Basic Administration

This page outlines some frequently used utilities and tools available within ClustrixDB.

- Load Balancing ClustrixDB
- Accessing ClustrixDB from the Command Line
- Loading Data onto ClustrixDB
- Web-based Cluster Administration
- Examining ClustrixDB Server Logs
- Using CLX - Cluster-Wide Command Line Utility
- Monitoring ClustrixDB Using statd
- Upgrading ClustrixDB
- Expanding/Reducing Your Cluster's Capacity
- Parallel Backup and Restore

Load Balancing ClustrixDB

ClustrixDB requires an external TCP load balancer to balance queries across a cluster. Clustrix recommends HAProxy as a load balancer. For more information, see Load Balancing ClustrixDB with HAProxy.

Accessing ClustrixDB from the Command Line

To access ClustrixDB from the command line, run MySQL from a command line prompt.

```
shell> mysql
sql> SELECT * FROM test.sample ORDER BY birth;

+------------+------------+------------+-----+
| name       | birth      | death      | sex |
+------------+------------+------------+-----+
| Rowinsky   | 1952-08-01 | NULL       | F   |
| Snyder     | 1955-11-26 | 2014-05-17 | M   |
| Thompson   | 1965-03-09 | NULL       | F   |
| Johnson    | 1980-06-11 | NULL       | M   |
+------------+------------+------------+-----+
4 rows in set (0.00 sec)
```

Because ClustrixDB is MySQL compatible, many third party tools such as MySQLWorkbench, TOAD, etc. will work seamlessly with ClustrixDB.

Loading Data onto ClustrixDB

Depending on your specific needs, there are numerous options for loading data onto ClustrixDB:

- Clustrix recommends using clustrix_import if your source database is MySQL or you have a mysqldump file available. This is the fastest method for loading data onto ClustrixDB from another system.
- To cut over a production environment from MySQL to ClustrixDB, you can use MySQL replication as outlined in Migrating to ClustrixDB from MySQL.
- To transfer user accounts and permissions, use clustrix_clone_users as documented in Migrating User Permissions.
- Load Data Infile can also be used to import data directly to ClustrixDB from a csv file.

See also Loading Data Onto ClustrixDB.

Web-based Cluster Administration

ClustrixDB includes ClustrixGUI Administration UI for cluster administration.

Examining ClustrixDB Server Logs

ClustrixDB's server logs enable users and Customer Support to understand the queries and events that occur on a ClustrixDB cluster. All the logs created by ClustrixDB are placed, by default, in the /data/clustrix/log directory (this is configurable) and there are a separate set of logs for each node of your cluster. The logs used most frequently are query.log and clustrix.log.

The query.log shows information for SQL statements used within ClustrixDB. From the query.log you can identify slow queries as well as query syntax errors. For information see Query Logging.

The clustrix.log is a chronology of events that occur on your cluster. It includes alerts that are also logged to a separate log file - alert.log. Please see Database Alerts to establish the necessary alerts for your cluster.

Using CLX - Cluster-Wide Command Line Utility

ClustrixDB includes a general-purpose command line utility called clx for monitoring and managing a cluster. A single clx command is automatically processed in parallel on all nodes of a cluster. It is located in the /opt/clustrix/bin/ directory.

The full list of clx commands and options are outlined in The CLX Command-Line Administration Tool. These commands are likely to be the most useful.

Useful CLX Commands	Description
clx status	Displays the status, IP, and disk utilization of each node in a cluster.
clx cmd	Runs Linux commands in parallel on all nodes of the cluster.
clx logdump	Consolidates and filters ClustrixDB logs so database issues can be researched and resolved.
clx dbstop	Stops ClustrixDB on all nodes.
clx dbstart	Starts ClustrixDB on all nodes.
clx dbrestart	Restarts ClustrixDB on all nodes. (Use in lieu of dbstop/dbstart when applicable.)

ClustrixDB's clx utility also communicates with another critical element of the system called nanny. The nanny is a database cluster process that provides monitoring to ensure that all processes vital to ClustrixDB node functionality are running. It is activated at system startup. If any critical ClustrixDB component fails at any point, for any reason, nanny restarts it automatically.

For a list of the processes that nanny monitors and controls, please refer to this Overview of Nanny.

Monitoring ClustrixDB Using statd

Performance statistics are automatically collected by statd in the clustrix_statd database. The resulting statistics can be assessed by querying the tables directly or by viewing them with the ClustrixGUI administrative tool.

For a recap of the statistics gathered and some potential ways to analyze them, see Monitoring Your Cluster Using statd.

Upgrading ClustrixDB

Clustrix regularly provides both major and minor releases to continuously improve its product's performance, functionality, scalability, and stability. Each release passes a thorough and rigorous quality assurance process before it is released to our customers.

For specifics on how to upgrade ClustrixDB, see Upgrading ClustrixDB.

Expanding/Reducing Your Cluster's Capacity

Your application and database may experience fluctuating transaction volumes due to seasonal changes. The underlying design of ClustrixDB enables it to easily and uniquely adapt to those changes. To increase your cluster's capacity, simply add more nodes. Conversely, if transaction volume subsequently declines, removing nodes to decrease costs is equally painless. See Flex Up and Flex Down for specifics.

Parallel Backup and Restore

ClustrixDB includes a fast parallel backup and restore feature. For backups, each ClustrixDB node sends its data directly to the backup target in parallel, allowing for concurrency and for backup performance to scale with cluster size. Similarly, the restore operation is initiated on one node and that node coordinates with all other nodes of the cluster to read from the backup file and restore data throughout the cluster in parallel.

The ClustrixDB database backups contain schema definitions, metadata relative to the backup, compacted row data, and information to ensure data consistency. For more information, see ClustrixDB Fast Backup and Restore.