

Allocating Disk Space for Fault Tolerance and Availability

A cluster must contain sufficient free disk space to automatically recover from node or zone failures. To calculate the maximum amount of disk space that can be utilized while still allowing ClustrixDB to fully reprotect data after a failure, you can use the following formula:

$$\text{Maximum Disk Utilization \%} = (\text{Total Nodes} - k) * 80 / \text{Total Nodes}$$

In the formula above, k represents one of the following (whichever is larger):

- Value of MAX_FAILURES global variable (default value = 1)
- The total number of nodes in a zone (if an entire zone were to fail). Refer to [Zones](#) for more information on how ClustrixDB works in zones.

To examine the disk space usage of your cluster, use [The CLX Command Line Administrative Tool](#).

If there is not enough free space in your cluster to reprotect data following a node or zone failure, your cluster will be at risk for data loss or cluster failure if another node or zone is lost.

80 is the default value for the databasefull_user_warn_percentage threshold. If your applications write data at a rate that fills the cluster aggressively, use a value that is less than 80% in your calculations. This will ensure that data can continue to be written while you are waiting for replacement node(s) to join the cluster and for data redistribution to complete.

To configure the databasefull_user_warn_percentage threshold, and others related to database space utilization, please see [Managing File Space and Database Capacity](#).

Samples

Using this sample chart of pre-calculated thresholds, a 9 node cluster (not deployed in zones) with MAX_FAILURES = 1 will require that the database not exceed 71.11% capacity to ensure successful completion of reprotect actions in the event of a node failure.

A similar 9 node cluster instead deployed in 3 zones of 3 nodes each, would require that no more than 53.33% of your cluster's space be used to ensure successful completion of reprotect actions in the event of a zone failure.

k	Total Cluster Nodes				
	3	6	9	16	32
1	53.33%	66.67%	71.11%	75.00%	77.50%
2	*	53.33%	62.22%	70.00%	75.00%
3	*	*	53.33%	65.00%	72.50%
4	*	*	44.44%	60.00%	70.00%
5	*	*	*	55.00%	67.50%

*

= Not applicable as the remaining number of nodes will not constitute a quorum.

Space Alerts

If the amount of free space in your cluster goes below the amount of space required to fully reprotect it in the event of a node or zone failure, ClustrixDB will send an email alert to the list of users configured in [Database Alerts](#). The email will include **[WARNING] Insufficient space for reprotection** and provide details on the amount of space required.

The same message will also appear in clustrix.log as an ERROR.

```
ERROR 1 (HY000): [32779] Not enough space to reprotect if another node is lost: 94.4255% usage (without softfailed nodes) is greater than max 80.0000%
```

You may also encounter this message when softfailing or removing nodes.