for ESGF tape integration (SODA) [Prashanth D./Per Lundqvist]
- How UKMO publishes to ESGF through BADC [Mark Elkington]

### Session 3 : The user perspective (14:00 - 17:00)

- User workflow in the context of ESGF
  - For both use cases what would the process look like from the user perspective?
  - Storyboarding the flow for the user
  - including notification, download.

## Tuesday September 23rd (09:00 - 15:00)

### Session 4 : Data manager / system admin perspective (09:00 - 11:00)

- What constraints are imposed by file format conversion Grib->NetCDF?
- Compare the MARS and UKMO use cases. Are they sufficiently similar to be implemented as a single architecture?
- What components will be required to implement: e.g. caches, conversion processes.

### Session 5 : Proposed architecture (11:30 - 13:00)

This session will develop the first draft of an ESGF deep storage architecture proposal drawing together the ideas from session 3 and 4.

### Session 6 : Review outcomes and actions (14:00 - 15:00)

# Meeting Summary

## Session 1

Stephen introduced the workshop by identifying the goals of CLIP-C Task 5.4. These were:

1. Enable publishing of the EURO4M 3D-VAR reanalyses to ESGF, currently held at LIU on their MARS system.

2. Enable on-demand delivery of datasets from the UKMO's MASS-R system to BADC for publishing to ESGF.
3. Enable ESGF users to request datasets that are not covered by current ESGF holdings but could be delivered if sufficiently popular.
4. Recommend an architecture and API for publishing data to ESGF that is held in tape archives.

Stephen described how the CLIP-C architecture envisages ESGF as a service back-plane for CLIP-C and how tape extraction fits into the overall architecture.

We identified that OPeNDAP support for ESGF datasets is an important element for visualisation of ESGF data within the CLIP-C portal and therefore particular attention should be given to how/whether it will be possible to expose tape-archived data as OPeNDAP endpoints at ESGF nodes.

**ACTION:** Stephen to follow-up the importance of OPeNDAP with the architecture group.

Sebastien presented details of the EURO4M reanalysis datasets at LIU and their interface through the MARS query language, MARS keywords used in archiving EURO4M and a possible DRS mapping.

Stephen presented material provided by Stephan Kindermann at DKRZ on their interface between ESGF and the DKRZ long term storage system. The storage system comprises the CERA database and a container-based tape storage system. Two approaches were described:

1. Wrapping CERA data access and metadata APIs in ESGF adapters.
2. FUSE based API to files on tape.

## Session 2

Per presented details of the LIU deployment of MARS and exploring implications for integrating with ESGF.

- MARS is deployed on top of the TSM tape system. TSM is a significant bottleneck for extractions so extract requests should be expected to take hours rather than minutes.
- Grib is field based, which aids subset and aggregation operations. ESGF files could be mapped to collections of Grib records.
- **NOTE:** It will not be feasible to construct checksums for inclusion in ESGF as the byte-representation of files will depend on the extract and conversion process.
- MARS does not include authentication/authorisation so is unsuitable as an API external to NSC.

Prashanth presented the LIU proposal for an ESGF API to offline data as the System for Offline Data Access (SODA). The design proposed a component shipped with an ESGF installation (SODA) and plugins that interact with back-end system. Metadata being published into index nodes, and download requests triggering backend request.

Plugins would be responsible for:

- validating a DRS query
- converting the DRS query to plugin-specific terms (e.g. MARS keywords)
  - Either call MARS binary

- Or call ECMWF Web API

SODA would be responsible for generating ESGF metadata and publishing.

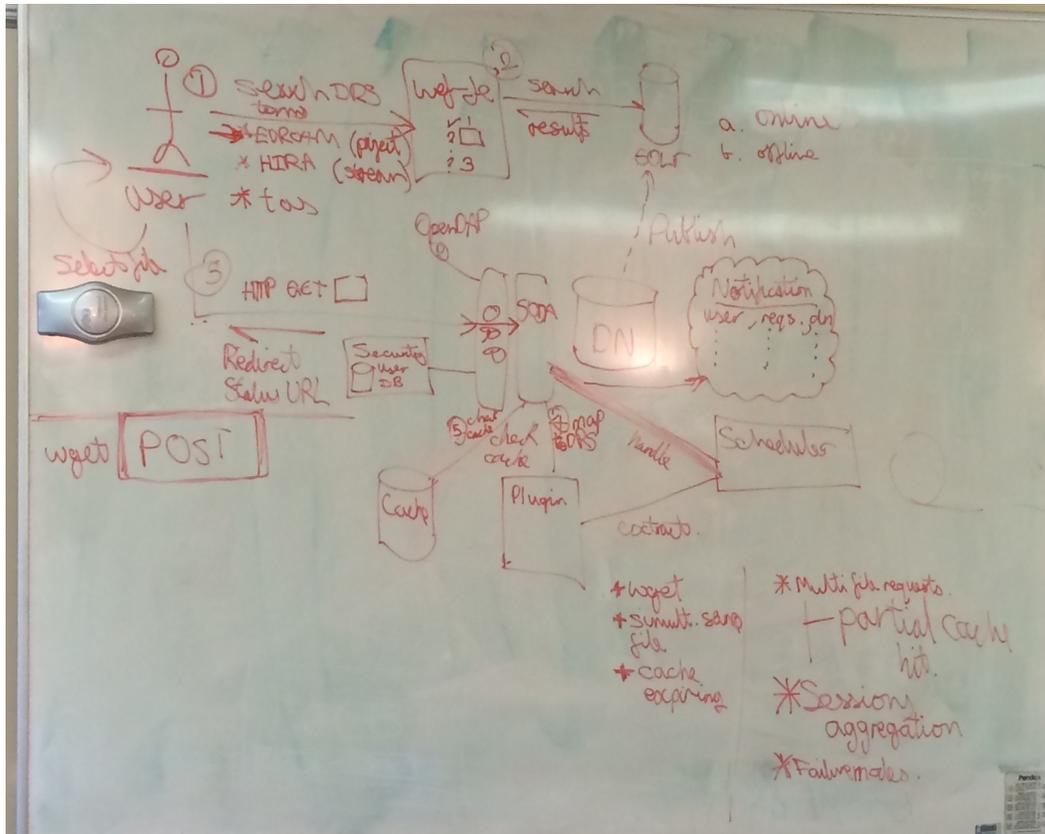
Mark Elkington presented the UKMO system which delivers data to the BADC for publication to ESGF. He identified 2 use cases which they would like to address through CLIP-C WP5.

1. Publication of CMIP5 and CORDEX has proceeded by first delivering data to BADC, waiting for confirmation of successful delivery and passing quality control, then UKMO archives the data in their MASS-R system whilst BADC publishes to ESGF. A more automated approach could be achieved if quality control is performed at UKMO and the data is archived in MASS-R first, then a synchronisation process occurs between MASS-R and BADC. This could include on-demand synchronisation of data that has been requested by users.
2. There are many cases where particular data has not been requested by MIPs and therefore UKMO has not performed the required post-processing for these data, however this data is requested by the community. It would be useful to have an automated mechanism for requesting data that could be provided if there is sufficient demand.

### Sessions 3

Working from LIU's SODA design we outlined the details of how such an API would interface to ESGF. Several key architectural features became apparent.

1. SODA will require a separate data request API to ensure that file requests are batched into suitable sizes for translation into queries of the back-end system.
2. Cache management will not be straightforward. When should a file be expired from the cache? Should all files pertaining to a query be expired at once or should each file be expired independently? Possible expiration algorithms are Least Frequently Used (LFU) or Least Recently Used (LRU).
3. When considering the asynchronous nature of a SODA request the core API begins to look very similar to the OGC Web Processing Service, in that:
  - i. An initial request results in a URL resolving to a status resource.
  - ii. Retrieval of the status URL returns a document listing the status (such as % complete) and lists any file which is available to download)
  - iii. SODA will have to coordinate the asynchronous execution of multiple jobs. Therefore it must be able to monitor multiple UNIX processes or interact with a scheduler.



Whiteboard notes illustrating how ESGF requests from the user would be translated into SODA API calls and requests to the tape back-end.

SODA datasets will need to be a new type of ESGF dataset, marked with an extra metadata item. Even if the data is available already as a standard the API will be different therefore the datasets will need to be different in SOLr. The user or agent makes a request for a bag of files which results in a handle (landing page URL) which will eventually return a document linking to the results. The User can poll the handle URL or will receive an email on completion.

Subsetting via OPeNDAP could be supported from the SODA cache but front-end systems need to be aware that a SODA dataset is not subsettable until it has been extracted to the cache. This is another reason to separate SODA datasets from standard datasets.

SODA should use an ESGF notification system to send emails when jobs are complete.

#### Session 4

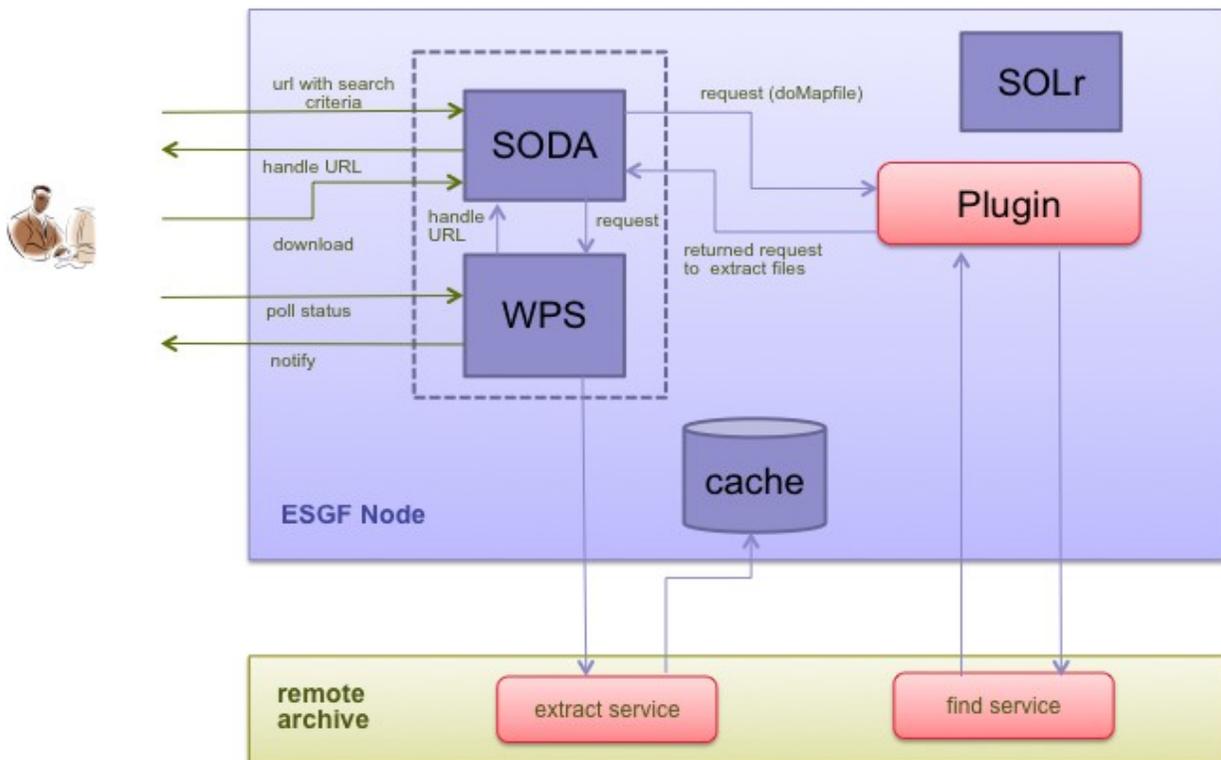
We discussed mapping Grib keywords to ESGF DRS terms. Most keys would map conveniently to a DRS structure except for forecast-time. We need to propose a mapping between the Grib keyword space to ESGF datasets and files (**ACTION:** Sebastien/Michael/Grisha).

We agreed that it was acceptable not to publish SODA datasets with checksums, as they couldn't be reliably calculated.

We discussed the division of responsibility between SODA and its plugins. Plugins would be responsible for mapping ESGF dataset/file metadata to and from processing backend queries (e.g. MARS queries). This is required for publishing and SODA query processing. Plugins could also be responsible for submitting to the processing backend (e.g. MARS) but this is to be clarified later. SODA plugins will enable publishing to ESGF via the new publishing API. We envisage a workflow similar to ESG Publisher, involving “mapfiles”.

It was not clear whether Plugins would communicate to ESGF directly or simply provide a source of metadata to SODA. This is to be addressed in the detailed SODA architecture (**ACTION:** Prashanth).

We discussed using SODA to implement the UKMO use cases. We agreed that it could fit well with their first use case (direct synchronisation of MASS-R with BADC).



## Actions

1. SebastienV+Michael+Grisha : estimate the quantity of virtual files and datasets that would need to be indexed to make EURO4M searchable within ESGF
  - a. Which DRS terms are used as facets?
  - b. How many datasets will result?
  - c. How many virtual files in each dataset?
  - d. How will datasets be divided into files (by level, by step, by time)[Estimates to be delivered by Oct. 6th for inclusion in EU Review]
2. Stephen : Clarify the requirements for SODA datasets with the CLIP-C architecture team. Do we need to support OPeNDAP?
3. Mark: Distribute architecture diagrams from the workshop.
4. Prashanth / Per: Write-up the SODA interface to the user and other ESGF components.
5. Prashanth : Decide on SODA implementation after investigation of WPS. Write an implementation plan.
6. Stephen : find out data access policy for EURO4M.
7. Stephen : discuss resourcing with WPL and STFC, LIU/SMHI, CMCC.
8. Telco for to coordinate on Oct 13th, 11:30 European time.