

Modifying Startup Configuration Options

As part of the ClustrixDB Installation process you made some configuration selections which are written to `clxnode.conf`. Post-installation, you can modify the configuration by editing the file on each node while ensuring that the options specified are the same (other than node specific items like `BACKEND_ADDR`).

The `clxnode.conf` is stored on each node in `/etc/clustrix/` and options specified during installation are written to this file. The following table summarizes these options and provides information on how they can be modified post-installation.

The steps outlined in this document apply to most installations of ClustrixDB, but do not apply to deployments based on the Clustrix AWS AMI.

- [Default Options for clxnode.conf](#)
- [Additional options for clxnode.conf](#)
- [Details of Modifying Startup Configuration Options](#)
 - [Stop/Start all ClustrixDB Services](#)
 - [CPU_CORES](#)
 - [BACKEND_ADDR](#)

ClustrixDB requires that all nodes have identical configurations (other than node specific items like `BACKEND_ADDR`). Modifications to any of the options described below must be made on every node in the cluster. Values commented out with a `#` will have no effect on the database. Changes to the configuration requires a restart of ClustrixDB services (see below) to take effect.

Default Options for clxnode.conf

Use standard bash conventions for `clxnode.conf`. The file should contain variable definitions, comments, and blank lines only. Option and value settings should not contain any extraneous spaces. For example `CPU_CORES=4`

Option	Description	Default Value	To Modify
<code>NODE_MEMORY</code>	Memory to use for ClustrixDB, in MiB	<i>auto detected</i>	<ol style="list-style-type: none">1. On every node in your cluster edit the <code>clxnode.conf</code> file and uncomment (if commented out) the line <code>#NODE_MEMORY</code> and change the value to the desired amount.2. Stop ClustrixDB, restart <code>hugetlb</code>, and then Start ClustrixDB on all nodes simultaneously by running the following command from a single node: <code>clx -P cmd 'systemctl stop clustrix; systemctl restart hugetlb; systemctl start clustrix'</code>
<code>MAX_REDO</code>	Maximum ClustrixDB Redo Space, in MiB	1024	Only modify this variable with guidance from Clustrix Support.
<code>CPU_CORES</code>	Number of CPU cores ClustrixDB will use	All	<ol style="list-style-type: none">1. Edit the <code>clxnode.conf</code> file for every node in your cluster to the desired value.2. Stop/Start all ClustrixDB services. <p>For additional info, please see section below on setting <code>CPU_CORES</code>.</p>
<code>DATA_PATH</code>	ClustrixDB Database Storage Path	<code>/data/clustrix</code>	<ol style="list-style-type: none">1. Stop all ClustrixDB services.2. On each node:<ol style="list-style-type: none">a. Move the directory to the desired location.b. Modify the value of <code>DATA_PATH</code> in <code>clxnode.conf</code> to reflect the new location.3. Start all ClustrixDB services.
<code>DATA_ALLOCATE</code>	Storage space to pre-allocate to ClustrixDB	<i>80% of free space on \$DATA_PATH volume</i>	This value here is used for initialization only. To allocate more space on a node that already has ClustrixDB installed, see ALTER CLUSTER RESIZE DEVICES
<code>LOG_PATH</code>	ClustrixDB Logs Path	<code>/data/clustrix/log</code>	<ol style="list-style-type: none">1. Stop all ClustrixDB services.2. On each node:<ol style="list-style-type: none">a. Move the directory to the desired location.b. Create a simlink in <code>/data/clustrix/</code> pointing to the new log folder. Example: <code>ln -s /log /data/clustrix/log/</code>3. Start all ClustrixDB services.

UI_LOGDIR	ClustrixDB WebUI Logs Path	/data/clustrix/log/clustrix_ui	<ol style="list-style-type: none"> 1. Stop all ClustrixDB services. 2. On each node: <ol style="list-style-type: none"> a. Move the directory to the desired location. b. Create a simlink in /data/clustrix/log pointing to the new log folder. Example: <code>ln -s /log/clustrix_ui /data/clustrix/log/clustrix_ui</code> 3. Start all ClustrixDB services.
UI_CACHEDIR	ClustrixDB WebUI Cache Path	/var/cache/clustrix/django	Only modify this variable with guidance from Clustrix Support.
UNIX_SOCKET_PATH	ClustrixDB MySQL Protocol Unix Socket	/var/lib/mysql/mysql.sock	Only modify this variable with guidance from Clustrix Support.
LISTEN_ADDR	ClustrixDB Database Listen Address (Front-End IP) Interface	0.0.0.0	
MYSQL_PORT	ClustrixDB Database MySQL Port	3306	
BACKEND_ADDR	ClustrixDB Private Back-End IP Interface	<i>auto detected</i>	See section below on changing BACKEND_ADDR.
BACKEND_PORT	ClustrixDB Back End Network Port	24378	It is not recommended to modify the default value for the ClustrixDB Back End Network port.
HEALTH_MONITOR_PORT	The port used for the ClustrixDB health monitor.	3581	<ol style="list-style-type: none"> 1. Edit value in <code>clxnode.conf</code> 2. Stop/Start all ClustrixDB services.
HTTP_PORT	ClustrixDB WebUI HTTP Port	80	<ol style="list-style-type: none"> 1. Edit value in <code>clxnode.conf</code> 2. Stop/Start all ClustrixDB services.
NANNY_PORT	ClustrixDB Nanny Port	2424	It is not recommended to modify the default value for the ClustrixDB nanny port.
CONTROL_PORT	ClustrixDB Control Port	2048	It is not recommended to modify the default value for the ClustrixDB control port.
CLX_USER	ClustrixDB OS User	root	Do not modify this value after installation.
MANAGER_USER	ClustrixDB Manager User	root	Do not modify this value after installation.
WRITE_HOSTS	Allow ClustrixDB to modify <code>sshd_config</code> and <code>/etc/hosts</code> . This used by root installations only, for internode communication for administrative tasks, including upgrades.	<i>#Commented out</i>	<ol style="list-style-type: none"> 1. On each node: <ol style="list-style-type: none"> a. Remove comment and set variable to <i>any string</i>. b. Stop/Start all ClustrixDB services.
HUGE_TLB_ENABLE	ClustrixDB enables HugeTLB memory allocation for faster startup. NOTE: This causes instability on some systems (particularly certain virtual machines). Required for non-root installations that have more than 20GiB of \$NODE_MEMORY Contact Clustrix Support before changing from default.	True	<ol style="list-style-type: none"> 1. On each node, comment out to disable this option. (i.e. <code>#HUGE_TLB_ENABLE</code>) 2. Stop/Start all ClustrixDB services.

Additional options for `clxnode.conf`

The following options are NOT included by default in `clxnode.conf`, but can be added to the end of that file to modify behavior from the default. ClustrixDB must be restarted for changes to take effect.

Option	Description	Default Value	To Modify
DISABLE_MULTIPORT	When this option is present, the multiport feature is disabled.	This option is not included in <code>clxnode.conf</code> , resulting in multiport being enabled by default.	<p>To disable multiport, add an entry for <code>DISABLE_MULTIPORT = Y</code>.</p> <p>To enable multiport, remove or comment out this entry from <code>clxnode.conf</code> on all nodes.</p>
MULTIPAGE_ALLOC	Bytes reserved for multi-page allocators.	The option is not included in <code>clxnode.conf</code> , resulting in a value of 1GiB	Add this entry and define the number of bytes that Clustrix Support recommends for your workload.

Details of Modifying Startup Configuration Options

ClustrixDB requires that all nodes must have the same values in `clxnode.conf`

Stop/Start all ClustrixDB Services

To apply the changes that have been made to `clxnode.conf`, the ClustrixDB services on all the nodes will need to be restarted:

```
shell> /opt/clustrix/bin/clx
dbrestart
```

CPU_CORES

The default setting of `CPU_CORES=All` will make use of all available cores in the system. If you would like to limit the number of cores utilized by ClustrixDB you must modify the value for `CPU_CORES` to specify the number of cores that should be used. ClustrixDB licensing will base its enforcement on the number of cores being utilized in the system as specified by `CPU_CORES`.

BACKEND_ADDR

1. Modify the `BACKEND_ADDR` entry in the `/etc/clustrix/clxnode.conf` file with the new IP.
2. Modify `/data/clustrix/broadcast_nodes` file with the new IP addresses on all nodes.
 - a. Note: Specifying port 24378 is not necessary unless you're using a non-default configuration.
3. Run the `clx` command to ensure it has cached the IP address of each node.

```
shell> /opt/clustrix/bin/clx cmd
'clx cache'
```

4. Stop all ClustrixDB services (per above).
5. Change the backend interface IP address on each node.
6. Modify iptables as needed.
7. Start all ClustrixDB services (per above).