

DATA ANALYTICS REFERENCE DOCUMENT	
Document Title:	MySQL deamon commandline help
Document No.:	1548325727
Author(s):	Gerhard van der Linde
Contributor(s):	

REVISION HISTORY

Revision	Details of Modification(s)	Reason for modification	Date	By
0	Draft release	MySQL deamon commandline help	2019/01/24 10:28	Gerhard van der Linde

MySQLD Help

```
mysqld.exe Ver 8.0.14 for Win64 on x86_64 (MySQL Community Server - GPL)
Copyright (c) 2000, 2019, Oracle and/or its affiliates. All rights reserved.
```

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

Starts the MySQL database server.

```
Usage: mysqld.exe [OPTIONS]
NT and Win32 specific options:
--install                         Install the default service (NT).
--install-manual                   Install the default service started manually (NT).
--install service_name              Install an optional service (NT).
--install-manual service_name      Install an optional service started manually (NT).
--remove                           Remove the default service from the service list (NT).
--remove service_name               Remove the service_name from the service list (NT).
--enable-named-pipe                Only to be used for the default server (NT).
--standalone                       Dummy option to start as a standalone server (NT).
```

Default options are read from the following files in the given order:

```
C:\WINDOWS\my.ini C:\WINDOWS\my.cnf C:\my.ini C:\my.cnf C:\Users\121988\Downloads\mysql-8.0.14-win64\my.ini C:\Users\121988\Downloads\mysql-8.0.14-win64\my.cnf
```

The following groups are read: mysqld server mysqld-8.0

The following options may be given as the first argument:

```
--print-defaults                  Print the program argument list and exit.
--no-defaults                     Don't read default options from any option file,
                                 except for login file.
--defaults-file=#                Only read default options from the given file #.
--defaults-extra-file=#          Read this file after the global files are read.
--defaults-group-suffix=#        Also read groups with concat(group, suffix)
--login-path=#                   Read this path from the login file.

--abort-slave-event-count=#
```

```
        Option used by mysql-test for debugging and testing of
        replication.

--activate-all-roles-on-login
        Automatically set all granted roles as active after the
        user has authenticated successfully.

--admin-address=name
        IP address to bind to for service connection. Address can
        be an IPv4 address, IPv6 address, or host name. Wildcard
        values *, ::, 0.0.0.0 are not allowed.

--admin-port=#    Port number to use for service connection, built-in
                  default (33062)

--allow-suspicious-udfs
        Allows use of UDFs consisting of only one symbol xxx()
        without corresponding xxx_init() or xxx_deinit(). That
        also means that one can load any function from any
        library, for example exit() from libc.so

-a, --ansi
        Use ANSI SQL syntax instead of MySQL syntax. This mode
        will also set transaction isolation level 'serializable'.

--archive[=name]
        Enable or disable ARCHIVE plugin. Possible values are ON,
        OFF, FORCE (don't start if the plugin fails to load).

--auto-generate-certs
        Auto generate SSL certificates at server startup if --ssl
        is set to ON and none of the other SSL system variables
        are specified and certificate/key files are not present
        in data directory.
        (Defaults to on; use --skip-auto-generate-certs to disable.)

--auto-increment-increment[=#]
        Auto-increment columns are incremented by this

--auto-increment-offset[=#]
        Offset added to Auto-increment columns. Used when
        auto-increment-increment != 1

--autocommit
        Set default value for autocommit (0 or 1)
        (Defaults to on; use --skip-autocommit to disable.)

--automatic-sp-privileges
        Creating and dropping stored procedures alters ACLs
        (Defaults to on; use --skip-automatic-sp-privileges to disable.)

--avoid-temporal-upgrade
        When this option is enabled, the pre-5.6.4 temporal types
        are not upgraded to the new format for ALTER TABLE
        requests ADD/CHANGE/MODIFY COLUMN, ADD INDEX or FORCE
        operation. This variable is deprecated and will be
        removed in a future release.

--back-log=#    The number of outstanding connection requests MySQL can
                  have. This comes into play when the main MySQL thread
                  gets very many connection requests in a very short time

-b, --basedir=name
        Path to installation directory. All paths are usually
        resolved relative to this

--big-tables
        Allow big result sets by saving all temporary sets on
        file (Solves most 'table full' errors)

--bind-address=name
        IP address(es) to bind to. Syntax: address[,address]...
        where address can be an IPv4 address, IPv6 address, host
        name or one of the wildcard values *, ::, 0.0.0.0. In
        case more than one address is specified in a
        comma-separated list, wildcard values are not allowed.

--binlog-cache-size=
        The size of the transactional cache for updates to
        transactional engines for the binary log. If you often
        use transactions containing many statements, you can
        increase this to get more performance

--binlog-checksum=name
        Type of BINLOG_CHECKSUM_ALG. Include checksum for log
        events in the binary log. Possible values are NONE and
        CRC32; default is CRC32.
```

```
--binlog-direct-non-transactional-updates
    Causes updates to non-transactional engines using
    statement format to be written directly to binary log.
    Before using this option make sure that there are no
    dependencies between transactional and non-transactional
    tables such as in the statement INSERT INTO t_myisam
    SELECT * FROM t_innodb; otherwise, slaves may diverge
    from the master.

--binlog-do-db=name Tells the master it should log updates for the specified
                     database, and exclude all others not explicitly
                     mentioned.

--binlog-encryption Enable/disable binary and relay logs encryption.

--binlog-error-action=name
    When statements cannot be written to the binary log due
    to a fatal error, the server can either ignore the error
    and let the master continue, or abort.

--binlog-expire-logs-seconds=# 
    If non-zero, binary logs will be purged after
    binlog_expire_logs_seconds seconds; If both this option
    and expire_logs_days are set to non-zero values, this
    option takes priority. Purges happen at startup and at
    binary log rotation.

--binlog-format=name
    What form of binary logging the master will use: either
    ROW for row-based binary logging, STATEMENT for
    statement-based binary logging, or MIXED. MIXED is
    statement-based binary logging except for those
    statements where only row-based is correct: those which
    involve user-defined functions (i.e. UDFs) or the UUID()
    function; for those, row-based binary logging is
    automatically used. If NDBCLUSTER is enabled and
    binlog-format is MIXED, the format switches to row-based
    and back implicitly per each query accessing an
    NDBCLUSTER table.

--binlog-group-commit-sync-delay=# 
    The number of microseconds the server waits for the
    binary log group commit sync queue to fill before
    continuing. Default: 0. Min: 0. Max: 1000000.

--binlog-group-commit-sync-no-delay-count=# 
    If there are this many transactions in the commit sync
    queue and the server is waiting for more transactions to
    be enqueued (as set using
    --binlog-group-commit-sync-delay), the commit procedure
    resumes.

--binlog-gtid-simple-recovery
    If this option is enabled, the server does not open more
    than two binary logs when initializing GTID_PURGED and
    GTID_EXECUTED, either during server restart or when
    binary logs are being purged. Enabling this option is
    useful when the server has already generated many binary
    logs without GTID events (e.g., having GTID_MODE = OFF).
    Note: If this option is enabled, GLOBAL.GTID_EXECUTED and
    GLOBAL.GTID_PURGED may be initialized wrongly in two
    cases: (1) All binary logs were generated by MySQL 5.7.5
    or older, and GTID_MODE was ON for some binary logs but
    OFF for the newest binary log. (2) The oldest existing
    binary log was generated by MySQL 5.7.5 or older, and SET
    GTID_PURGED was issued after the oldest binary log was
    generated. If a wrong set is computed in one of case (1)
    or case (2), it will remain wrong even if the server is
    later restarted with this option disabled.
    (Defaults to on; use --skip-binlog-gtid-simple-recovery to disable.)

--binlog-ignore-db=name
```

Tells the master that updates to the given database should not be logged to the binary log.

--binlog-max-flush-queue-time=#
The maximum time that the binary log group commit will keep reading transactions before it flush the transactions to the binary log (and optionally sync, depending on the value of sync_binlog).

--binlog-order-commits
Issue internal commit calls in the same order as transactions are written to the binary log. Default is to order commits.
(Defaults to on; use --skip-binlog-order-commits to disable.)

--binlog-rotate-encryption-master-key-at-startup
Force binlog encryption master key rotation at startup

--binlog-row-event-max-size=#
The maximum size of a row-based binary log event in bytes. Rows will be grouped into events smaller than this size if possible. The value has to be a multiple of 256.

--binlog-row-image=name
Controls whether rows should be logged in 'FULL', 'NOBLOB' or 'MINIMAL' formats. 'FULL', means that all columns in the before and after image are logged. 'NOBLOB', means that mysqld avoids logging blob columns whenever possible (eg, blob column was not changed or is not part of primary key). 'MINIMAL', means that a PK equivalent (PK columns or full row if there is no PK in the table) is logged in the before image, and only changed columns are logged in the after image. (Default: FULL).

--binlog-row-metadata=name
Controls whether metadata is logged using FULL or MINIMAL format. FULL causes all metadata to be logged; MINIMAL means that only metadata actually required by slave is logged. Default: MINIMAL.

--binlog-row-value-options=name
When set to PARTIAL_JSON, this option enables a space-efficient row-based binary log format for UPDATE statements that modify a JSON value using only the functions JSON_SET, JSON_REPLACE, and JSON_REMOVE. For such updates, only the modified parts of the JSON document are included in the binary log, so small changes of big documents may need significantly less space.

--binlog-rows-query-log-events
Allow writing of Rows_query_log events into binary log.

--binlog-stmt-cache-size=#
The size of the statement cache for updates to non-transactional engines for the binary log. If you often use statements updating a great number of rows, you can increase this to get more performance

--binlog-transaction-dependency-history-size=#
Maximum number of rows to keep in the writeset history.

--binlog-transaction-dependency-tracking=name
Selects the source of dependency information from which to assess which transactions can be executed in parallel by the slave's multi-threaded applier. Possible values are COMMIT_ORDER, WRITESET and WRITESET_SESSION.

--blackhole[=name] Enable or disable BLACKHOLE plugin. Possible values are ON, OFF, FORCE (don't start if the plugin fails to load).

--block-encryption-mode=name
mode for AES_ENCRYPT/AES_DECRYPT

--bulk-insert-buffer-size=#
Size of tree cache used in bulk insert optimisation. Note that this is a limit per thread!

```
--caching-sha2-password-auto-generate-rsa-keys
    Auto generate RSA keys at server startup if corresponding
    system variables are not specified and key files are not
    present at the default location.
    (Defaults to on; use --skip-caching-sha2-password-auto-generate-rsa-keys
to disable.)
--caching-sha2-password-private-key-path=name
    A fully qualified path to the private RSA key used for
    authentication.
--caching-sha2-password-public-key-path=name
    A fully qualified path to the public RSA key used for
    authentication.
--character-set-client-handshake
    Don't ignore client side character set value sent during
    handshake.
    (Defaults to on; use --skip-character-set-client-handshake to disable.)
--character-set-filesystem=name
    Set the filesystem character set.
-C, --character-set-server=name
    Set the default character set.
--character-sets-dir=name
    Directory where character sets are
--check-proxy-users If set to FALSE (the default), then proxy user identity
    will not be mapped for authentication plugins which
    support mapping from grant tables. When set to TRUE,
    users associated with authentication plugins which signal
    proxy user mapping should be done according to GRANT
    PROXY privilege definition.
-r, --chroot=name Chroot mysqld daemon during startup.
--collation-server=name
    Set the default collation.
--completion-type=name
    The transaction completion type, one of NO_CHAIN, CHAIN,
    RELEASE
--concurrent-insert[=name]
    Use concurrent insert with MyISAM. Possible values are
    NEVER, AUTO, ALWAYS
--connect-timeout=# The number of seconds the mysqld server is waiting for a
    connect packet before responding with 'Bad handshake'
    Write error output on screen; don't remove the console
    window on windows.
--core-file      Write core on errors.
--create-admin-listener-thread
    Use a dedicated thread for listening incoming connections
    on admin interface
--cte-max-recursion-depth=#
    Abort a recursive common table expression if it does more
    than this number of iterations.
-h, --datadir=name Path to the database root directory
--default-authentication-plugin=name
    The default authentication plugin used by the server to
    hash the password.
--default-password-lifetime=#
    The number of days after which the password will expire.
--default-storage-engine=name
    The default storage engine for new tables
--default-time-zone=name
    Set the default time zone.
--default-tmp-storage-engine=name
    The default storage engine for new explicit temporary
    tables
--default-week-format=#
    The default week format used by WEEK() functions
```

```
--delay-key-write[=name]
    Type of DELAY_KEY_WRITE
--delayed-insert-limit=#
    After inserting delayed_insert_limit rows, the INSERT
    DELAYED handler will check if there are any SELECT
    statements pending. If so, it allows these to execute
    before continuing. This variable is deprecated along with
    INSERT DELAYED.
--delayed-insert-timeout=#
    How long a INSERT DELAYED thread should wait for INSERT
    statements before terminating. This variable is
    deprecated along with INSERT DELAYED.
--delayed-queue-size=#
    What size queue (in rows) should be allocated for
    handling INSERT DELAYED. If the queue becomes full, any
    client that does INSERT DELAYED will wait until there is
    room in the queue again. This variable is deprecated
    along with INSERT DELAYED.
--disabled-storage-engines=name
    Limit CREATE TABLE for the storage engines listed
--disconnect-on-expired-password
    Give clients that don't signal password expiration
    support execution time error(s) instead of connection
    error
    (Defaults to on; use --skip-disconnect-on-expired-password to disable.)
--disconnect-slave-event-count=#
    Option used by mysql-test for debugging and testing of
    replication.
--div-precision-increment=#
    Precision of the result of '/' operator will be increased
    on that value
--early-plugin-load=name
    Optional semicolon-separated list of plugins to load
    before storage engine initialization, where each plugin
    is identified as name=library, where name is the plugin
    name and library is the plugin library in plugin_dir.
--end-markers-in-json
    In JSON output ("EXPLAIN FORMAT=JSON" and optimizer
    trace), if variable is set to 1, repeats the structure's
    key (if it has one) near the closing bracket
--enforce-gtid-consistency[=name]
    Prevents execution of statements that would be impossible
    to log in a transactionally safe manner. Currently, the
    disallowed statements include CREATE TEMPORARY TABLE
    inside transactions, all updates to non-transactional
    tables, and CREATE TABLE ... SELECT.
--eq-range-index-dive-limit=#
    The optimizer will use existing index statistics instead
    of doing index dives for equality ranges if the number of
    equality ranges for the index is larger than or equal to
    this number. If set to 0, index dives are always used.
--event-scheduler[=name]
    Enable the event scheduler. Possible values are ON, OFF,
    and DISABLED (keep the event scheduler completely
    deactivated, it cannot be activated run-time)
-T, --exit-info[=#] Used for debugging. Use at your own risk.
--expire-logs-days=#
    If non-zero, binary logs will be purged after
    expire_logs_days days; If this option alone is set on the
    command line or in a configuration file, it overrides the
    default value for binlog-expire-logs-seconds. If both
    options are set to nonzero values,
    binlog-expire-logs-seconds takes priority. Possible
```

```
purges happen at startup and at binary log rotation.
--explicit-defaults-for-timestamp
    This option causes CREATE TABLE to create all TIMESTAMP
    columns as NULL with DEFAULT NULL attribute. Without this
    option, TIMESTAMP columns are NOT NULL and have implicit
    DEFAULT clauses. The old behavior is deprecated. The
    variable can only be set by users having the SUPER
    privilege.
    (Defaults to on; use --skip-explicit-defaults-for-timestamp to disable.)
--external-locking
    Use system (external) locking (disabled by default).
    With this option enabled you can run myisamchk to test
    (not repair) tables while the MySQL server is running.
    Disable with --skip-external-locking.
--federated[=name]
    Enable or disable FEDERATED plugin. Possible values are
    ON, OFF, FORCE (don't start if the plugin fails to load).
--flush
    Flush MyISAM tables to disk between SQL commands
--flush-time=#
    A dedicated thread is created to flush all tables at the
    given interval
--ft-boolean-syntax=name
    List of operators for MATCH ... AGAINST ( ... IN BOOLEAN
    MODE)
--ft-max-word-len=#
    The maximum length of the word to be included in a
    FULLTEXT index. Note: FULLTEXT indexes must be rebuilt
    after changing this variable
--ft-min-word-len=#
    The minimum length of the word to be included in a
    FULLTEXT index. Note: FULLTEXT indexes must be rebuilt
    after changing this variable
--ft-query-expansion-limit=#
    Number of best matches to use for query expansion
--ft-stopword-file=name
    Use stopwords from this file instead of built-in list
--gdb
    Set up signals usable for debugging.
--general-log
    Log connections and queries to a table or log file.
    Defaults to logging to a file hostname.log, or if
    --log-output=TABLE is used, to a table mysql.general_log.
--general-log-file=name
    Log connections and queries to given file
--group-concat-max-len=#
    The maximum length of the result of function
    GROUP_CONCAT()
--group-replication-consistency[=name]
    Transaction consistency guarantee, possible values:
    EVENTUAL, BEFORE_ON_PRIMARY_FAILOVER, BEFORE, AFTER,
    BEFORE_AND_AFTER
--gtid-executed-compression-period[=#]
    When binlog is disabled, a background thread wakes up to
    compress the gtid_executed table every
    gtid_executed_compression_period transactions, as a
    special case, if variable is 0, the thread never wakes up
    to compress the gtid_executed table.
--gtid-mode=name
    Controls whether Global Transaction Identifiers (GTIDs)
    are enabled. Can be OFF, OFF_PERMISSIVE, ON_PERMISSIVE,
    or ON. OFF means that no transaction has a GTID.
    OFF_PERMISSIVE means that new transactions (committed in
    a client session using GTID_NEXT='AUTOMATIC') are not
    assigned any GTID, and replicated transactions are
    allowed to have or not have a GTID. ON_PERMISSIVE means
    that new transactions are assigned a GTID, and replicated
    transactions are allowed to have or not have a GTID. ON
    means that all transactions have a GTID. ON is required
    on a master before any slave can use
    MASTER_AUTO_POSITION=1. To safely switch from OFF to ON,
    first set all servers to OFF_PERMISSIVE, then set all
```

```
servers to ON_PERMISSIVE, then wait for all transactions
without a GTID to be replicated and executed on all
servers, and finally set all servers to GTID_MODE = ON.
-?, --help          Display this help and exit.
--histogram-generation-max-mem-size=#
                  Maximum amount of memory available for generating
                  histograms
--host-cache-size=# How many host names should be cached to avoid resolving.
--information-schema-stats-expiry=#
                  The number of seconds after which mysqld server will
                  fetch data from storage engine and replace the data in
                  cache.
--init-connect=name Command(s) that are executed for each new connection
--init-file=name    Read SQL commands from this file at startup
--init-slave=name   Command(s) that are executed by a slave server each time
                   the SQL thread starts
-I, --initialize    Create the default database and exit. Create a super user
                   with a random expired password and store it into the log.
--initialize-insecure
                  Create the default database and exit. Create a super user
                  with empty password.
--innodb           Deprecated option. Provided for backward compatibility
                   only. The option has no effect on the server behaviour.
                   InnoDB is always enabled. The option will be removed in a
                   future release.
--innodb-adaptive-flushing
                  Attempt flushing dirty pages to avoid IO bursts at
                  checkpoints.
                  (Defaults to on; use --skip-innodb-adaptive-flushing to disable.)
--innodb-adaptive-flushing-lwm=#
                  Percentage of log capacity below which no adaptive
                  flushing happens.
--innodb-adaptive-hash-index
                  Enable InnoDB adaptive hash index (enabled by default).
                  Disable with --skip-innodb-adaptive-hash-index.
                  (Defaults to on; use --skip-innodb-adaptive-hash-index to disable.)
--innodb-adaptive-hash-index-parts[=#]
                  Number of InnoDB Adapative Hash Index Partitions.
                  (default = 8).
--innodb-adaptive-max-sleep-delay=#
                  The upper limit of the sleep delay in usec. Value of 0
                  disables it.
--innodb-api-bk-commit-interval[=#]
                  Background commit interval in seconds
--innodb-api-disable-rowlock
                  Disable row lock when direct access InnoDB through InnoDB
                  APIs
--innodb-api-enable-binlog
                  Enable binlog for applications direct access InnoDB
                  through InnoDB APIs
--innodb-api-enable-mdl
                  Enable MDL for applications direct access InnoDB through
                  InnoDB APIs
--innodb-api-trx-level[=#]
                  InnoDB API transaction isolation level
--innodb-autoextend-increment=#
                  Data file autoextend increment in megabytes
--innodb-autoinc-lock-mode=#
                  The AUTOINC lock modes supported by InnoDB: 0 => Old
                  style AUTOINC locking (for backward compatibility); 1 =>
                  New style AUTOINC locking; 2 => No AUTOINC locking
                  (unsafe for SBR)
--innodb-buffer-pool-chunk-size=#
```

```
Size of a single memory chunk within each buffer pool
instance for resizing buffer pool. Online buffer pool
resizing happens at this granularity. 0 means disable
resizing buffer pool.

--innodb-buffer-pool-dump-at-shutdown
    Dump the buffer pool into a file named
    @innodb_buffer_pool_filename
    (Defaults to on; use --skip-innodb-buffer-pool-dump-at-shutdown to
disable.)

--innodb-buffer-pool-dump-now
    Trigger an immediate dump of the buffer pool into a file
    named @innodb_buffer_pool_filename

--innodb-buffer-pool-dump-pct=#
    Dump only the hottest N% of each buffer pool, defaults to
    25

--innodb-buffer-pool-filename=name
    Filename to/from which to dump/load the InnoDB buffer
    pool

--innodb-buffer-pool-in-core-file
    This option has no effect if @@core_file is OFF. If
    @@core_file is ON, and this option is OFF, then the core
    dump file will be generated only if it is possible to
    exclude buffer pool from it. As soon as it will be
    determined that such exclusion is impossible a warning
    will be emitted and @@core_file will be set to OFF to
    prevent generating a core dump. If this option is enabled
    (which is the default), then core dumping logic will not
    be affected.
    (Defaults to on; use --skip-innodb-buffer-pool-in-core-file to disable.)

--innodb-buffer-pool-instances=#
    Number of buffer pool instances, set to higher value on
    high-end machines to increase scalability

--innodb-buffer-pool-load-abort
    Abort a currently running load of the buffer pool

--innodb-buffer-pool-load-at-startup
    Load the buffer pool from a file named
    @innodb_buffer_pool_filename
    (Defaults to on; use --skip-innodb-buffer-pool-load-at-startup to
disable.)

--innodb-buffer-pool-load-now
    Trigger an immediate load of the buffer pool from a file
    named @innodb_buffer_pool_filename

--innodb-buffer-pool-size=#
    The size of the memory buffer InnoDB uses to cache data
    and indexes of its tables.

--innodb-change-buffer-max-size=#
    Maximum on-disk size of change buffer in terms of
    percentage of the buffer pool.

--innodb-change-buffering=name
    Buffer changes to reduce random access: OFF, ON,
    inserting, deleting, changing, or purging.

--innodb-checksum-algorithm=name
    The algorithm InnoDB uses for page checksumming. Possible
    values are CRC32 (hardware accelerated if the CPU
    supports it) write crc32, allow any of the other
    checksums to match when reading; STRICT_CRC32 write
    crc32, do not allow other algorithms to match when
    reading; INNODB write a software calculated checksum,
    allow any other checksums to match when reading;
    STRICT_INNODB write a software calculated checksum, do
    not allow other algorithms to match when reading; NONE
    write a constant magic number, do not do any checksum
    verification when reading; STRICT_NONE write a constant
```

```
magic number, do not allow values other than that magic
number when reading; Files updated when this option is
set to crc32 or strict_crc32 will not be readable by
MySQL versions older than 5.6.3
--innodb-cmp-per-index-enabled
    Enable INFORMATION_SCHEMA.innodb_cmp_per_index, may have
    negative impact on performance (off by default)
--innodb-commit-concurrency=#
    Helps in performance tuning in heavily concurrent
    environments.
--innodb-compression-failure-threshold-pct[=#]
    If the compression failure rate of a table is greater
    than this number more padding is added to the pages to
    reduce the failures. A value of zero implies no padding
--innodb-compression-level=#
    Compression level used for compressed row format. 0 is
    no compression, 1 is fastest, 9 is best compression and
    default is 6.
--innodb-compression-pad-pct-max[=#]
    Percentage of empty space on a data page that can be
    reserved to make the page compressible.
--innodb-concurrency-tickets=#
    Number of times a thread is allowed to enter InnoDB
    within the same SQL query after it has once got the
    ticket
--innodb-data-file-path=name
    Path to individual files and their sizes.
--innodb-data-home-dir=name
    The common part for InnoDB table spaces.
--innodb-deadlock-detect
    Enable/disable InnoDB deadlock detector (default ON). if
    set to OFF, deadlock detection is skipped, and we rely on
    innodb_lock_wait_timeout in case of deadlock.
    (Defaults to on; use --skip-innodb-deadlock-detect to disable.)
--innodb-dedicated-server
    Automatically scale innodb_buffer_pool_size and
    innodb_log_file_size based on system memory. Also set
    innodb_flush_method=0_DIRECT_NO_FSYNC, if supported
--innodb-default-row-format=name
    The default ROW FORMAT for all innodb tables created
    without explicit ROW_FORMAT. Possible values are
    REDUNDANT, COMPACT, and DYNAMIC. The ROW_FORMAT value
    COMPRESSED is not allowed
--innodb-directories=name
    List of directories 'dir1;dir2;..;dirN' to scan for
    tablespace files. Default is to scan
    'innodb-data-home-dir;innodb-undo-directory;datadir'
--innodb-disable-sort-file-cache
    Whether to disable OS system file cache for sort I/O
--innodb-doublewrite
    Enable InnoDB doublewrite buffer (enabled by default).
    Disable with --skip-innodb-doublewrite.
    (Defaults to on; use --skip-innodb-doublewrite to disable.)
--innodb-fast-shutdown[=#]
    Speeds up the shutdown process of the InnoDB storage
    engine. Possible values are 0, 1 (faster) or 2 (fastest -
    crash-like).
--innodb-file-per-table
    Stores each InnoDB table to an .ibd file in the database
    dir.
    (Defaults to on; use --skip-innodb-file-per-table to disable.)
--innodb-fill-factor=#
    Percentage of B-tree page filled during bulk insert
```

```
--innodb-flush-log-at-timeout[=#]
    Write and flush logs every (n) second.
--innodb-flush-log-at-trx-commit[=#]
    Set to 0 (write and flush once per second), 1 (write and
    flush at each commit), or 2 (write at commit, flush once
    per second).
--innodb-flush-method=name
    With which method to flush data
--innodb-flush-neighbors[=#]
    Set to 0 (don't flush neighbors from buffer pool), 1
    (flush contiguous neighbors from buffer pool) or 2 (flush
    neighbors from buffer pool), when flushing a block
--innodb-flush-sync Allow IO bursts at the checkpoints ignoring io_capacity
    setting.
    (Defaults to on; use --skip-innodb-flush-sync to disable.)
--innodb-flushing-avg-loops=#
    Number of iterations over which the background flushing
    is averaged.
--innodb-force-load-corrupted
    Force InnoDB to load metadata of corrupted table.
--innodb-force-recovery=#
    Helps to save your data in case the disk image of the
    database becomes corrupt.
--innodb-fsync-threshold=#
    The value of this variable determines how often InnoDB
    calls fsync when creating a new file. Default is zero
    which would make InnoDB flush the entire file at once
    before closing it.
--innodb-ft-aux-table
    FTS internal auxiliary table to be checked
--innodb-ft-cache-size=#
    InnoDB Fulltext search cache size in bytes
--innodb-ft-enable-diag-print
    Whether to enable additional FTS diagnostic printout
--innodb-ft-enable-stopword
    Create FTS index with stopword.
    (Defaults to on; use --skip-innodb-ft-enable-stopword to disable.)
--innodb-ft-max-token-size=#
    InnoDB Fulltext search maximum token size in characters
--innodb-ft-min-token-size=#
    InnoDB Fulltext search minimum token size in characters
--innodb-ft-num-word-optimize[=#]
    InnoDB Fulltext search number of words to optimize for
    each optimize table call
--innodb-ft-result-cache-limit=#
    InnoDB Fulltext search query result cache limit in bytes
--innodb-ft-server-stopword-table[=name]
    The user supplied stopword table name.
--innodb-ft-sort-poll-degree=#
    InnoDB Fulltext search parallel sort degree, will round
    up to nearest power of 2 number
--innodb-ft-total-cache-size=#
    Total memory allocated for InnoDB Fulltext Search cache
--innodb-ft-user-stopword-table[=name]
    User supplied stopword table name, effective in the
    session level.
--innodb-io-capacity=#
    Number of IOPs the server can do. Tunes the background IO
    rate
--innodb-io-capacity-max=#
    Limit to which innodb_io_capacity can be inflated.
--innodb-lock-wait-timeout=#
    Timeout in seconds an InnoDB transaction may wait for a
```

```
lock before being rolled back. Values above 100000000
disable the timeout.
--innodb-log-buffer-size=#
    The size of the buffer which InnoDB uses to write log to
    the log files on disk.
--innodb-log-checksums
    Whether to compute and require checksums for InnoDB redo
    log blocks
    (Defaults to on; use --skip-innodb-log-checksums to disable.)
--innodb-log-compressed-pages
    Enables/disables the logging of entire compressed page
    images. InnoDB logs the compressed pages to prevent
    corruption if the zlib compression algorithm changes.
    When turned OFF, InnoDB will assume that the zlib
    compression algorithm doesn't change.
    (Defaults to on; use --skip-innodb-log-compressed-pages to disable.)
--innodb-log-file-size=#
    Size of each log file (in bytes).
--innodb-log-files-in-group=#
    Number of log files (when multiplied by
    innodb_log_file_size gives total size of log files).
    InnoDB writes to files in a circular fashion.
--innodb-log-group-home-dir=name
    Path to InnoDB log files.
--innodb-log-spin-cpu-abs-lwm=#
    Minimum value of cpu time for which spin-delay is used.
    Expressed in percentage of single cpu core.
--innodb-log-spin-cpu-pct-hwm=#
    Maximum value of cpu time for which spin-delay is used.
    Expressed in percentage of all cpu cores.
--innodb-log-wait-for-flush-spin-hwm=#
    Maximum value of average log flush time for which
    spin-delay is used. When flushing takes longer, user
    threads no longer spin when waiting forflushed redo.
    Expressed in microseconds.
--innodb-log-write-ahead-size=#
    Log write ahead unit size to avoid read-on-write, it
    should match the OS cache block IO size.
--innodb-lru-scan-depth=#
    How deep to scan LRU to keep it clean
--innodb-max-dirty-pages-pct=#
    Percentage of dirty pages allowed in bufferpool.
--innodb-max-dirty-pages-pct-lwm=#
    Percentage of dirty pages at which flushing kicks in.
--innodb-max-purge-lag=#
    Desired maximum length of the purge queue (0 = no limit)
--innodb-max-purge-lag-delay=#
    Maximum delay of user threads in micro-seconds
--innodb-max-undo-log-size[=#]
    Maximum size of an UNDO tablespace in MB (If an UNDO
    tablespace grows beyond this size it will be truncated in
    due course).
--innodb-monitor-disable=name
    Turn off a monitor counter
--innodb-monitor-enable=name
    Turn on a monitor counter
--innodb-monitor-reset=name
    Reset a monitor counter
--innodb-monitor-reset-all=name
    Reset all values for a monitor counter
--innodb-old-blocks-pct=#
    Percentage of the buffer pool to reserve for 'old'
    blocks.
```

```
--innodb-old-blocks-time=#  
    Move blocks to the 'new' end of the buffer pool if the  
    first access was at least this many milliseconds ago. The  
    timeout is disabled if 0.  
--innodb-online-alter-log-max-size=#  
    Maximum modification log file size for online index  
    creation  
--innodb-open-files=#  
    How many files at the maximum InnoDB keeps open at the  
    same time.  
--innodb-optimize-fulltext-only  
    Only optimize the Fulltext index of the table  
--innodb-page-cleaners[=#]  
    Page cleaner threads can be from 1 to 64. Default is 4.  
--innodb-page-size[=#]  
    Page size to use for all InnoDB tablespaces.  
--innodb-parallel-read-threads=#  
    Number of threads to do parallel read.  
--innodb-print-all-deadlocks  
    Print all deadlocks to MySQL error log (off by default)  
--innodb-print-ddl-logs  
    Print all DDL logs to MySQL error log (off by default)  
--innodb-purge-batch-size[=#]  
    Number of UNDO log pages to purge in one batch from the  
    history list.  
--innodb-purge-rseg-truncate-frequency[=#]  
    Dictates rate at which UNDO records are purged. Value N  
    means purge rollback segment(s) on every Nth iteration of  
    purge invocation  
--innodb-purge-threads[=#]  
    Purge threads can be from 1 to 32. Default is 4.  
--innodb-random-read-ahead  
    Whether to use read ahead for random access within an  
    extent.  
--innodb-read-ahead-threshold=#  
    Number of pages that must be accessed sequentially for  
    InnoDB to trigger a readahead.  
--innodb-read-io-threads=#  
    Number of background read I/O threads in InnoDB.  
--innodb-read-only Start InnoDB in read only mode (off by default)  
--innodb-redo-log-encrypt  
    Enable or disable Encryption of REDO tablespace.  
--innodb-replication-delay=#  
    Replication thread delay (ms) on the slave server if  
    innodb_thread_concurrency is reached (0 by default)  
--innodb-rollback-on-timeout  
    Roll back the complete transaction on lock wait timeout,  
    for 4.x compatibility (disabled by default)  
--innodb-rollback-segments[=#]  
    Number of rollback segments per tablespace. This applies  
    to the system tablespace, the temporary tablespace & any  
    undo tablespace.  
--innodb-sort-buffer-size=#  
    Memory buffer size for index creation  
--innodb-spin-wait-delay[=#]  
    Maximum delay between polling for a spin lock (6 by  
    default)  
--innodb-stats-auto-recalc  
    InnoDB automatic recalculation of persistent statistics  
    enabled for all tables unless overridden at table level  
    (automatic recalculation is only done when InnoDB decides  
    that the table has changed too much and needs a new  
    statistics)
```

```
(Defaults to on; use --skip-innodb-stats-auto-recalc to disable.)  
--innodb-stats-include-delete-marked  
    Include delete marked records when calculating persistent  
    statistics  
--innodb-stats-method=name  
    Specifies how InnoDB index statistics collection code  
    should treat NULLs. Possible values are NULLS_EQUAL  
    (default), NULLS_UNEQUAL and NULLS_IGNORED  
--innodb-stats-on-metadata  
    Enable statistics gathering for metadata commands such as  
    SHOW TABLE STATUS for tables that use transient  
    statistics (off by default)  
--innodb-stats-persistent  
    InnoDB persistent statistics enabled for all tables  
    unless overridden at table level  
    (Defaults to on; use --skip-innodb-stats-persistent to disable.)  
--innodb-stats-persistent-sample-pages=#  
    The number of leaf index pages to sample when calculating  
    persistent statistics (by ANALYZE, default 20)  
--innodb-stats-transient-sample-pages=#  
    The number of leaf index pages to sample when calculating  
    transient statistics (if persistent statistics are not  
    used, default 8)  
--innodb-status-file  
    Enable SHOW ENGINE INNODB STATUS output in the  
    innodb_status.<pid> file  
--innodb-status-output  
    Enable InnoDB monitor output to the error log.  
--innodb-status-output-locks  
    Enable InnoDB lock monitor output to the error log.  
    Requires innodb_status_output=ON.  
--innodb-strict-mode  
    Use strict mode when evaluating create options.  
    (Defaults to on; use --skip-innodb-strict-mode to disable.)  
--innodb-sync-array-size[=#]  
    Size of the mutex/lock wait array.  
--innodb-sync-spin-loops=#  
    Count of spin-loop rounds in InnoDB mutexes (30 by  
    default)  
--innodb-table-locks  
    Enable InnoDB locking in LOCK TABLES  
    (Defaults to on; use --skip-innodb-table-locks to disable.)  
--innodb-temp-data-file-path=name  
    Path to files and their sizes making temp-tablespace.  
--innodb-temp-tablespaces-dir=name  
    Directory where temp tablespace files live, this path can  
    be absolute.  
--innodb-thread-concurrency=#  
    Helps in performance tuning in heavily concurrent  
    environments. Sets the maximum number of threads allowed  
    inside InnoDB. Value 0 will disable the thread  
    throttling.  
--innodb-thread-sleep-delay=#  
    Time of innodb thread sleeping before joining InnoDB  
    queue (usec). Value 0 disable a sleep  
--innodb-tmpdir[=name]  
    Directory for temporary non-tablespace files.  
--innodb-undo-directory=name  
    Directory where undo tablespace files live, this path can  
    be absolute.  
--innodb-undo-log-encrypt  
    Enable or disable Encrypt of UNDO tablespace.  
--innodb-undo-log-truncate
```

```
        Enable or Disable Truncate of UNDO tablespace.  
        (Defaults to on; use --skip-innodb-undo-log-truncate to disable.)  
--innodb-undo-tablespaces=#  
        Number of undo tablespaces to use. (deprecated)  
--innodb-use-native-aio  
        Use native AIO if supported on this platform.  
        (Defaults to on; use --skip-innodb-use-native-aio to disable.)  
--innodb-write-io-threads=#  
        Number of background write I/O threads in InnoDB.  
--interactive-timeout=#  
        The number of seconds the server waits for activity on an  
        interactive connection before closing it  
--internal-tmp-disk-storage-engine[=name]  
        The default storage engine for on-disk internal temporary  
        tables.  
--internal-tmp-mem-storage-engine=name  
        The default storage engine for in-memory internal  
        temporary tables.  
--join-buffer-size=#  
        The size of the buffer that is used for full joins  
--keep-files-on-create  
        Don't overwrite stale .MYD and .MYI even if no directory  
        is specified  
--key-buffer-size=# The size of the buffer used for index blocks for MyISAM  
        tables. Increase this to get better index handling (for  
        all reads and multiple writes) to as much as you can  
        afford  
--key-cache-age-threshold=#  
        This characterizes the number of hits a hot block has to  
        be untouched until it is considered aged enough to be  
        downgraded to a warm block. This specifies the percentage  
        ratio of that number of hits to the total number of  
        blocks in key cache  
--key-cache-block-size=#  
        The default size of key cache blocks  
--key-cache-division-limit=#  
        The minimum percentage of warm blocks in key cache  
--keyring-migration-destination=name  
        Keyring plugin to which the keys are migrated to. This  
        option must be specified along with  
        --keyring-migration-source.  
--keyring-migration-host=name  
        Connect to host.  
-p, --keyring-migration-password[=name]  
        Password to use when connecting to server during keyring  
        migration. If password value is not specified then it  
        will be asked from the tty.  
--keyring-migration-port=#  
        Port number to use for connection.  
--keyring-migration-socket=name  
        The socket file to use for connection.  
--keyring-migration-source=name  
        Keyring plugin from where the keys needs to be migrated  
        to. This option must be specified along with  
        --keyring-migration-destination.  
--keyring-migration-user=name  
        User to login to server.  
-L, --language=name Client error messages in given language. May be given as  
        a full path. Deprecated. Use --lc-messages-dir instead.  
--lc-messages=name Set the language used for the error messages.  
--lc-messages-dir=name  
        Directory where error messages are  
--lc-time-names=name
```

Set the language used for the month names and the days of the week.

--local-infile Enable LOAD DATA LOCAL INFILE

--lock-wait-timeout=# Timeout in seconds to wait for a lock before returning an error.

--log-bin[=name] Configures the name prefix to use for binary log files. If the --log-bin option is not supplied, the name prefix defaults to "binlog". If the --log-bin option is supplied without argument, the name prefix defaults to "HOSTNAME-bin", where HOSTNAME is the machine's hostname. To set a different name prefix for binary log files, use --log-bin=name. To disable binary logging, use the --skip-log-bin or --disable-log-bin option.

--log-bin-index=name File that holds the names for binary log files.

--log-bin-trust-function-creators If set to FALSE (the default), then when --log-bin is used, creation of a stored function (or trigger) is allowed only to users having the SUPER privilege and only if this stored function (trigger) may not break binary logging. Note that if ALL connections to this server ALWAYS use row-based binary logging, the security issues do not exist and the binary logging cannot break, so you can safely set this to TRUE

--log-bin-use-v1-row-events If equal to 1 then version 1 row events are written to a row based binary log. If equal to 0, then the latest version of events are written. This option is useful during some upgrades.

--log-error[=name] Error log file

--log-error-services=name Services that should be called when an error event is received

--log-error-suppression-list=name Comma-separated list of error-codes. Error messages corresponding to these codes will not be included in the error log. Only events with a severity of Warning or Information can be suppressed; events with System or Error severity will always be included. Requires the filter 'log_filter_internal' to be set in @global.log_error_services, which is the default.

--log-error-verbosity=# How detailed the error log should be. 1, log errors only. 2, log errors and warnings. 3, log errors, warnings, and notes. Messages sent to the client are unaffected by this setting.

--log-isam[=name] Log all MyISAM changes to file.

--log-output=name Syntax: log-output=value[,value...], where "value" could be TABLE, FILE or NONE

--log-queries-not-using-indexes Log queries that are executed without benefit of any index to the slow log if it is open

--log-raw Log to general log before any rewriting of the query. For use in debugging, not production as sensitive information may be logged.

--log-short-format Don't log extra information to update and slow-query logs.

--log-slave-updates Tells the slave to log the updates from the slave thread to the binary log.
(Defaults to on; use --skip-log-slave-updates to disable.)

--log-slow-admin-statements Log slow OPTIMIZE, ANALYZE, ALTER and other

```
--log-slow-extra      administrative statements to the slow log if it is open.
--log-slow-slave-statements  Print more attributes to the slow query log file. Has no
                            effect on logging to table.
--log-statements-unsafe-for-binlog   Log slow statements executed by slave thread to the slow
                            log if it is open.
--log-tc=name        Path to transaction coordinator log (used for
                            transactions that affect more than one storage engine,
                            when binary log is disabled).
--log-tc-size=#      Size of transaction coordinator log.
--log-throttle-queries-not-using-indexes=#  Log at most this many 'not using index' warnings per
                            minute to the slow log. Any further warnings will be
                            condensed into a single summary line. A value of 0
                            disables throttling. Option has no effect unless
                            --log_queries_not_using_indexes is set.
--log-timestamps=name    UTC to timestamp log files in zulu time, for more concise
                            timestamps and easier correlation of logs from servers
                            from multiple time zones, or SYSTEM to use the system's
                            local time. This affects only log files, not log tables,
                            as the timestamp columns of the latter can be converted
                            at will.
--long-query-time=#   Log all queries that have taken more than long_query_time
                            seconds to execute to file. The argument will be treated
                            as a decimal value with microsecond precision
--low-priority-updates  INSERT/DELETE/UPDATE has lower priority than selects
--lower-case-table-names[=#]  If set to 1 table names are stored in lowercase on disk
                            and table names will be case-insensitive. Should be set
                            to 2 if you are using a case insensitive file system
--mandatory-roles=name  All the specified roles are always considered granted to
                            every user and they can't be revoked. Mandatory roles
                            still require activation unless they are made into
                            default roles. The granted roles will not be visible in
                            the mysql.role_edges table.
--master-info-file=name  The location and name of the file that remembers the
                            master and where the I/O replication thread is in the
                            master's binlogs.
--master-info-repository=name  Defines the type of the repository for the master
                            information.
--master-retry-count=#   The number of tries the slave will make to connect to the
                            master before giving up. Deprecated option, use 'CHANGE
                            MASTER TO master_retry_count = <num>' instead.
--master-verify-checksum  Force checksum verification of logged events in binary
                            log before sending them to slaves or printing them in
                            output of SHOW BINLOG EVENTS. Disabled by default.
--max-allowed-packet=#   Max packet length to send to or receive from the server
--max-binlog-cache-size=#  Sets the total size of the transactional cache
--max-binlog-dump-events=#  Option used by mysql-test for debugging and testing of
```

```
replication.  
--max-binlog-size=# Binary log will be rotated automatically when the size  
exceeds this value. Will also apply to relay logs if  
max_relay_log_size is 0  
--max-binlog-stmt-cache-size=#  
Sets the total size of the statement cache  
--max-connect-errors=#  
If there is more than this number of interrupted  
connections from a host this host will be blocked from  
further connections  
--max-connections=# The number of simultaneous clients allowed  
--max-delayed-threads=#  
Don't start more than this number of threads to handle  
INSERT DELAYED statements. If set to zero INSERT DELAYED  
will be not used. This variable is deprecated along with  
INSERT DELAYED.  
--max-digest-length=#  
Maximum length considered for digest text.  
--max-error-count=# Max number of errors/warnings to store for a statement  
--max-execution-time=#  
Kill SELECT statement that takes over the specified  
number of milliseconds  
--max-heap-table-size=#  
Don't allow creation of heap tables bigger than this  
--max-join-size=# Joins that are probably going to read more than  
max_join_size records return an error  
--max-length-for-sort-data=#  
Max number of bytes in sorted records  
--max-points-in-geometry[=#]  
Maximum number of points in a geometry  
--max-prepared-stmt-count=#  
Maximum number of prepared statements in the server  
--max-relay-log-size=#  
If non-zero: relay log will be rotated automatically when  
the size exceeds this value; if zero: when the size  
exceeds max_binlog_size  
--max-seeks-for-key=#  
Limit assumed max number of seeks when looking up rows  
based on a key  
--max-sort-length=# The number of bytes to use when sorting long values with  
PAD SPACE collations (only the first max_sort_length  
bytes of each value are used; the rest are ignored)  
--max-sp-recursion-depth[=#]  
Maximum stored procedure recursion depth  
--max-user-connections=#  
The maximum number of active connections for a single  
user (0 = no limit)  
--max-write-lock-count=#  
After this many write locks, allow some read locks to run  
in between  
--memlock Lock mysqld in memory.  
--min-examined-row-limit=#  
Don't write queries to slow log that examine fewer rows  
than that  
--myisam-block-size=#  
Block size to be used for MyISAM index pages  
--myisam-data-pointer-size=#  
Default pointer size to be used for MyISAM tables  
--myisam-max-sort-file-size=#  
Don't use the fast sort index method to created index if  
the temporary file would get bigger than this  
--myisam-mmap-size=#  
Restricts the total memory used for memory mapping of
```

MySQL tables
--myisam-recover-options[=name]
 Syntax: myisam-recover-options[=option[,option...]],
 where option can be DEFAULT, BACKUP, FORCE, QUICK, or OFF
--myisam-repair-threads=#
 If larger than 1, when repairing a MyISAM table all
 indexes will be created in parallel, with one thread per
 index. The value of 1 disables parallel repair
--myisam-sort-buffer-size=#
 The buffer that is allocated when sorting the index when
 doing a REPAIR or when creating indexes with CREATE INDEX
 or ALTER TABLE
--myisam-stats-method=name
 Specifies how MyISAM index statistics collection code
 should treat NULLs. Possible values of name are
 NULLS_UNEQUAL (default behavior for 4.1 and later),
 NULLS_EQUAL (emulate 4.0 behavior), and NULLS_IGNORED
--myisam-use-mmap Use memory mapping for reading and writing MyISAM tables
--mysql-native-password-proxy-users
 If set to FALSE (the default), then the
 mysql_native_password plugin will not signal for
 authenticated users to be checked for mapping to proxy
 users. When set to TRUE, the plugin will flag associated
 authenticated accounts to be mapped to proxy users when
 the server option check_proxy_users is enabled.
--mysqlx[=name] Enable or disable mysqlx plugin. Possible values are ON,
 OFF, FORCE (don't start if the plugin fails to load).
--mysqlx-bind-address[=name]
 Address to which X Plugin should bind the TCP socket.
--mysqlx-cache-cleaner[=name]
 Enable or disable mysqlx_cache_cleaner plugin. Possible
 values are ON, OFF, FORCE (don't start if the plugin
 fails to load).
--mysqlx-connect-timeout[=#]
 Maximum allowed waiting time for connection to setup a
 session (in seconds).
--mysqlx-document-id-unique-prefix[=#]
 Unique prefix is a value assigned by InnoDB cluster to
 the instance, which is meant to make document id unique
 across all replicaset from the same cluster
--mysqlx-idle-worker-thread-timeout[=#]
 Time after which an idle worker thread is terminated (in
 seconds).
--mysqlx-interactive-timeout[=#]
 Default value for "mysqlx_wait_timeout", when the
 connection is interactive. The value defines number of
 seconds that X Plugin must wait for activity on
 interactive connection
--mysqlx-max-allowed-packet[=#]
 Size of largest message that client is going to handle.
--mysqlx-max-connections[=#]
 Maximum number of concurrent X protocol connections.
 Actual number of connections is also affected by the
 general max_connections.
--mysqlx-min-worker-threads[=#]
 Minimal number of worker threads.
--mysqlx-port[=#] Port on which X Plugin is going to accept incoming
 connections.
--mysqlx-port-open-timeout[=#]
 How long X Plugin is going to retry binding of server
 socket (in case of failure)
--mysqlx-read-timeout[=#]
 Number of seconds that X Plugin must wait for blocking

```
        read operation to complete
--mysqlx-socket[=name]
        X Plugin's unix socket for local connection.
--mysqlx-ssl-ca=name
        CA file in PEM format.
--mysqlx-ssl-capath=name
        CA directory.
--mysqlx-ssl-cert=name
        X509 cert in PEM format.
--mysqlx-ssl-cipher=name
        SSL cipher to use.
--mysqlx-ssl-crl=name
        Certificate revocation list.
--mysqlx-ssl-crlpath=name
        Certificate revocation list path.
--mysqlx-ssl-key=name
        X509 key in PEM format.
--mysqlx-wait-timeout[=#]
        Number or seconds that X Plugin must wait for activity on
        noninteractive connection
--mysqlx-write-timeout[=#]
        Number or seconds that X Plugin must wait for blocking
        write operation to complete
--named-pipe
        Enable the named pipe (NT)
--named-pipe-full-access-group=name
        Name of Windows group granted full access to the named
        pipe
--net-buffer-length=#
        Buffer length for TCP/IP and socket communication
--net-read-timeout=#
        Number of seconds to wait for more data from a connection
        before aborting the read
--net-retry-count=# If a read on a communication port is interrupted, retry
        this many times before giving up
--net-write-timeout=#
        Number of seconds to wait for a block to be written to a
        connection before aborting the write
-n, --new
        Use very new possible "unsafe" functions
--ngram[=name]
        Enable or disable ngram plugin. Possible values are ON,
        OFF, FORCE (don't start if the plugin fails to load).
--ngram-token-size=#
        InnoDB ngram full text plugin parser token size in
        characters
--no-dd-upgrade
        Abort restart if automatic upgrade or downgrade of the
        data dictionary is needed.
--no-monitor
        Disable monitor process.
--offline-mode
        Make the server into offline mode
--old
        Use compatible behavior
--old-alter-table
        Use old, non-optimized alter table
--old-style-user-limits
        Enable old-style user limits (before 5.0.3, user
        resources were counted per each user+host vs. per
        account).
--open-files-limit=#
        If this is not 0, then mysqld will use this value to
        reserve file descriptors to use with setrlimit(). If this
        value is 0 then mysqld will reserve max_connections*5 or
        max_connections + table_open_cache*2 (whichever is
        larger) number of file descriptors
--optimizer-prune-level=#
        Controls the heuristic(s) applied during query
        optimization to prune less-promising partial plans from
        the optimizer search space. Meaning: 0 - do not apply any
```

```
heuristic, thus perform exhaustive search; 1 - prune
plans based on number of retrieved rows
--optimizer-search-depth=#
    Maximum depth of search performed by the query optimizer.
    Values larger than the number of relations in a query
    result in better query plans, but take longer to compile
    a query. Values smaller than the number of tables in a
    relation result in faster optimization, but may produce
    very bad query plans. If set to 0, the system will
    automatically pick a reasonable value
--optimizer-switch=name
    optimizer_switch=option=val[,option=val...], where option
    is one of {index_merge, index_merge_union,
    index_merge_sort_union, index_merge_intersection,
    engine_condition_pushdown, index_condition_pushdown, mrr,
    mrr_cost_based, materialization, semijoin, loosescan,
    firstmatch, duplicateweedout,
    subquery_materialization_cost_based, skip_scan,
    block_nested_loop, batched_key_access,
    use_index_extensions, condition_fanout_filter,
    derived_merge} and val is one of {on, off, default}
--optimizer-trace=name
    Controls tracing of the Optimizer:
    optimizer_trace=option=val[,option=val...], where option
    is one of {enabled, one_line} and val is one of {on,
    default}
--optimizer-trace-features=name
    Enables/disables tracing of selected features of the
    Optimizer:
    optimizer_trace_features=option=val[,option=val...],
    where option is one of {greedy_search, range_optimizer,
    dynamic_range, repeated_subselect} and val is one of {on,
    off, default}
--optimizer-trace-limit=#
    Maximum number of shown optimizer traces
--optimizer-trace-max-mem-size=#
    Maximum allowed cumulated size of stored optimizer traces
--optimizer-trace-offset=#
    Offset of first optimizer trace to show; see manual
--parser-max-mem-size=#
    Maximum amount of memory available to the parser
--password-history=#
    The number of old passwords to check in the history. Set
    to 0 (the default) to turn the checks off
--password-require-current
    Current password is needed to be specified in order to
    change it
--password-reuse-interval=#
    The minimum number of days that need to pass before a
    password can be reused. Set to 0 (the default) to turn
    the checks off
--performance-schema
    Enable the performance schema.
    (Defaults to on; use --skip-performance-schema to disable.)
--performance-schema-accounts-size=#
    Maximum number of instrumented user@host accounts. Use 0
    to disable, -1 for automated scaling.
--performance-schema-consumer-events-stages-current
    Default startup value for the events_stages_current
    consumer.
--performance-schema-consumer-events-stages-history
    Default startup value for the events_stages_history
    consumer.
```

```
--performance-schema-consumer-events-stages-history-long
    Default startup value for the events_stages_history_long
    consumer.
--performance-schema-consumer-events-statements-current
    Default startup value for the events_statements_current
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-events-
statements-current to disable.)
--performance-schema-consumer-events-statements-history
    Default startup value for the events_statements_history
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-events-
statements-history to disable.)
--performance-schema-consumer-events-statements-history-long
    Default startup value for the
    events_statements_history_long consumer.
--performance-schema-consumer-events-transactions-current
    Default startup value for the events_transactions_current
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-events-
transactions-current to disable.)
--performance-schema-consumer-events-transactions-history
    Default startup value for the events_transactions_history
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-events-
transactions-history to disable.)
--performance-schema-consumer-events-transactions-history-long
    Default startup value for the
    events_transactions_history_long consumer.
--performance-schema-consumer-events-waits-current
    Default startup value for the events_waits_current
    consumer.
--performance-schema-consumer-events-waits-history
    Default startup value for the events_waits_history
    consumer.
--performance-schema-consumer-events-waits-history-long
    Default startup value for the events_waits_history_long
    consumer.
--performance-schema-consumer-global-instrumentation
    Default startup value for the global_instrumentation
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-global-
instrumentation to disable.)
--performance-schema-consumer-statements-digest
    Default startup value for the statements_digest consumer.
    (Defaults to on; use --skip-performance-schema-consumer-statements-digest
to disable.)
--performance-schema-consumer-thread-instrumentation
    Default startup value for the thread_instrumentation
    consumer.
    (Defaults to on; use --skip-performance-schema-consumer-thread-
instrumentation to disable.)
--performance-schema-digests-size=#
    Size of the statement digest. Use 0 to disable, -1 for
    automated sizing.
--performance-schema-error-size=#
    Number of server errors instrumented.
--performance-schema-events-stages-history-long-size=#
    Number of rows in EVENTS_STAGES_HISTORY_LONG. Use 0 to
    disable, -1 for automated sizing.
--performance-schema-events-stages-history-size=#
    Number of rows per thread in EVENTS_STAGES_HISTORY. Use 0
    to disable, -1 for automated sizing.
```

```
--performance-schema-events-statements-history-long-size=#  
    Number of rows in EVENTS_STATEMENTS_HISTORY_LONG. Use 0  
    to disable, -1 for automated sizing.  
--performance-schema-events-statements-history-size=#  
    Number of rows per thread in EVENTS_STATEMENTS_HISTORY.  
    Use 0 to disable, -1 for automated sizing.  
--performance-schema-events-transactions-history-long-size=#  
    Number of rows in EVENTS_TRANSACTIONS_HISTORY_LONG. Use 0  
    to disable, -1 for automated sizing.  
--performance-schema-events-transactions-history-size=#  
    Number of rows per thread in EVENTS_TRANSACTIONS_HISTORY.  
    Use 0 to disable, -1 for automated sizing.  
--performance-schema-events-waits-history-long-size=#  
    Number of rows in EVENTS_WAITS_HISTORY_LONG. Use 0 to  
    disable, -1 for automated sizing.  
--performance-schema-events-waits-history-size=#  
    Number of rows per thread in EVENTS_WAITS_HISTORY. Use 0  
    to disable, -1 for automated sizing.  
--performance-schema-hosts-size=#  
    Maximum number of instrumented hosts. Use 0 to disable,  
    -1 for automated scaling.  
--performance-schema-instrument[=name]  
    Default startup value for a performance schema  
    instrument.  
--performance-schema-max-cond-classes=#  
    Maximum number of condition instruments.  
--performance-schema-max-cond-instances=#  
    Maximum number of instrumented condition objects. Use 0  
    to disable, -1 for automated scaling.  
--performance-schema-max-digest-length=#  
    Maximum length considered for digest text, when stored in  
    performance_schema tables.  
--performance-schema-max-digest-sample-age=#  
    The time in seconds after which a previous query sample  
    is considered old. When the value is 0, queries are  
    sampled once. When the value is greater than zero,  
    queries are re sampled if the last sample is more than  
    performance_schema_max_digest_sample_age seconds old.  
--performance-schema-max-file-classes=#  
    Maximum number of file instruments.  
--performance-schema-max-file-handles=#  
    Maximum number of opened instrumented files.  
--performance-schema-max-file-instances=#  
    Maximum number of instrumented files. Use 0 to disable,  
    -1 for automated scaling.  
--performance-schema-max-index-stat=#  
    Maximum number of index statistics for instrumented  
    tables. Use 0 to disable, -1 for automated scaling.  
--performance-schema-max-memory-classes=#  
    Maximum number of memory pool instruments.  
--performance-schema-max-metadata-locks=#  
    Maximum number of metadata locks. Use 0 to disable, -1  
    for automated scaling.  
--performance-schema-max-mutex-classes=#  
    Maximum number of mutex instruments.  
--performance-schema-max-mutex-instances=#  
    Maximum number of instrumented MUTEX objects. Use 0 to  
    disable, -1 for automated scaling.  
--performance-schema-max-prepared-statements-instances=#  
    Maximum number of instrumented prepared statements. Use 0  
    to disable, -1 for automated scaling.  
--performance-schema-max-program-instances=#  
    Maximum number of instrumented programs. Use 0 to
```

```
        disable, -1 for automated scaling.
--performance-schema-max-rwlock-classes=#
    Maximum number of rwlock instruments.
--performance-schema-max-rwlock-instances=#
    Maximum number of instrumented RWLOCK objects. Use 0 to
    disable, -1 for automated scaling.
--performance-schema-max-socket-classes=#
    Maximum number of socket instruments.
--performance-schema-max-socket-instances=#
    Maximum number of opened instrumented sockets. Use 0 to
    disable, -1 for automated scaling.
--performance-schema-max-sql-text-length=#
    Maximum length of displayed sql text.
--performance-schema-max-stage-classes=#
    Maximum number of stage instruments.
--performance-schema-max-statement-classes=#
    Maximum number of statement instruments.
--performance-schema-max-statement-stack=#
    Number of rows per thread in EVENTS_STATEMENTS_CURRENT.
--performance-schema-max-table-handles=#
    Maximum number of opened instrumented tables. Use 0 to
    disable, -1 for automated scaling.
--performance-schema-max-table-instances=#
    Maximum number of instrumented tables. Use 0 to disable,
    -1 for automated scaling.
--performance-schema-max-table-lock-stat=#
    Maximum number of lock statistics for instrumented
    tables. Use 0 to disable, -1 for automated scaling.
--performance-schema-max-thread-classes=#
    Maximum number of thread instruments.
--performance-schema-max-thread-instances=#
    Maximum number of instrumented threads. Use 0 to disable,
    -1 for automated scaling.
--performance-schema-session-connect-attrs-size=#
    Size of session attribute string buffer per thread. Use 0
    to disable, -1 for automated sizing.
--performance-schema-setup-actors-size=#
    Maximum number of rows in SETUP_ACTORS. Use 0 to disable,
    -1 for automated scaling.
--performance-schema-setup-objects-size=#
    Maximum number of rows in SETUP_OBJECTS. Use 0 to
    disable, -1 for automated scaling.
--performance-schema-users-size=#
    Maximum number of instrumented users. Use 0 to disable,
    -1 for automated scaling.
--persist-only-admin-x509-subject[=name]
    The client peer certificate name required to enable
    setting all system variables via SET PERSIST[_ONLY]
--persisted-globals-load
    When this option is enabled, config file mysqld-auto.cnf
    is read and applied to server, else this file is ignored
    even if present.
    (Defaults to on; use --skip-persisted-globals-load to disable.)
--pid-file=name      Pid file used by safe_mysqld
--plugin-dir=name   Directory for plugins
--plugin-load=name   Optional semicolon-separated list of plugins to load,
                    where each plugin is identified as name=library, where
                    name is the plugin name and library is the plugin library
                    in plugin_dir.
--plugin-load-add=name
    Optional semicolon-separated list of plugins to load,
    where each plugin is identified as name=library, where
    name is the plugin name and library is the plugin library
```

```
        in plugin_dir. This option adds to the list specified by
        --plugin-load in an incremental way. Multiple
        --plugin-load-add are supported.
-P, --port=#          Port number to use for connection or 0 to default to,
                      my.cnf, $MYSQL_TCP_PORT, /etc/services, built-in default
                      (3306), whatever comes first
--port-open-timeout=#  Maximum time in seconds to wait for the port to become
                      free. (Default: No wait).
--preload-buffer-size=# The size of the buffer that is allocated when preloading
                      indexes
--profiling-history-size=# Limit of query profiling memory
--query-alloc-block-size=# Allocation block size for query parsing and execution
--query-prealloc-size=# Persistent buffer for query parsing and execution
--range-alloc-block-size=# Allocation block size for storing ranges during
                      optimization
--range-optimizer-max-mem-size=# Maximum amount of memory used by the range optimizer to
                      allocate predicates during range analysis. The larger the
                      number, more memory may be consumed during range
                      analysis. If the value is too low to completed range
                      optimization of a query, index range scan will not be
                      considered for this query. A value of 0 means range
                      optimizer does not have any cap on memory.
--read-buffer-size=#   Each thread that does a sequential scan allocates a
                      buffer of this size for each table it scans. If you do
                      many sequential scans, you may want to increase this
                      value
--read-only            Make all non-temporary tables read-only, with the
                      exception for replication (slave) threads and users with
                      the SUPER privilege
--read-rnd-buffer-size=# When reading rows in sorted order after a sort, the rows
                      are read through this buffer to avoid a disk seeks
--regexp-stack-limit=# Stack size limit for regular expressions matches
--regexp-time-limit=#  Timeout for regular expressions matches, in steps of the
                      match engine, typically on the order of milliseconds.
--relay-log=name       The location and name to use for relay logs
--relay-log-index=name File that holds the names for relay log files.
--relay-log-info-file=name The location and name of the file that remembers where
                      the SQL replication thread is in the relay logs
--relay-log-info-repository=name Defines the type of the repository for the relay log
                      information and associated workers.
--relay-log-purge      if disabled - do not purge relay logs. if enabled - purge
                      them as soon as they are no more needed
                      (Defaults to on; use --skip-relay-log-purge to disable.)
--relay-log-recovery  Enables automatic relay log recovery right after the
                      database startup, which means that the IO Thread starts
                      re-fetching from the master right after the last
                      transaction processed
--relay-log-space-limit=#
```

Maximum space to use for all relay logs

--replicate-do-db=name
Tells the slave thread to restrict replication to the specified database. To specify more than one database, use the directive multiple times, once for each database. Note that this will only work if you do not use cross-database queries such as UPDATE some_db.some_table SET foo='bar' while having selected a different or no database. If you need cross database updates to work, make sure you have 3.23.28 or later, and use replicate-wild-do-table=db_name.%.

--replicate-do-table=name
Tells the slave thread to restrict replication to the specified table. To specify more than one table, use the directive multiple times, once for each table. This will work for cross-database updates, in contrast to replicate-do-db.

--replicate-ignore-db=name
Tells the slave thread to not replicate to the specified database. To specify more than one database to ignore, use the directive multiple times, once for each database. This option will not work if you use cross database updates. If you need cross database updates to work, make sure you have 3.23.28 or later, and use replicate-wild-ignore-table=db_name.%.

--replicate-ignore-table=name
Tells the slave thread to not replicate to the specified table. To specify more than one table to ignore, use the directive multiple times, once for each table. This will work for cross-database updates, in contrast to replicate-ignore-db.

--replicate-rewrite-db=name
Updates to a database with a different name than the original. Example:
replicate-rewrite-db=master_db_name->slave_db_name.

--replicate-same-server-id
In replication, if set to 1, do not skip events having our server id. Default value is 0 (to break infinite loops in circular replication). Can't be set to 1 if --log-slave-updates is used.

--replicate-wild-do-table=name
Tells the slave thread to restrict replication to the tables that match the specified wildcard pattern. To specify more than one table, use the directive multiple times, once for each table. This will work for cross-database updates. Example:
replicate-wild-do-table=foo%.bar% will replicate only updates to tables in all databases that start with foo and whose table names start with bar.

--replicate-wild-ignore-table=name
Tells the slave thread to not replicate to the tables that match the given wildcard pattern. To specify more than one table to ignore, use the directive multiple times, once for each table. This will work for cross-database updates. Example:
replicate-wild-ignore-table=foo%.bar% will not do updates to tables in databases that start with foo and whose table names start with bar.

--report-host=name
Hostname or IP of the slave to be reported to the master during slave registration. Will appear in the output of SHOW SLAVE HOSTS. Leave unset if you do not want the slave to register itself with the master. Note that it is not sufficient for the master to simply read the IP of

```
the slave off the socket once the slave connects. Due to
NAT and other routing issues, that IP may not be valid
for connecting to the slave from the master or other
hosts
--report-password=name
    The account password of the slave to be reported to the
    master during slave registration
--report-port=# Port for connecting to slave reported to the master
    during slave registration. Set it only if the slave is
    listening on a non-default port or if you have a special
    tunnel from the master or other clients to the slave. If
    not sure, leave this option unset
--report-user=name The account user name of the slave to be reported to the
    master during slave registration
--require-secure-transport
    When this option is enabled, connections attempted using
    insecure transport will be rejected. Secure transports
    are SSL/TLS, Unix socket or Shared Memory (on Windows).
--rpl-read-size=# The size for reads done from the binlog and relay log. It
    must be a multiple of 4kb. Making it larger might help
    with IO stalls while reading these files when they are
    not in the OS buffer cache
--rpl-stop-slave-timeout=#
    Timeout in seconds to wait for slave to stop before
    returning a warning.
--safe-user-create Don't allow new user creation by the user who has no
    write privileges to the mysql.user table.
--schema-definition-cache=#
    The number of cached schema definitions
--secure-file-priv=name
    Limit LOAD DATA, SELECT ... OUTFILE, and LOAD_FILE() to
    files within specified directory
--server-id=# Uniquely identifies the server instance in the community
    of replication partners
--server-id-bits=# Set number of significant bits in server-id
--session-track-gtids=name
    Controls the amount of global transaction ids to be
    included in the response packet sent by the
    server.(Default: OFF).
--session-track-schema
    Track changes to the 'default schema'.
    (Defaults to on; use --skip-session-track-schema to disable.)
--session-track-state-change
    Track changes to the 'session state'.
--session-track-system-variables=name
    Track changes in registered system variables.
--session-track-transaction-info=name
    Track changes to the transaction attributes. OFF to
    disable; STATE to track just transaction state (Is there
    an active transaction? Does it have any data? etc.);  

    CHARACTERISTICS to track transaction state and report all
    statements needed to start a transaction with the same
    characteristics (isolation level, read only/read write,
    snapshot - but not any work done / data modified within
    the transaction).
--sha256-password-auto-generate-rsa-keys
    Auto generate RSA keys at server startup if corresponding
    system variables are not specified and key files are not
    present at the default location.
    (Defaults to on; use --skip-sha256-password-auto-generate-rsa-keys to
    disable.)
--sha256-password-private-key-path=name
    A fully qualified path to the private RSA key used for
```

```
        authentication
--sha256-password-proxy-users
    If set to FALSE (the default), then the sha256_password
    authentication plugin will not signal for authenticated
    users to be checked for mapping to proxy users. When set
    to TRUE, the plugin will flag associated authenticated
    accounts to be mapped to proxy users when the server
    option check_proxy_users is enabled.
--sha256-password-public-key-path=name
    A fully qualified path to the public RSA key used for
    authentication
--shared-memory      Enable the shared memory
--shared-memory-base-name=name
    Base name of shared memory
--show-create-table-verbosity
    When this option is enabled, it increases the verbosity
    of 'SHOW CREATE TABLE'.
--show-old-temporals
    When this option is enabled, the pre-5.6.4 temporal types
    will be marked in the 'SHOW CREATE TABLE' and
    'INFORMATION_SCHEMA.COLUMNS' table as a comment in
    COLUMN_TYPE field. This variable is deprecated and will
    be removed in a future release.
--show-slave-auth-info
    Show user and password in SHOW SLAVE HOSTS on this
    master.
--skip-grant-tables Start without grant tables. This gives all users FULL
    ACCESS to all tables.
--skip-host-cache   Don't cache host names.
--skip-name-resolve Don't resolve hostnames. All hostnames are IP's or
    'localhost'.
--skip-networking   Don't allow connection with TCP/IP
--skip-new          Don't use new, possibly wrong routines.
--skip-show-database
    Don't allow 'SHOW DATABASE' commands
--skip-slave-start  If set, slave is not autostarted.
--skip-stack-trace  Don't print a stack trace on failure.
--slave-allow-batching
    Allow slave to batch requests
--slave-checkpoint-group=#"
    Maximum number of processed transactions by
    Multi-threaded slave before a checkpoint operation is
    called to update progress status.
--slave-checkpoint-period="#"
    Gather workers' activities to Update progress status of
    Multi-threaded slave and flush the relay log info to disk
    after every #th milli-seconds.
--slave-compressed-protocol
    Use compression on master/slave protocol
--slave-exec-mode=name
    Modes for how replication events should be executed.
    Legal values are STRICT (default) and IDEMPOTENT. In
    IDEMPOTENT mode, replication will not stop for operations
    that are idempotent. In STRICT mode, replication will
    stop on any unexpected difference between the master and
    the slave
--slave-load-tmpdir=name
    The location where the slave should put its temporary
    files when replicating a LOAD DATA INFILE command
--slave-max-allowed-packet=#
    The maximum packet length to sent successfully from the
    master to slave.
--slave-net-timeout=#
```

Number of seconds to wait for more data from a master/slave connection before aborting the read

--slave-parallel-type=name
Specifies if the slave will use database partitioning or information from master to parallelize transactions.(Default: DATABASE).

--slave-parallel-workers=#
Number of worker threads for executing events in parallel

--slave-pending-jobs-size-max=#
Max size of Slave Worker queues holding yet not applied events.The least possible value must be not less than the master side max_allowed_packet.

--slave-preserve-commit-order
Force slave workers to make commits in the same order as on the master. Disabled by default.

--slave-rows-search-algorithms=name
Set of searching algorithms that the slave will use while searching for records from the storage engine to either updated or deleted them. Possible values are: INDEX_SCAN, TABLE_SCAN and HASH_SCAN. Any combination is allowed, and the slave will always pick the most suitable algorithm for any given scenario. (Default: INDEX_SCAN, HASH_SCAN).

--slave-skip-errors=name
Tells the slave thread to continue replication when a query event returns an error from the provided list

--slave-sql-verify-checksum
Force checksum verification of replication events after reading them from relay log. Note: Events are always checksum-verified by slave on receiving them from the network before writing them to the relay log. Enabled by default.
(Defaults to on; use --skip-slave-sql-verify-checksum to disable.)

--slave-transaction-retries=#
Number of times the slave SQL thread will retry a transaction in case it failed with a deadlock or elapsed lock wait timeout, before giving up and stopping

--slave-type-conversions=name
Set of slave type conversions that are enabled. Legal values are: ALL_LOSSY to enable lossy conversions, ALL_NON LOSSY to enable non-lossy conversions, ALL_UNSIGNED to treat all integer column type data to be unsigned values, and ALL_SIGNED to treat all integer column type data to be signed values. Default treatment is ALL_SIGNED. If ALL_SIGNED and ALL_UNSIGNED both are specified, ALL_SIGNED will take higher priority than ALL_UNSIGNED. If the variable is assigned the empty set, no conversions are allowed and it is expected that the types match exactly.

--slow-launch-time=#
If creating the thread takes longer than this value (in seconds), the Slow_launch_threads counter will be incremented

--slow-query-log
Log slow queries to a table or log file. Defaults logging to a file hostname-slow.log or a table mysql.slow_log if --log-output=TABLE is used. Must be enabled to activate other slow log options

--slow-query-log-file=name
Log slow queries to given log file. Defaults logging to hostname-slow.log. Must be enabled to activate other slow log options

--slow-start-timeout=#
Maximum number of milliseconds that the service control manager should wait before trying to kill the windows

```

          service during startup(Default: 15000).
--socket=name      Socket file to use for connection
--sort-buffer-size=#   Each thread that needs to do a sort allocates a buffer of
                      this size
--sporadic-binlog-dump-fail   Option used by mysql-test for debugging and testing of
                                replication.
--sql-mode=name     Syntax: sql-mode=mode[,mode[,mode...]]. See the manual
                      for the complete list of valid sql modes
--sql-require-primary-key   When set, tables must be created with a primary key, and
                                an existing primary key cannot be removed with 'ALTER
                                TABLE'. Attempts to do so will result in an error.
--ssl               Enable SSL for connection (automatically enabled with
                      other flags).
                      (Defaults to on; use --skip-ssl to disable.)
--ssl-ca=name       CA file in PEM format (check OpenSSL docs, implies --ssl)
--ssl-capath=name   CA directory (check OpenSSL docs, implies --ssl)
--ssl-cert=name     X509 cert in PEM format (implies --ssl)
--ssl-cipher=name   SSL cipher to use (implies --ssl)
--ssl-crl=name     CRL file in PEM format (check OpenSSL docs, implies
                      --ssl)
--ssl-crlpath=name CRL directory (check OpenSSL docs, implies --ssl)
--ssl-fips-mode=name   SSL FIPS mode (applies only for OpenSSL); permitted
                      values are: OFF, ON, STRICT
--ssl-key=name      X509 key in PEM format (implies --ssl)
--standalone        Dummy option to start as a standalone program (NT).
--stored-program-cache=#   The soft upper limit for number of cached stored routines
                                for one connection.
--stored-program-definition-cache=#   The number of cached stored program definitions
--super-read-only   Make all non-temporary tables read-only, with the
                      exception for replication (slave) threads. Users with
                      the SUPER privilege are affected, unlike read_only.
                      Setting super_read_only to ON also sets read_only to ON.
-s, --symbolic-links   Enable symbolic link support (deprecated and will be
                      removed in a future release).
--sync-binlog=#     Synchronously flush binary log to disk after every #th
                      write to the file. Use 0 to disable synchronous flushing
--sync-master-info=#   Synchronously flush master info to disk after every #th
                      event. Use 0 to disable synchronous flushing
--sync-relay-log=#   Synchronously flush relay log to disk after every #th
                      event. Use 0 to disable synchronous flushing
--sync-relay-log-info=#   Synchronously flush relay log info to disk after every
                                #th transaction. Use 0 to disable synchronous flushing
--sysdate-is-now    Non-default option to alias SYSDATE() to NOW() to make it
                      safe-replicable. Since 5.0, SYSDATE() returns a `dynamic'
                      value different for different invocations, even within
                      the same statement.
--table-definition-cache=#   The number of cached table definitions
--table-open-cache=#   The number of cached open tables (total for all table
                      cache instances)
--table-open-cache-instances=#   The number of table cache instances
--tablespace-definition-cache=#

```

```

The number of cached tablespace definitions
--tc-heuristic-recover=name
    Decision to use in heuristic recover process. Possible
    values are OFF, COMMIT or ROLLBACK.
--temptable-max-ram=#
    Maximum amount of memory (in bytes) the TempTable storage
    engine is allowed to allocate from the main memory (RAM)
    before starting to store data on disk.
--thread-cache-size=#
    How many threads we should keep in a cache for reuse
--thread-handling=name
    Define threads usage for handling queries, one of
    one-thread-per-connection, no-threads, loaded-dynamically
--thread-stack=#
    The stack size for each thread
--tls-version=name
    TLS version, permitted values are TLSv1, TLSv1.1, TLSv1.2
--tmp-table-size=#
    If an internal in-memory temporary table in the MEMORY
    storage engine exceeds this size, MySQL will
    automatically convert it to an on-disk table
-t, --tmpdir=name
    Path for temporary files. Several paths may be specified,
    separated by a semicolon (;), in this case they are used
    in a round-robin fashion
--transaction-alloc-block-size=#
    Allocation block size for transactions to be stored in
    binary log
--transaction-isolation=name
    Default transaction isolation level.
--transaction-prealloc-size=#
    Persistent buffer for transactions to be stored in binary
    log
--transaction-read-only
    Default transaction access mode. True if transactions are
    read-only.
--transaction-write-set-extraction[=name]
    This option is used to let the server know when to
    extract the write set which will be used for various
    purposes.
--updatable-views-with-limit=name
    YES = Don't issue an error message (warning only) if a
    VIEW without presence of a key of the underlying table is
    used in queries with a LIMIT clause for updating. NO =
    Prohibit update of a VIEW, which does not contain a key
    of the underlying table and the query uses a LIMIT clause
    (usually get from GUI tools)
-u, --user=name
    Run mysqld daemon as user.
--validate-user-plugins
    Turns on additional validation of authentication plugins
    assigned to user accounts.
    (Defaults to on; use --skip-validate-user-plugins to disable.)
-v, --verbose
    Used with --help option for detailed help.
-V, --version
    Output version information and exit.
--wait-timeout=#
    The number of seconds the server waits for activity on a
    connection before closing it
--windowing-use-high-precision
    For SQL window functions, determines whether to enable
    inversion optimization for moving window frames also for
    floating values.
    (Defaults to on; use --skip-windowing-use-high-precision to disable.)
```

Variables (--variable-name=value)
and boolean options {FALSE|TRUE}

Value (after reading options)

abort-slave-event-count	0
activate-all-roles-on-login	FALSE

admin-address	(No default value)
admin-port	33062
allow-suspicious-udfs	FALSE
archive	ON
auto-generate-certs	TRUE
auto-increment-increment	1
auto-increment-offset	1
autocommit	TRUE
automatic-sp-privileges	TRUE
avoid-temporal-upgrade	FALSE
back-log	151
basedir	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\	
big-tables	FALSE
bind-address	*
binlog-cache-size	32768
binlog-checksum	CRC32
binlog-direct-non-transactional-updates	FALSE
binlog-encryption	FALSE
binlog-error-action	ABORT_SERVER
binlog-expire-logs-seconds	2592000
binlog-format	ROW
binlog-group-commit-sync-delay	0
binlog-group-commit-sync-no-delay-count	0
binlog-gtid-simple-recovery	TRUE
binlog-max-flush-queue-time	0
binlog-order-commits	TRUE
binlog-rotate-encryption-master-key-at-startup	FALSE
binlog-row-event-max-size	8192
binlog-row-image	FULL
binlog-row-metadata	MINIMAL
binlog-row-value-options	
binlog-rows-query-log-events	FALSE
binlog-stmt-cache-size	32768
binlog-transaction-dependency-history-size	25000
binlog-transaction-dependency-tracking	COMMIT_ORDER
blackhole	ON
block-encryption-mode	aes-128-ecb
bulk-insert-buffer-size	8388608
caching-sha2-password-auto-generate-rsa-keys	TRUE
caching-sha2-password-private-key-path	private_key.pem
caching-sha2-password-public-key-path	public_key.pem
character-set-client-handshake	TRUE
character-set-filesystem	binary
character-set-server	utf8mb4
character-sets-dir	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\share\charsets\	
check-proxy-users	FALSE
chroot	(No default value)
collation-server	utf8mb4_0900_ai_ci
completion-type	NO_CHAIN
concurrent-insert	AUTO
connect-timeout	10
console	FALSE
create-admin-listener-thread	FALSE
cte-max-recursion-depth	1000
datadir	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\data\	
default-authentication-plugin	caching_sha2_password
default-password-lifetime	0
default-storage-engine	InnoDB
default-time-zone	(No default value)
default-tmp-storage-engine	InnoDB

default-week-format	0
delay-key-write	ON
delayed-insert-limit	100
delayed-insert-timeout	300
delayed-queue-size	1000
disabled-storage-engines	
disconnect-on-expired-password	TRUE
disconnect-slave-event-count	0
div-precision-increment	4
end-markers-in-json	FALSE
enforce-gtid-consistency	FALSE
eq-range-index-dive-limit	200
event-scheduler	ON
expire-logs-days	0
explicit-defaults-for-timestamp	TRUE
external-locking	FALSE
federated	ON
flush	FALSE
flush-time	0
ft-boolean-syntax	+ -><()~*:""&
ft-max-word-len	84
ft-min-word-len	4
ft-query-expansion-limit	20
ft-stopword-file	(No default value)
gdb	FALSE
general-log	FALSE
general-log-file	c:\Users\121988\Downloads\mysql-8.0.14-winx64\data\TEE121988.log
group-concat-max-len	1024
group-replication-consistency	EVENTUAL
gtid-executed-compression-period	1000
gtid-mode	OFF
help	TRUE
histogram-generation-max-mem-size	20000000
host-cache-size	279
information-schema-stats-expiry	86400
init-connect	
init-file	(No default value)
init-slave	
initialize	FALSE
initialize-insecure	FALSE
innodb-adaptive-flushing	TRUE
innodb-adaptive-flushing-lwm	10
innodb-adaptive-hash-index	TRUE
innodb-adaptive-hash-index-parts	8
innodb-adaptive-max-sleep-delay	150000
innodb-api-bk-commit-interval	5
innodb-api-disable-rowlock	FALSE
innodb-api-enable-binlog	FALSE
innodb-api-enable-mdl	FALSE
innodb-api-trx-level	0
innodb-autoextend-increment	64
innodb-autoinc-lock-mode	2
innodb-buffer-pool-chunk-size	134217728
innodb-buffer-pool-dump-at-shutdown	TRUE
innodb-buffer-pool-dump-now	FALSE
innodb-buffer-pool-dump-pct	25
innodb-buffer-pool-filename	ib_buffer_pool
innodb-buffer-pool-in-core-file	TRUE
innodb-buffer-pool-instances	0
innodb-buffer-pool-load-abort	FALSE
innodb-buffer-pool-load-at-startup	TRUE
innodb-buffer-pool-load-now	FALSE

innodb-buffer-pool-size	134217728
innodb-change-buffer-max-size	25
innodb-change-buffering	all
innodb-checksum-algorithm	crc32
innodb-cmp-per-index-enabled	FALSE
innodb-commit-concurrency	0
innodb-compression-failure-threshold-pct	5
innodb-compression-level	6
innodb-compression-pad-pct-max	50
innodb-concurrency-tickets	5000
innodb-data-file-path	(No default value)
innodb-data-home-dir	(No default value)
innodb-deadlock-detect	TRUE
innodb-dedicated-server	FALSE
innodb-default-row-format	dynamic
innodb-directories	(No default value)
innodb-disable-sort-file-cache	FALSE
innodb-doublewrite	TRUE
innodb-fast-shutdown	1
innodb-file-per-table	TRUE
innodb-fill-factor	100
innodb-flush-log-at-timeout	1
innodb-flush-log-at-trx-commit	1
innodb-flush-method	unbuffered
innodb-flush-neighbors	0
innodb-flush-sync	TRUE
innodb-flushing-avg-loops	30
innodb-force-load-corrupted	FALSE
innodb-force-recovery	0
innodb-fsync-threshold	0
innodb-ft-aux-table	(No default value)
innodb-ft-cache-size	8000000
innodb-ft-enable-diag-print	FALSE
innodb-ft-enable-stopword	TRUE
innodb-ft-max-token-size	84
innodb-ft-min-token-size	3
innodb-ft-num-word-optimize	2000
innodb-ft-result-cache-limit	2000000000
innodb-ft-server-stopword-table	(No default value)
innodb-ft-sort-pll-degree	2
innodb-ft-total-cache-size	640000000
innodb-ft-user-stopword-table	(No default value)
innodb-io-capacity	200
innodb-io-capacity-max	4294967295
innodb-lock-wait-timeout	50
innodb-log-buffer-size	16777216
innodb-log-checksums	TRUE
innodb-log-compressed-pages	TRUE
innodb-log-file-size	50331648
innodb-log-files-in-group	2
innodb-log-group-home-dir	(No default value)
innodb-log-spin-cpu-abs-lwm	80
innodb-log-spin-cpu-pct-hwm	50
innodb-log-wait-for-flush-spin-hwm	400
innodb-log-write-ahead-size	8192
innodb-lru-scan-depth	1024
innodb-max-dirty-pages-pct	90
innodb-max-dirty-pages-pct-lwm	10
innodb-max-purge-lag	0
innodb-max-purge-lag-delay	0
innodb-max-undo-log-size	1073741824
innodb-monitor-disable	(No default value)
innodb-monitor-enable	(No default value)

innodb-monitor-reset	(No default value)
innodb-monitor-reset-all	(No default value)
innodb-old-blocks-pct	37
innodb-old-blocks-time	1000
innodb-online-alter-log-max-size	134217728
innodb-open-files	0
innodb-optimize-fulltext-only	FALSE
innodb-page-cleaners	4
innodb-page-size	16384
innodb-parallel-read-threads	4
innodb-print-all-deadlocks	FALSE
innodb-print-ddl-logs	FALSE
innodb-purge-batch-size	300
innodb-purge-rseg-truncate-frequency	128
innodb-purge-threads	4
innodb-random-read-ahead	FALSE
innodb-read-ahead-threshold	56
innodb-read-io-threads	4
innodb-read-only	FALSE
innodb-redo-log-encrypt	FALSE
innodb-replication-delay	0
innodb-rollback-on-timeout	FALSE
innodb-rollback-segments	128
innodb-sort-buffer-size	1048576
innodb-spin-wait-delay	6
innodb-stats-auto-recalc	TRUE
innodb-stats-include-delete-marked	FALSE
innodb-stats-method	nulls_equal
innodb-stats-on-metadata	FALSE
innodb-stats-persistent	TRUE
innodb-stats-persistent-sample-pages	20
innodb-stats-transient-sample-pages	8
innodb-status-file	FALSE
innodb-status-output	FALSE
innodb-status-output-locks	FALSE
innodb-strict-mode	TRUE
innodb-sync-array-size	1
innodb-sync-spin-loops	30
innodb-table-locks	TRUE
innodb-temp-data-file-path	(No default value)
innodb-temp-tablespaces-dir	(No default value)
innodb-thread-concurrency	0
innodb-thread-sleep-delay	10000
innodb-tmpdir	(No default value)
innodb-undo-directory	(No default value)
innodb-undo-log-encrypt	FALSE
innodb-undo-log-truncate	TRUE
innodb-undo-tablespaces	2
innodb-use-native-aio	TRUE
innodb-write-io-threads	4
interactive-timeout	28800
internal-tmp-disk-storage-engine	InnoDB
internal-tmp-mem-storage-engine	TempTable
join-buffer-size	262144
keep-files-on-create	FALSE
key-buffer-size	8388608
key-cache-age-threshold	300
key-cache-block-size	1024
key-cache-division-limit	100
keyring-migration-destination	(No default value)
keyring-migration-host	(No default value)
keyring-migration-port	0
keyring-migration-socket	(No default value)

keyring-migration-source	(No default value)
keyring-migration-user	(No default value)
language	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\share\lc-messages	en_US
lc-messages-dir	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\share\lc-time-names	en_US
local-infile	FALSE
lock-wait-timeout	31536000
log-bin	binlog
log-bin-index	binlog.index
log-bin-trust-function-creators	FALSE
log-bin-use-vl-row-events	FALSE
log-error	stderr
log-error-services	log_filter_internal;
log_sink_internal	
log-error-suppression-list	
log-error-verbosity	1
log-isam	myisam.log
log-output	FILE
log-queries-not-using-indexes	FALSE
log/raw	FALSE
log-short-format	FALSE
log-slave-updates	TRUE
log-slow-admin-statements	FALSE
log-slow-extra	FALSE
log-slow-slave-statements	FALSE
log-statements-unsafe-for-binlog	TRUE
log-tc	tc.log
log-tc-size	24576
log-throttle-queries-not-using-indexes	0
log-timestamps	UTC
long-query-time	10
low-priority-updates	FALSE
lower-case-table-names	1
mandatory-roles	
master-info-file	master.info
master-info-repository	TABLE
master-retry-count	86400
master-verify-checksum	FALSE
max-allowed-packet	67108864
max-binlog-cache-size	18446744073709547520
max-binlog-dump-events	0
max-binlog-size	1073741824
max-binlog-stmt-cache-size	18446744073709547520
max-connect-errors	100
max-connections	151
max-delayed-threads	20
max-digest-length	1024
max-error-count	1024
max-execution-time	0
max-heap-table-size	16777216
max-join-size	18446744073709551615
max-length-for-sort-data	4096
max-points-in-geometry	65536
max-prepared-stmt-count	16382
max-relay-log-size	0
max-seeks-for-key	4294967295
max-sort-length	1024
max-sp-recursion-depth	0
max-user-connections	0
max-write-lock-count	4294967295

memlock	FALSE
min-examined-row-limit	0
myisam-block-size	1024
myisam-data-pointer-size	6
myisam-max-sort-file-size	2146435072
myisam-mmap-size	18446744073709551615
myisam-recover-options	OFF
myisam-repair-threads	1
myisam-sort-buffer-size	8388608
myisam-stats-method	nulls_unequal
myisam-use-mmap	FALSE
mysql-native-password-proxy-users	FALSE
mysqlx	ON
mysqlx-bind-address	*
mysqlx-cache-cleaner	ON
mysqlx-connect-timeout	30
mysqlx-document-id-unique-prefix	0
mysqlx-idle-worker-thread-timeout	60
mysqlx-interactive-timeout	28800
mysqlx-max-allowed-packet	67108864
mysqlx-max-connections	100
mysqlx-min-worker-threads	2
mysqlx-port	33060
mysqlx-port-open-timeout	0
mysqlx-read-timeout	30
mysqlx-socket	(No default value)
mysqlx-ssl-ca	(No default value)
mysqlx-ssl-capath	(No default value)
mysqlx-ssl-cert	(No default value)
mysqlx-ssl-cipher	(No default value)
mysqlx-ssl-crl	(No default value)
mysqlx-ssl-crlpath	(No default value)
mysqlx-ssl-key	(No default value)
mysqlx-wait-timeout	28800
mysqlx-write-timeout	60
named-pipe	FALSE
named-pipe-full-access-group	*everyone*
net-buffer-length	16384
net-read-timeout	30
net-retry-count	10
net-write-timeout	60
new	FALSE
ngram	ON
ngram-token-size	2
no-dd-upgrade	FALSE
no-monitor	FALSE
offline-mode	FALSE
old	FALSE
old-alter-table	FALSE
old-style-user-limits	FALSE
open-files-limit	10209
optimizer-prune-level	1
optimizer-search-depth	62
optimizer-switch	
index_merge=on, index_merge_union=on, index_merge_sort_union=on, index_merge_intersection=on, engine_condition_pushdown=on, index_condition_pushdown=on, mrr=on, mrr_cost_based=on, block_nested_loop=on, batched_key_access=off, materialization=on, semijoin=on, loosescan=on, firstmatch=on, duplicateweedout=on, subquery_materialization_cost_based=on, use_index_extensions=on, condition_fanout_filter=on, derived_merge=on, use_invisible_indexes=off, skip_scan=on	
optimizer-trace	
optimizer-trace-features	
greedy_search=on, range_optimizer=on, dynamic_range=on, repeated_subselect=on	1
optimizer-trace-limit	

optimizer-trace-max-mem-size	1048576
optimizer-trace-offset	-1
parser-max-mem-size	18446744073709551615
password-history	0
password-require-current	FALSE
password-reuse-interval	0
performance-schema	TRUE
performance-schema-accounts-size	-1
performance-schema-consumer-events-stages-current	FALSE
performance-schema-consumer-events-stages-history	FALSE
performance-schema-consumer-events-stages-history-long	FALSE
performance-schema-consumer-events-statements-current	TRUE
performance-schema-consumer-events-statements-history	TRUE
performance-schema-consumer-events-statements-history-long	FALSE
performance-schema-consumer-events-transactions-current	TRUE
performance-schema-consumer-events-transactions-history	TRUE
performance-schema-consumer-events-transactions-history-long	FALSE
performance-schema-consumer-events-waits-current	FALSE
performance-schema-consumer-events-waits-history	FALSE
performance-schema-consumer-events-waits-history-long	FALSE
performance-schema-consumer-global-instrumentation	TRUE
performance-schema-consumer-statements-digest	TRUE
performance-schema-consumer-thread-instrumentation	TRUE
performance-schema-digests-size	-1
performance-schema-error-size	4367
performance-schema-events-stages-history-long-size	-1
performance-schema-events-stages-history-size	-1
performance-schema-events-statements-history-long-size	-1
performance-schema-events-statements-history-size	-1
performance-schema-events-transactions-history-long-size	-1
performance-schema-events-transactions-history-size	-1
performance-schema-events-waits-history-long-size	-1
performance-schema-events-waits-history-size	-1
performance-schema-hosts-size	-1
performance-schema-instrument	
performance-schema-max-cond-classes	100
performance-schema-max-cond-instances	-1
performance-schema-max-digest-length	1024
performance-schema-max-digest-sample-age	60
performance-schema-max-file-classes	80
performance-schema-max-file-handles	32768
performance-schema-max-file-instances	-1
performance-schema-max-index-stat	-1
performance-schema-max-memory-classes	450
performance-schema-max-metadata-locks	-1
performance-schema-max-mutex-classes	300
performance-schema-max-mutex-instances	-1
performance-schema-max-prepared-statements-instances	-1
performance-schema-max-program-instances	-1
performance-schema-max-rwlock-classes	60
performance-schema-max-rwlock-instances	-1
performance-schema-max-socket-classes	10
performance-schema-max-socket-instances	-1
performance-schema-max-sql-text-length	1024
performance-schema-max-stage-classes	175
performance-schema-max-statement-classes	218
performance-schema-max-statement-stack	10
performance-schema-max-table-handles	-1
performance-schema-max-table-instances	-1
performance-schema-max-table-lock-stat	-1
performance-schema-max-thread-classes	100
performance-schema-max-thread-instances	-1
performance-schema-session-connect-attrs-size	-1

performance-schema-setup-actors-size	-1
performance-schema-setup-objects-size	-1
performance-schema-users-size	-1
persist-only-admin-x509-subject	
persisted-globals-load	TRUE
pid-file	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\data\TEE121988.pid	
plugin-dir	
c:\Users\121988\Downloads\mysql-8.0.14-winx64\lib\plugin\	
port	3306
port-open-timeout	0
preload-buffer-size	32768
profiling-history-size	15
query-alloc-block-size	8192
query-prealloc-size	8192
range-alloc-block-size	4096
range-optimizer-max-mem-size	8388608
read-buffer-size	131072
read-only	FALSE
read-rnd-buffer-size	262144
regexp-stack-limit	8000000
regexp-time-limit	32
relay-log	TEE121988-relay-bin
relay-log-index	TEE121988-relay-bin.index
relay-log-info-file	relay-log.info
relay-log-info-repository	TABLE
relay-log-purge	TRUE
relay-log-recovery	FALSE
relay-log-space-limit	0
replicate-same-server-id	FALSE
report-host	(No default value)
report-password	(No default value)
report-port	0
report-user	(No default value)
require-secure-transport	FALSE
rpl-read-size	8192
rpl-stop-slave-timeout	31536000
safe-user-create	FALSE
schema-definition-cache	256
secure-file-priv	NULL
server-id	1
server-id-bits	32
session-track-gtids	OFF
session-track-schema	TRUE
session-track-state-change	FALSE
session-track-system-variables	
time_zone,autocommit,character_set_client,character_set_results,character_set_connection	
session-track-transaction-info	OFF
sha256-password-auto-generate-rsa-keys	TRUE
sha256-password-private-key-path	private_key.pem
sha256-password-proxy-users	FALSE
sha256-password-public-key-path	public_key.pem
shared-memory	FALSE
shared-memory-base-name	MYSQL
show-create-table-verbosity	FALSE
show-old-temporals	FALSE
show-slave-auth-info	FALSE
skip-grant-tables	FALSE
skip-name-resolve	FALSE
skip-networking	FALSE
skip-show-database	FALSE
skip-slave-start	FALSE
slave-allow-batching	FALSE

slave-checkpoint-group	512
slave-checkpoint-period	300
slave-compressed-protocol	FALSE
slave-exec-mode	STRICT
slave-load-tmpdir	C:\Users\121988\AppData\Local\Temp
slave-max-allowed-packet	1073741824
slave-net-timeout	60
slave-parallel-type	DATABASE
slave-parallel-workers	0
slave-pending-jobs-size-max	134217728
slave-preserve-commit-order	FALSE
slave-rows-search-algorithms	INDEX_SCAN,HASH_SCAN
slave-skip-errors	(No default value)
slave-sql-verify-checksum	TRUE
slave-transaction-retries	10
slave-type-conversions	
slow-launch-time	2
slow-query-log	FALSE
slow-query-log-file	c:\Users\121988\Downloads\mysql-8.0.14-winx64\data\TEE121988-slow.log
slow-start-timeout	15000
socket	MySQL
sort-buffer-size	262144
sporadic-binlog-dump-fail	FALSE
sql-mode	
ONLY_FULL_GROUP_BY,STRICT_TRANS_TABLES,NO_ZERO_IN_DATE,NO_ZERO_DATE,ERROR_FOR_DIVISION_BY_ZERO,NO_ENGINE_SUBSTITUTION	
sql-require-primary-key	FALSE
ssl	TRUE
ssl-ca	(No default value)
ssl-capath	(No default value)
ssl-cert	(No default value)
ssl-cipher	(No default value)
ssl-crl	(No default value)
ssl-crlpath	(No default value)
ssl-fips-mode	OFF
ssl-key	(No default value)
stored-program-cache	256
stored-program-definition-cache	256
super-read-only	FALSE
symbolic-links	FALSE
sync-binlog	1
sync-master-info	10000
sync-relay-log	10000
sync-relay-log-info	10000
sysdate-is-now	FALSE
table-definition-cache	2000
table-open-cache	4000
table-open-cache-instances	16
tablespace-definition-cache	256
tc-heuristic-recover	OFF
temptable-max-ram	1073741824
thread-cache-size	9
thread-handling	one-thread-per-connection
thread-stack	286720
tls-version	TLSv1,TLSv1.1,TLSv1.2
tmp-table-size	16777216
tmpdir	C:\Users\121988\AppData\Local\Temp
transaction-alloc-block-size	8192
transaction-isolation	REPEATABLE-READ
transaction-prealloc-size	4096
transaction-read-only	FALSE
transaction-write-set-extraction	XXHASH64

updatable-views-with-limit	YES
validate-user-plugins	TRUE
verbose	TRUE
wait-timeout	28800
windowing-use-high-precision	TRUE

To see what values a running MySQL server is using, type
'mysqladmin variables' instead of 'mysqld --verbose --help'.

From:

<http://www.hdip-data-analytics.com/> - HDip Data Analytics

Permanent link:

<http://www.hdip-data-analytics.com/help/mysql/mysqld>

Last update: **2020/06/20 14:39**