

In-Memory Tables

- [What is an In-Memory Table?](#)
- [Why In-Memory Tables?](#)
- [How do In-Memory Tables Work?](#)
- [Summary of In-Memory Table Features](#)

What is an In-Memory Table?

ClustrixDB's In-Memory tables are kept in memory and are not written to persistent storage. They are not durable. In the event of power failure or controlled cluster shutdown, In-Memory tables will be lost.

In-Memory tables can be used in conjunction with persistent tables for queries and multi-statement transactions. They are designed for atomicity, consistency, and isolation, but not durability (they are not fully ACID compliant). In-Memory tables can be moved to persistent storage (and made durable) by using CREATE TABLE as SELECT FROM, Backup, Restore, or ALTER.

As with persistent tables, In-Memory tables are accessed via SQL. ClustrixDB stores multiple copies of In-Memory tables across the cluster to allow for concurrency and fault-tolerance. If you lose a single replica, ClustrixDB will reprotect the In-Memory tables the same as persistent tables. Queries for In-Memory tables use the same query planner and execution engine as other types of tables.

Why In-Memory Tables?

In-Memory tables can provide a performance advantage if:

- The data does not need to be durable and if the data is lost, it can re-loaded or reconstructed without serious impact to the application.
- The data fits in memory and does not grow aggressively.
- The data has a high write to read ratio - i.e. data loading or frequent updates.
- Updates use auto-committed, standalone transactions versus multi-statement transactions.

Data loading (both bulk loads and single-row data ingestion) is significantly faster into an In-Memory table versus a persistent table.

How do In-Memory Tables Work?

Differences from Tables Stored on Disk

In-Memory tables:

- Are stored non-durably in memory whereas persistent tables are stored on disk.
- Are limited to a row size of 32k (enforced at INSERT/UPDATE). Maximum row size for tables stored on persistent storage is 64MB.
- Include a memory overhead of 312 bytes per row, whereas the corresponding row overhead for persistent tables is much lower and is stored on disk.
- Use Optimistic MVCC, which means that colliding transactions will need to be reprocessed by your application (i.e. retried, rejected, other). Persistent tables use ClustrixDB's standard MVCC processing to facilitate [Concurrency Control](#).
- Queries to In-memory tables do not take locks. For example, queries that perform a SELECT ... FOR UPDATE will not block.

DDL Operations for In-Memory Tables

CREATE TABLE

ClustrixDB uses [skip lists](#) data structure for In-Memory tables.. Specify container = skiplist to create an In-Memory table..

```
SQL> CREATE TABLE sample
      (Id integer primary key,
       type varchar(5),
       description varchar(100)
      ) container=skiplist;
```

You can also use ENGINE=MEMORY to create an In-Memory table.

ALTER TABLE

To make a persistent table In-Memory:

```
SQL> ALTER TABLE sample container=skiplist;
```

To make an In-Memory table persistent:

```
SQL> ALTER TABLE sample container=layered;
```

You can also DROP and TRUNCATE In-Memory tables.

Fault Tolerance and Durability of In-Memory Tables

If a single node fails, In-Memory tables will continue to be operational and the ClustrixDB Rebalancer will work to make sure there are sufficient replicas of all tables.

In the event of a power failure or cluster restart, data stored in In-Memory tables will be lost, along with a portion of that table's metadata. The table's definition is preserved on disk, but the data, along with metadata for slices, replicas, and indices, are lost. Before any queries can be run against an In-Memory table, including a query to populate it, all metadata for the table must be restored.

If you query the table before this metadata exists, you will see this error:

```
SQL> select * from sample;
SQL> ERROR 1 (HY000): [6144] No distribution for representation: No final distribution for representation
"__idx_sample_PRIMARY"
```

To restore a table's metadata, perform one of the following:

Option 1: Drop the In-Memory table, then restore it from the most recent backup. See [ClustrixDB Fast Backup and Restore](#).

Option 2: Combine the DROP and CREATE TABLE to rebuild the table definition.

```
SQL> DROP TABLE sample;
SQL> CREATE TABLE sample
      (Id integer primary key,
       type varchar(5),
       description varchar(100)
      ) container=skiplist;
```

Optimistic MVCC

In-Memory tables use Optimistic MVCC. Persistent tables use ClustrixDB's standard Multi-Version [Concurrency Control](#) (MVCC).

The primary difference is that Optimistic MVCC does not lock table rows to prevent contention, but instead gracefully rejects colliding statements. With Optimistic MVCC, a statement within a transaction will be rejected if it attempts to write the same keyed row at the exact same time as another write transaction.

In-house tests on a small data set with no coordination and lots of concurrency produced collisions 2% - 5% of the time. Depending on your use case, these errors can be auto re-tried by your application. Setting the global variable autoretry to TRUE will cause a failed statement within a transaction to be automatically retried.

The possible errors are:

```
ERROR 1 (HY000): [5123] Container optimistic MVCC conflict: This transaction conflicted with another transaction
```

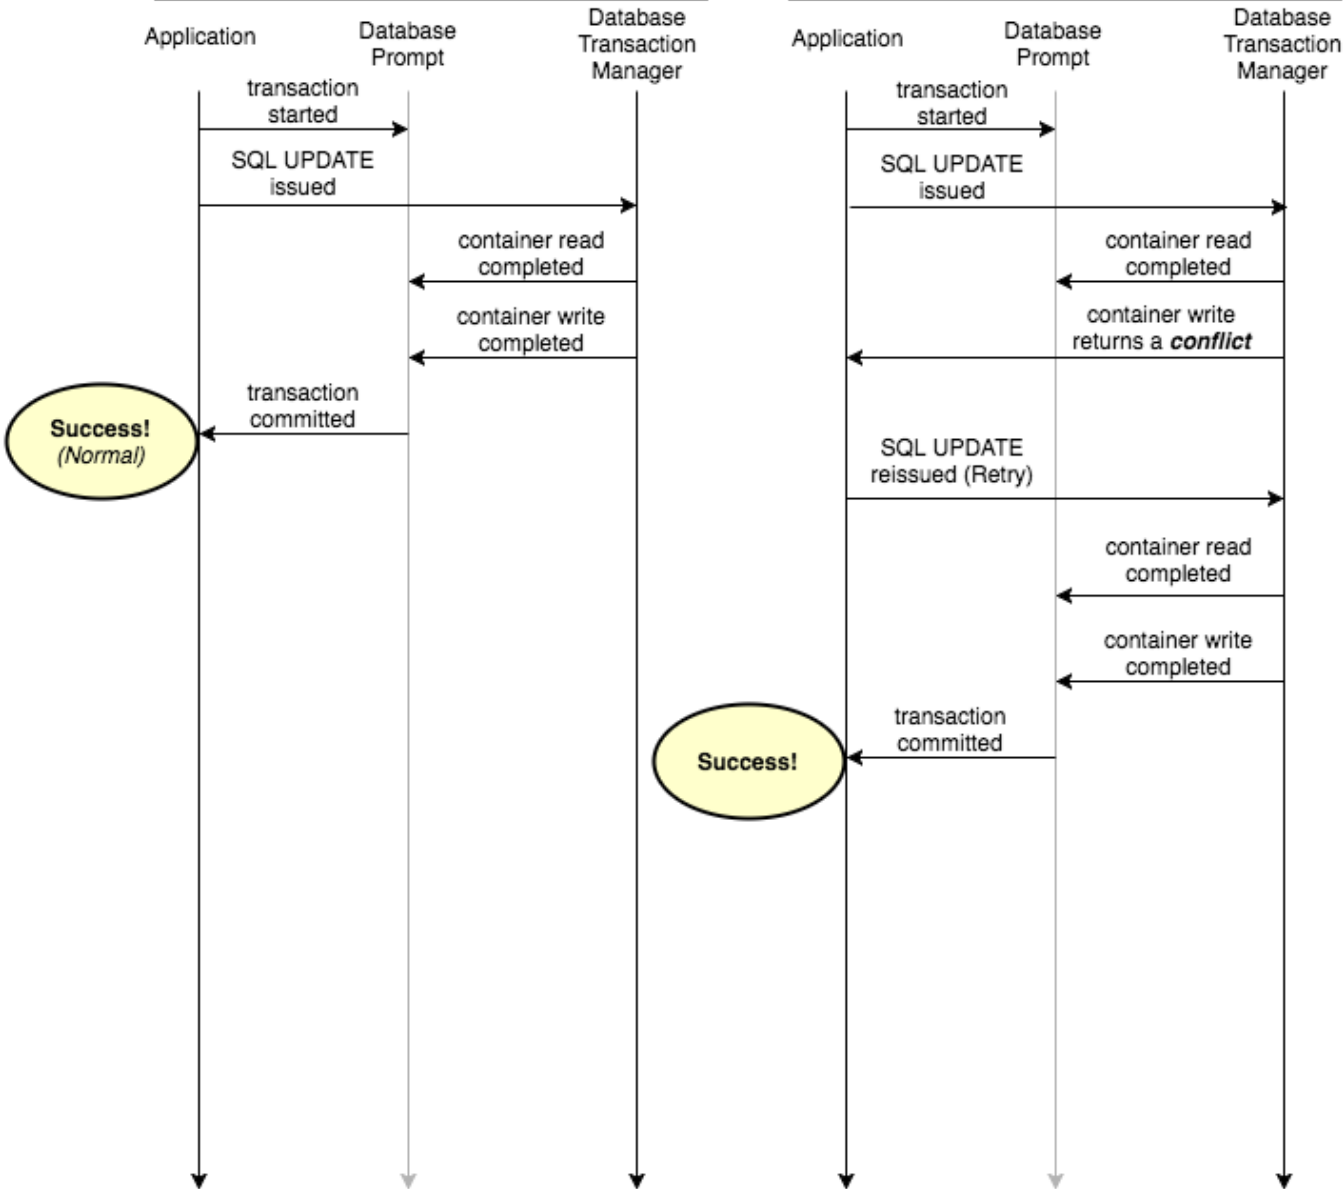
```
ERROR 1062 (23000): [5120] Duplicate key in container:
```

To illustrate a case when those errors could arise, let's compare Optimistic MVCC to standard MVCC. This sample shows two concurrent writers accessing different columns of the same row at the exact same time (rare). The left-hand diagram shows processing for Optimistic MVCC, as used for In-Memory Tables. The right-hand diagram shows similar processing using ClustrixDB's Standard MVCC processing, as used for Persistent Tables.

Scenario 1
Optimistic MVCC - In Memory Tables

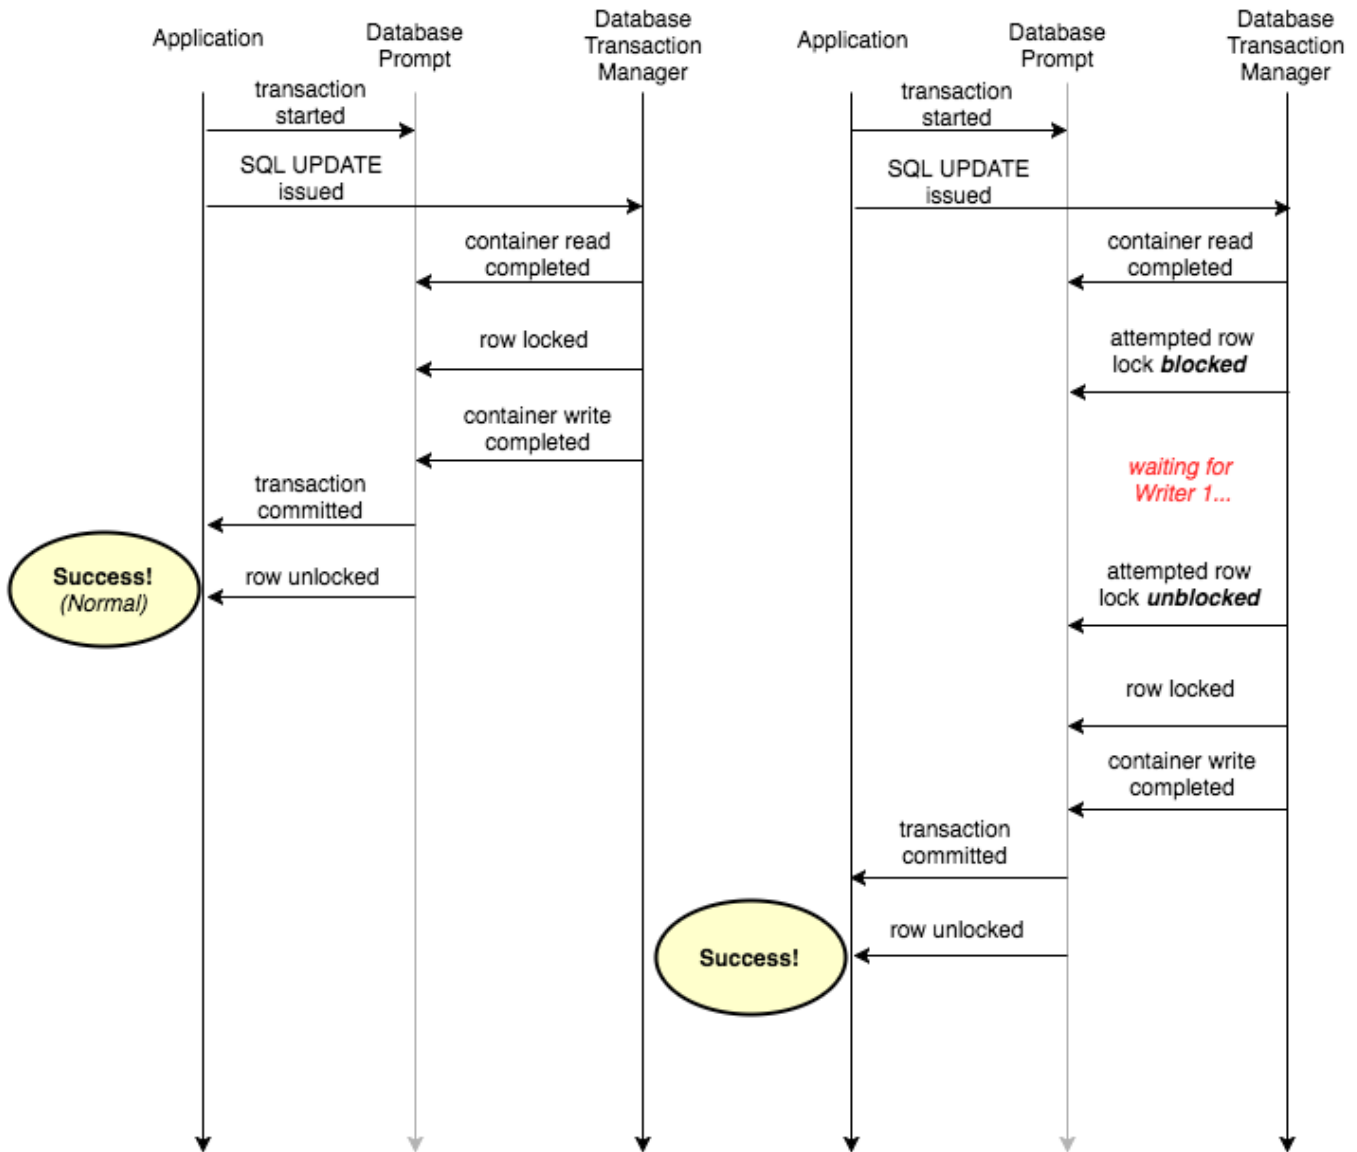
Writer 1

Writer 2
(Same exact row, same time, diff cols)



Scenario 2
Standard MVCC - Persistent Tables

Writer 1 **Writer 2**
(Same exact row, same time, diff cols)



Garbage Collection for In-Memory Tables

ClustrixDB maintains version history in undo logs for all table types. For In-Memory tables, this undo information is stored in a separate log file in memory. As with persistent tables, entries are purged from this log once no transactions reference them. You may observe that long-running transactions (i.e. ALTER TABLE or TRUNCATE TABLE) consume memory until they complete.

Rebalancer Details

As with persistent data, the ClustrixDB Rebalancer detects imbalances in data distribution and load for In-Memory tables and will redistribute data for those tables with a background process. If a node is lost, the Rebalancer will also reprotect In-Memory tables to ensure that the requisite number of replicas are maintained.

Global Variables for Memory Management

ClustrixDB provides the following controls for managing the amount of memory usable by In-Memory tables:

Variable	Description	Default
max_memory_table_limit_mb	Maximum amount of memory usable by in-memory tables.	16

The maximum amount of memory that can be allocated to In-Memory tables is up to 50% of available memory claimed by the ClustrixDB process.

These global variables control how memory limits are enforced and when alerts are sent if those limits are exceeded. For information on ClustrixDB's Alerter, please see [Database Alerts](#).

Variable	Description	Default
memory_table_system_full_error_percentage	Fail system queries when space usage for in-memory tables surpasses this percentage.	97
memory_table_system_full_warn_percentage	Warn about system queries when space usage for in-memory tables surpasses this percentage.	95
memory_table_user_full_error_percentage	Fail user queries when space usage for in-memory tables surpasses this percentage.	90
memory_table_user_full_warn_percentage	Warn about user queries when space usage for in-memory tables surpasses this percentage.	80

The [Health Dashboard of the ClustrixGUI Administrative Tool](#) also displays information about your cluster's memory utilization.

Summary of In-Memory Table Features

The following chart summarizes the differences between In-Memory and persistent tables.

Physical Properties	In-Memory	Persistent Storage Comparison
Storage	In-Memory (exclusively)	On disk (exclusively)
Row Size Limit	32K (enforced on INSERT/UPDATE)	64MB (enforced on INSERT/UPDATE)
Container Structure (data and indices)	skiplist	layered, btree
ClustrixDB Compatibility	In-Memory	Persistent Storage Comparison
Standard SQL DDL and DML Interface	Yes	Yes
MVCC	Optimistic MVCC	MVCC
Row Locking/Latches	No	Yes
Sliced	Yes	Yes
Replicated	Yes	Yes
Distributed/Balanced Cluster-wide	Yes	Yes
Fault Tolerant (if you lose a node, ClustrixDB reprotects)	No	Yes
Join In-Memory tables and persistent tables	Yes	Yes
Foreign Key Support	Yes	Yes
Utilizes Sierra Query Planner and Execution	Yes	Yes
Views	Yes	Yes
Indexes	Yes	Yes
Triggers	Yes	Yes
Stored Procedures	Yes	Yes
Partitioned Tables	Yes	Yes
Temporary Tables	Yes	Yes
ACID Compliance	In-Memory	Persistent Storage Comparison

Atomic, Consistent, and Isolated transactions cluster-wide	Yes	Yes
Durable data and transactions	No	Yes
Utilities/Other	In-Memory	Persistent Storage Comparison
clustrix_import	Yes	Yes
LOAD DATA INFILE	Yes	Yes
Backup and Restore	Yes	Yes
mysqldump	Yes	Yes
Online Schema Changes	Yes	Yes