

Expanding Your Cluster's Capacity - Flex Up

Congratulations on needing to expand your cluster on ClustrixDB!

- [Step 1. Prepare the Node\(s\)](#)
- [Step 2. Add Node\(s\) to Your Cluster](#)
- [Step 3. Verify Successful Node Addition](#)
- [Step 4. Monitor Rebalancer Activity](#)
- [Errors During Flex Up](#)

Your ClustrixDB is licensed for a maximum number of cores per node as well as a maximum number of nodes for the cluster. Contact [Clustrix Sales](#) to if you need an expanded license.

Step 1. Prepare the Node(s)

Provision nodes and [Install ClustrixDB](#) on each node, using the same version and installation configuration.

If [zones](#) are in use, ensure that the target configuration has an equal number of nodes in each zone.

If you are using [encrypted connections](#), ensure that certificates in keys are present in the same location with the correct permissions.

Step 2. Add Node(s) to Your Cluster

When adding nodes to your cluster, you may need to open ports. For a listing of all ports required by ClustrixDB, including those used for multiport, please see [Network Security with ClustrixDB](#).

Connect to one of the existing nodes of your cluster and run the following from a SQL prompt. The IP(s) specified are those for new nodes being added.

Add nodes to your cluster

```
sql> ALTER CLUSTER ADD '
ip' [, 'ip'] ...;
```

ClustrixDB will perform a [Group Change](#) to add new node(s).

Clustrix recommends running this command during non-peak periods or during a scheduled maintenance window.



There will be a short disruption of service while the node(s) are being added.

You may also notice a slight degradation of performance while the Rebalancer moves data to the new node(s).

(optional) Configure Zone for newly added Nodes

If zones are in use, the newly added nodes must be assigned to their respective zone using [ALTER CLUSTER ZONE](#). Clusters that are only partially configured for zones are not fault tolerant in the face of a zone loss.

Step 3. Verify Successful Node Addition

Run this query from a SQL prompt. New node(s) are automatically assigned consecutive node number(s).

```
sql> SELECT * FROM system.nodeinfo ORDER BY nodeid;
```

You can also use [The CLX Command-Line Administration Tool](#).

View your cluster's status by providing the following at a bash prompt

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

You should see that all nodes appear **OK** on the display. If zones are in use, all nodes should have a non-zero zone assigned. If zones are not in use, all nodes will be assigned to zone 0.

You may notice that the data distribution amongst your nodes is not yet balanced. Be patient. It will be, soon.

If the node you are trying to add does not appear in the list above, see section below on Errors during Flex Up.

Step 4. Monitor Rebalancer Activity

Your new node(s) have been successfully added to your ClustrixDB cluster but they do not yet contain data. The Rebalancer will now automatically work in the background to move data onto the new node(s). To monitor this process, refer to the instructions regarding [Managing the Rebalancer](#). Your cluster is fully functional and able to be used during this process.

Errors During Flex Up

As part of adding nodes to your cluster, ClustrixDB performs some checks to ensure the nodes have the same configuration. This section describes errors that can be encountered with ALTER CLUSTER ADD and how to resolve those issues.

SQL Errors

The following are errors you may encounter at the SQL prompt

- ERROR 1 (HY000): [26633] Bad parameter.: Host "xxx" already in cluster
 - This error occurs if you attempt to add a node that is already part of the cluster.
- ERROR 1 (HY000): [55300] Multiple nodes cannot be added when an existing addition is pending
 - A node addition could be in a pending state due to:
 - A node is still in the process of being added
 - A node running a different version of ClustrixDB was added to the cluster
 - A node with a different multiport setting was added to the cluster
 - See sections below on PROBLEM_NODES and PENDING_INVITES to resolve these errors.

PENDING_INVITES

List of nodes with pending invitations

```
sql> SELECT * FROM system.pending_invites;
```

The cluster periodically attempts to send invitations to nodes in system.pending_invites. For each invitation that is sent, there will be entries in clustrix.log:

```
sending invitation response(no error) to "10.2.13.68:24378"
```

Note: If the same node is in system.pending_invites and system.problem_nodes, you may want to remove the node from system.pending_invites before resolving the issue with system.problem_nodes. Doing this will prevent the node from being automatically added to the cluster once the problem is resolved, causing a premature group change. Instead, you may prefer to complete the node addition during off-peak hours.

Remove a pending node addition from PENDING_INVITES

```
sql> DELETE FROM system.pending_invites;
```

PROBLEM_NODES

Query system.problem_nodes to see why a node could not be added

```
sql> SELECT * FROM system.problem_nodes;
```

Here is the list of reasons provided in system.problem_nodes and how to resolve those issues:

Reason	Resolution
Software binaries differ	Make sure all nodes are running the same version of ClustrixDB
Multiport settings mismatched	See Modifying Startup Configuration Options for instructions on how to disable Multiport. Network Security with ClustrixDB contains information relative to enabling Multiport ports.