The LOFAR Cluster

This page describes the LOFAR cluster for (potential) users. The cluster is used to store the correlated data from the BlueGene and to run pipelines that do the "standard" reduction and calibration. After this pipeline processing, the results are stored in the Lofar Export (staging) Archive.

We welcome authorized users on this cluster. For the time being the users are bound to only a part of the cluster. This part is called a subcluster. Various users groups have been granted access to a specific subcluster. Users can utilized a standard login environment (see below) and they can access the cluster resources using one of the two frontends "lfe001" and "lfe002".

LOFAR Cluster layout (brief)

The Lofar cluster is devided into 8 subclusters. Each subcluster is a processing cluster for a specific commsissioning group. There are 72 compute nodes (named `lfexxx`) and 24 storage nodes (named `lsexxx`) in total. Each subcluster has 9 compute nodes and 3 storage nodes of 24TB raw capacity. The storage nodes have 4 RAID5 partitions each. Each partition holds a single XFS filesystem. Each filesystem is NFS mounted on all 9 compute nodes in the subcluster (but NOT in other subclusters). So 1 compute node has 12 NFS data volumes of approx. 1.5 Tbyte mounted.

Frontend

A frontend has 2 Intel Xeon L5420 quad core processors, 16GB internal meomory, 2 GbE interfaces and 2TB disks in RAID5 configuration. There are actually two identical frontends. Both of them serve a specific group of subclusters. The frontends are used to build the software and regulate the workload on the subclusters. There is a page that describes the frontend [in detail](#)

Processing units

The compute elements have 2 Intel Xeon L5420 quad core processors, 16GB internal meomory, 2 GbE interfaces and 1TB disks in RAID0 configuration. They can be accessed by secure shell and they are grouped.

Storage Units

The storage nodes are HP DL180G5 boxes, having 2 Intel Xeon L5420 quad core processors, 16GB internal memory, 6 GbE network interfaces and 24TB disks. The disks are devided into 4 partitions of 6 disks each, set up in RAID5 configuration. The XFS filesystems are called "/data1" till "/data4".

Available Subclusters

The current subcluster assignment is:

```
lfe001:~> showsub

This script shows the subcluster definitions

sub   lce-nodes       lse-nodes       cexec-lce       cexec-lse       In use for:
====  =========       =========       =========       =========       ===========
sub1  lce001-lce009   lse001-lse003   lce:0-8         lse:0-2         production
group
sub2  lce010-lce018   lse004-lse006   lce:9-17        lse:3-5         no power
sub3  lce019-lce027   lse007-lse009   lce:18-26       lse:6-8         imaging group
sub4  lce028-lce036   lse010-lse012   lce:27-35       lse:9-11        no power
sub5  lce037-lce045   lse013-lse015   lce:36-44       lse:12-14       pulsar group
sub6  lce046-lce054   lse016-lse018   lce:45-53       lse:15-17       no power
sub7  lce055-lce063   lse019-lse021   lce:54-62       lse:18-20       developers
group
sub8  lce064-lce072   lse022-lse024   lce:63-71       lse:21-23       no power
```

The lce-nodes ar ethe compute/processing nodes, the lse-nodes are the associated storage nodes. Only use the storage nodes that are associated with your processing node!

User Access

Access through portal to Cluster frontend

You can access the Lofar cluster through the portal: "**ssh -X portal.lofar.eu**" We maintain a ssh whitelist, so only known institutes are able to login. Please send an email to grit at astron.nl or h.paas at rug.nl to add your institute or personal IP number. Once logged in, you'll find an empty home directory. You'll have to login at one of the two frontends "lfe001" of "lfe002" using "ssh -X". When you don't know which one, please use "lfe001" by default.

LOGIN environment

See The Lofar Login Environment page for information on your default login environment. Only the default environment is supported, so if you defer from this, you're on your own..!

Do and Don'ts

Old Stuff

Offline processing clusters

The disks of the offline storage nodes are NFS mounted on all these offline nodes. These systems are in use for offline processing (i.e. when data has been transfered from the central storage facility).

The clusters are described on more detail on this page. See this page for more information on the login environment.

Offline storage cluster

This cluster acts as a temporary data buffer. Data is copied from the online storage systems to these systems, so that the online systems can be used to store new data. These systems are accessable from the offline processing clusters, allowing post-processing of the data. More information can be found on this page.

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

Permanent link:
https://www.astron.nl/lofarwiki/doku.php?id=public:lofar_cluster&rev=1255435306

Last update: **2009-10-13 12:01**