
2025-10-30 19:23 1/2 cs1 storage

Retrieving CS1 data

So where does all the data go?

In the "offline" domain, the main central storage facility consists of twelve storage servers with a storage capacity of 1.2 Tbyte each. These systems are called lifs001 through lifs012. All of the offline storage directories are mounted on all of the offline processing servers (lioff001 ... lioff018). They are also directly accessible on the LISTFEN frontend server as automounted directories /lifs001 ... /lifs012.

The diagram below shows the, for MOST users, relevant systems for the CS1 data streams.



First, it is important to realize that there is no such thing as THE data. There is observation data, monitoring data, configuration data, etc. Next, it is also important to be aware that, just as for the future "normal" LOFAR operations, the bulk of the data will not be stored in a single central location indefinitely. Most data will be stored centrally for a limitted period after which it (or at least a selection) will be archived elsewhere according to the CS1 data policy.

Observation datasets

To find the most interesting datasets, check the Observation catalog. The "MS-name" column provides the name of the dataset. The "stored" column provides information on where the associated datasets can be found.

For example: "lifs001 - lifs003, lifs005" in the location columnn indicates that the measurement is to be found in parts divided over four of the offline storage nodes. On each node you will find the datasets in the /data/<MS-name> directory. Note that the /data directories are mounted on all offline processing servers, as well as on the listfen frontend system. It is therefore in practice not necessary (and even not recommended) to actually log in to one of the storage nodes!

An entry "archived" indicates that the dataset has moved to long term storage on the GRID where it is located in the lfc directory /grid/lofar/ops/cs1_data. Check the CS1 GRID UI page for information on how to retrieve data from the GRID. NB: in general, datasets should not be transfered from the GRID based archive back to the LOFAR cluster. Rather, transfer data to your own system or analyse it using GRID based processing resources.

Station (auto) correlation data

Station data is copied (manually) from the station SCU's to a machine (lofar15) in Dwingeloo. The oldest datasets are archived, as a temporary solution, on adam. Contact Rob Millenaar or Hanno Holties for details. The intention for the future is to create and store station data (at least as long as is deemed necessary for commissioning purposes) together with the observation datasets.

Configuration data

The intention is to store all system configuration (a.k.a. "parset") files together with the observations (TBD). These files are used throughout the system to configure observation processes. The files are generated manually or, in the near future, automatically but he SAS/MAC system and distributed to the appropriate processes.

CS1 Data storage and archiving policy

The CS1 data storage and archiving policy is provided here.

From:

https://www.astron.nl/lofarwiki/ - LOFAR Wiki

Permanent link:

https://www.astron.nl/lofarwiki/doku.php?id=public:cs1 storage&rev=1190364527

Last update: **2007-09-21 08:48**

