

Reducing Your Cluster's Capacity - Flex Down

On occasion, you may need to reduce your cluster's capacity:

- To reduce operating costs following a peak event (i.e. following Cyber-Monday).
- To allocate servers for other purposes.
- To remove failing hardware. (See [ALTER CLUSTER DROP](#) to drop a permanently failed node.)

The process to downsize your cluster within Xpand is a simple:

- [Step 1: Initiate SOFTFAIL](#)
- [Step 2: Monitor the SOFTFAIL Process](#)
- [Step 3: Remove Softfailed Node\(s\) from Your Cluster](#)

Xpand recommends running this process while logged on to a node other than one you wish to drop.

Review target cluster configuration

- Xpand requires a minimum of three nodes to support production systems. Going from three or more nodes to a single node is not supported via the steps outlined on this page.
- When [Zones](#) are configured, Xpand requires a minimum of 3 zones.
- For clusters deployed in zones, Xpand requires an equal number of nodes in each zone.
- Ensure that the target cluster configuration has sufficient space. See [Allocating Disk Space for Fault Tolerance and Availability](#).

Flex Down

Step 1: Initiate SOFTFAIL

Marking a node as softfailed directs the Xpand Rebalancer to move all data from the node(s) specified to others within the cluster. The Rebalancer proceeds in the background while the database continues to serve your ongoing production needs.

If necessary, determine the nodeid assigned to a given IP or hostname by running the following SQL select.

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

To initiate a SOFTFAIL, using [ALTER CLUSTER](#).

```
ALTER CLUSTER SOFTFAIL  
nodeid [, nodeid] ...
```

The SOFTFAIL operation will issue an error if there is not sufficient space to complete the softfail or if the softfail would leave the cluster unable to protect data should an additional node be lost.

To cancel a SOFTFAIL process before it completes, use the following syntax. Your system will be restored to its prior configuration.

```
ALTER CLUSTER UNSOFTFAIL  
nodeid [, nodeid] ...
```

Step 2: Monitor the SOFTFAIL Process

Once marked as softfailed, the Rebalancer moves data from the softfailed node(s). The Rebalancer process runs in the background while foreground processing continues to serve your production workload.

To monitor the progress of the SOFTFAIL:

Verify that the node(s) you specified are indeed marked for removal.

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

The system.softfailing_containers tables will show the list of containers that are slated to be moved as part of the SOFTFAIL operation. When this query returns 0 rows, the data migration is complete.

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

This query shows the list of softfailed node(s) that are ready for removal.

```
sql> SELECT * FROM system.softfailed_nodes
      WHERE nodeid NOT IN
          (SELECT DISTINCT nodeid
           FROM system.softfailing_containers);
```

Once the SOFTFAIL is complete for all nodes, the clustrix.log file will contain the following message:

```
softfailing nodes are ready to be removed: <list of node ids>
```

Step 3: Remove Softfailed Node(s) from Your Cluster

To remove the softfailed node(s) from the cluster, issue the following SQL command.

```
sql> ALTER CLUSTER REFORM;
```

There will be a brief interruption of service while the node(s) are removed.