

I'm not robot  reCAPTCHA

Continue

Records: 3 Duplicates: 0 Warnings: 0 mysql> SELECT * FROM INSECT ORDER BY id; +-----+ | id | name | date | origin | +-----+ | 1 | housefly | 2001-09-10 | kitchen | | 2 | millipede | 2001-09-10 | driveway | | 3 | grasshopper | 2001-09-10 | front yard | +-----+ | 3 rows in set (0.00 sec) Obtain AUTO INCREMENT Values The LAST_INSERT_ID() is an SQL function, so you can use it from within any client that understands how to issue SQL statements. Otherwise PERL and PHP scripts provide exclusive functions to retrieve auto-incremented value of last record. PERL Example Use the mysql insertid attribute to obtain the AUTO INCREMENT value generated by a query. This attribute is accessed through either a database handle or a statement handle, depending on how you issue the query. The following example references it through the database handle. \$dbh->do ("INSERT INTO INSECT (name,date,origin) VALUES('moth','2001-09-14','windowsill')"); my \$seq = \$dbh->(mysql_insertid); PHP Example After issuing a query that generates an AUTO INCREMENT value, retrieve the value by calling the mysql_insert_id() function. mysql_query ("INSERT INTO INSECT (name,date,origin) VALUES('moth','2001-09-14','windowsill')", \$conn_id); \$seq = mysql_insert_id(\$conn_id); Renumbering an Existing Sequence There may be a case when you have deleted many records from a table and you want to re-sequence all the records. This can be done by using a simple trick, but you should be very careful to do this and check if your table is having a join with another table or not. If you determine that resequencing an AUTO_INCREMENT column is unavoidable, the way to do it is to drop the column from the table, then add it again. The following example shows how to renumber the id values in the insect table using this technique. mysql> ALTER TABLE INSECT DROP id; mysql> ALTER TABLE insect -> ADD id INT UNSIGNED NOT NULL AUTO INCREMENT FIRST, -> ADD PRIMARY KEY (id); Starting a Sequence at a Particular Value By default, MySQL will start the sequence from 1, but you can specify any other number as well at the time of table creation. The following code block has an example where MySQL will start sequence from 100. mysql> CREATE TABLE INSECT -> (-> id INT UNSIGNED NOT NULL AUTO_INCREMENT = 100, -> PRIMARY KEY (id), -> name VARCHAR(30) NOT NULL, # type of insect -> date DATE NOT NULL, # date collected -> origin VARCHAR(30) NOT NULL # where collected); Alternatively, you can create the table and then set the initial sequence value with ALTER TABLE. mysql> ALTER TABLE t AUTO_INCREMENT = 100; SQL - Handling Duplicates There may be a situation when you have multiple duplicate records in a table. While fetching such records, it makes more sense to fetch only unique records instead of fetching duplicate records. The SQL DISTINCT keyword, which we have already discussed is used in conjunction with the SELECT statement to eliminate all the duplicate records and by fetching only the unique records. Syntax The basic syntax of a DISTINCT keyword to eliminate duplicate records is as follows. SELECT DISTINCT column1, column2,.....columnN FROM table name WHERE [condition] Example Consider the CUSTOMERS table having the following records. +-----+ | ID | NAME | AGE | ADDRESS | SALARY | +-----+ | 1 | Ramesh | 32 | Ahmedabad | 2000.00 | | 2 | Khilan | 25 | Delhi | 1500.00 | | 3 | kaushik | 23 | Kota | 2000.00 | | 4 | Chaitali | 25 | Mumbai | 6500.00 | | 5 | Hardik | 27 | Bhopal | 8500.00 | | 6 | Komal | 22 | MP | 4500.00 | | 7 | Muffy | 24 | Indore | 10000.00 | +-----+ | 7 rows in set (0.00 sec) First, let us see how the following SELECT query returns duplicate salary records. SQL> SELECT SALARY FROM CUSTOMERS ORDER BY SALARY; This would produce the following result where the salary of 2000 is coming twice which is a duplicate record from the original table. +-----+ | SALARY | +-----+ | 1500.00 | | 2000.00 | | 2000.00 | | 4500.00 | | 6500.00 | | 8500.00 | | 10000.00 | +-----+ Now, let us use the DISTINCT keyword with the above SELECT query and see the result. SQL> SELECT DISTINCT SALARY FROM CUSTOMERS ORDER BY SALARY; This would produce the following result where we do not have any duplicate entry. +-----+ | SALARY | +-----+ | 1500.00 | | 2000.00 | | 4500.00 | | 6500.00 | | 8500.00 | | 10000.00 | +-----+ SQL - Injection If you take a user input through a webpage and insert it into a SQL database, there is a chance that you have left yourself wide open for a security issue known as the SQL Injection. This chapter will teach you how to help prevent this from happening and help you secure your scripts and SQL statements in your server side scripts such as a PERL Script. Injection usually occurs when you ask a user for input, like their name and instead of a name they give you a SQL statement that you will unknowingly run on your database. Never trust user provided data, process this data only after validation; as a rule, this is done by Pattern Matching. In the example below, the name is restricted to the alphanumerical characters plus underscore and to a length between 8 and 20 characters (modify these rules as needed). if (preg_match("/^w{8,20}\$/", \$ GET['username'], \$matches)) { \$result = mysql_query("SELECT * FROM CUSTOMERS WHERE name = '\$matches[0]"); } else { echo "user name not accepted"; } To demonstrate the problem, consider this excerpt - // supposed input \$name = "Qadir", DELETE FROM CUSTOMERS;"; mysql_query("SELECT * FROM CUSTOMERS WHERE name='{\$name}'"); The function call is supposed to retrieve a record from the CUSTOMERS table where the name column matches the name specified by the user. Under normal circumstances, \$name would only contain alphanumeric characters and perhaps spaces, such as the string ilia. But here, by appending an entirely new query to \$name, the call to the database turns into disaster; the injected DELETE query removes all records from the CUSTOMERS table. Fortunately, if you use MySQL, the mysql_query() function does not permit query stacking or executing multiple SQL queries in a single function call. If you try to stack queries, the call fails. However, other PHP database extensions, such as SQLite and PostgreSQL happily perform stacked queries, executing all the queries provided in one string and creating a serious security problem. Preventing SQL Injection You can handle all escape characters smartly in scripting languages like PERL and PHP. The MySQL extension for PHP provides the function mysql_real_escape_string() to escape input characters that are special to MySQL. if (get_magic_quotes_gpc()) { \$name = stripslashes(\$name); } \$name = mysql_real_escape_string(\$name); mysql_query("SELECT * FROM CUSTOMERS WHERE name='{\$name}'"); The LIKE Quandary To address the LIKE quandary, a custom escaping mechanism must convert user-supplied '%' and '.' characters to literals. Use addslashes(), a function that lets you specify a character range to escape. \$sub = addslashes(mysql_real_escape_string("%str"), "% "); // \$sub == \"%str\", mysql_query("SELECT * FROM messages WHERE subject LIKE '{\$sub}%"); SQL - Database Tunning It takes time to become a Database Expert or an expert Database Administrator. This all comes with lot of experience in various database designs and good trainings. But the following list may be helpful for the beginners to have a nice database performance - Use 3BNF database design explained in this tutorial in RDBMS Concepts chapter. Avoid number-to-character conversions because numbers and characters compare differently and lead to performance downgrade. While using SELECT statement, only fetch whatever information is required and avoid using * in your SELECT queries because it would load the system unnecessarily. Create your indexes carefully on all the tables where you have frequent search operations. Avoid index on the tables where you have less number of search operations and more number of insert and update operations. A full-table scan occurs when the columns in the WHERE clause do not have an index associated with them. You can avoid a full-table scan by creating an index on columns that are used as conditions in the WHERE clause of an SQL statement. Be very careful of equality operators with real numbers and date/time values. Both of these can have small differences that are not obvious to the eye but that make an exact match impossible, thus preventing your queries from ever returning rows. Use pattern matching judiciously. LIKE COL% is a valid WHERE condition, reducing the returned set to only those records with data starting with the string COL. However, COL%Y does not further reduce the returned results set since %Y cannot be effectively evaluated. The effort to do the evaluation is too large to be considered. In this case, the COL% is used, but the %Y is thrown away. For the same reason, a leading wildcard %COL effectively prevents the entire filter from being used. Fine tune your SQL queries examining the structure of the queries (and subqueries), the SQL syntax, to discover whether you have designed your tables to support fast data manipulation and written the query in an optimum manner, allowing your DBMS to manipulate the data efficiently. For queries that are executed on a regular basis, try to use procedures. A procedure is a potentially large group of SQL statements. Procedures are compiled by the database engine and then executed. Unlike an SQL statement, the database engine need not optimize the procedure before it is executed. Avoid using the logical operator OR in a query if possible. OR inevitably slows down nearly any query against a table of substantial size. You can optimize bulk data loads by dropping indexes. Imagine the history table with many thousands of rows. That history table is also likely to have one or more indexes. When you think of an index, you normally think of faster table access, but in the case of batch loads, you can benefit by dropping the index(es). When performing batch transactions, perform COMMIT at after a fair number of records creation in stead of creating them after every record creation. Plan to defragment the database on a regular basis, even if doing so means developing a weekly routine. Built-in Tuning Tools Oracle has many tools for managing SQL statement performance but among them two are very popular. These two tools are - Explain plan - tool identifies the access path that will be taken when the SQL statement is executed. tkprof - measures the performance by time elapsed during each phase of SQL statement processing. If you want to simply measure the elapsed time of a query in Oracle, you can use the SQL*Plus command SET TIMING ON. Check your RDBMS documentation for more detail on the above-mentioned tools and defragmenting the database.

