

## mysqldxchk Documentation

mysqldxchk (MySQL Index Checker) checks MySQL databases/tables for unused indexes. To accomplish this task, mysqldxchk requires one or more log files which can be any number and combination of [slow](#), [general](#), and "raw" log files. mysqldxchk then compares the indexes used by the queries from the log files to all existing indexes in the database schema and reports which indexes in the database schema are not used.

The completeness of mysqldxchk's unused index report is related to the representativeness of the queries from the log files for the entire database schema. For example, if no queries use database x, then mysqldxchk *may* report that all indexes for all tables in database x are unused. These details are discussed further in the [Guide To Using mysqldxchk](#).

### A Note About Options

Technically, command line options are in the form `--option`, but `-option` works too. All options can be abbreviated if the abbreviation is unique. For example, option `--host` can be abbreviated `--ho` but not `--h` because `--h` is ambiguous: it could mean `--host` or `--help`.

### Options

<code>--user USER</code> <code>--password</code> <code>--host ADDRESS</code> <code>--port PORT</code> <code>--socket SOCKET</code> <code>--no-mycnf</code> <code>--help (-?)</code>	These options mimic most standard applications. <code>--password</code> can take the password on the command line like <code>"--password FOO"</code> . Using <code>--password</code> alone without giving a password on the command line causes mysqldxchk to prompt for a password. <code>--no-mycnf</code> makes mysqldxchk not read <code>~/my.cnf</code> which it does by default otherwise. <code>--user</code> and <code>--password</code> always override values from <code>~/my.cnf</code> .
<code>--debug</code>	Using <code>--debug</code> causes mysqldxchk to print a <i>lot</i> of extra debugging information about what it is doing internally. In release versions of mysqldxchk, debugging information does not include log parsing information. For more information about release versus debug versions of mysqldxchk, read <a href="#">mysqldxchk Bugs &amp; Problems</a> .
<code>--general (-g) LOG</code> <code>--slow (-s) LOG</code> <code>--raw (-r) LOG</code>	mysqldxchk can read and combine multiple MySQL slow and general logs, as well as "raw" logs. Raw logs are simply files with <code>;\n</code> (semi-colon new-line) terminated SQL statements. At least one log file must be given. Multiple log files can be specified like: <code>"-g log1,log2,log3 -s slow_queries.log"</code> .  By default, only SELECT, UPDATE and USE db statements are read; all other statements are discarded.  General logs and raw logs are more reliable than slow logs because general logs always contain an equivalent "USE db" statement, and a "USE db" statement can and should be added to raw logs. Slow logs, however, may or may not contain a "USE db." If not, mysqldxchk tries to discover which database a query belongs to. The database discovery process is discussed in the <a href="#">Guide To Using mysqldxchk</a> .
<code>--grep P</code>	grep and save only those statements which match Perl regex pattern P.
<code>--databases (-D) D</code>	By default mysqldxchk uses and reports on every index in every table in every database. However, there are times when you may wish to only consider a limited number of databases. The <code>--databases</code> option causes mysqldxchk to only report on the given databases D (where D is a comma-separated list of database names without any spaces). Queries which use databases not given in D are silently discarded.
<code>--show-used (-su)</code>	By default mysqldxchk only reports unused indexes. <code>--show-used</code> causes mysqldxchk to show all indexes both unused and used. Used indexes have a number in parenthesis beside them, like (7), where '7' is the number of unique queries that use that index.
<code>--ignore-update (-iu)</code>	mysqldxchk transforms UPDATE statements into SELECT statements by dropping the SET clause and adding "SELECT * FROM". For example, "UPDATE table SET col = 1 WHERE something IS NULL" becomes "SELECT * FROM table WHERE something IS NULL". <code>--ignore-update</code> causes mysqldxchk to silently discard all UPDATE statements.
<code>--no-db-discovery (-ndd)</code>	If a query does not have a database, mysqldxchk attempts to discover the query's appropriate database. <code>--no-db-discovery</code> prevents mysqldxchk from this attempt. Queries without databases will not be EXPLAINED.  The database discovery process is discussed in the <a href="#">Guide To Using mysqldxchk</a> .
<code>--no-discovery-report (-ndr)</code>	While attempting to discover databases for queries without databases, mysqldxchk prints extra information about its successes and failures. <code>--no-discovery-report</code> suppress this extra information. Database discovery will still be attempted with this option.

### What To Do About Bug and Errors

If mysqldxchk will not work for you, please read the brief guide to submitting bug and problem reports : [mysqldxchk Bugs & Problems](#)

(Doc rev: Mar 29 2007)

Last updated: 2009-09-30