

# Cluster Definition File

As described in the page on [running DiFX](#) you need to provide mpifxcorr with machines and threads files and an appropriate mpirun command. There are utilities to automate this. Currently the primary options are [genmachines/startdifax](#), [espresso](#) and [startcorr.pl](#).

The tools above all need a basic definition of your cluster in order to automatically set up the required MPI files. A standard cluster definition file format has been agreed to facilitate this, and will be used by all the above tools starting from the DiFX-2.3 release.

## Cluster Definition File Format

# is used to indicate a comment. The # and all subsequent characters on the line are not parsed. Empty lines and comment-only lines are permitted (and ignored).

The first line of the file gives the version of the cluster file *format* in the form `version = <INTEGER>`. This version number is to differentiate between different versions of the cluster definition file format, in case changes to the format are required in the future. Currently only `version = 1` is valid.

Subsequent lines in the file define the nodes that comprise the cluster. Each node definition line contains a comma separated list of values that define the relevant features of one or more nodes. If a node name appears more than once in the file, the last entry for that node is used (later entries supersede earlier ones). The meanings and allowed formats of each column are given in the following table:

Column number	Format	Meaning
1	String, with zero-padded numeric expansion of numbers contained in square brackets. E.g. <code>cuppa[01-20]</code> expands to include nodes <code>cuppa01</code> , <code>cuppa02</code> , <code>cuppa03</code> ... <code>cuppa20</code> . For expansions within the square brackets, the number of digits must be the same (with zero padding used where appropriate). Only one set of square brackets can be used on a single line.	The name of the cluster node(s). Multiple identical nodes can be described by a single line using the '[' ]' notation.
2	integer	0 means the node is disabled, 1 means the node is enabled, 2 means the node is enabled and eligible to be the master node
3	integer	the maximum number of compute threads to be used on this node. Nodes which are only to be used as datastream or master nodes may be given the value 0 to avoid using them for compute

Column number	Format	Meaning
4	space separated urls. file://<path>, mark5://, mark6://, network://<IP>	a space separated list of urls for data sources on this node. Allowed url formats are: file://<path> for directories containing files, mark5:// for Mark5 machines, network://<IP> for an eVLBI data source (<IP> is the external IP address of the node), mark6: // for Mark6 machines


An example cluster definition file:

```

version = 1      # version number is an integer
# node, enabled/disabled (2=head node), number compute threads, space
separated list of urls for data
# Later lines supersede earlier lines
cuppa01,        2, 7, file:///exports/xraid01/l_1/corr
file:///exports/xraid01/r_1/corr # possible master node
cuppa02,        1, 7, file:///exports/xraid02/l_1/corr
file:///exports/xraid02/r_1/corr
cuppa03,        1, 7, file:///exports/xraid03/l_1/corr
file:///exports/xraid03/r_1/corr file:///arch/corr/bbdata/corr
cuppa05,        1, 7, file:///exports/xraid05/l_1/corr
file:///exports/xraid05/r_1/corr
cuppa[07-15],   1, 7          # zero-padded numeric ranges are allowed
cuppa12,        0, 7          # disable cuppa12 - supersedes previous line
in expanded range.
cuppa16,        1, 7, file:///mnt/disk1/corr file:///mnt/disk2/corr
cuppa17,        1, 7, file:///mnt/disk1/corr
cuppa18,        1, 7, file:///mnt/disk1/corr
cuppa19,        1, 7
cuppa2[1-4],    1, 0, file:///mnt/raid/corr      # datastream only nodes - 0
compute threads.
mark5b-1,       1, 0, mark5:// network://202.8.37.0 network://202.8.37.1
file:///data    # a Mark5 with a linux partition, and also used for eVLBI
mark6-01,       1, 0, mark6:// file:///fuse_mark6-01      # a mark6 (native
playback) or file-based playback from /fuse_mark6-01
    
```

## Notes for using genmachines with the cluster definition file

- genmachines evaluates the \$DIFX\_MACHINES environment variable to point at the cluster definition file
- multiple nodes can be assigned to the same file: *data source*. In this case the associated nodes get assigned in a round-robin fashion if file: data sources appears more than once in the current job.



From:

<http://www.atnf.csiro.au/vlbi/dokuwiki/> - **ATNF VLBI Wiki**

Permanent link:

<http://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/clusterdef?rev=1536155488>

Last update: **2018/09/05 23:51**

