

Using AUTO_UNIQUE

AUTO_INCREMENT

Many tables have surrogate keys that are auto-populated using AUTO_INCREMENT. This MySQL feature is also supported by ClustrixDB and creates a unique ID as records are inserted into a table. These generated IDs increase monotonically.

Unique Identifiers are valuable in many applications, however generating strictly consecutive IDs is often unnecessary. To guarantee that an AUTO_INCREMENT value is generated consecutively, the database must find the most recent value, increment that value, and then return the next value. This process is repeated for each subsequent insert so that each transaction gets a unique value. Since this process must be serialized, AUTO_INCREMENT becomes a bottleneck for workloads with high concurrency.

Replacing AUTO_INCREMENT with AUTO_UNIQUE

To circumvent the AUTO_INCREMENT bottleneck, ClustrixDB has added a feature called AUTO_UNIQUE. Use AUTO_UNIQUE to generate unique identifiers for IDs when sequential ordering of IDs is not an application requirement.

We recommend reviewing usage of AUTO_INCREMENT values and replacing AUTO_INCREMENT with AUTO_UNIQUE when feasible.

Alter a column to use AUTO_UNIQUE

```
sql> ALTER TABLE table1 MODIFY COLUMN id bigint(0) unsigned auto_unique;
```

AUTO_UNIQUE columns must be of type BIGINT(0) unsigned.