Migrating a MySQL Database to SQL Anywhere 10

A whitepaper from Sybase iAnywhere

CONTENTS

INTRODUCTION

Migrating data from MySQL to SQL Anywhere can be a straightforward process if there are not a lot of MySQL extensions in use within your database and application. SQL Anywhere simplifies migration by including built-in tools that facilitate a smooth transition from MySQL (and other RDBMS's) to SQL Anywhere.

The first part of this document discusses in detail differences between SQL Anywhere and MySQL, including data type differences, feature differences, and syntax differences. Some of the features that are unique to MySQL can hinder migration. Approaches to how you might choose to deal with these issues are provided. The second part of this document includes a systematic explanation of how to migrate data from a MySQL database into a SQL Anywhere database using the Sybase Central Data Migration wizard. Finally, the third part of this document supplies an example of how you might migrate an existing application running against MySQL to one that runs against SQL Anywhere.

This document was written for SQL Anywhere Studio version 10 and later, and MySQL version 4.1 and later. For SQL Anywhere version 9, see http://www.sybase.com/detail?id=1030300.

DIFFERENCES BETWEEN MYSQL 5 AND SQL ANYWHERE 10

The following sections describe some of the differences between MySQL and SQL Anywhere that you may encounter during migration, along with some suggested solutions that can be used as starting points to resolve any issues that arise during migration. There are many ways to optimize your code with SQL Anywhere features that are missing from MySQL.

It is highly recommended that you review the SQL Anywhere documentation as well as the developer resources, including samples and technical documents, available on the iAnywhere Solutions website at http://www.sybase.com/developer/library/sql-anywhere-techcorner when moving to SQL Anywhere.

DATA TYPES

In most cases, the MySQL data types can map directly to SQL Anywhere data types. The following table lists some examples:

MySQL data type	Equivalent SQL Anywhere data type	Notes
SMALLINT [UNSIGNED]	[UNSIGNED] SMALLINT	
TINYINT UNSIGNED	TINYINT	
BOOL/BOOLEAN	TINYINT or BIT	
DOUBLE(m,n)	DOUBLE(n)	The 'm' argument in MySQL represents the total number of digits.
FLOAT(m,n) UNSIGNED	FLOAT(n)	The 'm' argument in MySQL represents the total number of digits.
DECIMAL(m,d), FIXED(m,d)	DECIMAL(m,d)	
VARCHAR(32) BINARY	BINARY(32)	
CHAR/VARCHAR	CHAR/VARCHAR	
TEXT	VARCHAR	
TINYBLOB	BINARY(255)	
BLOB/MEDIUMBLOB/LONGBLOB	LONG BINARY	SQL Anywhere stores only what is required to hold the BLOB value, so

	the extra requirement of specifying BLOB size via different BLOB types is not required

Note: In addition to the differences in data types themselves, there is also a difference in the declaration of data types. MySQL provides an optional parameter for its numeric types that allows you to specify the maximum display width for integer types. For example, an INT(4) column would return the value '1' as '<s><s><s>1', where <s> is a space. The optional ZEROFILL modifier on the type definition would replace the spaces in the previous example with zeros and add the UNSIGNED attribute to the column. For example, '1' is returned as '0001'. The merge of display format and data values in the type definition is not supported by SQL Anywhere. The CAST and CONVERT functions, along with the various string manipulation functions are available to format data values when they are retrieved from the database.

The following data types differ from SQL Anywhere more substantially than by syntax:

MEDIUMINT: These are 3-byte integer values. They can easily be simulated using an INTEGER (4 bytes) or SMALLINT (2 bytes) in SQL Anywhere, depending on the expected range of values for the column.

YEAR: Year is a 2 or 4 digit value. The SQL Anywhere DATE data type can be used to hold year values, but uses slightly more storage space. Date arithmetic and conversion can be performed using the SQL Anywhere built-in functions listed under "Date and Time Functions" in the "SQL Functions" chapter of the "SQL Anywhere Server - SQL Reference" manual.

The following data types do not match exactly, and will require some work to migrate to SQL Anywhere:

NCHAR/NVARCHAR: As of MySQL 4.1, an NCHAR value is stored in MySQL using the UTF8 character set. SQL Anywhere supports a variety of character sets, including UTF8. With a database created using the proper collation, the use of a special data type to store international values is not required, though SQL Anywhere does support the NCHAR datatype. To learn more about the latest international character set support in SQL Anywhere, see the chapter "International Language and Character Sets" in the "SQL Anywhere Server - Database Administration" manual.

ENUM: An ENUM value is a string object whose value must be chosen from a list of supplied values enumerated in the column definition when a table is created. The enumerated values can also be inserted or retrieved by their index position in the ENUM definition. The index value 0 is reserved for the empty string. The ENUM data type is represented in SQL Anywhere by a TINYINT column. There are a few options to accomplish the same behavior as the MySQL ENUM, but changes to the client application will almost certainly be required. Some options are:

- Altering the client side application to remove the need for ENUM values
- Translating the ENUM values on the client side
- Adding some logic to the server side to attempt to mimic the MySQL behavior of ENUM values by using stored procedures, triggers, computed columns, view, and/or a mapping table for the ENUM types

For example, a view could be created on the table containing the ENUM fields to allow for the return of the values as a string, while a regular SELECT could be used to return them as a number. Here is an example of a view that could be used:

```
CREATE TABLE enumtbl (pkey INTEGER NOT NULL PRIMARY KEY, enumval
TINYINT);

CREATE VIEW v_enumtable AS
  SELECT pkey
```

```
                CASE WHEN 0 then ''
                     WHEN 1 then 'val1'
                     WHEN 2 then 'val2'
                     WHEN 3 then 'val3'
                ELSE NULL
          END
      FROM enumtbl;
```

Then, a query may look something like this:

```
SELECT pkey, enumval FROM v_enumtable;
```

Alternatively, a mapping table could be created for the ENUM vales and whenever you retrieve data from `enumtbl`, a join can be made to the mapping table containing the ENUM strings.

```
CREATE TABLE enummap( enumval TINAY INT NOT NULL PRIMARY KEY, enumstr
CHAR(16));
```

Then a query may look something like this:

```
SELECT pkey, enumval FROM enumtbl, enummap
WHERE enumtbl.enumval = enummap.enumval;
```

An insert on the table can be done directly if you are using the index values of the ENUM; otherwise, a stored procedure could be used to insert a row into any table containing an ENUM. The stored procedure would contain the logic to decode the ENUM values. Following is a sample stored procedure implementation to deal with an ENUM column equivalent in SQL Anywhere (using the same table definition as above):

```
CREATE PROCEDURE sp_insert_enumval (IN pkeyval int, IN enum CHAR(16))
BEGIN
  DECLARE enum_map TINYINT;
  IF enum IS NOT NULL THEN
     CASE enum
        WHEN '' THEN SET enum_map = 0
        WHEN 'val1' THEN SET enum_map = 1
        WHEN 'val2' THEN SET enum_map = 2
        WHEN 'val3' THEN SET enum_map = 3
     ELSE SET enum_map = 0
     END CASE
  END IF;


  INSERT INTO enumtbl VALUES( pkeyval, enum_map);
END
```

SET: A SET value is a string object whose value must be chosen from a list of values supplied when the column is defined. It is different from the ENUM type in that 0 or more values from the list can be combined to create a valid value for the column. Each value in the set is assigned a binary value and data can be assigned or retrieved by using a number representing the combination of values to be set. For example, specifying a value of 9 would insert the first and fourth value from the set into the column. Depending on how many values are in the set (64 is maximum), anything from a TINYINT to a BIGINT is required to map a SET value from MySQL to SQL Anywhere. To achieve the same behavior as MySQL, methods similar to those demonstrated above with the ENUM data type can be used.

MYSQL FUNCTION MAPPINGS TO SQL ANYWHERE

Many of the functions in both MySQL and SQL Anywhere have the same name. Most MySQL functions that have different names have an equivalent SQL Anywhere version. MySQL contains a few built-in functions that do not exist in SQL Anywhere. Most of these functions can be created in SQL

Anywhere as user-defined functions that perform the same activity. If you give these functions the same name in the SQL Anywhere database, you will not need to modify the existing client application's SQL statements. Here are some examples of how SQL Anywhere user-defined functions can supply the same functionality as their MySQL built-in counterparts:

```
CREATE FUNCTION FROM_UNIXTIME (IN fromdt bigint default 0, IN fmt
varchar(32) default 'Mmm dd, yyyy hh:mm:ss' ) RETURNS datetime
BEGIN
  RETURN(dateformat(dateadd(second, fromdt, '1970/01/01 00:00:00',
fmt))
END;

CREATE FUNCTION SEC_TO_TIME (IN sec bigint default 0) RETURNS time
BEGIN
  RETRUN (dateadd(second, sec, '1970/01/01 00:00:00') )
END;
```

The following sections detail many of the MySQL functions along with their SQL Anywhere equivalents. The list is extensive, but not exhaustive, as the list of function in both SQL Anywhere and MySQL changes with each release.

STRING FUNCTIONS

MySQL function	Equivalent SQL Anywhere function	Notes
IFNULL(a,b)	IFNULL(a,b,a) or ISNULL(a,b)	The SQL Anywhere IFNULL function behaves slightly different from the MySQL version.
IF(cond, a, b)	IF cond THEN a ELSE b	Both MySQL and SQL Anywhere support the IF statement in procedural and embedded logic.
CONCAT(a, b, …)	STRING (a, b, …)	In MySQL, if any of a, b … is NULL the return value is NULL, SQL Anywhere returns NULL only if all arguments are NULL. Otherwise NULLS are treated as empty strings for the purpose of concatenation.
CONCAT_WS(sep, str1, str2 …)	STRING(str1, sep, str2, sep …)	See CONCAT note.
CONV(N, frombase, tobase)	INTTOHEX(N), HEXTOINT(N)	SQL Anywhere only provides functions that allow you to convert to and from hexadecimal. Other conversions have to be manually implemented using user defined functions.
HEX(arg)	INTTOHEX(numeric arg), HEXTOINT (string arg)	
CHAR(N, …)	CHAR(N)	SQL Anywhere function only supports on argument.
STRCMP(expr1, expr2)	COMPARE(expr1, expr2)	
LENGTH(str)	BYTE_LENGTH(str)	The SQL Anywhere LENGTH() function returns character length, not byte length.
BIT_LENGTH(str)	BYTE_LENGTH(str) * 8	BIT_LENGTH for SQL Anywhere

		returns the number of bits stored in a bit array. In SQL Anywhere, BIT_LENGTH('1001') would be 4. In MySQL, BIT_LENGTH('1001') would be 32.
LOCATE(substr, str [,pos])	LOCATE (str, substr [,pos])	The order of the arguments differs in SQL Anywhere.
POSITION(substr IN str)	LOCATE(str, substr)	
INSTR(str, substr)	LOCATE(str, substr)	
SUBSTINRG(str FROM pos [FOR len])	SUBSTIRNG(str, pos [,len])	
TRIM(str)	TRIM(str)	SQL Anywhere does not support the other forms of the MySQL TRIM() function.
INSERT(str, pos, len, newstr)	STUFF(str, pos, len, newstr)	
ELT(N, str1, str2, …)	ARGN(N, str1, str2,…)	

NUMERIC FUNCTIONS

MySQL function	**Equivalent SQL Anywhere function**	**Notes**
ROUND(x)	ROUND(x,0)	
x DIV y	FLOOR (x/y)	
LN(x)	LOG(x)	
LOG(x,b)	LOG(x)/LOG(b)	
LOG2(x)	LOG(x)/LOG(2)	
ATAN(x,y)	ATAN2(x,y)	

DATE AND TIME FUNCTIONS

MySQL function	**Equivalent SQL Anywhere function**	**Notes**
TIME(expr)	CAST(expr as TIME)	The CONVERT function could also be used.
TIMESTAMP(expr)	DATETIME(expr)	
DAYOFWEEK(expr)	DOW(expr)	
WEEKOFYEAR(expr)	DATEPART (Calweekofyear,expr)	
PERIOD_ADD(expr,N)	DATEFORMAT(DATEADD (month,N,expr || '01'), 'YYYYMM')	
PERIOD_DIFF(P1,P2)	ABS(DATEDIFF(month, P1||'01', P2||'01'))	
ADDATE(date, numdays)	DATEADD(day, numdays, date)	
SUBDATE(date, numdays)	DATEADD(day, -numdays, date)	
EXTRACT(type FROM date)	DATEPART(type, date)	The 'type' argument differs between SQL Anywhere and MySQL and must be adjusted accordingly.

TO_DAYS(date)	DAYS(date) – 58	SQL Anywhere measures differences in date from '0000/02/29' instead of '0000/01/01'
MAKEDATE(year, dayofyear)	YMD(year, 0, dayofyear)	
UNIX_TIMESTAMP(date)	DATEDIFF(second, '1979/01/01', date)	

SYNTAX MAPPINGS

   Most of the syntax features of MySQL are available in SQL Anywhere, but occasionally the syntax for accessing those features is different. The following charts detail many of these statements along with their SQL Anywhere equivalents. For specific examples of SQL Anywhere syntax listed below, see the "SQL Statements" chapter of the "SQL Anywhere Server - SQL Reference" manual.

OPERATORS

MySQL operator	Equivalent SQL Anywhere operator	Notes
!=	<>	
<=>	(expr1 = expr2 OR ((expr1 IS NULL) AND (expr2 IS NULL)))	The <=> operator represents equality including NULL values (NULL=NULL is true).
ISNULL(expr)	IS NULL expr	
INTERVAL(N, N1, N2,…)	None built in	A user defined function could easily be used to achieve the same functionality. For example: if (N < N1) then 0 elseif (N < N2) then 1 elseif…
!	NOT	
&&	AND	
||	OR	
A XOR B	((a AND (NOT b)) OR ((NOT a) AND b))	The SQL Anywhere expression is complex for large numbers of XOR arguments, so an alternative migration technique is recommended dependent on the application scenario.

DATA MANIPULATION LANGUAGE

MySQL statement	Equivalent SQL Anywhere statement	Notes
INSERT …	INSERT …	
ON DUPLICATE KEY UPDATE	ON EXISTING UPDATE	SQL Anywhere also offers the options ERROR and SKIP for existing rows.
SELECT …. INTO OUTFILE	UNLOAD SELECT DBISQL OUTPUT TO	
SELECT/UPDATE/DELETE … LIMIT	FIRST or TOP n	
DEFAULT '0' NOT NULL auto_increment	NOT NULL DEFAULT AUTOINCREMENT	
LIMIT offset, numRows	TOP numRows START AT offset	

Insert IGNORE	INSERT … ON EXISTING SKIP	
Replace …	INSERT … ON EXISTING UPDATE	
GROUP_CONCAT	LIST	
INSERT INTO … DEFAULT VALUES	INSERT INTO … VALUES (DEFAULT)	
LOAD DATA INFILE	LOAD TABLE	

MISCELLANEOUS SYNTAX

   The following is a miscellaneous list of compatibility items that do not fit into the aforementioned categories. It also includes mappings between functions that are not exactly the same, but are designed to provide the same functionality.

MySQL syntax	Equivalent SQL Anywhere syntax	Notes
VERSION()	@@version global variable	
mysql_insert_id()	@@identity global variable	
LAST_INSERT_ID variable	@@identity global variable	
mysql_affected_rows()	@@rowcount global variable	
ANALYZE TABLE	sa_table_page_usage, sa_table_fragmentation	SQL Anywhere also offers access to other properties via the property() function.
OPTIMIZE TABLE	CREATE STATISTICS	SQL Anywhere has a self-tuning optimizer that automatically maintains statistics, so statistics do not need to be updated manually.
CHECK TABLE	sa_validate () procedure	
USE database-name		There is no equivalent in SQL Anywhere. Each database running on a server requires its own connection.
LOCK TABLES (name) WRITE	LOCK TABLES table-name IN EXCLUSIVE MODE	SQL Anywhere supports row-level locking, so table locks are generally not required.
UNLOCK TABLES	COMMIT	A COMMIT releases all locks, unless a cursor is opened using the WITH HOLD clause.
Create table(KEY …)	CREATE TABLE… CREATE INDEX	SQL Anywhere requires two statements.
DO	CALL	
FLUSH/RESET	sa_flush_cache, sa_flush_statistics	Most of the other flushable elements in MySQL are automatically managed by SQL Anywhere and do not need to be flushed.
REGEXP/RLIKE	SIMILAR	SIMILAR works differently from the mysql REGEX syntax, but performs the same function. It may suit the needs where the MySQL REGEXP expression is being used.

BINARY str	CAST str AS BINARY	
CURDATE() | CURRENT_DATE()	CURRENT DATE	
CURTIME() | CURRENT_TIME()	CURRENT TIME	
SYSDATE() | LOCALTIME() | CURRENT_TIMESTAMP() | NOW()	NOW(), CURRENT TIMESTAMP	
UTC_DATE()	CURRENT UTC TIMESTAMP	
DATABASE()	CURRENT DATABASE	
LOAD_FILE(file)	xp_read_file(file)	In SQL Anywhere, the contents of a file are returned as a long binary field, while in MySQL they are returned as a string.
CONNECTION_ID()	CONNECTION_PROPERTY ('Number')	

OTHER MIGRATION ISSUES

The following is a list of miscellaneous notes to keep in mind when migrating from MySQL to SQL anywhere:

- The identifiers in MySQL are optionally enclosed with the back quote (`), while SQL Anywhere uses the double quote (") or , alternatively, square brackets ([]).
- Some words that are keywords in SQL Anywhere are not in MySQL, such as 'comment' and 'session'. These keywords must be enclosed in double quotes in order to be used with SQL Anywhere. Alternatively, you can use the SQL Anywhere NON_KEYWORDS option to change the list of recognized keywords. For information about the NON_KEYWORDS option, see "NON_KEYWORDS option [compatibility]" in the "Database Options" chapter of the "SQL Anywhere Server - Database Administration" manual.
- The minimum timestamp value in SQL Anywhere is '0001-01-01 00:00:00', while it is '0000-0000-00 00:00:00' in MySQL.
- Timestamps in MySQL have the format of YYYY-MM-DD hh:mm:ss. SQL Anywhere includes fractions of a second as part of the timestamp value. The TIME_FORMAT option allows you to specify the exact format used to return datetime value. For information about the TIME_FORMAT option, see "TIME_FORMAT option [compatibility]" in the "Database Options" chapter of the "SQL Anywhere Server - Database Administration" manual.
- While MySQL allows the use of single or double quotes around string literals, by default single quotes must be used to enclose string values in SQL Anywhere and double quotes signify the use of a database object identifier. This behavior can be changed by setting the QUOTED_IDENTIFIER option in the database. For information about the QUOTED_INDENTIFIER option, see "QUOTED_IDENTIFIER option [compatibility]" in the "Database Options" chapter of the "SQL Anywhere Server - Database Administration" manual.

MIGRATING A MYSQL DATABASE TO A SQL ANYWHERE DATABASE

Migrating data from MySQL to SQL Anywhere is a straightforward process, with minor issues occurring only if you are using the MySQL-specific data types mentioned previously. Data migration can be accomplished using the Data Migration wizard that is part of Sybase Central. Alternatively, a

more customized migration can be done using the sa_migrate set of stored procedures in SQL Anywhere. The mysqldump utility, coupled with the SQL Anywhere LOAD TABLE statement, could also be used to migrate the data. Note that if the MySQL SET or ENUM data types are used in the MySQL database, you may have some additional considerations when migrating your MySQL database to SQL Anywhere. For information about these data types and the differences from SQL Anywhere, see "Data types" on page 3.

REQUIREMENTS

- This document assumes you have a MySQL database running on any of its supported platforms and SQL Anywhere 10 installed on any of the supported Windows platforms.
- If you have not created a MySQL database, you can create a few tables in the MySQL test database to walk through the migration steps.
- The MySQL ODBC 3.51 (or later) driver must also be installed on the computer running the SQL Anywhere database.

CREATING A SQL ANYWHERE DATABASE

You must first create a SQL Anywhere database to migrate the MySQL database to. The following steps explain how to create a new database using Sybase Central.

1. Start Sybase Central. From the Start Menu, choose Program Files → SQL Anywhere 10→ Sybase Central.
2. Create a new SQL Anywhere 10 database. Choose Tools → SQL Anywhere 10 → Create Database. Follow the instruction in the wizard to create a new database.

CREATING A DATA SOURCE FOR THE MYSQL DATABASE

The migration process requires an ODBC connection to the source database. Therefore, you need to create an ODBC data source (DSN) for the MySQL database.

1. Download and install the MySQL ODBC 3.51 driver if you have not already done so. The most recent driver is located at http://dev.mysql.com/downloads/connector/odbc/3.51.html.
2. From Sybase Central, choose Tools → SQL Anywhere 10 → Open ODBC Administrator. The ODBC Data Source Administration dialog appears.
3. Click Add to add a new DSN. The Create New Data Source wizard appears.
4. Select the MySQL ODBC 3.51 Driver from the list of available drivers and then click Finish. The MySQL ODBC 3.51 Driver - DSN Configuration dialog appears.
5. Type a name for the data source in the Data Source Name field. For example, name the data source 'MySQL migrate'.
6. Supply the server name, user ID, password, and database name on the logic page of the MySQL ODBC Connector dialog.
7. Click the Advanced tab and ensure that the options 'Don't Optimize Column Width' and 'Return Matching Rows' are selected.
8. Click the Test Data Source button to ensure you have configured the data source correctly.
9. Click OK.

MIGRATING THE MYSQL DATABASE TO SQL ANYWHERE

The steps to migration are outlined below.

Connecting to the SQL Anywhere Database

In order to migrate the new SQL Anywhere database, you must first connect to the SQL Anywhere database.

1. If you are not already connected, from Sybase Central, choose Connects → Connect with SQL Anywhere 10. The Connect dialog appears.
2. On the Identification tab, type a valid User ID and password for your database. By default, there is a user 'DBA' with password 'SQL'.
3. On the Database tab, click the Browse button and then select the SQL Anywhere database file you created.
4. Click OK. The SQL Anywhere database server starts automatically.

Creating a Remote Server and External Login

The next step is to tell Sybase Central where to find the MySQL database. This is done by creating a remote server.
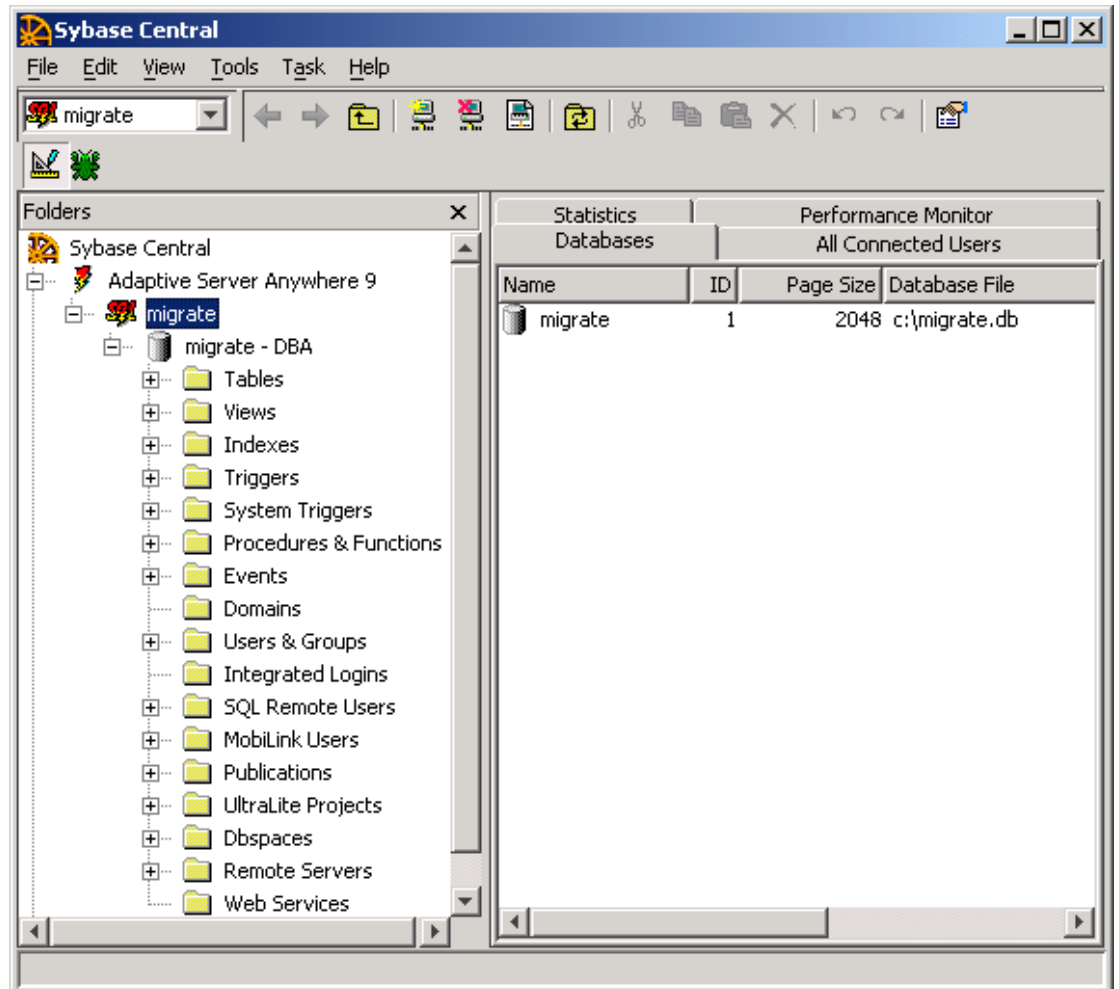

1. In the left pane of Sybase Central, expand your database server and database icons. In the example below, the database named migrate is running on a database server that is also named migrate.
2. In Sybase Central, select the Remote Servers folder in the left pane.
3. From the File men, choose New → Remote Server. The Remote Server Creation wizard appears.
4. Follow the instructions in the wizard to create a remote server that connects to your MySQL database.
   a. On the first page of the wizard, type a name for the remote server, for example 'MySQL migrate' and then click Next.

> b. Choose Generic as the type of remote server. Click Next.
>
> c. Select the Open Database Connectivity (ODBC) Option and type the name of the ODBC data source for your MySQL database in the connection information field. For example, if you named your ODBC data source 'MySQL migrate' when you created it, type 'MySQL migrate' in the connection information field.

5. Click Next.
6. Do not choose to make the server read only. Click Next.
7. If the remote server does not define a user that is the same as the user ID you are connected to the SQL Anywhere database with, you must create an external login for your current user. For example, if you connected to the SQL Anywhere database with the user ID DBA, and your MySQL database does not contain a user ID DBA, then you must create an external login. Type a user name from the MySQL database in the Login Name field. Type the password for this user in the Password and Confirm Password fields.
8. Use the "Test Connection" button to ensure you can connect. Then, click Finish.
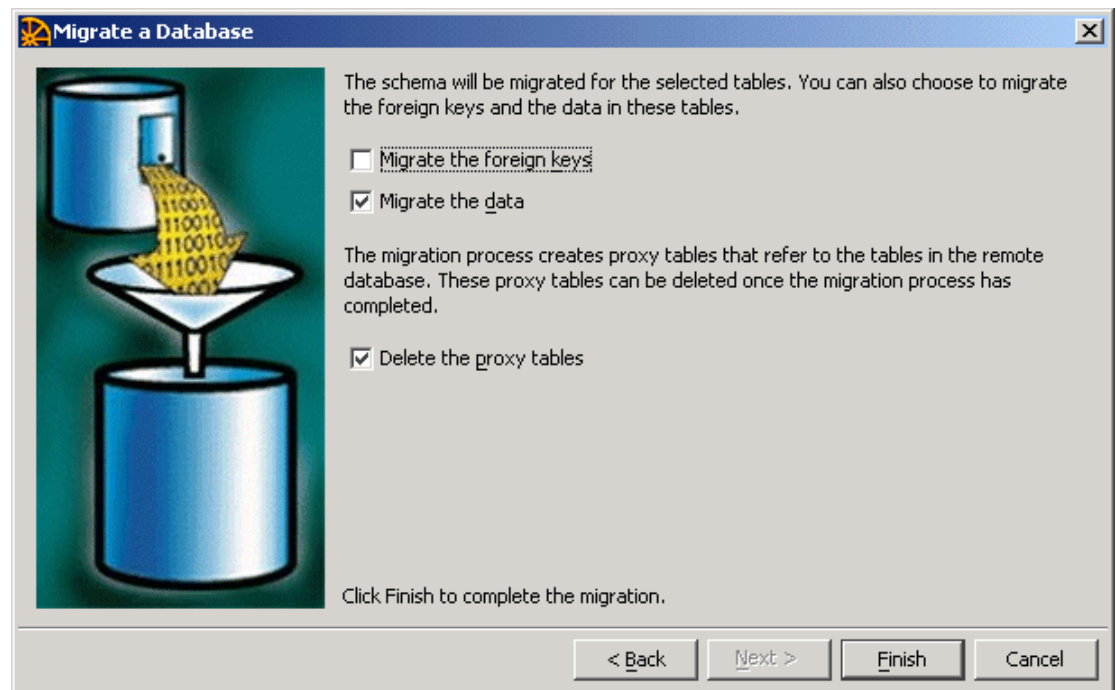
Migrating the MySQL database

Now you are ready to migrate your MySQL database: SQL Anywhere is running, connected and able to communicate to the MySQL database via ODBC. The next step is to use the Migration wizard to perform the migration.

1. From Sybase Central, choose Tools → SQL Anywhere 10 → Migrate database. The Database Migration Wizard appears.
2. Click Next on the introductory page.
3. Select the current database and then click Next.
4. Select the MySQL remote server you create, for example, MySQL migrate, and then click Next.
5. Click the Add All button and then click Next to migrate all the MySQL tables to the SQL Anywhere database.
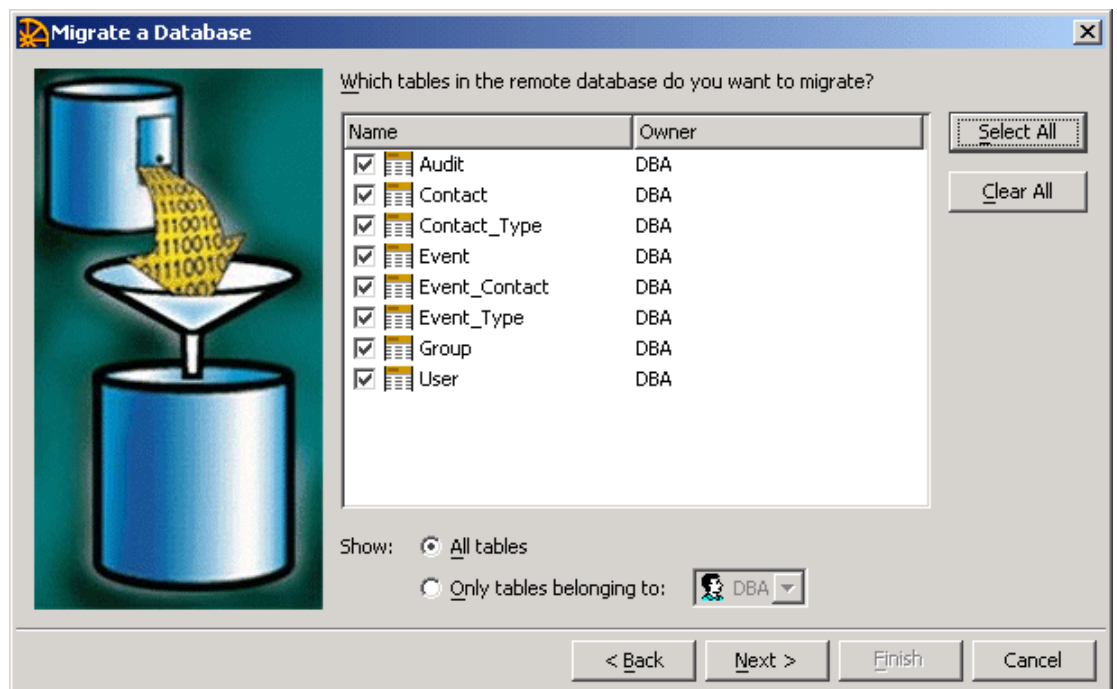
6. Select the SQL Anywhere database user you wish to own the tables. Click Next.
7. Select the options you wish to use. **Note**: Even for InnoDB users, the MySQL ODBC driver does not fully support inspecting foreign key relationships. For this reason, choose not to migrate foreign keys. After the data has been migrated, restore referential integrity as explained in 'Tweaking the New SQL Anywhere Database'.

8. Click Finish to start the migration. The Migration Database window appears. You can close this window when the status changes to 'Completed'.

TWEAKING THE NEW SQL ANYWHERE DATABASE

Now that you have migrated the MySQL schema and data to the SQL Anywhere database, you can start enjoying the benefit SQL Anywhere brings. One immediate benefit is transactional support.

Since not all MySQL tables support referential integrity, your MySQL schema may not have foreign keys. Even with InnoDB, in step 7 of 'Migrating the MySQL Database' we chose to ignore foreign key relationships. To add referential integrity support:
1. List the foreign keys in my MySQL database by issuing the following SQL statement against the MySQL database:

        SHOW TABLE STATUS FROM database_name

   Alternatively, `SHOW CREATE TABLE table_name` will also reveal any foreign key relationships. The referential constraints are listed under the comment column for each table in the form.

        (column_name) REFER ref_db_name/ref_table_name(ref_column_name)
2. Specify referential integrity constraints:
   a. You can use Sybase Central to add foreign keys to your database.
   b. Alternatively, for each of the foreign keys, issue the following SQL statement against the SQL Anywhere database (using the Interactive SQL utility: dbisql):

        ALTER TABLE "table-name"
            ADD FOREIGN KEY "foreign_key_name" ("column_name")

```
                         REFERENCES "ref_table_name" ("ref_column_name");
```
With the new foreign key constraints in place, the SQL Anywhere database checks for referential integrity automatically and greatly improves data integrity.

Properly placed indexes improve database performance significantly, while poorly placed ones hinder performance with equal significance. SQL Anywhere 10 offers the Index Consultant that inspects database usage and workload and recommends changes to the indexing structure as needed. MySQL dictates that foreign key columns must have indexes explicitly defined, but this is not the case in SQL Anywhere. Also, for each Primary key, MySQL creates a primary index that is redundant in SQL Anywhere. The Index Consultant will likely recommend removing the redundant indexes that are copied from the MySQL database during the migration process. The Index Consultant can prove to be a useful tool to boost the performance of the migrated SQL Anywhere database. For information about optimizing your schema, refer to your SQL Anywhere Studio documentation and the iAnywhere developer resources available online at www.ianywhere.com/developer.

MIGRATING APPLICATIONS FROM MYSQL TO SQL ANYWHERE

Application migration from MySQL to SQL Anywhere depends on the interface used to access your MySQL application.   The following are some of the more popular interfaces that should required only minimal work to migrate:

- **ODBC**: Both SQL Anywhere and MySQL support the ODBC 3.51 API specification. Generally, migration of these applications involves changes the ODBC data source to point to SQL Anywhere instead of MySQL. There may be some specific differences in terms of the implementations of creating API functions, but given the maturity of the ODBC specification, these should be minor.
- **JDBC:** MySQL has a type 4 JDBC Driver (100% Java implementation). To migrate to the SQL Anywhere equivalent, the Sybase jConnect driver should be used. However, to achieve the maximum performance benefits of SQL Anywhere, it is recommended that you use the iAnywhere JDBC driver. The iAnywhere JDBC driver is a type 2 JDBC driver. The SQL Anywhere JDBC drivers support all of the core elements of the JDBC 2.0 specification and some of the optional ones.
- **Perl:** For information about migrating Perl applications, see "Migrating a Perl application from MySQL to SQL Anywhere" on page 22.
- **PHP:** For information about migrating PHP applications, see "Migrating a PHP application from MySQL to SQL Anywhere" on page 23.

Applications written using the other interfaces supported by MySQL will require more work to migrate as there is no support for these drivers in SQL Anywhere. This includes the MySQL C/C++ API and the Python, Tcl, and Eiffel access drivers. In some cases, a third-party driver may be found that allows you to bridge to ODBC and natively access SQL Anywhere.

MIGRATING A PERL APPLICATION FROM MYSQL TO SQL ANYWHERE

Migrating Perl applications from MySQL to SQL Anywhere is very simple. You have the option of using ODBC to connect using the DBD::ODBC driver or using the native SQL Anywhere driver (called DBD:SQLAnywhere) that is included with SQL Anywhere.

If you are already using the DBD::ODBC driver, application migration is simply a matter of changing your connection string to refer to SQL Anywhere. Once that is complete, there may be some minor tweaks required to deal with the differences between SQL Anywhere and MySQL as discussed in pervious sections of this paper, but minimal work is required to complete the migration.

Some MySQL-specific methods can be migrated to SQL Anywhere equivalents by using queries or standard DBD functionality. For example:

MySQL	Equivalent SQL Anywhere	Note
mysql_insertid	SELECT @@identity	
is_blob, is_num, is_not_null, length, name, table, type	NAME, TYPE, SCALE, PRECISION, NULLABLE	All of these property items are DBD standard elements.
is_key, is_pri_key	SELECT .. FROM syscolumn WHERE	Detection of indexes/keys can be done by looking at the table and column definitions in the system tables.

MIGRATING A PHP APPLICATION FROM MYSQL TO SQL ANYWHERE

Migrating a PHP application from MySQL to SQL anywhere is simple. You have the option of using ODBC to connect to SQL Anywhere or using the SQL Anywhere PHP module.

Windows users may prefer to migrate to the ODBC API. Setting up a DSN in Windows for use with ODBC is simple. In addition, the Windows binary for PHP already has built-in ODBC support.

Linux users, on the other hand, may find the PHP module more convenient to set up. SQL Anywhere support can be compiled into PHP using the –with sqlanywhere=[path_to_sa] flag when calling the configure script. Details about the module can be found in the "SQL Anywhere PHP API" chapter of the "SQL Anywhere Server – Programming" manual.

If the PHP application is already using ODBC to connect to the MySQL database, then there is no need to change the function calls. You can skip the section below and go directly to "PHP migration notes" on page 25.

Function mapping

The MySQL, ODBC and SQL Anywhere APIs are very similar. It is often possible to map one function directly to another. Sometimes, when a function has no equivalent counterpart, you must be creative and write alternative code that achieves the same result. In certain cases, you may be better off rewriting small portions of the code to take advantage of advanced features provided by SQL Anywhere. For example, with transaction support, the application can efficiently maintain atomicity and easily ensure data integrity.

The following table lists some commonly used MySQL functions and their ODBC and SQL Anywhere equivalents:

MySQL	Equivalent SQL Anywhere (ODBC)	Equivalent SQL Anywhere (PHP Module)
mysql_close	odbc_close	sqlanywhere_disconnet
mysql_connect	odbc_connect	sqlanywhere_connect
mysql_errno	odbc_error	See "mysql_errno" note.
mysql_error	odbc_errormsg	none
mysql_escape_string	See "mysql_escape_string" note.	See "mysql_escape_string" note.
mysql_fetch_row	odbc_fetch_row	sqlanywhere_fetch_row
mysql_insert_id	See "mysql_insert_id" note.	See "mysql_insert_id" note.
mysql_num_fields	odbc_num_fields	sqlanywhere_num_fields
mysql_num_rows	odbc_num_rows	sqlanywhere_num_rows
mysql_query	odbc_exec	sqlanywhere_query
mysql_select_db	none	none

Notes:

mysql_connect: Connecting via ODBC is straight forward. The odbc_connect function takes, at a minimum, the DSN, user ID, and password. Connecting via the PHP module requires a SQL Anywhere connection string. Usually this can be done with the following function call:

```
sqlanywhere_connect("uid=DBA;pwd=SQL;eng=eng_name")
```

mysql_errno: This function returns the error number of the previous query. The same result can be obtained by issuing the following SQL statement:

```
SELECT @@error
```

mysql_escape_string: Neither ODBC nor the PHP module provides a way to escape a SQL string. However, this can be easily done by replacing each single quote with two single quotes.

mysql_insert_id: This function returns the last inserted ID of an autoincrement column. The same result can be obtained by issuing the following SQL statement:

```
SELECT @@identity
```

As you can see, many MySQL functions translate directly into ODBC and SQL Anywhere calls. For the remaining functions, simple alternatives can be found. As with any migration job, there are, unfortunately, differences between MySQL and SQL Anywhere that require more attention. These points are discussed in the following section.

PHP migration notes

These are subtle differences in the way SQL strings are treated by the various database vendors. For example, timestamps in MySQL have the format YYYY-MM-DD hh:mm:ss, while SQL Anywhere supports timestamps with fractions of a second. The strtotime function in PHP fails to recognize SQL Anywhere timestamps. Extra work must be done to remove the fractional second portion of the SQL Anywhere timestamp.

SQL Anywhere via ODBC also provides support for transactions and prepared statements. The odbc_commit and odbc_rollback functions terminate a transaction as you would expect. One point to notice is that PHP defaults to autocommit, meaning every statement is committed as soon as it is successfully executed. The odbc_autocommit function is used to set the autocommit behavior to enable the use of large transactions. Prepared statements are useful if the same queries, possibly with different parameters, are to be executed many times. This can help create efficiency as each dynamic SQL statement is built within the engine once only. The odbc_prepare and odbc_execute functions are used to execute prepared statements.

SUMMARY

Migrating from MySQL to SQL Anywhere involves migrating the database, changing MySQL function calls to SQL Anywhere calls, and tweaking the schema and SQL statements to resolve any differences between the databases. Typically, some performance gains can be achieved by utilizing advanced features available in SQL Anywhere.