

# CMIP5 Dataset Version Directory Structure

Stephen Pascoe, Version 0.5, 2010-06-04

This document recommends a directory layout for supporting multiple versions of datasets within the CMIP5 archive. In drawing up this proposal the following factors have been considered:

1. The structure should conform as closely as possible to the CMIP5 Data Reference Syntax document [DRS].
2. The structure should be workable with the current ESG publisher software.
3. The structure should present a consistent user interface for downloading data files.
4. The structure should be manageable by operators of ESGF data nodes.

To give some background each of considerations are discussed below then the proposed directory structure is described in detail.

## 1 Discussion of Considerations

The current plans for version control in the CMIP5 archive derive from 2 independent sources: The DRS document and the ESG publisher software. **[MORE]**

### 1.1 DRS Document

The DRS document divides the CMIP5 archive into a set of *atomic datasets* defined terms of several components including a version number, therefore it implies datasets are immutable with a fixed version number. The assumption is that when data within an atomic dataset is superseded a new dataset is generated with an incremented version number; the original dataset remains unchanged and available to users.

The DRS document recommends version should be part of the URL syntax used to allow users to download data and also part of the directory structure used to store the data at the data-node.

ESGF has decided to publish CMIP5 data at the DRS realm level. Due to the way the ESG publisher manages versions this implies that version numbers are associated with realms within an experiment rather than individual atomic datasets. Therefore a change to the DRS directory hierarchy is recommended.

**Recommendation:** The DRS component *version* should be moved in the suggested directory/URL hierarchy from just below ensemble-member to just below realm. The DRS file hierarchy then becomes:

*activity/product/institute/model/experiment/frequency/realm/**version**/variable/ensemble-member/files*

### 1.2 ESG Publisher

ESG publisher uses the terms *dataset* and *dataset version* to represent a set of files published in one operation. A dataset version is an immutable object which can represent a 'DRS dataset including version number'. The published 'dataset version' itself has an identifier which typically consists of dataset\_id+version number; this appears in the THREDDS catalog of the dataset version. For each dataset version a new THREDDS catalogue is produced. ESG Publisher also detects when files change such that an updated file will be recorded as a new version of that file. In this case the ESG publisher database keeps track of the metadata of previous file versions but does not manage

previous versions of the actual files. The paths to files within a dataset can be changed with the *rename* operation, however this operation does not increment the version number.

### **1.3 User Interface to Data Files**

It is expected that users will visit an ESG Gateway to generate wget or DML scripts containing lists of URLs to data files. These URLs should follow the URL syntax suggested in the DRS document. Importantly the path portion of the URL identifies the atomic dataset and therefore, from the user's perspective, these URLs will be the primary identifiers of both atomic datasets and files in the archive. The version directory structure has been designed such that URLs have the following user-centric features, assuming URLs follow reflect the directory structure:

1. The user has a record of which version of the data they downloaded from the version numbers embedded in the URLs
2. Running the script still retrieves the original data even after a new version is published (except for “latest”, see below)
3. URLs can be edited to retrieve other versions of an atomic dataset if the user knows that version exists.
4. When a new dataset version is published it should be possible for a user to download only those files that have changed in the new version. This may require tool support.
5. A special DRS version “latest” can be used to retrieve the latest version of a dataset. Downloading “latest” has the disadvantage that it is not immediately clear to the user which version you are downloading. Additional tools should be made available to deduce the exact DRS version of a set of files based on their *tracking\_id* or checksum
6. A possible desirable feature is for wget scripts to support recursive download below the atomic dataset level so that if the user doesn't know the exact filenames in a dataset version they can still retrieve it.

### **1.4 Data Management**

Maintaining previous versions of *atomic datasets* must not put too high a burden on the data management practices of data-node administrators.

1. If a file is unchanged between versions there should be no need to store the file twice. Using symbolic links is recommended for solving this problem.
2. Filenames should remain unchanged when a version is superseded.
3. Mixing data files and symbolic links within one directory should be avoided as the asymmetry complicates data management.
4. It should be possible to remove the data of previous versions without breaking symbolic links. This will probably require tool support.

## 2 Proposed Version Directory Structure

We can think of the directory structure containing a datanode's data as a collection of realm directories. These realm directories are arranged in the DRS directory hierarchy relative to the CMIP5 activity directory as follows:

*/activity/product/institution/model/experiment/frequency/realm*

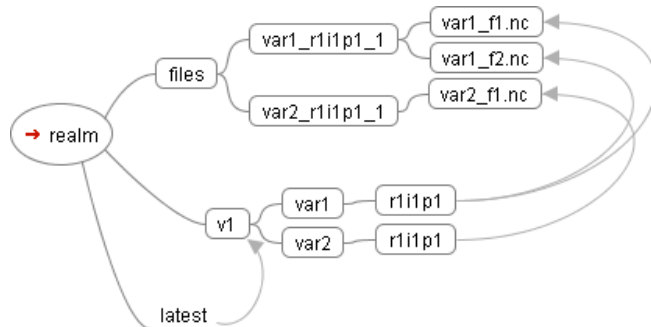
Each realm directory contains all versions of the data related to that realm and each of these versions are published in the ESG system as individual dataset versions.

The versioned directory structure below each realm directory is divided into a **files** directory containing all data files, the directories **v<n>** containing links to files at a particular version and a link **latest** to the most recent **v<n>** directory.

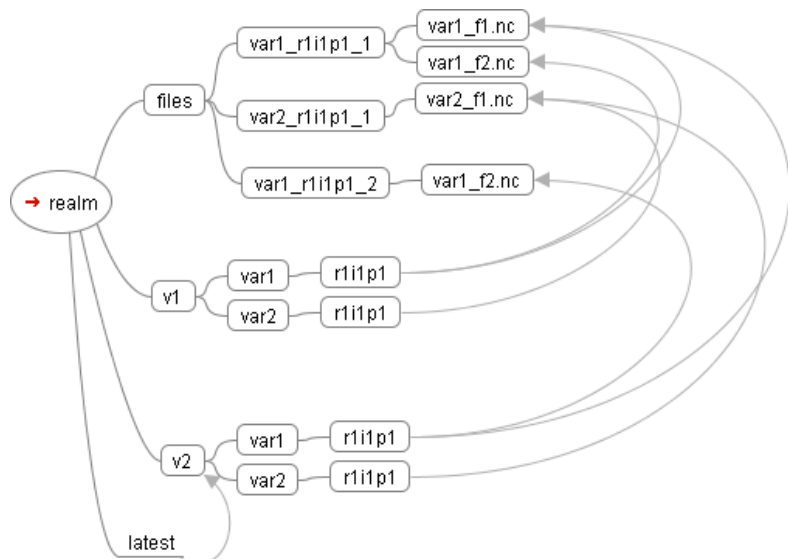
Within **files** a series of subdirectories

**files/<var>\_<emember>\_<n>** contain files changed or added at each version. Using a compound naming scheme for these directories keeps the tree relatively flat without affecting the DRS hierarchy which is accessible through the **v<n>** directories.

This is illustrated with the following scenario of an atomic dataset at 2 successive versions.



*Version 1: Realm of 2 variables containing 3 files*



*Version 2: Replace **var1\_f2.nc***

### 3 Transition of data between version and publication states

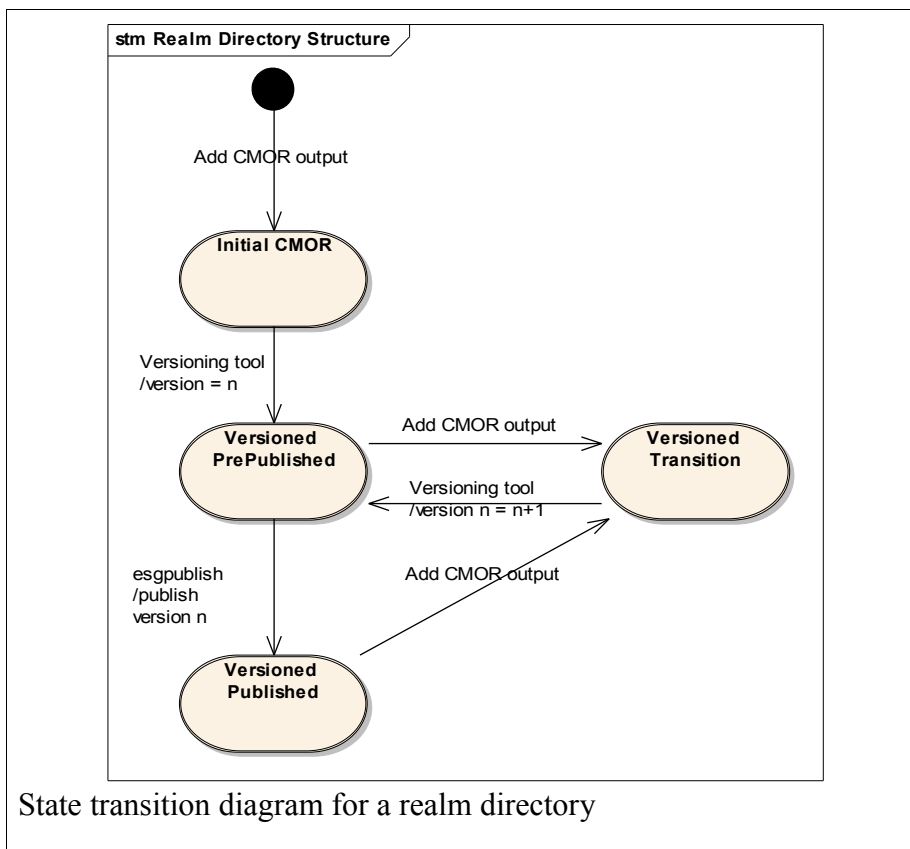
When a realm dataset is first placed on a datanode it is assumed the directory structure will be that created by CMOR. CMOR outputs files into the DRS hierarchy without the version component:


*/activity/product/institution/model/experiment/frequency/realm/variable/emember/file.nc*

In order to store multiple versions of a realm dataset on a single filesystem this directory structure must be transformed into one that contains the current version whilst maintaining previous versions. Therefore during the data delivery and publication lifecycle a realm directory will move between 4 states:


1. **Initial CMOR:** structure reflects exactly what CMOR produces. No **files** or **v<n>** subdirectories are present.
2. **Versioned PrePublished:** The realm directory has been transformed into the Version Directory Structure and all files moved to their appropriate locations. The latest version has not been published with ESG publisher.
3. **Versioned Published:** The ESG publisher has been run on all version subdirectories of the realm dataset such that each version is published on the datanode.
4. **Versioned Transition:** New data files, to be published in the next version, are present in the CMOR directory structure alongside the published versions.

Transitions between these 4 states are managed by a combination of the ESG publisher and a version maintenance tool (to be written).






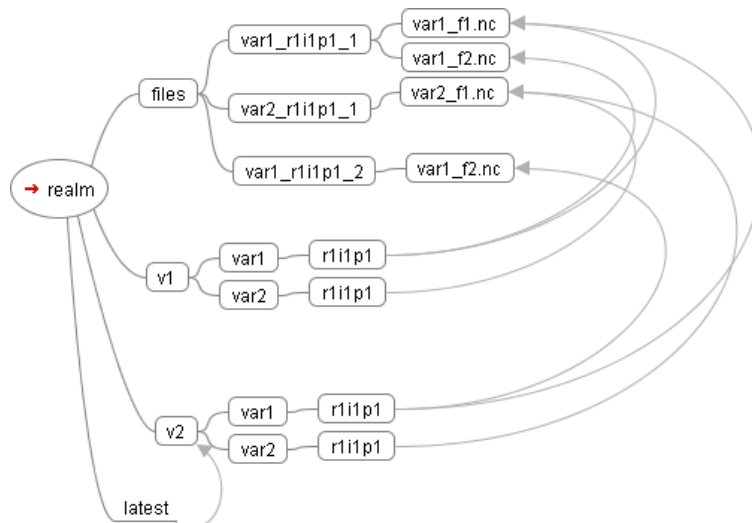
State **Initial CMOR**



State **Versioned PrePublished**. Version 1 ready to publish



State **Versioned Transition**. Files have been added prior to publishing version 2



State **Version PrePublished**. Version 2 ready to publish.

### **3.1 User download experience**

When a new dataset version is published users will wish to download only those parts of the dataset that have changed. This information is contained in the **files** subdirectories that are not expected to be visible to the user. We therefore need a service to provide users with wget scripts to download these changes.

## **4 Citations**

[DRS] CMIP5 Data Reference Syntax (DRS) and Controlled Vocabularies. [http://cmip-pcmdi.llnl.gov/cmip5/docs/cmip5\\_data\\_reference\\_syntax.pdf](http://cmip-pcmdi.llnl.gov/cmip5/docs/cmip5_data_reference_syntax.pdf)