

FLEX UP AND FLEX DOWN FOR SEASONAL DEMAND

'Right Size' Your Platform to Customer Demand—All Year Long

Introduction

As your e-commerce site grows, your hosted infrastructure will need to scale to keep up with customer demand. Your website's database is particularly vulnerable to a move—especially if it relies on a single server, as MySQL does. Moving to another server can involve considerable downtime as you dump the data off the system to be replaced and restore it to the new server with more capacity.

As your business continues to grow, downtime becomes increasingly disruptive and expensive. Moreover, as you approach the capacity limits of the largest server, you must consider workarounds to offload processing from the overburdened MySQL server. Managing these workarounds is not simple and can lead to additional costs in both people and resources. It can also lead to errors and downtime.

ClustrixDB Overview

ClustrixDB is a distributed SQL database that is designed to seamlessly scale horizontally—by adding cores and servers. ClustrixDB uses patented data distribution algorithms to distribute load and data across a cluster of nodes dynamically, all the while as a single logical database. ClustrixDB is designed to be a drop-in replacement for MySQL and includes support for key features like stored procedures, triggers, and MySQL replication.

A key advantage of ClustrixDB is the ability to add capacity ('flex up') and reduce capacity ('flex down') while maintaining high availability. This enables customers to tailor the amount of hosting capacity they acquire to their short-term needs, such as the period leading up to and including Black Friday.

This document describes in more detail how Flex Up and Flex Down work on ClustrixDB.

Flex Up

When your business hits its peak season, you will want extra database capacity to handle increased traffic and orders. If you use MySQL and run out of database capacity, you know what happens: page load times slow, your database starts locking up, and ultimately your site goes down. Capacity problems like these often show up at the worst time, when demand is high and your revenue should be spiking. When capacity problems appear, they can escalate quickly. It's dangerous to think that you will have enough time to react.

With ClustrixDB, Adding Capacity is Fast and Easy

To add capacity to your database or prepare for your peak period, start by assessing your workload needs and allocate a number of nodes to add to your ClustrixDB cluster. Once you have determined the number of nodes to add, simply make sure they are on the same network as your Clustrix database and record the IP addresses.

First Step: Run the Flex Up Installer

Choose a node in your existing ClustrixDB system and run the Flex Up Installer with the IPs of your new nodes. This is done with a simple one-line command for each ClustrixDB instance to be created.

The Flex Up Installer automatically:

- Determines your cluster's configuration and generates an installation package to expand your cluster
- Deploys that installation package to your new node and performs the installation, creating a new ClustrixDB instance

Once you run the Flex Up Installer for each node to be added—creating new ClustrixDB instances—you add the new nodes to your cluster with another simple command, this time at the ClustrixDB prompt.

```
clustrix> ALTER CLUSTER ADD 10.2.3.4, 10.4.7.8, 10.2.0.1
```

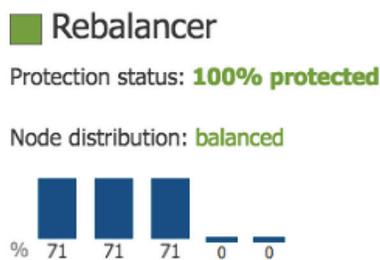
This results in a short interruption of service (measured in seconds), after which the new nodes are part of your database cluster.

At this point, the added instances are fully functional and work can be directed to them by your load balancer.

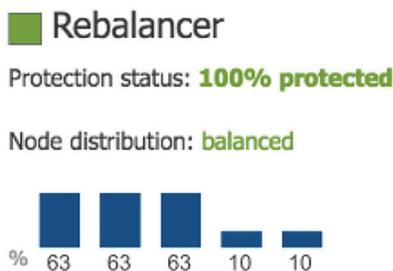
Second Step: Rebalance the Data

The new instances are now a working part of your cluster, but they will not provide full value until the database has been rebalanced across the new instances.

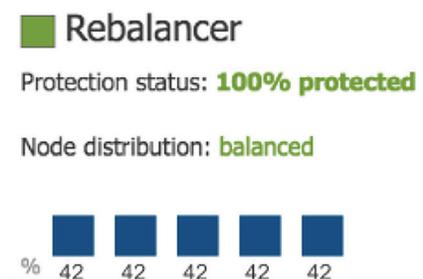
Below, you see that the new instances are part of your cluster, but do not yet contain data:



Clustrix designed its Rebalancer to redistribute data to the new instances in the background while continuing to serve your production workload. It is easy to monitor the progress of the rebalancer while data is being transferred. Below, you can see that the new nodes are at 10% of capacity as data is being offloaded from the three original nodes.



Once the rebalancing is complete, data and load are distributed evenly across the cluster:



If you underestimate your database capacity needs or find that your traffic is exceeding your projections, you can repeat the steps above to add nodes at any time.

Flex Down

Once your peak period has ended, you may want to 'flex down' to save on operational costs. To do this, you simply mark the nodes to be decommissioned as 'sod failed'. This tells ClustrixDB that it should move all data off the 'sod failed' nodes to the remaining nodes in the cluster. Once again, the redistribution of data proceeds in the background, while the database continues to serve your production load.

When the soft fail process has completed, you can remove the instances from the cluster with a simple operation. This results in a short interruption of service (again, measured in seconds), after which your database will be operational, but with fewer instances. The nodes can now be decommissioned.

Conclusion

ClustrixDB can be scaled both up and down to match your customer demand. The ability to add and remove nodes easily and redistribute data dynamically is a significant advantage of the ClustrixDB architecture.

There is no need to pay for more hosting infrastructure than you're likely to need. When peak season is over, the extra capacity is an unnecessary expense. ClustrixDB is designed to replace MySQL simply and easily. Once ClustrixDB is installed, adding and removing platform capacity is fast and easy.

Call Clustrix to find out how we can help you save money by 'right-sizing' your database platform to customer demand—all year long.

Clustrix

Clustrix provides the leading scale-out SQL database engineered for the cloud. With ClustrixDB you can build innovative business critical applications that deliver real-time analytics on live operational data with massive transactional volume. Our exceptional customer service team supports more than one trillion transactions per month across a range of industry segments including Ad Tech, e-commerce, and social analytics. Clustrix customers include AOL, engage:BDR, MedExpert, Photobox, Rakuten, Symantec, and Twoo.com. Headquartered in San Francisco, Clustrix is funded by HighBAR Partners, Sequoia Capital, U.S. Venture Partners, Don Listwin, and ATA Ventures. ClustrixDB is available as a free community edition for developers, a software download that runs on any cloud, and on the AWS marketplace.

To learn more about Clustrix, visit us at www.clustrix.com

