

CS1 Data storage and archiving policy

16/07/2007, H.A. Holties

Background

As for the future LOFAR system, the CS1 storage capacity is too small to store all data in a central location indefinitely. Moreover, the central storage location needs to provide storage not only for observation datasets, but also for (temporary) commissioning and analysis data. The policy below has been defined to allow efficient use of the offline storage cluster while providing as much support to people involved in the commissioning of LOAR CS1 as possible.

NB: The policy as provided here applies principally to the data in the “offline” domain. In this domain, the main central storage facility consists of five storage systems, each with a storage capacity of 1.2 Tbyte. A more complete description of the CS1 offline systems can be found [here](#).

CS1 observation data policy

The datasets created with the LOFAR CS1 system are copied to the offline storage cluster after observations have finished. The available datasets and their locations can be found in the [observation catalog](#). Data is kept on the central storage system for a period of at least a week. After the retention period, the data is either deleted or archived, depending on its status in the “Archive” column in the observation catalog. The oldest datasets will be (re) moved first. The exact time at which deletion or archiving is carried out will depend on the storage capacity required for near future observations and commissioning activities.

Archived data will be moved to GRID based storage servers. Commissioning data will be owned by the /lofar/ops group and stored under the lfc directory /grid/lofar/ops/cs1_data as a tarball. The current policy for this data is that it will be archived on tape in two copies after an unknown period. Please beware that GRID based storage is experimental. At present, no guarantees can be given about data lifetime and retrieval although these issues are being addressed and data security is likely orders of magnitude better than on the central CS1 storage facilities.

CS1 commissioning data (user data) policy

To allow commissioning analysis of datasets and testing of offline processing software, users may request usage of central storage capacity. Requests should be directed to the CS1 system administrator (Teun Grit). The request should be accompanied by the capacity needed. Further information may be requested by the CS1 system administrator (such as motivation, expected period for which storage capacity is needed, etc).

The CS1 system administrator will confirm the allocated storage capacity and the period for which it can be used. If demanded so for operational or commissioning purposes, the CS1 system administrator can ask users to (re) move their data. An overview of data capacity allocated to users is

given in the following table:

Useraccount	Storage directory	Allocated capacity	Used capacity (if larger)	Allocated until
brentjens	/lifs001/brentjens	100 GB (example only)	650 GB	CS1 EOL

NB1: The CS1 offline storage servers provide storage on “best effort” basis. Data is stored on disks in a RAID striping configuration (for optimal performance) and is not backed up. Any data that is important enough to keep for any period should be copied to other systems that provide better data security.

NB2: There is no software used for accounting storage capacity. Please check the used capacity yourself. Manual checks will be carried out by the CS1 system administrator as well, in particular in situations where the available storage capacity becomes too small to allow planned operations and commissioning. Offenders will be notified and prompt action is expected in that case. In case there is no response, corrective actions may be taken by the CS1 system administrator. Offending users may be “named and shamed”.

From:

<https://www.astron.nl/lofarwiki/> - **LOFAR Wiki**

Permanent link:

https://www.astron.nl/lofarwiki/doku.php?id=public:cs1_data_policy&rev=1184580517

Last update: **2007-07-16 10:08**

