

Monitoring Cluster Status

ClustrixDB provides all of the information necessary for monitoring the cluster. The system database contains low-level information about all portions of the cluster, and proprietary SQL commands return high-level summary information. The following sections describe these features in detail.

As a best practice, unix sysadmin tools such as `/usr/bin/top` should not be used on ClustrixDB to monitor node status, as their results may not be reliable.

Listing Current Sessions

The following query will list all sessions in the cluster:

```
sql> select * from system.sessions\G
***** 1. row *****
      nodeid: 1
      session_id: 1560577
      source_ip: NULL
      source_port: NULL
      local_ip: NULL
      local_port: NULL
      user: 4106
      database: NULL
      trx_state: closed
      statement_state: finished
          xid: NULL
          cpu: 3
      isolation: REPEATABLE-READ
      last_statement: delete from system.rebalancer_activity_log where finished < subdate(now(), interval @global.
rebalancer_activity_log_hours * 60 minute)
      time_in_state_s: 161
      created: 2016-01-11 20:31:41
      heap_id: 288230379201751147
      trx_age: NULL
      trx_mode: autocommit
      trx_counter_select: 0
      trx_counter_insert: 0
      trx_counter_update: 0
      trx_counter_delete: 1
          trx_is_writer: 1
1 row in set (0.01 sec)
```

Monitoring CPU Load

To display CPU load percentages (0 to 100) for CPU utilization, disk read/write utilization, and buffer cache miss rates, issue the `SHOW LOAD` command. In the cluster, CPU core 0 is used for specific tasks and therefore omitted from the load average when calculating the overall load.

Disk read and write values are calculated as a percentage using estimated maximums for solid state drives and might exceed 100% under certain conditions.

Monitoring Space Utilization per Table

The `system.table_sizes` table shows the size in bytes of each table in every database, including indexes, replicas, and row data. The value in the `size` column represents total storage on disk, rather than just the size of the data.

Monitoring Data Rebalancing Activity

ClustrixDB continuously monitors how data is distributed in the cluster. When it detects an imbalance, it rebalances data by moving, copying, and splitting as required. Data is also copied for protective purposes, such as when a disk or node is lost.

See [ClustrixGUI Administration UI](#) for more information on monitoring your cluster.