


Global Variables

This section describes the ClustrixDB global and session variables. All variables are global. Some variables may also be set at the session level. Session variables override the global setting for the duration of a specific session and are noted below.

- [Displaying Variable Values](#)
 - [Finding Non-Default Globals](#)
- [Setting Variable Values](#)
- [Global Variables](#)

For the majority of workloads, Clustrix recommends retaining the default values for global variables.

 Please contact [Clustrix Support](#) with specific questions about modifying any of the default values as the product does not warn of inadvisable settings.

Displaying Variable Values

To display values for global or session variables, use the following syntax:

```
SHOW [GLOBAL | SESSION] VARIABLES [LIKE
pattern]
SHOW GLOBAL          VARIABLES [LIKE
pattern] [WITH DEFAULTS]
```

The WITH DEFAULTS option shows whether the variable is DEFAULTED, which specifies that the global variable will always be set to the default value recommended by Clustrix. If Clustrix determines that a variable's value should be set to a different default value, variables with DEFAULTED = 1 will automatically be modified as part of an upgrade.

Finding Non-Default Globals

The following query will list of variables whose values differ from the default, excluding variables whose values are not subject to defaults:

```
sql> SELECT name, value, default_value
FROM system.global_variables
JOIN system.global_variable_definitions USING (name)
WHERE value != default_value
AND name NOT IN
('cluster_id', 'cluster_name', 'clustrix_version', 'customer_name', 'format_version',
'global_variables_ignored_version', 'license', 'mysql_port', 'server_id',
'ssl_cert', 'ssl_key', 'view_strmaps_upgraded');
```

Setting Variable Values

To set a global or session variable to a specific value:

```
SET [GLOBAL | SESSION] variable_name = desired_value;
```

To modify a global variable to always use the Clustrix-recommended default value (DEFAULTED = 1):

```
sql> SET GLOBAL
variable_name =
DEFAULT;
```

To set a global variable to a default value with DEFAULTED = 0:

```
SET GLOBAL variable_name = actual_default_value;
```

If a session system variable is modified, the value remains in effect within your session until the session ends or the variable is set to another value. No other sessions are affected.

If you change a global variable, the value is applied to new sessions, but not sessions that are already open (including the session where the SET GLOBAL statement occurred).

Global Variables

Name	Description	Default Value	Session Variable
------	-------------	---------------	------------------

autocommit	Each statement will be its own transaction, and automatically applied to the database.	true	✓
autoretry	If enabled, when a transaction is interrupted by a group change or encounters a retrieable error, the database will automatically retry some in-process transactions. Only transactions that were submitted with autocommit = 1 or the first statement of an explicit transaction are retried. Stored procedure and function calls are never retried. If the retried statements are not executed successfully, the application will receive an error.	true	✓
auto_increment_increment	Amount in which auto increment values increase by, by default.	1	✓
auto_increment_offset	Value where auto increment values start at by default.	1	✓
backup_backup_concurrency	The number of tables that can be backed up simultaneously.	1	
backup_restore_concurrency	The maximum number of slices restored concurrently on each node.	16	
backup_write_compression_level	Compression level from 1 (fastest) to 9 (best compression)	6	
binlog_checksum	Always NONE. Clustrix masters do not support generating event checksums.	NONE	
binlog_format	Force all binlogs to log in this format, unless set to 'DEFAULT'. Valid values : statement, row. Read-only for global	DEFAULT	✓
character_set_client		latin1	✓
character_set_connection		latin1	✓
character_set_database	Dummy variable for compatibility. Must be utf8.	utf8	✓
character_set_results		latin1	✓
character_set_server	The default character set for databases when they are created	utf8	✓
cluster_id	64-bit cluster id	<auto populated>	
cluster_name	Name of the cluster	<auto populated>	
clustrix_version	Clustrix software version	<auto populated>	
collation_connection		latin1_swedish_ci	✓
collation_database	The collation used by the default database. This value cannot be modified.	utf8_general_ci	✓
collation_server	The default collation for databases when they are created	utf8_general_ci	✓
consistent_order	Force consistent ordering. See Guide.	false	✓
customer_name	This should be set on cluster formation. Used in Alerts / Warnings.	<auto populated>	
databasefull_message_interval_s	Database almost full message interval in seconds.	120	
databasefull_system_error_percentage	Fail system queries when space usage surpasses this percentage.	97	
databasefull_system_warn_percentage	Warn about system queries when space usage surpasses this percentage.	95	
databasefull_user_error_percentage	Fail user queries when space usage surpasses this percentage.	90	
databasefull_user_warn_percentage	Warn about user queries when space usage surpasses this percentage.	80	
debug_deadlocks	Attempt to get and log conflicting transaction session and statement information on distributed deadlock detection.	false	

device_auto_resize_to_largest	Automatically resize all (online) devices in the cluster to match the largest device	true	
device_temporary_space_limit_bytes	Maximum number of bytes allowed to be used for temporary containers.	5368709120	
foreign_key_cascade_limit	Limit of nested or cyclic foreign key cascading	3	
foreign_key_checks	Enable/Disable foreign key checks.	true	✓
format_version	A version string indicating the version of the software when the cluster was initially created.	<auto populated>	
gtid_mode	Always OFF. Clustrix masters do not support generating GTID events.	OFF	
gtid_purged	Dummy variable for compatibility. (Clustrix does not support replication with Global Transaction Identifiers.)		
gtm_schedule_til	Enable the Completely Fair Scheduler.	true	
gtm_schedule_til_batch_rows	Rows to process before rescheduling.	100	
hash_dist_min_slices	The default number of slices used when a table or index is created. Set this to 0 to automatically equal the number of the nodes in the cluster.	0	✓
have_query_cache	Dummy variable for compatibility.	NO	
hostname	Dummy variable for compatibility.		✓
idle_trx_timeout_s	Maximum allowed age for idle transactions. Specify 0 for no timeout.	120	✓
innodb_flush_log_at_trx_commit	Dummy variable for compatibility.	1	
interactive_timeout	Dummy variable for compatibility.	28800	✓
internode_latency_warning_us	If internode latency exceeds this, send a warning to clustrix.log. Setting to 0 turns warnings off.	0	
jdbcCompliantTruncation	JDBC compliant truncation check	false	✓
language	Dummy variable for compatibility.	/usr/local/mysql/share/mysql/english	
last_auto_increment_relation	The relation to which we last inserted an auto_increment value	0	✓
lc_time_names	Dummy variable for compatibility.	en_US	✓
license	The license string for the cluster		
lockman_max_locks	The maximum number of locks the lock manager will hold on each node in the cluster.	5000000	
lockman_max_transaction_locks	The maximum number of locks a single transaction can hold on each node in the cluster.	1000000	
lock_on_insert_select	Acquire a read lock on the source data when using INSERT INTO...SELECT FROM statements. This is necessary for correct statement based replication.	false	
lock_wait_timeout_ms	Milliseconds a query waits for a lock before timing out.	300000	
log_bin	Dummy variable for compatibility.	ON	
lower_case_table_names	Table names are stored in the case specified in the CREATE TABLE statement and name comparisons are not case sensitive. The value of this variable does NOT correspond to MySQL.	1	
master_status_binlog	Binlog used in SHOW MASTER STATUS when used without specifying a binlog.		✓
max_allowed_packet	Maximum allowed query size	16777216	
max_connections	The maximum number of connections allowed per node	500	

max_failures	Number of nodes or zones that can fail simultaneously without losing data or the ability to resolve transactions	1	
max_memlog_keep_files	number of old log files the memlog should keep	30	
max_memlog_message_bytes	maximum number of bytes to log in a single log message	1048576	
max_memory_table_limit_mb	Maximum amount of memory usable by in-memory tables.	16	
max_sierra_opt_mem_MiB	Cease Sierra planner optimization attempts once we have used this many total MiB and attempt to continue with the best plan found so far.	80	
max_sierra_parse_mem_MiB	Halt the Sierra planner and return an error to the user if this many MiB have been consumed during the initial parse phase. This variable is independent of the other max_sierra_%_mem_MiB variables.	256	
max_sierra_plans	Stop the sierra planner once the planner has been working this long on a plan. Specify 0 to disable this check.	10	
max_sierra_total_mem_MiB	Halt the Sierra planner and return an error to the user once this many total MiB have been used. This value should always be at least 50MiB greater than max_sierra_opt_mem_MiB.	160	
max_sierra_working_mem_MiB	Halt the Sierra planner and return an error to the user if the working set memory exceeds this many MiB limit. This variable is independent of the other max_sierra_%_mem_MiB variables.	512	
max_slices	The maximum allowed number of slices for a representation.	2000	
max_tables	The maximum allowed number of tables. Do not change.	2000	
memlog_rollover_hours	memlog will start a new file after this many hours (0 disables time-based rollover)	0	
memlog_rollover_size_MiB	memlog will start a new file if the current one exceeds this size (mb)	1024	
memory_table_system_full_error_percentage	Fail system writes when memory usage for in-memory tables surpasses this percentage.	97	
memory_table_system_full_warn_percentage	Warn about system writes when memory usage for in-memory tables surpasses this percentage.	95	
memory_table_user_full_error_percentage	Fail user writes when memory usage for in-memory tables surpasses this percentage.	90	
memory_table_user_full_warn_percentage	Warn about user writes when memory usage for in-memory tables surpasses this percentage.	80	
mysql_default_db_replication_policy	Replicate databases not specified in mysql_slave_db_replication_policy.	true	
mysql_default_table_replication_policy	Replicate tables not specified in mysql_table_replication_policy.	true	
mysql_float_format		mysql51	✓
mysql_port	Default port for mysql access to Clustrix.	3306	
mysql_relay_log_bytes	Maximum size of relay log in bytes a slave process is allowed to create.	67108864	
mysql_slave_batch_kb_limit		512	
mysql_version	The reported MySQL server version	5.0.45	
net_buffer_length	Dummy variable for compatibility.	16384	
net_write_timeout	Timeout in seconds if no data is received from a client to close the connection.	60	✓
port	Default port for mysql access to Clustrix.	3306	
qrc_enabled	Enable the Query Results Cache.	false	
query_cache_size	Dummy variable for compatibility.	0	

query_cache_type	Dummy variable for compatibility.	OFF	✓
query_fanout	Enable query fanout. This takes precedence over all other fanout variables.	true	✓
query_fanout_all_writes	Enable fanout for INSERT, UPDATE, and DELETE queries. Order of writes is not guaranteed.	false	✓
query_fanout_insert_select	Enable fanout for INSERT INTO ... SELECT FROM ... queries. This takes precedence over query_fanout_all_writes.	true	✓
read_only	Enable/Disable read only mode.	false	
rebalancer_global_task_limit	Maximum number of simultaneous rebalancer operations.	16	
rebalancer_rebalance_task_limit	Maximum number of operations that rebalancer_imbalanced and rebalancer_rebalance_distribution will each schedule at once.	2	
rebalancer_rebalance_threshold	Minimum coefficient of overall write load variation that will trigger rebalance activity.	0.05	
rebalancer_reprotect_queue_interval_s	Queued replicas count as healthy for this many seconds, to give missing nodes the chance to come back online before rebalancer_reprotect starts copying.	600	
rebalancer_split_threshold_kb	Default size at which the rebalancer splits slices.	8388608	
rebalancer_vdev_task_limit	Maximum number of simultaneous rebalancer operations targeting one device.	1	
server_id	Server ID for Clustrix as Replication Master.	1	
session_id	Session ID for client connections.	0	✓
session_log_bad_queries	Log BAD queries to the query.log.	false	✓
session_log_ddl	Log DDL statements to query.log.	true	
session_log_error_queries	Log ERROR statements to query.log.	true	
session_log_slow_queries	Log SLOW statements to query.log.	true	
session_log_slow_threshold_ms	Query duration threshold in milliseconds before logging this query.	10000	✓
session_log_users	Log users and LOGIN/LOGOUT to user.log.	false	
sigma_skiplist	Enable skiplist containers for aggregates and sorting.	false	✓
slave_max_authorized_packet	Maximum allowed packet size in bytes for the slave.	16777216	
sql_auto_increment_is_null	If true, and <col> is an auto_increment column, testing for <col> IS NULL will match the last row inserted	false	✓
sql_log_bin	Log statements to binary logs. This variable can be set to FALSE on a per-session basis.	true	✓
sql_mode	ClustrixDB provides limited support for SQL_MODE.	STRICT_TRANS_TABLES	✓
sql_notes	Dummy variable for compatibility	false	✓
sql_quote_show_create		true	✓
sql_safe_updates	Dummy variable for compatibility.	false	✓
sql_select_limit		18446744073709551615	✓
ssl_cert	SSL public key certificate file	server-cert.pem	
ssl_enabled	SSL is enabled	false	

ssl_key	SSL private key file	server-key.pem	
sync_binlog	Dummy variable for compatibility.	0	
system_time_zone	Time Zone. Must be set in Olson time zone format. See Guide.	UTC	
task_rebalancer_rebalance_distribution_interval_ms	Milliseconds between runs of periodic task "rebalancer_rebalance_distribution". Specify 0 to disable periodic task.	30000	
task_rebalancer_rebalance_interval_ms	Milliseconds between runs of periodic task "rebalancer_rebalance". Specify 0 to disable periodic task.	30000	
task_rebalancer_reprotect_interval_ms	Milliseconds between runs of periodic task "rebalancer_reprotect". Specify 0 to disable periodic task.	15000	
task_rebalancer_split_interval_ms	Milliseconds between runs of periodic task "rebalancer_split". Specify 0 to disable periodic task.	30000	
task_rebalancer_zone_balance_interval_ms	Milliseconds between runs of periodic task "rebalancer_zone_balance". Specify 0 to disable periodic task.	60000	
task_rebalancer_zone_missing_interval_ms	Milliseconds between runs of periodic task "rebalancer_zone_missing". Specify 0 to disable periodic task.	300000	
trxshoot_disk_min_pct	Kill the oldest transaction when available disk space goes below this percentage and available undo space goes below trxshoot_undo_min_pct.	5	
trxshoot_undo_min_pct	Kill the oldest transaction when available undo space goes below this percentage and available disk space goes below trxshoot_disk_min_pct.	2	
trx_timeout_s	Maximum allowed age for transactions. Specify 0 for no timeout.	0	✓
tx_sync_commit	Controls when the client is notified of successful commit. 'RELAXED' is lower latency, but may result in 'committed' transactions being lost.	STRICT	✓
unique_checks		true	✓
version	Concatenation of mysql_version and clustrix_version.	<auto populated>	
wait_timeout	Timeout in seconds after data is sent to a client to close the connection.	28800	✓
write_trx_timeout_s	Maximum allowed age for write transactions. Specify 0 for no timeout.	0	✓