

MySQL Compatibility

ClustrixDB built from the ground up to be a distributed, shared-nothing database and does not have any source code from other projects in its code base. Any application that uses a MySQL client or JDBC/ODBC drivers (e.g. Ruby on Rails or Hibernate) can connect to ClustrixDB without issue and use standard MySQL constructs to interact with ClustrixDB.

General Differences

- ClustrixDB has a different query planner and query plans may vary from those of MySQL. As part of your migration testing, we recommend reviewing query plans for slow queries to see if they might benefit from additional tuning. See [Understanding the ClustrixDB Explain Output](#) for more information.
- ClustrixDB executable comments are of the form /*\$ */. ClustrixDB ignores other types of executable comments.
- ClustrixDB error codes do not match MySQL
- ClustrixDB implements a subset of MySQL's global variables, but behavior for those global variables may differ.

Schema Limitations

Limit	Description
Number of Tables	ClustrixDB supports a maximum of approximately 2,000 tables. The precise limit for your installation depends on the number and type of indexes.
Number of Columns	ClustrixDB recommends that each table have no more than 200 columns (including system-defined hidden columns), as higher column counts can impact performance and scalability.
Row Size	In ClustrixDB, the maximum row size for persistent tables (stored on disk) is 64MB. The maximum row size when using RBR replication is 32MB.

The maximum length for object names is 64 characters, which matches MySQL.

On ClustrixDB, object name comparisons are case insensitive and table names are stored in the case specified in the CREATE statement. The `lower_case_table_names` global variable has a value of 1 and cannot be changed.

Reserved Keywords

The following are reserved keywords for ClustrixDB but not MySQL:

- complement
- full
- legacy
- signed
- temporary

ClustrixDB Internals:

ClustrixDB is a clustered, distributed database built from the ground up for scale, concurrency, and performance as well as ease of use. While ClustrixDB seeks to be MySQL compatible, it contains no actual MySQL source code.

Here is more on ClustrixDB internals:

- Written primarily in the C programming language
- Distributed, shared-nothing architecture with [automatic data distribution across nodes](#) and a [distributed query processing](#).
- Automatically retains multiple copies of data to provide redundancy and fault tolerance

See also

- [Unsupported Features](#)
- [ClustrixDB Version and MySQL Version Compatibility](#)
- [ClustrixDB DML](#)
 - [Administrative SQL statements](#)
 - [Function and Operators](#)
 - [START TRANSACTION](#)
- [ClustrixDB DDL](#)
 - [Foreign Keys](#)
 - [Invisible Indexes](#)
 - [Partitioned Tables](#)
 - [Managing Partitions](#)
 - [Partitioned Tables Limitations](#)
 - [Why Partition a Table](#)
 - [Generated Columns](#)
 - [Online Schema Changes](#)