



UltraLite® Database Management and Reference

Version 12.0.1

January 2012

Version 12.0.1
January 2012

Copyright © 2012 iAnywhere Solutions, Inc. Portions copyright © 2012 Sybase, Inc. All rights reserved.

This documentation is provided AS IS, without warranty or liability of any kind (unless provided by a separate written agreement between you and iAnywhere).

You may use, print, reproduce, and distribute this documentation (in whole or in part) subject to the following conditions: 1) you must retain this and all other proprietary notices, on all copies of the documentation or portions thereof, 2) you may not modify the documentation, 3) you may not do anything to indicate that you or anyone other than iAnywhere is the author or source of the documentation.

iAnywhere®, Sybase®, and the marks listed at <http://www.sybase.com/detail?id=1011207> are trademarks of Sybase, Inc. or its subsidiaries. ® indicates registration in the United States of America.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Contents

About this book	vii
UltraLite overview	1
UltraLite architecture	1
UltraLite synchronization client features	2
UltraLite supported platforms	3
UltraLite, UltraLite Java edition, and SQL Anywhere feature comparisons	3
UltraLite and UltraLite Java edition database limitations	8
UltraLite solution considerations for Windows Mobile	12
UltraLite database creation	17
Creating an UltraLite database with the Create Database Wizard	17
Create an UltraLite database at a command prompt	17
Create an UltraLite database from a MobiLink synchronization model	18
Create an UltraLite database using central administration of remote databases	18
Creating an UltraLite database from a SQL Anywhere reference database	19
Creating an UltraLite database from an XML file	20
Create an UltraLite database on a first connection	21
Specify UltraLite creation parameters	22
UltraLite character sets	23
UltraLite database security	26
UltraLite database connections	29
UltraLite connection strings and parameters	29
UltraLite connection parameters and the ULSQLCONNECT environment variable	31
UltraLite file path formats in connection parameters	31

- UltraLite database tasks and features 33**
 - Reading database properties 33
 - Accessing database options 34
 - UltraLite event notifications 34
 - UltraLite support for spatial data 37
 - Change the isolation level 40
 - Validating an UltraLite database 42
 - UltraLite and UltraLite Java edition database back up and recovery 43

- UltraLite database schemas 45**
 - UltraLite tables and columns 46
 - UltraLite indexes 54
 - UltraLite users 59

- UltraLite as a MobiLink client 67**
 - UltraLite clients 67
 - ActiveSync with UltraLite on Windows Mobile 85
 - UltraLite synchronization parameters and network protocol options 87

- UltraLite deployment 115**
 - UltraLite application build and deployment specifications 115
 - Deploy an UltraLite and UltraLite Java edition database to devices 122
 - Deploying UltraLite database schema upgrades 123
 - Start a deployed UltraLite engine 125
 - Deploying the ActiveSync provider for UltraLite 126
 - Registering applications with the ActiveSync Manager 127

- Tutorial: Building the UltraLite CustDB sample application 129**
 - CustDB sample overview 130
 - Lesson 1: Building and running the CustDB application 131
 - Lesson 2: Logging in and populating the UltraLite remote database 132
 - Lesson 3: Using the CustDB client application 133

Lesson 4: Synchronizing with the CustDB consolidated database	134
Lesson 5: Browsing MobiLink synchronization scripts	136
Building your own application	137
UltraLite database reference	139
UltraLite creation parameters	139
UltraLite connection parameters	165
UltraLite database properties	188
UltraLite database options	192
UltraLite utilities	197
UltraLite system tables	231
UltraLite Java edition database properties	235
UltraLite Java edition database options	237
UltraLite Java edition utilities	243
UltraLite Java edition system tables	251
UltraLite SQL reference	259
UltraLite SQL elements	259
UltraLite SQL functions	299
UltraLite SQL statements	410
UltraLite performance tips	455
Adjusting the cache size for an UltraLite database	455
UltraLite Java edition database cache sizes	456
Query performance tips	457
Insert and Update performance tips	468
UltraLite benchmark tips	472
UltraLite troubleshooting	479
Unable to start the UltraLite engine	479
Unable to connect to databases after upgrade	479
UltraLite database corruption	480
Database size not stabilizing	481

Importing ASCII data into a new UltraLite database	481
Utilities still running as the previous version	482
Result set changes unpredictably	483
UltraLite engine client fails with error -764	483
Index	485

About this book

This book introduces the UltraLite database system for small devices.

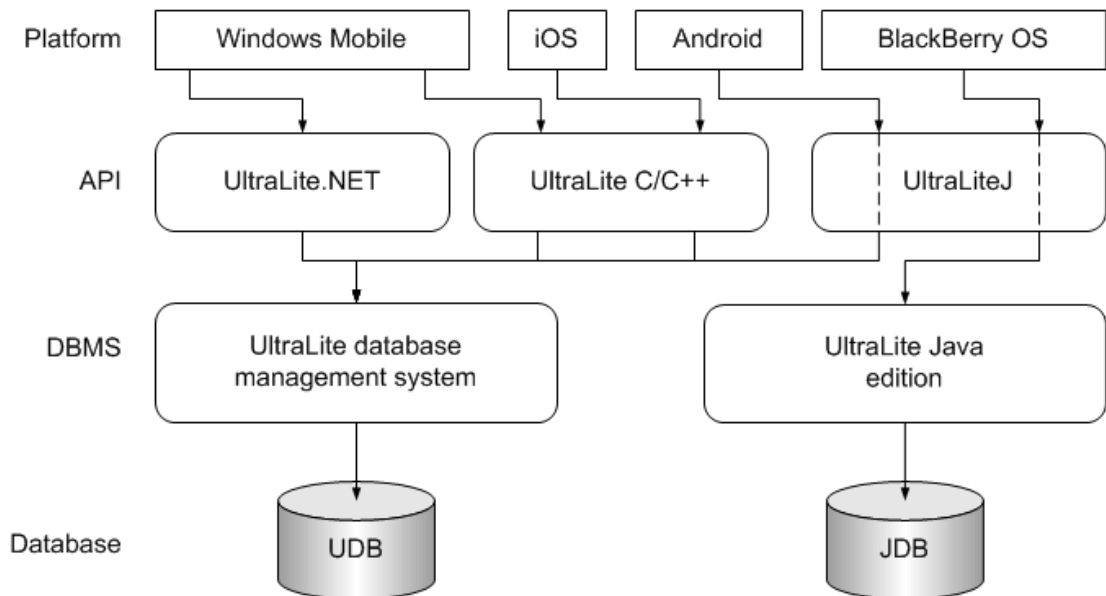
UltraLite overview

UltraLite is a compact relational database management system with many of the same features as SQL Anywhere. It can be used to create mobile databases for small-footprint devices such as smartphones, handheld computers, and tablet PCs.

UltraLite includes a built-in synchronization client that tracks changes made in UltraLite and UltraLite Java edition databases, and exchanges updates with a MobiLink server over a network. As a MobiLink client, UltraLite ensures that mobile applications can stay synchronized with a central database and with other UltraLite databases.

UltraLite architecture

UltraLite supports a multitude of mobile platforms and consists of API development, database management, and database layers, as represented in the following diagram:



- **Mobile platform support** Your target mobile platform determines which UltraLite API is available for application development. For more information about supported platforms and devices, see <http://www.sybase.com/detail?id=1061806>.
- **API development layer** Refer to the diagram above to determine which API to use for your target mobile platform. For more information about API development, see:
 - “UltraLite C++ application development” [*UltraLite - C and C++ Programming*]
 - “UltraLite.NET application development” [*UltraLite - .NET Programming*]
 - “UltraLiteJ application development” [*UltraLite - Java Programming*]

- **Database management layer and synchronization client** UltraLite provides two data management systems—the native **UltraLite database management system**, and the **UltraLite Java edition**.
 - **UltraLite database management system** Most mobile platforms use the UltraLite APIs to interface with the UltraLite database management system. This system allows you to create and connect to an **UltraLite database**.

A comprehensive set of administration tools is provided to help you maintain your UltraLite project. You can run these tools as either command line utilities or wizards in the UltraLite plug-in for Sybase Central. For more information, see [“UltraLite utilities” on page 197](#).

- **UltraLite Java edition** The BlackBerry OS platform requires use of the UltraLiteJ API to interface with the UltraLite Java edition to create and connect to an **UltraLite Java edition database**. These databases are not interchangeable with UltraLite databases and have some restrictions. For more information, see [“UltraLite and UltraLite Java edition database limitations” on page 8](#).

UltraLite provides a set of Java edition utilities to perform additional tasks on UltraLite Java edition databases. For more information, see [“UltraLite Java edition utilities” on page 243](#).

For more information about the UltraLite synchronization client, see [“UltraLite synchronization client features” on page 2](#).

- **Database layer** This layer is the local data repository stored as a file. UltraLite databases are stored as UDB files, and UltraLite Java edition databases are stored as JDB files. UDB files are portable across all mobile platforms except for BlackBerry OS.

UltraLite synchronization client features

UltraLite includes a built-in bi-directional synchronization client that causes all data in an UltraLite database to be synchronized by default. Users new to MobiLink synchronization may use this default behavior until business requirements necessitate a custom synchronization design to alter what UltraLite data gets synchronized to the consolidated database. Unlike SQL Anywhere remote databases, you do not need to increase the size of the UltraLite footprint to include synchronization functionality.

Important synchronization features built into the UltraLite runtime include a row state tracking mechanism and a synchronization state tracking mechanism.

The row state tracking mechanism

Tracking the state of tables and rows is particularly important for data synchronization. Each row in an UltraLite database has an associated row state structure. In addition to synchronization, UltraLite also uses the row states to control transaction processing and data recovery.

Synchronization state tracking

UltraLite uses a progress counter to ensure robust synchronization. Each upload is given a unique number to identify it. This allows UltraLite to determine whether an upload was successful when a communication error occurs.

When you first create a new database, UltraLite always sets the synchronization progress counter to zero. A progress counter value of zero identifies the database as a new UltraLite database, which tells the MobiLink server to reset its state information for this client.

Caution

Because UltraLite increments the progress counter each time a synchronization occurs, you cannot synchronize an UltraLite database to different consolidated databases. If the progress counter value is not zero and does not match that sequence number stored in the consolidated database, MobiLink synchronization reports an offset mismatch and synchronization fails. You cannot replace an UltraLite database with a backup copy if the progress counter is older than the current value.

See also

- “UltraLite database row state management” on page 469
- “MobiLink synchronization” [*MobiLink - Getting Started*]
- “UltraLite as a MobiLink client” on page 67
- “Tutorial: Using MobiLink with a SQL Anywhere consolidated database” [*MobiLink - Getting Started*]

UltraLite supported platforms

For more information about the devices, platforms, and network protocols supported by UltraLite, see <http://www.sybase.com/detail?id=1061806>.

See also

- “UltraLite as a MobiLink client” on page 67
- “Network protocol options for UltraLite synchronization streams” on page 112
- “UltraLite Synchronization utility (ulsync)” on page 221

UltraLite, UltraLite Java edition, and SQL Anywhere feature comparisons

The following table illustrates feature comparisons between UltraLite and SQL Anywhere. UltraLite features are supported by UltraLite Java edition unless otherwise specified.

Note

The UltraLite database management system adds 750-1000 KB to the size of your application while UltraLite Java edition adds 500 KB. The SQL Anywhere database, database server, and synchronization client add approximately 6 MB.

Feature	SQL Anywhere	UltraLite	Considerations
Transaction processing, and multi-table joins	X	X	
Triggers, stored procedures, and views	X		
External stored procedures (callable external DLLs)	X		
Built-in referential and entity integrity	X	X	UltraLite Java edition does not enforce foreign key constraints. See “Avoiding synchronization issues with foreign key cycles” on page 77.
Cascading updates and deletes	X	Limited ¹	Declarative referential integrity, where deletes and updates are cascaded, is a feature that is not supported in UltraLite databases, except during synchronization when deletes are cascaded for this purpose.
Dynamic, multiple database support	X	X	
Multi-threaded application support	X	X	
Row-level locking	X	X	

Feature	SQL Anywhere	UltraLite	Considerations
XML unload and load utilities		X	<p>UltraLite uses separate administration tools to complete XML load and unloads. These features are not built into the runtime.</p> <p>For UltraLite, see:</p> <ul style="list-style-type: none"> • “UltraLite Load XML to Database utility (ulload)” on page 217 • “UltraLite Database Unload utility (ulunload)” on page 227 <p>For UltraLite Java edition, see:</p> <ul style="list-style-type: none"> • “UltraLite Java Edition Database Load utility (ULjLoad)” on page 244 • “UltraLite Java Edition Database Unload utility (ULjUnload)” on page 246
XML export and import utilities	X		
SQLX functionality	X		
SQL functions	X	X	Not all SQL functions are available for use in UltraLite applications. If you use an unsupported function, you trigger an error. See “UltraLite SQL functions” on page 299 .
SQL statements	X	Limited	The scope of SQL statements are limited in UltraLite compared to SQL Anywhere. See “UltraLite SQL statements” on page 410 .
Integrated HTTP server	X		
Strong encryption for database files and network communications	X	X	
Event scheduling and handling	X	X ¹	An UltraLite event model differs from SQL Anywhere.

Feature	SQL Anywhere	UltraLite	Considerations
High-performance, self-tuning, cost-based query optimizer	X		UltraLite has a query optimizer that is not as extensive as that of SQL Anywhere. Therefore, UltraLite does not perform as highly on complex queries as SQL Anywhere.
Choice of several thread-safe APIs	X	X	UltraLite gives application developers a uniquely flexible architecture that allows for the creation of applications for changing and/or varied deployment environments. See “Choose an UltraLite API for Windows Mobile” on page 13.
Cursor support	X	X	See “UltraLite and UltraLite Java edition database limitations” .
Dynamic cache sizing	X	X ¹	<p>UltraLite allows you to set an initial, minimum, and/or maximum cache size for a database. The size of the cache is optimized by UltraLite on an ongoing basis, up to the maximum size (if specified). See:</p> <ul style="list-style-type: none"> • “UltraLite CACHE_SIZE connection parameter” on page 169 • “UltraLite CACHE_MIN_SIZE connection parameter” on page 168 • “UltraLite CACHE_MAX_SIZE connection parameter” on page 167 <p>UltraLite Java edition databases support fixed cache sizes only.</p>
Database recovery after system or application failure	X	X	
Binary Large Object (BLOB) support	X	X	UltraLite cannot index or compare BLOBs.
Windows Performance Monitor integration	X		

Feature	SQL Anywhere	UltraLite	Considerations
Online table and index defragmentation	X		
Online backup	X		
Direct device connections to a Windows Mobile device from the desktop.		X ¹	SQL Anywhere databases need a database server before allowing desktop connections to the database that you deploy on a Windows Mobile device. On UltraLite, you prefix the connection string with WCE:\ . See “Windows Mobile” on page 32 .
High-performance updates and retrievals through the use of indexes	X	X	UltraLite uses a mechanism to determine whether each table is searched using an index or by scanning the rows directly. Additionally, you can hash indexes to speed up data retrieval. See “UltraLite max_hash_size creation parameter” on page 149 .
Synchronizing to Oracle, DB2, Sybase Adaptive Server Enterprise, Microsoft SQL Server, MySQL, or SQL Anywhere	X	X	
Built-in synchronization		X	Unlike SQL Anywhere deployments, UltraLite does not require a client agent for synchronization. Synchronization is built into the UltraLite runtime to minimize the components you need to deploy. See “UltraLite clients” on page 67 .
In-process execution		X	
Computed columns	X		
Declared temporary tables/global temporary tables	X		

Feature	SQL Anywhere	UltraLite	Considerations
System functions	X		
Timestamp columns	X	X	SQL Anywhere Transact-SQL timestamp columns are created with the DEFAULT TIMESTAMP default. UltraLite timestamp columns are created with the DEFAULT CURRENT TIMESTAMP default. Therefore, UltraLite does not automatically update the timestamp when the row is updated.
User-based permission scheme to determine object-based ownership and access	X		UltraLite is primarily designed for single user databases in which an authorization system is not needed. However, you can include up to four user IDs and passwords, which are used for authentication purposes only. These users have access to all database objects. See “UltraLite users” on page 59 .
Spatial data	X	Limited	UltraLite supports point data only.
Full text data	X		

¹ Not available for UltraLite Java edition.

UltraLite and UltraLite Java edition database limitations

The following table illustrates hard limits that apply to UltraLite and UltraLite Java edition databases. In many cases, the limits are beyond the maximum capabilities of mobile devices. Performance considerations and device capabilities impose stricter limitations.

Item	UltraLite database limitations	UltraLite Java edition database limitations
Database and file size	1 file per database. The file size is limited by the operating system and file system.	1 file per database. The file size is limited by the operating system and file system.

Item	UltraLite database limitations	UltraLite Java edition database limitations
Temporary file size	Limited by the operating system.	Limited by the operating system.
Cache size	Limited by the available memory on the device.	Limited by the available memory on the device.
Dynamic cache sizing	<p>UltraLite allows you to set an initial, minimum, and/or maximum cache size for a database. The size of the cache is optimized by UltraLite on an ongoing basis, up to the maximum size (if specified). See:</p> <ul style="list-style-type: none"> ● “UltraLite CACHE_SIZE connection parameter” on page 169 ● “UltraLite CACHE_MIN_SIZE connection parameter” on page 168 ● “UltraLite CACHE_MAX_SIZE connection parameter” on page 167 	UltraLite Java edition databases support fixed cache sizes only.
Maximum number of concurrent open connections supported by a database	Up to 14.	No limit.
Maximum number of concurrent open connections to all databases	Up to 16 on a mobile device, and up to 64 on a desktop.	No limit.
Maximum number of databases that can run concurrently	The UltraLite engine supports up to 8 on a mobile device, and up to 32 on a desktop.	No limit.
Direct device connections to a Windows Mobile device from the desktop.	You prefix the connection string with WCE:\ before allowing desktop connections to the database that you deploy on a Windows Mobile device. See “Windows Mobile” on page 32 .	Not applicable.
SQL communication areas (SQLCA)	Up to 63.	Not applicable.
File-based persistent store (database size)	2 GB file or operating system limit on file size.	Limited by the operating system.

Item	UltraLite database limitations	UltraLite Java edition database limitations
Rows per table	<p>Up to 16 million.</p> <p>Sometimes changes to the row (deletes and updates) and other state information are maintained with the row data. This information allows those changes to be synchronized. So, the actual row limit can be smaller than 16 million, depending on the number of transactions on a table between synchronization, or whether the table is synchronized at all. See “UltraLite transaction processing” on page 470.</p>	<p>Up to 16 million.</p> <p>Sometimes changes to the row (deletes and updates) and other state information are maintained with the row data. This information allows those changes to be synchronized. So, the actual row limit can be smaller than 16 million, depending on the number of transactions on a table between synchronization, or whether the table is synchronized at all. See “UltraLite transaction processing” on page 470.</p>
Row Size	<p>The length of each packed row must not exceed the page size. See “Row packing and table definitions” on page 46.</p> <p>Character strings are stored without padding when they are shorter than the column size. This restriction excludes columns declared as LONG BINARY and LONG VARCHAR as these are stored separately.</p>	<p>Row contents (after possible compression) must not exceed the database page size.</p>
Rows per database	Limited by the persistent store.	Limited by the persistent store.
Table size	Limited by the database size.	Limited by the database size.
Tables per database	Limited by the database size.	Up to 32000.
Columns per table	<p>Row size is limited by page size, so the practical limit on the number of columns per table is derived from this size. Typically, this practical limit is much less than 4000.</p>	<p>Row size is limited by page size, so the practical limit on the number of columns per table is derived from this size. Typically, this practical limit is much less than 4000.</p>
Indexes per table	Limited by the database size.	Limited by the database size.
Tables referenced per transaction	No limit.	No limit.

Item	UltraLite database limitations	UltraLite Java edition database limitations
Stored procedure length	Not applicable.	Not applicable.
Stored procedures per database	Not applicable.	Not applicable.
Triggers per database	Not applicable.	Not applicable.
Nesting	Not applicable.	Not applicable.
Number of publications	Up to 63.	Up to 63.
Number of database pages	Limited by file size.	Up to 2 ¹⁶ .
Database page size	16 KB.	Minimum 256 bytes; up to 16 KB.
Cursors per connection	The number of maximum allowable cursors on a given connection to an UltraLite database is 64 (all platforms).	The number of maximum allowable cursors on a given connection to an UltraLite database is 64 (all platforms).
Strings	The row must fit on a page.	The row must fit on a page.
Binary data types	The row must fit on a page.	The row must fit on a page.
Long binary/long varchar size	Limited only by database size.	Limited only by database size.
Blob size	Limited by file size.	Up to 2 ²⁴ bytes.
Available Isolation levels	0 (read uncommitted) or 1 (read committed).	0 (read uncommitted).
Cascading updates and deletes	Declarative referential integrity, where deletes and updates are cascaded, is a feature that is not supported in UltraLite databases, except during synchronization when deletes are cascaded for this purpose.	Not supported.
Event scheduling and handling	An UltraLite event model differs from SQL Anywhere.	Not supported.

Item	UltraLite database limitations	UltraLite Java edition database limitations
UltraLite Java edition database compatibility	<p>UltraLite databases are not interchangeable with UltraLite Java edition databases. An UltraLite database can be converted into an UltraLite Java edition database and vice versa using the respective load and unload utilities. See:</p> <ul style="list-style-type: none"> ● “UltraLite Database Unload utility (ulunload)” on page 227 ● “UltraLite Load XML to Database utility (ulload)” on page 217 	<p>UltraLite databases are not interchangeable with UltraLite Java edition databases. An UltraLite database can be converted into an UltraLite Java edition database and vice versa using the respective load and unload utilities. See:</p> <ul style="list-style-type: none"> ● “UltraLite Java Edition Database Unload utility (ULjUnload)” on page 246 ● “UltraLite Java Edition Database Load utility (ULjLoad)” on page 244

Object store limitations for BlackBerry

On BlackBerry smartphones, the size of the database store is limited by the number of object handles available. The number of available object handles is a BlackBerry feature and is determined by the size of the flash memory.

UltraLite requires one or more object handles, depending on the handle type, to store a database value in memory. For example, a table row with ten columns and two indexes requires a minimum of twelve object handles.

To permit larger database stores, the UltraLiteJ API allows users to limit the number of rows kept in memory. Stored database rows are combined on a database page. The UltraLiteJ API requires only one persistent object handle for each database page.

See also

- [“SQL Anywhere size and number limitations” \[SQL Anywhere Server - Database Administration\]](#)

UltraLite solution considerations for Windows Mobile

This section outlines the UltraLite design options that are available for Windows mobile development, and the implementation advantages for each option. The following selections are described:

- **UltraLite API selection** The benefits of using each of the following APIs for Windows Mobile development are described:
 - UltraLite C/C++
 - UltraLite embedded SQL
 - UltraLite.NET
- **Data management component selection** The benefits of using each of the following data management components for Windows Mobile development are described:
 - UltraLite in-process runtime environment
 - UltraLite database engine

Choose an UltraLite API for Windows Mobile

The UltraLite C/C++, embedded SQL, and .NET APIs offer several data access models, including a simple table-based data access interface and dynamic SQL for more complex queries. By combining these benefits, UltraLite gives application developers a flexible architecture for creating applications for their varied deployment environments.

UltraLite.NET API benefits

The UltraLite.NET API is usually recommended for Windows Mobile development because the SQL Anywhere .NET API provides common programming models that are shared between UltraLite components and SQL Anywhere, and because of the .NET programming compared C/C++.

UltraLite C/C++ and embedded SQL API benefits

While UltraLite provides high performance in a variety of environments and use cases, embedded SQL and the UltraLite C/C++ API are the lowest level APIs and generally deliver the highest performance.

Use the UltraLite C/C++ API when you are trying to create the smallest application footprint. These applications typically yield the best performance and still maintain a small application file size.

See also

- [“.NET application programming” \[SQL Anywhere Server - Programming\]](#)
- [“UltraLite C++ application development” \[UltraLite - C and C++ Programming\]](#)
- [“UltraLite C++ application development using embedded SQL” \[UltraLite - C and C++ Programming\]](#)
- [“UltraLite.NET application development” \[UltraLite - .NET Programming\]](#)

Choose an UltraLite data management component for Windows Mobile

UltraLite allows you to build a small-footprint relational database solution without requiring the additional overhead of setting up a separate database server. Instead, UltraLite programming interfaces

use one of two approaches—the **UltraLite in-process runtime library** and the **UltraLite database engine**. Both approaches control connection and data access requests.

Both components include a built-in bi-directional synchronization client that links UltraLite databases with the MobiLink synchronization server. For more information about the UltraLite synchronization client, see [“UltraLite synchronization client features” on page 2](#).

UltraLite in-process runtime library

The UltraLite in-process runtime is recommended when only one application needs to access a database at a time.

The runtime and the application are part of the same process, which makes the database specific to the application. The runtime library manages UltraLite databases and built-in synchronization operations.

Linking to the runtime requires a different import library/DLL pair from that of the engine.

UltraLite supports both static and dynamic linkage.

- **Static linking** Static linking requires less device memory and is more effective when only a single UltraLite application is used on the device.
- **Dynamic linking** Dynamic linking may be more economical with device memory when multiple UltraLite applications are used on the device.

UltraLite database engine (the *uleng12.exe* utility)

The UltraLite engine is only available for Windows desktop and Windows Mobile platforms. The engine is a separate executable that supports concurrent access from multiple applications. Each application must use a client library to communicate with the UltraLite engine.

The UltraLite engine requires more system resources than the UltraLite runtime and may yield lower performance when large amounts of data are moved between the client and database.

Connecting to the engine requires that you specify a different import library/DLL pair than that of the runtime.

The UltraLite engine is required under the following conditions:

- Multiple processes access the same database file at potentially the same time (same time means multiple processes have connections open to the same database at the same time).
- Central Administration is used to manage the UltraLite application database.

See also

- [“UltraLite application build and deployment specifications” on page 115](#)
- [“Build and deploy UltraLite C++ applications” \[*UltraLite - C and C++ Programming*\]](#)
- [“Start a deployed UltraLite engine” on page 125](#)
- [“UltraLite Engine utility \(uleng12\)” on page 206](#)
- [“UltraLite concurrency” on page 468](#)
- [“UltraLite, UltraLite Java edition, and SQL Anywhere feature comparisons” on page 3](#)
- [“UltraLite and UltraLite Java edition database limitations” on page 8](#)

UltraLite database creation

There are three common types of database creation methods:

- Desktop creation methods with UltraLite administration tools designed for this purpose.
- On-device creation methods with UltraLite APIs. On-device creation methods are primarily described in each API specific UltraLite programming book.
- A Central Administration remote task, configured to create an UltraLite database on a device.

Once the database is created, you can connect to it and build tables and other database objects.

See also

- [“Specify UltraLite creation parameters” on page 22](#)
- [“Accessing database options” on page 34](#)

Creating an UltraLite database with the Create Database Wizard

You can create a database in Sybase Central using the **Create Database Wizard**.

Choose this method if you want help navigating the available database creation parameters. This wizard simplifies your choices by restricting what you can configure based on the target platform(s) you select. Once the database is created, it displays the command line syntax that you can record and use with the ulinit utility.

Create a new UltraLite database (Sybase Central)

1. Click **Start » Programs » SQL Anywhere 12 » Administration Tools » Sybase Central**.
2. Click **Tools » UltraLite 12 » Create Database**.
3. Follow the instructions in the **Create Database Wizard**.

Create an UltraLite database at a command prompt

Choose any of the following utilities to create a database from a command prompt:

- Use the ulinit utility if you want to create a new, empty UltraLite database or one sourced from a SQL Anywhere reference database schema. With this utility, you can include utility options to configure the database. See [“Creating an UltraLite database from a SQL Anywhere reference database” on page 19](#).

- Use the uload utility if you have an XML file that will serve as the source point for the schema and/or data of your new UltraLite database. See [“Creating an UltraLite database from an XML file” on page 20](#).
- **Central Administration:** Choose this method if you have a deployment where the MobiLink Agent is configured on all your deployed devices, or you are unable to deploy your initial UltraLite database with your application. You can configure a remote task to create a new UltraLite database on the device. This database can then be managed centrally by an administrator. See [“Central administration of remote databases” \[MobiLink - Server Administration\]](#).

Create a new UltraLite database (command line)

Run the ulinit utility specifying the new UltraLite database file if you want to accept the defaults:

```
ulinit test.udb
```

For information about the supported options, see [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#).

Create an UltraLite database from a MobiLink synchronization model

To simplify development, MobiLink includes a **Create Synchronization Model Wizard** to create your UltraLite database and server-side synchronization logic.

Choose this method if you are creating a synchronization system with remote UltraLite databases and a centralized consolidated database.

Once you have created your model, you can work in MobiLink Model mode in Sybase Central to customize your synchronization model before you deploy it. When the model is ready, you can then deploy it to generate the scripts and tables required for your synchronization application.

See also

- [“MobiLink plug-in for Sybase Central” \[MobiLink - Getting Started\]](#)

Create an UltraLite database using central administration of remote databases

MobiLink provides a create database command that allows you to create an UltraLite database.

See also

- [“Create database command” \[MobiLink - Server Administration\]](#)
- [“Central administration of remote databases” \[MobiLink - Server Administration\]](#)

Creating an UltraLite database from a SQL Anywhere reference database

A reference database is a SQL Anywhere database that serves as a template for the UltraLite database you are creating. Your UltraLite database is a subset of the columns, tables, and indexes in the reference database. You select these objects as part of a publication in the reference database. You can also choose to include data from your SQL Anywhere database in your new UltraLite database.

Creating a database from a reference database may be useful if you want to first model your data with an architecture tool such as Sybase PowerDesigner.

To create a database from a reference database, use the ulinit utility. See [“Creating an UltraLite database from a SQL Anywhere reference database” on page 19](#) and [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#).

Initialize/extract a new UltraLite database from a reference database (command line)

1. Create a new SQL Anywhere database as the reference database.

You can create a new SQL Anywhere database with the dbinit utility or Sybase Central. You can also create a SQL Anywhere database from non-SQL Anywhere databases, by migrating data from these third-party files.

See [“SQL Anywhere database creation” \[SQL Anywhere Server - Database Administration\]](#).

Configure the database with UltraLite usage in mind The UltraLite database is generated with the same settings as those in the reference database. By setting the following options in the reference database, you also control the behavior of your UltraLite database:

- Date format (see [“UltraLite date_format creation parameter” on page 144](#))
- Date order (see [“UltraLite date_order creation parameter” on page 147](#))
- Nearest century (see [“UltraLite nearest_century creation parameter” on page 151](#))
- Precision (see [“UltraLite precision creation parameter” on page 155](#))
- Scale (see [“UltraLite scale creation parameter” on page 156](#))
- Time format (see [“UltraLite time_format creation parameter” on page 157](#))
- Timestamp format (see [“UltraLite timestamp_format creation parameter” on page 159](#))
- Timestamp with timezone (see [“UltraLite timestamp_with_time_zone_format creation parameter” on page 163](#))

2. Prepare the reference database by adding objects required by the UltraLite database:

- **Tables and keys** Add the tables and remember to set primary keys as they are required by UltraLite. If you need to, you can also assign foreign key relationships that you need within your UltraLite application. You can use Sybase Central, Sybase PowerDesigner Physical Data Model, or another database design tool. See [“UltraLite tables and columns” on page 46](#).
- **Indexes** An index can improve performance dramatically, particularly on slow devices. Note that primary keys are automatically indexed, but other columns are not. See [“When to use an index” on page 56](#).

- **Publications** If you want to synchronize different tables at different times, use publications. You can use multiple UltraLite publications to define table subsets and set synchronization priority with them. See [“Publishing data in UltraLite” on page 76](#).

Tip

If your UltraLite applications frequently retrieve information in a particular order, consider adding an index to your reference database specifically for this purpose. See [“Use index scans” on page 457](#).

3. Run the ulinit utility, including any necessary options:

```
ulinit -a DBF=MySource.db customer.udb -n Pub1 -s Pub2
```

In this example, MySource.db is the SQL Anywhere reference database and customer.udb is the UltraLite database that gets created. The tables and indexes will match those contained in the Pub1 and Pub2 publications. A publication for the UltraLite database is created for Pub2.

See also

- [“Creating an UltraLite database with the Create Database Wizard” on page 17](#)
- [“Create an UltraLite database at a command prompt” on page 17](#)
- [“Creating an UltraLite database from an XML file” on page 20](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“Specify UltraLite creation parameters” on page 22](#)
- [“UltraLite upgrades” \[SQL Anywhere 12 - Changes and Upgrading\]](#)

Creating an UltraLite database from an XML file

You can use XML as an intermediate format for managing your UltraLite database if the format follows the requirements for UltraLite usage. You can use XML as follows:

- Load data into a new database with a different set of database properties/options.
- Upgrade the schema from a database created by a previous version of UltraLite.
- Create a text version of your UltraLite database.

Create an UltraLite database from an XML file

UltraLite cannot use an arbitrary XML file. The `%SQLANY12%\Bin32` and `%SQLANY12%\Bin64` directories contains a `usm.xsd` file, containing the schema definition. Use this file to review the XML format.

1. Save the XML file to a directory of your choosing. You can either:
 - Export/unload a database to an XML file. If you are unloading a SQL Anywhere database, use any of the supported export methods.

- Take XML output from another source—that source could be another relational database or even a web site where transactions are recorded to a file. You must always ensure that the format of the XML meets the UltraLite requirements.
2. Run the ulload utility, including any necessary options.

For example, to create a new UltraLite database in the file *sample.udb* from the table formats and data in *sample.xml*:

```
ulload -c DBF=sample.udb sample.xml
```

See also

- “Exporting relational data as XML” [*SQL Anywhere Server - SQL Usage*]
- “Creating an UltraLite database with the Create Database Wizard” on page 17
- “Create an UltraLite database at a command prompt” on page 17
- “Creating an UltraLite database from a SQL Anywhere reference database” on page 19
- “UltraLite upgrades” [*SQL Anywhere 12 - Changes and Upgrading*]
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “Specify UltraLite creation parameters” on page 22

Create an UltraLite database on a first connection

You can program your application to create a new UltraLite database if one cannot be detected at connection time. The application can then use SQL to create tables, indexes, foreign keys, and so on. To populate the database, synchronize with a consolidated database.

Considerations

By adding the additional database creation and SQL code, your application size can grow considerably. However, this option can simplify deployment because you only need to deploy the application to the device. In some pre-production development cycles, you may want to delete and reconstruct the database on your device for testing purposes.

See also

- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Creating or connecting to an UltraLite or UltraLite Java edition database” [*UltraLite - Java Programming*]

Specify UltraLite creation parameters

Creation parameters are used to configure the UltraLite database when you first create it. You can only change these settings by recreating the database.

There are several options you can control when creating your UltraLite database. These are designed to help with the wide variety of UltraLite uses. Note that most parameters specified at creation time cannot be changed later.

You can specify creation parameters when creating a database using the `ulinit` or `ulload` utility, and from the supported client interfaces.

Boolean creation parameters are turned on with `YES`, `Y`, `ON`, `TRUE`, `T`, or `1`, and are turned off with any of `NO`, `N`, `OFF`, `FALSE`, `F`, and `0`. The parameters are case insensitive.

UltraLite creation parameters are specified in a semicolon separated string when creating a database from a programming interfaces and as a separate command line parameters when using a command line utility. For example:

```
ulinit --case --utf8_encoding=1 test.udb
```

Alternatively, you can specify multiple `-c` options.

You can further configure other aspects of your database with database options and connection parameters.

Accessing creation parameter values

You cannot change creation parameter values after you have created a database. However, you can view the corresponding database properties in Sybase Central.

For the UltraLiteJ API, you can use the `getCreationString` method to view the creation string registered with the `setCreationString` method.

For other UltraLite APIs, you can access the database properties programmatically from your UltraLite application by calling the `GetDatabaseProperty` function appropriate to the UltraLite API.

See also

- “UltraLite database options” on page 192
- “UltraLite database connections” on page 29
- “UltraLite connection parameters” on page 165
- “Reading database properties” on page 33
- “ULConnection.GetDatabaseProperty method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseSchema.GetDatabaseProperty method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseSchema.getDatabaseProperty method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “ConfigPersistent.getCreationString method [Android] [UltraLiteJ]” [*UltraLite - Java Programming*]
- “ConfigPersistent.setCreationString method [Android] [UltraLiteJ]” [*UltraLite - Java Programming*]

UltraLite character sets

The results of comparisons on strings, and the sort order of strings, in part depends on the character set, collation, and encoding properties of the database.

Choosing the correct character set, collation, and encoding properties for your database is primarily determined by:

- The sort order you require. You should choose the collation that best sorts the characters you intend to store in your database.
- The platform of your device. Requirements among supported devices can vary, and some require that you use UTF-8 to encode your characters. If you need to support multiple devices, you need to determine whether a database can be shared.
- If you are synchronizing data, which languages and character sets are supported by the consolidated database. You must ensure that the character sets used in the UltraLite database and the consolidated database are compatible. Otherwise, data could be lost or become altered in unexpected ways if characters in one database's character set do not exist in the other's character set. If you have deployed UltraLite in a multilingual environment, you should also use UTF-8 to encode your UltraLite database.

When you synchronize, the MobiLink server converts characters as follows:

1. The UltraLite database characters are converted to Unicode.
2. The Unicode characters are converted into the consolidated database's character set.

Character sets and collations for UltraLite Java edition

UltraLite Java edition databases are encoded as UTF-8 (Unicode). The collation is the default sort order for Java and is equivalent to the UTF8BIN collation supported by SQL Anywhere. During synchronization with a MobiLink server, UltraLite notifies MobiLink that it uses the UTF8 character set and collation.

See also

- [“UltraLite platform requirements for character set encoding” on page 24](#)
- [“UltraLite collation creation parameter” on page 143](#)
- [“UltraLite utf8_encoding creation parameter” on page 164](#)
- [“Character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite connection parameters” on page 165](#)
- [“Character set considerations” \[MobiLink - Server Administration\]](#)
- [“UltraLite case creation parameter” on page 141](#)
- [“UltraLite database security” on page 26](#)

UltraLite platform requirements for character set encoding

Each platform has specific character set and encoding requirements.

Windows desktop and Windows Mobile

When using a UTF-8 encoded database on Windows, you should pass wide characters to the database. If you use UTF-8 encoding on these platforms, UltraLite expects that non-wide string parameters are UTF-8 encoded, which is not a natural character set to use on Windows. The exception is for connection strings, where string parameters are expected to be in the active code page. However, by using wide characters, you can avoid this complication.

See also

- [“UltraLite utf8_encoding creation parameter” on page 164](#)
- [“Character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite connection parameters” on page 165](#)
- [“Character set considerations” \[MobiLink - Server Administration\]](#)
- [“UltraLite database security” on page 26](#)

UltraLite supported collations

The following table lists the supported CHAR collations in UltraLite. You can also generate the list by executing the following command:

```
ulinit -Z
```

UltraLite Java edition databases only supports the UTF8BIN collation.

Collation label	Description
1250LATIN2	Code Page 1250, Windows Latin 2, Central/Eastern European
1250POL	Code Page 1250, Windows Latin 2, Polish
1251CYR	Code Page 1251, Windows Cyrillic

Collation label	Description
1252LATIN1	Code Page 1252, Windows Latin 1, Western
1252NOR	Code Page 1252, Windows Latin 1, Norwegian
1252SPA	Code Page 1252, Windows Latin 1, Spanish
1252SWEFIN	Code Page 1252, Windows Latin 1, Swedish/Finnish
1253ELL	Code Page 1253, Windows Greek, ISO8859-7 with extensions
1254TRK	Code Page 1254, Windows Turkish, ISO8859-9 with extensions
1254TRKALT	Code Page 1254, Windows Turkish, ISO8859-9 with extensions, I-dot e als I-no-dot
1255HEB	Code Page 1255, Windows Hebrew, ISO8859-8 with extensions
1256ARA	Code Page 1256, Windows Arabic, ISO8859-6 with extensions
1257LIT	Code Page 1257, Windows Baltic Rim, Lithuanian
874THAIBIN	Code Page 874, Windows Thai, ISO8859-11, binary ordering
932JPN	Code Page 932, Japanese Shift-JIS with Microsoft extensions
936ZHO	Code Page 936, Simplified Chinese, PRC GBK
949KOR	Code Page 949, Korean KS C 5601-1987 Encoding, Wansung
950ZHO_HK	Code Page 950, Traditional Chinese, Big 5 Encoding with HKSCS
950ZHO_TW	Code Page 950, Traditional Chinese, Big 5 Encoding
EUC_CHINA	Simplified Chinese, GB 2312-80 Encoding
EUC_JAPAN	Japanese EUC JIS X 0208-1990 and JIS X 0212-1990 Encoding
EUC_KOREA	Code Page 1361, Korean KS C 5601-1992 8-bit Encoding, Johab
EUC_TAIWAN	Code Page 964, EUC-TW Encoding
ISO1LATIN1	ISO8859-1, ISO Latin 1, Western, Latin 1 Ordering
ISO9LATIN1	ISO8859-15, ISO Latin 9, Western, Latin 1 Ordering
ISO_1	ISO8859-1, ISO Latin 1, Western
ISO_BINENG	Binary ordering, English ISO/ASCII 7-bit letter case mappings

Collation label	Description
UTF8BIN	UTF-8, 8-bit multibyte encoding for Unicode, binary ordering

UltraLite database security

You have the option to encrypt or obfuscate your UltraLite databases. Encryption provides secure representation of the data in the database whereas obfuscation only prevents casual observation of the contents of the database.

By default, UltraLite databases are not encrypted or obfuscated. When using a viewing tool such as a hex editor, text and binary columns can be read. Consider the following options if you do not want your data stored as plain text:

- **Obfuscation** Also known as simple encryption, this option provides protection against casual attempts to access data in the database. It does not provide as much security as strong encryption. Obfuscation has a minimal performance impact. You enable obfuscation with the obfuscate creation parameter. End users do not need to supply a corresponding connection parameter. You do not need any special configuration to use simple obfuscation on your device. See [“UltraLite obfuscate creation parameter” on page 152](#).
- **AES 256-bit encryption** UltraLite databases can be strongly encrypted using the AES 256-bit algorithm. Strong encryption provides security against skilled and determined attempts to gain access to the data, but has a significant performance impact. You set encryption in the Sybase Central wizards by clicking the **Encrypt the Database** option and then clicking **Use Strong Encryption and AES**. Using a command line utility, you set the encryption key with the DBKEY connection parameter. This parameter value is used by end users when connecting to the database. You do not need any special configuration to use AES encryption on your device. See [“UltraLite DBKEY connection parameter” on page 176](#).
- **AES FIPS 140-2 certified encryption (Windows and Windows Mobile only)** UltraLite provides encryption libraries certified with the FIPS 140-2 US and Canadian government standard (using a Certicom certified cryptographic module). You set FIPS-certified encryption with the fips creation parameter. The user must supply the required key in their connection string. AES FIPS-certified encryption requires that you configure your device appropriately. See [“UltraLite fips creation parameter” on page 148](#).

Note

Both the FIPS-certified and AES database encryption types use 256-bit AES.

Consider the effects of database cache size when choosing to encrypt or obfuscate UltraLite databases. There is an increase in overhead of between 5-10% as a result, resulting in decreased performance. The exact effect on performance depends on the size of your cache. If your cache is too small, encryption can add significant overhead. However, if your cache is sufficiently large, you may not see any difference at all. For more information about cache sizes, see [“Adjusting the cache size for an UltraLite database” on page 455](#).

Obfuscating an UltraLite database

To obfuscate data in an UltraLite database, specify `obfuscate=1` as a database creation parameter when you create your database. For more information, see [“UltraLite database creation” on page 17](#), and [“Specify UltraLite creation parameters” on page 22](#).

To obfuscate data in an UltraLite database using the UltraLiteJ API, you can use the `enableObfuscation` method during database creation. See [“ConfigPersistent.enableObfuscation method \[Android\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Encrypting an UltraLite database

To create an UltraLite database with strong encryption, set an encryption key when creating the database using the `ulinit` utility, or use the `DBKEY` creation parameter when creating the database using an UltraLite API. For more information, see:

- [“UltraLite DBKEY connection parameter” on page 176](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“Specify UltraLite creation parameters” on page 22](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“ConfigPersistent.setCreationString method \[Android\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

The encryption key should contain a combination of characters, numbers, and special symbols to be effective. Using a long encryption key reduces the chances of someone guessing the key.

In the UltraLiteJ API, you can specify the database key using the `enableAesDBEncryption` and `setDatabaseKey` methods as an alternative to specifying the `DBKEY` creation parameter. For more information, see [“ConfigPersistent.enableAesDBEncryption method \[Android\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#), and [“ConfigPersistent.setDatabaseKey method \[Android\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

After the database is encrypted, the encryption key cannot be recovered. In the UltraLiteJ API, you can get the key set by the `setDatabaseKey` method by using the `getDatabaseKey` method. For more information, see [“ConfigPersistent.getDatabaseKey method \[Android\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Connections to the database must specify the correct encryption key; otherwise, the connections fail. Encryption keys should be treated as sensitive information.

Caution

You can change the encryption key after the database has been created, but only under extreme caution. See:

- [“ULConnection.ChangeEncryptionKey method \[UltraLite C++\]”](#) [*UltraLite - C and C++ Programming*]
- [“ULConnection.ChangeEncryptionKey method \[UltraLite.NET\]”](#) [*UltraLite - .NET Programming*]
- [“Connection.changeEncryptionKey method \[UltraLite for M-Business Anywhere\]”](#) [*UltraLite - M-Business Anywhere Programming (deprecated)*]

This operation is costly and is non-recoverable: if your operation terminates mid-course, you lose your database entirely.

For strongly encrypted databases, store a copy of the key in a safe location. If you lose the encryption key there is no way to access the data, even with the assistance of technical support. The database must be discarded and you must create a new database.

Encrypting an UltraLite Java edition database

To create an UltraLite Java edition database with strong encryption, you must implement your own custom encryption control in your Java application, and use the `setEncryption` method of the `UltraLiteJ` API before creating or connecting to the database. For more information, see [“Encrypting data in an UltraLite Java edition database”](#) [*UltraLite - Java Programming*].

See also

- [“UltraLite fips creation parameter”](#) on page 148
- [“UltraLite obfuscate creation parameter”](#) on page 152
- [“UltraLite DBKEY connection parameter”](#) on page 176
- [“Specify UltraLite creation parameters”](#) on page 22
- [“Adjusting the cache size for an UltraLite database”](#) on page 455

UltraLite database connections

Applications that use a database must establish a connection to that database before transactions can occur. An application can be an UltraLite utility, a connection window, or your own custom application.

By connecting to an UltraLite database, you form a channel through which all activity from the application takes place. Each connection attempt creates a database specific SQL transaction.

See also

- [“UltraLite connection parameters” on page 165](#)
- UltraLite C/C++: [“Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- UltraLite.NET: [“Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- UltraLiteJ: [“Creating or connecting to an UltraLite or UltraLite Java edition database” \[UltraLite - Java Programming\]](#)

UltraLite connection strings and parameters

A **connection string** is a set of connection parameters that is passed from an application so that a database connection can be defined and established. Some parameters are always required to open a connection while others are used to adjust database features for a single connection.

Connection strings are defined as *keyword=value* pairs in a semicolon delimited list. The following example illustrates a connection string fragment that specifies a database file name, user ID, and password:

```
DBF=myULdb.udb;UID=JDoe;PWD=token
```

Methods of supplying these parameters to a database can vary depending on whether you are connecting from an UltraLite utility or an UltraLite application.

UltraLite command line utilities typically use a connection string if a connection to a database is required.

UltraLite applications can be developed using an UltraLite API to read connection parameter values from a stored file or in the application code. You can supply fixed connection strings when user authentication is not required, or prompt users to supply parameter values at connection time.

When a connection string has been assembled, it is passed to the UltraLite runtime for processing. The connection to the database is granted when the connection attempt is validated. Connection failures can occur if the database file does not exist, or the authentication was not successful.

UltraLite generates an error when it encounters an unrecognized connection parameter.

Precedence of connection parameters for UltraLite administration tools

UltraLite administration tools follow a specific order of connection parameter precedence:

- OS-specific options take precedence over nonspecific options. For example: CE_DBF takes precedence over DBF on Windows CE devices.
- If specified, the CE_FILE, desktop, device, and NT_FILE parameters always take precedence over DBF.
- If you supply duplicate parameters in a connection string, the last one supplied is used. All others are ignored.
- Parameters in the connection string take precedence over those supplied in the ULSQLCONNECT environment variable or a connection object.
- If no value is supplied for *both* UID and PWD in either the connection string or ULSQLCONNECT, the defaults of **UID=DBA** and **PWD=sql** are assumed.

Limitations

Any leading and/or trailing spaces in connection string parameter values are ignored. Connection parameter values cannot include leading single quotes ('), leading double quotes ("), or semicolons (;).

See also

- [“UltraLite connection parameters” on page 165](#)
- UltraLite C/C++: [“Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- UltraLite.NET: [“Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- UltraLiteJ: [“Creating or connecting to an UltraLite or UltraLite Java edition database” \[UltraLite - Java Programming\]](#)

Example

You can type a connection string on a single line with the parameter names and values separated by semicolons:

```
parameter1=value1;parameter2=value2
```

The UltraLite runtime ensures that the parameters are assembled into a connection string before establishing a connection with it. For example, if you use the ulload utility, the following connection string is used to load new XML data into an existing database. You cannot connect to the database file until you supply this string:

```
ulload -c "DBF=sample.udb;UID=DBA;PWD=sql" sample.xml
```

UltraLite connection parameters and the ULSQLCONNECT environment variable

Use the ULSQLCONNECT environment variable to avoid having to supply the same connection parameters repeatedly to Interactive SQL during development. Interactive SQL is the only tool that supports the ULSQLCONNECT environment variable; you cannot use the variable for custom applications.

The ULSQLCONNECT environment variable is optional, and is not specified by default. You can set the ULSQLCONNECT environment variable to contain a list of parameters defined as *keyword=value* pairs in a semicolon delimited list.

The supplied values become defaults for the UltraLite desktop administration tools. If Interactive SQL requires any additional parameters or if you need to override default values set with this environment variable, ensure you set these values. User supplied values always take precedence over this environment variable.

Caution

Do not use the pound character (#) as an alternative to the equal sign; the pound character is ignored by dbisql. All platforms supported by UltraLite allow you to use = inside an environment variable setting.

See also

- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)

Example

In this example, you use the ULSQLCONNECT environment variable to connect to a file named *c:\database\myfile.udb* and authenticate the user **demo** with the password **test**.

Set the following variable in your ULSQLCONNECT environment variable:

```
set ULSQLCONNECT="DBF=c:\database\myfile.udb;UID=demo;PWD=test"
```

By setting this environment variable, you no longer need to use the `dbisql -c` connection option for these defaults values—unless you need to override these values.

UltraLite file path formats in connection parameters

The physical storage of your device determines whether the database is saved as a file and what naming conventions you must follow when identifying your database.

Note

Use absolute file paths when using the UltraLite engine to support multi-process access to a database since the engine may be started in different locations.

The DBF parameter is most appropriate when targeting a single deployment platform or when using UltraLite desktop administration tools. For example:

```
ulload -c DBF=sample.udb sample.xml
```

Tip

You can use the UltraLite administration tools to administer databases already deployed to an attached device. See [“Windows Mobile” on page 32](#).

Otherwise, if you are writing a cross-platform application, use the platform-specific (CE_DBF or NT_DBF) file connection parameters to construct a universal connection string. For example:

```
Connection = DatabaseManager::OpenConnection("UID=JDoe;PWD=ULdb;CE_DBF=\database\MyCEDB.udb;NT_FILE=MyDB.udb")
```

Desktop

Desktops allow either absolute or relative paths.

Windows Mobile

Windows Mobile devices require absolute paths.

You can administer a Windows Mobile database on either the desktop or the attached device. To administer a database on a Windows Mobile device, ensure you prefix the absolute path with **wce:**. For example, using the `ulunload` utility:

```
ulunload -c DBF=wce:\UltraLite\myULdb.udb c:\out\ce.xml
```

In this example, UltraLite unloads the database from the Windows Mobile device to the `ce.xml` file in the Windows desktop folder of `c:\out`.

If you are using the `ulunloadold` or `ulunload` utilities to administer a database on the Windows Mobile device directly, UltraLite cannot back up the database before the unload or action occurs. You must perform this action manually before running these utilities.

See also

- [“UltraLite DBF connection parameter” on page 174](#)
- [“UltraLite NT_FILE connection parameter” on page 181](#)
- [“UltraLite CE_FILE connection parameter” on page 171](#)

UltraLite database tasks and features

Reading database properties

UltraLite provides a set of database properties that you can retrieve for a database.

You can change the settings of any database property that does not correspond to a database creation parameter.

Access UltraLite database properties in Sybase Central

You can access database properties with Sybase Central.

Browse UltraLite database properties (Sybase Central)

1. Connect to the database.
2. Right-click the database and click **Properties**.

In the **Database Properties** window, database properties are listed on the **General** and **Synchronization Information** tabs. On the **Synchronization Information** tab, the database properties are listed alphabetically by the property name. To sort database properties by the value, click the **Value** column.

Access UltraLite database properties using an UltraLite API

Each UltraLite API contains a `GetDatabaseProperty` method that you can use in your applications to access database properties.

Access UltraLite database properties with SQL

You can access database properties by calling the `DB_PROPERTY` SQL function.

See also

- “`ULConnection.GetDatabaseProperty` method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “`ULConnection.GetDatabasePropertyInt` method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “`ULDatabaseSchema.GetDatabaseProperty` method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “`DatabaseSchema.getDatabaseProperty` method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “`Connection.getDatabaseProperty` method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “`DB_PROPERTY` function [System]” on page 334

Accessing database options

Options are used to configure database behavior. Database options can be set or modified at any time. Temporary database options only persist while the database is running.

Option values are set by using the SET OPTION SQL statement. For example, the following statement sets the `global_database_id` option to 100:

```
SET OPTION global_database_id=100;
```

Browse or modify persistent UltraLite database options (Sybase Central)

You can view and change the setting of persistent database options from Sybase Central. Temporary UltraLite database options *cannot* be viewed or set from Sybase Central.

1. Connect to the database.
2. Right-click the database and click **Options**.
3. If you want to set or reset an option, type a new value in the **Value** field.
4. Click **Set Now** or **Reset Now** to commit the change.

See also

- [“UltraLite database options” on page 192](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.SetDatabaseOptionInt method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseSchema.setDatabaseOption method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Reading database properties” on page 33](#)

UltraLite event notifications

UltraLite now supports events and notifications. A notification is a message that is sent when an event occurs, also providing additional parameter information. UltraLite has system events and events can also be user-defined.

Event notifications allow you to provide coordination and signaling between connections or applications connected to the same database. Notifications are managed in queues: either a connection's default queue or, optionally, queues that are explicitly created and named. When an event occurs, notifications are sent to registered queues (or connections).

Each connection manages its own notification queues. Named queues can be created for any connection.

Using predefined system events this feature also provides "triggers" for changes to data, such as when a change is made to a table, for example, or signaling when a synchronization has occurred. Predefined events include:

- Commit
- SyncComplete
- TableModified

User events may also be defined and triggered by an application.

APIs for events and notifications are provided in each supported language. Additionally, a SQL function is provided to access the API functionality.

Events

Event	Occurrence
Commit	Signaled upon completion of a commit.
SyncComplete	Signaled upon completion of a sync.
TableModified	<p>Triggered when rows in a table are inserted, updated, or deleted. One event is signaled per request, no matter how many rows were affected by the request when registering for the event.</p> <p>The <i>object_name</i> parameter specifies the table to monitor. A value of "*" means all tables in the database.</p> <p>The <i>table_name</i> notification parameter is the name of the modified table.</p>

```

note_info.event_name = "SyncComplete";
note_info.event_name_len = 12;
note_info.parms_type = ul_ev_note_info::P_NONE;

note_info.event_name = "TableModified";
note_info.event_name_len = 13;
note_info.parms_type = ul_ev_note_info::P_TABLE_NAME;
note_info.parms = table->name->data;
note_info.parms_len = table->name->len;

```

Working with queues

Queues can be created and destroyed.

CreateNotificationQueue creates an event notification queue for the current connection. Queue names are scoped per-connection, so different connections can create queues with the same name. When an event notification is sent, all queues in the database with a matching name receive a separate instance of the notification. Names are case insensitive. A default queue is created on demand for each connection if no queue is specified. This call fails with an error if the name already exists for the connection or isn't valid.

DestroyNotificationQueue destroys the given event notification queue. A warning is signaled if unread notifications remain in the queue. Unread notifications are discarded. A connection's default event queue, if created, is destroyed when the connection is closed.

Working with events

DeclareEvent declares an event which can then be registered for and triggered. UltraLite predefines some system events triggered by operations on the database or the environment. The event name must be unique and names are case insensitive. Returns true if the event was declared successfully, false if the name is already used or invalid.

RegisterForEvent registers a queue to receive notifications of an event. If no queue name is supplied, the default connection queue is implied, and created if required. Certain system events allow specification of an object name to which the event applies. For example, the TableModified event can specify the table name. Unlike SendNotification, only the specific queue registered will receive notifications of the event; other queues with the same name on different connections will not (unless they are also explicitly registered). Returns true if the registration succeeded, false if the queue or event does not exist.

TriggerEvent triggers an event and sends a notification to all registered queues. Returns the number of event notifications sent. Parameters may be supplied as name=value; pairs.

Working with notifications

SendNotification sends a notification to all queues in the database matching the given name (including any such queue on the current connection). This call does not block. Use the special queue name "*" to send to all queues. Returns the number of notifications sent (the number of matching queues). Parameters may be supplied as name=value; pairs.

GetNotification reads an event notification. This call blocks until a notification is received or until the given wait period expires. To cancel a wait, send another notification to the given queue or use CancelGetNotification. After reading a notification, use ReadNotificationParameter to retrieve additional parameters. Returns true if an event was read, false if the wait period expired or was canceled.

GetNotificationParameter gets a named parameter for the event notification just read by GetNotification. Only the parameters from the most-recently read notification on the given queue are available. Returns true if the parameter was found, false if the parameter was not found.

CancelGetNotification cancels any pending GetNotification calls on all queues matching the given name. Returns the number of affected queues (not necessarily the number of blocked reads).

Other considerations

- Notification queue and event names are limited to 32 characters.
- To govern system resources, the number of notifications is limited. When this limit is exceeded, SQLE_EVENT_NOTIFICATION_QUEUE_FULL is signaled and the pending notification is discarded.

UltraLite support for spatial data

This section introduces UltraLite spatial support and explains its purpose, describes the supported data types, and explains how to generate and analyze spatial data.

The spatial data documentation assumes you already have some familiarity with spatial reference systems and with the spatial data you intend to work with.

Spatial data is data that describes the position, shape, and orientation of objects in a defined space. UltraLite provides storage and data management features for spatial data, in the form of points, allowing you to store information such as geographic locations and routing information, for instance. Points are defined using a **spatial type**, `ST_Geometry`. You use functions and constructors to access and manipulate the spatial data. UltraLite also provides a set of SQL spatial functions designed for compatibility with other products.

A point defines a single location in space. A point geometry does not have length or area. A point always has an X and Y coordinate.

In GIS data, points are typically used to represent locations such as addresses, or geographic features such as a mountain.

See also

- “Recommended reading on spatial topics” [*SQL Anywhere Server - Spatial Data Support*]
- “`ST_Geometry` type” on page 38

Compliance with spatial standards

UltraLite spatial support complies with the following standards:

- **International Organization for Standardization (ISO)** UltraLite geometries conform to the ISO standards for defining spatial user-types, routines, schemas, and for processing spatial data. UltraLite conforms to the specific recommendations made by the International Standard ISO/IEC 13249-3:2006. See http://www.iso.org/iso/catalogue_detail.htm?csnumber=38651.
- **Open Geospatial Consortium (OGC) Geometry Model** UltraLite geometries conform to the OGC OpenGIS Implementation Specification for Geographic information - Simple feature access - Part 2: SQL option version 1.2.0 (OGC 06-104r3). See <http://www.opengeospatial.org/standards/sfs>.

UltraLite uses the standards recommended by the OGC to ensure that spatial information can be shared between different vendors and applications.

To ensure compatibility with UltraLite spatial geometries, it is recommended that you adhere to the standards specified by the OGC.

- **SQL Multimedia (SQL/MM)** UltraLite follows the SQL/MM standard, and uses the prefix `ST_` for all function names.

SQL/MM is an international standard that defines how to store, retrieve, and process spatial data using SQL. Spatial type hierarchies such as ST_Geometry are one of the functions used to retrieve spatial data. The ST_Geometry hierarchy includes a number of subtypes such as ST_Point, ST_Curve, and ST_Polygon. With the SQL/MM standard, every spatial value included in a query must be defined in the same spatial reference system.

ST_Geometry type

The ST_Geometry type is used to store spatial data in the form of points.

Casts

Geometry objects can be explicitly and implicitly cast to any character or binary type. In the case of a character type, the geometry will be presented in EWKT format. In the case of a binary type, the geometry will be presented in WKB format. Character and binary types can also be explicitly and implicitly cast as a geometry value. For binary values, the value must represent a valid geometry in WKB format. For character values, the value must represent a valid geometry in either WKT or EWKT format. If casting from WKB or WKT to a geometry value, the default value SRID of 0 is assigned to the object.

Column and object definitions

UltraLite provides a fixed set of three different reference systems that can be attributed to a column during its creation. Individual geometry objects can be associated with any SRID value except the undefined reference system, and can only be stored in a column associated with a matching SRID value or the undefined reference system.

The predefined reference systems are:

- **Undefined or "null" reference system** This is the default reference system if no SRID value is provided. It allows contained geometry values to be in any valid reference system. This allows for "catch-all" columns that do not enforce any reference system consistency among their geometry objects.
- **Default planar reference system** Defined by specifying a SRID value of 0 during column creation, this column can contain only geometry values associated with this reference system. The values are treated as being in 2D planar space.
- **WGS 84 Geodetic Reference System** Defined by specifying a SRID value of 4326 during column creation, this column can only contain geometry values associated with this reference system. The values are treated as being on the Earth's surface and operations are applied accordingly.

Note

A point in SRID 4326 can be stored in a column with the WGS 84 reference system or with the undefined reference system, but not in the default planar system.

No transformations between reference systems is supported.

Example

```
create table T1 (  
  V1 Integer primary key,  
  V2 ST_Geometry(SRID=0),  
  V3 ST_Geometry);
```

creates a table with one column associated with the default planar reference system and one with an undefined reference system.

Note

ST_Geometry columns cannot be primary keys.

Functions for spatial data

UltraLite supports the following spatial functions:

- ST_AsBinary
- ST_AsExtText
- ST_AsText
- ST_Distance
- ST_Equals
- ST_IntersectsRect
- ST_Point function
- ST_PointFromExtText
- ST_PointFromText
- ST_PointFromWKB
- ST_SRID function
- ST_X function
- ST_Y function

See also

- [“ST_AsBinary function \[Spatial\]”](#)
- [“ST_AsExtText function \[Spatial\]”](#)
- [“ST_AsText function \[Spatial\]”](#)
- [“ST_Distance function \[Spatial\]”](#)
- [“ST_Equals function \[Spatial\]”](#)
- [“ST_IntersectsRect function \[Spatial\]”](#)
- [“ST_Point function \[Spatial\]”](#)
- [“ST_PointFromExtText function \[Spatial\]”](#)
- [“ST_PointFromText function \[Spatial\]”](#)
- [“ST_PointFromWKB function \[Spatial\]”](#)
- [“ST_SRID function \[Spatial\]”](#)
- [“ST_X function \[Spatial\]”](#)
- [“ST_Y function \[Spatial\]”](#)

Change the isolation level

Isolation levels define the degree to which the operations in one transaction are visible to the operations in other concurrent transactions. UltraLite uses the default isolation level, `read_committed`, for connections. The default UltraLite isolation level aids data consistency by isolating uncommitted rows.

This isolation level...	Has these characteristics...
0—read uncommitted	<ul style="list-style-type: none"> • Allow dirty reads, non-repeatable reads, and phantom rows. • No guarantee that concurrent transactions will not modify row or roll back changes to rows
1—read committed	<ul style="list-style-type: none"> • Allow non-repeatable reads and phantom rows • Prevent dirty reads • No guarantee that query results will not change during transactions

You can change the isolation level from `read_committed` to `read_uncommitted` using one of following approaches:

- Use the `SET OPTION SQL` statement and the `isolation_level` database option.
- For the UltraLite C++ API, use the `ULConnection.SetDatabaseOption` method to change the isolation level.

For the UltraLite.NET API, use the `ULConnection.BeginTransaction` or `ULDatabaseSchema.SetDatabaseOption` methods to create a transaction with the `read_committed` isolation level.

For the UltraLiteJ API, use the `Connection.setOption` method.

Note

UltraLite Java edition databases only support the `read_uncommitted` isolation level. See [“Data synchronization on a BlackBerry smartphone” \[UltraLite - Java Programming\]](#).

Concurrency and locking for UltraLite Java edition databases

- **Locking** Two different connections cannot modify the same row concurrently. If two connections attempt to operate on the same row, the second connection receives an error and is prohibited from modifying the row until the first connection commits or rolls back its current transaction.
- **Visibility** One connection's operation on the database is immediately visible to other connections.

See also

- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“UltraLite isolation_level option” on page 196](#)
- [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.BeginTransaction method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

Characteristics of the read_uncommitted isolation level

The following side effects are possible when UltraLite operates at an isolation_level of 0 (read_uncommitted):

- Applications can read uncommitted data (dirty reads). In this scenario, transactions may access rows in the database that are not committed and may still get rolled back by another transaction. This phenomena can result in phantom rows (rows that get added after the original query, making the result set returned in a repeated, duplicate query different).

For a tutorial that demonstrates the effects of dirty reads, see [“Tutorial: Understanding dirty reads” \[SQL Anywhere Server - SQL Usage\]](#).

For a tutorial that demonstrates a phantom row, see [“Tutorial: Understanding phantom rows” \[SQL Anywhere Server - SQL Usage\]](#).

- Applications can perform non-repeatable reads. In this scenario, an application reads a row from the database, and then goes on to perform other operations. Then a second application updates/deletes the row and commits the change. If the first application attempts to re-read the original row, it receives either the updated information or discovers that the original row was deleted.

For a tutorial that demonstrates the effects of non-repeatable reads, see [“Tutorial: Understanding non-repeatable reads” \[SQL Anywhere Server - SQL Usage\]](#).

See also

- [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.BeginTransaction method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“UltraLite isolation_level option” on page 196](#)

Example

Consider two connections, A and B, each with their own transactions.

1. As connection A works with the result set of a query, UltraLite **fetches** a copy of the current row into a buffer.

Note

Reading or fetching a row does not lock the row. If connection A fetches but does not modify a row, connection B can still modify the row.

2. As A modifies the current row, it changes the copy in the buffer. The copy in the buffer is written back into the database when connection A calls an Update method or closes the result set.
3. A write lock is placed on the row to prevent other transactions from modifying it. This modification is uncommitted, until connection A performs a commit.
4. Depending on the modification, if connection B fetches the current row, it may experience the following:

Connection A's modification	Result ¹
Row has been deleted.	Connection B gets the next row in the result set.
Row has been modified.	Connection B gets the latest copy of the row.

¹ Queries used by Connection A and B do not contain temporary tables. Temporary tables can cause other side effects.

Validating an UltraLite database

Caution

Database validation should be performed while no connections are making changes to the database; otherwise, errors indicating database corruption might be reported even though no corruption actually exists.

You can validate an UltraLite database using any of the following methods:

- The **Validate Database Wizard** in Sybase Central.
- The ulvalid command line utility.
- The ValidateDatabase method in your UltraLite API. (if available)

UltraLite Java edition database validation is not supported.

Validate an UltraLite database (Sybase Central)

1. In the left pane of Sybase Central, click the UltraLite database.
2. Click **File » Validate Database**.
3. Follow the instructions in the **Validate Database Wizard**.

See also

- [“ULConnection.ValidateDatabase method \[UltraLite C++\]”](#) [*UltraLite - C and C++ Programming*]
- [“ULConnection.ValidateDatabase method \[UltraLite.NET\]”](#) [*UltraLite - .NET Programming*]
- [“UltraLite Validate Database utility \(ulvalid\)”](#) on page 230

UltraLite and UltraLite Java edition database back up and recovery

If an application using an UltraLite or UltraLite Java edition database stops unexpectedly, the database automatically recovers to a consistent state when the application is restarted. All committed transactions flushed to memory before the unexpected failure are present in the database. All transactions not flushed at the time of the failure are rolled back.

An UltraLite database does not use a transaction log to perform recovery. Instead, UltraLite stores state information for every row to determine the fate of a row when recovering.

Backups

UltraLite provides protection against system failures, but not from media failures. The best way to make a backup of an UltraLite application is to synchronize with a consolidated database. To restore an UltraLite or UltraLite Java edition database, start with an empty database and populate it from the consolidated database through synchronization.

See also

- [“UltraLite database row state management”](#) on page 469
- [“Flush single or grouped transactions”](#) on page 471
- [“UltraLite as a MobiLink client”](#) on page 67

UltraLite database schemas

The logical framework of the database is known as a **schema**.

UltraLite database schemas

You can upgrade the schema of an UltraLite database with the appropriate Data Definition Language (DDL) statements or by using the ALTER DATABASE SCHEMA FROM FILE statement to modify the schema definition using a SQL script.

Schema changes can take a considerable amount of time. For example, all rows in the associated table must be updated when the column type is changed. DDL statements successfully execute when there are not any:

- Uncommitted transactions.
- Other active uses of the database, such as synchronization, prepared but unreleased statements, or executing database operations.

When the DDL statement is executing, any other attempt to use the database is blocked until the DDL statement completes the schema change.

UltraLite Java edition database schemas

The UltraLite Java edition database schema is maintained as a catalog of system tables that hold the metadata for the UltraLite Java edition database. System table metadata include:

- **Table definitions** Stored in the systable system table.
- **Column definitions** Stored in the syscolumn system table.
- **Index definitions** Stored in the sysindex and sysindexcolumn system tables.
- **Publication definitions** Stored in the syspublications and sysarticles system tables.
- **Foreign key definitions** Stored in the sysforeignkey and sysfkcol system tables.
- **User names and passwords** Stored in the sysuldata system tables.

See also

- [“ALTER DATABASE SCHEMA FROM FILE statement \[UltraLite\]” on page 412](#)
- [“Deploying UltraLite database schema upgrades” on page 123](#)
- [“systable system table” on page 251](#)
- [“syscolumn system table” on page 252](#)
- [“sysindex system table” on page 253](#)
- [“sysindexcolumn system table” on page 254](#)
- [“syspublications system table” on page 255](#)
- [“sysarticles system table” on page 255](#)
- [“sysforeignkey system table” on page 256](#)
- [“sysfkcol system table” on page 256](#)
- [“sysuldata system table” on page 257](#)

UltraLite tables and columns

Tables are used to store data and define the relationships for data in them. Tables consist of rows and columns. Each column carries a particular kind of information, such as a phone number or a name, while each row specifies a particular entry.

When you first create an UltraLite database, the only tables you will see are the system tables. System tables hold the UltraLite schema. You can hide or show these tables from Sybase Central as needed.

You can then add new tables as required by your application. You can also browse data in those tables, and copy and paste data among existing tables in the source database or even among other open destination databases.

Row packing and table definitions

UltraLite works with rows in two formats:

- **Unpacked rows** are the uncompressed format. Each row must be unpacked before individual column values can be read or written.
- **Packed rows** are the compressed representation of the unpacked row, where each of the column values is compressed so that the entire row takes up as little memory as possible. The size of a packed row depends entirely on the values in each column: for example, two rows can belong to the same table, but can differ significantly in their packed size. Note also that LONG BINARY and LONG VARCHAR columns are stored separate from the packed row.

UltraLite has a limitation that a packed row must fit on a database page. Since LONG BINARY and LONG VARCHAR columns are not stored with the packed row, they can exceed the page size.

It is important to understand that table definitions describe the row *before* the UltraLite runtime packs the data. Because the size of a packed row depends on the values in each column, you cannot readily pre-determine from the table definition whether the packed row requirement is satisfied. For this reason, UltraLite allows you to define a table where an unpacked row would not fit on a page. To know if a row fits on a page, you must try inserting or updating the row itself; if a row does not fit, UltraLite detects and reports this error.

Note

You cannot declare tables to be any large size you require. UltraLite maintains a declared table row size limit of 64 KB. If you try to define a table where an unpacked row can exceed this maximum, UltraLite generates a SQL error code of SQLE_MAX_ROW_SIZE_EXCEEDED (-1132).

See also

- [“UltraLite page_size creation parameter” on page 153](#)
- [“Database tables” \[SQL Anywhere 12 - Introduction\]](#)
- [“SQL Anywhere database creation” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite system tables” on page 231](#)

Creating UltraLite tables

You can create new tables to hold your relational data, either with SQL statements in Interactive SQL or with Sybase Central.

In UltraLite, you can only create base tables, which you declare to hold persistent data. The table and its data continue to exist until you explicitly delete the data or drop the table. UltraLite does not support global temporary or declared temporary tables.

Note

Tables in UltraLite applications must include a primary key. Primary keys are also required during MobiLink synchronization, to associate rows in the UltraLite database with rows in the consolidated database.

Sybase Central

In Sybase Central, you can perform these tasks while working with a selected database.

Create an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, right-click **Tables** and click **New » Table**.
3. In the **What Do You Want To Name The New Table** field, type the new table name.
4. Click **Finish**.
5. From the **File** menu, click **Save**.

Interactive SQL

In Interactive SQL, you can declare columns while creating a new table.

Create an UltraLite table (Interactive SQL)

1. Connect to the UltraLite database.

2. Execute a CREATE TABLE statement.

For example, the following statement creates a new table to describe the various skills and professional aptitudes of employees within a company. The table has columns to hold an identifying number, a name, and a type (for example, technical, or administrative) for each skill.

```
CREATE TABLE Skills (  
    SkillID INTEGER PRIMARY KEY,  
    SkillName CHAR( 20 ) NOT NULL,  
    SkillType CHAR( 20 ) NOT NULL  
);
```

See also

- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)
- [“Adding a column to an UltraLite table” on page 48](#)

Adding a column to an UltraLite table

You can add a new column easily if the table is empty. However, if the table already holds data, you can only add a column if the column definition includes a default value or allows NULL values.

You can use either Sybase Central or execute a SQL statement (for example, Interactive SQL) to perform this task.

Sybase Central

In Sybase Central, you can perform this task while working with a selected table.

Add a new column to an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, double-click **Tables**.
3. Double-click a table.
4. Click the **Columns** tab, right-click the white space below the table and click **New » Column**.
5. Set the attributes for the new column.
6. From the **File** menu, click **Save Table**.

Interactive SQL

In Interactive SQL, you can only declare columns while creating or altering a table.

Add columns to a new UltraLite table (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute a CREATE TABLE statement or ALTER TABLE, and define the columns by declaring the name, and other attributes.

The following example creates a table for a library database to hold information about borrowed books. The default value for `date_borrowed` indicates that the book is borrowed on the day the entry is made. The `date_returned` column is NULL until the book is returned.

```
CREATE TABLE borrowed_book (  
  loaner_name CHAR(100) PRIMARY KEY,  
  date_borrowed DATE NOT NULL DEFAULT CURRENT DATE,  
  date_returned DATE,  
  book CHAR(20)  
);
```

The following example modifies the customer table to now include a column for addresses that can hold up to 50 characters:

```
ALTER TABLE customer  
ADD address CHAR(50);
```

See also

- “Object name considerations” [[SQL Anywhere Server - Database Administration](#)]
- “Data types in UltraLite” on page 265
- “Column data type considerations” [[SQL Anywhere Server - Database Administration](#)]
- “CREATE TABLE statement [UltraLite]” on page 425
- “ALTER TABLE statement [UltraLite]” on page 415

Altering UltraLite column definitions

You can change the structure of column definitions for a table by altering various column attributes, or even deleting columns entirely. The modified column definition must suit the requirements of any data already stored in the column. For example, you cannot alter a column to disallow NULL if the column already has a NULL entry.

You can use either Sybase Central or Interactive SQL to perform this task.

Sybase Central

In Sybase Central, you can perform these tasks while working with a selected table.

Alter an existing UltraLite column (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, double-click **Tables**.
3. Double-click a table.
4. Click the **Columns** tab and alter the column attributes.
5. From the **File** menu, click **Save Table**.

Interactive SQL

In Interactive SQL, you can perform these tasks with the ALTER TABLE statement.

Alter an existing UltraLite column (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute an ALTER TABLE statement.

The following examples show how to change the structure of the database. In all these cases, the statement is committed immediately. So, once you make the change, any item referring to this table may no longer work.

The following statement shortens the SkillDescription column from a maximum of 254 characters to a maximum of 80:

```
ALTER TABLE Skills
MODIFY SkillDescription CHAR( 80 );
```

The following statement deletes the Classification column:

```
ALTER TABLE Skills
DROP Classification;
```

The following statement changes the name of the entire table:

```
ALTER TABLE Skills
RENAME Qualification;
```

See also

- “Object name considerations” [[SQL Anywhere Server - Database Administration](#)]
- “Data types in UltraLite” on page 265
- “Column data type considerations” [[SQL Anywhere Server - Database Administration](#)]
- “ALTER TABLE statement [UltraLite]” on page 415

Deleting UltraLite tables

You can drop any table if the table:

- Is not being used as an article in a publication.
- Does not have any columns that are referenced by another table's foreign key.

In these cases, you must change the publication or delete the foreign key *before* you can successfully delete the table.

You can use either Sybase Central or Interactive SQL to perform this task.

Sybase Central

In Sybase Central, you can perform these tasks while working with a selected table.

Delete an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.

2. In the left pane, double-click **Tables**.
3. Right-click a table and click **Delete**.
4. Click **Yes**.

Interactive SQL

In Interactive SQL, deleting a table is also called dropping it. You can drop a table by executing a DROP TABLE statement.

Delete an UltraLite table (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute a DROP TABLE statement.

For example, the following DROP TABLE statement deletes all the records in the Skills table and then removes the definition of the Skills table from the database:

```
DROP TABLE Skills;
```

Like the CREATE statement, the DROP statement automatically executes a COMMIT statement before and after dropping the table. This statement makes all changes to the database since the last COMMIT or ROLLBACK permanent. The DROP statement also drops all indexes on the table.

See also

- [“DROP TABLE statement \[UltraLite\]” on page 435](#)

Browsing the information in UltraLite tables

You can use Sybase Central or Interactive SQL to browse the data held within the tables of an UltraLite database. Tables can be user tables or system tables. You can filter tables by showing and hiding system tables from your current view of the database. Because UltraLite does not have a concept of ownership, all users can browse all tables.

Sybase Central

In Sybase Central, you can perform these tasks while working with a selected database.

Browse an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. To view a list of tables, double-click **Tables**.
3. To view table data, double-click a table and click the **Data** tab in the right pane.

Interactive SQL

In Interactive SQL, you can perform these tasks with the SELECT statement.

Browse UltraLite user tables (Interactive SQL)

1. Connect to a database.
2. Execute a SELECT statement, specifying the user table you want to browse.

Browse UltraLite system tables (Interactive SQL)

1. Connect to a database.
2. Execute a SELECT statement, by the system table you want to browse.

For example, to display the contents of the table systable on the **Results** tab in the **Results** pane in Interactive SQL, execute the following command:

```
SELECT * FROM SYSTABLE;
```

See also

- [“UltraLite system tables” on page 231](#)

Data copying and pasting to or from UltraLite databases

With Sybase Central you can copy and paste and drag and drop. This data transferral allows you to share or move objects among one or more databases. By copying and pasting or dragging and dropping you can share data as described by the table that follows.

Target	Result
Another UltraLite or SQL Anywhere database.	A new object is created, and the original object's code is copied to the new object.
The same UltraLite database.	A copy of the object is created; you must rename the new object.

Note

You can copy data from a database opened in MobiLink and paste it into an UltraLite database. However, you cannot paste UltraLite data into a database opened in MobiLink.

Sybase Central

When you copy any of the following objects in the UltraLite plug-in, the SQL for the object is also copied to the clipboard. You can paste this SQL into other applications, such as Interactive SQL or a text editor.

For example, if you copy an index in Sybase Central and paste it into a text editor, the CREATE INDEX statement for that index appears. You can copy the following objects in the UltraLite plug-in:

- Articles
- Columns
- Foreign keys
- Indexes
- Publications
- Tables
- Unique constraints

Interactive SQL

With Interactive SQL you can also copy data from a result set into another object.

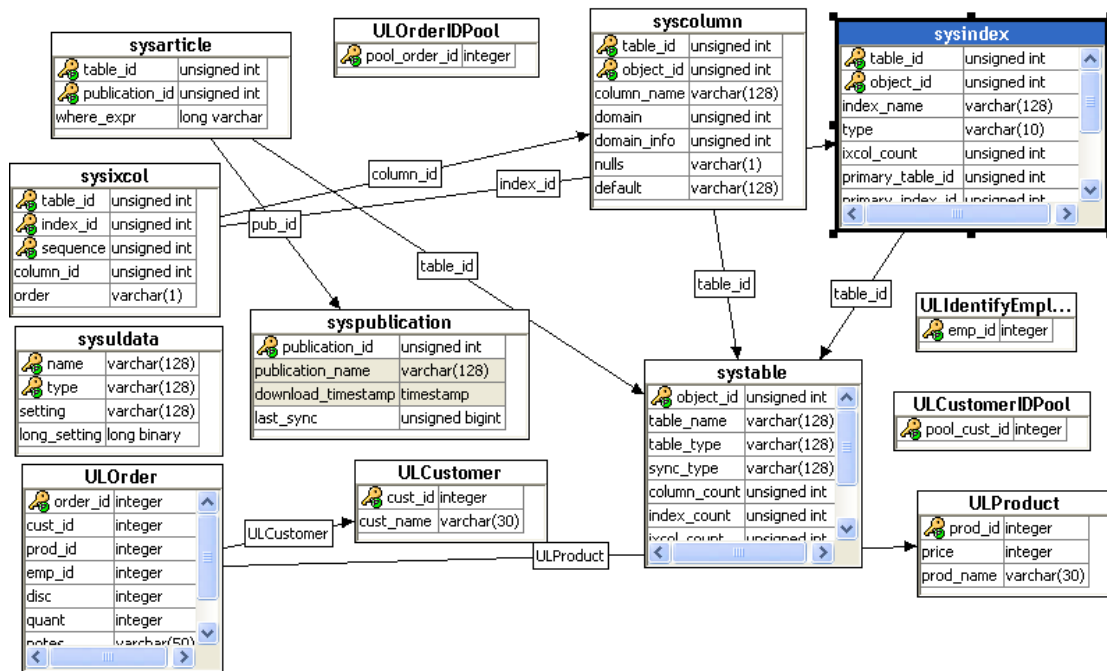
- Use the SELECT statement results into a named object.
- Use the INSERT statement to insert a row or selection of rows from elsewhere in the database into a table.

See also

- [“Database objects in the SQL Anywhere 12 plug-in” \[SQL Anywhere Server - Database Administration\]](#)
- [“INSERT statement \[UltraLite\]” on page 438](#)
- [“SELECT statement \[UltraLite\]” on page 444](#)

Entity-relationship diagrams from the UltraLite plug-in

When you are connected to a database from the UltraLite plug-in, you can view an entity-relationship diagram of the tables in the database. When you have the database selected, click the **ER Diagram** tab in the right pane to see the diagram.



When you rearrange objects in the diagram, the changes persist between Sybase Central sessions. Double-clicking a table takes you to the column definitions for that table.

See also

- “SQL Anywhere database creation” [[SQL Anywhere Server - Database Administration](#)]

UltraLite indexes

An index is a set of pointers to rows in a table based on the order of the values of data in one or more table columns. The index is a database object that is maintained automatically by UltraLite after it has been created. When UltraLite optimizes a query, it scans existing indexes to see if one exists for the table(s) named in the query. If it can help UltraLite return rows more quickly, the index is used. If you are using the UltraLite Table API in your application, you can specify an index that helps determine the order in which rows are traversed.

Tip

Indexes can improve the performance of a query—especially for large tables. To see whether a query is using a particular index, you can check the execution plan with Interactive SQL.

Alternatively, your UltraLite applications can include PreparedStatement objects which have a method to return plans.

UltraLite supports the following indexes. These indexes can be single or multi-column (also known as composite indexes). You cannot index LONG VARCHAR or LONG BINARY columns.

Index	Characteristics
Primary key	Required. An instance of a unique key. You can only have one primary key. Values in the indexed column or columns must be unique and cannot be NULL.
Foreign key ¹	Optional. Values in the indexed column or columns can be duplicated. Nullability depends on whether the column was created to allow NULL. Values in the foreign key columns must exist in the table being referenced
Unique key ²	Optional. Values in the indexed column or columns must be unique and cannot be NULL.
Non-unique index	Optional. Values in the indexed column or columns can be duplicated and can be NULL.
Unique index	Optional. Values in the indexed column or columns cannot be duplicated and can be NULL.

¹ A foreign key can reference either a primary key or a unique key.

² Also known as a unique constraint.

About composite indexes

Multi-column indexes are sometimes called composite indexes. Additional columns in an index can allow you to narrow down your search, but having a two-column index is not the same as having two separate indexes. For example, the following statement creates a two-column composite index:

```
CREATE INDEX name
ON Employees ( Surname, GivenName );
```

A composite index is useful if the first column alone does not provide high selectivity. For example, a composite index on Surname and GivenName is useful when many employees have the same surname. A composite index on EmployeeID and Surname would not be useful because each employee has a unique ID, so the column Surname does not provide any additional selectivity.

See also

- “Use index scans” on page 457
- “Execution plans in UltraLite” on page 462
- “Composite indexes” [*SQL Anywhere Server - SQL Usage*]
- UltraLite.NET: “Data creation and modification using the ULTable class” [*UltraLite - .NET Programming*]
- “ULCommand.Prepare method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- UltraLite for C++: “Data creation and modification using the ULTable class” [*UltraLite - C and C++ Programming*]
- “ULPreparedStatement class [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “Data creation and modification using the Table API” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “PreparedStatement class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

When to use an index

Use an index when:

- **You want UltraLite to maintain referential integrity** An index also affords UltraLite a means of enforcing a uniqueness constraint on the rows in a table. You do not need to add an index for data that is very similar.
- **The performance of a particular query is important to your application** If an index improves performance of a query and the performance of that query is important to your application and is used frequently, then you want to maintain that index. Unless the table in question is extremely small, indexes can improve search performance dramatically. Indexes are typically recommended whenever you search data frequently.
- **You have complicated queries** More complicated queries, (for example, those with JOIN, GROUP BY, and ORDER BY clauses), can yield substantial improvements when an index is used—though it may be harder to determine the degree to which performance has been enhanced. Therefore, test your queries both with and without indexes, to see which yields better performance.
- **The size of an UltraLite table is large** The average time to find a row increases with the size of the table. Therefore, to increase searchability in a very large table, consider using an index. An index allows UltraLite to find rows quickly—but only for columns that are indexed. Otherwise, UltraLite must search every row in the table to see if the row matches the search condition, which can be time consuming in a large table.
- **The UltraLite client application is not performing a large amount of insert, update, or delete operations** Because UltraLite maintains indexes along with the data itself, an index in this context will have an adverse effect on the performance of database operations. For this reason, you should restrict the use of indexes to data that will be queried regularly as described in the point above. Maintaining the UltraLite default indexes (indexes for primary keys and for unique constraints) may be enough.

- **Use indexes on columns involved in WHERE clauses and/or ORDER BY clause** These indexes can speed the evaluation of these clauses. In particular, an index helps optimize a multi-column ORDER BY clause—but only when the placement of columns in the index and ORDER BY clauses are exactly the same.

Index types

UltraLite supports different types of indexes: unique keys, unique indexes, and non-unique indexes. What differentiates one from the others is what is allowed in that index.

Index characteristic	Unique keys	Unique indexes	Non-unique indexes
Allows duplicate index entries for rows that have the same values in indexed columns.	no	no	yes
Allows null values in index columns.	no	yes	yes

Note

You can create foreign keys to unique keys, but not to unique indexes.

Also, manually creating an index on a key column is not necessary and generally not recommended. UltraLite creates and maintains indexes for unique keys automatically.

See also

- [“Adding an UltraLite index” on page 57](#)

Adding an UltraLite index

You can use either Sybase Central or Interactive SQL to perform this task.

Note

UltraLite does not detect duplicate or redundant indexes. As indexes must be maintained with the data in your database, add your indexes carefully.

Sybase Central

In Sybase Central, you can perform this task while working with a selected database.

Create a new index for a given UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click **Indexes**, and click **New » Index**.

3. Follow the instructions in the wizard.

Interactive SQL

In Interactive SQL, you can perform this task with the CREATE INDEX statement.

Create a new index for a given UltraLite table (Interactive SQL)

1. Connect to an UltraLite database.
2. Execute a CREATE INDEX statement.

This statement creates an index with the default maximum hash size you have configured. To create an index that overrides the default, ensure you use the WITH MAX HASH SIZE *value* clause to set a new value for this index instance.

For example, to speed up a search on employee surnames in a database that tracks employee information, and tune the performance of queries against this index, you could create an index called EmployeeNames and increase the hash size to 20 bytes with the following statement:

```
CREATE INDEX EmployeeNames
ON Employees (Surname, GivenName)
WITH MAX HASH SIZE 20;
```

See also

- [“CREATE INDEX statement \[UltraLite\]” on page 421](#)

Dropping an UltraLite index

Dropping an index deletes it from the database.

You can use either Sybase Central or Interactive SQL to perform this task.

Sybase Central

In Sybase Central, you can perform this task while working with a selected database.

Drop an UltraLite index (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, double-click **Indexes**.
3. Right-click an index and then click **Delete**.
4. Click **Yes**.

Interactive SQL

In Interactive SQL, deleting an index is also called dropping it. You can perform this task with the DROP INDEX statement.

Drop an UltraLite index (Interactive SQL)

1. Connect to a database.
2. Execute a DROP INDEX statement.

For example, the following statement removes the EmployeeNames index from the database:

```
DROP INDEX EmployeeNames;
```

See also

- [“DROP INDEX statement \[UltraLite\]” on page 433](#)

UltraLite users

A typical UltraLite database contains one user ID and password. UltraLite databases are created with a default user ID of **DBA** and default password of **sql** unless otherwise specified. UltraLite Java edition databases are created with a default password of **DBA**.

Changing the user schema is optional and **not** required. Many applications do not need database-level authentication and assume that a device level password is sufficient authentication to access an application and its data.

Common reasons for not authenticating users may be because the deployment is to a single-user device, or that it is too awkward to prompt a user each time they start the application.

You do not need to include a user ID or password in the database connection string if you do not need database-level authentication. The simplest UltraLite connection string is **DBF=filename**. The remainder of this section explains how UltraLite implements user IDs and describes how to use them for cases where you need explicit user authentication.

When developing an UltraLite application with a custom user authentication interface, you can effectively use the UltraLite user IDs and password hashes stored in an UltraLite database to validate user-supplied credentials and avoid creating your own password hashing algorithm. By adding users to your UltraLite database, you store their user IDs and password hashes. You can then validate the user-supplied credentials in your application by attempting to connect to the database with the UID and PWD connection parameters, where **UID=username** and **PWD=password**. A successful UltraLite database connection indicates that the user is authentic.

Caution

Unlike SQL Anywhere users, UltraLite users are created and managed solely for authentication and not for object ownership or specific database permissions. Once users are authenticated, they gain full access to the database.

By creating user IDs and passwords, you control connections to the UltraLite database but do not secure the data in the database file. The contents are stored as plain text and can be read directly.

To secure the database contents, it is recommended that you encrypt the file. When you encrypt the file, you can authenticate users with an encryption key rather than a user ID and password.

You can obfuscate the file to alter the storage so that data is not stored as plain text, but this approach does not secure the data.

For more information, see [“UltraLite database security” on page 26](#) and [“UltraLite DBKEY connection parameter” on page 176](#).

Note

UltraLite user IDs are different from MobiLink user names.

Limitations

User IDs and passwords are not stored in UltraLite Java edition databases. Those databases are protected with a single password and do not require user creation or authentication. For more information, see [“Creating or connecting to an UltraLite or UltraLite Java edition database” \[UltraLite - Java Programming\]](#).

The following limitations apply to UltraLite user IDs:

- UltraLite supports up to four unique user IDs per UltraLite database.
- User IDs and passwords can be changed using Sybase Central, SQL statements, or UltraLite API methods in your application.
- User IDs have a 31-character limit.
- User IDs cannot include leading single quotes('), leading double quotes ("), or semicolons(;).
- User IDs are always case insensitive and passwords are always case sensitive.
- User IDs cannot be renamed. You can only add new user IDs and delete existing ones from an existing database connection.
- Users can not be listed programmatically using the UltraLite APIs. You can only use database tools to list existing users in the database.
- When connecting to an UltraLite database for the first time, the **UID** and **PWD** are the same values that were set when the database was created. UltraLite attempts to connect with the **DBA** user ID and **sql** password when these connection parameters are not specified. You do not need to supply a username or

password when connecting to the database if you did not explicitly set a username and password during its creation.

See also

- [“UltraLite UID connection parameter” on page 187](#)
- [“UltraLite PWD connection parameter” on page 183](#)

Manage UltraLite users with connection parameters

You can use the **UID** and **PWD** connection parameters to create or authenticate users in an UltraLite database.

Note

As an alternative to connection parameters, you can use the following UltraLite API methods in your application to grant or revoke user access to an UltraLite database:

- [“ULConnection.GrantConnectTo method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.RevokeConnectFrom method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULGrantConnectTo method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULRevokeConnectFrom method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.GrantConnectTo method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.RevokeConnectFrom method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Connection.grantConnectTo method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.revokeConnectFrom method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)

Grant and revoke methods are not available in the UltraLiteJ API.

Using connection parameters with UltraLite databases

For most UltraLite APIs, the `createDatabase` method of a `DatabaseManager` object can be used to create a new UltraLite database with the specified connection and creation parameters.

The following example illustrates how to create a default user for a new UltraLite database by passing the **UID** and **PWD** parameters to the `CreateDatabase` method in the UltraLite C++ API:

```
ULConnection * conn;
ULError ulerr;
ULDatabaseManager::CreateDatabase("dbf=sample.udb;uid=default-
name;pwd=default-password", &ulerr);
```

The following example illustrates how to authenticate a user in an existing UltraLite database by passing the **UID** and **PWD** parameters to the `OpenConnection` method in the UltraLite C++ API:

```
ULConnection * conn;
ULError ulerr;
```

```
ULDatabaseManager::OpenConnection("dbf=sample.udb;uid=test-name;pwd=test-  
password", &ulerr);
```

Using connection parameters with UltraLite databases on Android smartphones

For the UltraLiteJ API, you use the `setConnectionString` method of a Configuration object in the UltraLiteJ API to create or authenticate users.

The following example illustrates how to create a default user for a new UltraLite database by passing the **UID** and **PWD** parameters to the `createDatabase` method in the UltraLiteJ API:

```
ConfigFile config =  
    DatabaseManager.createConfigurationFileAndroid("DBname.udb",  
    getApplicationContext());  
config.setConnectionString("uid=default-name;pwd=default-password");  
Connection conn = DatabaseManager.createDatabase(config);
```

The following example illustrates how to authenticate a user in an existing UltraLite database by passing the **UID** and **PWD** parameters and the `connect` method in the UltraLiteJ API:

```
ConfigFile config =  
    DatabaseManager.createConfigurationFileAndroid("DBname.udb",  
    getApplicationContext());  
config.setConnectionString("uid=test-name;pwd=test-password");  
Connection conn = DatabaseManager.connect(config);
```

As an alternative to the `setConnectionString` method, you can use the `setPassword` or `setUserName` methods to create or authenticate a user, respectively.

Creating or authenticating users on an UltraLite Java edition database

UltraLite Java edition databases do not support user authentication. Instead, you set a database password using the `Configuration.setPassword` method.

The following example illustrates how to set the password for an existing UltraLite Java edition database:

```
ConfigNonPersistent config =  
    DatabaseManager.createConfigurationNonPersistent("DBname.ulj");  
config.setPassword("my_password");  
Connection conn = DatabaseManager.connect(config);
```

See also

- [“ConfigPersistent.setUserName method \[Android\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Configuration.setPassword method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“UltraLite UID connection parameter”](#) on page 187
- [“UltraLite PWD connection parameter”](#) on page 183
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]”](#) [*UltraLite - C and C++ Programming*]
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]”](#) [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]”](#) [*UltraLite - .NET Programming*]
- [“ConfigPersistent.setConnectionString method \[Android\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“DatabaseManager.createDatabase method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

Manage UltraLite users with SQL statements

You can use the CREATE USER, ALTER USER, and DELETE USER statements to manage users in an UltraLite database.

You cannot use these SQL statements with UltraLite Java edition databases.

Note

As an alternative to SQL statements, you can use the following UltraLite API methods in your application to grant or revoke user access to an UltraLite database:

- “ULConnection.GrantConnectTo method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULConnection.RevokeConnectFrom method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULGrantConnectTo method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULRevokeConnectFrom method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULConnection.GrantConnectTo method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ULConnection.RevokeConnectFrom method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “Connection.grantConnectTo method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Connection.revokeConnectFrom method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Grant and revoke methods are not available in the UltraLiteJ API.

Example

The following example illustrates how to connect to an existing UltraLite database using the UltraLiteJ API, and use the CREATE USER statement to create a new user:

```
ConfigFile config =
    DatabaseManager.createConfigurationFileAndroid("DBname.udb",
getApplicationContext());
Connection conn = DatabaseManager.connect(config);

String sql_string = "CREATE USER test-user IDENTIFIED BY test-password";
PreparedStatement authenticator = conn.prepareStatement(sql_string);
authenticator.execute();
authenticator.close();
```

See also

- [“CREATE USER statement \[UltraLite\]” on page 431](#)
- [“ALTER USER statement \[UltraLite\]” on page 419](#)
- [“DROP USER statement \[UltraLite\]” on page 436](#)
- [“ULPreparedStatement class \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.PrepareStatement method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“PreparedStatement class \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.prepareStatement method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“ULCommand class \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULCommand.Prepare method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“PreparedStatement interface \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.prepareStatement method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

Managing UltraLite users with Sybase Central

You can use Sybase Central to explicitly create or delete users from an UltraLite database.

Sybase Central is not compatible with UltraLite Java edition databases.

Note

As an alternative to Sybase Central, you can use the following UltraLite API methods in your application to grant or revoke user access to an UltraLite database:

- [“ULConnection.GrantConnectTo method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.RevokeConnectFrom method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULGrantConnectTo method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULRevokeConnectFrom method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.GrantConnectTo method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.RevokeConnectFrom method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Connection.grantConnectTo method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.revokeConnectFrom method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)

Grant and revoke methods are not available in the UltraLiteJ API.

Create a new UltraLite user in Sybase Central

1. Connect to the UltraLite database.
2. Right-click the **Users** folder, and click **New » User**.
3. Follow the instructions in the wizard.

Delete an existing UltraLite user in Sybase Central

1. Connect to the UltraLite database.
2. In the left pane, double-click the **Users** folder.
3. Right-click the user and click **Delete**.

UltraLite as a MobiLink client

This section contains material that describes how to set up and run UltraLite clients for MobiLink synchronization.

UltraLite clients

Synchronizing an UltraLite database requires your application to set synchronization parameters identifying the address of the MobiLink server and other required information, and calling a synchronization function or executing the SYNCHRONIZE SQL statement. The option you chose depends on the API you are using.

UltraLite client synchronization behavior customization

Adding custom synchronization support to UltraLite can involve up to three tasks:

- **Maintain primary key uniqueness in synchronization models that include more than one remote client** Required. In a synchronization system, the primary key is the only way to identify the same row in different databases (remote and consolidated) and the only way to detect conflicts. Therefore, multiple clients must adhere to the following rules:
 - Every table that is to be synchronized must have a primary key.
 - Never update the values of primary keys.
 - Primary keys must be unique across all synchronized databases.
- **Ensure your date columns are set up so that fractional data is not lost** For a SQL Anywhere consolidated database this is not typically an issue. However, for databases like Oracle, there may be compatibility issues that you need to consider. For example, UltraLite and Oracle databases must share the same timestamp precision. Additionally, you should also add a `TIMESTAMP` to the Oracle database to avoid losing fractional second data when the UltraLite remote databases uploads data to the consolidated database.
- **Describe what data subsets you want to upload to the consolidated database** Optional. You only need to do this when you do not want to synchronize all data by default. To target what data you want to synchronize, use one or more subsetting techniques.

For example, you may want to create a publication for high-priority data. The application could then synchronize this data over wireless networks. Because wireless networks can have high usage costs associated with them, you may want to limit these usage fees to those that are business critical. You can then synchronize less time-sensitive data from a cradle at a later time.

- **Initialize synchronization from your UltraLite application and supply the parameters that describe the session** Required. Programming synchronization has two parts: describing the session, and then initiating the synchronization operation.

Describing the session primarily involves choosing a synchronization communication stream (also known as a network protocol), and the parameters for that stream, setting the version of your synchronization scripts, and identifying the MobiLink user. However, there are other parameters you can set: for example, use the `upload_only` and `download_only` parameters to change the default bi-directional synchronization to one-way only.

All other important synchronization behaviors are controlled at the MobiLink server with MobiLink synchronization scripts. These include:

- What data is downloaded as updates or inserts to tables in the UltraLite remote.
- What processing is required on uploaded changes from a remote database.

This means that you can write your synchronization scripts so that data is partitioned among remote databases in an appropriate manner.

See also

- [“Unique primary keys”](#) [*MobiLink - Server Administration*]
- [“Primary key uniqueness in UltraLite”](#) on page 68
- [“Oracle consolidated database”](#) [*MobiLink - Server Administration*]
- [“UltraLite precision creation parameter”](#) on page 155
- [“UltraLite client synchronization design”](#) on page 72
- [“Adding synchronization to your UltraLite application”](#) on page 78
- [“MobiLink consolidated databases”](#) [*MobiLink - Server Administration*]
- [“Synchronization script writing”](#) [*MobiLink - Server Administration*]
- [“Direct row handling”](#) [*MobiLink - Server Administration*]
- [“Partitioning rows among remote databases”](#) [*MobiLink - Server Administration*]

Primary key uniqueness in UltraLite

UltraLite can maintain primary key uniqueness using any of the techniques supported by MobiLink.

One of these methods is to use a GLOBAL AUTOINCREMENT column. GLOBAL AUTOINCREMENT is similar to AUTOINCREMENT, except that the domain is partitioned. UltraLite supplies column values only from the partition assigned to the database's global database ID. Each UltraLite database is assigned a unique integer global database ID.

A second method is to use a UUID primary key column. A UUID requires more data, but needs no distinct database identifier.

See also

- [“Unique primary keys”](#) [*MobiLink - Server Administration*]
- [“UltraLite global_database_id option”](#) on page 195

Using GLOBAL AUTOINCREMENT in UltraLite

You can declare the default value of a column in an UltraLite database to be of type GLOBAL AUTOINCREMENT. However, before you can autoincrement these column IDs, you must first set the global database ID for the UltraLite database.

Caution

GLOBAL AUTOINCREMENT column values downloaded via MobiLink synchronization do not update the GLOBAL AUTOINCREMENT value counter. As a result, an error can occur should one MobiLink client insert a value into another client's partition. To avoid this problem, ensure that each copy of your UltraLite application inserts values only in its own partition.

Declare GLOBAL AUTOINCREMENT columns in your UltraLite database

1. Assign each copy of the database a unique global ID number.

The `global_database_id` database option sets the value in your UltraLite database. When deploying UltraLite, you must assign a different identification number to each database.

2. Allow UltraLite to supply default values for the column using the partition uniquely identified by the UltraLite database's number. UltraLite follows these rules:

- If the column contains no values in the current partition, the first default value is $pn + 1$. p represents the partition size and n represents the global ID number.
- If the column contains values in the current partition, but all are less than $p(n + 1)$, the next default value will be one greater than the previous maximum value in this range.
- Default column values are not affected by values in the column outside the current partition; that is, by numbers less than $pn + 1$ or greater than $p(n + 1)$. Such values may be present if they have been replicated from another database via MobiLink synchronization.

For example, if you assigned your UltraLite database a global ID of 1 and the partition size is 1000, then the default values in that database would be chosen in the range 1001-2000. Another copy of the database, assigned the identification number 2, would supply default values for the same column in the range 2001-3000.

- Because you cannot set the global ID number to negative values, the values UltraLite chooses for GLOBAL AUTOINCREMENT columns are always positive. The maximum identification number is restricted only by the column data type and the partition size.
 - If you do not set a global ID value, or if you exhaust values from the partition, a NULL value is inserted into the column. Should NULL values not be permitted, the attempt to insert the row causes an error.
3. If you exhaust or will soon exhaust available values for columns declared as GLOBAL AUTOINCREMENT, you need to set a new global database ID. UltraLite chooses GLOBAL AUTOINCREMENT values from the partition identified by the global ID number, but only until the maximum value is reached. If you exceed values, UltraLite begins to generate NULL values. By assigning a new global database ID number, you allow UltraLite to set appropriate values from another partition.

One method of choosing a new global database ID is to maintain a pool of unused global database ID values. This pool is maintained in the same manner as a pool of primary keys.

Tip

UltraLite APIs provide means of obtaining the proportion of numbers that have been used. The return value is a SHORT in the range 0-100 that represents the percent of values used so far. For example, a value of 99 indicates that very few unused values remain and the database should be assigned a new identification number. The method of setting this identification number varies according to the programming interface you are using.

See also

- “UltraLite global_database_id option” on page 195
- “Primary key pools” [*SQL Remote*]
- “Overriding partition sizes for autoincremented columns” on page 71
- “UltraLite global_database_id option” on page 195
- “ULCommand.Connection property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ULConnection.Synchronize method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “Connection.setDatabaseID method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “ULSetDatabaseID method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULConnection.GlobalAutoIncrementUsage property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “Connection.getGlobalAutoIncrementUsage method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “ULGlobalAutoincUsage method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “Connection.setDatabaseId method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “Connection.getDatabaseId method [BlackBerry] [UltraLiteJ]” [*UltraLite - Java Programming*]

Determine the most recently assigned GLOBAL AUTOINCREMENT value

You can retrieve the GLOBAL AUTOINCREMENT value that was chosen during the most recent insert operation. Since these values are often used for primary keys, knowing the generated value may let you more easily insert rows that reference the primary key of the first row. You can check the value with:

- **UltraLite for C/C++** Use the GetLastIdentity function on the ULConnection object.
- **UltraLite.NET** Use the LastIdentity property on the ULConnection class.
- **UltraLite for M-Business Anywhere** Use the GetLastIdentity method on the Connection class.
- **API UltraLiteJ** Use the getLastIdentity method on the Connection interface.

The returned value is an unsigned 64-bit integer, database data type UNSIGNED BIGINT. Since this statement only allows you to determine the most recently assigned default value, you should retrieve this value soon after executing the insert statement to avoid spurious results.

Note

Occasionally, a single INSERT statement may include more than one column of type GLOBAL AUTOINCREMENT. In this case, the return value is one of the generated default values, but there is no reliable means to determine which one. For this reason, you should design your database and write your INSERT statements in a way that avoids this situation.

See also

- “ULConnection.GetLastIdentity method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULConnection.LastIdentity property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “Connection.getLastIdentity method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Connection.getLastIdentity method [UltraLiteJ]” [*UltraLite - Java Programming*]

Overriding partition sizes for autoincremented columns

The partition size is any positive integer, although the partition size is generally chosen so that the supply of numbers within any one partition will rarely, if ever, be exhausted.

For columns of type INT or UNSIGNED INT, the default partition size is $2^{16} = 65536$; for columns of other types the default partition size is $2^{32} = 4294967296$. Since these defaults may be inappropriate, it is best to specify the partition size explicitly.

Default partition sizes for some data types are different in UltraLite applications than in SQL Anywhere databases. Declare the partition size explicitly if you want different databases to remain consistent.

Override UltraLite partition values (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the selected column and click **Properties**.
3. Click the **Value** tab.
4. Enter any positive integer in the **Partition Size** field.

Declare autoincrement columns in UltraLite (SQL)

1. Connect to the UltraLite database.
2. Execute a CREATE TABLE or ALTER TABLE statement with a DEFAULT GLOBAL AUTOINCREMENT clause with the partition size specified in parentheses.

For example, the following statement creates a simple reference table with two columns: an integer that holds a customer identification number and a character string that holds the customer's name. A partition size of 5000 is required for this table.

```
CREATE TABLE customer (
  id      INT          DEFAULT GLOBAL AUTOINCREMENT (5000),
```

```

        name VARCHAR(128) NOT NULL,
        PRIMARY KEY (id)
    );

```

See also

- “CREATE TABLE statement [UltraLite]” on page 425
- “ALTER TABLE statement [UltraLite]” on page 415
- “ULTableSchema.GetColumnPartitionSize method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ULTableSchema.GetGlobalAutoincPartitionSize method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “TableSchema.getColumnPartitionSize method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “ULGlobalAutoincUsage method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “Connection.getLastIdentity method [UltraLiteJ]” [*UltraLite - Java Programming*]

UltraLite client synchronization design

All data in an UltraLite database is synchronized by default. If you are new to deploying UltraLite as a MobiLink remote database, plan to use the default behavior initially.

Once you become comfortable with the synchronization process, you may decide to customize the behavior of the synchronization operation to capture more complex business logic. Designing custom synchronization behavior requires that you ask the following questions. If your business requirements are simple, you may only need to use a single synchronization feature. However, in very complex deployments, you may need to use multiple synchronization features to configure the synchronization behavior you require.

Design question	If you answer yes, use the following
Do you want to download changes from the consolidated database but not have local changes uploaded to the database?	The download_only table name suffix allows you to identify any tables for which the synchronization should be download only. Changes made to the local tables are not uploaded to the consolidated database.
Do you want to exclude tables from synchronization?	The nosync table name suffix allows you to identify any tables that you do not want to synchronize.
Do you only want to synchronize entire tables even when data hasn't changed?	The allsync table name suffix allows you to synchronize the entire table, even when no changes are detected.

Design question	If you answer yes, use the following
Do you want to synchronize an entire table or just rows that meet specific conditions? Does some of the data require synchronization priority due to its importance or time-sensitivity?	<p>A publication includes articles that list the tables that require synchronization. An article can include a WHERE clause that specifies the rows to upload based on whether the rows meet the defined criteria.</p> <p>Multiple publications can address priority issues that require certain UltraLite data be uploaded before others.</p>
Do you require a table order for synchronization because you have cycles of foreign keys?	The Table Order synchronization parameter allows you to determine the order of synchronization operations when you have foreign key cycles. However, foreign key cycles are generally not recommended for UltraLite.
Do you want to control synchronization behavior? For example, do you need downloads to occur at the same time as uploads? Or do you want to change bi-directional synchronization to one-way only?	<p>Use the appropriate synchronization parameter as part of:</p> <ul style="list-style-type: none"> • Your application's synchronization structure (or the synchronization enumeration). • The ulsync utility's -e option.
Do you want your UltraLite client to be TLS-enabled?	What encryption algorithm you choose determines how your device must be set up according to the platform that runs on that device.

See also

- [“UltraLite download-only tables” on page 74](#)
- [“UltraLite non-synchronizing tables” on page 73](#)
- [“UltraLite synchronize-all tables” on page 75](#)
- [“Publishing data in UltraLite” on page 76](#)
- [“Table order in UltraLite” on page 77](#)
- [“UltraLite synchronization parameters and network protocol options” on page 87](#)
- [“The synchronization process” \[MobiLink - Getting Started\]](#)
- [“The upload and the download” \[MobiLink - Getting Started\]](#)

UltraLite non-synchronizing tables

By creating the table using SYNCHRONIZE OFF, you control when to exclude the entire table from the upload operation. You can use these non-synchronizing tables for client-specific persistent data that is not required in the consolidated database. Other than being excluded from synchronization, you can use these tables in exactly the same way as other tables in the UltraLite database.

Note

The synchronization type for a table can only be changed if it does not have any unsynchronized changes that need to be uploaded.

If you create a table with a `_nosync` suffix, you can only rename that table so it retains the `_nosync` suffix. For example, the following `ALTER TABLE` statement with a `rename` clause is not allowed because the new name no longer ends in `nosync`:

```
ALTER TABLE purchase_comments_nosync
RENAME comments;
```

To correct this, the statement must be rewritten to include this suffix:

```
ALTER TABLE purchase_comments_nosync
RENAME comments_nosync;
```

You can alternatively use publications to achieve the same effect.

Note

As an alternative to creating or altering a table with the `SYNCHRONIZE OFF` clause, you can append the phrase `_nosync` to the table name to turn it into a non-synchronizing table.

See also

- [“Download Only synchronization parameter” on page 92](#)
- [“Upload Only synchronization parameter” on page 109](#)
- [“Additional Parameters synchronization parameter” on page 88](#)
- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)
- [“ALTER TABLE statement \[UltraLite\]” on page 415](#)
- [“Publications” \[MobiLink - Client Administration\]](#)
- [“UltraLite publications” on page 82](#)
- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)
- [“Synchronization script writing” \[MobiLink - Server Administration\]](#)

UltraLite download-only tables

By creating the table using `SYNCHRONIZE DOWNLOAD`, you exclude entire tables from the upload operation. You can use these tables for data that should not be synchronized to the consolidated database. Other than being excluded from synchronization, you can use these tables in exactly the same way as other tables in the UltraLite database.

You can alternatively use publications to achieve the same effect.

Note

The synchronization type for a table can only be changed if it does not have any unsynchronized changes that need to be uploaded.

Note

As an alternative to creating or altering a table with the SYNCHRONIZE DOWNLOAD clause, you can append the phrase **_download_only** to the table name to turn it into a download-only table.

See also

- “Download Only synchronization parameter” on page 92
- “Upload Only synchronization parameter” on page 109
- “Additional Parameters synchronization parameter” on page 88
- “CREATE TABLE statement [UltraLite]” on page 425
- “ALTER TABLE statement [UltraLite]” on page 415
- “Publications” [*MobiLink - Client Administration*]
- “UltraLite publications” on page 82
- “CREATE PUBLICATION statement [UltraLite]” on page 423
- “Synchronization script writing” [*MobiLink - Server Administration*]

UltraLite synchronize-all tables

By creating or altering a table using SYNCHRONIZE ALL you control whether to change the synchronization behavior during upload so that it synchronizes all table data, even if nothing has changed since the previous synchronization session.

Note

The synchronization type for a table can only be changed if it does not have any unsynchronized changes that need to be uploaded.

Some UltraLite applications require user/client-specific data that you can store in a SYNCHRONIZE ALL TABLES. You can upload the data in the table to a temporary table in the consolidated database, use the data to control synchronization by your other scripts without having the data maintained in the consolidated database. For example, you may want your UltraLite applications to indicate which channels or topics they are interested in, and use this information to download the appropriate rows.

Note

As an alternative to creating or altering a table with the SYNCHRONIZE ALL clause, you can append the phrase **_allsync** to the table name to turn it into a synchronize-all table.

See also

- “Download Only synchronization parameter” on page 92
- “Upload Only synchronization parameter” on page 109
- “Additional Parameters synchronization parameter” on page 88
- “CREATE TABLE statement [UltraLite]” on page 425
- “ALTER TABLE statement [UltraLite]” on page 415
- “Publications” [*MobiLink - Client Administration*]
- “UltraLite publications” on page 82
- “CREATE PUBLICATION statement [UltraLite]” on page 423
- “Synchronization script writing” [*MobiLink - Server Administration*]

Publishing data in UltraLite

Publications define a set of articles that describe the data to be synchronized. Each article can be a whole table, or can define a subset of the data in a table. You can include an optional predicate (a WHERE clause) if you want to define a subset of rows from a given table.

Publications are more flexible than creating tables with SYNCHRONIZE OFF. To synchronize data subsets of an UltraLite database *separately*, use multiple publications. You can then combine publications with upload-only or download-only synchronization parameters to synchronize high-priority changes efficiently.

Note

The maximum number of user publications in UltraLite is 63.

Adding publications

You can add publications to an UltraLite database with Sybase Central, or using SQL. For UltraLite synchronization, each article in a publication may include either a complete table, or may include a WHERE clause.

Note

UltraLite publications do not support the definition of column subsets, nor the SUBSCRIBE BY clause that is available in SQL Anywhere. If columns in an UltraLite table do not exactly match tables in a consolidated database, use MobiLink scripts to resolve those differences.

You do not need to set a table synchronization order in a publication. If table order is important for your deployment, you can set the table order when you synchronize the UltraLite database by setting the Table Order synchronization parameter.

Publish data from an UltraLite database (Sybase Central)

1. Connect to the UltraLite database using the UltraLite plug-in.
2. Right-click the **Publications** folder and click **New » Publication**.
3. Enter a name for the new publication. Click **Next**.

4. On the **Tables** tab, click a table in the **Matching Tables** list. Click **Add**.

The table appears in the **Selected Tables** list on the right.

5. Add additional tables.
6. If necessary, click the **Where** tab to specify the rows to be included in the publication. You cannot specify column subsets.
7. Click **Finish**.

Publish data from an UltraLite database (SQL)

1. Connect to the UltraLite database.
2. Execute a CREATE PUBLICATION statement that specifies the name of the new publication and the tables you want to publish.

See also

- [“UltraLite publications” on page 82](#)
- [“Download Only synchronization parameter” on page 92](#)
- [“Upload Only synchronization parameter” on page 109](#)
- [“Additional Parameters synchronization parameter” on page 88](#)
- [“Publications” \[MobiLink - Client Administration\]](#)
- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)
- [“Synchronization script writing” \[MobiLink - Server Administration\]](#)

Table order in UltraLite

By setting the Table Order synchronization parameter you can control the order of synchronization operations. If you want to specify a table order for synchronization, you can use the TableOrder parameter programmatically or as part of the ulsync utility during testing. The TableOrder parameter specifies the order of tables that are to be uploaded.

You only need to explicitly set the table order if your UltraLite database has:

- Foreign key cycles. You must then list all tables that are part of a cycle.
- Different foreign key relationships from those used in the consolidated database.

Avoiding synchronization issues with foreign key cycles

Table order is particularly important for UltraLite databases that use foreign key cycles. A cycle occurs when you link a series of tables together such that a circle is formed. However, due to complexities that arise when cycles between the consolidated database and the UltraLite remote database differ, foreign key cycles are not recommended.

With foreign key cycles, you should order your tables so that operations for a primary table come before the associated foreign table. A Table Order parameter ensures that the insert in the foreign table will have its foreign key referential integrity constraint satisfied (likewise for other operations like delete).

In addition to table ordering, another method you can use to avoid synchronization issues is to postpone the checking of referential integrity until the transaction is committed. If your consolidated database is a SQL Anywhere database, set one of the foreign keys to **check on commit**. This ensures that foreign key referential integrity is checked during the commit phase rather than when the operation is initiated. For example:

```
CREATE TABLE c (  
    id INTEGER NOT NULL PRIMARY KEY,  
    c_pk INTEGER NOT NULL  
);  
CREATE TABLE p (  
    pk INTEGER NOT NULL PRIMARY KEY,  
    c_id INTEGER NOT NULL,  
    FOREIGN KEY p_to_c (c_id) REFERENCES c(id)  
);  
ALTER TABLE c  
    ADD FOREIGN KEY c_to_p (c_pk)  
    REFERENCES p(pk)  
    CHECK ON COMMIT;
```

If your consolidated database is from another database vendor, check to see if the database has similar methods of checking referential integrity. If so, you should implement this method. Otherwise, you must redesign table relationships to eliminate all foreign key cycles.

See also

- [“Additional Parameters synchronization parameter”](#) on page 88
- [“Referential integrity and synchronization”](#) [*MobiLink - Getting Started*]

Adding synchronization to your UltraLite application

In UltraLite, synchronization begins by opening a specific connection with the MobiLink server over the configured communication stream (also known as a network protocol). In addition to synchronization support for direct network connections, Windows Mobile devices also support ActiveSync synchronization.

Defining the connection

Each UltraLite remote database that synchronizes with a MobiLink server does so over a network protocol. You set the network protocol with the synchronization stream parameter. Supported network protocols include TCP/IP, HTTP, HTTPS, and TLS. For the protocol you choose, you also need to supply stream parameters that define other required connection information like the MobiLink server host and the port. You must also supply the MobiLink user information and the synchronization script version.

Defining the synchronization behavior

You can control synchronization behavior by setting various synchronization parameters. The way you set parameters depends on the specific UltraLite interface you are using.

Important behaviors to consider include:

- **Synchronization direction** By default, synchronization is bi-directional. If you require one-way synchronizations only, remember to use the appropriate `upload_only` or `download_only` parameter. By performing one-way synchronizations, you minimize the synchronization time required. Also, with download-only synchronization, you do not have to commit all changes to the UltraLite database before synchronization. Uncommitted changes to tables not involved in synchronization are not uploaded, so incomplete transactions do not cause problems.

To use download-only synchronization, you must ensure that rows overlapping with the download are not changed locally. If any data is changed locally, synchronization fails in the UltraLite application with a `SQLE_DOWNLOAD_CONFLICT` error.

- **Concurrent changes during synchronization** During the upload phase, UltraLite applications can access UltraLite databases in a read-only fashion. During the download phase, read-write access is permitted, but if an application changes a row that the download then attempts to change, the download will fail and roll back. You can disable concurrent access to data during synchronization by setting the `DisableConcurrency` synchronization parameter.

Add synchronization code to your UltraLite application

1. Supply the necessary synchronization parameters and protocol options you require for the session as fields of a synchronization information structure.

For example, using the C/C++ API, you add synchronization to the UltraLite application by setting appropriate values in the `ul_sync_info` structure:

```
ul_sync_info info;
    // define a sync structure named "info"
    ULEnableTcpipSynchronization( &sqlca );
    // use a TCP/IP stream
    conn->InitSynchInfo( &info );
    // initialize the structure
    info.stream = ULStream();
    // specify the Socket Stream
    info.stream_parms= UL_TEXT( "host=myMLserver;port=2439" );
    // set the MobiLink host information
    info.version = UL_TEXT( "custdb 11.0" );
    // set the MobiLink version information
    info.user_name = UL_TEXT( "50" );
    // set the MobiLink user name
    info.download_only =ul_true;
    // make the synchronization download-only
```

2. Initialize synchronization.

For direct synchronization, you would call an API-specific synchronization function. These functions return a boolean indicating success or failure of the synchronization operation. If the synchronization fails, you can examine detailed error status fields in another structure to get additional error information.

For ActiveSync synchronization, you must catch the synchronization message from the ActiveSync provider and use the `DoSync` function to call `ULSynchronize`.

3. Use an observer callback function if you want to report the progress of the synchronization to the user.

Tip

If you have an environment where DLLs fail either because the DLL is very large or the network connection is unreliable, you may want to implement resumable downloads. See “Failed downloads” [[MobiLink - Server Administration](#)] and “Resumption of failed downloads” [[MobiLink - Server Administration](#)].

See also

- “ActiveSync with UltraLite on Windows Mobile” on page 85
- “Upload-only and download-only synchronizations” [[MobiLink - Server Administration](#)]
- “The upload and the download” [[MobiLink - Getting Started](#)]
- “UltraLite synchronization parameters and network protocol options” on page 87
- UltraLite.NET: “MobiLink data synchronization” [[UltraLite - .NET Programming](#)]
- UltraLite C/C++: “MobiLink data synchronization” [[UltraLite - C and C++ Programming](#)]
- UltraLite C/C++: “Add ActiveSync synchronization to your application” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “MobiLink data synchronization” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- UltraLite for embedded SQL: “Adding synchronization to your application” [[UltraLite - C and C++ Programming](#)]
- UltraLiteJ API: “MobiLink data synchronization” [[UltraLite - Java Programming](#)]

MobiLink file transfers

UltraLite supports the ability to transfer files with the MobiLink server. M-Business Anywhere does not need this functionality because it has its own mechanism for file deployments or transfers (called **channel synchronization**).

For all other APIs, use the MobiLink file transfer mechanism when:

- You have multiple files that you need to deploy to multiple devices, particularly when corporate firewalls are used as a security measure. Because MobiLink is already configured to handle synchronization through these firewalls, the MLFileTransfer mechanism makes device provisioning for upgrades and other types of file transfers very convenient.
- You have files that you want to target to a specific MobiLink user ID. This requires that you create one or more user-specific directories on the MobiLink server for each user ID you require. Otherwise, if you only have a single version of the file, you can use a default directory.

How file transfers work

You can employ one of two MobiLink-initiated file transfer mechanisms to download files to a device: run the mlfiletransfer utility for desktop transfers, or call the appropriate function for the API you are using to code your UltraLite application. Both approaches require that you:

1. Describe the transfer destination.

Whether you use the `mlfiletransfer` utility from the desktop, or whether you use the function appropriate to your API, you must set the local path and file name of the file on the target device or desktop computer. If none are supplied in the application or by the end user, then the source file name is assumed and the file is stored in the current directory.

The destination directory of the target can vary depending on the device's operating system:

- On Windows Mobile, if the destination is `NULL`, the file is stored in the root directory (`\`).
The file name must follow file name conventions for Windows Mobile. See [“Windows Mobile” on page 32](#).
- On the desktop, if the destination is `NULL`, the file is stored in the current directory.
The file name must follow file name conventions for the desktop system. See [“Desktop” on page 32](#).
- On iPhone, you should store files in your application's document directory. You can get the location of the document directory by calling the `NSSearchPathForDirectories/uDomains` using the `NSDocumentDirectory` parameter.
- On BlackBerry smartphones, set the location using the `FileTransfer` object.

2. Set the MobiLink user credentials that allow the user to be identified and the correct file(s) to be downloaded.

This user name and password are separate from any database user ID and password, and serve to identify and authenticate the application to the MobiLink server.

3. Set the stream type you want to use, and define the parameters for the stream you require. These are the same parameters supported by UltraLite for MobiLink synchronization.

Most synchronization streams require parameters to identify the MobiLink server address and control other behavior. If you set the stream type to a value that is invalid for the platform, the stream type is set to TCP/IP, except for UltraLite Java edition databases which supports only HTTP.

4. Describe the required behavior for the transfer mechanism.

For example, you can set properties that allow this mechanism to force a download even when the file already exists on the target and has not changed, or that allow partial downloads to be resumed. You can also set whether you want the progress to be monitored and reported upon.

5. Ensure the MobiLink server is running and has been started with the `-ftr` option.
6. Start the transfer, and, if applicable, monitor the download progress.

By displaying the download progress, the user can cancel and resume the download at a later time.

Note

UltraLite for M-Business: Not supported.

See also

- “UltraLite synchronization parameters and network protocol options” on page 87
- “-ftr mlsrv12 option” [*MobiLink - Server Administration*]
- “MobiLink File Transfer utility (mlfiletransfer)” [*MobiLink - Client Administration*]
- “MLFileDownload method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “MLFileUpload method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULFileTransfer class [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “FileTransfer interface [UltraLiteJ]” [*UltraLite - Java Programming*]

UltraLite publications

A publication is a database object that identifies the data that is to be synchronized. If you want to synchronize all tables and all rows of those tables in your UltraLite database, do not create any publications.

A publication consists of a set of articles. Each article may be an entire table, or may be rows in a table. You can define this set of rows with a **WHERE** clause.

Each database can have multiple publications, depending on the synchronization logic you require. For example, you may want to create a publication for high-priority data. The user can synchronize this data over high-speed wireless networks. Because wireless networks can have usage costs associated with them, you would want to limit these usage fees to those that are business-critical only. All other less time-sensitive data could be synchronized from a cradle at a later time.

You create publications using Sybase Central or with the **CREATE PUBLICATION** statement. In Sybase Central, all publications and articles appear in the **Publications** folder.

Usage notes

- UltraLite publications do not support the definition of column subsets, nor the **SUBSCRIBE BY** clause. If columns in an UltraLite table do not exactly match tables in a SQL Anywhere consolidated database, use MobiLink scripts to resolve those differences.
- Columns are always sent in the order in which they were defined in the **CREATE TABLE** statement.
- You do not need to set a table synchronization order in a publication. If table order is important for your deployment, you can set the table order when you synchronize the UltraLite database by setting the **Table Order** synchronization parameter.
- Because object ownership is not supported in UltraLite, any user can delete a publication.

See also

- “Table order in UltraLite” on page 77
- “Publications” [*MobiLink - Client Administration*]
- “UltraLite client synchronization design” on page 72
- “Synchronization script writing” [*MobiLink - Server Administration*]

Publishing whole tables in UltraLite

The simplest publication you can make consists of a single article, which consists of all rows and columns of a table.

You can use either Sybase Central or Interactive SQL to perform this task.

Sybase Central

In Sybase Central, you can perform this task while working with the connected database.

Publish one or more whole UltraLite tables (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the **Publications** folder, and click **New » Publication**.
3. In the **What Do You Want To Name The New Publication** field, type a name for the new publication. Click **Next**.
4. On the **Tables** tab, click tables in the **Available Tables** list. Click **Add**.
5. Click **Finish**.

Interactive SQL

In Interactive SQL, you can perform this task with the CREATE PUBLICATION statement.

Publish one or more whole UltraLite tables (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute a CREATE PUBLICATION statement that specifies the name of the new publication and the table you want to publish.

For example, the following statement creates a publication that publishes the whole customer table:

```
CREATE PUBLICATION pub_customer (  
    TABLE customer  
);
```

See also

- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)
- [“UltraLite clients” on page 67](#)

Publishing a subset of rows from an UltraLite table

A publication can only contain specific table rows. In Sybase Central or Interactive SQL, a WHERE clause limits the rows that are uploaded to those that have changed and satisfy a search condition in the WHERE clause.

To upload all changed rows, do not specify a WHERE clause.

Sybase Central

In Sybase Central, you can perform this task while working with the connected database.

Publish only some rows in an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the **Publications** folder, and click **New » Publication**.
3. In the **What Do You Want To Name The New Publication** field, type a name for the new publication.
4. Click **Next**.
5. In the **Available Tables** list, click a table and click **Add**.
6. Click the **WHERE Clauses** tab, and click the table from the **Articles** list. Optionally, you can use the Insert window to assist you in formatting the search condition.
7. Click **Finish**.

Interactive SQL

In Interactive SQL, you can perform this task with the CREATE PUBLICATION statement.

Create a publication in UltraLite using a WHERE clause (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute a CREATE PUBLICATION statement that includes the tables you want to include in the publication and a WHERE condition.

For example, the following example creates a single-article publication that includes all sales order information for sales rep number 856:

```
CREATE PUBLICATION pub_orders_samuel_singer
( TABLE SalesOrders
  WHERE SalesRepresentative = 856 );
```

See also

- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)
- [“UltraLite clients” on page 67](#)

Dropping a publication for UltraLite

You can drop a publication using either Sybase Central or Interactive SQL.

Sybase Central

In Sybase Central, you can perform this task while working with the connected database.

Drop a publication (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, double-click the **Publications** folder.
3. Right-click the publication and click **Delete**.
4. Click **Yes**.

Interactive SQL

In Interactive SQL, deleting a publication is also called dropping it. You can perform this task with the DROP PUBLICATION statement.

Drop a publication (Interactive SQL)

1. Connect to the UltraLite database.
2. Execute a DROP PUBLICATION statement.

For example, the following statement drops the publication named pub_orders:

```
DROP PUBLICATION pub_orders;
```

See also

- [“DROP PUBLICATION statement \[UltraLite\]” on page 434](#)
- [“UltraLite clients” on page 67](#)

ActiveSync with UltraLite on Windows Mobile

While you can synchronize data from a Windows Mobile device over an Ethernet or Wi-Fi connection, this section describes how to configure your desktop and device to use ActiveSync synchronization so that your UltraLite database is synchronized at the same time as other ActiveSync operations. If you want to synchronize directly using one of the other alternative methods, you need to program your application to do so using an appropriate synchronize function.

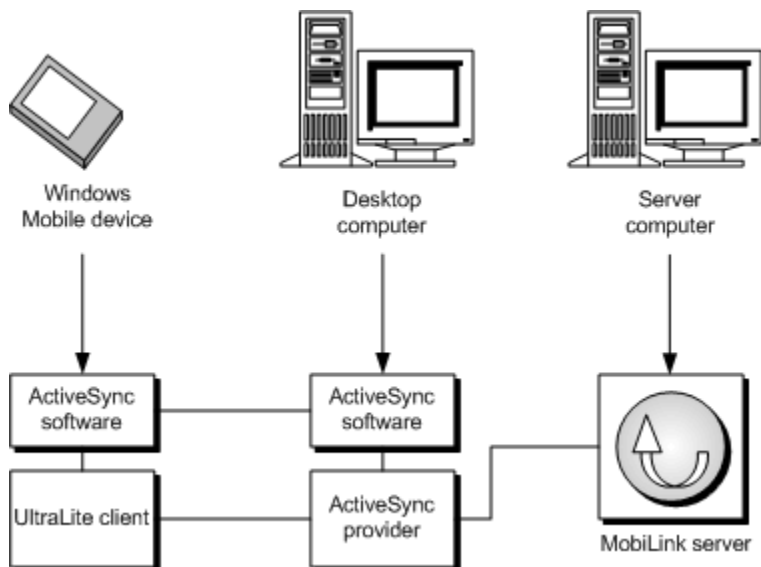
To use ActiveSync initiated synchronization requires that you:

- Register all applications that need to use ActiveSync initiated synchronization with ActiveSync.
- Have the ActiveSync provider installed on your desktop, and deployed to your device.

To determine which platforms the provider is supported on, see <http://www.sybase.com/detail?id=1002288>.

The ActiveSync architecture

The following diagram shows the computing layers required by the ActiveSync architecture.



Notice that you must install the ActiveSync provider on your device in addition to your desktop. You can only have a single ActiveSync provider on a single computer. However, if you have more than one UltraLite application installed on a Windows Mobile device, you can register them with the same provider so they are synchronized simultaneously.

See also

- UltraLite C/C++: “Add ActiveSync synchronization to your application” [[UltraLite - C and C++ Programming](#)]
- UltraLite for embedded SQL: “Adding synchronization to your application” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “MobiLink data synchronization” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]

ActiveSync synchronization overview

1. ActiveSync begins a synchronization session.
2. The ActiveSync provider sends a synchronize notification message to the first registered application on the device. The application is started if it is not yet running.
3. WndProc is invoked for each registered application.
4. Once the application has determined that this is the synchronize notification message from ActiveSync, the application calls `ULIsSynchronizeMessage` to invoke the database synchronization procedure.

5. Once synchronization is complete, the application calls `ULSignalSyncIsComplete` to let the provider know that it has finished synchronizing.
6. Steps two-five are repeated for each application that has been registered with the provider.

UltraLite synchronization parameters and network protocol options

Synchronization parameters for UltraLite

Synchronization parameters control the synchronization between an UltraLite database and the MobiLink server. The way you set parameters depends on the specific UltraLite interface you are using. This section describes the effects of the parameters, and provides links to other locations for information about to set them.

Note

The parameters described in this section only apply to UltraLite remote databases. To synchronize SQL Anywhere remote databases, see [“MobiLink SQL Anywhere client utility \(dbmlsync\)”](#) [*MobiLink - Client Administration*].

Required parameters

The following parameters are required:

- Stream Type
- User Name
- Version

The synchronization function throws an exception, such as `SQLCode.SQLLE_SYNC_INFO_INVALID` or its equivalent, if you do not set these parameters.

Conflicting parameters

You can specify at most one of these parameters:

- Download Only
- Ping
- Upload Only

The synchronization function throws an exception, such as `SQLCode.SQLLE_SYNC_INFO_INVALID` or its equivalent, if you set more than one of these parameters to true.

See also

- [“ULSynchronize method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULSyncParms class \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Synchronize method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“SyncParms class \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“SyncParms class \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Stream Type synchronization parameter” on page 105](#)
- [“User Name synchronization parameter” on page 110](#)
- [“Version synchronization parameter” on page 111](#)
- [“Download Only synchronization parameter” on page 92](#)
- [“Ping synchronization parameter” on page 99](#)
- [“Upload Only synchronization parameter” on page 109](#)

Additional Parameters synchronization parameter

This synchronization parameter allows an application to supply additional parameters that can not be readily specified using any other predefined parameters. Some parameters that are seldom used are specified in this parameter field.

The additional parameters are supplied as a string of keyword=value settings, separated with a semicolon.

Syntax

The syntax varies depending on the API you use. It is not available in UltraLiteJ.

Allowed values

The following properties can be specified as part of the additional parameters setting:

Property name	Description
AllowDownloadDupRows	Prevents errors from being raised when a synchronization encounters downloaded rows with duplicate primary keys. Set this property to 0 to raise errors and roll back the download; otherwise, set to 1 to raise warnings and continue the download. This property is only available in UltraLite C/C++.
CheckpointStore	Adds additional checkpoints of the database during synchronization to limit database growth during the synchronization process. Set this property to 1 to enable this feature, which is beneficial for large downloads with many updates but slows down synchronization; otherwise, set to 0 , which is the default.

Property name	Description
DisableConcurrency	<p>Disallows database access from other threads during synchronization during the upload phase.</p> <p>Set this property to 0 to allow concurrent database access; otherwise, set to 1. By default, this property is set to 0.</p>
TableOrder	<p>Sets the table order required for priority synchronization if the UltraLite default table ordering is not suitable for your deployment.</p> <p>Set this property to a list of table names, arranged in the desired order for upload. For UltraLite, use a comma delimited list; for ulsync, use a semicolon delimited list. By default, the order is based on foreign key relationships. Typically, the default is acceptable when the foreign keys on your consolidated database match the UltraLite remote database and there are no foreign key cycles.</p> <p>Quote tables names with either single or double quotes. For example, "Customer,Sales" and 'Customer,Sales' are both supported in UltraLite.</p> <p>If you include tables that are not included in the synchronization, they are ignored. Any tables that you do not list are appropriately sorted based on the foreign keys defined in the remote database.</p> <p>The order of tables on the download is the same as those you define for upload.</p> <p>You only need to explicitly set the table order if your UltraLite tables:</p> <ul style="list-style-type: none"> • Are part of foreign key cycles. You must then list all tables that are part of a cycle. • Have different foreign key relationships in the consolidated database.

See also

- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.AdditionalParms property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms.setAdditionalParms method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Example

UltraLite for C/C++ applications can set additional parameters as follows:

```
ul_sync_info info;
// ...
info.additional_parms = UL_TEXT(
    "AllowDownloadDupRows=1;
```

```
CheckpointStore=1;  
DisableConcurrency=1;  
TableOrder=Customer,Sales"  
);
```

Authentication Parameters synchronization parameter

Supplies parameters to authentication parameters in MobiLink events.

Syntax

The syntax varies depending on the API you use.

Remarks

Parameters may be a user name and password, for example.

If you use this parameter, you must also supply the number of parameters.

Allowed values

An array of strings. Null is not allowed as a value for any of the strings, but you can supply an empty string.

See also

- “Number of Authentication Parameters parameter” on page 96
- “Authentication parameters” [*MobiLink - Server Administration*]
- “authenticate_parameters connection event” [*MobiLink - Server Administration*]
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.AuthenticationParms property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms.getAuthenticationParms method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParms.setAuthenticationParms method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can set the parameters as follows:

```
ul_char * Params[ 3 ] = { UL_TEXT( "parm1" ),  
                          UL_TEXT( "parm2" ),  
                          UL_TEXT( "parm3" ) };  
  
// ...  
info.num_auth_parms = 3;  
info.auth_parms = Params;
```

Authentication Status synchronization parameter

This field is set by a synchronization to report the status of MobiLink user authentication. The MobiLink server provides this information to the client.

Syntax

The syntax varies depending on the API you use.

Allowed values

The allowed values are held in an interface-specific enumeration. For example, for C/C++ applications the enumeration is as follows.

Constant	Value	Description
UL_AUTH_STATUS_UNKNOWN	0	Authorization status is unknown, possibly because the connection has not yet synchronized.
UL_AUTH_STATUS_VALID	1	User ID and password were valid at the time of synchronization.
UL_AUTH_STATUS_VALID_BUT_EXPIRES_SOON	2	User ID and password were valid at the time of synchronization but will expire soon.
UL_AUTH_STATUS_EXPIRED	3	Authorization failed: user ID or password have expired.
UL_AUTH_STATUS_INVALID	4	Authorization failed: bad user ID or password.
UL_AUTH_STATUS_IN_USE	5	Authorization failed: user ID is already in use.

Remarks

If a custom **authenticate_user** synchronization script at the consolidated database returns a different value, the value is interpreted according to the rules given in an `authenticate_user` connection event.

If you are implementing a custom authentication scheme, the `authenticate_user` or `authenticate_user_hashed` synchronization script must return one of the allowed values of this parameter.

The parameter is set by the MobiLink server, and so is read-only.

See also

- “`authenticate_user` connection event” [*MobiLink - Server Administration*]
- “MobiLink users” [*MobiLink - Client Administration*]
- “`ul_sync_info` structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “`ULSyncResult.AuthStatus` property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “`SyncResult` class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “`SyncResult.getAuthStatus` method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can access the parameter as follows:

```
ul_sync_info info;
// ...
returncode = info.auth_status;
```

Authentication Value synchronization parameter

This field is set by a synchronization to report results of a custom MobiLink user authentication script. The MobiLink server provides this information to the client.

Syntax

The syntax varies depending on the API you use. It is not available in UltraLiteJ.

Remarks

The values set by the default MobiLink user authentication mechanism are described in the `authenticate_user` connection event and Authentication Status synchronization parameter.

The parameter is set by the MobiLink server, and so is read-only.

See also

- “`authenticate_user` connection event” [*MobiLink - Server Administration*]
- “`authenticate_user_hashed` connection event” [*MobiLink - Server Administration*]
- “Authentication Status synchronization parameter” on page 90
- “`ul_sync_info` structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “`ULSyncResult.AuthValue` property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “`SyncResult` class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “`SyncResult.getAuthValue` method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can access the parameter as follows:

```
ul_sync_info info;
// ...
returncode = info.auth_value;
```

Download Only synchronization parameter

Prevents changes from being uploaded from the UltraLite database during this synchronization.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Default

False

Allowed values

Boolean

Conflicts with

Ping and Upload Only

Remarks

When you have remotes that are synchronized by download-only synchronization, you should regularly do a full synchronization to reduce the amount of log that is scanned by the download-only synchronization. Otherwise, the download-only synchronizations will take an increasingly long time to complete.

For `ulsync` When download-only synchronization occurs, `ulsync` does not upload any changes to the data. Instead, it:

- Uploads information about the schema and the value stored in the progress counter.
- Ensures that changes on the remote database are not overwritten during download-only synchronization.

`ulsync` performs these actions by scanning the UltraLite database log to watch for rows with pending operations on the consolidated database. If `ulsync` detects a conflict, the download is rolled back and the synchronization fails. You must then do a full synchronization (that is an upload and a download) to correct this conflict.

See also

- [“Upload Only synchronization parameter” on page 109](#)
- [“Synchronization state tracking” on page 2](#)
- [“UltraLite Synchronization utility \(`ulsync`\)” on page 221](#)
- [“`ul_sync_info` structure \[UltraLite C and Embedded SQL datatypes\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“`ULSyncParms.DownloadOnly` property \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“`SyncParms` class \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“`SyncParms.setDownloadOnly` method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)

Example

`ulsync` supports this parameter as an extended synchronization parameter:

```
ulsync -c DBF=myuldb.udb  
"MobiLinkId=remoteA;ScriptVersion=2;DownloadOnly=ON;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.download_only = ul_true;
```

Ignored Rows synchronization parameter

This field is set by a synchronization to indicate that rows were ignored by the MobiLink server during synchronization because of absent scripts.

Syntax

The syntax varies depending on the API you use.

Allowed values

Boolean

Remarks

The parameter is read-only.

See also

- “[ul_sync_info structure \[UltraLite C and Embedded SQL datatypes\]](#)” [*UltraLite - C and C++ Programming*]
- “[ULSyncResult.IgnoredRows property \[UltraLite.NET\]](#)” [*UltraLite - .NET Programming*]
- “[SyncResult class \[UltraLite for M-Business Anywhere\]](#)” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “[SyncResult.getIgnoredRows method \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can access the parameter as follows:

```
ul_sync_info info;  
// ...  
res = info.ignored_rows;
```

Keep Partial Download synchronization parameter

Controls whether UltraLite holds on to the partial download rather than rolling back the changes, when a download fails because of a communications error during synchronization.

Syntax

The syntax varies depending on the API you use. It is not available in UltraLiteJ.

Default

False, which indicates that UltraLite rolls back all changes after a failed download.

Allowed values

Boolean

See also

- “Resumption of failed downloads” [*MobiLink - Server Administration*]
- “Resume Partial Download synchronization parameter” on page 101
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.KeepPartialDownload property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Example

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.keep_partial_download = ul_true;
```

New Password synchronization parameter

Sets a new MobiLink password associated with the user name.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Allowed values

String

Remarks

The parameter is optional.

See also

- “MobiLink users” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.NewPassword property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParms.setPassword method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

`ulsync` can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb
"MobiLinkId=remoteA;ScriptVersion=2;NewMobiLinkPwd=mynewpassword;Stream=http
"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.new_password = UL_TEXT( "mlnewpass" );
```

Number of Authentication Parameters parameter

Supplies the number of authentication parameters being passed to authentication parameters in MobiLink events.

Syntax

The syntax varies depending on the API you use. Not required for UltraLiteJ.

Default

No parameters passed to a custom authentication script.

Remarks

The parameter is used together with Authentication Parameters to supply information to custom authentication scripts.

See also

- “Authentication Parameters synchronization parameter” on page 90
- “authenticate_parameters connection event” [*MobiLink - Server Administration*]
- “Authentication parameters” [*MobiLink - Server Administration*]
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.AuthenticationParms property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Example

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.num_auth_parms = 3;
```

Observer synchronization parameter

Specifies a pointer to a callback function or event handler that monitors synchronization. The signature of the callback function that you need to implement to use is of the type **ul_sync_observer_fn**:


```
typedef void(UL_CALLBACK_FN *ul_sync_observer_fn)( ul_sync_status * status );
```

Syntax

The syntax varies depending on the API you use.

See also

- “User Data synchronization parameter” on page 109
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncProgressListener interface [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “Connection.synchronizeWithParm method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParams.setSyncObserver method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.observer=callfunction;
```

Partial Download Retained synchronization parameter

This field is set by a synchronization to indicate whether UltraLite applied those changes that were downloaded rather than rolling back the changes when a download fails because of a communications error during synchronization.

Syntax

The syntax varies depending on the API you use. Not supported by UltraLite Java edition databases.

Allowed values

Boolean

Remarks

The parameter is set during synchronization if a download error occurs and a partial download was retained.

Partial downloads are retained only if Keep Partial Download is set to true.

See also

- “Keep Partial Download synchronization parameter” on page 94
- “Resumption of failed downloads” [*MobiLink - Server Administration*]
- “Resume Partial Download synchronization parameter” on page 101
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncResult.PartialDownloadRetained property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncResult class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Example

Access the parameter as follows:

```
ul_sync_info info;  
// ...  
returncode=info.partial_download_retained;
```

Password synchronization parameter

Specifies the MobiLink password associated with the user name.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Allowed values

String

Remarks

The parameter is optional.

This MobiLink user name and password are different than any database user ID and password, and serve to only identify and authenticate the application to the MobiLink server.

If the MobiLink client already has a password, use the New Password parameter to change it.

See also

- “User Name synchronization parameter” on page 110
- “New Password synchronization parameter” on page 95
- “MobiLink users” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (`ulsync`)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.Password property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParms.setPassword method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

ulsync can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb  
"MobiLinkId=remoteA;ScriptVersion=2;MobiLinkPwd=mypassword;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.password = UL_TEXT( "mypassword" );
```

Ping synchronization parameter

Confirms communications between the UltraLite client and the MobiLink server. When this parameter is set to true, no synchronization takes place.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with ulsync.

Default

False

Allowed values

Boolean

Remarks

When the MobiLink server receives a ping request, it connects to the consolidated database, authenticates the user, and then sends the authenticating user status and value back to the client.

If the ping succeeds, the MobiLink server issues an information message. If the ping does not succeed, it issues an error message.

If the MobiLink user ID cannot be found in the ml_user system table and the MobiLink server is running with the command line option -zu+, the MobiLink server adds the user to ml_user.

The MobiLink server may execute the following scripts, if they exist, for a ping request:

- begin_connection
- authenticate_user
- authenticate_user_hashed
- authenticate_parameters
- end_connection

See also

- “-pi dbmlsync option” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParams.PingOnly property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParams class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParams.setPingOnly method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

ulsync can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb
"MobiLinkUid=remoteA;ScriptVersion=2;Ping=True;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.ping = ul_true;
```

Publications synchronization parameter

Specifies the publications to be synchronized.

Syntax

The syntax varies depending on the API you use. You can also use this parameter with ulsync.

Default

Synchronize all publications.

Remarks

When synchronizing in C/C++, set the publications synchronization parameter to a **publication list**: a comma-separated list of publication names.

See also

- “Publishing data in UltraLite” on page 76
- “UltraLite publications” on page 82
- “UltraLite Synchronization utility (ulsync)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULConnection.SYNC_ALL_DB field [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ULConnection.SYNC_ALL_PUBS field [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “PublicationSchema class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParams.setPublications method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

ulsync can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb
"MobiLinkId=remoteA;ScriptVersion=2;Publications=UL_PUB_MYPUB1,UL_PUB_MYPUB2
;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.publications = UL_TEXT( "Pubs1,Pubs3" );
```

Resume Partial Download synchronization parameter

Resumes a failed download.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with ulsync. Not supported by UltraLite Java edition databases.

Default

False

Allowed values

Boolean

Remarks

The synchronization does not upload changes; it only downloads those changes that were to be downloaded in the failed download.

See also

- [“Resumption of failed downloads” \[MobiLink - Server Administration\]](#)
- [“Keep Partial Download synchronization parameter” on page 94](#)
- [“Partial Download Retained synchronization parameter” on page 97](#)
- [“UltraLite Synchronization utility \(ulsync\)” on page 221](#)
- [“ul_sync_info structure \[UltraLite C and Embedded SQL datatypes\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULSyncParms.ResumePartialDownload property \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“SyncParms class \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)

Example

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.resume_partial_download = ul_true;
```

Send Column Names synchronization parameter

Specifies that column names should be sent in the upload.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Default

False

Allowed values

Boolean

Remarks

This option is used by the MobiLink server for direct row handling. When using direct row handling, you should enable this option. Otherwise, it has no effect.

If `ml_add_column` names are used, then names sent up from the client are ignored.

See also

- “Direct row handling” [*MobiLink - Server Administration*]
- “`ml_add_column` system procedure” [*MobiLink - Server Administration*]
- “UltraLite Synchronization utility (`ulsync`)” on page 221
- “`ul_sync_info` structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “`ULSyncParms.SendColumnNames` property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “`SyncParms` class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “`SyncParms.setSendColumnNames` method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

`ulsync` can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb
"MobiLinkUid=remoteA;ScriptVersion=2;SendColumnNames=true;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.send_column_names = ul_true;
```

Send Download Acknowledgement synchronization parameter

Instructs the MobiLink server that the client will provide a download acknowledgement.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Default

False

Allowed values

Boolean

See also

- [“UltraLite Synchronization utility \(ulsync\)” on page 221](#)
- [“ul_sync_info structure \[UltraLite C and Embedded SQL datatypes\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULSyncParms.SendDownloadAck property \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“SyncParms class \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“SyncParms.setAcknowledgeDownload method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)

Example

`ulsync` can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb
"MobiLinkId=remoteA;ScriptVersion=2;SendDownloadACK=true;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;
// ...
info.send_download_ack = ul_true;
```

Stream Error synchronization parameter

Provides a structure to hold communications error reporting information.

Syntax

The syntax varies depending on the API you use.

Applies to

This parameter applies only to C/C++ interfaces.

Allowed values

The parameter has no default value, and must be explicitly set using one of the supported fields. The `ul_stream_error` fields are as follows:

- **stream_error_code** For the error code suffixes, see `%SQLANY12%\SDK\Include\sserror.h`.

- **system_error_code** A system-specific error code. For more information about the error code, you must look at your platform documentation. For Windows platforms, this is the Microsoft Developer Network documentation.

The following are common system errors on Windows:

- **10048 (WSAADDRINUSE)** Address already in use.
 - **10053 (WSAECONNABORTED)** Software caused connection abort.
 - **10054 (WSAECONNRESET)** The other side of the communication closed the socket.
 - **10060 (WSAETIMEDOUT)** Connection timed out.
 - **10061 (WSAECONNREFUSED)** Connection refused. Typically, this means that the MobiLink server is not running or is not listening on the specified port. See <http://msdn2.microsoft.com/en-us/library/ms740668.aspx>.
- **error_string** A string with additional information, if available, for the stream_error_code. The string may or may not be empty. A non-empty error_string value provides information in addition to the stream_error_code value. For example, for a write error (error code 9) the error string is a number showing how many bytes it was trying to write.

Remarks

UltraLite applications other than the UltraLite C++ Component receive communications error information as part of the Sync Result parameter.

The stream_error field is a structure of type ul_stream_error.

```
typedef struct {
    ss_error_code stream_error_code;
    asa_uint16 alignment;
    asa_int32 system_error_code;
    char error_string[UL_STREAM_ERROR_STRING_SIZE];
} ul_stream_error, * p_ul_stream_error;
```

The structure is defined in `%SQLANYI2%\SDK\Include\sserror.h`.

Check for a `SQLE_MOBILINK_COMMUNICATIONS_ERROR`:

```
ULConnection * conn;
ul_sync_info info;
...
conn->InitSynchInfo( &info );
if( !conn->Synchronize( &info ) ) {
    ULError const * error = conn->GetLastError();
    char buf[256];
    if( error->GetSQLCode() == SQLE_MOBILINK_COMMUNICATIONS_ERROR ) {
        error->GetString( buf, sizeof(buf) );
        printf( "%s\n", buf );
        // more handling for communication error
    }
}
```


See also

- “MobiLink communication error messages” [*Error Messages*]
- “Sync Result synchronization parameter” on page 107
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]

Stream Type synchronization parameter

Sets the MobiLink network protocol to use for synchronization.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Remarks

This parameter is required. It has no default value.

Most network protocols require protocol options to identify the MobiLink server address and other behavior. These options are supplied in the Stream Parameters parameter.

When the network protocol requires an option, pass that option using the Stream Parameters parameter; otherwise, set the Stream Parameters parameter to null.

The following stream types are available, but not all are available on all target platforms:

Network protocol	Description
HTTP	Synchronize over HTTP.
HTTPS	Synchronize over HTTPS. The HTTPS protocol uses TLS as its underlying security layer. It operates over TCP/IP.
TCP/IP	Synchronize over TCP/IP. This protocol is not supported by UltraLite Java edition databases.
TLS	Synchronize over TCP/IP with transport-layer security (TLS). TLS secures client/server communications using digital certificates and public-key cryptography.

For a list of supported platforms, see <http://www.sybase.com/detail?id=1061806>.

See also

- “Stream Parameters synchronization parameter” on page 106
- “Certificate Creation utility (createcert)” [*SQL Anywhere Server - Database Administration*]
- “Certificate Viewer utility (viewcert)” [*SQL Anywhere Server - Database Administration*]
- “Transport-layer security” [*SQL Anywhere Server - Database Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 221
- “Network protocol options for UltraLite synchronization streams” on page 112
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.Stream property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

Example

For UltraLite for C/C++ applications, set the parameter as follows:

```
Connection conn;
ul_sync_info info;
...
conn.InitSynchInfo( &info );
info.stream = "http";
```

Stream Parameters synchronization parameter

Sets options to configure the network protocol.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with ulsync.

Default

Null

Allowed values

String

Remarks

This parameter is optional. It accepts a semicolon separated list of network protocol options. Each option is of the form *keyword=value*, where the allowed sets of keywords depends on the network protocol.

See also

- “UltraLite Synchronization utility (ulsync)” on page 221
- “Network protocol options for UltraLite synchronization streams” on page 112
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.StreamParms property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “StreamHTTPParms interface [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.stream_parms= UL_TEXT( "host=myserver;port=2439" );
```

Sync Result synchronization parameter

Reports the status of a synchronization.

Syntax

The syntax varies depending on the API you use.

Remarks

The parameter is set by UltraLite, and is read-only.

The C/C++ interface receives this information in separate parameters as part of a ul_sync_info struct. Otherwise, this information is defined as a compound parameter containing a variety of information in separate fields:

- **Authentication Status** Reports success or failure of authentication.
- **Ignored Rows** Reports the number of ignored rows.
- **Stream Error information** The Stream Error information includes a Stream Error Code, Stream Error Context, Stream Error ID, and Stream Error System.
- **Upload OK** Reports the success or failure of the upload phase.

See also

- “Authentication Status synchronization parameter” on page 90
- “Ignored Rows synchronization parameter” on page 94
- “Stream Error synchronization parameter” on page 103
- “Upload OK synchronization parameter” on page 108
- “ULSyncParms class [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ul_sync_result structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “ULGetSyncResult method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “SyncResult class [UltraLiteJ]” [*UltraLite - Java Programming*]

Upload OK synchronization parameter

This field is set by a synchronization to report the status of data uploaded to the MobiLink server.

Syntax

The syntax varies depending on the API you use.

Remarks

The parameter is set by UltraLite, and so is read-only.

After synchronization, the parameter holds **true** if the upload was successful, and **false** otherwise. You can check this parameter if there was a synchronization error, to know whether data was successfully uploaded before the error occurred.

See also

- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncResult.UploadOK property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncResult class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncResult.isUploadOK method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

UltraLite for C/C++ applications can access the parameter as follows:

```
ul_sync_info info;  
// ...  
returncode = info.upload_ok;
```

Upload Only synchronization parameter

Indicates that there should be no downloads in the current synchronization, which can save communication time, especially over slow communication links.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Default

False

Allowed values

Boolean

Conflicts with

Download Only, Ping, and Resume Partial Download

Remarks

When set to true, the client waits for the upload acknowledgement from the MobiLink server, after which it terminates the synchronization session successfully.

See also

- [“UltraLite client synchronization design” on page 72](#)
- [“Download Only synchronization parameter” on page 92](#)
- [“UltraLite Synchronization utility \(ulsync\)” on page 221](#)
- [“ul_sync_info structure \[UltraLite C and Embedded SQL datatypes\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULSyncParms.UploadOnly property \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“SyncParms class \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“SyncParms.isUploadOnly method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)

Example

`ulsync` can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb  
"MobiLinkUid=remoteA;ScriptVersion=2;UploadOnly=True;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.upload_only = ul_true;
```

User Data synchronization parameter

Makes application-specific information available to the synchronization observer.

Applies to

C/C++ applications only. Other components, such as UltraLite.NET, do not require a separate parameter to handle user data and so have no User Data parameter.

Syntax

The syntax varies depending on the API you use.

Remarks

When implementing the synchronization observer callback function or event handler, you can make application-specific information available by providing information using the User Data parameter.

See also

- [“Observer synchronization parameter” on page 96](#)
- [“ul_sync_info structure \[UltraLite C and Embedded SQL datatypes\]” \[UltraLite - C and C++ Programming\]](#)

User Name synchronization parameter

Required. A string that the MobiLink server uses for authentication purposes.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with `ulsync`.

Remarks

This parameter is required. Empty strings and NULL strings are universally rejected.

The parameter has no default value, and must be explicitly set.

The user name does not have to be unique when a remote ID is used.

This MobiLink user name and password are separate from any database user ID and password, and serves only to identify and authenticate the application to the MobiLink server.

For a user to be part of a synchronization system, you must register the user name with the MobiLink server. The user name is stored in the `name` column of the `ml_user` MobiLink system table in the consolidated database.

See also

- “Remote IDs” [*MobiLink - Client Administration*]
- “Password synchronization parameter” on page 98
- “MobiLink users” [*MobiLink - Client Administration*]
- “UltraLite user authentication” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.UserName property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParms.setUserName method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

ulsync can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb "MobiLinkUid=remoteA;ScriptVersion=2;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.user_name= UL_TEXT( "remoteA" );
```

Version synchronization parameter

Defines the consolidated database version.

Syntax

The syntax varies depending on the API you use. You can also set this parameter with ulsync.

Allowed values

String

Remarks

This parameter is required. Empty strings and NULL strings are universally rejected.

Each synchronization script in the consolidated database is marked with a version string. For example, there may be two different download_cursor scripts, identified by different version strings.

See also

- “Script versions” [*MobiLink - Server Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 221
- “ul_sync_info structure [UltraLite C and Embedded SQL datatypes]” [*UltraLite - C and C++ Programming*]
- “ULSyncParms.Version property [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “SyncParms class [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “SyncParms.setVersion method [UltraLiteJ]” [*UltraLite - Java Programming*]

Example

ulsync can set this parameter as an extended synchronization parameter as follows:

```
ulsync -c DBF=myuldb.udb "MobiLinkUid=remoteA;ScriptVersion=2;Stream=http"
```

UltraLite for C/C++ applications can set the parameter as follows:

```
ul_sync_info info;  
// ...  
info.version = UL_TEXT( "default" );
```

Network protocol options for UltraLite synchronization streams

You must set the network protocol in your application. Each UltraLite database that synchronizes with a MobiLink server does so over a network protocol. Available network protocols include TCP/IP, HTTP, HTTPS, and TLS. Support is also provided for ActiveSync notification on Windows Mobile.

For the network protocol you set, you can choose from a set of corresponding protocol options to ensure that the UltraLite application can locate and properly communicate with the MobiLink server. The network protocol options provide information such as addressing information (host and port) and protocol-specific information. Refer to the table below to determine which options you can use for the stream type you are using.

See also

- “MobiLink client network protocol options” [*MobiLink - Client Administration*]
- “Configuring UltraLite clients to use transport-layer security” [*SQL Anywhere Server - Database Administration*]
- “MobiLink client network protocol options” [*MobiLink - Client Administration*]
- “Stream Parameters synchronization parameter” on page 106
- -x option in “UltraLite Synchronization utility (ulsync)” on page 221

Set the synchronization stream and options

You can provide the information needed to locate the MobiLink server in your application by setting the Stream Parameters parameter.

See also

- [“Stream Parameters synchronization parameter” on page 106](#)
- [“ULSyncParms.StreamParms property \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“SyncParms.setStreamParms method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)

UltraLite deployment

In the majority of cases, development occurs on a Windows desktop or Mac OS X with the final release target for UltraLite being the mobile device. However, depending on your deployment environment, you can use various deployment mechanisms to install UltraLite.

UltraLite application projects may evolve with different iterations of the same UltraLite database: a development database, a test database, and a deployed production database. During the lifetime of a deployed database application, changes and improvements are first made in the development database, then propagated to the test database, before finally being distributed to the production database.

The modules you need to use for your UltraLite application depend on the platform you are targeting, the interface you are using, and the functionality you want to use.

See also

- [“Choose an UltraLite data management component for Windows Mobile” on page 13](#)
- [“Build and deploy UltraLite C++ applications” \[*UltraLite - C and C++ Programming*\]](#)
- [“Windows Mobile and Windows desktops deployment” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“UltraLite.NET application development” \[*UltraLite - .NET Programming*\]](#)
- [“UltraLiteJ application deployment” \[*UltraLite - Java Programming*\]](#)

UltraLite application build and deployment specifications

This section contains tables that show the build and deployment requirements for UltraLite applications.

Note

There may be versions of the UltraLite engine located in directories that contains the `_dev` suffix, such as the `x86_dev` directory. These versions contain development-time logging functionality that can be used to diagnose problems on platforms where it is difficult to debug. The logging output is intended for use by Sybase engineers rather than customers. For production systems, use a version of the engine that is not in a `_dev` directory.

Build and deployment requirements for UltraLite applications and encryption

The following table describes the minimum requirements to build and deploy an UltraLite application for all supported platforms and devices, including the requirements for UltraLite database encryption.

Platform or device	Minimum requirements	AES encryption requirements	FIPS 140-2 AES encryption requirements
<p>Windows Mobile and desktop (UltraLite C/C++ using static linkage)</p>	<p>Link against:</p> <ul style="list-style-type: none"> • <i>ulrt.lib</i>¹ • <i>ulbase.lib</i>¹ 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Call the EnableAesDBEncryption method.</p>	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Set the creation parameter fips=yes when creating the database.</p> <p>Call the EnableAesFipsDBEncryption method.</p> <p>Deploy:</p> <ul style="list-style-type: none"> • <i>ulfips12.dll</i>² • <i>sbgse2.dll</i>²
<p>Windows Mobile and desktop (UltraLite C/C++ using dynamic linkage)</p>	<p>Link against:</p> <ul style="list-style-type: none"> • <i>ulimp.lib</i>^{1, 10} • <i>ulbase.lib</i>¹ <p>Deploy:</p> <ul style="list-style-type: none"> • <i>ulrt12.dll</i>¹ 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Call the EnableAesDBEncryption method.</p>	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Set the creation parameter fips=yes when creating the database.</p> <p>Call the EnableAesFipsDBEncryption method.</p> <p>Deploy:</p> <ul style="list-style-type: none"> • <i>ulfips12.dll</i>² • <i>sbgse2.dll</i>²

Platform or device	Minimum requirements	AES encryption requirements	FIPS 140-2 AES encryption requirements
Windows Mobile and desktop (UltraLite C/C++ with the UltraLite engine)	Link against: <ul style="list-style-type: none"> ● <i>ulrtc.lib</i>¹ ● <i>ulbase.lib</i>¹ Deploy: <ul style="list-style-type: none"> ● <i>uleng12.exe</i>² 	Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.	Set the DBKEY creation parameter to the encryption key when creating and connecting to the database. Set the creation parameter fips=yes when creating the database. Deploy: <ul style="list-style-type: none"> ● <i>ulfips12.dll</i>² ● <i>sbgse2.dll</i>²
Windows Mobile and desktop (UltraLite.NET)	Add references to: <ul style="list-style-type: none"> ● iAnywhere.Data.UltraLite ● iAnywhere.Data.UltraLite.resources Deploy: <ul style="list-style-type: none"> ● <i>iAnywhere.Data.UltraLite.dll</i>⁷ ● <i>iAnywhere.Data.UltraLite.resources.dll</i>⁸ ● <i>ulnet12.dll</i>⁶ 	Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.	Set the DBKEY creation parameter to the encryption key when creating and connecting to the database. Set the creation parameter fips=yes when creating the database. Deploy: <ul style="list-style-type: none"> ● <i>ulfips12.dll</i>² ● <i>sbgse2.dll</i>²

Platform or device	Minimum requirements	AES encryption requirements	FIPS 140-2 AES encryption requirements
Windows Mobile and desktop (UltraLite.NET with the UltraLite engine)	<p>Add references to:</p> <ul style="list-style-type: none"> • iAnywhere.Data.UltraLite • iAnywhere.Data.UltraLite.resources <p>Deploy:</p> <ul style="list-style-type: none"> • <i>iAnywhere.Data.UltraLite.dll</i>⁷ • <i>iAnywhere.Data.UltraLite.resources.dll</i>⁸ • <i>ulnetclient12.dll</i>⁶ • <i>uleng12.exe</i>² 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p>	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Set the creation parameter fips=yes when creating the database.</p> <p>Deploy:</p> <ul style="list-style-type: none"> • <i>ulfips12.dll</i>² • <i>sbgse2.dll</i>²
Mac OS X and iOS (UltraLite C/C++)	<p>Add to your Xcode project:</p> <ul style="list-style-type: none"> • <i>libulrt.a</i>⁹ • <i>libulbase.a</i>⁹ (Mac OS X only) 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Call the EnableAesDBEncryption method.</p>	Not applicable
Linux (UltraLite C/C++)	<p>Link against:</p> <ul style="list-style-type: none"> • <i>libulrt.a</i>³ • <i>libulbase.a</i>³ 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p> <p>Call the EnableAesDBEncryption method.</p>	Not applicable
Android (UltraLiteJ)	<p>Add to your Android project:</p> <ul style="list-style-type: none"> • <i>UltraLiteJNI12.jar</i>⁵ • <i>libultralitej12.so</i>⁴ 	<p>Set the DBKEY creation parameter to the encryption key when creating and connecting to the database.</p>	Not applicable
BlackBerry (UltraLiteJ)	<p>Deploy:</p> <ul style="list-style-type: none"> • <i>UltraLiteJ12.cod</i>¹¹ 	Not applicable	Not applicable

¹ For Windows Mobile, this file is located in %SQLANY12%\UltraLite\CE\Arm.50\Lib. For Windows, it is located in %SQLANY12%\UltraLite\Windows\x64\Lib\VS8 or %SQLANY12%\UltraLite\Windows\x86\Lib\VS8.

² For Windows Mobile, this file is located in %SQLANY12%\UltraLite\CE\Arm.50. For Windows, it is located in %SQLANY12%\UltraLite\Windows\x64 or %SQLANY12%\UltraLite\Windows\x86.

³ This file is located in /opt/sqlanywhere12/ultralite/linux/x86/586/lib.

⁴ This file is located in %SQLANY12%\UltraLite\UltraLiteJ\Android\ARM.

⁵ This file is located in %SQLANY12%\UltraLite\UltraLiteJ\Android.

⁶ For Windows Mobile, this file is located in %SQLANY12%\UltraLite\UltraLite.NET\CE\Arm.50. For Windows, it is located in %SQLANY12%\UltraLite\UltraLite.NET\x64 or %SQLANY12%\UltraLite\UltraLite.NET\win32.

⁷ This file is located in %SQLANY12%\UltraLite\UltraLite.NET\Assembly\V2.

⁸ This file is located in %SQLANY12%\UltraLite\UltraLite.NET\Assembly\V2\en.

⁹ For Mac OS X, this file is located in /Applications/SQLAnywhere12/System/ultralite/macosx/x86_64. For iOS, UltraLite runtimes must be built after installation. Follow the instructions provided in *install-dir/ultralite/iphone/readme.txt*.

¹⁰ When linking against this library, define the UL_USE_DLL preprocessor macro when compiling. For example, specify the following:

¹¹ This file is located in %SQLANY12%\UltraLite\UltraLiteJ\BlackBerry4.2.

```
-DUL_USE_DLL
```

Note

To use UltraLite database obfuscation on any platform or device, you must specify the creation parameter **obfuscate=1** when creating the database.

Additional build and deployment requirements for synchronization and compression

The following table describes the stream, protocol option, and code requirements for building and deploying an UltraLite application that uses synchronization:

Synchronization type	Stream option specification	Protocol option requirements	Method call requirements for UltraLite C and C++
TCP/IP	"tcpip"	None	<ul style="list-style-type: none"> ● EnableTcpipSynchronization
HTTP	"http"	None	<ul style="list-style-type: none"> ● EnableHttpSynchronization
RSA TLS	"tls"	None	<ul style="list-style-type: none"> ● EnableTlsSynchronization ● EnableRsaSyncEncryption

Synchronization type	Stream option specification	Protocol option requirements	Method call requirements for UltraLite C and C++
RSA HTTPS	"https"	None	<ul style="list-style-type: none"> • EnableHttpsSynchronization • EnableRsaSyncEncryption
ECC TLS	"tls"	tls_type=ecc	<ul style="list-style-type: none"> • EnableTlsSynchronization • EnableEccSyncEncryption
ECC HTTPS	"https"	tls_type=ecc	<ul style="list-style-type: none"> • EnableHttpsSynchronization • EnableEccSyncEncryption
RSA FIPS 140-2 TLS	"tls"	fips=yes	<ul style="list-style-type: none"> • EnableTlsSynchronization • EnableRsaFipsSyncEncryption
RSA FIPS 140-2 HTTPS	"https"	fips=yes	<ul style="list-style-type: none"> • EnableHttpsSynchronization • EnableRsaFipsSyncEncryption

The following table describes additional protocol option and code requirements for building and deploying an UltraLite application that uses compression or end-to-end encryption:

Compression and encryption options	Protocol option requirements	Method call requirements for UltraLite C and C++
ZLIB compression	<ul style="list-style-type: none"> • compression=zlib 	<ul style="list-style-type: none"> • EnableZlibSyncCompression
RSA E2EE	<ul style="list-style-type: none"> • e2ee_public_key=key-file 	<ul style="list-style-type: none"> • EnableRsaE2ee
ECC E2EE	<ul style="list-style-type: none"> • e2ee_public_key=key-file • e2ee_type=ecc 	<ul style="list-style-type: none"> • EnableEccE2ee
RSA FIPS 140-2 E2EE	<ul style="list-style-type: none"> • e2ee_public_key=key-file • fips=yes 	<ul style="list-style-type: none"> • EnableRsaFipsE2ee

The following table illustrates additional build and deployment requirements for compression and encrypted synchronization.

Note

There are no additional build and deployment requirements for TCP/IP and HTTP synchronization.

Platform or device	ZLIB compression requirements	RSA TLS, RSA HTTPS, and RSA E2EE requirements	ECC TLS, ECC HTTPS, and ECC E2EE requirements	RSA FIPS 140-2 TLS, RSA FIPS 140-2 ECC, and RSA FIPS E2EE requirements
Windows Mobile and desktop (UltraLite C/C++ using static linkage)	None	Link against: <ul style="list-style-type: none"> • <i>ulrsa.lib</i>¹ 	Link against: <ul style="list-style-type: none"> • <i>ulecc.lib</i>¹ 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa-fips12.dll</i>² • <i>sbgse2.dll</i>²
Windows Mobile and desktop (UltraLite C/C++ using dynamic linkage)	Deploy: <ul style="list-style-type: none"> • <i>mlczlib12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcecc12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa-fips12.dll</i>² • <i>sbgse2.dll</i>²
Windows Mobile and desktop (UltraLite C/C++ with the UltraLite engine)	Deploy: <ul style="list-style-type: none"> • <i>mlczlib12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcecc12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa-fips12.dll</i>² • <i>sbgse2.dll</i>²
Windows Mobile and desktop (UltraLite.NET)	Deploy: <ul style="list-style-type: none"> • <i>mlczlib12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcecc12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa-fips12.dll</i>² • <i>sbgse2.dll</i>²
Windows Mobile and desktop (UltraLite.NET with the UltraLite engine)	Deploy: <ul style="list-style-type: none"> • <i>mlczlib12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcecc12.dll</i>² 	Deploy: <ul style="list-style-type: none"> • <i>mlcrsa-fips12.dll</i>² • <i>sbgse2.dll</i>²
Mac OS X and iOS (UltraLite C/C++)	None	None	Not applicable	Not applicable
Linux (UltraLite C/C++)	None	Link against: <ul style="list-style-type: none"> • <i>libulrsa.a</i>³ 	Link against: <ul style="list-style-type: none"> • <i>libulecc.a</i>³ 	Not applicable
Android (UltraLiteJ)	None	Deploy: <ul style="list-style-type: none"> • <i>libmlcrsa12.so</i>⁴ 	Not applicable	Not applicable

Platform or device	ZLIB compression requirements	RSA TLS, RSA HTTPS, and RSA E2EE requirements	ECC TLS, ECC HTTPS, and ECC E2EE requirements	RSA FIPS 140-2 TLS, RSA FIPS 140-2 ECC, and RSA FIPS E2EE requirements
BlackBerry (UltraLiteJ)	Not applicable	Not applicable	Not applicable	Not applicable

¹ For Windows Mobile, this file is located in %SQLANY12%\UltraLite\CE\Arm.50\Lib. For Windows, it is located in %SQLANY12%\UltraLite\Windows\x64\Lib\VS8 or %SQLANY12%\UltraLite\Windows\x86\Lib\VS8.

² For Windows Mobile, this file is located in %SQLANY12%\CE\Arm.50. For Windows, it is located in %SQLANY12%\Windows\x64 or %SQLANY12%\Windows\x86.

³ This file is located in /opt/sqlanywhere12/ultralite/linux/x86/586/lib.

⁴ This file is located in %SQLANY12%\UltraLite\UltraLiteJ\Android\ARM.

See also

- “UltraLite C++ application deployment” [[UltraLite - C and C++ Programming](#)]
- “UltraLite.NET application deployment” [[UltraLite - .NET Programming](#)]
- “UltraLiteJ application deployment” [[UltraLite - Java Programming](#)]
- “MobiLink client network protocol options” [[MobiLink - Client Administration](#)]

Deploy an UltraLite and UltraLite Java edition database to devices

You can use one of the following techniques to get the initial database file onto the device:

- Use the UltraLite API in your application to create the initial database file.
- Create a schema file from a SQL script and use the ALTER DATABASE SCHEMA FROM FILE statement.¹
- Use UltraLite FileTransfer methods to download the initial database file if it does not already exist on the device.
- Bundle the initial database with the application.¹
- When deploying to a Windows or Windows Mobile device, use central administration to send down the initial UDB file, or send a command to create the initial database file and execute SQL to give it schema.¹

¹ Not supported for UltraLite Java edition databases.

See also

- “Deploying UltraLite database schema upgrades” on page 123
- “ALTER DATABASE SCHEMA FROM FILE statement [UltraLite]” on page 412
- “Central administration of remote databases” [*MobiLink - Server Administration*]

Deploying UltraLite database schema upgrades

UltraLite database schema upgrades can be deployed using one of the following techniques:

- **Individual DDL statements** For example, in UltraLite C, you can run the following statement to create a new publication:

```
dbconnection->ExecuteStatement("CREATE PUBLICATION p (table t)");
```

- **The ALTER DATABASE SCHEMA FROM FILE statement** This statement can be used to perform schema upgrades when you do not know the DDL statement requirements, or do not want to specify the individual DDL statements. The SQL file you use must contain the entire new schema.

Caution

Do not reset a device during a schema upgrade. If you reset the device during a schema upgrade, data will be lost and the UltraLite database marked as "bad."

UltraLite executes the following steps when you upgrade an UltraLite database schema with the ALTER DATABASE SCHEMA FROM FILE statement:

1. Both the new and existing database schemas are compared to see what differs.
2. The schema of the existing database is altered.
3. Rows that do not fit the new schema are dropped. For example:
 - If you add a uniqueness constraint to a table and there are multiple rows with the same values, all but one row will be dropped.
 - If you try to change a column domain and a conversion error occurs, then that row will be dropped. For example, if you have a VARCHAR column and convert it to an INT column and the a row has the value ABCD, then that row is dropped.
 - If your new schema has new foreign keys where the foreign row does not have a matching primary row, these rows are dropped.
4. When rows are dropped, a SQLE_ROW_DROPPED_DURING_SCHEMA_UPGRADE (130) warning is raised.

Upgrade an UltraLite database schema with the ALTER DATABASE SCHEMA FROM FILE statement

1. Create a SQL script of DDL statements to create a completely new schema.

You can keep a master schema on your development machine and update the schema as your application changes.

Use either the `ulinit` or `ulunload` utilities to extract the DDL statements required for your script. By using these utilities with the following options, you ensure that the DDL statements are syntactically correct:

- For an UltraLite database, use the `ulunload` utility with the `-n` and `-s [schema-file]` options. For example:

```
ulunload -c dbf=mydatabase.udb -n -s MySchema.sql
```

- For a SQL Anywhere database, use the `ulinit` utility with the `-a` and `-l [schema-file]` option. For example:

```
ulinit -a "dsn=mysqlanywheredatabase" -l MySchema.sql
```

If you do not use the `ulunload` or `ulinit` utilities, review the script and ensure the following:

- The script declares the entire desired schema with `CREATE` statements.
 - Tables, columns, and publications are not renamed. The `RENAME` operation is not supported. Renamed tables are processed as a `DROP TABLE` and `CREATE TABLE` operation.
 - There are no non-DDL statements, including non-DDL statements that may not have the effect you expect.
 - Words in the SQL statement are separated by spaces.
 - Only one SQL statement appears in each line.
 - Comments are prepended with double hyphens (`-`), and only occur at the start of a line.
 - Each statement is separated by a line containing exactly the word **GO**.
2. Deploy the new SQL script file.
 3. Ensure that the database is synchronized.
 4. Run the new statement on the device. For example:

```
ALTER DATABASE SCHEMA FROM FILE 'MySchema.sql' ;
```

Error notification

Because UltraLite error callback is active during the upgrade process, you are notified of errors during the conversion process. For example, `SQLE_CONVERSION_ERROR` reports all values that could not be converted in its parameters. Errors do *not* mean the process failed. The final SQL code after the statement returns is a 130 warning in this case. These warnings describe operations of the conversion process and do not stop the upgrade process.

See also

- [“UltraLite database schemas” on page 45](#)
- [“ALTER DATABASE SCHEMA FROM FILE statement \[UltraLite\]” on page 412](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Database Unload utility \(ulunload\)” on page 227](#)

Start a deployed UltraLite engine

When using the UltraLite engine to manage data on a Windows or Windows Mobile device, your UltraLite application starts the engine automatically unless the application needs to explicitly provide the directory location of the engine.

When an UltraLite application attempts to start the UltraLite engine, the application searches the following directories:

Client platform	Directory locations
Windows desktop	<ol style="list-style-type: none"> 1. The directory of the application that is starting it 2. The current working directory 3. The system path 4. The SQL Anywhere install directory (either under <i>bin32</i> or <i>bin64</i>), depending on whether the client is 32-bit or 64-bit
Windows Mobile/CE	<ol style="list-style-type: none"> 1. <i>\Windows\</i> 2. <i>\</i> (the root directory) 3. <i>\UltraLiteDB\</i>
Linux	<ol style="list-style-type: none"> 1. The directory of the application that is auto-starting it 2. <i>%%SQLANY12%/bin32</i>

If the UltraLite engine is stored in a different location, start the engine by specifying the START connection parameter.

For example, a connection string to the database or connection code for a Windows Mobile client application might use the following START parameter value:

```
"START=\Program Files\MyApp\uleng12.exe"
```

See also

- [“UltraLite application build and deployment specifications” on page 115](#)
- [“UltraLite START connection parameter” on page 185](#)

Deploying the ActiveSync provider for UltraLite

The UltraLite ActiveSync provider is a software module that let users gain access to their devices from the desktop. Like other software components, you need to deploy the necessary files to the device to ensure that UltraLite operates with Windows Mobile ActiveSync.

When you deploy UltraLite for the end user, you must manually install and register the ActiveSync provider on the end user's computer. This requirement ensures that ActiveSync knows when to call a specific instance of a provider for a specific application.

- ***mlasinst.exe*** Installs the ActiveSync provider and registers it with the ActiveSync Manager. This utility also registers applications with the ActiveSync provider for synchronization.
- ***mlasdesk.dll*** The DLL that is loaded by the ActiveSync Manager on the desktop. *mlasinst.exe* registers the location of this file with the ActiveSync Manager.
- ***mlasdev.dll*** The DLL that is loaded by the ActiveSync Manager on the device. *mlasinst.exe* deploys this file to the correct location on the device.
- ***dbigen12.dll*** The language resource library.

For a list of supported provider platforms, see <http://www.sybase.com/detail?id=1002288>.

Install ActiveSync applications

1. Ensure that the end-user has:
 - The ActiveSync Manager installed.
 - The ActiveSync provider files copied from a development computer to the user's hard drive.
2. Run *mlasinst* to install a provider for ActiveSync. You can also use it to register and deploy the UltraLite application to the user's Windows Mobile device—depending on the command line syntax you use. If your UltraLite application uses multiple files, you must manually copy the required files.

The following example assumes that both *mlasdesk.dll* and *mlasdev.dll* are in the current directory. The *-k* and *-v* options are used. The *-p* and *-x* options are command line options for the application when it is started by ActiveSync.

```
mlasinst "C:\My Files\myULapp.exe" "\Program Files\myULapp.exe"  
"My Application" MYAPP -p -x -v -k
```

If you were to use this utility to deploy a pre-compiled CustDB for the ARM 5.0 processor, the command line would be similar to the following one:

```
mlasinst -v "%SQLANY12%\UltraLite\ce\arm.50"  
"%SQLANY12%\UltraLite\ce\arm.50\custdb.exe" custdb.exe CustDB CUSTDBDEMO
```

Note

You can also use the ActiveSync to register your UltraLite application at a later time if you choose. See “[Registering applications with the ActiveSync Manager](#)” on page 127.

3. Restart your computer so ActiveSync can recognize the new provider.
4. Enable the MobiLink provider.
 - a. From the ActiveSync window, click **Options**.
 - b. Check **MobiLink Clients** in the list and click **OK** to activate the provider.
 - c. To see a list of registered applications, click **Options**, click **MobiLink Clients**, and click **Settings**.

See also

- [“Registering applications with the ActiveSync Manager” on page 127](#)
- [“Microsoft ActiveSync Provider Installation utility \(mlasinst\)” \[MobiLink - Client Administration\]](#)

Registering applications with the ActiveSync Manager

Applications that use ActiveSync synchronization must be registered with ActiveSync and copied to the device.

You can register your application for use with ActiveSync either by using the ActiveSync Provider Installation utility or using the ActiveSync Manager itself. This section describes how to use the ActiveSync Manager.

Register an application for use with the ActiveSync Manager

1. Launch ActiveSync.
2. From the ActiveSync window, click **Options**.
3. From the list of information types, click **MobiLink Clients** and click **Settings**.
4. In the **MobiLink Synchronization** window, click **New**.
5. Enter the following information for your application:
 - **Application name** A name identifying the application that appears in the ActiveSync user interface.
 - **Class name** The registered class name for the application.
 - **Path** The location of the application on the device.
 - **Arguments** Any command line arguments to be used when ActiveSync starts the application.
6. Click **OK** to register the application.

See also

- [“Assigning class names for applications”](#) [*UltraLite - C and C++ Programming*]
- [“Microsoft ActiveSync Provider Installation utility \(mlasinst\)”](#) [*MobiLink - Client Administration*]

Tutorial: Building the UltraLite CustDB sample application

The CustDB sample is installed with SQL Anywhere. It is a multi-tiered database management solution that implements MobiLink synchronization with a SQL Anywhere consolidated database.

CustDB consists of the following:

- A **consolidated** SQL Anywhere database. The database is pre-populated with sales status data.
- A **remote** UltraLite database. This database is initially empty.
- An UltraLite client application.
- MobiLink server synchronization scripts.

Different versions of the application code exist for each supported programming interface and platform. However, the tutorial references the compiled version of the application for Windows desktops only. Remember that each version varies to conform to the conventions of each platform.

Note

You can only run one instance of CustDB at a time. Trying to run more than one instance brings the first instance to the foreground.

CustDB allows sales personnel to track and monitor transactions and then pool information from two types of users:

- Sales personnel that authenticate with user IDs 51, 52, and 53.
- Mobile managers that authenticate with user ID 50.

Information gathered by these different users can be synchronized with the consolidated database.

After following each lesson you will know how to:

- Run the MobiLink server to carry out data synchronization between the consolidated database and the UltraLite remote.
- Use Sybase Central to browse the data in the UltraLite remote.
- Manage UltraLite databases with UltraLite command line utilities.

See also

- [“CustDB sample overview” on page 130](#)
- [“CustDB Scenario” \[MobiLink - Getting Started\]](#)
- [“Users in the CustDB sample” \[MobiLink - Getting Started\]](#)
- [“Tables in the CustDB databases” \[MobiLink - Getting Started\]](#)

CustDB sample overview

SQL Anywhere CustDB database

This is the consolidated database. During installation, an ODBC data source called SQL Anywhere 12 CustDB is created for this database. The database file is located at `%SQLANYWHERE%\UltraLite\CustDB`. You can recreate this database at any time using `makedbs.cmd` (`makedbs.sh` for Mac or Linux).

You can erase changes that were synchronized into the consolidated `CustDB.db` file, so you have a clean version to work with using this script: `C:\Program Files\SQL Anywhere 12\UltraLite\CustDB\newdb.bat`

For more information about the schema of this file, see [“CustDB sample for MobiLink” \[MobiLink - Getting Started\]](#).

The UltraLite CustDB database

This is the remote version of the consolidated database that contains only a subset of the information, depending on which user synchronizes the database.

The file name and location can vary depending on the platform, programming language, or even device.

- For UltraLite.NET: `C:\Program Files\SQL Anywhere 12\UltraLite.NET\CustDB\Common\custdb.udb`
- For all other platforms and APIs: `C:\Program Files\SQL Anywhere 12\UltraLite\CustDB\custdb.udb`

The UltraLite database is also recreated by the `makedbs.cmd` script (`makedbs.sh` for Mac or Linux).

RDBMS-specific build scripts

The SQL scripts that rebuild a CustDB consolidated database for any one of the supported RDBMSs.

In the `C:\Program Files\SQL Anywhere 12\MobiLink\CustDB` directory, you can find the following files:

- For SQL Anywhere: `custdb.sql`
- For Adaptive Server Enterprise: `custase.sql`
- For Microsoft SQL Server: `custmss.sql`
- For Oracle: `custora.sql`
- For IBM DB2: `custdb2.sql`
- For MySQL: `custmys.sql`

For more information about setting up a consolidated database, see [“Setting up the CustDB consolidated database” \[MobiLink - Getting Started\]](#).

UltraLite CustDB client applications and ReadMe files

These are the end-user applications that provide a user-friendly interface to the UltraLite remote database. There is a sample client installed for each supported platform.

Each client application also contains a `ReadMe.html` or `ReadMe.txt` file. Although the contents of these files vary, they all include an outline the steps required to build and run the sample.

The location the applications and its ReadMe depends on your development environment. See “[Lesson 1: Building and running the CustDB application](#)” on page 131.

Synchronization logic

The UltraLite database SQL statements and synchronization calls are located in *custdbcpp.cpp* for the C++ API and in *custdb.sql* for ESQL.

See also

- “MobiLink consolidated databases” [*MobiLink - Server Administration*]

Lesson 1: Building and running the CustDB application

The CustDB application is built for many development environments.

Build and run the CustDB application

1. Build the CustDB application:
 - a. Open a CustDB project file in the appropriate environment.
 - b. Compile the source code.
2. Run the CustDB application:
 - a. Deploy the CustDB executable file to the mobile device.
 - b. Deploy the UltraLite CustDB database to the mobile device.
 - c. Run the CustDB executable file.

UltraLite for Windows 32-bit desktop

You do not need to build the CustDB application before running it.

You can find the CustDB executable file in the *%SQLANY12%\UltraLite\Windows\x86* directory.

UltraLite for C/C++

- **All versions of C/C++** You can find multiple versions of the C/C++ CustDB project file because of the many C/C++ development environments. Most versions make use of the generic files. These files are located in the *C:\Program Files\SQL Anywhere 12\UltraLite\Custdb* directory.

For information about all versions of C/C++ CustDB applications, see *C:\Program Files\SQL Anywhere 12\UltraLite\Custdb\readme.txt*.

- **Visual Studio** You can find project files in the *C:\Program Files\SQL Anywhere 12\UltraLite\CustDB\vs9* and *C:\Program Files\SQL Anywhere 12\UltraLite\CustDB\vs8* directories depending on your version of Visual Studio. To build and run the CustDB application, follow the instructions given in the beginning of the lesson.

- **Xcode for iPhone** For development on Mac OS X and iPhone, see *samples/ultralite/custdb/iphone*.

UltraLite.NET

You can find project files specific to Microsoft Visual Studio in the *C:\Program Files\SQL Anywhere 12\UltraLite.NET\CustDB* directory.

For instructions on building the CustDB application for Windows Mobile using Microsoft Visual Studio, refer to *C:\Program Files\SQL Anywhere 12\UltraLite.NET\CustDB\CE\ReadMe.html*.

To obtain deployment directory information for Microsoft Windows desktop, and information about where to download additional UltraLite.NET samples, see *C:\Documents and Settings\All Users\Documents\SQL Anywhere 12\Samples\UltraLite.NET\CustDB\Desktop\ReadMe.html*.

UltraLite for M-Business Anywhere

You can find project files specific to M-Business Anywhere in the *C:\Program Files\SQL Anywhere 12\UltraLiteForMBusinessAnywhere\CustDB* directory.

For more information about building the CustDB application using M-Business Anywhere on Windows and Windows Mobile, see “[Quick Start to UltraLite for M-Business Anywhere](#)” [*UltraLite - M-Business Anywhere Programming (deprecated)*].

Lesson 2: Logging in and populating the UltraLite remote database

This lesson demonstrates how to:

- Start the sample MobiLink server.
- Start the sample UltraLite client application.
- Log into UltraLite.

In this tutorial, the sample application is running on the same desktop computer as the MobiLink server. However, you can deploy a client application to the device and achieve the same result.

Start and synchronize the sample application

1. Click **Start » Programs » SQL Anywhere 12 » MobiLink » Synchronization Server Sample**. Or, execute the following command:

```
mksrv12 -c "DSN=SQL Anywhere 12 CustDB" -vcrs
```

(Use *mobilink.sh* on Mac OS X or Linux.)

The window displays messages about the MobiLink server's status.

2. Click **Start » Programs » SQL Anywhere 12 » UltraLite » Windows Sample Application**.
3. On the **File** menu, click **Synchronize**.

The application synchronizes and the MobiLink server messages window displays messages showing the synchronization taking place.

The synchronization script determines which subset of customers, products, and orders is downloaded to the application when user 50 logs in. In this case, only orders that have not yet been approved are downloaded.

4. Confirm that the company name and a sample order appear in the application window.

Lesson 3: Using the CustDB client application

Both the consolidated and remote databases contain a table named ULOrder. While the consolidated database holds all orders (approved and those pending approval), the UltraLite remote database only displays a subset of rows according to the user that has authenticated.

Columns in the table appear as fields in the client application. When you add an order, you must populate the Customer, Product, Quantity, Price, and Discount fields. You can also append other details such as Status or Notes. The timestamp column identifies whether the row needs to be synchronized.

Browse orders

1. Browsing orders is accomplished in a similar method for each version of the UltraLite client application. By browsing an order, you are scrolling through the data in your local UltraLite database. Because customers are sorted alphabetically, you can easily scroll through the list and locate a customer by name.

To scroll down the list of customers, click **Next**.

2. To scroll up through the list of customers, click **Previous**.

Add an order

1. Adding an order is carried out in a similar way in each version of the UltraLite client application. By adding an order, you have modified the data in your local UltraLite database. This data is not shared with the consolidated database until you synchronize.

Click **Order » New**.

2. In the **Customer** list, click **Basements R Us**.
3. In the **Product** list, click **Screwmaster Drill**. The price of this item is automatically entered in the **Price** field.
4. In the **Quantity** field, type **20**.
5. In the **Discount** field, type **5** (percent) and click **OK**.

Approve, deny, and delete orders

1. Because you have authenticated your identity as user ID 50, you are a manager that can perform all the same tasks as a sales person, but you have the added ability to accept or reject orders. By accepting or rejecting an order, you are changing the status of it and adding an additional note for the sales person to review. However, the data in the consolidated database is unchanged until you synchronize.

Approve the order for **Apple Street Builders**.

- a. To locate the customer, click **Previous**.
- b. To approve the order, click **Order** and then **Approve**.
- c. In the **Note** list, click **Good**.
- d. Click **OK**.

The order appears with a status of **Approved**.

2. Deny the order for **Art's Renovations**.
 - a. Go to the next order in the list, which is from **Art's Renovations**.
 - b. To deny the order, click **Order** and then **Deny**.
 - c. In the **Note** list, click **Discount Is Too High**.
 - d. Click **OK**.

The order appears with a status of Denied.

3. Delete the order for **Awnings R Us**.
 - a. Go to the next order in the list, which is from **Awnings R Us**.
 - b. Delete this order by choosing **Order » Delete**.
Click **OK** to confirm the deletion.

The order is marked as deleted. However, the current data remains in the UltraLite remote database until you synchronize changes to the consolidated database.

See also

- [“Tables in the CustDB databases” \[MobiLink - Getting Started\]](#)

Lesson 4: Synchronizing with the CustDB consolidated database

For synchronization to take place, the MobiLink server must be running. If you have shut down your MobiLink server, you need to restart it. See [“Lesson 2: Logging in and populating the UltraLite remote database” on page 132](#).

The synchronization process for the sample application removes approved orders from your database.

You can use Interactive SQL or Sybase Central to connect to the consolidated database and confirm that your changes were synchronized.

Synchronize the UltraLite remote

1. To synchronize your data, from the **File** menu click **Synchronize Database**.
2. Confirm that synchronization took place.
 - At the remote database, you can confirm that all required transactions occurred by checking that the order for **Awnings R Us** is now deleted. Perform this action by browsing the orders to confirm the absence of this entry.
 - At the consolidated database, you can also confirm that all required actions occurred by checking data in the consolidated database.

Confirm the synchronization (Sybase Central)

1. Click **Start » Programs » SQL Anywhere 12 » Administration Tools » Sybase Central**.
2. Click **Connections » Connect With SQL Anywhere 12**.
3. In the **Action** drop-down, click **Connect with an ODBC Data Source**.
4. Click **ODBC Data Source Name**.
5. Click **Browse** and click **SQL Anywhere 12 CustDB**.
6. Click **OK**.
7. Click **Connect**.
8. Double-click **Tables**.
9. Double-click **ULOrder**.
10. Click the **Data** tab and verify that order 5100 is approved, order 5101 is denied, and order 5102 is deleted.

Confirm the synchronization (Interactive SQL)

1. Connect to the consolidated database from Interactive SQL.
 - a. Click **Start » Programs » SQL Anywhere 12 » Administration Tools » Interactive SQL**.
 - b. In the **Action** drop-down, click **Connect with an ODBC Data Source**.
 - c. Click **ODBC Data Source Name** and click **SQL Anywhere 12 CustDB**.
2. To confirm that the approval and denial have been synchronized, execute the following statement:

```
SELECT order_id, status
FROM ULOrder
WHERE status IS NOT NULL;
```

The results show that order 5100 is approved, and 5101 is denied.

3. The deleted order has an order_id of 5102. The following query returns no rows, demonstrating that the order has been removed from the system.

```
SELECT *
FROM ULOrder
WHERE order_id = 5102;
```

Lesson 5: Browsing MobiLink synchronization scripts

The synchronization logic for CustDB is held in the consolidated database as MobiLink synchronization scripts. Synchronization logic allows you to determine how much of the consolidated database you need to download and/or upload. You can download complete tables or partial tables (with either row or column subsets) using such techniques as timestamp-based synchronization or snapshot synchronization.

In addition to the tables, users, and publications, you can also use Sybase Central to browse the synchronization scripts that are stored in the consolidated database. Sybase Central is the primary tool for adding these scripts to the database.

The *custdb.sql* file adds each synchronization script to the consolidated database by calling `ml_add_connection_script` or `ml_add_table_script`. Connection scripts control high level events that are not associated with a particular table. Use these events to perform global tasks that are required during every synchronization. Table scripts allow actions at specific events relating to the synchronization of a specific table, such as the start or end of uploading rows, resolving conflicts, or selecting rows to download.

For more information about the synchronization logic used in CustDB, see [“Synchronization logic source code” \[MobiLink - Getting Started\]](#).

For more information about the implementation of synchronization in CustDB, see [“Synchronization design” \[MobiLink - Getting Started\]](#).

Browse the synchronization scripts

1. Click **Start** » **Programs** » **SQL Anywhere 12** » **Administration Tools** » **Sybase Central**.
2. In the left pane of **Sybase Central**, right-click **MobiLink 12** and then **Open Project**. Navigate to `%SQLANYXSAMP12%\MobiLink\CustDB` and click *project.mlp*. Double-click **Consolidated Databases** and then double-click the CustDB consolidated database. You are connected to the consolidated database based on the connection information that was provided when you added the consolidated database to your project.
3. Double-click **Connection Scripts**.

The right pane lists a set of synchronization scripts and a set of events with which these scripts are associated. As the MobiLink server carries out the synchronization process, it triggers a sequence of

events. Any synchronization script associated with an event is run at that time. By writing synchronization scripts and assigning them to the synchronization events, you can control the actions that are carried out during synchronization.

4. Click **Synchronized Tables**.
5. In the right pane, double-click **ULCustomer**.

A set of scripts specific to this table, and their corresponding events appears. These scripts control the way that data in the ULCustomer table is synchronized with the remote databases.

See also

- [“Synchronization script writing” \[MobiLink - Server Administration\]](#)
- [“UltraLite clients” on page 67](#)
- [“Connection scripts” \[MobiLink - Server Administration\]](#)
- [“Table scripts” \[MobiLink - Server Administration\]](#)

Building your own application

Use one of the supported interfaces to build your own application. For more information, see:

- [“Tutorial: Building a Windows application using the C++ API” \[UltraLite - C and C++ Programming\]](#)
- [“Tutorial: Building a Windows Mobile application using UltraLite.NET” \[UltraLite - .NET Programming\]](#)
- [“Tutorial: Building an M-Business Anywhere application” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)

UltraLite database reference

This section provides a reference for UltraLite database properties, options, connection parameters, and utilities.

UltraLite creation parameters

This section describes the UltraLite creation parameters that are available when you create a new UltraLite database.

In addition to creation parameters, the following UltraLite connection parameters can be used as creation parameters:

- [“UltraLite CE_FILE connection parameter” on page 171](#)
- [“UltraLite DBF connection parameter”](#)
- [“UltraLite DBKEY connection parameter” on page 176](#)
- [“UltraLite desktop connection parameter” on page 177](#)
- [“UltraLite device connection parameter” on page 179](#)
- [“UltraLite MIRROR_FILE connection parameter” on page 180](#)
- [“UltraLite NT_FILE connection parameter” on page 181](#)
- [“UltraLite PWD connection parameter” on page 183](#)
- [“UltraLite RESERVE_SIZE connection parameter” on page 184](#)
- [“UltraLite UID connection parameter” on page 187](#)

Members

UltraLite supports the following creation parameters:

Name	Description
case	Sets the case-sensitivity of string comparisons in the UltraLite database. See “UltraLite case creation parameter” on page 141 .
checksum_level	Sets the level of checksum validation in the database. See “UltraLite checksum_level creation parameter” on page 142 .

Name	Description
collation	Sets the collation sequence used by the UltraLite database. Setting this property with or without the UTF-8 property determines the character set of the database. See “UltraLite character sets” on page 23 and “UltraLite collation creation parameter” on page 143 and “UltraLite utf8_encoding creation parameter” on page 164 .
date_format	Sets the default string format in which dates are retrieved from the database. See “UltraLite date_format creation parameter” on page 144 .
date_order	Controls the interpretation of date ordering of months, days, and years. See “UltraLite date_order creation parameter” on page 147 .
fips	Controls AES FIPS-certified encryption by using a Certicom certified cryptographic algorithm. See “UltraLite database security” on page 26 and “UltraLite fips creation parameter” on page 148 .
max_hash_size	Sets the default index hash size in bytes. See “UltraLite max_hash_size creation parameter” on page 149 .
nearest_century	Controls the interpretation of two-digit years in string-to-date conversions. See “UltraLite nearest_century creation parameter” on page 151 .
obfuscate	Controls whether data in the database is obfuscated. Obfuscation is a form of simple encryption. See “UltraLite database security” on page 26 and “UltraLite obfuscate creation parameter” on page 152 .
page_size	Defines the database page size. See “UltraLite page_size creation parameter” on page 153 .
precision	Specifies the maximum number of digits in decimal point arithmetic results. See “UltraLite precision creation parameter” on page 155 .
scale	Specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision. See “UltraLite scale creation parameter” on page 156 .
time_format	Sets the format for times retrieved from the database. See “UltraLite time_format creation parameter” on page 157 .
timestamp_format	Sets the format for timestamps retrieved from the database. See “UltraLite timestamp_format creation parameter” on page 159 .
timestamp_increment	Determines how the timestamp is truncated in UltraLite. See “UltraLite timestamp_increment creation parameter” on page 162 .

Name	Description
time-stamp_with_time_zone_format	This option sets the format for <code>TIMESTAMP WITH TIME ZONE</code> values retrieved from the database. See “UltraLite timestamp_with_time_zone_format creation parameter” on page 163.
utf8_encoding	Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode. See “UltraLite character sets” on page 23 and “UltraLite utf8_encoding creation parameter” on page 164.

See also

- [“Accessing creation parameter values”](#) on page 22
- [“UltraLite Initialize Database utility \(ulinit\)”](#) on page 210
- [“UltraLite Load XML to Database utility \(uload\)”](#) on page 217

UltraLite case creation parameter

Sets the case sensitivity of string comparisons in the UltraLite database. Pass in **case=respect** to the creation string parameter of the `CreateDatabase` method in your programming interface (or **case=ignore** for a case-insensitive database).

Syntax

```
ulinit --case=value database.udb
```

Allowed values

Ignore, Respect

Default

Ignore

Remarks

The case sensitivity of data is reflected in tables, indexes, and so on. By default, UltraLite databases perform case-insensitive comparisons, although data is always held in the case in which you enter it. Identifiers (such as table and column names) and user IDs are always case insensitive, regardless of the database case sensitivity. Passwords are always case sensitive, regardless of the case sensitivity of the database.

The results of comparisons on strings, and the sort order of strings, depend in part on the case sensitivity of the database.

There are some collations where particular care is required when assuming case insensitivity of identifiers. In particular, Turkish collations have a case-conversion behavior that can cause unexpected and subtle errors. The most common error is that a system object containing a letter `i` or `I` is not found.

You cannot change the case of an existing database. Instead, you must create a new database.

From Sybase Central, you can set the case sensitivity in any wizard that creates a database. On the **New Database Collation And Character Set** page, click the **Use Case-sensitive String Comparisons** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“Strings in UltraLite” on page 260](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

UltraLite checksum_level creation parameter

Sets the level of checksum validation for the database.

Syntax

```
ulinit --checksum_level=value
```

Allowed values

0, 1, 2

Default

0

Remarks

Checksums are used to detect offline corruption on pages stored to disk, flash, or memory, which can help reduce the chances of other data being corrupted as the result of a bad critical page. Depending on the level you choose, UltraLite calculates and records a checksum for each database page before it writes the page to storage.

If the calculated checksum does not match the stored checksum for a page read from storage, the page has been modified or became corrupted during the storage/retrieval of the page. If a checksum validation fails, when the database loads a page, UltraLite stops the database and reports a fatal error. This error cannot be corrected; you must re-create your UltraLite database and report the database failure to iAnywhere.

If you unload and reload an UltraLite database with checksums enabled, the checksum level is preserved and restored.

The following values are supported for the checksum_level:

- **0** Do not add checksums to database pages.
- **1** Add checksums to important database pages, such as indexes and synchronization status pages, but not row pages.
- **2** Add checksums to all database pages.

From Sybase Central, you can configure the use of checksums in any wizard that creates a database. On the **New Database Storage Settings** page, click the **Checksum Level For Database Pages** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- checksum_level database property: “UltraLite database properties” on page 188
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “UltraLite performance tips” on page 455
- “UltraLite page_size creation parameter” on page 153
- “UltraLite database connections” on page 29

UltraLite collation creation parameter

Sets the database collation.

Syntax

```
ulinit --collation=value
```

Allowed values

String

Default

1252Latin1

Remarks

You can also view a list of supported collations in UltraLite by executing the following command:

```
ulinit -Z
```

From Sybase Central, you can set the collation in any wizard that creates a database. On the **New Database Collation And Character Set** page, click either the default collation (1252Latin1), or an alternate one from the list.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite supported collations” on page 24](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“UltraLite character sets” on page 23](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

UltraLite date_format creation parameter

Sets the format for dates retrieved from the database.

For Android smartphones, you can use `Connection.setOption(OPTION_DATE_FORMAT, value)` as an alternative to setting this creation parameter. See [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#) and [“Connection.OPTION_DATE_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Syntax

```
ulinit --date_format=value
```

Allowed values

String

Default

YYYY-MM-DD (this corresponds to ISO date format specifications)

Remarks

DATE data type values are represented in a format set by the `date_format` creation parameter. Date values can, however, also be represented by strings. Before the value can be retrieved, it must be assigned to a string.

UltraLite builds a date from date parts. Date parts can include the year, the month, the day of the month, the day of the week, the day of the year, the hour, the minute, and the second (and parts thereof).

ISO (YYYY-MM-DD) is the default date format and order. For example, "7th of January 2006" in this international format is written: 2006-01-07. If you do not want to use the default ISO date format and order, you must specify a different format and order for these date parts.

The format is a string using the following symbols:

Symbol	Description
<i>yy</i>	Two digit year.
<i>yyyy</i>	Four digit year.
<i>mm</i>	Two digit month, or two digit minutes if following a colon (as in <i>hh:mm</i>).
<i>mmm</i> [<i>m...</i>]	Character short form for months—as many characters as there are "m"s. An uppercase M causes the output to be made uppercase.
<i>d</i>	Single digit day of week, (0 = Sunday, 6 = Saturday).
<i>dd</i>	Two digit day of month. A leading zero is not required.
<i>ddd</i> [<i>d...</i>]	Character short form for day of the week. An uppercase D causes the output to be made uppercase.
<i>hh</i>	Two digit hours. A leading zero is not required.
<i>nn</i>	Two digit minutes. A leading zero is not required.
<i>ss</i> [<i>.ss..</i>]	Seconds and parts of a second.
<i>aa</i>	Use 12 hour clock. Indicate times before noon with AM.
<i>pp</i>	Use 12 hour clock. Indicate times after noon with PM.
<i>jjj</i>	Day of the year, from 1 to 366.

You cannot change the date format of an existing database. Instead, you must create a new database.

Allowed values are constructed from the symbols listed in the table above. Each symbol is substituted with the appropriate data for the date that is being formatted.

For the character short forms, the number of letters specified is counted. The A.M. or P.M. indicator (which could be localized) is also truncated, if necessary, to the number of bytes corresponding to the number of characters specified.

Controlling output case For symbols that represent character data (such as *mmm*), you can control the case of the output as follows:

- Type the symbol in uppercase to have the format appear in uppercase. For example, MMM produces JAN.

- Type the symbol in lowercase to have the format appear in lowercase. For example, mmm produces jan.
- Type the symbol in mixed case to have UltraLite choose the appropriate case for the language that is being used. For example, in English, typing Mmm produces May, while in French it produces mai.

Controlling zero-padding For symbols that represent numeric data, you can control zero-padding with the case of the symbols:

- Type the symbol in same-case (such as MM or mm) to allow zero padding. For example, yyyy/mm/dd could produce 2002/01/01.
- Type the symbol in mixed case (such as Mm) to suppress zero padding. For example, yyyy/Mm/Dd could produce 2002/1/1.

From Sybase Central, you can set the date format in any wizard that creates a database. On the **New database creation parameters** page, click the **Date Format** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite date_order creation parameter” on page 147](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

Example

The following table illustrates date_format settings, together with the output from a SELECT CURRENT DATE statement, executed on Thursday May 21, 2001.

date_format syntax used	Result returned
yyyy/mm/ddd/ddd	2001/05/21/thu
jjj	141
mmm yyyy	may 2001
mm-yyyy	05-2001

UltraLite date_order creation parameter

Controls the interpretation of date formats.

For Android smartphones, you can use `Connection.setOption(OPTION_DATE_ORDER, value)` as an alternative to setting this creation parameter. See “[Connection.setOption method \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*] and “[Connection.OPTION_DATE_ORDER variable \[BlackBerry\] \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*].

Syntax

```
ulinit --date_order=value
```

Allowed values

MDY, YMD, DMY

Default

YMD (this corresponds to ISO date format specifications)

Remarks

DATE data type values are represented in a format set by the `date_format` creation parameter. Date values can, however, also be represented by strings. Before the value can be retrieved, it must be assigned to a string.

UltraLite builds a date from date parts. Date parts can include the year, the month, the day of the month, the day of the week, the day of the year, the hour, the minute, and the second (and parts thereof).

ISO (YYYY-MM-DD) is the default date format and order. For example, "7th of January 2006" in this international format is written: 2006-01-07. If you do not want to use the default ISO date format and order, you must specify a different format and order for these date parts.

You cannot change the date order of an existing database. Instead, you must create a new database.

From Sybase Central, you can set the date order in any wizard that creates a database. On the **New database creation parameters** page, click the **Date Order** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- “UltraLite date_format creation parameter” on page 144
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Accessing creation parameter values” on page 22

Example

Different values determine how the date of 10/11/12 is translated:

Syntax used	Translation
MDY	Oct 11 1912
YMD	Nov 12 1910
DMY	Nov 10 1912

UltraLite fips creation parameter

Controls whether the new database should be encrypted using AES or AES_FIPS strong encryption. This parameter is not supported by UltraLiteJ, or UltraLite for iPhone.

Syntax

```
{ulinit -a | ulload -c } --fips=value --key=value
```

Allowed values

Yes (use AES_FIPS), No (use AES)

Default

Yes

Remarks

The only way to change the type of database encryption is to recreate the database with the appropriate fips or obfuscate creation parameter. You can change the database encryption key by specifying a new encryption key on the Connection object. Users connecting to the database must supply the key each time they connect.

From Sybase Central, you can configure encryption in any wizard that creates a database. On the **New Database Storage Settings** page, click the **AES FIPS Algorithm** option. You must also set and confirm the encryption key.

Set this parameter as one of the creation parameters for the create database method on the database manager class.

To deploy a database with the fips option turned on, copy all appropriate libraries for your platform.

See also

- “Strong encryption” [*SQL Anywhere Server - Database Administration*]
- “UltraLite database security” on page 26
- “UltraLite obfuscate creation parameter” on page 152
- “UltraLite DBKEY connection parameter” on page 176
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULChangeEncryptionKey method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “ULConnection.ChangeEncryptionKey method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Connection.changeEncryptionKey method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Accessing creation parameter values” on page 22

UltraLite max_hash_size creation parameter

Sets the maximum default index hash size in bytes.

Syntax

```
ulinit --max_hash_size=value
```

Allowed values

0 to 32 bytes

Default

4 bytes

Remarks

A hash is an optional part of an index entry that is stored in the index page. The hash transforms the actual row values for the indexed columns into a numerical equivalent (a key), while still preserving ordering for

that index. The size of the key, and how much of the actual value UltraLite hashes, is determined by the hash size you set.

A row ID allows UltraLite to locate the row for the actual data in the table. A row ID is always part of an index entry. If you set the hash size to 0 (disable index hashing), then the index entry only contains this row ID. For all other hash sizes, the hash key, which can contain all or part of the transformed data in that row, is stored along with the row ID in the index page. You can improve query performance on these indexed columns because UltraLite may not always need to find, load, and unpack data before it can compare actual row values.

Determining an appropriate default database hash size requires that you evaluate the trade-off between query efficiency and database size: the higher the maximum hash value, the larger the database size grows.

UltraLite only uses as many bytes as required for the data type(s) of the column(s), up to the maximum value specified by this parameter. The default hash size is only used if you do not set a size when you create the index. If you set the default hash size to 0, UltraLite does not hash row values.

You cannot change the hash size for an existing index. When creating a new index, you can override the default value with the UltraLite **Create Index Wizard** in Sybase Central, or with the WITH MAX SIZE clause of a CREATE INDEX or a CREATE TABLE statement.

If you declare your columns as DOUBLE, FLOAT, or REAL, no hashing is used. The hash size is always ignored.

From Sybase Central, you can set the maximum hash size in any wizard that creates a database. On the **New Database Storage Settings** page, click the **Maximum Hash Size For Indexes** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite performance tips” on page 455](#)
- [“UltraLite indexes” on page 54](#)
- [“Choose and set the maximum hash size” on page 459](#)
- [“CREATE INDEX statement \[UltraLite\]” on page 421](#)
- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

UltraLite nearest_century creation parameter

Controls the interpretation of two-digit years in string-to-date conversions

For Android smartphones, you can use `Connection.setOption(OPTION_NEAREST_CENTURY, value)` as an alternative to setting this creation parameter. See “[Connection.setOption method \[UltraLiteJ\]](#)” [[UltraLite - Java Programming](#)] and “[Connection.OPTION_NEAREST_CENTURY variable \[BlackBerry\] \[UltraLiteJ\]](#)” [[UltraLite - Java Programming](#)].

Syntax

```
ulinit --nearest_century=value
```

Allowed values

Integer, between 0 and 100, inclusive

Default

50

Remarks

UltraLite automatically converts a string into a date when a date value is expected, even if the year is represented in the string by only two digits. For a two-digit date, you need to set the appropriate rollover value. Two digit years less than the value are converted to 20yy, while years greater than or equal to the value are converted to 19yy.

Choosing an appropriate rollover value typically is determined by:

- **The use of two-digit dates** Otherwise, nearest century conversion isn't applicable. Two-digit years less than the nearest_century value you set are converted to 20yy, while years greater than or equal to the value are converted to 19yy.

It is recommended that you store four-digit dates to avoid issues with incorrect conversions.

- **Consolidated database compatibility** For example, the historical SQL Anywhere behavior is to add 1900 to the year. Adaptive Server Enterprise behavior is to use the nearest century, so for any year where value yy is less than 50, the year is set to 20yy.
- **What the date represents: past event or future event** Birth years are typically those that would require a lower rollover value since they occur in the past. So for any year where yy is less than 20, the year should be set to 20yy. However, if the date is used as an expiry date, then having a higher value would be a logical choice, since the date is occurring in the future.

You cannot change the nearest century of an existing database. Instead, you must create a new database.

From Sybase Central, you can configure the nearest century setting in any wizard that creates a database. On the **New database creation parameters** page, click the **Nearest Century** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Accessing creation parameter values” on page 22

UltraLite obfuscate creation parameter

Controls obfuscation of data in the database. Obfuscation is a form of simple encryption.

Syntax

```
{ulinit -a | ulload -c } obfuscate=value
```

Allowed values

Boolean.

Default

0 (databases are not obfuscated)

Remarks

Simple encryption is equivalent to obfuscation and makes it more difficult for someone using a disk utility to look at the file to decipher the data in your database. Simple encryption does not require a key to encrypt the database.

If you want to make the database inaccessible without the correct encryption key, you must use strong encryption.

From Sybase Central, you can set obfuscation in any wizard that creates a database. On the **New Database Storage Settings** page, click the **Use Simple Encryption (Obfuscation)** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- “UltraLite fips creation parameter” on page 148
- “Using simple encryption” [*SQL Anywhere Server - Database Administration*]
- “UltraLite database security” on page 26
- “UltraLite fips creation parameter” on page 148
- “UltraLite DBKEY connection parameter” on page 176
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Accessing creation parameter values” on page 22

UltraLite page_size creation parameter

Defines the database page size in kilobytes.

Syntax

```
{ulinit | ulload } -c --page_size=sizeK
```

Allowed values

1K, 2K, 4K, 8K, 16K

Default

4K

Remarks

The page size must be entered with a K or k after the digit, or, alternatively, the equivalent number of bytes (1024, 2048, 4096, 8192, or 16384).

UltraLite databases are stored in pages, and all I/O operations are carried out a page at a time. The page size you choose can affect the performance or size of the database.

If you use any value other than those listed, the size is changed to the next larger size. If you do not specify a unit, bytes are assumed.

If your platform has limited dynamic memory, consider using a smaller page size to limit the effect on synchronization memory requirements.

When choosing a page size, you should keep the following guidelines in mind:

- **Database size** Larger databases usually benefit from a larger page size. Larger pages hold more information and therefore use space more effectively—particularly if you insert rows that are slightly more than half a page in size. The larger the page, the less page swapping that is required.

- **Number of rows** Because a row (excluding BLOBs) must fit on a page, the page size determines how large the largest packed row can be, and how many rows you can store on each page. Sometimes reading one page to obtain the values of one row may have the side effect of loading the contents of the next few rows into memory.
- **Query types** In general, smaller page sizes are likely to benefit queries that retrieve a relatively small number of rows from random locations. By contrast, larger pages tend to benefit queries that perform sequential table scans.
- **Cache size** Large page sizes may require larger cache sizes. With dynamic cache sizing, UltraLite grows the cache as required.
- **Index entries** Page size also affects indexes. The larger the database page, the more index entries it can hold.
- **Device memory** Small pages are particularly useful if your database must run on small devices with limited memory. For example, 1 MB of memory can hold 1000 pages that are each 1 KB in size, but only 250 pages that are 4 KB in size.

You cannot change the page size of an existing database. Instead, you must create a new database.

From Sybase Central, you can set the page size in any wizard that creates a database. On the **New Database Storage Settings** page, click the appropriate byte value.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“Row packing and table definitions” on page 46](#)
- [“UltraLite indexes” on page 54](#)
- [“UltraLite case creation parameter” on page 141](#)
- [“UltraLite CACHE_SIZE connection parameter” on page 169](#)
- [“UltraLite CACHE_MIN_SIZE connection parameter” on page 168](#)
- [“UltraLite CACHE_MAX_SIZE connection parameter” on page 167](#)
- [“UltraLite RESERVE_SIZE connection parameter” on page 184](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“Accessing creation parameter values” on page 22](#)

Example

To set the page size of the database to 8 KB, specify **page_size=8k** or **page_size=8192**:

```
ulinit test.udb -a --page_size=8k
```

UltraLite precision creation parameter

Specifies the maximum number of digits in decimal point arithmetic results.

For Android smartphones, you can use `Connection.setOption(OPTION_PRECISION, value)` as an alternative to setting this creation parameter. See “[Connection.setOption method \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*] and “[Connection.OPTION_PRECISION variable \[BlackBerry\] \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*].

Syntax

```
ulinit precision=value
```

Allowed values

Integer, between 1 and 127, inclusive

Default

30

Remarks

The position of the decimal point is determined by the precision and the scale of the number: precision is the total number of digits to the left and right of the decimal point; scale is the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision.

Choosing an appropriate decimal point position is typically determined by:

- **The type of arithmetic procedures you perform** Multiplication, division, addition, subtraction, and aggregate functions can all have results that exceed the maximum precision.

For example, when a `DECIMAL(8,2)` is multiplied with a `DECIMAL(9,2)`, the result could require a `DECIMAL(17,4)`. If precision is 15, only 15 digits are kept in the result. If scale is 4, the result is a `DECIMAL(15,4)`. If scale is 2, the result is a `DECIMAL(15,2)`. In both cases, there is a possibility of an overflow error.

- **The relationship between scale and precision values** The scale sets the number of digits in the fractional part of the number, and cannot be negative or greater than the precision.

You cannot change the precision of an existing database. Instead, you must create a new database.

If you are using an Oracle database as the consolidated database, all UltraLite remotes and the Oracle consolidated database must have the same precision value.

From Sybase Central, you can set the precision in any wizard that creates a database. On the **New database creation parameters** page, click the **Precision** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Initialize Database utility (ulinit)” on page 210
- “UltraLite Load XML to Database utility (ulload)” on page 217
- “ULCreateDatabase method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.CreateDatabase method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseManager.createDatabase method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “Accessing creation parameter values” on page 22

UltraLite scale creation parameter

Specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision.

For Android smartphones, you can use `Connection.setOption(OPTION_SCALE, value)` as an alternative to setting this creation parameter. See “[Connection.setOption method \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*] and “[Connection.OPTION_SCALE variable \[BlackBerry\] \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*].

Syntax

```
ulinit --scale=value
```

Allowed values

Integer, between 0 and 127, inclusive

Default

6

Remarks

The position of the decimal point is determined by the precision and the scale of the number: precision is the total number of digits to the left and right of the decimal point; scale is the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision.

Choosing an appropriate decimal point position is typically determined by:

- **The type of arithmetic procedures you perform** Multiplication, division, addition, subtraction, and aggregate functions can all have results that exceed the maximum precision.

For example, when a `DECIMAL(8,2)` is multiplied with a `DECIMAL(9,2)`, the result could require a `DECIMAL(17,4)`. If precision is 15, only 15 digits are kept in the result. If scale is 4, the result is a `DECIMAL(15,4)`. If scale is 2, the result is a `DECIMAL(15,2)`. In both cases, there is a possibility of an overflow error.

- **The relationship between scale and precision values** The scale sets the number of digits in the fractional part of the number, and cannot be negative or greater than the precision.

You cannot change the scale of an existing database. Instead, you must create a new database.

From Sybase Central, you can set the scale in any wizard that creates a database. On the **New database creation parameters** page, click the **Scale** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite precision creation parameter” on page 155](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(uload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

Example

When a DECIMAL(8,2) is multiplied with a DECIMAL(9,2), the result could require a DECIMAL(17,4). If precision is 15, only 15 digits are kept in the result. If scale is 4, the result is DECIMAL(15,4). If scale is 2, the result is a DECIMAL(15,2). In both cases, there is a possibility of overflow.

UltraLite time_format creation parameter

Sets the format for times retrieved from the database.

For Android smartphones, you can use `Connection.setOption(OPTION_TIME_FORMAT, value)` as an alternative to setting this creation parameter. See [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#) and [“Connection.OPTION_TIME_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Syntax

```
ulinit --time_format=value
```

Allowed values

String (composed of the symbols listed below)

Default

```
HH:NN:SS.sss
```

Remarks

UltraLite writes times from time parts you set with the `time_format` creation parameter. Time parts can include hours, minutes, seconds, and milliseconds.

Time values can also be represented by strings. Before a time value can be retrieved, it must be assigned to a string variable.

ISO (HH:MM:SS) is the default time format. For example, "midnight" in this international format is written: 00:00:00. If you do not want to use the default ISO time format, you must specify a different format and order for these time parts.

The format is a string using the following symbols:

Symbol	Description
<i>HH</i>	Two digit hours (24 hour clock).
<i>NN</i>	Two digit minutes.
<i>MM</i>	Two digit minutes if following a colon (as in <i>hh:mm</i>).
<i>SS[.s...]</i>	Two digit seconds plus optional fraction.

You cannot change the time format of an existing database. Instead, you must create a new database.

Each symbol is substituted with the appropriate data for the time that is being formatted. Any format symbol that represents character rather than digit output can be put in uppercase, which causes the substituted characters to be in uppercase. For numbers, using mixed case in the format string suppresses leading zeros.

You can control zero-padding with the case of the symbols:

- Type the symbol in same-case (such as HH or hh) to allow zero padding. For example, HH:NN:SS could produce 01:01:01.
- Type the symbol in mixed case (such as Hh or hH) to suppress zero padding. For example, Hh:Nn:Ss could produce 1:1:1.

From Sybase Central, you can set the time format in any wizard that creates a database. On the **New database creation parameters** page, click the **Time Format** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite timestamp_format creation parameter”](#) on page 159
- [“UltraLite Initialize Database utility \(ulinit\)”](#) on page 210
- [“UltraLite Load XML to Database utility \(uload\)”](#) on page 217
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]”](#) [*UltraLite - C and C++ Programming*]
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]”](#) [*UltraLite - C and C++ Programming*]
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]”](#) [*UltraLite - .NET Programming*]
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]”](#) [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- [“Accessing creation parameter values”](#) on page 22

Example

If a transaction was executed at 3:30 P.M. and you used the default time_format syntax of *HH:NN:SS.sss*, the result would be:

15:30:55.0

UltraLite timestamp_format creation parameter

Sets the format for timestamps that are retrieved from the database.

For Android smartphones, you can use `Connection.setOption(OPTION_TIMESTAMP_FORMAT, value)` as an alternative to setting this creation parameter. See [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*] and [“Connection.OPTION_TIME_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*].

Syntax

```
ulinit --timestamp_format=value
```

Allowed values

String

Default

YYYY-MM-DD HH:NN:SS.SSS

Remarks

UltraLite creates a timestamp from date and time parts that you set with the `date_format` and `time_format` creation parameters. Together, date and time total seven parts (year, month, day, hour, minute, second, and millisecond).

Timestamp values can also be represented by strings. Before it can be retrieved, a timestamp value must be assigned to a string variable.

Typically timestamp columns ensure that data integrity is maintained when synchronizing with a consolidated database. Timestamps help identify when concurrent data updates have occurred among multiple remote databases by tracking the last time that each user synchronized.

Tip

Ensure that the consolidated database and the UltraLite remote database maintain timestamps and timestamp increments to the same resolution. By setting these creation parameters to match that of the consolidated database, you can help avoid spurious inequalities.

The format is a string using the following symbols:

Symbol	Description
<i>YY</i>	Two digit year.
<i>YYYY</i>	Four digit year.
<i>MM</i>	Two digit month, or two digit minutes if following a colon (as in <i>hh:mm</i>).
<i>MMM[m...]</i>	Character short form for months—as many characters as there are "m"s. An uppercase M causes the output to be made uppercase.
<i>D</i>	Single digit day of week, (0 = Sunday, 6 = Saturday).
<i>DD</i>	Two digit day of month. A leading zero is not required.
<i>DDD[d...]</i>	Character short form for day of the week. An uppercase D causes the output to be made uppercase.
<i>HH</i>	Two digit hours. A leading zero is not required.
<i>NN</i>	Two digit minutes. A leading zero is not required.
<i>SS[.ss.]</i>	Seconds and parts of a second.
<i>AA</i>	Use 12 hour clock. Indicate times before noon with AM.
<i>PP</i>	Use 12 hour clock. Indicate times after noon with PM.
<i>JJJ</i>	Day of the year, from 1 to 366.

You cannot change the timestamp format of an existing database. Instead, you must create a new database.

Allowed values are constructed from the symbols listed in the table above. Each symbol is substituted with the appropriate data for the date that is being formatted.

For the character short forms, the number of letters specified is counted. The A.M. or P.M. indicator (which could be localized) is also truncated, if necessary, to the number of bytes corresponding to the number of characters specified.

For symbols that represent character data (such as *mmm*), you can control the case of the output as follows:

- Type the symbol in all uppercase to have the format appear in all uppercase. For example, MMM produces JAN.
- Type the symbol in all lowercase to have the format appear in all lowercase. For example, mmm produces jan.
- Type the symbol in mixed case to have UltraLite choose the appropriate case for the language that is being used. For example, in English, typing Mmm produces May, while in French it produces mai.

For symbols that represent numeric data, you can control zero-padding with the case of the symbols:

- Type the symbol in same-case (such as MM or mm) to allow zero padding. For example, yyyy/mm/dd could produce 2002/01/01.
- Type the symbol in mixed case (such as Mm) to suppress zero padding. For example, yyyy/Mm/Dd could produce 2002/1/1.

From Sybase Central, you can set the timestamp format in any wizard that creates a database. On the **New database creation parameters** page, click the **Timestamp Format** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite timestamp_increment creation parameter” on page 162](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)
- [“Implementing timestamp-based downloads” \[MobiLink - Server Administration\]](#)
- [“UltraLite concurrency” on page 468](#)

Example

If a transaction was executed on Friday May 12, 2006 at 3:30 PM and you used the default timestamp_format syntax of *YYYY-MM-DD HH:NN:SS.SSS*, the result would be:

2006-05-12 15:30:55.0

UltraLite timestamp_increment creation parameter

Limits the resolution of timestamp values. As timestamps are inserted into the database, UltraLite truncates them to match this increment.

For Android smartphones, you can use `Connection.setOption(OPTION_TIMESTAMP_INCREMENT, value)` as an alternative to setting this creation parameter. See [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#) and [“Connection.OPTION_TIMESTAMP_INCREMENT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Syntax

```
ulinit timestamp_increment=value
```

Allowed values

1 to 60000000 microseconds

Default

1 microsecond

Remarks

1000000 microseconds equals 1 second.

You cannot change the timestamp increment of an existing database. Instead, you must create a new database.

This increment is useful when a DEFAULT TIMESTAMP column is being used as a primary key or row identifier.

From Sybase Central, you can set the timestamp increment in any wizard that creates a database. On the **New database creation parameters** page, click the **Timestamp Increment** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite timestamp_format creation parameter” on page 159](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULCreateDatabase method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)
- [“Implementing timestamp-based downloads” \[MobiLink - Server Administration\]](#)
- [“UltraLite concurrency” on page 468](#)

Example

To store a value such as '2000/12/05 10:50:53:700', set this creation parameter to 100000. This value truncates the timestamp after the first decimal place in the seconds component.

UltraLite timestamp_with_time_zone_format creation parameter

Sets the format for `TIMESTAMP WITH TIME ZONE` values retrieved from the database.

Syntax

```
ulinit --timestamp_with_time_zone_format=value
```

Allowed values

String (composed of the symbols listed below)

Default

YYYY-MM-DD HH:NN:SS.SSS+HH:NN

Remarks

The format is a string using the following symbols:

Symbol	Description
YY	Two digit year
YYYY	Four digit year
MM	Two digit month, or two digit minutes if following a colon (as in 'HH:MM')
MMM[m...]	Character short form for months—as many characters as there are "m"s
DD	Two digit day of month
DDD[d...]	Character short form for day of the week
HH	Two digit hours
NN	Two digit minutes
SS.SSSSSS	Seconds and fractions of a second, up to six decimal places. Not all platforms support timestamps to a precision of six places.
AA	A.M. or P.M. (12 hour clock)—omit AA and PP for 24 hour time

Symbol	Description
PP	P.M. if needed (12 hour clock)—omit AA and PP for 24 hour time
HH	Two digit hours (time zone offset)
NN	Two digit minutes (time zone offset)

Each symbol is substituted with the appropriate data for the date that is being formatted.

For symbols that represent character data (such as *MMM*), you can control the case of the output as follows:

- Type the symbol in all uppercase to have the format appear in all uppercase. For example, *MMM* produces JAN.
- Type the symbol in all lowercase to have the format appear in all lowercase. For example, *mmm* produces jan.
- Type the symbol in mixed case to have UltraLite choose the appropriate case for the language that is being used. For example, in English, typing *Mmm* produces May, while in French it produces mai.

If the character data is multibyte, the length of each symbol reflects the number of characters, not the number of bytes. For example, the *mmm* symbol specifies a length of three characters for the month.

For symbols that represent numeric data, you can control zero-padding with the case of the symbols:

- Type the symbol in same-case (such as *MM* or *mm*) to allow zero padding. For example, *yyyy/mm/dd* could produce 2002/01/01.
- Type the symbol in mixed case (such as *Mm*) to suppress zero padding. For example, *yyyy/Mm/Dd* could produce 2002/1/1.

Note

If you change the setting for `timestamp_with_time_zone_format` option in a way that re-orders the date format, be sure to change the `date_order` option to reflect the same change, and vice versa. See “[date_order option](#)” [*SQL Anywhere Server - Database Administration*].

See also

- “[TIMESTAMP WITH TIME ZONE data type](#)” [*SQL Anywhere Server - SQL Reference*]

UltraLite utf8_encoding creation parameter

Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode.

Syntax

```
ulload -c --utf8_encoding
```

Values

Boolean.

Default

1 (databases are UTF-8 encoded)

Remarks

UTF-8 characters are represented by one to four bytes. For other multibyte collations, one or two bytes are used. For all provided multibyte collations, characters of two or more bytes are considered to be alphabetic. This means that you can use these characters in identifiers without requiring double quotes.

Characters in an UltraLite database are either from the codepage implicit in the chosen collation, or are UTF8 encoded. UltraLite databases that use the UTF8BIN collation are automatically UTF8 encoded. If the operating system to which you are deploying your UltraLite application uses UTF8 or Unicode (like most Linux distributions, Windows Mobile and iPhone) or if you plan to store characters from multiple languages in your database, you should create your database using a UTF8 encoding. If you try synchronizing UTF-8 encoded characters into a consolidated table that does not support Unicode, a user error is reported.

From Sybase Central, you can choose UTF-8 encoding in any wizard that creates a database. On the **New database collation and character set** page, click the **Yes, use UTF-8 as the database character set** option.

From a client application, set this parameter as one of the creation parameters for the create database method on the database manager class.

See also

- [“UltraLite platform requirements for character set encoding” on page 24](#)
- [“UltraLite character sets” on page 23](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 217](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.CreateDatabase method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseManager.createDatabase method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Accessing creation parameter values” on page 22](#)

UltraLite connection parameters

This section describes the UltraLite connection parameters that are available when you connect to an UltraLite database.

Members

UltraLite supports the following connection parameters:

Parameter name	Description
CACHE_MAX_SIZE, CACHE_MIN_SIZE, CACHE_SIZE	Defines the size of the database cache. See “UltraLite CACHE_SIZE connection parameter” on page 169, “UltraLite CACHE_MIN_SIZE connection parameter” on page 168, and “UltraLite CACHE_MAX_SIZE connection parameter” on page 167.
COMMIT_FLUSH	Determines when committed transactions are flushed to storage after a commit call. See “UltraLite COMMIT_FLUSH connection parameter” on page 172.
CON	Specifies a name of the current connection. See “UltraLite CON connection parameter” on page 173.
CE_FILE, DBF, desktop, device, and NT_FILE	<p>At creation time, these parameters set the location of the database.</p> <p>For subsequent connections, they tell UltraLite where to find the file.</p> <p>You can use DBF if you are creating a single-platform application or are connecting to an UltraLite administration tool. Use the other platform-specific versions if you are programming an UltraLite client that connects to different platform-specific databases. See:</p> <ul style="list-style-type: none"> • “UltraLite DBF connection parameter” on page 174 • “UltraLite CE_FILE connection parameter” on page 171 • “UltraLite desktop connection parameter” on page 177 • “UltraLite device connection parameter” on page 179 • “UltraLite NT_FILE connection parameter” on page 181
DBKEY	<p>At creation-time, this parameter sets the encryption key used to encrypt the database.</p> <p>For subsequent connections, it specifies the encryption key used to encrypt the database.</p> <p>See “UltraLite DBKEY connection parameter” on page 176.</p>
DBN	Identifies a running database by name rather than file name. See “UltraLite DBN connection parameter” on page 176.
MIRROR_FILE	Specifies the name of a database mirror file. See “UltraLite MIRROR_FILE connection parameter” on page 180.
PWD	<p>At creation-time, this parameter sets the initial password for a user.</p> <p>For subsequent connections, it supplies the password for the user ID.</p> <p>See “UltraLite PWD connection parameter” on page 183.</p>

Parameter name	Description
RESERVE_SIZE	Pre-allocates the file system space required for your UltraLite database without actually inserting any data. See “UltraLite RESERVE_SIZE connection parameter” on page 184.
START	Specifies the location of the UltraLite engine executable. See “UltraLite START connection parameter” on page 185.
TEMP_DIR	Specifies the name of the directory (which must already exist) into which UltraLite will place a temporary file (with a name derived from the database name). See “UltraLite TEMP_DIR connection parameter” on page 186.
UID	At creation time, this parameter sets the initial user ID. For subsequent connections, it identifies a user to the database. The user ID must be one of up to four user IDs stored in the UltraLite database. See “UltraLite UID connection parameter” on page 187.

See also

- [“UltraLite connection strings and parameters”](#) on page 29

UltraLite **CACHE_MAX_SIZE** connection parameter

Defines the maximum size of the database cache. UltraLite manages the cache size automatically, so setting this parameter should not be necessary.

Syntax

CACHE_MAX_SIZE=*number*{ **k** | **m** }

Default

The default maximum cache size is 20 MB for devices and 50 MB for desktops.

Remarks

The `cache_max_size` connection parameter specifies the maximum amount of memory to allocate for the file cache. By default, the size is in bytes. Use `k` or `m` to specify units of kilobytes or megabytes.

If you exceed the maximum cache size, your platform's upper cache size limit is used instead. Note that UltraLite does not grow the cache size beyond the actual file size of the database.

Any leading or trailing spaces in connection parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- “UltraLite CACHE_MIN_SIZE connection parameter” on page 168
- “UltraLite CACHE_SIZE connection parameter” on page 169
- “UltraLite cache_allocation option” on page 192
- “UltraLite connection strings and parameters” on page 29
- “UltraLite page_size creation parameter” on page 153
- “UltraLite RESERVE_SIZE connection parameter” on page 184
- UltraLite for C/C++: “Connecting to an UltraLite database” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- UltraLite.NET: “Connecting to an UltraLite database” [*UltraLite - .NET Programming*]
- “ULConnection.Open method [UltraLite.NET]” [*UltraLite - .NET Programming*]

Example

The following connection string fragment sets the maximum cache size to 100 MB.

```
"CACHE_MAX_SIZE=100m"
```

UltraLite CACHE_MIN_SIZE connection parameter

Defines the minimum size of the database cache. UltraLite manages the cache size automatically, so setting this parameter should not be necessary.

Syntax

```
CACHE_MIN_SIZE=number{ k | m }
```

Default

The default cache size for devices is 256 KB. The default cache size for desktops is 512 KB.

Remarks

The cache_min_size connection parameter specifies the minimum amount of memory to allocate for the file cache. By default, the size is in bytes. Use k or m to specify units of kilobytes or megabytes.

If you set the minimum cache size to be greater than the maximum cache size, UltraLite returns an error message and the connection fails.

Any leading or trailing spaces in connection parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- “UltraLite CACHE_MAX_SIZE connection parameter” on page 167
- “UltraLite CACHE_SIZE connection parameter” on page 169
- “UltraLite cache_allocation option” on page 192
- “UltraLite connection strings and parameters” on page 29
- “UltraLite page_size creation parameter” on page 153
- “UltraLite RESERVE_SIZE connection parameter” on page 184
- UltraLite for C/C++: “Connecting to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [[UltraLite - C and C++ Programming](#)]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- UltraLite.NET: “Connecting to an UltraLite database” [[UltraLite - .NET Programming](#)]
- “ULConnection.Open method [UltraLite.NET]” [[UltraLite - .NET Programming](#)]

Example

The following connection string fragment sets the minimum cache size to 1 MB.

```
"CACHE_MIN_SIZE=1m"
```

UltraLite CACHE_SIZE connection parameter

Defines the initial size of the database cache. UltraLite manages the cache size automatically, so setting this parameter should not be necessary.

For Android smartphones, you can use Configuration.setPageSize as an alternative to setting this connection parameter. See “[Configuration.setPageSize method \[UltraLiteJ\]](#)” [[UltraLite - Java Programming](#)].

Syntax

```
CACHE_SIZE=number{ k | m }
```

Default

The default initial cache size is determined by the amount of memory available on your system and the size of the database.

Remarks

The cache_size connection parameter specifies the initial amount of memory to allocate for the file cache. This cache is used to hold recently-used pages from the database file in memory so they can be accessed quickly when needed again, and also to collect multiple modifications to a page before writing it back to storage. Accessing a page from the cache is many times faster than reading from storage. Writing to storage is more expensive, so grouping multiple modifications in a single write is important for

performance. Encrypted databases also benefit from the cache because decryption occurs only when the page is loaded into the cache, and encryption occurs before the page is written back to storage. If the cache is sufficiently large, the overhead of encryption becomes negligible.

As an example of cache usage, consider synchronization. While UltraLite is receiving a download, the rows are inserted into the database, and referential integrity checks are performed. When inserted, the rows are also indexed; they are added to each index on the table. So, while synchronizing, the cache tends to hold the pages where the new rows are stored, as well as the index pages for the current table. Synchronization performance depends greatly on whether the cache is large enough to contain an appropriate working set of pages for a table being synchronized. If the cache is too small, row inserts may require repeated reads of index pages from storage, incurring a noticeable performance penalty over the case when the required index pages fit in the cache.

By default, the size is in bytes. Use k or m to specify units of kilobytes or megabytes.

If you exceed the permissible maximum cache size, it is automatically replaced with your platform's upper cache size limit.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite CACHE_MAX_SIZE connection parameter” on page 167](#)
- [“UltraLite CACHE_MIN_SIZE connection parameter” on page 168](#)
- [“UltraLite cache_allocation option” on page 192](#)
- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite page_size creation parameter” on page 153](#)
- [“UltraLite RESERVE_SIZE connection parameter” on page 184](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[*UltraLite - .NET Programming*\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)

Example

The following connection string fragment sets the cache size to 20 MB.

```
"CACHE_SIZE=20m"
```

UltraLite CE_FILE connection parameter

When creating a new UltraLite database, this connection parameter names the new database file.

When opening a connection to an existing database, it identifies the database.

For Android smartphones, you can use `Configuration.setDatabaseName` as an equivalent to setting this connection parameter. See “[Configuration.setDatabaseName method \[UltraLiteJ\]](#)” [*UltraLite - Java Programming*].

Syntax

`CE_FILE=path\ce-db`

Default

DBF connection parameter.

Behavior

1. If DBF is specified, look for a matching database (identical filename) and connect if found, proceed to auto-start if not.
2. A database is auto-started when required if DBF is specified.

Remarks

You should use the CE_FILE connection parameter for UltraLite client applications that use the same connection string to connect to a Microsoft Windows Mobile device, and other platforms.

The CE_FILE connection parameter takes precedence over the DBF connection parameter. If you are connecting from an UltraLite administration tool, or your connection object only connects to a Windows Mobile database, use the DBF connection parameter.

The value of CE_FILE must meet the file name requirements for Windows Mobile. If you include an absolute path to the database, then all directories must exist before setting the path to this file. UltraLite does not create them automatically.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- “UltraLite connection strings and parameters” on page 29
- “UltraLite file path formats in connection parameters” on page 31
- “Precedence of connection parameters for UltraLite administration tools” on page 30
- “UltraLite DBF connection parameter” on page 174
- UltraLite for C/C++: “Connecting to an UltraLite database” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- UltraLite.NET: “Connecting to an UltraLite database” [*UltraLite - .NET Programming*]
- “ULConnection.Open method [UltraLite.NET]” [*UltraLite - .NET Programming*]

Example

The following example creates a new connection and identifies different database files for the Windows desktop and Windows Mobile platforms:

```
Set Connection = DatabaseMgr.OpenConnection("DBF=d:\Dbfile.udb;CE_FILE=\myapp\nMyDB.udb")
```

UltraLite COMMIT_FLUSH connection parameter

Determines when committed transactions are flushed to storage after a commit call. If no calls to commit are made by the UltraLite application, no flush can occur.

Syntax

```
COMMIT_FLUSH={ immediate | grouped | on_checkpoint }
```

Default

immediate

Remarks

This connection parameter defines which transactions are recovered following a hardware failure or crash. You can group logical autocommit operations as a single recovery point.

By grouping these operations, you can improve UltraLite performance, but at the expense of data recoverability. There is a slight chance that a transaction may be lost—even though it has been committed—if a hardware failure or crash occurs after a commit, but before the transaction is flushed to storage.

The following parameters are supported:

- **immediate** Committed transactions are flushed to storage immediately upon a commit call before the commit operation completes.

- **grouped** Committed transactions are flushed to storage on a commit call, but only after a threshold you configure has been reached. You can configure either a transaction count threshold with the `commit_flush_count` database option or a time-based threshold with the `commit_flush_timeout` database option.

If set, both the `commit_flush_count` and the `commit_flush_timeout` options act as possible triggers for the commit flush; the first threshold that is met triggers the flush. When the flush occurs, UltraLite sets the counter and the timer back to 0. Then, both the counter and timer are monitored, until one of these thresholds is reached again.

- **on_checkpoint** Committed transactions are flushed to storage on a checkpoint operation. You can perform a checkpoint with any of the following:
 - The CHECKPOINT statement. APIs that do not have a checkpoint method must use this SQL statement.
 - The `ULCheckpoint` function for UltraLite embedded SQL.
 - The Checkpoint method on a connection object in a C++ component.

See also

- [“Flush single or grouped transactions” on page 471](#)
- [“UltraLite `commit_flush_count` option \[temporary\]” on page 193](#)
- [“UltraLite `commit_flush_timeout` option \[temporary\]” on page 194](#)
- [“CHECKPOINT statement \[UltraLite\]” on page 420](#)
- [“`ULCheckpoint` method \[UltraLite Embedded SQL\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“`ULConnection.Checkpoint` method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)

UltraLite CON connection parameter

Names a connection so that switching to it is easier in multi-connection applications.

Syntax

CON=*name*

Default

No connection name.

Remarks

The CON connection parameter is global to the application.

Do not use this connection parameter unless you are going to establish and switch between two or more concurrent connections.

The connection name is not the same as the database name.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite DBN connection parameter” on page 176](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

The following connection string fragment sets the first connection name to MyFirstCon.

```
"CON=MyFirstCon"
```

UltraLite DBF connection parameter

When creating a new UltraLite database, this connection parameter names the new database file.

When opening a connection to an existing database, it indicates which database file you want to load and connect to.

Syntax

```
DBF=ul-db
```

Behavior

1. On connect, look to see if the database is already running. If DBN is specified, look for a matching database and connect if found, proceed to auto-start if not.
2. If DBF is specified, look for a matching database (identical filename) and connect if found, proceed to auto-start if not.
3. If neither DBN nor DBF is specified, and a single database is running, connect to it.
4. A database is auto-started when required if DBF is specified. If DBN is also specified, it becomes the name of the running database, otherwise a name is generated from the base filename.

Remarks

Because they are aliases, if DBF is used concurrently, the last one specified takes precedence.

If you are connecting to multiple databases on different devices from a single connection string, you can use the following parameters to name platform-specific alternates:

- CE_FILE
- desktop
- device
- NT_FILE

If specified, these platform-specific connection parameters take precedence over DBF.

The value of DBF must meet the file name requirements for the platform.

Windows Mobile If you are deploying to a Windows Mobile device, UltraLite utilities and wizards can administer an UltraLite database on an attached Windows Mobile device. To identify a file on a Windows Mobile device, you must specify the required absolute path, and use the **wce:** prefix.

Any leading or trailing spaces in parameter values are ignored. The value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite file path formats in connection parameters” on page 31](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)
- [“UltraLite DBN connection parameter” on page 176](#)
- [“UltraLite CE_FILE connection parameter” on page 171](#)
- [“UltraLite NT_FILE connection parameter” on page 181](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[*UltraLite - .NET Programming*\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)

Examples

To connect to the database, *MyULdb.udb*, installed in the desktop directory *c:\mydb*, use the following connection string:

```
"DBF=c:\mydb\MyULdb.udb"
```

To connect to the same database that is deployed to the *UltraLite* folder of the attached Windows Mobile device, use the following connection string:

```
"DBF=wce:\UltraLite\MyULdb.udb"
```

UltraLite DBKEY connection parameter

When creating a new UltraLite database, this connection parameter provides an encryption key for the database.

When opening a connection to an existing database, it provides the encryption key for the database.

Syntax

DBKEY=*string*

Default

No key is provided.

Remarks

If you do not specify the correct encryption key for the database, the connection fails.

If a database is created using an encryption key, the database file is strongly encrypted using either the AES 256-bit or AES FIPS-certified algorithm. By using strong encryption, you have increased security against skilled and determined attempts to gain access to the data. However, the use of strong encryption has a significant performance impact.

Any leading or trailing spaces in parameter values are ignored. The value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite database security” on page 26](#)
- [“UltraLite obfuscate creation parameter” on page 152](#)
- UltraLite for C/C++: [“Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- UltraLite.NET: [“Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

UltraLite DBN connection parameter

Differentiates databases by name when applications connect to more than one database.

Syntax

DBN=*db-name*

Default

None.

Behavior

1. On connect, look to see if the database is already running. If DBN is specified, look for a matching database and connect if found, proceed to auto-start if not.
2. If DBF is specified, look for a matching database (identical filename) and connect if found, proceed to auto-start if not.
3. If neither DBN nor DBF is specified, and a single database is running, connect to it.
4. A database is auto-started when required if DBF is specified. If DBN is also specified, it becomes the name of the running database, otherwise a name is generated from the base filename.

Remarks

UltraLite sets the database name after the database has been opened. Client applications can then connect to this database via its name instead of its file.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite DBF connection parameter” on page 174](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

Use the following parameters to connect to the running UltraLite database named Kitchener:

```
DBN=Kitchener;DBF=cities.udb
```

UltraLite desktop connection parameter

When creating a new UltraLite database, this connection parameter names the new database file.

When opening a connection to an existing database, it identifies the database.

Syntax

desktop:DBF=*path\db* | **temp_dir=***\Temp*

Remarks

You should use the desktop connection parameter for UltraLite client applications that use the same connection string to connect to a Microsoft Windows or Mac device.

Colon (:) is considered to be a separator in addition to underscore (_). Options with a prefix take precedence over options without a prefix.

A **temp_dir** connection parameter is now available. This must name a directory (which already exists). UL will place the temporary file (with name still derived from the database name) in the specified directory, rather than beside the database file (the default and previous behavior). Specifying a temporary directory with faster I/O characteristics can improve the performance of things like temporary tables which are large relative to the cache size. Long-running transactions can also consume noticeable space in the temp file.

The value of desktop must meet the file name requirements for Windows or Mac. If you include an absolute path to the database, then all directories must exist before setting the path to this file. UltraLite does not create them automatically.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

Behavior

1. If DBF is specified, look for a matching database (identical filename) and connect if found, proceed to auto-start if not.
2. A database is auto-started when required if DBF is specified.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite file path formats in connection parameters” on page 31](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)
- [“UltraLite DBF connection parameter” on page 174](#)
- UltraLite for C/C++: [“Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- UltraLite.NET: [“Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

The following example creates a new connection and identifies different database files for the Windows desktop and Windows Mobile platforms:

```
"desktop:DBF=C:\dir\db.udb; device:DBF=\SD Card\db.udb; device:temp_dir=\Temp; device:cache_size=4M"
```

UltraLite device connection parameter

When creating a new UltraLite database, this connection parameter names the new database file.

When opening a connection to an existing database, it identifies the database.

Syntax

```
device:DBF=path\db | temp_dir=\Temp
```

Remarks

You should use the device connection parameter for UltraLite client applications that use the same connection string to connect to a Microsoft Windows CE or iPhone device.

Colon (:) is considered to be a separator in addition to underscore (_). Options with a prefix take precedence over options without a prefix.

A **temp_dir** connection parameter is now available. This must name a directory (which already exists). UL will place the temporary file (with name still derived from the database name) in the specified directory, rather than beside the database file (the default and previous behavior). Specifying a temporary directory with faster I/O characteristics can improve the performance of things like temporary tables which are large relative to the cache size. Long-running transactions can also consume noticeable space in the temp file.

The value of device must meet the file name requirements for Windows CE or iPhone. If you include an absolute path to the database, then all directories must exist before setting the path to this file. UltraLite does not create them automatically.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite file path formats in connection parameters” on page 31](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)
- [“UltraLite DBF connection parameter” on page 174](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[*UltraLite - .NET Programming*\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)

Example

The following example creates a new connection and identifies different database files for the Windows desktop and Windows Mobile platforms:

```
"desktop:DBF=C:\dir\db.udb; device:DBF=\SD Card\db.udb; device:temp_dir=\Temp; device:cache_size=4M"
```

UltraLite MIRROR_FILE connection parameter

Specifies the name of the database mirror file to which all database writes will be issued (at the same time as they are to the main database file).

Syntax

```
MIRROR_FILE=path\mirrorfile-db
```

Default

None.

Remarks

UltraLite provides basic database file mirroring to improve fault tolerance on potentially unreliable storage systems. This is accomplished using the mirror file. All database writes are issued to the mirror file at the same time as they are to the main database file (write overhead is therefore doubled; read overhead is not affected). If a corrupt page is read from the database file, the page is recovered by reading from the mirror file.

Mirroring is supported on all platforms using a file-based store.

When the **mirror_file=** option is specified when you start the database, UltraLite will open the named file and verify that it matches the main database file before continuing. If the mirror file does not exist, it is

created at that point by copying the main file. If the mirror is not a database file, or is corrupt, an error is reported and the database will not start until the file is removed or a different mirror is specified. If the mirror does not match the database, `SQLE_MIRROR_FILE_MISMATCH` is generated and the database will not start. When a corrupt page is recovered, the warning `SQLE_CORRUPT_PAGE_READ_RETRY` is generated. (Without mirroring, or if the mirror file is also corrupt, the error `SQLE_DEVICE_ERROR` is generated and the database is halted.)

To effectively protect against media failures, page checksums must be enabled when you use a mirror file. (With or without mirroring, page checksums allow UltraLite to detect page corruption as soon as the page is loaded and avoid referencing corrupt data.) Specify the `checksum_level` database creation option to enable checksums. UltraLite will generate the warning `SQLE_MIRROR_FILE_REQUIRES_CHECKSUMS` if checksums are not enabled when using a mirror file.

Note that because the mirror is an exact copy of the database file, it can be started directly as a database. The `ulvalid` utility will report corrupt pages.

See also

- [“UltraLite `checksum_level` creation parameter” on page 142](#)
- [“UltraLite Validate Database utility \(`ulvalid`\)” on page 230](#)
- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite file path formats in connection parameters” on page 31](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [“`ULDatabaseManager.OpenConnection` method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“`DatabaseManager.openConnection` method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[*UltraLite - .NET Programming*\]](#)
- [“`ULConnection.Open` method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)

Example

The following example creates a new connection and creates a mirror file:

```
Connection = DatabaseMgr.OpenConnection("DBF=c:\Dbfile.udb;
UID=JDoe;PWD=ULdb;
MIRROR_FILE=c:\test\MyMirrorDB.udb")
```

UltraLite `NT_FILE` connection parameter

When creating a new UltraLite database, this connection parameter names the new database file.

When opening a connection to an existing database, it identifies the database.

For Android smartphones, you can use `Configuration.setDatabaseName` as an equivalent to setting this connection parameter. See [“Configuration.setDatabaseName method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Syntax

`NT_FILE=path\nt-db`

Default

DBF connection parameter.

Remarks

You should use the `NT_FILE` connection parameter for UltraLite client applications that use the same connection string to connect to a desktop database, and a database on other platforms.

This connection parameter takes precedence over the `DBF` parameter. If you are connecting from an UltraLite administration tool, or your connection object only connects to a desktop database, use the `DBF` connection parameter.

The value of `NT_FILE` must meet the file name requirements for Windows desktop platforms.

The path can be absolute or relative. If you include a directory as part of the file name, then all directories must exist before setting the path to this file. UltraLite does not create them automatically.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite file path formats in connection parameters” on page 31](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 30](#)
- [“UltraLite DBF connection parameter” on page 174](#)
- UltraLite for C/C++: [“Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- UltraLite.NET: [“Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

The following example creates a new connection and identifies different database files for the desktop and Windows Mobile platforms:

```
Connection = DatabaseMgr.OpenConnection("UID=JDoe;PWD=ULdb;  
NT_FILE=c:\test\MyTestDB.udb;CE_FILE=\database\MyCEDB.udb")
```

UltraLite PWD connection parameter

When creating a new UltraLite database, this connection parameter sets the password for the default user.

When connecting to an existing database, it defines the password for a user ID that is used for authentication.

For Android smartphones, you can use `Configuration.setPassword` as an alternative to setting this connection parameter. See [“Configuration.setPassword method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

Syntax

PWD=*password*

Default

If you do not set both the UID and PWD, UltraLite opens connections with **UID=DBA** and **PWD=sql**.

Remarks

Every user of a database has a password. UltraLite supports up to four user ID/password combinations.

You can set passwords to NULL or an empty string.

A random 4-byte salt value is generated when a new user is created or an existing user changes their password. The salt value is appended to the user's password when calculating the password hash and is stored in the database along with the hash. Salting significantly decreases vulnerability to dictionary attacks and also ensures that users with the same password will have different password hashes.

This connection parameter is not encrypted. However, UltraLite hashes the password before saving it, so you can only modify a password using Sybase Central.

See also

- [“Users” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite users” on page 59](#)
- [“UltraLite UID connection parameter” on page 187](#)
- [UltraLite for C/C++: “User authentication” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “User authentication” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Examples

The following partial connection string supplies the user ID DBA and password sql:

```
"UID=DBA;PWD=sql"
```

The following partial connection string supplies the user ID DBA and an empty password:

```
"UID=DBA;PWD= ' ' "
```

UltraLite RESERVE_SIZE connection parameter

Pre-allocates the file system space required for your UltraLite database, without actually inserting any data. Reserving the file system space prevents the space from being used up by other files.

Syntax

```
RESERVE_SIZE= number{ k | m | g }
```

Default

0 (no reserve size).

Remarks

The value you supply can be any value from 0 to your maximum database size. Use k, m, or g to specify units of kilobytes, megabytes, or gigabytes, respectively. If you do not specify a unit, bytes are assumed by default.

You should run the database with test data and observe the database size and choose a reserve size that suits your UltraLite deployment.

If the RESERVE_SIZE value is smaller than the database size, UltraLite ignores the parameter.

Reserving file system space can improve performance slightly because it may:

- Reduce the degree of file fragmentation compared to growing incrementally.
- Prevent out-of-storage memory failures.

Because an UltraLite database consists of data and metadata, the database size grows only when required (when the application updates the database).

See also

- “UltraLite connection strings and parameters” on page 29
- “UltraLite CACHE_SIZE connection parameter” on page 169
- “UltraLite page_size creation parameter” on page 153
- UltraLite for C/C++: “Connecting to an UltraLite database” [*UltraLite - C and C++ Programming*]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- UltraLite.NET: “Connecting to an UltraLite database” [*UltraLite - .NET Programming*]
- “ULConnection.Open method [UltraLite.NET]” [*UltraLite - .NET Programming*]

Example

The following connection string fragment sets the reserve size to 128 KB so the system reserves that much system space for the database upon startup.

```
"RESERVE_SIZE=128K"
```

UltraLite START connection parameter

Starts the UltraLite engine executable. This parameter is not supported for UltraLite for Android. This parameter is only required if the engine is not in one of the expected locations.

Syntax

```
START=path\uleng12.exe
```

Remarks

Only supply a StartLine (START) connection parameter if you are connecting to an engine that is not currently running.

Paths with spaces require quotes. Otherwise, the client returns SQLE_UNABLE_TO_CONNECT_OR_START.

See also

- “Start a deployed UltraLite engine” on page 125
- “UltraLite Engine utility (uleng12)” on page 206
- “Choose an UltraLite data management component for Windows Mobile” on page 13
- UltraLite for C/C++: “Connecting to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [[UltraLite - C and C++ Programming](#)]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- UltraLite.NET: “Connecting to an UltraLite database” [[UltraLite - .NET Programming](#)]
- “ULConnection.Open method [UltraLite.NET]” [[UltraLite - .NET Programming](#)]

Example

The following command starts the UltraLite engine that is located in the *Program Files* directory:

```
Start="Program Files\uleng12.exe"
```

An alternative way to define this path is to put the entire string in single quotes:

```
Start='"Program Files\uleng12.exe"'
```

UltraLite TEMP_DIR connection parameter

Specifies the name of the directory (which must already exist) into which UltraLite will place the temporary file (with a name derived from the database name).

Syntax

```
TEMP_DIR=\path
```

Remarks

In addition to the database file, UltraLite creates and maintains a temporary file during database operation. You do not need to work with or maintain the file in any way.

By default, UltraLite maintains its temporary file in the same folder (if one exists) as the UltraLite database itself. The temporary file has the same file name as the database, but for file-based platforms the tilde is included in the extension of the file. For example, if you run the *CustDB.udb* sample database, the temporary file called *CustDB.~db* is maintained in the same directory as the database file.

Specifying a temporary directory with faster I/O characteristics can improve the performance of things like temporary tables which are large relative to the cache size. Long-running transactions can also consume noticeable space in the temp file.

Paths with spaces require quotes. Otherwise, the client returns `SQL_UNABLE_TO_CONNECT_OR_START`.

See also

- UltraLite for C/C++: “Connecting to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- “ULDatabaseManager.OpenConnection method [UltraLite C++]” [[UltraLite - C and C++ Programming](#)]
- UltraLite for embedded SQL: “Connect to an UltraLite database” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Connect to an UltraLite database” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- “DatabaseManager.openConnection method [UltraLite for M-Business Anywhere]” [[UltraLite - M-Business Anywhere Programming \(deprecated\)](#)]
- UltraLite.NET: “Connecting to an UltraLite database” [[UltraLite - .NET Programming](#)]
- “ULConnection.Open method [UltraLite.NET]” [[UltraLite - .NET Programming](#)]

Example

The following connection string fragment puts the temp file in the `\Temp` directory:

```
temp_dir=\Temp;
```

UltraLite UID connection parameter

When creating a new UltraLite database, this connection parameter sets the default user ID for the database.

When connecting to an existing database, it specifies the user ID with which you connect to the database. The value must be an authenticated user for the database.

Syntax

UID=*user*

Default

If you do not set the UID and PWD when connecting, UltraLite opens connections with **UID=DBA** and **PWD=sql**.

Remarks

Every user of a database has a user ID. UltraLite supports up to four user ID/password combinations.

UltraLite user IDs are separate from MobiLink user names and from other SQL Anywhere user IDs. You cannot change a user ID once it is created. Instead, you must delete the user ID and then add a new one.

You cannot set the UID to NULL or an empty string. The maximum length for a user ID is 31 characters. User IDs are case insensitive.

Any leading or trailing spaces in parameter values are ignored. This connection parameter's value cannot include leading single quotes, leading double quotes, or semicolons.

See also

- [“UltraLite connection strings and parameters” on page 29](#)
- [“UltraLite users” on page 59](#)
- [UltraLite for C/C++: “User authentication” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for C/C++: “Connecting to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseManager.OpenConnection method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “User authentication” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “Connect to an UltraLite database” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Connect to an UltraLite database” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“DatabaseManager.openConnection method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [UltraLite.NET: “Connecting to an UltraLite database” \[UltraLite - .NET Programming\]](#)
- [“ULConnection.Open method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

The following connection string fragment supplies the user ID DBA and password sql for a database:

```
"UID=DBA;PWD=sql"
```

UltraLite database properties

UltraLite database property values are defined when the database is first created. They can be changed by re-creating the UltraLite database or editing their corresponding database option, if available.

UltraLite supports the following database properties:

Property	Description
cache_allocation	Returns the current cache size as a percentage of the minimum and maximum settings. See “UltraLite cache_allocation option” on page 192 .
CaseSensitive	Returns the status of the case sensitivity feature. Returns On if the database is case sensitive. Otherwise, it returns Off. See “UltraLite case creation parameter” on page 141 .

Property	Description
CharSet	<p>Returns the CHAR character set of the database. The character set used by the database is determined by the database's collation sequence and whether the data is UTF-8 encoded.</p> <p>See also:</p> <ul style="list-style-type: none"> • “UltraLite utf8_encoding creation parameter” on page 164 • “UltraLite collation creation parameter” on page 143
ChecksumLevel	<p>Returns the level of checksum validation in the database, one of 0 (do not add checksums), 1 (add checksums only to important pages), or 2 (add checksums to all pages). See “UltraLite checksum_level creation parameter” on page 142.</p>
Collation	<p>Returns the name of the database's collation sequence. See “UltraLite collation creation parameter” on page 143.</p>
commit_flush_count	<p>Returns the value of the commit_flush_count option that sets a commit count threshold. See “UltraLite commit_flush_count option [temporary]” on page 193.</p>
commit_flush_timeout	<p>Returns the value of the commit_flush_timeout option that sets a time interval threshold. See “UltraLite commit_flush_timeout option [temporary]” on page 194.</p>
ConnCount	<p>Returns the number of connections to the database. The value is dynamic: it can vary depending on how many connections currently exist. UltraLite supports up to fourteen concurrent database connections.</p>
date_format	<p>Returns the date format the database uses for string conversions. See “UltraLite date_format creation parameter” on page 144.</p>
date_order	<p>Returns the date order the database uses for string conversions. See “UltraLite date_order creation parameter” on page 147.</p>

Property	Description
Encryption	<p>Returns the type of database encryption, one of None, Simple, AES, or AES_FIPS.</p> <p>The encryption used by the database is determined by whether or not you have configured strong encryption (AES or AES_FIPS) and the DBKEY creation parameter, or obfuscation (simple encryption).</p> <p>The only time this property can change is when the value is originally None (that is, neither fips nor obfuscation is used) and you then change the encryption key by specifying a new encryption key on the Connection object by calling the correct function or method for your API. In this case, the value would change to AES because the fips creation parameter cannot be set after the database has been created. See:</p> <ul style="list-style-type: none"> • “ULChangeEncryptionKey method [UltraLite Embedded SQL]” [UltraLite - C and C++ Programming] • “ULConnection.ChangeEncryptionKey method [UltraLite.NET]” [UltraLite - .NET Programming] • “Connection.changeEncryptionKey method [UltraLite for M-Business Anywhere]” [UltraLite - M-Business Anywhere Programming (deprecated)] • “UltraLite database security” on page 26 • “UltraLite fips creation parameter” on page 148 • “UltraLite obfuscate creation parameter” on page 152 • “UltraLite DBKEY connection parameter” on page 176
File	<p>Returns the name of the database root file for the current connection, including the path. This is the value specified in the DBF connection parameter value. See:</p> <ul style="list-style-type: none"> • “UltraLite DBF connection parameter” on page 174
global_database_id	<p>Returns the value of the global_database_id option used for global autoincrement columns. See “UltraLite global_database_id option” on page 195.</p>
isolation_level	<p>Returns the current isolation level of the database. The value can either be read_committed or read_uncommitted. See “Change the isolation level” on page 40.</p>
MaxHashSize	<p>Returns the default number of maximum bytes to use for index hashing. This property can be set on a per-index basis. See “UltraLite max_hash_size creation parameter” on page 149.</p>

Property	Description
ml_remote_id	Returns the value of the <code>ml_remote_id</code> option that uniquely identifies the database for MobiLink synchronization. See “UltraLite ml_remote_id option” on page 196.
Name	Returns the name (or alias) of the database for the current connection. The name returned matches the DBN connection parameter value. If you did not use the DBN connection parameter, the name returned is the database file without the path and extension. See also: <ul style="list-style-type: none"> • “UltraLite DBN connection parameter” on page 176 • “UltraLite DBF connection parameter” on page 174
nearest_century	Returns the nearest century the database uses for string conversions. See “UltraLite nearest_century creation parameter” on page 151.
PageSize	Returns the page size of the database, in bytes. See “UltraLite page_size creation parameter” on page 153.
precision	Returns the floating-point precision the database uses for string conversions. See “UltraLite precision creation parameter” on page 155.
scale	Returns the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum PRECISION during string conversions by the database. See “UltraLite scale creation parameter” on page 156.
time_format	Returns the time format the database uses for string conversions. See “UltraLite time_format creation parameter” on page 157.
timestamp_format	Returns the timestamp format the database uses for string conversions. See “UltraLite timestamp_format creation parameter” on page 159.
timestamp_increment	Returns the minimum difference between two unique timestamps, in microseconds. See “UltraLite timestamp_increment creation parameter” on page 162.
timestamp_with_time_zone	Returns the timestamp format for <code>TIMESTAMP WITH TIME ZONE</code> values. See “UltraLite timestamp_with_time_zone_format creation parameter” on page 163.

See also

- [“Reading database properties” on page 33](#)
- [“Accessing database options” on page 34](#)
- [“UltraLite creation parameters” on page 139](#)
- [“UltraLite database options” on page 192](#)

UltraLite database options

This section describes the UltraLite database options that are available.

UltraLite database option values are defined when the database is first created and can be altered while connected to the database.

See also

- [“Accessing database options” on page 34](#)
- [“Reading database properties” on page 33](#)
- [“UltraLite creation parameters” on page 139](#)

UltraLite cache_allocation option

Explicitly resizes the cache. The value is a percentage of the minimum-to-maximum range. A value of zero means the minimum size, and a value of 100 means the maximum size.

Allowed values

Integer

Default

None.

Remarks

The cache_allocation property returns the current cache size as a percentage value of the minimum and maximum cache size.

See also

- “UltraLite database properties” on page 188
- “Accessing database options” on page 34
- “SET OPTION statement [UltraLite]” on page 446
- “UltraLite CACHE_SIZE connection parameter” on page 169
- “UltraLite CACHE_MIN_SIZE connection parameter” on page 168
- “UltraLite CACHE_MAX_SIZE connection parameter” on page 167
- “ULConnection.SetDatabaseOptionInt method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “ULSetDatabaseOptionULong method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “ULDatabaseSchema.SetDatabaseOption method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “DatabaseSchema.setDatabaseOption method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

UltraLite commit_flush_count option [temporary]

Sets a commit count threshold, after which a commit flush is performed.

Allowed values

Integer

Default

10

Remarks

Use 0 to disable the transaction count. When the transaction count is disabled, the number of commits is unlimited when a flush is triggered.

Both `commit_flush_count` and `commit_flush_timeout` are temporary database options. You must set these options each time you start a database. They persist as long as the database continues to run. They are only required when you set `COMMIT_FLUSH=grouped` as part of a connection string.

When you set this option and set the `COMMIT_FLUSH` connection parameter to `grouped` in your connection string, either threshold triggers a flush. When the flush occurs, UltraLite sets the counter *and* the timer back to 0. Then, both the counter and timer are monitored until one of these thresholds is subsequently reached.

An important consideration for setting the commit flush options is how much the delay to flush committed transactions poses a risk to the recoverability of your data. There is a slight chance that a transaction may be lost, even though it has been committed. If a serious hardware failure occurs after a commit, but before the transaction is flushed to storage, the transaction is rolled back on recovery. A longer delay can increase UltraLite performance. You must choose an appropriate count threshold with care.

To set the `commit_flush_count` option from a client application, set the option using the `set database` option function for the programming interface you are using or use the `SET OPTION SQL` statement.

See also

- “Accessing database options” on page 34
- “`SET OPTION` statement [UltraLite]” on page 446
- “Flush single or grouped transactions” on page 471
- “UltraLite `commit_flush_timeout` option [temporary]” on page 194
- “UltraLite `COMMIT_FLUSH` connection parameter” on page 172
- “`ULConnection.SetDatabaseOptionInt` method [UltraLite C++]” [*UltraLite - C and C++ Programming*]
- “`ULSetDatabaseOptionULong` method [UltraLite Embedded SQL]” [*UltraLite - C and C++ Programming*]
- “`ULDatabaseSchema.SetDatabaseOption` method [UltraLite.NET]” [*UltraLite - .NET Programming*]
- “`DatabaseSchema.setDatabaseOption` method [UltraLite for M-Business Anywhere]” [*UltraLite - M-Business Anywhere Programming (deprecated)*]

UltraLite `commit_flush_timeout` option [temporary]

Sets a time interval threshold, after which a grouped commit flush is performed.

Allowed values

Integer, in milliseconds

Default

10000 milliseconds

Remarks

Use 0 to disable the time threshold.

Both `commit_flush_count` and `commit_flush_timeout` are temporary database options. You must set these options each time you start a database. They persist as long as the database continues to run. They are only required when you set `COMMIT_FLUSH=grouped` as part of a connection string.

If you set this option in addition to the `commit_flush_timeout` option and if you have set the `COMMIT_FLUSH` connection parameter to `grouped`, either threshold triggers a flush. When the flush occurs, UltraLite sets the counter *and* the timer back to 0. Then, both the counter and timer are monitored until one of these thresholds is subsequently reached.

An important consideration for setting the commit flush options is how much the delay to flush committed transactions poses a risk to the recoverability of your data. There is a slight chance that a transaction may be lost, even though it has been committed. If a serious hardware failure occurs after a commit, but before the transaction is flushed to storage, the transaction is rolled back on recovery. A longer delay can increase UltraLite performance. You must choose an appropriate timeout threshold with care.

To set the `commit_flush_timeout` option from a client application, set it using the set database option function for the programming interface you are using or use the SET OPTION SQL statement.

See also

- [“Accessing database options” on page 34](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“UltraLite commit_flush_count option \[temporary\]” on page 193](#)
- [“UltraLite COMMIT_FLUSH connection parameter” on page 172](#)
- [“ULConnection.SetDatabaseOptionInt method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULSetDatabaseOptionULong method \[UltraLite Embedded SQL\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“DatabaseSchema.setDatabaseOption method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)

UltraLite global_database_id option

Sets the database identification number.

Allowed values

Unique, non-negative integer

Default

The range of default values for a particular global autoincrement column is $pn + 1$ to $p(n + 1)$, where p is the partition size of the column and n is the global database identification number.

Remarks

To maintain primary key uniqueness when synchronizing with a MobiLink server, the global ID sets a starting value for GLOBAL AUTOINCREMENT columns. The global ID must be set before default values can be assigned. If a row is added to a table and does not have a value set already, UltraLite generates a value for the column by combining the `global_database_id` value and the partition size.

When this option is set, UltraLite performs a commit.

When deploying an application, you must assign a different identification number to each database for synchronization with the MobiLink server. You can change the global ID of an existing database at any time.

To set the `global_database_id` option from a client application, use the set database option function for the programming interface you are using or use the SET OPTION SQL statement.

See also

- [“Using GLOBAL AUTOINCREMENT” \[MobiLink - Server Administration\]](#)
- [“Accessing database options” on page 34](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“Using GLOBAL AUTOINCREMENT in UltraLite” on page 69](#)
- [“ULConnection.SetDatabaseOptionInt method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULSetDatabaseID method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.DatabaseID property \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Connection.setDatabaseID method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.setDatabaseId method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_DATABASE_ID variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

Example

To autoincrement UltraLite database columns from 3001 to 4000, set the global ID to 3.

```
SET OPTION global_database_id=3;
```

UltraLite isolation_level option

Isolation levels define the degree to which the operations in one transaction are visible to the operations in other concurrent transactions. UltraLite uses the default isolation level, read_committed, for connections in auto-commit mode.

See also

- [“Accessing database options” on page 34](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“Change the isolation level” on page 40](#)
- [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.BeginTransaction method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“Data synchronization on a BlackBerry smartphone” \[UltraLite - Java Programming\]](#)

UltraLite ml_remote_id option

The **remote ID** is a unique identifier for an UltraLite database that is used by MobiLink to identify the database for synchronization. The remote ID can be any string that has meaning to you, as long as the string remains unique among all remote MobiLink clients. The ID can also be set to NULL (NULL is the initial value). During synchronization, if the remote ID is NULL, UltraLite will assign it to a generated GUID.

If you prepopulate an UltraLite database using synchronization for distribution to multiple devices, you must reset the remote ID to NULL before distribution to ensure that each database has a unique remote

ID. Upon distribution, a new unique remote ID can be set explicitly or it can be left as NULL so that UltraLite will automatically generate a new unique value.

Allowed values

Any value that uniquely identifies the database for MobiLink synchronization.

Default

Null

Remarks

MobiLink uses the remote ID to store the synchronization information for the remote database. Given the remote ID, MobiLink user names are no longer required to be unique. The remote ID becomes particularly useful when you have multiple MobiLink users synchronizing the same UltraLite database. In this case, your synchronization scripts should reference the remote ID and not just the user name.

When this option is set, UltraLite performs a commit.

To set the `ml_remote_id` option from a client application, set it using the `set database option` function for the programming interface you are using or use the `SET OPTION SQL` statement.

See also

- [“Accessing database options” on page 34](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“Remote IDs” \[MobiLink - Client Administration\]](#)
- [“User Name synchronization parameter” on page 110](#)
- [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULSetDatabaseOptionString method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)
- [“DatabaseSchema.setDatabaseOption method \[UltraLite for M-Business Anywhere\]” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_ML_REMOTE_ID variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

UltraLite utilities

UltraLite includes a set of utilities that are designed to perform basic database administration activities at a command prompt. Many of these utilities share a similar functionality to the SQL Anywhere Server utilities. However, the way options are used can vary. Always refer to the UltraLite reference documentation for the UltraLite implementation of these options.

Note

Options for the utilities documented in this section are case sensitive, unless otherwise noted. Type options *exactly* as they are displayed.

Supported exit codes

The ulload, ulsync, and ulunload utilities return exit codes to indicate whether the operation a utility attempted to complete was successful. 0 indicates a successful operation. Any other value indicates that the operation failed.

Exit code	Status	Description
0	EXIT_OKAY	Operation successful.
1	EXIT_FAIL	Operation failure.
3	EXIT_FILE_ERROR	Database cannot be found.
4	EXIT_OUT_OF_MEMORY	Exhausted the dynamic memory of the device.
6	EXIT_COMMUNICATIONS_FAIL	Communications error generated while talking to the UltraLite engine.
9	EXIT_UNABLE_TO_CONNECT	Invalid UID or PWD provided, therefore cannot connect to the database.
12	EXIT_BAD_ENCRYPT_KEY	Missing or invalid encryption key.
13	EXIT_DB_VER_NEWER	Detected that the database version is incompatible. The database must be upgraded to a newer version.
255	EXIT_USAGE	Invalid command line options.

Interactive SQL for UltraLite utility (dbisql)

Executes SQL statements and runs script files against a database.

Syntax

```
dbisql -c "connection-string" [ options ] [ dbisql-statement | dbisql-script-file ]
```

```
dbisql -c "connection-string" -ul [ options ] [ dbisql-statement | dbisql-script-file ]
```

dbisql-statement: A SQL statement or a series of sql statements separated by a command-delimiter.

Option	Description
<i>@data</i>	<p>Reads in options from the specified environment variable or configuration file.</p> <p>If both the environment variable and configuration file exist with the same name, the environment variable is used. See “Configuration files” [SQL Anywhere Server - Database Administration].</p> <p>If you want to protect passwords or other information in the configuration file, you can use the File Hiding utility to obfuscate the contents of the configuration file. See “File Hiding utility (dbfhide)” [SQL Anywhere Server - Database Administration].</p>
-c <i>"keyword=value; ..."</i>	<p>Specifies connection parameters. If Interactive SQL cannot connect, you are presented with a window where you can enter the connection parameters. If you do not specify both a user ID and a password, the default UID of DBA and PWD of sql are assumed. See “Connection parameters” [SQL Anywhere Server - Database Administration].</p>
-d <i>delimiter</i>	<p>Specifies a command delimiter. Quotation marks around the delimiter are optional, but are required when the command shell itself interprets the delimiter in some special way.</p> <p>This option overrides the setting of the <code>command_delimiter</code> option. See “command_delimiter option [Interactive SQL]” [SQL Anywhere Server - Database Administration].</p>
-d1	<p>Echoes all statements explicitly executed by the user to the command window (STDOUT). This can provide useful feedback for debugging SQL scripts, or when Interactive SQL is processing a long SQL script. (The final character is a number 1, not a lowercase L). This option is only available when you run Interactive SQL as a command line program.</p>
-datasource <i>DSN-name</i>	<p>Specifies an ODBC data source to connect to.</p>

Option	Description
-f <i>filename</i>	<p>Opens (but does not run) the file called <i>filename</i> in the SQL Statements pane.</p> <p>If the -f option is given, the -c option is ignored; that is, no connection is made to the database.</p> <p>The file name can be enclosed in quotation marks, and <i>must</i> be enclosed in quotation marks if the file name contains a space.</p> <p>If the file does not exist, or if it is really a directory instead of a file, Interactive SQL prints an error message and then quits.</p> <p>If the file name does not include a full drive and path specification, it is assumed to be relative to the current directory.</p> <p>This option is only supported when Interactive SQL is run as a windowed application.</p>
-host <i>hostname</i>	<p>Specifies the <i>hostname</i> or IP address of the computer on which the database server is running. You can use the name localhost to represent the current computer.</p>
-nogui	<p>Runs Interactive SQL as a console application, with no windowed user interface. This is useful for batch operations.</p> <p>If you specify either <i>dbisql-statement</i> or <i>dbisql-script-file</i>, then -nogui is assumed.</p> <p>In this mode, Interactive SQL sets the program exit code to indicate success or failure. On Windows operating systems, the environment variable ERRORLEVEL is set to the program exit code. See “Software component exit codes” [SQL Anywhere Server - Programming].</p>

Option	Description
<p>-onerror { continue exit }</p>	<p>Controls what happens if an error is encountered while reading statements from a script file. It is useful when using Interactive SQL in batch operations. This option overrides the <code>on_error</code> setting. See “on_error option [Interactive SQL]” [SQL Anywhere Server - Database Administration].</p> <p>Define one of the following supported <i>behavior</i> values:</p> <ul style="list-style-type: none"> ● Continue The error is ignored and Interactive SQL continues executing statements. ● Exit Interactive SQL terminates.
<p>-q</p>	<p>Suppresses output messages. Sets the utility to run in quiet mode. This is useful only if you start Interactive SQL with a statement or script file. Specifying this option does not suppress error messages, but it does suppress the following:</p> <ul style="list-style-type: none"> ● warnings and other non-fatal messages ● the printing of result sets
<p>-ul</p>	<p>Specifies that UltraLite databases are the default. Interactive SQL customizes the options available to you depending on the type of database you are connected to.</p> <p>By default, Interactive SQL assumes that you are connecting to SQL Anywhere databases. When you specify the <code>-ul</code> option, the default changes to UltraLite databases. Regardless of the type of database set as the default, you can connect to either SQL Anywhere or UltraLite databases by choosing the database type from the Change Database Type dropdown list on the Connect window.</p> <p>For more information about connecting to UltraLite databases from Interactive SQL, see “Interactive SQL for UltraLite utility (dbisql)” on page 198.</p>

Option	Description
-version	Displays the version number of Interactive SQL. You can also view the version number from within Interactive SQL; from the Help menu, click About Interactive SQL .
-x	Scans statements but does not execute them. This is useful for checking long script files for syntax errors. For detailed descriptions of SQL statements and Interactive SQL statements, see “ SQL language elements ” [SQL Anywhere Server - SQL Reference].
<i>dbisql-statement</i> <i>dbisql-script-file</i>	Execute the SQL statement or execute the specified <i>dbisql-script-file</i> . If you do not specify a <i>dbisql-statement</i> or <i>dbisql-script-file</i> , Interactive SQL enters interactive mode, where you can type a statement into a command window.

Remarks

Interactive SQL allows you to browse the database, execute SQL statements, and run script files. It also provides feedback about:

- the number of rows affected
- the time required for each statement
- the execution plan of queries
- any error messages

You can connect to both SQL Anywhere and UltraLite databases. For information about connecting to SQL Anywhere databases, see “[Interactive SQL utility \(dbisql\)](#)” [[SQL Anywhere Server - Database Administration](#)].

Interactive SQL is supported on Windows, Solaris, Linux, and Mac OS X. See <http://www.sybase.com/detail?id=1061806>.

For Windows, there are two executables:

1. Batch scripts should call *dbisql* or *dbisql.com*, not *dbisql.exe*. The *dbisql.com* executable is linked as a console application.
2. The *dbisql.exe* executable is linked as a windowed application and does not block the command shell from which it was started. If *dbisql.exe* is run from a batch file, you won't see any output sent to the standard output or standard error files.

The default encoding for Interactive SQL can also be temporarily set using the `default_isql_encoding` option. See “[default_isql_encoding option \[Interactive SQL\]](#)” [*SQL Anywhere Server - Database Administration*].

You can specify the encoding to use when reading or writing files using the `ENCODING` clause of the `INPUT`, `OUTPUT`, or `READ` statement. See:

- “[INPUT statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[OUTPUT statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[READ statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]

Exit codes are 0 (success) or non-zero (failure). Non-zero exit codes are set only when you run Interactive SQL in batch mode (with a command line that contains a SQL statement or the name of a script file). See “[Software component exit codes](#)” [*SQL Anywhere Server - Programming*].

In command-prompt mode, Interactive SQL sets the program exit code to indicate success or failure. On Windows operating systems, the environment variable `ERRORLEVEL` is set to the program exit code.

When executing a *reload.sql* file with Interactive SQL, you must specify the encryption key as a parameter. If you do not provide the key in the `READ` statement, Interactive SQL prompts for the key.

You can start Interactive SQL in the following ways:

- From Sybase Central, by clicking **File** » **Open Interactive SQL**.
- From the **Start** menu by clicking **Start** » **Programs** » **SQL Anywhere 12** » **Administration Tools** » **Interactive SQL**.
- Using the `dbisql` command at a command prompt.

See also

- “[SQL statements for Interactive SQL](#)” [*SQL Anywhere Server - Database Administration*]
- “[Interactive SQL](#)” [*SQL Anywhere Server - Database Administration*]
- “[Configuration files](#)” [*SQL Anywhere Server - Database Administration*]
- “[File Hiding utility \(dbfhide\)](#)” [*SQL Anywhere Server - Database Administration*]
- “[INPUT statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[OUTPUT statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[READ statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[UltraLite connection parameters](#)” on page 165
- “[Supported exit codes](#)” on page 198

Example

The following command runs the script file *mycom.sql* against the *CustDB.udb* database for UltraLite. Because a user ID and password are not defined, the default user ID **DBA** and password **sql** are assumed. The `-onerror` option is defined as `Exit`; so, if there is an error in the script file, the process terminates.

```
dbisql -ul -c DBF=CustDB.udb -onerror exit mycom.sql
```

SQL preprocessor for UltraLite utility (sqlpp)

Preprocesses a C/C++ program that contains embedded SQL (ESQL), so that code required for that program can be generated before you run the compiler. The table below describes the entire set of options for completeness, but the only relevant options for UltraLite are **-eu** and **-wu**.

Syntax

sqlpp -u [*options*] *esql-filename* [*output-filename*]

Option	Description
-d	Generate code that reduces data space size, but increases code size. Data structures are reused and initialized at execution time before use.
-e flag	<p>This option flags as an error any static embedded SQL that is not part of a specified standard. The <i>level</i> value indicates the standard to use. For example, <code>sqlpp -e c03 . . .</code> flags any syntax that is not part of the core SQL/2003 standard.</p> <p>The allowed values of <i>level</i> are:</p> <ul style="list-style-type: none"> ● c03 Flag syntax that is not core SQL/2003 syntax ● p03 Flag syntax that is not full SQL/2003 syntax ● c99 Flag syntax that is not core SQL/1999 syntax ● p99 Flag syntax that is not full SQL/1999 syntax ● e92 Flag syntax that is not entry-level SQL/1992 syntax ● i92 Flag syntax that is not intermediate-level SQL/1992 syntax ● f92 Flag syntax that is not full-SQL/1992 syntax ● t Flag non-standard host variable types ● u Flag syntax that is not supported by UltraLite <p>For compatibility with previous SQL Anywhere versions, you can also specify e, I, and f, which correspond to e92, i92, and f92, respectively.</p>
-h width	Limits the maximum length of split lines output by sqlpp to <i>width</i> in the <i>.c</i> file. Backslash characters are added to the end of split lines, so that a C compiler can parse the split lines as one continuous line. The default value is no maximum line length (output lines are not split by default).
-k	Notify the preprocessor that the program to be compiled includes a user declaration of SQLCODE.

Option	Description
-m mode	Cursor updatability mode. Either HISTORICAL or READONLY .
-n	Generate line number information in the C file by using <code>#line</code> directives in the appropriate places in the generated code. Use this option to the report source errors and to debug source on line numbers in the <i>esql-filename</i> file, rather than in the <i>output-filename</i> file.
-o O/S spec	Not applicable to UltraLite.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-r	Not applicable to UltraLite.
-s string-length	Set the maximum size string that the preprocessor will put into the C file. Strings longer than this value are initialized using a list of characters ('a','b','c', and so on). Most C compilers have a limit on the size of string literal they can handle. This option is used to set that upper limit. The default value is 500.
-u	Required for UltraLite. Generate output specifically required for UltraLite databases.
-w level	Flag non-conforming SQL syntax as a warning. The <i>level</i> value indicates the standard to use. For example, <code>sqlpp -w c03 . . .</code> flags any SQL syntax that is not part of the core SQL/2003 syntax. The allowed values of <i>level</i> are: <ul style="list-style-type: none"> • c03 Flag syntax that is not core SQL/2003 syntax • p03 Flag syntax that is not full SQL/2003 syntax • c99 Flag syntax that is not core SQL/1999 syntax • p99 Flag syntax that is not full SQL/1999 syntax • e92 Flag syntax that is not entry-level SQL/1992 syntax • i92 Flag syntax that is not intermediate-level SQL/1992 syntax • f92 Flag syntax that is not full-SQL/1992 syntax • t Flag non-standard host variable types • u Flag syntax that is not supported by UltraLite For compatibility with previous SQL Anywhere versions, you can also specify <i>e</i> , <i>I</i> , and <i>f</i> , which correspond to <i>e92</i> , <i>i92</i> , and <i>f92</i> , respectively.

Option	Description
-x	Change multibyte strings to escape sequences, so that they can be passed through a compiler.
-z <i>collation-sequence</i>	Specify the collation sequence.

Remarks

This preprocessor translates the SQL statements in the input-file into C/C++. It writes the result to the *output-filename*. The normal extension for source files containing embedded SQL is *sqlc*. The default *output-filename* is the *esql-filename* base name with an extension of *c*. However, if the *esql-filename* already has the *.c* extension, the default output extension is *.cc*.

The collation sequence is used to help the preprocessor understand the characters used in the source code of the program. For example, in identifying alphabetic characters suitable for use in identifiers. In UltraLite, collations include a code page plus a sort order. If you do not specify **-z**, the preprocessor attempts to determine a reasonable collation to use based on the operating system.

To see a list of supported collations (and their corresponding codepages), run the command **ulinit -Z**.

Tip

The SQL preprocessor (sqlpp) has the ability to flag static SQL statements in an embedded SQL application at compile time. This feature can be especially useful when developing an UltraLite application, to verify SQL statements for UltraLite compatibility. You can test compatibility of SQL for both SQL Anywhere and UltraLite applications by using either **-e** and/or **-w** options. For an overview of the SQL Flagger, see “SQL compliance testing using the SQL Flagger” [[SQL Anywhere Server - SQL Usage](#)].

See also

- “Embedded SQL” [[SQL Anywhere Server - Programming](#)]
- “UltraLite character sets” on page 23

Example

The following command preprocesses the *srcfile.sqlc* embedded SQL file in quiet mode for an UltraLite application.

```
sqlpp -u -q MyEsqFile.sqlc
```

UltraLite Engine utility (uleng12)

Manages concurrent UltraLite database connections from applications on 32- and 64-bit Windows Desktop, 32-bit Linux and Windows Mobile. Also allows the UltraLite engine to run as a daemon using the **-ud** option.

Syntax

```
uleng12 [ -ud ] [ db-file-name ]
```

Option	Description
-ud <i>db-file-name</i>	Lets you run the engine so that it continues running after the current user session ends. When you start the daemon directly using the -ud option, the uleng12 command creates the daemon process and returns immediately (exiting and allowing the next command to be executed) before the daemon initializes itself or attempts to open any of the databases specified in the command.

Remarks

The UltraLite engine does not display a messages window on startup.

The UltraLite engine should be used by an application in scenarios where multiple processes could be accessing the same database at the same time. The engine is installed in the SQL Anywhere *bin32* or *bin64* directory because the UltraLite desktop administration tools use the engine to connect to databases.

Using the -ud option, you can run the UltraLite engine so that when you log off the computer, the database engine remains running. (Normally when you log off the computer, all applications associated with the session shut down.)

See also

- [“UltraLite deployment” on page 115](#)
- [“Build and deploy UltraLite C++ applications” \[*UltraLite - C and C++ Programming*\]](#)
- [“Choose an UltraLite data management component for Windows Mobile” on page 13](#)
- [“UltraLite Engine Stop utility \(ulstop\)” on page 207](#)
- [“UltraLite START connection parameter” on page 185](#)

UltraLite Engine Stop utility (ulstop)

Stops the UltraLite engine on 32-bit Windows desktops and Windows Mobile.

Syntax

```
ulstop
```

Remarks

Use ulstop during development to shut down the engine manually. You typically do not require ulstop in live deployments.

See also

- [“Choose an UltraLite data management component for Windows Mobile” on page 13](#)
- [“UltraLite Engine utility \(uleng12\)” on page 206](#)

UltraLite Erase utility (ulerase)

Erases an UltraLite database.

Syntax

ulerase [*options*] [*db-file-name*]

Option	Description
@ <i>data</i>	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.
-k <i>key</i> OR --ek= <i>key</i>	Specify the encryption key for an encrypted database.
-p OR --ep	Specify that you want to be prompted for the encryption key.
--log	Log operations to the specified file.
-q OR --quiet	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-u <i>uid,pwd</i> OR --dba= <i>uid,pwd</i>	Specify the userid and password required to access the database.
-?	Displays utility usage information and exits.
<i>db-file-name</i>	Erase the specified database.

Remarks

The database must be accessible. The user ID and password combination must allow a connection, otherwise the database is not erased.

Encrypted databases require a key provided in the connection string, or using one of **-k** *key* or **-p**.

See also

- “Configuration files” [[SQL Anywhere Server - Database Administration](#)]

UltraLite Information utility (ulinfo)

Displays information about an UltraLite database.

Syntax

ulinfo -c options

Option	Description
@data	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.
-c "keyword=value;..." OR --connect"keyword=value;..."	Supply database connection parameters. Required.
-q OR --quiet	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
--log=filename	Log operations to the specified file.
-? OR --help	Displays utility usage information and exits.

Remarks

Warning messages generated when opening an UltraLite database are always displayed unless you use the -q option.

See also

- “Configuration files” [[SQL Anywhere Server - Database Administration](#)]
- “UltraLite connection parameters” on page 165
- “UltraLite global_database_id option” on page 195
- “UltraLite ml_remote_id option” on page 196

Example

Show basic database internals for a file named *cv_dbattr.udb* that has already been synchronized:

```
ulinfo -c DBF=cv_dbattr.udb
```

Show database internals for a file named *CustDB.udb* and display database properties by enabling verbose messaging:

```
ulinfo -c DBF=CustDB.udb -v
```

UltraLite Initialize Database utility (ulinit)

Creates a new UltraLite database. This utility will either create an empty database with characteristics specified with the command line arguments ("empty mode"), or it can create a database based on a SQL Anywhere database "extract mode". In the latter case, an initial schema is created that matches tables and indexes in the SQL Anywhere reference database. Moreover, many of the SQL Anywhere reference database characteristics will be used for the new UltraLite database.

In the usage below, some options are only permissible in "extract mode," while others are only permissible in "empty mode." If no mode is stated in the description, then the option can be used in either mode.

Syntax

ulinit *options dbname*

Option	Description
@ <i>data</i>	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used. See “Configuration files” [SQL Anywhere Server - Database Administration] .
-a " <i>keyword=value;...</i> " OR --SAconnect=" <i>keyword=value;...</i> "	Sets the utility to "extract" mode and connects to an existing database using the specified connection parameters (extract mode). If this option is not present, the utility creates a new database using the specified connection parameters (empty mode).
-c OR --case	Empty mode. Enforce case sensitivity on all string comparisons.

Option	Description
-d OR --datacopy	Extract mode. For each table in the new UltraLite database, copy data from the corresponding table in the SQL Anywhere database. By default, this data will not be uploaded in subsequent synchronizations. To include the data in the next upload synchronization, use -i with -d .
--date_format=format	Empty mode. Sets the format for dates retrieved from the database. See “UltraLite date_format creation parameter” on page 144.
--date_order=date-format-interpretation	Empty mode. Sets the interpretation of the date format. See “UltraLite date_order creation parameter” on page 147.
-e value OR --fips=value	Empty mode. On or off, 1 or 0, and so on. This option controls AES FIPS-certified encryption by using a Certicom certified cryptographic algorithm. See “UltraLite database security” on page 26 and “UltraLite fips creation parameter” on page 148.
-f OR --exactschema	Extract mode. Fail if exact schema is not supported in UltraLite; otherwise, warnings will appear if schema differs.
-g id OR --databaseid=id	Set the initial database ID to the INTEGER value you assign. This initial value is used with a partition size for new rows that have global autoincrement columns. When deploying an application, you must assign a different range of identification numbers to each database for synchronization with the MobiLink server. See “UltraLite global_database_id option” on page 195.

Option	Description
-i OR --insertforupload	Extract mode. Use with -d . Include inserted rows in the next upload synchronization. By default, rows inserted by this utility are not uploaded during synchronization.
--identity-file = <i>file</i>	Specify the file containing the client TLS identity. See “ identity ” [<i>MobiLink - Client Administration</i>].
--identity-password = <i>password</i>	Specify the password for the client TLS identity. See “ identity_password ” [<i>MobiLink - Client Administration</i>].
-k <i>key</i> OR --key = <i>key</i>	Extract mode. Specify the encryption key for an encrypted database.
-K OR --prompt	Empty mode. Specify that you want to be prompted for the encryption key.
-l <i>logfile</i> OR --sql = <i>logfile</i>	Extract mode. Log DDL database schema creation SQL statements, as executed, to <i>logfile</i> .
--log = <i>filename</i>	Empty mode. Log operations to the specified file.
-m <i>filename</i> OR --mirror_file = <i>filename</i>	Extract mode. Specify the database mirror file. See “ UltraLite MIRROR_FILE connection parameter ” on page 180.
--max_hash_size = <i>size</i>	Empty mode. Sets the default index hash size in bytes. See “ UltraLite max_hash_size creation parameter ” on page 149.

Option	Description
<p>-n <i>pubname</i></p> <p>OR</p> <p>--publication=<i>pubname</i></p>	<p>Extract mode.</p> <p>Required. Add tables to the UltraLite database schema.</p> <p><i>pubname</i> specifies a publication in the reference database. Tables in the publication are added to the UltraLite database.</p> <p>Specify the option multiple times to add tables from multiple publications to the UltraLite database. To add all tables in the reference database to the UltraLite database, specify -n*.</p>
<p>--nearest_century=<i>yy</i></p>	<p>Empty mode.</p> <p>Controls the interpretation of two-digit years in string-to-date conversions. See “UltraLite nearest_century creation parameter” on page 151.</p>
<p>-o <i>value</i></p> <p>OR</p> <p>--obfuscate=<i>value</i></p>	<p>Empty mode.</p> <p>On or off, 1 or 0, and so on. Controls whether data in the database is obfuscated. Obfuscation is a form of simple encryption. See “UltraLite database security” on page 26 and “UltraLite obfuscate creation parameter” on page 152.</p>
<p>-p <i>size</i></p> <p>OR</p> <p>--page_size=<i>size</i></p>	<p>Empty mode.</p> <p>Specify the database page size.</p>
<p>--precision=<i>precision</i></p>	<p>Empty mode.</p> <p>Specifies the maximum number of digits in decimal point arithmetic results. See “UltraLite precision creation parameter” on page 155.</p>
<p>-q</p> <p>OR</p> <p>--quiet</p>	<p>Empty mode and extract mode.</p> <p>Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages and warnings are still displayed, however.</p>

Option	Description
<p>-r <i>size</i></p> <p>OR</p> <p>--reserve_size</p>	<p>Database connection only.</p> <p>Reserve size. See “UltraLite RESERVE_SIZE connection parameter” on page 184.</p>
<p>-s <i>pubname</i></p> <p>OR</p> <p>--sync_publication</p>	<p>Extract mode.</p> <p>Create a publication in the UltraLite database with the same definition as <i>pubname</i> in the reference database. Publications are used to configure synchronization. Supply more than one -s option to name more than one synchronization publication.</p> <p>Note that the tables in this publication must be included in a publication listed by the -n option.</p> <p>If -s is not supplied, the UltraLite remote database has no named publications.</p> <p>For more information about how to create publications for MobiLink synchronization, see “Publishing data in UltraLite” on page 76.</p>
<p>-S <i>checksum_level</i></p> <p>OR</p> <p>--checksum_level=checksum_level</p>	<p>Empty mode.</p> <p>0, 1, or 2. Specifies the checksum level validation on database pages. See “UltraLite checksum_level creation parameter” on page 142.</p>
<p>--scale=scale</p>	<p>Empty mode.</p> <p>Specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision. See “UltraLite scale creation parameter” on page 156.</p>
<p>-t <i>file</i></p> <p>OR</p> <p>--rootcert=file</p>	<p>Empty mode and extract mode.</p> <p>Specify the file containing the trusted root certificate. This certificate is required for server authentication.</p>

Option	Description
--time_format= <i>format</i>	Empty mode. Sets the format for times retrieved from the database. See “UltraLite time_format creation parameter” on page 157.
--timestamp_format= <i>format</i>	Empty mode. Sets the format for timestamps retrieved from the database. See “UltraLite timestamp_format creation parameter” on page 159.
--timestamp_increment= <i>increment</i>	Empty mode. Determines how the timestamp is truncated in UltraLite. See “UltraLite timestamp_increment creation parameter” on page 162.
--timestamp_with_time_zone_format= <i>format</i>	Empty mode. This option sets the format for TIMESTAMP WITH TIME ZONE values retrieved from the database. See “UltraLite timestamp_with_time_zone_format creation parameter” on page 163.
-u <i><uid></i> , <i><pwd></i> OR --dba= <i><uid></i> , <i><pwd></i>	Database connection only. Specify the userid and password.
--utf8_encoding= <i>value</i>	Empty mode. On or off, 1 or 0, and so on. Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode. See “UltraLite character sets” on page 23 and “UltraLite utf8_encoding creation parameter” on page 164.
-w OR --nowarnings	Extract mode. Do not display warnings.

Option	Description
-x <i>table</i> OR --exclude	Extract mode. Exclude the tables named in the list.
-y OR --overwrite	Empty mode and extract mode. Over-write the existing database file.
-z <i>collation-sequence</i> OR --collation=collation-sequence	Empty mode. Specify the collation sequence.
-Z OR --listcollation	Empty mode. List the available collation sequences and exit.
-? OR --help	Display utility usage and exit.

Remarks

The SQL Anywhere reference database, when used, acts as the source for:

- database configuration (for example, the collation sequence used)
- table definitions
- synchronization publications

Together they help create the UltraLite schema—information that defines the structure of the new UltraLite database. However, the new database you create is initially empty unless you use the **-d** option.

When run in extract mode, ULINIT will attempt to create an UltraLite database that matches, as closely as possible, the SQL Anywhere database. For example, if a column in the SQL Anywhere database includes a clause that UltraLite does not support, the default value will be ignored and the UltraLite default used instead. A warning will be generated and creation will continue. This supports the case where SQL Anywhere tables cannot be modified, but a reasonable UltraLite alternative is available. To enforce an exact schema match, use the **-f** option. If the schema will not support a reasonable UltraLite alternative, ulinit will fail.

If you want to initialize an UltraLite database from an RDBMS other than SQL Anywhere, use the **Create Synchronization Model Wizard** in Sybase Central. When you run the wizard, you are prompted to connect to a consolidated database to obtain schema information.

UltraLite uses the name of the collation sequence that was defined in the reference database. However, you can still choose to use UTF-8 to encode the database, by setting the `utf8_encoding` property.

To see a list of supported collations (and corresponding codepages), run **ulinit -Z** at a command prompt. If your collation sequence is not supported by UltraLite, you should change it to one that is. For example, if your reference database collation is the UCA collation, you should:

1. Unload the reference database and then reload it with a different collation.
2. Run `ulinit` on this new version of the database.

See also

- [“Creating an UltraLite database from a SQL Anywhere reference database” on page 19](#)
- [“Synchronization models” \[MobiLink - Getting Started\]](#)
- [“UltraLite connection parameters” on page 165](#)

Examples

Create a file called *customer.udb* that contains the tables defined in *TestPublication*:

```
ulinit -a "DSN=MySADb;UID=JimmyB;PWD=secret" -n TestPublication -k mykey
customer.udb
```

This example connects to a SQL Anywhere database defined in the *MySADb* datasource. It creates an UltraLite database with all the database options from that database and all the tables contained in the *TestPublication* publication. The new UltraLite database is called *customer.udb* and is encrypted with the key *mykey*.

Create a file called *customer.udb* that contains two distinct publications. Specifically, *Pub1* may contain a small subset of data for priority synchronization, while *Pub2* could contain the bulk of the data:

```
ulinit -a "DSN=MySADb;UID=JimmyB;PWD=secret" --exactschema -n Pub1 -n Pub2 -
s Pub1 -s Pub2 customer.udb
```

This example connects to a SQL Anywhere database defined in the *MySADb* datasource. It creates an UltraLite database with all the database options from that database and all the tables contained in the publications *Pub1* and *Pub2*. The new UltraLite database is also created with the publications *Pub1* and *Pub2*. Since the `--exactschema` option is set, `ulinit` will fail if it cannot extract the all precise schema.

Create a new blank database that overwrites another *customer.udb* file if it already exists. The new database has no schema and all the database options are set to default values.

```
ulinit -y customer.udb
```

UltraLite Load XML to Database utility (ulload)

Loads data from an XML file into a new or existing database.

Syntax

```
ulload -c "connection-string" [ options ] xml-file
```

Option	Description
@data	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used. See “ Configuration files ” [<i>SQL Anywhere Server - Database Administration</i>].
-a OR --append	Add data and schema definitions into an existing database.
-c "keyword=value;..." OR --connect "keyword=value;..."	Supply the database connection parameters.
-d OR --dataonly	Load data only, ignoring any schema metadata in the XML file input. -d or --dataonly switches can only be used when -a is specified (because it is loading data only, the UDB it is loading the data into must exist with a schema that supports the data being loaded into it).
-e value OR --fips= value	Specify on or off, 1 or 0, and so on. This option controls AES FIPS-certified encryption, by using a Certicom certified cryptographic algorithm. See “ UltraLite database security ” on page 26 and “ UltraLite fips creation parameter ” on page 148.
-E behavior OR --onerror=behavior	Control what happens if an error is encountered while reading data from the XML file. Specify one of the following supported <i>behavior</i> values: <ul style="list-style-type: none"> ● continue ulload ignores the error and continues to load XML. ● prompt ulload prompts you to see if you want to continue. ● quit ulload stops loading the XML and terminates with an error. This behavior is the default behavior if none is specified. ● exit ulload exits.

Option	Description
-f <i>directory</i> OR --filedir = <i>directory</i>	Set the directory that contains files with additional data to load. See “UltraLite Database Unload utility (ulunload)” on page 227 .
-g <i>ID</i> OR --databaseid = <i>ID</i>	Set the initial database ID to the INTEGER value you assign. This initial value is used with a partition size for new rows that have global autoincrement columns. When deploying an application, you must assign a different range of identification numbers to each database for synchronization with the MobiLink server. See “UltraLite global_database_id option” on page 195 .
-i OR --insertforsync	Include inserted rows in the next upload synchronization. By default, rows inserted by this utility are not uploaded during synchronization.
--identity-file = <i>file</i>	Specify the file containing the client TLS identity. See “identity” [MobiLink - Client Administration] .
--identity-password = <i>password</i>	Specify the password for the client TLS identity. See “identity_password” [MobiLink - Client Administration] .
-l <i>filename</i> OR --log = <i>filename</i>	Log operations to the specified file.
-n OR --schemaonly	Load schema metadata only, ignoring any data in the XML input file.
-o <i>value</i> OR --obfuscate = <i>value</i>	On or off, 1 or 0, and so on. Controls whether data in the database is obfuscated. Obfuscation is a form of simple encryption. See “UltraLite database security” on page 26 and “UltraLite obfuscate creation parameter” on page 152 .
-p <i>page-size</i> OR --page_size = <i>page-size</i>	Defines the database page size. See “UltraLite page_size creation parameter” on page 153 .

Option	Description
-q OR --quiet	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-s file OR --sql=file	Log the SQL statements used to load the database into the specified <i>file</i> .
-t file OR --rootcert=file	Specify the file containing the trusted root certificate. This certificate is required for server authentication.
--utf8_encoding=value	On or off, 1 or 0, and so on. Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode. See “UltraLite character sets” on page 23 and “UltraLite utf8_encoding creation parameter” on page 164 .
-v OR --verbose	Print verbose messages.
-y OR --overwrite	Overwrite the database file without confirmation. This only applies when you use ulload to create a new database.
-? OR --help	Display the utility usage and exit.

Remarks

The ulload utility takes an input XML file generated by ulunload, ulunloadold (provided with SQL Anywhere 10), or ulxml (in UltraLite versions 8 and 9). When used along with ulunload this utility provides you with the ability to rebuild a database. An alternative method to rebuild a database is using ulunload to generate SQL statements and then use DBISQL to read them into a new database.

The XML file can contain metadata for the schema and/or metadata for the database data. **-d** ignores the schema metadata, only adding data to the *.udb* file. **-n** ignores the data and the metadata, only adding the schema to the *.udb* file.

Setting an option or specifying a certificate on the command line overrides any settings in the *xml-file* that is processed by ulload.

The ulload utility restores any synchronization profiles to the database when reading the XML.

This utility returns error codes. Any value other than 0 means that the operation failed.

See also

- [“UltraLite connection parameters” on page 165](#)
- [“UltraLite Database Unload utility \(ulunload\)” on page 227](#)
- [“Supported exit codes” on page 198](#)
- [“UltraLite global_database_id option” on page 195](#)

Example

Create a new UltraLite database file, *sample.udb*, and load it with data in *sample.xml*:

```
ulload -c DBF=sample.udb sample.xml
```

Load the data from *sample.xml* into the existing database *sample.udb*, and if an error occurs, prompt for action:

```
ulload -d -c DBF=sample.udb --onerror=prompt sample.xml
```

Create the schema and data stored in *test_data.xml* in the *sample.udb* database. Since the **-a** switch is specified, *sample.udb* must exist prior to running this command. Moreover, any schema or data that conflicts with what is already in *sample.udb* will mean the ULLOAD command will fail.

```
ulload -c DBF=sample.udb -a test_data.xml
```

UltraLite Synchronization utility (ulsync)

Synchronizes an UltraLite database with a MobiLink server. This tool can be used for testing synchronization during application development.

Syntax

```
ulsync -c [ options ] [ synchronization parameters]
```

Option	Description
<i>@data</i>	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used. See “Configuration files” [SQL Anywhere Server - Database Administration] .
-c " <i>connection-string</i> "	Required. Connect to the database as identified in the DBF or file_name parameter of your <i>connection-string</i> . If you do not specify both a user ID and a password, the default UID of DBA and PWD of sql are assumed.
--log	Log operations to the specified file.
-p <i>profile-name</i> OR --profile=profile	Synchronize using the named sync profile, equivalent to: <pre>SYNCHRONIZE <i>profileName</i> MERGE <i>syncOptions</i></pre> where sync options are taken from the trailing ulsync options. For example: <pre>ulsync -p <i>profileName</i> "MobiLinkId=ml;ScriptVersion=Version001...<i>syncOptions</i>"</pre> See “Synchronization profile options” on page 224 .
-q OR --quiet	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-r OR --result	Display last synchronization results and exit.
-v OR --verbose	Display synchronization progress messages. This also determines whether progress is displayed for any synchronization, whether using the C++ API or the SQL SYNCHRONIZE PROFILE statement. See “CREATE SYNCHRONIZATION PROFILE statement [UltraLite]” on page 424 .

Remarks

When a certificate file is specified either with the **trusted_certificate** or the **e2ee_public_key** option, the UltraLite runtime only looks for these files in the main resource bundle, which is a part of every iPhone application deployment package. You add items to this bundle by including them in the */Resources* folder in your Xcode project. This is not applicable to certificates that are stored in the UltraLite database, and does not affect Mac OS clients (only iPhone). See “[trusted_certificate](#)” [*MobiLink - Client Administration*] and “[e2ee_type](#)” [*MobiLink - Client Administration*].

The following options that were valid for versions 10 and earlier are no longer supported: *-a authenticate-parameters*, *-e sync-parms*, *-k stream-type*, *-n* (no sync), and *-x protocol options*. *-e keyword=value* is now part of the sync parameters string and *-k* and *-x* are now part of the **Stream= stream{stream-parms}** sync parameters string.

ulsync can be considered to be equivalent to one the following SQL statements, depending on usage:

```
ulsync -p profile "parms"
```

is equivalent to:

```
SYNCHRONIZE PROFILE profile MERGE
parms
```

and

```
ulsync "parms"
```

is equivalent to:

```
SYNCHRONIZE USING <parms>
```

For secure synchronization, the UltraLite application must have access to the public certificate. You can reference a certificate by:

- Incorporating the certificate information into the UltraLite database at creation time with the **-t file** option using `ulinit` or `ulload`.
- Referencing an external certificate file at synchronization time with the **trusted_certificate=file** stream option.

This utility returns error codes. Any value other than 0 means that the operation failed.

See also

- “[Synchronization profile options](#)” on page 224
- “[End-to-end encryption](#)” [*SQL Anywhere Server - Database Administration*]
- “[trusted_certificate](#)” [*MobiLink - Client Administration*]
- “[UltraLite connection parameters](#)” on page 165
- “[UltraLite clients](#)” on page 67
- “[Supported exit codes](#)” on page 198
- “[MobiLink File Transfer utility \(mlfiletransfer\)](#)” [*MobiLink - Client Administration*]

Examples

The following command synchronizes a database file called *myuldb.udb* for a MobiLink user called **remoteA**.

```
ulsync -c DBF=myuldb.udb "MobiLinkUid=remoteA;Stream=http;ScriptVersion=2"
```

The following command synchronizes a database file called *myuldb.udb* over HTTPS with the *c:\certs\rsa.crt* certificate. The **trusted_certificate=file** option must be used because the trusted certificate file was not added to the database when the database was created. Additionally, the MobiLink user name is **remoteB**.

```
ulsync -c DBF=myuldb.udb "Stream=https{trusted_certificate=c:\certs\rsa.crt};MobiLinkUid=remoteB;ScriptVersion=2;UploadOnly=ON"
```

The following command displays the last synchronization results for a database file named *synced.udb*.

```
ulsync -r -c dbf=synced.udb
```

The previous synchronization results are listed as follows:

```
SQL Anywhere UltraLite Database Synchronize Utility Version XX.X
Results of last synchronization:
Succeeded
  Download timestamp: 2006-07-25 16:39:36.708000
  Upload OK
  No ignored rows
  Partial download retained
  Authentication value: 1000 (0x3e8)
```

The following example shows the command line used to synchronize the CustDB database with a user name of 50 over TCP/IP on a port of 2439. It uses verbose progress messages.

```
ulsync -c "dbf=C:\Documents and Settings\All Users\Documents\SQL Anywhere
12\Samples\UltraLite\SyncEncrypt\custdb.udb"
"MobiLinkUid=50;ScriptVersion=custdb 12.0;Stream=tcPIP{port=2439}"
```

The following command illustrates how to use TLS encryption with E2EE:

```
ulsync -c "uid=dba;pwd=sql;dbf=myuldb.db"
"MobiLinkUid=rem1;MobiLinkPwd=password;ScriptVersion=v1;Stream=tls{host=mySer
ver;port=2439;trusted_certificate=c:
\clientcert.pem;e2ee_type=rsa;e2ee_public_key=c:\e2eepublic.pem}"
```

Synchronization profile options

You specify synchronization profile options with the `ulsync` utility on the command line after you have defined all other command line options you want to use. The keywords are case insensitive.

Synchronization profile option	Valid values	Description
AllowDownloadDupRows	Boolean	This option prevents errors from being raised when multiple rows are downloaded that have the same primary key. This can be used to allow inconsistent data to be synchronized without causing the synchronization to fail. The default value is "no." See “Additional Parameters synchronization parameter” on page 88
AuthParms	String (comma separated)	Specifies the list of authentication parameters sent to the MobiLink server. You can use authentication parameters to perform custom authentication in MobiLink scripts. See “Authentication Parameters synchronization parameter” on page 90.
CheckpointStore	Boolean	Adds additional checkpoints of the database during synchronization to limit database growth during the synchronization process. See “Additional Parameters synchronization parameter” on page 88.
ContinueDownload	Boolean	Restarts a previously failed download. When continuing a download, only the changes that were selected to be downloaded with the failed synchronization are received. By default, UltraLite does not continue downloads. See “Resumption of failed downloads” [<i>MobiLink - Server Administration</i>].
DisableConcurrency	Boolean	Disallow database access from other threads during synchronization. See “Additional Parameters synchronization parameter” on page 88.
DownloadOnly	Boolean	Performs a download-only synchronization. See “Download Only synchronization parameter” on page 92.
KeepPartialDownload	Boolean	Controls whether UltraLite keeps a partial download if a communication error occurs. By default, UltraLite does not roll back partially downloaded changes. See “Keep Partial Download synchronization parameter” on page 94.
MobiLinkPwd	String	Specifies the existing MobiLink password associated with the user name. See “MobiLinkPwd (mp) extended option” [<i>MobiLink - Client Administration</i>].
MobiLinkUid	String	Specifies the MobiLink user name. See “-u dbmsync option (deprecated)” [<i>MobiLink - Client Administration</i>] and “-mn dbmsync option” [<i>MobiLink - Client Administration</i>].
NewMobiLinkPwd	String	Supplies a new password for the MobiLink user. Use this option when you want to change an existing password. See “-mn dbmsync option” [<i>MobiLink - Client Administration</i>].

Synchronization profile option	Valid values	Description
Ping	Boolean	Confirms communications with the server only; no synchronization is performed. See “Ping synchronization parameter” on page 99.
Publications	String (comma separated)	Specifies the publications(s) to synchronize. The publications determine the tables on the remote database that are involved in synchronization. If this parameter is blank (the default) then all tables are synchronized. If the parameter is an asterisk (*) then all publications are synchronized. See “Publishing data in UltraLite” on page 76.
ScriptVersion	String	Specifies the MobiLink script version. The script version determines which scripts are run by MobiLink on the consolidated database during synchronization. If you do not specify a script version, 'default' is used. See “ScriptVersion (sv) extended option” [<i>MobiLink - Client Administration</i>].
SendColumnNames	String	Specifies that column names should be sent to the MobiLink server as part of the upload file when synchronizing. By default, column names are not sent. See “Send Column Names synchronization parameter” on page 102.
SendDownloadACK	Boolean	Specifies that a download acknowledgement should be sent from the client to the server. By default, the MobiLink server does not provide a download acknowledgement. See “Send Download Acknowledgement synchronization parameter” on page 102.
Stream	String (with sub-list)	Specifies the MobiLink network synchronization protocol. See “Stream Type synchronization parameter” on page 105.
TableOrder	String (comma separated)	Specifies the order of tables in the upload. By default, UltraLite selects an order based on foreign key relationships. See “Additional Parameters synchronization parameter” on page 88.
UploadOnly	String	Specifies that synchronization will only include an upload, and no download will occur. See “Upload Only synchronization parameter” on page 109.

The Boolean values can be specified as Yes/No, 1/0, True/False, On/Off. In all the Boolean cases, the default is No. For all other values, the default is simply unspecified.

See also

- “ALTER SYNCHRONIZATION PROFILE statement [UltraLite]” on page 413
- “DROP SYNCHRONIZATION PROFILE statement [UltraLite]” on page 434
- “SYNCHRONIZE statement [UltraLite]” on page 448
- “UltraLite creation parameters” on page 139

UltraLite Database Unload utility (ulunload)

Unloads any of the following, depending on the options used:

- An entire UltraLite database to XML or SQL.
- All or part of UltraLite data only to XML or SQL.

Syntax

ulunload -c "*connection-string*" [*options*] *output-file*

Option	Description
@ <i>data</i>	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.
-b <i>max-size</i> OR --maxblob = <i>max-size</i>	Set the maximum size of column data to be stored in the XML file. The default is 10 KB. To store all data in the XML file (no maximum size), use -b -1 .
-c " <i>keyword=value;...</i> " OR --connect =" <i>keyword=value;...</i> "	Required. Connect to the database as identified in the DBF or file_name parameter of your <i>connection-string</i> . If you do not specify both a user ID and a password, the default UID of DBA and PWD of sql are assumed.
-d OR --dataonly	Only unload the data from the database to the output file. Do not unload any schema information.
-e <i>table,...</i> OR --exclude = <i>table,...</i>	Exclude the named <i>table</i> when unloading the database. You can name multiple tables in a comma-separated list. For example: <code>-e mydbtable1,mydbtable5</code>

Option	Description
-f <i>directory</i> OR --filedir = <i>directory</i>	Set the directory to store data larger than the maximum size specified by -b. The default is the same directory as the output file.
-l <i>filename</i> OR --log = <i>filename</i>	Unload schema only, ignoring any data in the database.
-n OR --schemaonly	Unload schema only, ignoring any data in the database.
-q OR --quiet	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-s OR --sql	Unload as SQL Anywhere-compatible SQL statements. SQL file output can be read by UltraLite or SQL Anywhere using DBISQL.
-t <i>table,...</i> OR --include = <i>table,...</i>	Unload data in the named <i>table</i> only. You can name multiple tables in a comma separated list. For example: <pre>-t mydbtable2,mydbtable6</pre>
-v OR --verbose	Print verbose messages.
-x <i>owner</i> OR --owner = <i>owner</i>	Output tables so they are owned by a specific user ID. You can use this option with the -s option.

Option	Description
-y OR --overwrite	Overwrite <i>output-file</i> without confirmation.
<i>output-file</i>	Required. Set the name of the file that the database is unloaded into. If you use the -s option, database is unloaded as SQL statements. Otherwise, the database is unloaded as XML.

Remarks

By default, ulunload outputs XML that describes the schema and data in the database. You can use the output for archival purposes, or to keep the UltraLite database portable across all releases.

Saving a database with a synchronization profile results in XML that is incompatible with earlier versions of the UltraLite utilities. A workaround is to edit the XML and remove the text section marked with

```
<syncprofiles>...</syncprofiles>
```

Unloading a database does not preserve:

- Synchronization state, stored synchronization counts, and row deletions. Ensure you synchronize the database before unloading it.
- UltraLite user entries.

To confirm what database options or properties have been preserved, run ulinfo after you have reloaded your database with the ulload utility.

If column data exceeds the maximum size you specified with -b, the overflow is saved to a *.bin file in either:

- the same directory as the XML file
- the directory specified by -f.

The file follows this naming convention:

```
tablename-columnname-rownumber.bin
```

The -x option allows you to assign ownership to UltraLite tables. You only need to assign an owner to a table if you intend to use the resulting SQL statements for creating or modifying a SQL Anywhere database. When read by UltraLite, the owner names are silently ignored.

This utility returns error codes. Any value other than 0 means that the operation failed.

If you are using this utility to unload a database on the Windows Mobile device directly, UltraLite cannot back up the database before the unload or action occurs. You must perform this action manually before running these wizards.

See also

- [“Configuration files” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite connection parameters” on page 165](#)
- [“Supported exit codes” on page 198](#)
- [“UltraLite Load XML to Database utility \(uload\)” on page 217](#)
- [“UltraLite Information utility \(ulinfo\)” on page 209](#)

Example

Unload the *sample.udb* database into the *sample.xml* file.

```
ulunload -c DBF=sample.udb sample.xml
```

Unload the data from the *sample.udb* database into a SQL file called *sample1.sql*. Overwrite the SQL file if it exists.

```
ulunload -c DBF=sample.udb -d -y -s sample.sql
```

UltraLite Validate Database utility (ulvalid)

Performs a full (normal) validation of an UltraLite database.

Syntax

```
ulvalid -c "connection-string" [ options ]
```

Option	Description
@data	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.
-c "connection-string"	Required. Connect to the database as identified in <i>connection-string</i> . If you do not specify both a user ID and a password, the default UID of DBA and PWD of sql are assumed.
-e	Express validation. Only perform table validation. This option provides a faster validation than normal validation.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-v	Print verbose messages.
-?	Displays utility usage information and exits.

Remarks

Validating a database verifies the accuracy of the table metadata and ensures the file has not been corrupted.

The validation includes:

- **Database pages** Validate all database pages, using checksums when enabled. Note that certain critical pages always have checksums and even pages without checksums undergo a basic validity check.
- **Tables** Validate table(s) by checking that the table row count matches the count in each index.
- **Indexes** Validate indexes by checking that entries refer to valid rows. `ulvalid -e` performs an express check, which includes only table validation.

See also

- [“UltraLite checksum_level creation parameter” on page 142](#)
- [“Configuration files” \[SQL Anywhere Server - Database Administration\]](#)
- [“Validating an UltraLite database” on page 42](#)

Example

An example of an express validation of a database named *sample.udb* run in quiet mode.

```
ulvalid -c DBF=sample.udb -e -q
```

UltraLite system tables

The schema of an UltraLite database is stored in a proprietary format. Earlier versions of UltraLite databases were stored in several system tables. These system tables can still be queried for backward compatibility (they are in essence system views), but they only contain information about user schema (like tables, columns, indexes) not system schema. For example, you cannot query `systable` to find the properties of `systable` itself. You can only query `systable` to find the properties of user-created tables.

Each UltraLite programming API supports objects and methods that can be used to query the database about its schema. It is recommended that you use these objects and APIs to explore schema rather than querying the system views.

Note that all queries performed on these system views are equivalent to full table scans. Index scans are not supported on these system views.

systable system table

Each row in the `systable` system table describes one table in the database.

Column name	Column type	Description
<code>column_count</code>	UNSIGNED INT	The number of columns in the table.

Column name	Column type	Description
index_count	UNSIGNED INT	The number of indexes in the table.
ixcol_count	UNSIGNED INT	The total number of columns in all indexes in the table.
table_name	VARCHAR(128)	The name of the table.
object_id	UNSIGNED INT	A unique identifier for that table.
sync_type	VARCHAR(32)	Used for MobiLink synchronization. Can be one of either no_sync for no synchronization, all_sync to synchronize every row, or normal_sync for synchronize changed rows only.
table_type	VARCHAR(32)	user to indicate user-created tables.

Constraints

PRIMARY KEY (object_id)

syscolumn system table

Each row in the syscolumn system table describes one column.

Column name	Column type	Description
column_name	VARCHAR(128)	The name of the column.
default	VARCHAR(128)	The default value for this column. For example, autoincrement.
domain	UNSIGNED INT	The column domain, which is an enumerated value indicating the domain of the column.
domain_info	UNSIGNED INT	Used with a variable sized domain.
nulls	VARCHAR(1)	Determines if the column allows nulls default.
object_id	UNSIGNED INT	A unique identifier for that column.
table_id	UNSIGNED INT	The identifier of the table to which the column belongs.

Constraints

PRIMARY KEY(table_id, object_id)

FOREIGN KEY (table_id) REFERENCES systable (object_id)

sysindex system table

Each row in the sysindex system table describes one index in the database.

Column name	Column type	Description
check_on_commit	BIT	Indicates when referential integrity is checked to ensure there is a matching primary row for every foreign key. It is only required if type is foreign .
index_name	VARCHAR(128)	The name of the index.
ixcol_count	UNSIGNED INT	The number of columns in the index.
nullable	BIT	Only required if type is foreign . Indicates if nulls are allowed.
object_id	UNSIGNED INT	A unique identifier for an index.
primary_index_id	UNSIGNED INT	Only required if type is foreign . Lists the identifier of the primary index.
primary_table_id	UNSIGNED INT	Only required if type is foreign . Lists the identifier of the primary table.
root_handle	UNSIGNED INT	For internal use only.
table_id	UNSIGNED INT	A unique identifier for the table to which the index applies.
type	VARCHAR(10)	The type of index. Can be one of: <ul style="list-style-type: none"> ● primary ● foreign ● key ● unique ● index
hash_size	UNSIGNED SHORTINT	Stores the hash size used for index hashing.

Constraints

PRIMARY KEY (table_id, object_id)

FOREIGN KEY(table_id) REFERENCES systable(object_id)

See also

- [“sysixcol system table” on page 234](#)

sysixcol system table

Each row in the sysixcol system table describes one column of an index listed in sysindex.

Column name	Column type	Description
column_id	UNSIGNED INT	A unique identifier for the column being indexed.
index_id	UNSIGNED INT	A unique identifier for the index that this index-column belongs to.
order	VARCHAR(1)	Indicates whether the column in the index is kept in ascending (A) or descending (D) order.
sequence	UNSIGNED INT	The order of the column in the index.
table_id	UNSIGNED INT	A unique identifier for the table to which the index applies.

Constraints

PRIMARY KEY(table_id, index_id, sequence)

FOREIGN KEY(table_id, index_id) REFERENCES sysindex(table_id, object_id)

FOREIGN KEY(table_id, column_id) REFERENCES syscolumn(table_id, object_id)

See also

- [“sysindex system table” on page 233](#)

syspublication system table

Each row in the syspublication system table describes a publication.

Column name	Column type	Description
download_timestamp	TIMESTAMP	The time of the last download.
last_sync	UNSIGNED BIGINT	Used to keep track of upload progress.
publication_id	UNSIGNED INT	A unique identifier for the publication.
publication_name	VARCHAR(128)	The name of the publication.

Constraints

PRIMARY KEY (publication_id)

See also

- [“sysarticle system table” on page 235](#)

sysarticle system table

Each row in the sysarticle system table describes a table that belongs to a publication.

Column name	Column type	Description
publication_id	UNSIGNED INT	An identifier for the publication that this article belongs to.
table_id	UNSIGNED INT	The identifier of the table that belongs to the publication.
where_expr	VARCHAR(256)	An optional predicate to filter rows.

Constraints

PRIMARY KEY (publication_id, table_id)

FOREIGN KEY (publication_id) REFERENCES syspublication (publication_id)

FOREIGN KEY (table_id) REFERENCES systable (object_id)

See also

- [“syspublication system table” on page 234](#)

UltraLite Java edition database properties

UltraLite Java edition database property values are defined when the database is first created. They can be changed by re-creating the UltraLite Java edition database or editing their corresponding database option, if available.

UltraLite Java edition supports the following database properties:

Property	Description
blob_file_base_dir	Returns the base directory the database uses to search for BLOB files. See also: <ul style="list-style-type: none"> • “UltraLite Java edition blob_file_base_dir option” on page 238 • “UltraLite Java Edition Database Load utility (ULjLoad)” on page 244
date_format	Returns the date format the database uses for string conversions. See “UltraLite date_format creation parameter” on page 144 .
date_order	Returns the date order the database uses for string conversions. See “UltraLite date_order creation parameter” on page 147 .
global_database_id	Returns the value of the global_database_id option used for global autoincrement columns. See “UltraLite global_database_id option” on page 195 .
ml_remote_id	Returns the value of the ml_remote_id option that uniquely identifies the database for MobiLink synchronization. See “UltraLite ml_remote_id option” on page 196 .
database_name	Returns the name (or alias) of the database for the current connection. The name returned matches the DBN connection parameter value. If you did not use the DBN connection parameter, the name returned is the database file without the path and extension. See also: <ul style="list-style-type: none"> • “UltraLite DBN connection parameter” on page 176 • “UltraLite DBF connection parameter” on page 174
nearest_century	Returns the nearest century the database uses for string conversions. See “UltraLite nearest_century creation parameter” on page 151 .
page_size	Returns the page size of the database, in bytes. See “UltraLite page_size creation parameter” on page 153 .
precision	Returns the floating-point precision the database uses for string conversions. See “UltraLite precision creation parameter” on page 155 .

Property	Description
scale	Returns the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision during string conversions by the database. See “UltraLite scale creation parameter” on page 156.
time_format	Returns the time format the database uses for string conversions. See “UltraLite time_format creation parameter” on page 157.
timestamp_format	Returns the timestamp format the database uses for string conversions. See “UltraLite timestamp_format creation parameter” on page 159.
timestamp_increment	Returns the minimum difference between two unique timestamps, in microseconds. See “UltraLite timestamp_increment creation parameter” on page 162.
timestamp_with_time_zone	Returns the timestamp format for <code>TIMESTAMP WITH TIME ZONE</code> values. See “UltraLite timestamp_with_time_zone_format creation parameter” on page 163.

See also

- [“Reading database properties”](#) on page 33
- [“Accessing database options”](#) on page 34
- [“UltraLite creation parameters”](#) on page 139
- [“UltraLite Java edition database options”](#) on page 237

UltraLite Java edition database options

UltraLite Java edition database option values are defined when the database is first created. They can be altered while connected to the database.

This section describes the UltraLite database options that are available. In addition to these options, the following UltraLite database options are supported by UltraLite Java edition databases:

- `global_database_id`
- `ml_remote_id`

See also

- [“UltraLite global_database_id option”](#) on page 195
- [“UltraLite ml_remote_id option”](#) on page 196
- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“UltraLite creation parameters”](#) on page 139

UltraLite Java edition blob_file_base_dir option

Sets the base directory the database uses to search for BLOB files.

Allowed values

String

Default

The default for BlackBerry smartphones is `"file:///SDCard"`. The default for other platforms is `""`.

Remarks

The table column you used to store the file name is stored as the string `<base_dir>/<reference_to_external_file>`, where `<base_dir>` is the base directory where the blobfile is located, and `<reference_to_external_file>` is the exact string found in the contents column in the XML file.

The base directory can be set initially using the `-f` option of the UltraLite Java edition database load utility. Relative filenames are resolved with the `blob_file_base_dir` option value. If the file name does not begin with the prefix `file://`, UltraLite then prepends the file name with the option value before attempting to open it.

See also

- [“UltraLite Java Edition Database Load utility \(ULjLoad\)” on page 244](#)
- [“Accessing database options” on page 34](#)
- [“Reading database properties” on page 33](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)
- [“Connection.OPTION_BLOB_FILE_BASE_DIR variable \[BlackBerry\] \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)

UltraLite Java edition date_format option

Sets the date format the database uses for string conversions.

Allowed values

String

Default

`"YYYY-MM-DD"`

Remarks

For more information, see [“UltraLite date_format creation parameter” on page 144](#).

See also

- “Accessing database options” on page 34
- “Reading database properties” on page 33
- “SET OPTION statement [UltraLite]” on page 446
- “Connection.getDatabaseProperty method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “Connection.setOption method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “Connection.OPTION_DATE_FORMAT variable [BlackBerry] [UltraLiteJ]” [*UltraLite - Java Programming*]

UltraLite Java edition date_order option

Sets the date order the database uses for string conversions.

Allowed values

String

Default

"YMD"

Remarks

For more information, see “UltraLite date_order creation parameter” on page 147.

See also

- “Accessing database options” on page 34
- “Reading database properties” on page 33
- “SET OPTION statement [UltraLite]” on page 446
- “Connection.getDatabaseProperty method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “Connection.setOption method [UltraLiteJ]” [*UltraLite - Java Programming*]
- “Connection.OPTION_DATE_ORDER variable [BlackBerry] [UltraLiteJ]” [*UltraLite - Java Programming*]

UltraLite Java edition nearest_century option

Sets the nearest century the database uses for string conversions.

Allowed values

String

Default

"50"

Remarks

For more information, see [“UltraLite nearest_century creation parameter”](#) on page 151.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_NEAREST_CENTURY variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition precision option

Sets the floating-point precision the database uses for string conversions.

Allowed values

String

Default

"30"

Remarks

For more information, see [“UltraLite precision creation parameter”](#) on page 155.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_PRECISION variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition scale option

Sets the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision during string conversions by the database.

Allowed values

String

Default

"6"

Remarks

For more information, see [“UltraLite scale creation parameter”](#) on page 156.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_SCALE variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition time_format option

Sets the time format the database uses for string conversions.

Allowed values

String

Default

"HH:NN:SS.SSS"

Remarks

For more information, see [“UltraLite time_format creation parameter”](#) on page 157.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_TIME_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition timestamp_format option

Sets the timestamp format the database uses for string conversions.

Allowed values

String

Default

"YYYY-MM-DD HH:NN:SS.SSS"

Remarks

For more information, see [“UltraLite timestamp_format creation parameter”](#) on page 159.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_TIMESTAMP_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition timestamp_increment option

Sets the minimum difference between two unique timestamps, in microseconds.

Allowed values

String

Default

"1"

Remarks

For more information, see [“UltraLite timestamp_increment creation parameter”](#) on page 162.

See also

- [“Accessing database options”](#) on page 34
- [“Reading database properties”](#) on page 33
- [“SET OPTION statement \[UltraLite\]”](#) on page 446
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.setOption method \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]
- [“Connection.OPTION_TIMESTAMP_INCREMENT variable \[BlackBerry\] \[UltraLiteJ\]”](#) [*UltraLite - Java Programming*]

UltraLite Java edition timestamp_with_time_zone_format option

Sets the timestamp format for `TIMESTAMP WITH TIME ZONE` values.

Allowed values

String

Default

"YYYY-MM-DD HH:NN:SS.SSS+HH:NN"

Remarks

For more information, see [“UltraLite timestamp_with_time_zone_format creation parameter” on page 163.](#)

See also

- [“Accessing database options” on page 34](#)
- [“Reading database properties” on page 33](#)
- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- [“Connection.getDatabaseProperty method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.setOption method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_TIMESTAMP_WITH_TIME_ZONE_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

UltraLite Java edition utilities

UltraLite utilities are supplied to perform maintenance and administration tasks on UltraLite Java edition databases.

UltraLite Java Edition Database Information utility (ULjInfo)

Displays information about an existing UltraLite Java edition database.

Syntax

```
ULjInfo -c filename -p password [ options ]
```

Option	Description
-c filename	Required. The filename of the UltraLite Java edition database to examine.
-p password	The password to connect to the UltraLite Java edition database. The default is sql .
-?	Display command line usage information.

Example

The following is an example of the output from the ULjInfo program:

```
C:\ULj\bin>ULjInfo.cmd -c ..\Samples\Demol.ulj -p sql
SQL Anywhere UltraLite J Database Information Utility
Database name: ..\Samples\Demol.ulj
Disk file: '..\Samples\Demol.ulj'
Database ID: 0
Page size: 1024
0 rows for next upload
Date format: YYYY-MM-DD
Date order: YMD
Nearest century: 50
Numeric precision: 30
Numeric scale: 6
Time format: HH:NN:SS.SSS
Timestamp format: YYYY-MM-DD HH:NN:SS.SSS
Timestamp increment: 1
Number of tables: 1
Number of columns: 2
Number of publications: 0
Number of tables that will always be uploaded: 0
Number of tables that are never synchronized: 0
Number of primary keys: 1
Number of foreign keys: 0
Number of indexes: 0
Last download occurred on Thu Jul 05 11:31:05 EDT 2007
Upload OK: true
```

UltraLite Java Edition Database Load utility (ULjLoad)

Provides the capability to load an UltraLite Java edition database from an XML source file. The XML file is often produced by the ULjUnload utility and is customizable.

Note
 Database configuration parameters cannot be set in the command line. ULjLoad uses the default configuration when a new database is created. If you need to create a database with different configuration parameters, such as a different page size, you must create the database in Java first and then use **ULjLoad -a** to load the XML content into the database you created.

Syntax

ULjLoad -c filename -p password [options] inputfile

Option	Description
-a	Adds information from the XML file to an existing database. If this option is not specified, a new database is created.
-c filename	Required. Specifies the name of the database file.
-d	Loads data only; ignore schema information.
-f directory	Specifies the directory to retrieve data for columns larger than max blob size specified with -b option during ULjUnload.

Option	Description
-i	Inserts rows for upload synchronization.
-n	Loads schema information only; ignores row data.
-p <i>password</i>	Optional. Specifies the password. The default password is sql .
-q	Runs the utility in quiet mode—does not display messages.
-v	Displays verbose messages.
-y	Overwrites the output file if it exists (and -a option is not specified).
-z <i>pagesize</i>	Specifies the page size, in bytes, to create the database with.
-?	Displays command line usage information.
<i>inputfile</i>	Specifies the input file that contains XML statements.

Example for loading blobfile types

This example assumes that you have applied the following SQL statement to your UltraLite Java edition database and unloaded the file using the UltraLite Java edition database unload utility.

```
CREATE TABLE blobfile_example
  file_name CHAR(size) DEFAULT AUTOFILENAME( prefix, extension ),
  file_contents LONG BINARY STORE AS FILE( file_name ) [ CASCADE DELETE ]
```

ULLoad, while encountering a file exported by ULjUnload, treats the file_name column as a CHAR column and the file_contents column as a LONG BINARY column.

ULjLoad, while encountering the same XML file, reconstructs the blobfile type. The file_name column stores the string **<base_dir>/<reference_to_external_file>**, where **<base_dir>** is the exact string in the -f option, and **<reference_to_external_file>** is the exact string found in the file_contents column in the XML file. The newly constructed database then contains valid references to the external files.

The load fails with an error message prompting the user to increase the size of the file_name column if the generated file name is too long.

See also

- [“UltraLite Java Edition Database Unload utility \(ULjUnload\)” on page 246](#)

Example

The -? option displays command line information.

```
uljload -?
```

When the ULjLoad utility command line includes the -? option, the following usage information appears:

```
SQL Anywhere UltraLiteJ Database Load Utility
Usage: uljload [options] <XML file>
```

Create and load data into a new UltraLiteJ database from <XML file>.

Options:

```
-a      Add to existing database.
-c <file> Database file.
-d      Data only -- ignore schema.
-f <directory>
        Directory for loading .File paths.
-i      Insert rows for upload synchronization.
-n      Schema only -- ignore data.
-p      Password to connect to database.
-q      Quiet: do not print messages.
-v      Verbose messages.
-z <page size>
        Page size to create the database with.
-y      Overwrite file if it already exists.
```

UltraLite Java Edition Database Unload utility (ULjUnload)

Provides the capability to unload an UltraLite Java edition database—either the data, the schema, or both—to an XML file.

Syntax

ULjUnload -c *filename* -p *password* [*options*] *outputfile*

Options	Description
-b <i>max-blob-size</i>	Maximum size (in bytes) of blob/char data output to XML.
-c <i>filename</i>	Required. Name of database file to unload.
-d	Unload only data; do not output schema information.
-e <i>table, ...</i>	Exclude data for tables named in list.
-f <i>directory</i>	Directory to store data for columns larger than max blob size specified with -b option.
-n	Unload schema information only; do not output data.
-p <i>password</i>	Password to connect to database. The default is sql .
-q	Run in quiet mode—do not display messages.
-t <i>table, ...</i>	Only output data for tables named in list.
-v	Display verbose messages.
-y	Overwrite output file if it already exists.
-?	Display option usage/help information.

Options	Description
<i>outputfile</i>	Output file name (this file contains xml statements that describe the database contents).

Example for unloading blobfile types

Apply the following sample SQL statement to your database to unload a blobfile type:

```
CREATE TABLE blobfile_example
  file_name CHAR(size) DEFAULT AUTOFILENAME( prefix, extension ),
  file_contents LONG BINARY STORE AS FILE( file_name ) [ CASCADE DELETE ]
```

When applying this example and using this utility, ULjUnload exports the file_name column as a regular CHAR but with an additional attribute, default_autofilename, which stores the prefix and extension strings in the form 'prefix', 'extension'. The file_contents column is exported as a LONG BINARY but with an additional attribute, filename_col, which stores the name of the referenced file_name column. In the row containing the blobfile columns, the content of the file_name column is unchanged. The file_contents column behaves like an externally-stored blob column and has the following form:

```
file_contents.File="tablename-columnname-rownumber.bin"
```

The contents of the file_contents column are saved as bin files to the location specified by the -f option.

See also

- [“UltraLite Java Edition Database Unload utility \(ULjUnload\)” on page 246](#)

Example XML file contents

```
<?xml version="1.0" encoding="utf-8" standalone="no"?>
<ul:ulschema xmlns:ul="urn:ultralite">
  <collation name="1252LATIN1" case_sensitive="no"/>
  <options>
    <option name="dateformat" value="YYYY-MM-DD"/>
    <option name="dateorder" value="YMD"/>
    <option name="nearestcentury" value="50"/>
    <option name="precision" value="30"/>
    <option name="scale" value="6"/>
    <option name="timeformat" value="HH:NN:SS.SSS"/>
    <option name="timestampformat" value="YYYY-MM-DD HH:NN:SS.SSS"/>
    <option name="timestampincrement" value="1"/>
  </options>
  <tables>
    <table name="ULCustomer" sync="changes">
      <columns>
        <column name="cust_id" type="integer" null="no"/>
        <column name="cust_name" type="char(30)" null="yes"/>
      </columns>
      <primarykey>
        <primarycolumn name="cust_id" direction="asc"/>
      </primarykey>
      <indexes/>
    </table>
  </tables>
  <uldata>
    <table name="ULCustomer">
      <row cust_id="2000" cust_name="Apple St. Builders"/>
      <row cust_id="2001" cust_name="Art's Renovations"/>
    </table>
  </uldata>
</ul:ulschema>
```

```
<row cust_id="2002" cust_name="Awnings R Us"/>
<row cust_id="2003" cust_name="Al's Interior Design"/>
<row cust_id="2004" cust_name="Alpha Hardware"/>
<row cust_id="2005" cust_name="Ace Properties"/>
<row cust_id="2006" cust_name="Al Contracting"/>
<row cust_id="2007" cust_name="Archibald Inc."/>
<row cust_id="2008" cust_name="Acme Construction"/>
<row cust_id="2009" cust_name="ABCXYZ Inc."/>
<row cust_id="2010" cust_name="Buy It Co."/>
<row cust_id="2011" cust_name="Bill's Cages"/>
<row cust_id="2012" cust_name="Build-It Co."/>
<row cust_id="2013" cust_name="Bass Interiors"/>
<row cust_id="2014" cust_name="Burger Franchise"/>
<row cust_id="2015" cust_name="Big City Builders"/>
<row cust_id="2016" cust_name="Bob's Renovations"/>
<row cust_id="2017" cust_name="Basements R Us"/>
<row cust_id="2018" cust_name="BB Interior Design"/>
<row cust_id="2019" cust_name="Bond Hardware"/>
<row cust_id="2020" cust_name="Cat Properties"/>
<row cust_id="2021" cust_name="C & C Contracting"/>
<row cust_id="2022" cust_name="Classy Inc."/>
<row cust_id="2023" cust_name="Cooper Construction"/>
<row cust_id="2024" cust_name="City Schools"/>
<row cust_id="2025" cust_name="Can Do It Co."/>
<row cust_id="2026" cust_name="City Corrections"/>
<row cust_id="2027" cust_name="City Sports Arenas"/>
<row cust_id="2028" cust_name="Cantaloupe Interiors"/>
<row cust_id="2029" cust_name="Chicken Franchise"/>
</table>
</uldata>
</ul:ulschema>
```

UltraLite Java Edition Database Transfer utility (ULjDbT)

The ULjDbT utility provides the capability to transfer an UltraLite database from a BlackBerry smartphone to an external device, such as a desktop, laptop, or server. In addition, you can delete a database, display database information, or view or email the database transfer log. The utility consists of two applications that must run simultaneously—the UltraLite Java Edition Database Transfer desktop application (ULjDbT) and the BlackBerry smartphone client application (*ULjDatabaseTransfer.cod*).

The UltraLiteJ Database Transfer desktop application

The desktop application receives UltraLite Java edition databases using a USB or HTTP connection method. When you start the server application, it waits for a BlackBerry smartphone to transfer the database through the specified connection with the client application. The connection is closed either manually through the application interface, when the application times out, or when the transfer is complete.

The BlackBerry Smartphone client application

The BlackBerry Smartphone client application sends UltraLite Java edition databases through a USB cable or a specified TCP port to the desktop application.

In addition, you can delete a database, display database information, or view or email the database transfer log.

Start the client application

The client application is a signed file located in the *UltraLite\UltraLiteJ\BlackBerry4.2* directory of your SQL Anywhere installation.

1. Load *ULjDatabaseTransfer.cod* from the *UltraLite\UltraLiteJ\BlackBerry4.2* directory of your SQL Anywhere installation.

The client application icon appears in your list of applications.

2. Start the application and press the trackwheel.
3. On the **Database Connection** screen, complete the following fields:
 - **Database Name** The name of the database to transfer to the external device. If the name starts with *file://* (this is case sensitive), then the client application tries to find the database in the file system; otherwise, it finds the database in the object store.
 - **Database Password** The database password used to allow data transfer. If you leave **Database Password** blank, the default password is used.
4. Click **Next**.

The **Action** screen appears. This is the screen from which you can access all of the client application functionality.

Transfer a database using the BlackBerry Smartphone client application

1. On the **Action** screen, click the desired connection method (USB or HTTP).
2. For a USB transfer, click **USB database transfer**. For an HTTP transfer, proceed to step 4.
3. Follow the directions to start the database transfer desktop application.

Note

To ensure a successful database transfer, make sure that the device or simulator is connected to the BlackBerry Device Manager. For a simulator, make sure that a USB connection is simulated using **USB Cable Connected**.

- a. Click **Next** on the client application.
 - b. On the desktop application, make sure **USB** is selected and click **Start**.

The BlackBerry smartphone starts transferring the database to the external device. Progress information appears on the desktop application.
 - c. Click **OK** on both the client and desktop applications to close them.
4. For an HTTP transfer, click **HTTP database transfer**.
 - a. On the **HTTP Transfer** screen, click **Next**.
 - b. Specify the following values:
 - **Host** The IP address of your desktop.

- **Port** The port specified in **Connection Properties** on the desktop application.
 - **URL Suffix** The hostname of the server receiving the transfer, including the **http://** suffix (this is required).
- c. Click **Next**.

Note

To transfer through the BES on a non-ident device, leave **Suffix** empty. On ident devices, use the suffix **;deviceside=false**.

To transfer through Direct TCP use the suffix **;deviceside=true**. Not all carriers support this feature.

The carrier's WAP gateway can be used if you know the APN information for it. You have to append that information to the suffix as well. Note that even if you are going through a BES, there may be a firewall between the BES and the machine you are running the UltraLite Java Edition Database Transfer utility on. In this case you need to use an SSL tunneler. On the **HTTP Transfer Parmss** screen, specify the port and name or IP address of the SSL server running on the BES-side of the firewall. You also need to specify the port the SSL client is mapped to the transfer application.

If you are transferring a database from a BlackBerry simulator, you need to have a BlackBerry MDS simulator running or specify a URL suffix of **;deviceside=true** to the UltraLite Java Edition Database Transfer utility running in the BlackBerry simulator.

- d. On the desktop application, make sure **HTTP** is selected and click **Start**.
- e. Click **Next** on the client application.
- The BlackBerry smartphone starts transferring the database to the external device. Progress information appears on the desktop application.
- f. Click **OK** on both the client and desktop applications to close them.

Receive a database using the UltraLiteJ Database Transfer application

1. Run *ULjDbTServ.cmd* from the *Bin32* directory of your SQL Anywhere installation.

The UltraLite Java Edition Database Transfer application loads.

2. On the **Connect** tab, click the desired **Connection Method**.
3. Under **Connection Properties**, specify the following values:
 - **Port** This field only applies to HTTP connections. Type the TCP port number that you want the BlackBerry smartphone to connect to. Usually, this port number matches the port number specified to the UltraLite Java Edition Database Transfer utility running on the BlackBerry smartphone; however, if you are using SSL, then this number can be different.
 - **BlackBerry Password** This field only applies to USB connections. Type the password used to access the connected BlackBerry smartphone when it is locked. Leave this field blank if there is no password.

- **Timeout** The number of idle minutes before the server application times out and closes the connection.
 - **Output** Specify a file name and location to which to save the transferred database.
4. Click **Start** to open a connection to the BlackBerry smartphone.

The server application waits until it either times out or establishes a connection. If you have specified an existing file, you are asked whether you want to overwrite it.

The **Logs** tab provides details on the server status and transfer progress, including error messages.

Delete a database

1. On the **Action** screen, click **Delete The Database**.
2. In the confirmation window, click **Delete** to delete the database.
3. In the **Database Deleted** window, click **OK** to close the client.

View database information

1. On the **Action** screen, click **View Database Info**. Scroll down to view all the database information.
2. Click **Back** to return to the **Action** screen.

View the log file

1. On the **Database Connection** screen, display the menu.
2. Click **Log**. The log screen appears.
3. To email the log file, enter the email address to send the log to and click **Send Email**. To return to the previous screen, press enter.

See also

- [“Receive a database using the UltraLiteJ Database Transfer application” on page 250](#)

UltraLite Java edition system tables

This section introduces the system tables that are available to UltraLite Java edition database.

systable system table

Each row in the systable system table describes a table in the database.

Column name	Column type	Description
table_id	INTEGER	A unique identifier for the table.
table_name	VAR-CHAR(128)	The name of the table.
table_flags	UNSIGNED SHORT	A bitwise combination of one of the following flags: <ul style="list-style-type: none"> • TableSchema.TABLE_IS_SYSTEM • TableSchema.TABLE_IS_NO_SYNC • TableSchema.TABLE_IS_DOWNLOAD_ONLY
table_data	UNSIGNED INTEGER	Internal use only.
table_partition_size	UNSIGNED INTEGER	The partition size of the table, if specified or if the table contains a default value; otherwise, null. This value is 0 if the partition size is out of range.

Constraints

PRIMARY INDEX (table_id)

syscolumn system table

Each row in the syscolumn system table describes a column.

Column name	Column type	Description
table_id	UNSIGNED INTEGER	The identifier of the table to which the column belongs.
column_id	UNSIGNED INTEGER	A unique identifier for the column.
column_name	VARCHAR(128)	The name of the column. See “Domain interface [UltraLiteJ]” [UltraLite - Java Programming] .
column_flags	TINY	A bitwise combination of the following flags describing the attributes: <ul style="list-style-type: none"> • 0x01 Column is in the primary key. • 0x02 Column is nullable.

Column name	Column type	Description
column_domain	TINY	The column domain, which is an enumerated value indicating the domain of the column in the low-order 6 bits. The remaining bits are used internally.
column_length	UNSIGNED SHORT	The column length. For VARCHAR and BINARY type columns, which are defined in the Domain interface, this is the maximum length in bytes. For NUMERIC type columns, the precision is stored in the first byte, and the scale is stored in the second byte.
column_default_value	VARCHAR(128)	The default value for this column, which is specified by one of the COLUMN_DEFAULT values in the ColumnSchema interface. For example, COLUMN_DEFAULT_AUTOINC denotes an auto-incrementing default value. If a varchar column has DEFAULT AUTOFILENAME specified, then it stores the parameters prefix and extension using the encoding prefix extension.
filename_colid	UNSIGNED INTEGER	Stores the column ID of the column that contains the file name for a LONG BINARY column declared as STORE AS FILE; otherwise, this column is null.

Constraints

PRIMARY KEY (table_id, column_id)

sysindex system table

Each row in the sysindex system table describes an index in the database.

Column name	Column type	Description
table_id	UNSIGNED INTEGER	A unique identifier for the table to which the index applies.
index_id	UNSIGNED INTEGER	A unique identifier for an index.
index_name	VARCHAR(128)	The name of the index.

Column name	Column type	Description
index_flags	TINY	A bitwise combination of the following flags denoting the type of index and its persistence: <ul style="list-style-type: none"> ● 0x01 Unique key. ● 0x02 Unique index. ● 0x04 Index is persistent. ● 0x08 Primary key.
index_data	UNSIGNED INTEGER	Internal use only.

Constraints

PRIMARY KEY (table_id, index_id)

sysindexcolumn system table

Each row in the sysindexcolumn system table describes a column of an index listed in sysindex.

Column name	Column type	Description
table_id	UNSIGNED INTEGER	A unique identifier for the table to which the index applies.
index_id	UNSIGNED INTEGER	A unique identifier for the index that this index-column belongs to.
order	UNSIGNED INTEGER	The order of the column in the index.
column_id	UNSIGNED SHORT	A unique identifier for the column being indexed.
index_column_flag	TINY	An indication of where the column in the index is kept, either in ascending (1) or descending (0) order.

Constraints

PRIMARY KEY (table_id, index_id, order)

syspublications system table

Each row in the syspublications system table describes a publication.

Column name	Column type	Description
publication_id	UNSIGNED INTEGER	A unique identifier for the publication.
publication_name	VARCHAR(128)	The name of the publication.
download_timestamp	TIMESTAMP	The time of the last download.
last_sync_sent	UNSIGNED INTEGER	An integer that tracks an upload sent to MobiLink.
last_sync_confirmed	UNSIGNED INTEGER	An integer that tracks an upload confirmed as being received by MobiLink.

Constraints

PRIMARY KEY (publication_id)

sysarticles system table

Each row in the sysarticles system table describes a table that belongs to a publication.

Column name	Column type	Description
publication_id	UNSIGNED INTEGER	An identifier for the publication that this article belongs to.
table_id	UNSIGNED INTEGER	The identifier of the table that belongs to the publication.
predicate	VARCHAR(256)	The predicate expression specified by the associated WHERE clause of a CREATE PUBLICATION or ALTER PUBLICATION statement.

Constraints

PRIMARY KEY (publication_id, table_id)

sysforeignkey system table

Each row in the sysforeignkey system table describes a foreign key that belongs to a table.

Column name	Column type	Description
table_id	UNSIGNED INTEGER	The identifier of the table to which the foreign key belongs.
foreign_table_id	UNSIGNED SHORT	The identifier of the table to which this foreign-key-column refers to.
foreign_key_id	UNSIGNED INTEGER	The identifier of the foreign key.
name	VARCHAR(128)	The name of the foreign key.
index_name	VARCHAR(128)	The name of the index the foreign key refers to.

Constraints

PRIMARY KEY (table_id, foreign_key_id)

sysfkcol system table

Each row in the sysfkcol system table describes a foreign key column.

Column name	Column type	Description
table_id	UNSIGNED INTEGER	A unique identifier for the table to which the foreign key applies.
foreign_key_id	UNSIGNED INTEGER	A unique identifier for the foreign key that this column belongs to.
item_no	UNSIGNED SHORT	The order of the column in the foreign key.
column_id	UNSIGNED SHORT	A unique identifier of the table column that refers to the foreign column.
foreign_column_id	UNSIGNED SHORT	A unique identifier of the table column that is being referred to.

Constraints

PRIMARY KEY (table_id, foreign_key_id, item_no)

sysuldata system table

Each row in the sysuldata system table names value pairs of options and properties.

Synchronization profiles are stored in the sysuldata system table with **name** set to the profile's name, **type** set to **ulsync**, and **long_setting** set to the UTF-8 encoding of the profile string. Exactly one of the column settings and **long_setting** has an assigned value depending on usage.

Column name	Column type	Description
long_setting	LONGBINARY	A BLOB for long values.
name	VARCHAR(128)	The name of the property.
setting	VARCHAR(128)	The value of the property.
type	VARCHAR(128)	One of either opt for options, or prop for properties

Constraints

PRIMARY KEY (name, type)

See also

- [“Connection.OPTION_BLOB_FILE_BASE_DIR variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_DATABASE_ID variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_DATE_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_DATE_ORDER variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_ML_REMOTE_ID variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_NEAREST_CENTURY variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_PRECISION variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_SCALE variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_TIME_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_TIMESTAMP_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_TIMESTAMP_INCREMENT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.OPTION_TIMESTAMP_WITH_TIME_ZONE_FORMAT variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.PROPERTY_DATABASE_NAME variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)
- [“Connection.PROPERTY_PAGE_SIZE variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

UltraLite SQL reference

This section provides a reference for UltraLite SQL. UltraLite SQL is a unique subset of the SQL supported by SQL Anywhere databases.

UltraLite SQL elements

Keywords in UltraLite

Each SQL statement contains one or more keywords. SQL keywords are case insensitive, but throughout the documentation, keywords are indicated in uppercase. Some keywords cannot be used as identifiers without surrounding them in double quotes. These are called **reserved words**.

Note

UltraLite only supports a subset of SQL Anywhere keywords. However, to avoid potential problems in future releases, you should assume that all the reserved words for SQL Anywhere apply to UltraLite as well.

See also

- “Reserved words” [[SQL Anywhere Server - SQL Reference](#)]

Identifiers in UltraLite

Identifiers are names of objects in the database, such as user IDs, tables, and columns. Identifiers have a maximum length of 128 bytes.

You must enclose identifiers in double quotes if any of the following conditions are true:

- The identifier contains spaces.
- The first character of the identifier is not an alphabetic character. The database collation sequence dictates which characters are considered alphabetic or digit characters.
- The identifier contains a reserved word.
- The identifier contains characters other than alphabetic characters and digits.

You can only use a single backslash in an identifier if it is used as an escape character.

See also

- “Reserved words” [[SQL Anywhere Server - SQL Reference](#)]

Strings in UltraLite

Strings are used to hold character data in the database. UltraLite supports the same rules for strings as SQL Anywhere. The results of comparisons on strings, and the sort order of strings, depends on the case sensitivity of the database, the character set, and the collation sequence. These properties are set when the database is created.

See also

- “Strings” [[SQL Anywhere Server - SQL Reference](#)]
- “UltraLite character sets” on page 23

Comments in UltraLite

Comments are used to attach explanatory text to SQL statements or statement blocks. The UltraLite runtime does not execute comments.

The following comment indicators are available in UltraLite:

- **-- (Double hyphen)** The database server ignores any remaining characters on the line. This indicator is the SQL/2003 comment indicator.
- **// (Double slash)** The double slash has the same meaning as the double hyphen.
- **/* ... */ (Slash-asterisk)** Any characters between the two comment markers are ignored. The two comment markers may be on the same or different lines. Comments indicated in this style can be nested. This style of commenting is also called C-style comments.

Note

The percent sign (%) is not supported in UltraLite.

Examples

- The following example illustrates the use of double-hyphen comments:

```
CREATE TABLE borrowed_book (  
    loaner_name CHAR(100) PRIMARY KEY,  
    date_borrowed DATE NOT NULL DEFAULT CURRENT DATE,  
    date_returned DATE,  
    book CHAR(20)  
    FOREIGN KEY book REFERENCES library_books (isbn),  
);  
--This statement creates a table for a library database to hold  
information on borrowed books.  
--The default value for date_borrowed indicates that the book is borrowed  
on the day the entry is made.  
--The date_returned column is NULL until the book is returned.
```

- The following example illustrates the use of C-style comments:

```
CREATE TABLE borrowed_book (  
    loaner_name CHAR(100) PRIMARY KEY,
```

```
    date_borrowed DATE NOT NULL DEFAULT CURRENT DATE,  
    date_returned DATE,  
    book          CHAR(20)  
    FOREIGN KEY book REFERENCES library_books (isbn),  
);  
/* This statement creates a table for a library database to hold  
information on borrowed books.  
The default value for date_borrowed indicates that the book is borrowed on  
the day the entry is made.  
The date_returned column is NULL until the book is returned. */
```

Numbers in UltraLite

Numbers are used to hold numerical data in the database. A number can:

- be any sequence of digits
- be appended with decimal parts
- include an optional negative sign (-) or a plus sign (+)
- be followed by an e and then a numerical exponent value

For example, all numbers shown below are supported by UltraLite:

42

-4.038

.001

3.4e10

1e-10

The NULL value in UltraLite

As with SQL Anywhere, NULL is a special value that is different from any valid value for any data type. However, the NULL value is a legal value in any data type. NULL is used to represent unknown (no value) or inapplicable information.

See also

- “NULL special value” [[SQL Anywhere Server - SQL Reference](#)]

Special values in UltraLite

You can use special values in expressions, and as column defaults when you create tables.

CURRENT DATE special value

Returns the current year, month, and day.

Data type

DATE

Remarks

The returned date is based on a reading of the system clock when the SQL statement is executed by the UltraLite runtime. If you use CURRENT DATE with any of the following, all values are based on separate clock readings:

- CURRENT DATE multiple times within the same statement
- CURRENT DATE with CURRENT TIME or CURRENT TIMESTAMP within a single statement
- CURRENT DATE with the NOW function or GETDATE function within a single statement

See also

- [“Expressions in UltraLite” on page 280](#)
- [“GETDATE function \[Date and time\]” on page 339](#)
- [“NOW function \[Date and time\]” on page 364](#)

CURRENT TIME special value

The current hour, minute, second, and fraction of a second.

Data type

TIME

Remarks

The fraction of a second is stored to 6 decimal places. The accuracy of the current time is limited by the accuracy of the system clock.

The returned date is based on a reading of the system clock when the SQL statement is executed by the UltraLite runtime. If you use CURRENT TIME with any of the following, all values are based on separate clock readings:

- CURRENT TIME multiple times within the same statement
- CURRENT TIME with CURRENT DATE or CURRENT TIMESTAMP within a single statement
- CURRENT TIME with the NOW function or GETDATE function within a single statement

See also

- [“Expressions in UltraLite” on page 280](#)
- [“GETDATE function \[Date and time\]” on page 339](#)
- [“NOW function \[Date and time\]” on page 364](#)

CURRENT TIMESTAMP special value

Combines CURRENT DATE and CURRENT TIME to form a TIMESTAMP value containing the year, month, day, hour, minute, second, and fraction of a second.

Data type

TIMESTAMP

Remarks

The fraction of a second is stored to 3 decimal places. The accuracy is limited by the accuracy of the system clock.

Columns declared with DEFAULT CURRENT TIMESTAMP do not necessarily contain unique values.

The information CURRENT TIMESTAMP returns is equivalent to the information returned by the GETDATE and NOW functions.

CURRENT_TIMESTAMP is equivalent to CURRENT TIMESTAMP.

The returned date is based on a reading of the system clock when the SQL statement is executed by the UltraLite runtime. If you use CURRENT TIMESTAMP with any of the following, all values are based on separate clock readings:

- CURRENT TIMESTAMP multiple times within the same statement
- CURRENT TIMESTAMP with CURRENT DATE or CURRENT TIME within a single statement
- CURRENT TIMESTAMP with the NOW function or GETDATE function within a single statement

See also

- [“CURRENT TIME special value” on page 262](#)
- [“Expressions in UltraLite” on page 280](#)
- [“NOW function \[Date and time\]” on page 364](#)
- [“GETDATE function \[Date and time\]” on page 339](#)
- [“NOW function \[Date and time\]” on page 364](#)

CURRENT UTC TIMESTAMP special value

Returns a TIMESTAMP WITH TIME ZONE value that reflects the current UTC time containing the year, month, and day.

Data type

DATE

Remarks

The returned date is based on a reading of the system clock when the SQL statement is executed by the UltraLite runtime.

- CURRENT DATE multiple times within the same statement
- CURRENT DATE with CURRENT TIME or CURRENT TIMESTAMP within a single statement
- CURRENT DATE with the NOW function or GETDATE function within a single statement

See also

- [“Expressions in UltraLite” on page 280](#)
- [“GETDATE function \[Date and time\]” on page 339](#)
- [“CURRENT TIMESTAMP special value” on page 263](#)
- [“NOW function \[Date and time\]” on page 364](#)

SQLCODE special value

Current SQLCODE value at the time the special value was evaluated. Not supported by UltraLite Java edition databases.

Data type

String

Remarks

The SQLCODE value is set after each statement. You can check the SQLCODE to determine if the statement succeeded.

See also

- [“Expressions in UltraLite” on page 280](#)
- [“Error Messages”](#)

Example

Use a SELECT statement to produce an error code for each attempt to fetch a new row from the result set. For example: `SELECT a, b, SQLCODE FROM MyTable.`

Dates and times in UltraLite

Many of the date and time functions use dates built from date and time parts. UltraLite and SQL Anywhere support the same date parts.

See also

- [“Specifying date parts” on page 301](#)

Data types in UltraLite

Available data types in UltraLite SQL include:

- Integer
- Decimal
- Floating-point
- Character
- Binary
- Date/time

Note

Domains (user-defined data types) are not supported in UltraLite SQL.

Note

You cannot concatenate LONGVARCHAR and LONGBINARY data types. See [“String operators” on page 297](#).

You can create a host variable with any one of the supported types. UltraLite supports a subset of the data types available in SQL Anywhere. The following are the SQL data types supported in UltraLite databases.

Data type	Description
BIT	Boolean values (0 or 1). See “BIT data type” [SQL Anywhere Server - SQL Reference] .
{ CHAR CHARACTER } (<i>max-length</i>)	Character data of <i>max-length</i> , in the range of 1-32767 bytes. See “CHAR data type” [SQL Anywhere Server - SQL Reference] . When evaluating expressions, the maximum length for a temporary character value is 2048 bytes.
VARCHAR (<i>max-length</i>)	VARCHAR is used for variable-length character data of <i>max-length</i> . See “VARCHAR data type” [SQL Anywhere Server - SQL Reference] .

Data type	Description
LONG VARCHAR	Arbitrary length character data. Conditions in SQL statements (such as in the WHERE clause) cannot operate on LONG VARCHAR columns. The only operations allowed on LONG VARCHAR columns are to insert, update, or delete them, or to include them in the SELECT list of a query. See “LONG VARCHAR data type” [SQL Anywhere Server - SQL Reference] . You can cast strings to/from LONGVARCHAR data.
[UNSIGNED] BIGINT	An integer requiring 8 bytes of storage. See “BIGINT data type” [SQL Anywhere Server - SQL Reference] .
{ DECIMAL DEC NUMERIC } (<i>precision</i> , <i>scale</i>)]	The representation of a decimal number using two parts: <i>precision</i> (total digits) and <i>scale</i> (digits that follow a decimal point). See “DECIMAL data type” [SQL Anywhere Server - SQL Reference] , “NUMERIC data type” [SQL Anywhere Server - SQL Reference] , “UltraLite precision creation parameter” on page 155 , and “UltraLite scale creation parameter” on page 156 .
DOUBLE [PRECISION]	A double-precision floating-point number. In this data type PRECISION is an optional part of the DOUBLE data type name. See “DOUBLE data type” [SQL Anywhere Server - SQL Reference] .
FLOAT [(<i>precision</i>)]	A floating-point number, which may be single or double precision. See “FLOAT data type” [SQL Anywhere Server - SQL Reference] .
[UNSIGNED] { INT INTEGER }	An unsigned integer requiring 4 bytes of storage. See “INTEGER data type” [SQL Anywhere Server - SQL Reference] .
REAL	A single-precision floating-point number stored in 4 bytes. See “REAL data type” [SQL Anywhere Server - SQL Reference] .
[UNSIGNED] SMALLINT	An integer requiring 2 bytes of storage. See “SMALLINT data type” [SQL Anywhere Server - SQL Reference] .
[UNSIGNED] TINYINT	An integer requiring 1 byte of storage. See “TINYINT data type” [SQL Anywhere Server - SQL Reference] .
DATE	A calendar date, such as a year, month, and day. See “DATE data type” [SQL Anywhere Server - SQL Reference] .
TIME	The time of day, containing hour, minute, second, and fraction of a second. See “TIME data type” [SQL Anywhere Server - SQL Reference] .

Data type	Description
DATETIME	Identical to TIMESTAMP . See “ DATETIME data type ” [<i>SQL Anywhere Server - SQL Reference</i>].
TIMESTAMP	A point in time, containing year, month, day, hour, minute, second, and fraction of a second. See “ TIMESTAMP data type ” [<i>SQL Anywhere Server - SQL Reference</i>].
VARBINARY (<i>max-length</i>)	Identical to BINARY . See “ VARBINARY data type ” [<i>SQL Anywhere Server - SQL Reference</i>].
BINARY (<i>max-length</i>)	Binary data of maximum length <i>max-length</i> bytes. The maximum length should not exceed 2048 bytes. See “ BINARY data type ” [<i>SQL Anywhere Server - SQL Reference</i>].
LONG BINARY	Arbitrary length binary data. Conditions in SQL statements (such as in the WHERE clause) cannot operate on LONG BINARY columns. The only operations allowed on LONG BINARY columns are to insert, update, or delete them, or to include them in the SELECT list of a query. See “ LONG BINARY data type ” [<i>SQL Anywhere Server - SQL Reference</i>]. You can cast values to/from LONGBINARY data.
UNIQUEIDENTIFIER	Typically used for a primary key or other unique column to hold UUID (Universally Unique Identifier) values that uniquely identify rows. UltraLite provides functions that generate UUID values. These values are generated so that a value produced on one computer does not match a UUID produced on another computer. UNIQUEIDENTIFIER values generated in this way can therefore be used as keys in a synchronization environment. See “ UNIQUEIDENTIFIER data type ” [<i>SQL Anywhere Server - SQL Reference</i>].

User-defined data types and their equivalents

Unlike SQL Anywhere databases, UltraLite does not support user-defined data types. The following table lists UltraLite data type equivalents to built-in SQL Anywhere aliases:

SQL Anywhere data type	UltraLite equivalent
MONEY	NUMERIC(19,4)
SMALLMONEY	NUMERIC(10,4)
TEXT	LONG VARCHAR

SQL Anywhere data type	UltraLite equivalent
XML	LONG VARCHAR

Converting data types explicitly

UltraLite allows you to request data type conversions explicitly, by using either the CAST or CONVERT function.

Note

Self-casting usually has no effect on operations. However, self-casts to CHAR/VARCHAR, BINARY/VARBINARY and NUMERIC are not no-op procedures.

You can CAST or CONVERT most combinations of data types, as illustrated by the table that follows.

The ability to convert or not is contingent upon the value used in the conversion. As the **Value-dependent** column shows, the value must be compatible with the new data type to avoid generating a specific type of conversion error. For example:

- If you cast **varchar "1234"** to **long**, this conversion is supported. However, if you cast **varchar "hello"** to **long**, then this conversion generates a **SQLLE_CONVERSION_ERROR** error because **hello** is not a number.
- If you cast **long 1234** to **short**, this conversion is supported. However, if you cast **long 1000000** to **short**, then this conversion generates a **SQLLE_OVERFLOW_ERROR** error, because **1000000** is beyond the range of numbers a short can hold.

From	Always	Never	Value-dependent
BINARY or VARBI- NARY	CHAR or VARCHAR BINARY LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT BIGINT SIGNED BIG	LONG VARCHAR REAL TIME TIMESTAMP DOUBLE DATE	NUMERIC UID ¹

From	Always	Never	Value-dependent
LONG BINARY	BINARY LONG BINARY	BIT CHAR or VARCHAR LONG VARCHAR TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL DOUBLE NUMERIC DATE TIME TIMESTAMP UID	N/A

From	Always	Never	Value-dependent
BIT	CHAR or VARCHAR BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT REAL SIGNED BIG DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	N/A

From	Always	Never	Value-dependent
CHAR or VAR- CHAR	BINARY or VARBINARY CHAR or VARCHAR LONG VARCHAR	LONG BINARY	BIT TINYINT SIGNED SHORT SHORT INT LONG INT SIGNED LONG BIGINT SIGNED BIG DOUBLE NUMERIC REAL DATE TIME TIMESTAMP UID

From	Always	Never	Value-dependent
LONG VAR- CHAR	CHAR or VARCHAR LONG VARCHAR	BINARY or VARBINARY LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL NUMERIC DATE TIME TIMESTAMP DOUBLE UID	

From	Always	Never	Value-dependent
TINYINT	BINARY or VARBINARY CHAR or VARCHAR TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	
SHORT INT	BINARY or VARBINARY CHAR or VARCHAR SHORT INT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	BIT TINYINT SIGNED SHORT

From	Always	Never	Value-dependent
SIGNED SHORT	BINARY or VARBINARY CHAR or VARCHAR SIGNED SHORT SIGNED LONG SIGNED BIG REAL DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	SHORT INT LONG INT BIGINT BIT TINYINT
LONG INT	BINARY or VARBINARY CHAR or VARCHAR LONG INT BIGINT SIGNED BIG REAL DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	BIT TINYINT SHORT INT SIGNED SHORT SIGNED LONG
SIGNED LONG	BINARY or VARBINARY CHAR or VARCHAR SIGNED LONG SIGNED BIG REAL DOUBLE NUMERIC DATE TIMESTAMP	LONG VARCHAR LONG BINARY TIME UID	BIT TINYINT SHORT INT SIGNED SHORT LONG INT BIGINT

From	Always	Never	Value-dependent
BIGINT	BINARY or VARBINARY CHAR or VARCHAR BIGINT REAL DOUBLE NUMERIC	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG SIGNED BIG
SIGNED BIG	BINARY or VARBINARY CHAR or VARCHAR SIGNED BIG REAL DOUBLE NUMERIC DATE TIMESTAMP	LONG VARCHAR LONG BINARY TIME UID	BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT
REAL	CHAR or VARCHAR REAL DOUBLE NUMERIC	LONG VARCHAR BINARY or VARBINARY LONG BINARY DATE TIME TIMESTAMP UID	BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG

From	Always	Never	Value-dependent
DOUBLE	CHAR or VARCHAR DOUBLE NUMERIC	LONG VARCHAR BINARY or VARBINARY LONG BINARY DATE TIME TIMESTAMP UID	BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL
NUMERIC	CHAR or VARCHAR REAL NUMERIC DOUBLE	LONG VARCHAR LONG BINARY DATE TIME TIMESTAMP UID	BINARY or VARBINARY ² BIT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG TINYINT

From	Always	Never	Value-dependent
DATE	CHAR or VARCHAR SIGNED LONG SIGNED BIG DATE TIMESTAMP	LONG VARCHAR LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT BIGINT REAL DOUBLE NUMERIC TIME BINARY or VARBINARY UID	

From	Always	Never	Value-dependent
TIME	CHAR or VARCHAR TIME TIMESTAMP	LONG VARCHAR LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL DOUBLE NUMERIC DATE BINARY or VARBINARY UID	

From	Always	Never	Value-dependent
TIME-STAMP	CHAR or VARCHAR SIGNED LONG SIGNED BIG DATE TIME TIMESTAMP	LONG VARCHAR LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT BIGINT REAL DOUBLE NUMERIC BINARY or VARBINARY UID	

From	Always	Never	Value-dependent
UID	CHAR or VARCHAR UID	LONG VARCHAR LONG BINARY BIT TINYINT SHORT INT SIGNED SHORT LONG INT SIGNED LONG BIGINT SIGNED BIG REAL DOUBLE NUMERIC DATE TIME TIMESTAMP	BINARY or VARBINARY ¹

¹ The BINARY value must be a 16 byte length to be compatible with a UUID.

² Only works if the source NUMERIC value is able to cast as a BIGINT.

Expressions in UltraLite

Expressions are formed by combining data, often in the form of column references, with operators or functions.

Syntax

```
expression:
  case-expression
  | constant
  | [correlation-name.]column-name
  | - expression
```


| *expression operator expression*
 | (*expression*)
 | *function-name* (*expression*, ...)
 | *if-expression*
 | *special value*
 | *input-parameter*

Parameters

case-expression:

CASE *expression*
WHEN *expression*
THEN *expression*,...
 [**ELSE** *expression*]
END

alternative form of case-expression:

CASE
WHEN *search-condition*
THEN *expression*,...
 [**ELSE** *expression*]
END

constant:

integer | *number* | *string* | *host-variable*

special-value:

CURRENT { **DATE** | **TIME** | **TIMESTAMP** }
 | **NULL**
 | **SQLCODE**
 | **SQLSTATE**

if-expression:

IF *condition*
THEN *expression*
 [**ELSE** *expression*]
ENDIF

input-parameter:

{ ? | :*name* [: *indicator-name*] }

operator:

{ + | - | * | / | || | % }

See also

- [“Constants in expressions” on page 282](#)
- [“Special values in UltraLite” on page 261](#)
- [“Column names in expressions” on page 282](#)
- [“UltraLite SQL functions” on page 299](#)
- [“Subqueries in expressions” on page 285](#)
- [“Search conditions in UltraLite” on page 287](#)
- [“Data types in UltraLite” on page 265](#)
- [“CASE expressions” on page 283](#)
- [“Input parameters” on page 286](#)

Constants in expressions

In UltraLite, constants are numbers or string literals.

Syntax

`' constant '`

Usage

String constants are enclosed in single quotes (').

An apostrophe is represented inside a string by two single quotes in a row (").

See also

- [“Escape sequences” \[SQL Anywhere Server - SQL Reference\]](#)

Example

To use a possessive phrase, type the string literal as follows:

```
'John''s database'
```

Column names in expressions

An identifier in an expression.

Syntax

`correlation-name.column-name`

Remarks

A column name is preceded by an optional correlation name, which typically is the name of a table.

If a column name is a keyword or has characters other than letters, digits and underscore, it must be surrounded by quotation marks (" "). For example, the following are valid column names:

```
Employees.Name  
address  
"date hired"  
"salary"."date paid"
```

See also

- [“FROM clause \[UltraLite\]” on page 436](#)

IF expressions

Sets a search condition to return a specific subset of data.

Syntax 1

```

IF search-condition
THEN expression1
[ ELSE expression2 ]
ENDIF

```

Remarks

For compatibility reasons, this expression can end in either ENDIF or END IF.

This expression returns the following:

- If *search-condition* is TRUE, the IF expression returns *expression1*.
- If *search-condition* is FALSE and an ELSE clause is specified, the IF expression returns *expression2*.
- If *search-condition* is FALSE, and there is no *expression2*, the IF expression returns NULL.
- If *search-condition* is UNKNOWN, the IF expression returns NULL.

See also

- “NULL special value” [[SQL Anywhere Server - SQL Reference](#)]
- “Search conditions” [[SQL Anywhere Server - SQL Reference](#)]

CASE expressions

Provides conditional SQL expressions.

Syntax 1

```

CASE expression1
WHEN expression2 THEN expression3, ...
[ ELSE expression4 ]
END

```

```

SELECT id,
       ( CASE name
         WHEN 'Tee Shirt' THEN 'Shirt'
         WHEN 'Sweatshirt' THEN 'Shirt'
         WHEN 'Baseball Cap' THEN 'Hat'
         ELSE 'Unknown'
       END ) as Type
FROM Product;

```

Syntax 2

```

CASE
WHEN search-condition
THEN expression1, ...
[ ELSE expression2 ]
END

```

Remarks

For compatibility reasons, you can end this expression with either ENDCASE or END CASE.

You can use case expressions anywhere you can use regular expression.

Syntax 1 If the expression following the CASE keyword is equal to the expression following the first WHEN keyword, then the expression following the associated THEN keyword is returned. Otherwise the expression following the ELSE keyword is returned, if specified.

For example, the following code uses a case expression as the second clause in a SELECT statement. It selects a row from the Product table where the name column has a value of Sweatshirt.

Syntax 2 If the search-condition following the first WHEN keyword is TRUE, the expression following the associate THEN keyword is returned. Otherwise the expression following the ELSE clause is returned, if specified.

NULLIF function for abbreviated CASE expressions The NULLIF function provides a way to write some CASE statements in short form. The syntax for NULLIF is as follows:

NULLIF (*expression-1*, *expression-2*)

NULLIF compares the values of the two expressions. If the first expression equals the second expression, NULLIF returns NULL. If the first expression does not equal the second expression, NULLIF returns the first expression.

Example

The following statement uses a CASE expression as the third clause of a SELECT statement to associate a string with a search condition. If the name column's value is **Tee Shirt**, this query returns **Sale**. And if the name column's value is not **Tee Shirt** and the quantity is greater than fifty, it returns **Big Sale**. However, for all others, the query then returns **Regular price**.

```
SELECT id, name,
( CASE
  WHEN name='Tee Shirt' THEN 'Sale'
  WHEN quantity >= 50 THEN 'Big Sale'
  ELSE 'Regular price'
END ) as Type
FROM Product;
```

Aggregate expressions

Performs an aggregate computation that the UltraLite runtime does not provide.

Syntax

SUM(*expression*)

Remarks

An aggregate expression calculates a single value from a range of rows.

An aggregate expression is one in which either an aggregate function is used, or in which one or more of the operands is an aggregate expression.

When a `SELECT` statement does not have a `GROUP BY` clause, the expressions in the `SELECT` list must either contain all aggregate expressions or no aggregate expressions. When a `SELECT` statement does have a `GROUP BY` clause, any non-aggregate expression in the `SELECT` list must appear in the `GROUP BY` list.

Example

For example, the following query computes the total payroll for employees in the employee table. In this query, `SUM(salary)` is an aggregate expression:

```
SELECT SUM( salary )
FROM employee;
```

Subqueries in expressions

A `SELECT` statement that is nested inside another `SELECT` statement.

Syntax

A subquery is structured like a regular query.

Remarks

In UltraLite, you can only use subquery references in the following situations:

- As a table expression in the `FROM` clause. This form of table expression (also called **derived tables**) must have a derived table name and column names in which values in the `SELECT` list are fetched.
- To supply values for the `EXISTS`, `ANY`, `ALL`, and `IN` search conditions.

You can write subqueries about names that are specified before (to the left of) the subquery, sometimes known as outer references to the left. However, you cannot have references to items within subqueries (sometimes known as inner references).

See also

- [“SELECT statement \[UltraLite\]” on page 444](#)
- [“Use of subqueries” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Search conditions in UltraLite” on page 287](#)

Example

The following subquery is used to list all product IDs for items that are low in stock (that is, less than 20 items).

```
FROM SalesOrderItems
( SELECT ID
  FROM Products
  WHERE Quantity < 20 );
```

Input parameters

Acts as placeholders to allow end-users to supply values to a prepared statement. These user-supplied values are then used to execute the statement.

Syntax

```
{ ? | :name [ : indicator-name ] }
```

Remarks

Use the placeholder character of ? or the named form in expressions. You can use input parameters whenever you can use a column name or constant.

The precise mechanism used to supply the values to the statement are dependent upon the API you use to create your UltraLite client.

Using the named form The named form of an input parameter has special meaning. In general, *name* is always used to specify multiple locations where an actual value is supplied.

For embedded SQL applications only, the *indicator-name* supplies the variable into which the null indicator is placed. If you use the named form with the other components, *indicator-name* is ignored.

Deducing data types The data type of the input parameter is deduced when the statement is prepared from one of the following patterns:

- **CAST (? AS *type*)**

In this case, *type* is a database type specification such as CHAR(32).

- Exactly one operand of a binary operator is an input parameter. The type is deduced to be the type of the operand.

If the type cannot be deduced, UltraLite generates an error. For example:

- **-?:** the operand is unary.
- **? + ?:** both are input parameters.

See also

- [“Host variables” \[UltraLite - C and C++ Programming\]](#)
- [“Prepared statements” \[SQL Anywhere Server - Programming\]](#)
- UltraLite C/C++: [“Modifying data using INSERT, UPDATE and DELETE” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“Modifying data using INSERT, UPDATE, and DELETE” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business: [“Modifying data using INSERT, UPDATE, and DELETE” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)

Example

The following embedded SQL statement has two input parameters:

```
INSERT INTO MyTable VALUES ( :v1, :v2, :v1);
```

The first instance of v1 supplies its value to both the v2 and v1 locations in the statement.

Search conditions in UltraLite

A search condition is the criteria for a WHERE clause, a HAVING clause, an ON phrase in a join, or an IF expression. A search condition is also called a **predicate**.

Syntax

```
search-condition:
expression comparison-operator expression
| expression IS [ NOT ] NULL
| expression [ NOT ] BETWEEN expression AND expression
| expression [ NOT ] IN ( expression, ... )
| expression [ NOT ] IN ( subquery )
| expression [ NOT ] { ANY | ALL } ( subquery )
| expression [ NOT ] EXISTS ( subquery )
| expression [ NOT ] LIKE ( pattern )
NOT search-condition
search-condition AND search-condition
search-condition OR search-condition
( search-condition IS [ NOT ] {TRUE | FALSE | UNKNOWN }

```

comparison-operator :

```
=
<
>
^
<=
>=
<>
!=
<^
>^
v
```

Parameters

The different types of search conditions supported by UltraLite include:

- ALL condition
- ANY condition
- BETWEEN condition
- EXISTS condition
- IN condition
- LIKE condition

Remarks

In UltraLite, search conditions can appear in the:

- WHERE clause
- HAVING clause

- ON phrase
- SQL queries

Search conditions can be used to choose a subset of the rows from a table in a FROM clause in a SELECT statement, or in expressions such as an IF or CASE to select specific values. In UltraLite, every condition evaluates as one of three states: TRUE, FALSE, or UNKNOWN. When combined, these states are referred to as **three-valued logic**. The result of a comparison is UNKNOWN if either value being compared is the NULL value. Search conditions are satisfied only if the result of the condition is TRUE.

See also

- [“Search conditions in UltraLite” on page 287](#)
- [“Queries” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Pattern matching character strings in the WHERE clause” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Comparison operators” on page 288](#)
- [“Three-valued logic” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Subqueries in expressions” on page 285](#)

Comparison operators

Any operator that allows two or more expressions to be compared with in a search condition.

Syntax

expression operator expression

Parameters

Operator	Interpretation
=	equal to
>	greater than
<	less than
>=	greater than or equal to
<=	less than or equal to
!=	not equal to
<>	not equal to
!>	not greater than
!<	not less than

Remarks

Comparing dates In comparing dates, < means earlier and > means later.

Comparing LONG VARCHAR or LONG BINARY values UltraLite does not support comparisons using LONG VARCHAR or LONG BINARY values.

Case-sensitivity In UltraLite, comparisons are carried out with the same attention to case as the database on which they are operating. By default, UltraLite databases are created as case insensitive.

NOT operator The NOT operator negates an expression.

See also

- [“Logical operators” on page 289](#)
- [“Search conditions in UltraLite” on page 287](#)

Example

Either of the following two queries will find all Tee shirts and baseball caps that cost \$10 or less. However, note the difference in position between the negative logical operator (NOT) and the negative comparison operator (!>).

```
SELECT ID, Name, Quantity
FROM Products
WHERE (name = 'Tee Shirt' OR name = 'BaseBall Cap')
AND NOT UnitPrice > 10;
```

```
SELECT ID, Name, Quantity
FROM Products
WHERE (name = 'Tee Shirt' OR name = 'BaseBall Cap')
AND UnitPrice !> 10;
```

Logical operators

Does any of the following:

- Compare conditions (AND, OR, and NOT).
- Test the truth or NULL value nature of the expressions (IS).

Syntax 1

condition1 logical-operator condition2

Syntax 2

NOT *condition*

Syntax 3

expression IS [NOT] { truth-value | NULL }

Remarks

Search conditions can be used to choose a subset of the rows from a table in a FROM clause in a SELECT statement, or in expressions such as an IF or CASE to select specific values. In UltraLite, every condition evaluates as one of three states: TRUE, FALSE, or UNKNOWN. When combined, these states are referred to as **three-valued logic**. The result of a comparison is UNKNOWN if either value being compared is the NULL value. Search conditions are satisfied only if the result of the condition is TRUE.

AND The combined condition is TRUE if both conditions are TRUE, FALSE if either condition is FALSE, and UNKNOWN otherwise.

condition1 **AND** *condition2*

OR The combined condition is TRUE if either condition is TRUE, FALSE if both conditions are FALSE, and UNKNOWN otherwise.

NOT The NOT condition is TRUE if *condition* is FALSE, FALSE if *condition* is TRUE, and UNKNOWN if *condition* is UNKNOWN.

IS The condition is TRUE if the *expression* evaluates to the supplied *truth-value*, which must be one of TRUE, FALSE, or UNKNOWN. Otherwise, the value is FALSE.

See also

- [“Three-valued logic” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Comparison operators” on page 288](#)
- [“Search conditions in UltraLite” on page 287](#)

Example

The IS NULL condition is satisfied if the column contains a NULL value. If you use the IS NOT NULL operator, the condition is satisfied when the column contains a value that is not NULL. This example shows an IS NULL condition: WHERE paid_date IS NULL.

ALL search condition

Use the ALL condition in conjunction with comparison operators to compare a single value to the data values produced by the subquery.

Syntax

expression compare [**NOT**] **ALL** (*subquery*)

Parameters

compare:
= | > | < | >= | <= | <> | != | !< | !>

Remarks

UltraLite uses the specified comparison operator to compare the test value to each data value in the result set. If all the comparisons yield TRUE results, the ALL test returns TRUE.

See also

- “Subqueries and the ALL test” [[SQL Anywhere Server - SQL Usage](#)]
- “Comparison operators” on page 288

Example

Find the order and customer IDs of those orders placed after all products of order #2001 were shipped.

```
SELECT ID, CustomerID
FROM SalesOrders
WHERE OrderDate > ALL (
  SELECT ShipDate
  FROM SalesOrderItems
  WHERE ID=2001);
```

ANY search condition

Use the ANY condition in conjunction with comparison operators to compare a single value to the column of data values produced by the subquery.

Syntax 1

expression compare [NOT] ANY (*subquery*)

Syntax 2

expression = ANY (*subquery*)

Parameters

compare:

= | > | < | >= | <= | <> | != | !< | !>

Remarks

UltraLite uses the specified comparison operator to compare the test value to each data value in the column. If any of the comparisons yields a TRUE result, the ANY test returns TRUE.

Syntax 1 is TRUE if *expression* is equal to any of the values in the result of the subquery, and FALSE if the expression is not NULL and does not equal any of the values returned by the subquery. The ANY condition is UNKNOWN if *expression* is the NULL value, unless the result of the subquery has no rows, in which case the condition is always FALSE.

See also

- “Subqueries and the ANY test” [[SQL Anywhere Server - SQL Usage](#)]
- “Comparison operators” on page 288

Example

Find the order and customer IDs of those orders placed after the first product of the order #2005 was shipped.

```
SELECT ID, CustomerID
FROM SalesOrders
```

```
WHERE OrderDate > ANY (
  SELECT ShipDate
  FROM SalesOrderItems
  WHERE ID=2005);
```

BETWEEN search condition

Specifies an inclusive range, in which the lower value and the upper value and the values they delimit are searched for.

Syntax

```
expression [ NOT ] BETWEEN start-expression AND end-expression
```

Remarks

The BETWEEN condition can evaluate to TRUE, FALSE, or UNKNOWN. Without the NOT keyword, the condition evaluates as TRUE if *expression* is between *start-expression* and *end-expression*. The NOT keyword reverses the meaning of the condition, but leaves UNKNOWN unchanged.

The BETWEEN condition is equivalent to a combination of two inequalities:

```
[ NOT ] ( expression >= start-expression
          AND expression <= end-expression )
```

Example

List all the products less expensive than \$10 or more expensive than \$15.

```
SELECT Name, UnitPrice
FROM Products
WHERE UnitPrice NOT BETWEEN 10 AND 15;
```

EXISTS search condition

Checks whether a subquery produces any rows of query results

Syntax

```
[ NOT ] EXISTS ( subquery )
```

Remarks

The EXISTS condition is TRUE if the subquery result contains at least one row, and FALSE if the subquery result does not contain any rows. The EXISTS condition cannot be UNKNOWN.

You can reverse the logic of the EXISTS condition by using the NOT EXISTS form. In this case, the test returns TRUE if the subquery produces no rows, and FALSE otherwise.

Example

List the customers who placed orders after July 13, 2001.

```

SELECT GivenName, Surname
FROM Customers
WHERE EXISTS (
  SELECT *
  FROM SalesOrders
  WHERE (OrderDate > '2001-07-13') AND
        (Customers.ID = SalesOrders.CustomerID));

```

IN search condition

Checks membership by searching a value from the main query with another value in the subquery.

Syntax

```

expression [ NOT ] IN
{ ( subquery ) | ( value-expr, ... ) }

```

Parameters

value-expr are expressions that take on a single value, which may be a string, a number, a date, or any other SQL data type.

Remarks

An IN condition, without the NOT keyword, evaluates according to the following rules:

- TRUE if *expression* is not NULL and equals at least one of the values.
- UNKNOWN if *expression* is NULL and the values list is not empty, or if at least one of the values is NULL and *expression* does not equal any of the other values.
- FALSE if *expression* is NULL and *subquery* returns no values; or if *expression* is not NULL, none of the values are NULL, and *expression* does not equal any of the values.

You can reverse the logic of the IN condition by using the NOT IN form.

The following search condition *expression* **IN** (*values*) is identical to the search condition *expression* = **ANY** (*values*). The search condition *expression* **NOT IN** (*values*) is identical to the search condition *expression* <> **ALL** (*values*).

Example

Select the company name and state for customers who live in the following Canadian provinces: Ontario, Manitoba, and Quebec.

```

SELECT CompanyName , Province
FROM Customers
WHERE State IN( 'ON', 'MB', 'PQ' );

```

LIKE search condition

Syntax

The syntax for the LIKE search condition is as follows:

```
expression [ NOT ] LIKE pattern
```

Parameters

- **expression** The string to be searched.
- **pattern** The pattern to search for within *expression*.

Remarks

The LIKE search condition attempts to match *expression* with *pattern* and evaluates to TRUE, FALSE, or UNKNOWN. The search condition evaluates to TRUE if *expression* matches *pattern* (assuming NOT was not specified). If either *expression* or *pattern* is the NULL value, the search condition evaluates to UNKNOWN.

The NOT keyword reverses the meaning of the search condition, but leaves UNKNOWN unchanged.

expression and *pattern* are interpreted as CHAR strings. *pattern* can contain any number of the supported wildcards from the following table:

Wildcard	Matches
_ (underscore)	Any one character. For example, a_ matches ab and ac, but not a.
% (percent)	Any string of zero or more characters. For example, bl% matches bl and bla.
[]	Any single character in the specified range or set. For example, T[oi]m matches Tom or Tim.
[^]	Any single character <i>not</i> in the specified range or set. For example, M[^c] matches Mb and Md, but not Mc.

Example

The following search condition returns TRUE for any row where column-name starts with the letter a and has the letter b as its second last character:

```
SELECT * FROM table-name WHERE column-name LIKE 'a%b_'
```

Different ways to use the LIKE search condition

To search for	Example	Additional information
One of a set of characters	LIKE 'sm[iy]th'	A set of characters to look for is specified by listing the characters inside square brackets. In this example, the search condition matches <i>smith</i> and <i>smyth</i> .
One of a range of characters	LIKE '[a-r]ough'	<p>A range of characters to look for is specified by giving the ends of the range inside square brackets, separated by a hyphen. In this example, the search condition matches <i>bough</i> and <i>rough</i>, but not <i>tough</i>.</p> <p>The range of characters [a-z] is interpreted as "greater than or equal to a, and less than or equal to z", where the greater than and less than operations are carried out within the collation of the database.</p> <p>The lower end of the range must precede the higher end of the range. For example, [z-a] does not match anything because no character matches the [z-a] range.</p>
Ranges and sets combined	... LIKE '[a-rt]ough'	<p>You can combine ranges and sets within square brackets. In this example, ... LIKE '[a-rt]ough' matches <i>bough</i>, <i>rough</i>, and <i>tough</i>.</p> <p>The pattern [a-rt] is interpreted as exactly one character that is either in the range a to r inclusive, or is t.</p>
One character not in a range	... LIKE '[^a-r]ough'	<p>The caret character (^) is used to specify a range of characters that is excluded from a search. In this example, LIKE '[^a-r]ough' matches the string <i>tough</i>, but not the strings <i>rough</i> or <i>bough</i>.</p> <p>The caret negates the rest of the contents of the brackets. For example, the bracket [^a-rt] is interpreted as exactly one character that is not in the range a to r inclusive, and is not t.</p>
Search patterns with trailing blanks	'90 ', '90[]' and '90_'	When your search pattern includes trailing blanks, the database server matches the pattern only to values that contain blanks—it does not blank pad strings. For example, the patterns '90 ', '90[]', and '90_' match the expression '90 ', but do not match the expression '90', even if the value being tested is in a CHAR or VARCHAR column that is three or more characters in width.

Special cases of ranges and sets

Any single character in square brackets means that character. For example, [a] matches just the character a. [^] matches just the caret character, [%] matches just the percent character (the percent

character does not act as a wildcard in this context), and [_] matches just the underscore character. Also, [[] matches just the character [.

- The pattern [a-] matches either of the characters a or -.
- The pattern [] is never matched and always returns no rows.
- The patterns [or [a**p**-q return syntax errors because they are missing the closing bracket.
- You cannot use wildcards inside square brackets. The pattern [a%b] finds one of a, %, or b.
- You cannot use the caret character to negate ranges except as the first character in the bracket. The pattern [a^b] finds one of a, ^, or b.

Case sensitivity and how comparisons are performed

If the database collation is case sensitive, the search condition is also case sensitive. To perform a case insensitive search with a case sensitive collation, you must include upper and lower characters. For example, the following search condition evaluates to true for the strings Bough, rough, and TOUGH:

```
LIKE '[a-zA-Z][oO][uU][gG][hH]'
```

Operators in UltraLite

Operators are used to compute values, which may in turn be used as operands in a higher-level expression.

UltraLite SQL supports the following types of operators:

- Comparison operators evaluate and return a result using one (unary) or two (binary) comparison operands. Comparisons result in the usual three logical values: true, false, and unknown.
- Arithmetic operators evaluate and return a result set for all floating-point, decimal, and integer numbers.
- String operators concatenate two string values together. For example, "my" + "string" returns the string "my string".
- Bitwise operators evaluate and turn specific bits on or off within the internal representation of an integer.
- Logical operators evaluate search conditions. Logical evaluations result in the usual three logical values: true, false, and unknown.

The normal precedence of operations applies.

See also

- “Operator precedence” on page 298
- “Comparison operators” on page 288
- “Arithmetic operators” on page 297
- “String operators” on page 297
- “Bitwise operators” on page 298
- “Logical operators” on page 289

Arithmetic operators

Arithmetic operators allow you to perform calculations.

expression + expression Addition. If either expression is NULL, the result is NULL.

expression - expression Subtraction. If either expression is NULL, the result is NULL.

- expression Negation. If the expression is NULL, the result is NULL.

expression * expression Multiplication. If either expression is NULL, the result is NULL.

expression / expression Division. If either expression is NULL or if the second expression is 0, the result is NULL.

expression % expression Modulo finds the integer remainder after a division involving two whole numbers. For example, $21 \% 11 = 10$ because 21 divided by 11 equals 1 with a remainder of 10. If either expression is NULL, the result is NULL.

See also

- “Arithmetic operations” [[SQL Anywhere Server - SQL Usage](#)]

String operators

String operators allow you to concatenate strings—except for LONGVARCHAR and LONGBINARY data types.

expression || expression String concatenation (two vertical bars). If either string is NULL, it is treated as the empty string for concatenation.

expression + expression Alternative string concatenation. When using the + concatenation operator, you must ensure the operands are explicitly set to character data types rather than relying on implicit data conversion.

For example, the following query returns the integer value **579**:

```
SELECT 123 + 456;
```

However, the following query returns the character string **123456**:

```
SELECT '123' + '456';
```

You can use the CAST or CONVERT functions to explicitly convert data types.

Bitwise operators

Bitwise operators perform bit manipulations between two expressions. The following operators can be used on integer data types in UltraLite.

Operator	Description
&	bitwise AND
	bitwise OR
^	bitwise exclusive OR
~	bitwise NOT

The bitwise operators &, |, and ~ are not interchangeable with the logical operators AND, OR, and NOT. The bitwise operators operate on integer values using the bit representation of the values.

Example

The following statement selects rows in which the specified bits are set.

```
SELECT *  
FROM tableA  
WHERE (options & 0x0101) <> 0;
```

Operator precedence

The precedence of operators in expressions is as follows. Expressions in parentheses are evaluated first, then multiplication and division before addition and subtraction. String concatenation happens after addition and subtraction. The operators at the top of the list are evaluated before those at the bottom of the list.

Tip

Make the order of operation explicit in UltraLite, rather than relying on an operator precedence. That means, when you use more than one operator in an expression you should order operations clearly with parentheses.

1. names, functions, constants, IF expressions, CASE expressions
2. ()
3. unary operators (operators that require a single operand): +, -

4. ~
5. &, |, ^
6. *, /, %
7. +, -
8. ||
9. comparisons: >, <, <>, !=, <=, >=, [NOT] BETWEEN, [NOT] IN, [NOT] LIKE
10. comparisons: IS [NOT] TRUE, FALSE, UNKNOWN
11. NOT
12. AND
13. OR

Variables in UltraLite

You cannot use SQL variables (including global variables) in UltraLite applications.

UltraLite SQL functions

Functions are used to return information from the database. They are allowed anywhere an expression is allowed.

Unless otherwise specified in the documentation, NULL is returned for a function if any argument is NULL.

Functions use the same syntax conventions used by SQL statements.

See also

- “Syntax conventions” [[SQL Anywhere Server - SQL Reference](#)]

Function types

This section groups the available function by type.

UltraLite supports a subset of the same functions documented for SQL Anywhere, and sometimes with a few differences.

Note

Unless otherwise stated, any function that receives NULL as a parameter returns NULL.

See also

- [“Functions for spatial data” on page 39](#)

UltraLite aggregate functions

Aggregate functions summarize data over a group of rows from the database. The groups are formed using the GROUP BY clause of the SELECT statement. Aggregate functions are allowed only in the SELECT list and in the HAVING and ORDER BY clauses of a SELECT statement.

List of functions

The following aggregate functions are available:

- [“AVG function \[Aggregate\]”](#)
- [“COUNT function \[Aggregate\]”](#)
- [“COUNT_UPLOAD_ROWS function \[Aggregate\]”](#)
- [“LIST function \[Aggregate\]”](#)
- [“MAX function \[Aggregate\]”](#)
- [“MIN function \[Aggregate\]”](#)
- [“SUM function \[Aggregate\]”](#)

UltraLite data type conversion functions

Data type conversion functions are used to convert arguments from one data type to another, or to test whether they can be converted.

List of functions

The following data type conversion functions are available:

- [“CAST function \[Data type conversion\]”](#)
- [“CONVERT function \[Data type conversion\]”](#)
- [“HEXTOINT function \[Data type conversion\]”](#)
- [“INTTOHEX function \[Data type conversion\]”](#)
- [“ISDATE function \[Data type conversion\]”](#)

UltraLite date and time functions

Date and time functions perform operations on DATE, TIME, TIMESTAMP, and TIMESTAMP WITH TIME ZONE data types.

SQL Anywhere includes compatibility support for Transact-SQL date and time types, including DATETIME and SMALLDATETIME. These Transact-SQL data types are implemented as domains over the native SQL Anywhere TIMESTAMP data type.

Specifying date parts

Many of the date functions use dates built from **date parts**. The following table displays allowed values of date parts.

When using date and time functions, you can specify a minus sign to subtract from a date or time. For example, to get a timestamp from 31 days ago, you can execute the following:

```
SELECT DATEADD(day, -31, NOW());
```

Date part	Abbreviation	Values
Year	yy	1-9999
Quarter	qq	1-4
Month	mm	1-12
Week	wk	1-54. Weeks begin on Sunday. A 54-week year occurs in leap years that start on a Saturday.
Day	dd	1-31
Dayofyear	dy	1-366
Weekday	dw	1-7 (Sunday = 1, ..., Saturday = 7)
Hour	hh	0-23
Minute	mi	0-59
Second	ss	0-59
Millisecond	ms	0-999
Microsecond	mcs or us	0-999999
Calyearofweek	cyr	1-9999. The year in which the week begins. The week containing the first few days of the year may have started in the previous year, depending on the weekday on which the year started. Years starting on Monday through Thursday have no days that are part of the previous year, but years starting on Friday through Sunday start their first week on the first Monday of the year.

Date part	Abbreviation	Values
Calweekofyear	cwk	1-53. The week number within the year that contains the specified date. For more information about the ISO week system and the ISO 8601 date and time standard, see http://en.wikipedia.org/wiki/ISO_week_date .
Caldayofweek	cdw	1-7. (Monday = 1, ..., Sunday = 7)
TZOffset	tz	-840 to 840

List of date and time functions

The following date and time functions are available:

- “DATE function [Date and time]”
- “DATEADD function [Date and time]”
- “DATEDIFF function [Date and time]”
- “DATEFORMAT function [Date and time]”
- “DATENAME function [Date and time]”
- “DATEPART function [Date and time]”
- “DATETIME function [Date and time]”
- “DAY function [Date and time]”
- “DAYNAME function [Date and time]”
- “DAYS function [Date and time]”
- “DOW function [Date and time]”
- “GETDATE function [Date and time]”
- “HOUR function [Date and time]”
- “HOURS function [Date and time]”
- “MINUTE function [Date and time]”
- “MINUTES function [Date and time]”
- “MONTH function [Date and time]”
- “MONTHNAME function [Date and time]”
- “MONTHS function [Date and time]”
- “NOW function [Date and time]”
- “QUARTER function [Date and time]”
- “SECOND function [Date and time]”
- “SECONDS function [Date and time]”
- “SWITCHOFFSET function [Date and time]”
- “TODAY function [Date and time]”
- “TODATETIMEOFFSET function [Date and time]”
- “WEEKS function [Date and time]”
- “YEAR function [Date and time]”
- “YEARS function [Date and time]”
- “YMD function [Date and time]”

See also

- “Date and time data types” [[SQL Anywhere Server - SQL Reference](#)]
- “Data types in UltraLite” on page 265

UltraLite miscellaneous functions

Miscellaneous functions perform operations on arithmetic, string, or date/time expressions, including the return values of other functions.

List of functions

The following miscellaneous functions are available:

- “[ARGN function \[Miscellaneous\]](#)”
- “[COALESCE function \[Miscellaneous\]](#)”
- “[EXPLANATION function \[Miscellaneous\]](#)”
- “[GREATER function \[Miscellaneous\]](#)”
- “[IFNULL function \[Miscellaneous\]](#)”
- “[ISNULL function \[Miscellaneous\]](#)”
- “[LESSER function \[Miscellaneous\]](#)”
- “[NEWID function \[Miscellaneous\]](#)”
- “[NULLIF function \[Miscellaneous\]](#)”

UltraLite numeric functions

Numeric functions perform mathematical operations on numerical data types or return numeric information.

List of functions

The following numeric functions are available:

- “ABS function [Numeric]”
- “ACOS function [Numeric]”
- “ASIN function [Numeric]”
- “ATAN function [Numeric]”
- “ATAN2 function [Numeric]”
- “CEILING function [Numeric]”
- “COS function [Numeric]”
- “COT function [Numeric]”
- “DEGREES function [Numeric]”
- “EXP function [Numeric]”
- “FLOOR function [Numeric]”
- “LOG function [Numeric]”
- “LOG10 function [Numeric]”
- “MOD function [Numeric]”
- “PI function [Numeric]”
- “POWER function [Numeric]”
- “RADIANS function [Numeric]”
- “REMAINDER function [Numeric]”
- “ROUND function [Numeric]”
- “SIGN function [Numeric]”
- “SIN function [Numeric]”
- “SQRT function [Numeric]”
- “TAN function [Numeric]”
- “TRUNCNUM function [Numeric]”

UltraLite spatial functions

Spatial data is data that describes the position, shape, and orientation of objects in a defined space.

UltraLite provides storage and data management features for spatial data, in the form of points, allowing you to store information such as geographic locations and routing information, for instance. UltraLite provides a set of SQL spatial functions designed for compatibility with other products. You use these functions and constructors to access and manipulate the spatial data.

List of functions

The following spatial functions are available:

- “ST_AsBinary function [Spatial]”
- “ST_AsExtText function [Spatial]”
- “ST_AsText function [Spatial]”
- “ST_Distance function [Spatial]”
- “ST_Equals function [Spatial]”
- “ST_IntersectsRect function [Spatial]”
- “ST_Point function [Spatial]”
- “ST_PointFromExtText function [Spatial]”
- “ST_PointFromText function [Spatial]”
- “ST_PointFromWKB function [Spatial]”
- “ST_SRID function [Spatial]”
- “ST_X function [Spatial]”
- “ST_Y function [Spatial]”

UltraLite string functions

String functions perform conversion, extraction, or manipulation operations on strings, or return information about strings.

When working in a multibyte character set, check carefully whether the function being used returns information concerning characters or bytes.

List of functions

The following string functions are available:

- “ASCII function [String]”
- “BYTE_LENGTH function [String]”
- “BYTE_SUBSTR function [String]”
- “CHAR function [String]”
- “CHARINDEX function [String]”
- “CHAR_LENGTH function [String]”
- “DIFFERENCE function [String]”
- “INSERTSTR function [String]”
- “LCASE function [String]”
- “LEFT function [String]”
- “LENGTH function [String]”
- “LOCATE function [String]”
- “LOWER function [String]”
- “LTRIM function [String]”
- “PATINDEX function [String]”
- “REPEAT function [String]”
- “REPLACE function [String]”
- “REPLICATE function [String]”
- “RIGHT function [String]”
- “RTRIM function [String]”
- “SIMILAR function [String]”
- “SOUNDEX function [String]”
- “SPACE function [String]”
- “STR function [String]”
- “STRING function [String]”
- “STRTOUUID function [String]”
- “STUFF function [String]”
- “SUBSTRING function [String]”
- “TRIM function [String]”
- “UCASE function [String]”
- “UPPER function [String]”
- “UIDTOSTR function [String]”

UltraLite system functions

System functions return system information.

List of functions

The following system functions are available in UltraLite:

- “DB_PROPERTY function [System]”
- “ML_GET_SERVER_NOTIFICATION function [System]”
- “SYNC_PROFILE_OPTION_VALUE function [System]”

Functions

Each function is listed, and the function type (numeric, character, and so on) is indicated next to it.

For links to all functions of a given type, see [“Function types” on page 299](#).

ABS function [Numeric]

Returns the absolute value of a numeric expression.

Syntax

ABS(*numeric-expression*)

Parameters

- ***numeric-expression*** The number whose absolute value is to be returned.

Returns

An absolute value of the numeric expression.

Numeric-expression data type	Returns
INT	INT
FLOAT	FLOAT
DOUBLE	DOUBLE
NUMERIC	NUMERIC

Standards and compatibility

- **SQL/2008** The ABS function is part of optional SQL/2008 language feature T441.

Example

The following statement returns the value 66.

```
SELECT ABS( -66 );
```

ACOS function [Numeric]

Returns the arc-cosine, in radians, of a numeric expression. Not supported by UltraLite Java edition databases.

Syntax

ACOS(*numeric-expression*)

Parameters

- ***numeric-expression*** The cosine of the angle.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“ASIN function \[Numeric\]” on page 309](#)
- [“ATAN function \[Numeric\]” on page 310](#)
- [“ATAN2 function \[Numeric\]” on page 311](#)
- [“COS function \[Numeric\]” on page 322](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the arc-cosine value for 0.52.

```
SELECT ACOS( 0.52 );
```

ARGN function [Miscellaneous]

Returns a selected argument from a list of arguments.

Syntax

```
ARGN(integer-expression, expression [ , ... ])
```

Parameters

- ***integer-expression*** The position of an argument within the list of expressions.
- ***expression*** An expression of any data type passed into the function. All supplied expressions must be of the same data type.

Returns

Using the value of the *integer-expression* as n, returns the nth argument (starting at 1) from the remaining list of arguments.

Remarks

While the expressions can be of any data type, they must all be of the same data type. The integer expression must be from one to the number of expressions in the list or NULL is returned. Multiple expressions are separated by a comma.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 6.

```
SELECT ARGN( 6, 1,2,3,4,5,6 );
```

ASCII function [String]

Returns the integer ASCII value of the first byte in a string-expression.

Syntax

ASCII(*string-expression*)

Parameters

- ***string-expression*** The string.

Returns

SMALLINT

Remarks

If the string is empty, then ASCII returns zero. Literal strings must be enclosed in quotes. If the database character set is multibyte and the first character of the parameter string consists of more than one byte, the result is NULL.

See also

- [“CHAR function \[String\]” on page 316](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 90.

```
SELECT ASCII( 'Z' );
```

ASIN function [Numeric]

Returns the arc-sine, in radians, of a number.

Syntax

ASIN(*numeric-expression*)

Parameters

- ***numeric-expression*** The sine of the angle.

Returns

DOUBLE

Remarks

The SIN and ASIN functions are inverse operations.

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“ACOS function \[Numeric\]” on page 307](#)
- [“ATAN function \[Numeric\]” on page 310](#)
- [“ATAN2 function \[Numeric\]” on page 311](#)
- [“SIN function \[Numeric\]” on page 382](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the arc-sine value for 0.52.

```
SELECT ASIN( 0.52 );
```

ATAN function [Numeric]

Returns the arc-tangent, in radians, of a number. Not supported by UltraLite Java edition databases.

Syntax

```
ATAN( numeric-expression )
```

Parameters

- ***numeric-expression*** The tangent of the angle.

Remarks

The ATAN and TAN functions are inverse operations.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“ACOS function \[Numeric\]” on page 307](#)
- [“ASIN function \[Numeric\]” on page 309](#)
- [“ATAN2 function \[Numeric\]” on page 311](#)
- [“TAN function \[Numeric\]” on page 399](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the arc-tangent value for 0.52.

```
SELECT ATAN( 0.52 );
```

ATAN2 function [Numeric]

Returns the arc-tangent, in radians, of the ratio of two numbers. Not supported by UltraLite Java edition databases.

Syntax

ATAN2 (*numeric-expression-1*, *numeric-expression-2*)

Parameters

- ***numeric-expression-1*** The numerator in the ratio whose arc-tangent is calculated.
- ***numeric-expression-2*** The denominator in the ratio whose arc-tangent is calculated.

Returns

DOUBLE

Remarks

This function converts its arguments to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“ACOS function \[Numeric\]” on page 307](#)
- [“ASIN function \[Numeric\]” on page 309](#)
- [“ATAN function \[Numeric\]” on page 310](#)
- [“TAN function \[Numeric\]” on page 399](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the arc-tangent value for the ratio 0.52 to 0.60.

```
SELECT ATAN2( 0.52, 0.60 );
```

AVG function [Aggregate]

Computes the average, for a set of rows, of a numeric expression or of a set of unique values.

Syntax 1

```
AVG( [ DISTINCT ] numeric-expression )
```

Parameters

- **[ALL] *numeric-expression*** The expression whose average is calculated over the rows in each group.
- **DISTINCT clause** Computes the average of the unique numeric values in each group.

Returns

Returns the NULL value for a group containing no rows.

Returns DOUBLE if the argument is DOUBLE, otherwise NUMERIC.

Remarks

This average does not include rows where the *numeric-expression* is the NULL value.

This function can generate an overflow error, resulting in an error being returned. You can use the CAST function on *numeric-expression* to avoid the overflow error.

See also

- [“SUM function \[Aggregate\]” on page 397](#)
- [“COUNT function \[Aggregate\]” on page 323](#)

Standards and compatibility

- **SQL/2008** AVG is a core feature of the SQL/2008 standard.

See also

- [“Aggregate functions and outer references” \[SQL Anywhere Server - SQL Usage\]](#)

Example

The following statement returns the value 49988.623200.

```
SELECT AVG( Salary ) FROM Employees;
```

The following statement returns the average product price from the Products table:

```
SELECT AVG( DISTINCT UnitPrice ) FROM Products;
```


BYTE_LENGTH function [String]

Returns the number of bytes in a string.

Syntax

BYTE_LENGTH(*string-expression*)

Parameters

- ***string-expression*** The string whose length is to be calculated.

Returns

INT

Remarks

Trailing white space characters in the *string-expression* are included in the length returned.

The return value of a NULL string is NULL.

If the string is in a multibyte character set, the **BYTE_LENGTH** value may differ from the number of characters returned by **CHAR_LENGTH**.

See also

- [“CHAR_LENGTH function \[String\]” on page 317](#)
- [“DATALENGTH function \[System\]” on page 325](#)
- [“LENGTH function \[String\]” on page 349](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension. The equivalent function in the SQL/2008 standard is the **OCTET_LENGTH** function.

Example

The following statement returns the value 12.

```
SELECT BYTE_LENGTH( 'Test Message' );
```

BYTE_SUBSTR function [String]

Returns a substring of a string. The substring is calculated using bytes, not characters.

Syntax

BYTE_SUBSTR(*string-expression*, *start* [, *length*])

Parameters

- ***string-expression*** The string from which the substring is taken.

- **start** An integer expression indicating the start of the substring. A positive integer starts from the beginning of the string, with the first character being position 1. A negative integer specifies a substring starting from the end of the string, the final character being at position -1.
- **length** An integer expression indicating the length of the substring. A positive *length* specifies the number of bytes to be taken *starting* at the start position. A negative *length* returns at most *length* bytes up to, and including, the starting position, from the left of the starting position.

Returns

BINARY, or VARCHAR. The value returned depends on the type of *string-expression*. Also, the arguments you specify determine if the returned value is LONG. For example, LONG is not returned when you specify a constant < 32K for length.

Remarks

If *length* is specified, the substring is restricted to that number of bytes. Both *start* and *length* can be either positive or negative. Using appropriate combinations of negative and positive numbers, you can get a substring from either the beginning or end of the string.

If *start* is zero and length is non-negative, a *start* value of 1 is used. If *start* is zero and *length* is negative, a start value of -1 is used.

See also

- [“SUBSTRING function \[String\]” on page 395](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value Test.

```
SELECT BYTE_SUBSTR( 'Test Message', 1, 4 );
```

CAST function [Data type conversion]

Returns the value of an expression converted to a supplied data type.

The CAST, CONVERT, HEXTOINT, and INTTOHEX functions can be used to convert to and from hexadecimal values. For more information about using these functions, see [“Converting to and from hexadecimal values” \[SQL Anywhere Server - SQL Reference\]](#).

Syntax

```
CAST( expression AS datatype )
```

Parameters

- **expression** The expression to be converted.

- ***datatype*** The target data type.

Returns

Depends on the data type requested.

Remarks

If you do not indicate a length for character string types, the database server chooses an appropriate length. If neither precision nor scale is specified for a DECIMAL conversion, the database server selects appropriate values.

See also

- “Data type conversions” [[SQL Anywhere Server - SQL Reference](#)]
- “CONVERT function [Data type conversion]” on page 319
- “LEFT function [String]” on page 348

Standards and compatibility

- **SQL/2008** The CAST function is a core feature of the SQL/2008 standard. However, in SQL Anywhere CAST supports a number of data type conversions that are not permitted by the SQL standard. For example, in SQL Anywhere you can CAST an integer value to a DATE type, whereas in the SQL standard this type conversion is not permitted. For more information, see “Data type conversions” [[SQL Anywhere Server - SQL Reference](#)].

Example

The following function ensures a string is used as a date:

```
SELECT CAST( '2000-10-31' AS DATE );
```

The value of the expression 1 + 2 is calculated, and the result is then cast into a single-character string.

```
SELECT CAST( 1 + 2 AS CHAR );
```

You can use the CAST function to shorten strings

```
SELECT CAST ( 'Surname' AS CHAR(5) );
```

CEILING function [Numeric]

Returns the first integer that is greater or equal to a given value. For positive numbers, this is known as rounding up.

Syntax

```
{ CEILING | CEIL } ( numeric-expression )
```

Parameters

- ***numeric-expression*** The number whose ceiling is to be calculated.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“FLOOR function \[Numeric\]” on page 339](#)

Standards and compatibility

- **SQL/2008** The CEILING function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The following statement returns the value 60.

```
SELECT CEILING( 59.84567 );
```

CHAR function [String]

Returns the character with the ASCII value of a number.

Syntax

CHAR(*integer-expression*)

Parameters

- ***integer-expression*** The number to be converted to an ASCII character. The number must be in the range 0 to 255, inclusive.

Returns

VARCHAR

Remarks

The character returned corresponds to the supplied numeric expression in the current database character set, according to a binary sort order.

CHAR returns NULL for integer expressions with values greater than 255 or less than zero.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value Y.

```
SELECT CHAR( 89 );
```

CHAR_LENGTH function [String]

Returns the number of characters in a string.

Syntax

```
CHAR_LENGTH ( string-expression )
```

Parameters

- ***string-expression*** The string whose length is to be calculated.

Returns

INT

Remarks

Trailing white space characters are included in the length returned.

The return value of a NULL string is NULL.

If the string is in a multibyte character set, the value returned by the CHAR_LENGTH function may differ from the number of bytes returned by the BYTE_LENGTH function.

Note

You can use the CHAR_LENGTH function and the LENGTH function interchangeably for CHAR, VARCHAR, and LONG VARCHAR data types. However, you must use the LENGTH function for BINARY and bit array data types.

See also

- [“BYTE_LENGTH function \[String\]” on page 313](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** CHAR_LENGTH is a core feature of the SQL/2008 standard.

Example

The following statement returns the value 8.

```
SELECT CHAR_LENGTH( 'Chemical' );
```

CHARINDEX function [String]

Returns the position of one string in another.

Syntax

```
CHARINDEX( string-expression-1, string-expression-2 )
```

Parameters

- ***string-expression-1*** The string for which you are searching.
- ***string-expression-2*** The string to be searched.

Returns

INT

Remarks

The first character of *string-expression-1* is identified as 1. If the string being searched contains more than one instance of the other string, then the CHARINDEX function returns the position of the first instance.

If the string being searched does not contain the other string, then the CHARINDEX function returns 0.

If any of the arguments are NULL, the result is NULL.

See also

- [“SUBSTRING function \[String\]” on page 395](#)
- [“REPLACE function \[String\]” on page 373](#)
- [“LOCATE function \[String\]” on page 351](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns last and first names from the Surname and GivenName columns of the Employees table, but only when the last name includes the letter K:

```
SELECT Surname, GivenName
FROM Employees
WHERE CHARINDEX( 'K', Surname ) = 1;
```

Results returned:

Surname	GivenName
Klobucher	James
Kuo	Felicia

Surname	GivenName
Kelly	Moira

COALESCE function [Miscellaneous]

Returns the first non-NULL expression from a list. This function is identical to the ISNULL function.

Syntax

COALESCE(*expression*, *expression* [, ...])

Parameters

- ***expression*** Any expression.

At least two expressions must be passed into the function, and all expressions must be comparable.

Returns

The return type for this function depends on the expressions specified. That is, when the database server evaluates the function, it first searches for a data type in which all the expressions can be compared. When found, the database server compares the expressions and then returns the result in the type used for the comparison. If the database server cannot find a common comparison type, an error is returned.

Remarks

The result is NULL only if all the arguments are NULL.

The parameters can be of any scalar type, but not necessarily same type.

For a more detailed description of how the database server processes this function, see [“ISNULL function \[Miscellaneous\]” on page 346](#).

See also

- [“ISNULL function \[Miscellaneous\]” on page 346](#)

Standards and compatibility

- **SQL/2008** Core feature.

Example

The following statement returns the value 34.

```
SELECT COALESCE( NULL, 34, 13, 0 );
```

CONVERT function [Data type conversion]

Returns an expression converted to a supplied data type.

The CAST, CONVERT, HEXTOINT, and INTTOHEX functions can be used to convert to and from hexadecimal values. For more information about using these functions, see [“Converting to and from hexadecimal values” \[SQL Anywhere Server - SQL Reference\]](#).

Syntax

CONVERT(*datatype*, *expression* [, *format-style*])

Parameters

- **datatype** The data type to which the expression is converted.
- **expression** The expression to be converted.
- **format-style** The style code to apply to the output value. Use this parameter when converting strings to date or time data types, and vice versa. The table below shows the supported style codes, followed by a representation of the output format produced by that style code. The style codes are separated into two columns, depending on whether the century is included in the output format (for example, 06 versus 2006).

Without century (yy) style codes	With century (yyyy) style codes	Output format
-	0 or 100	Mmm dd yyyy hh:nnAA
1	101	mm/dd/yy[yy]
2	102	[yy]yy.mm.dd
3	103	dd/mm/yy[yy]
4	104	dd.mm.yy[yy]
5	105	dd-mm-yy[yy]
6	106	dd Mmm yy[yy]
7	107	Mmm dd, yy[yy]
8	108	hh:nn:ss
-	9 or 109	Mmm dd yyyy hh:nn:ss:sssAA
10	110	mm-dd-yy[yy]
11	111	[yy]yy/mm/dd
12	112	[yy]yyymmdd
-	13 or 113	dd Mmm yyyy hh:nn:ss:sss (24 hour clock, Europe default + milliseconds, 4-digit year)

Without century (yy) style codes	With century (yyyy) style codes	Output format
-	14 or 114	hh:nn:ss:sss (24 hour clock)
-	20 or 120	yyyy-mm-dd hh:nn:ss (24-hour clock, ODBC canonical, 4-digit year)
-	21 or 121	yyyy-mm-dd hh:nn:ss.sss (24 hour clock, ODBC canonical with milliseconds, 4-digit year)

Returns

Depends on the data type specified.

Remarks

If no *format-style* argument is provided, style code 0 is used.

The CONVERT function can be used to convert a string to a DATE, TIME, or TIMESTAMP data type, provided that there is no ambiguity when parsing the string. If *format-style* is specified, the database server may use it as a hint on how to parse the string. An error is returned if the database server cannot parse the string unambiguously.

For information about the styles produced by each output symbol (such as Mmm), see [“UltraLite date_format creation parameter” on page 144](#).

See also

- [“CAST function \[Data type conversion\]” on page 314](#)
- [“CSCONVERT function \[String\]” \[SQL Anywhere Server - SQL Reference\]](#)

Standards and compatibility

- **SQL/2008** Vendor extension. The CONVERT function is defined in the SQL/2008 standard. However, in the SQL standard the purpose of CONVERT is to perform a transcoding of the input string expression to a different character set, which is implemented in SQL Anywhere as the CSCONVERT function.

Example

The following statements illustrate the use of format style.

```
SELECT CONVERT( CHAR( 20 ), OrderDate, 104 ) FROM SalesOrders;
```

OrderDate
16.03.2000
20.03.2000
23.03.2000

OrderDate
25.03.2000
...

```
SELECT CONVERT( CHAR( 20 ), OrderDate, 7 ) FROM SalesOrders;
```

OrderDate
Mar 16, 00
Mar 20, 00
Mar 23, 00
Mar 25, 00
...

The following statement illustrates conversion to an integer, and returns the value 5.

```
SELECT CONVERT( integer, 5.2 );
```

COS function [Numeric]

Returns the cosine of the angle in radians given by its argument.

Syntax

```
COS( numeric-expression )
```

Parameters

- ***numeric-expression*** The angle, in radians.

Returns

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

See also

- [“ACOS function \[Numeric\]” on page 307](#)
- [“COT function \[Numeric\]” on page 323](#)
- [“SIN function \[Numeric\]” on page 382](#)
- [“TAN function \[Numeric\]” on page 399](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value of the cosine of an angle 0.52 radians.

```
SELECT COS( 0.52 );
```

COT function [Numeric]

Returns the cotangent of the angle in radians given by its argument.

Syntax

```
COT( numeric-expression )
```

Parameters

- ***numeric-expression*** The angle, in radians.

Returns

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

See also

- [“COS function \[Numeric\]” on page 322](#)
- [“SIN function \[Numeric\]” on page 382](#)
- [“TAN function \[Numeric\]” on page 399](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the cotangent value of 0.52.

```
SELECT COT( 0.52 );
```

COUNT function [Aggregate]

Counts the number of rows in a group depending on the specified parameters.

Syntax 1

```
COUNT( [ * | [ DISTINCT ] expression ] )
```

Parameters

- ***** Return the number of rows in each group. COUNT(*) and COUNT() are semantically equivalent.
- ***expression*** Return the number of rows in each group where the value of *expression* is not null.

- **DISTINCT *expression*** Return the number of distinct values of *expression* for all of the rows in each group where *expression* is not NULL.

Returns

The COUNT function returns a value of type INT.

COUNT never returns the value NULL. If a group contains no rows, or if there are no non-NULL values of *expression* in a group, then COUNT returns 0.

Remarks

The COUNT function returns a maximum value of 2147483647.

See also

- [“AVG function \[Aggregate\]” on page 312](#)
- [“SUM function \[Aggregate\]” on page 397](#)

Standards and compatibility

- **SQL/2008** Core feature.

See also

- [“Aggregate functions and outer references” \[SQL Anywhere Server - SQL Usage\]](#)

Example

The following statement returns each unique city, and the number of employees working in that city.

```
SELECT City, COUNT( * ) FROM Employees GROUP BY City;
```

COUNT_UPLOAD_ROWS function [Aggregate]

Returns a count of the number of rows that will be uploaded in the next synchronization.

Syntax

```
COUNT_UPLOAD_ROWS( pubs,threshold )
```

Parameters

- ***pubs*** A comma-separated list of publications to check for rows.
- ***threshold*** The maximum number of rows to count (a value of 0 corresponds to the maximum limit).

Returns

INT

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following returns the total number of rows to upload in *mypub1* and *mypub2*:

```
SELECT COUNT_UPLOAD_ROWS( 'mypub1,mypub2' , 0 );
```

DATALENGTH function [System]

Returns the length, in bytes, of the underlying storage for the result of an expression.

Syntax

```
DATALENGTH( expression )
```

Parameters

- **expression** Usually a column name. If *expression* is a string constant, you must enclose it in quotes.

Returns

UNSIGNED INT

Remarks

The return values of the DATALENGTH function are as follows:

Data type	DATALENGTH
SMALLINT	2
INTEGER	4
DOUBLE	8
CHAR	Length of the data
BINARY	Length of the data

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the length of the longest string in the *CompanyName* column.

```
SELECT MAX( DATALENGTH( CompanyName ) )
FROM Customers;
```

The following statement returns the length of the string '8sdofinsv8s7a7s7gehe4h':

```
SELECT DATALENGTH( '8sdofinsv8s7a7s7gehe4h' );
```

DATE function [Date and time]

Converts the expression into a date, and removes any hours, minutes, or seconds.

For information about controlling the interpretation of date formats, see [“UltraLite date_order creation parameter” on page 147](#).

Syntax

DATE(*expression*)

Returns

DATE

Parameters

- **expression** The value to be converted to date format, typically a string.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 1999-01-02 as a date.

```
SELECT DATE( '1999-01-02 21:20:53' );
```

The following statement returns the create dates of all the objects listed in the SYSOBJECT system view:

```
SELECT DATE( creation_time ) FROM SYSOBJECT;
```

DATEADD function [Date and time]

Returns a **TIMESTAMP** or **TIMESTAMP WITH TIME ZONE** value produced by adding a date part to its argument.

Syntax

DATEADD(*date-part*, *integer-expression*, *timestamp-expression*)

date-part :

```
year  
| quarter  
| month  
| week  
| day  
| dayofyear  
| hour  
| minute  
| second  
| millisecond  
| microsecond
```

Parameters

- **date-part** The date part that *integer-expression* represents.

For a complete listing of allowed date parts, see [“Specifying date parts” on page 301](#).

- **integer-expression** The number of *date-part* values to be added to *timestamp-expression*. Note that *integer-expression* can be any numeric type, but its value is truncated to an INTEGER.
- **timestamp-expression** The TIMESTAMP or TIMESTAMP WITH TIME ZONE value to be modified.

Returns

TIMESTAMP WITH TIME ZONE if *timestamp-expression* is a TIMESTAMP WITH TIME ZONE; otherwise TIMESTAMP.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the TIMESTAMP value 1995-11-02 00:00:00.000.

```
SELECT DATEADD( month, 102, '1987/05/02' );
```

The following statement returns the TIMESTAMP value 1987-05-02 04:00:00.000.

```
SELECT DATEADD( hour, 4, '1987/05/02' );
```

The following statement returns the TIMESTAMP WITH TIME ZONE value 1987-05-06 11:33:00.000+04:00

```
SELECT DATEADD( day, 4, CAST( '1987/05/02 11:33:00.000000+04:00' as
TIMESTAMP WITH TIME ZONE ));
```

DATEDIFF function [Date and time]

Returns the interval between two dates.

Syntax

```
DATEDIFF( date-part, date-expression-1, date-expression-2 )
```

```
date-part :
year
| quarter
| month
| week
| day
| dayofyear
| hour
| minute
| second
```

| millisecond
| microsecond

Parameters

- **date-part** Specifies the date part in which the interval is to be measured.

Choose one of the date objects listed above. For a complete list of date parts, see [“Specifying date parts” on page 301](#).

- **date-expression-1** The starting date for the interval. This value is subtracted from *date-expression-2* to return the number of *date-parts* between the two arguments.
- **date-expression-2** The ending date for the interval. *Date-expression-1* is subtracted from this value to return the number of *date-parts* between the two arguments.

Returns

INT with year, quarter, month, week, day, and dayofyear. BIGINT with hour, minute, second, millisecond, and microsecond.

Remarks

This function calculates the number of date parts between two specified dates. The result is a signed integer value equal to (date-expression-2 - date-expression-1), in date parts.

The DATEDIFF function results are truncated, not rounded, when the result is not an even multiple of the date part.

When you use **day** as the date part, the DATEDIFF function returns the number of midnights between the two times specified, including the second date but not the first.

When you use **month** as the date part, the DATEDIFF function returns the number of first-of-the-months between two dates, including the second date but not the first.

When you use **week** as the date part, the DATEDIFF function returns the number of Sundays between the two dates, including the second date but not the first.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns 1.

```
SELECT DATEDIFF( hour, '4:00AM', '5:50AM' );
```

The following statement returns 102.

```
SELECT DATEDIFF( month, '1987/05/02', '1995/11/15' );
```

The following statement returns 0.

```
SELECT DATEDIFF( day, '00:00', '23:59' );
```


The following statement returns 4.

```
SELECT DATEDIFF( day,
  '1999/07/19 00:00',
  '1999/07/23 23:59' );
```

The following statement returns 0.

```
SELECT DATEDIFF( month, '1999/07/19', '1999/07/23' );
```

The following statement returns 1.

```
SELECT DATEDIFF( month, '1999/07/19', '1999/08/23' );
```

DATEFORMAT function [Date and time]

Returns a string representing a date expression in the specified format.

Syntax

```
DATEFORMAT( datetime-expression, string-expression )
```

Parameters

- ***datetime-expression*** The datetime to be converted.
- ***string-expression*** The format of the converted date.

For information about date format descriptions, see [“UltraLite date_format creation parameter” on page 144](#).

Returns

VARCHAR

Remarks

Any allowable date format can be used for the string-expression.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value Jan 01, 1989.

```
SELECT DATEFORMAT( '1989-01-01', 'Mmm dd, yyyy' );
```

DATENAME function [Date and time]

Returns the name of the specified part (such as the month June) of a TIMESTAMP or TIMESTAMP WITH TIME ZONE value, as a character string.

Syntax

DATENAME(*date-part*, *timestamp-expression*)

Parameters

- ***date-part*** The date part to be named.

For a complete listing of allowed date parts, see [“Specifying date parts” on page 301](#).

- ***timestamp-expression*** The **TIMESTAMP** or **TIMESTAMP WITH TIME ZONE** value for which the date part name is to be returned. For meaningful results, *timestamp-expression* should contain the requested *date-part*.

Returns

VARCHAR

Remarks

The DATENAME function returns a string, even if the result is numeric, such as 23 for the day.

See also

- [“DATEPART function \[Date and time\]” on page 330](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value May.

```
SELECT DATENAME( month, '1987/05/02' );
```

DATEPART function [Date and time]

Syntax

DATEPART(*date-part*, *timestamp-expression*)

Parameters

- ***date-part*** The date part to be returned.

For a complete listing of allowed date parts, see [“Specifying date parts” on page 301](#).

- ***timestamp-expression*** The **TIMESTAMP** or **TIMESTAMP WITH TIME ZONE** value for which the part is to be returned.

Returns

INT

Remarks

For meaningful results *timestamp-expression* should contain the required *date-part* portion.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 5.

```
SELECT DATEPART( month , '1987/05/02' );
```

The following example creates a table, TableStatistics, and inserts into it the total number of sales orders per year as stored in the SalesOrders table:

```
CREATE TABLE TableStatistics (
  ID INTEGER NOT NULL DEFAULT AUTOINCREMENT,
  Year INT,
  NumberOrders INT );
INSERT INTO TableStatistics ( Year, NumberOrders )
SELECT DATEPART( Year, OrderDate ), COUNT(*)
FROM SalesOrders
GROUP BY DATEPART( Year, OrderDate );
```

DATETIME function [Date and time]

Converts an expression into a **TIMESTAMP** value.

Syntax

```
DATETIME( expression )
```

Parameters

- ***expression*** The expression to be converted. It is generally a string.

Returns

TIMESTAMP

Remarks

Attempts to convert numerical values return an error.

See also

- [“CAST function \[Data type conversion\]” on page 314](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns a timestamp with value 1998-09-09 12:12:12.000.

```
SELECT DATETIME( '1998-09-09 12:12:12.000' );
```

DAY function [Date and time]

Returns the day of the month of its argument as an integer between 1 and 31.

Syntax

```
DAY( date-expression )
```

Parameters

- ***date-expression*** The date as a DATE data type.

Returns

SMALLINT

Remarks

The DAY function returns an integer between 1 and 31, corresponding to the day of the month in the argument.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 12.

```
SELECT DAY( '2001-09-12' );
```

DAYNAME function [Date and time]

Returns the name of the day of the week from a date.

Syntax

```
DAYNAME( date-expression )
```

Parameters

- ***date-expression*** The date.

Returns

VARCHAR

Remarks

The English names are returned as: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value Saturday.

```
SELECT DAYNAME ( '1987/05/02' );
```

DAYS function [Date and time]

The DAYS function manipulates a **TIMESTAMP**, or returns the number of days between two **TIMESTAMP** values. For specific details, see the Remarks section below.

Syntax 1

```
DAYS( timestamp-expression )
```

Syntax 2

```
DAYS( timestamp-expression, timestamp-expression )
```

Syntax 3

```
DAYS( timestamp-expression, integer-expression )
```

Parameters

- ***timestamp-expression*** A **TIMESTAMP** value.
- ***integer-expression*** The number of days to be added to the *timestamp-expression*. If the *integer-expression* is negative, the appropriate number of days is subtracted from *timestamp-expression*. If you supply an integer expression, the *timestamp-expression* must be explicitly cast as a **TIME**, **DATE** or **TIMESTAMP**. If *timestamp-expression* is a **TIME** value, the current date is assumed.

For information about casting data types, see “[CAST function \[Data type conversion\]](#)” on page 314.

Returns

INTEGER with Syntax 1 or Syntax 2.

TIMESTAMP with Syntax 3.

Remarks

The result of the DAYS function depends on its arguments. The DAYS function ignores hours, minutes, and seconds in its arguments.

- **Syntax 1** If you pass a single *timestamp-expression* to the DAYS function, it will return the number of days between 0000-02-29 and *timestamp-expression* as an **INTEGER**.

Note

0000-02-29 is not meant to imply an actual date; it is the default date used by the DAYS function.

- **Syntax 2** If you pass two `TIMESTAMP` values to the `DAYS` function, the function returns the integer number of days between them.
- **Syntax 3** If you pass a `TIMESTAMP` value and an integer to the `DAYS` function, the function returns the `TIMESTAMP` result of adding the integer number of days to the *timestamp-expression* argument.

Instead of Syntax 2, use the `DATEDIFF` function. Instead of Syntax 3, use the `DATEADD` function.

See also

- [“DATEDIFF function \[Date and time\]” on page 327](#)
- [“DATEADD function \[Date and time\]” on page 326](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the integer 729889.

```
SELECT DAYS( '1998-07-13 06:07:12' );
```

The following statements return the integer value -366, indicating that the second `DATE` value is 366 days before the first. It is recommended that you use the second example (`DATEDIFF`).

```
SELECT DAYS( '1998-07-13 06:07:12',  
            '1997-07-12 10:07:12' );
```

```
SELECT DATEDIFF( day,  
                '1998-07-13 06:07:12',  
                '1997-07-12 10:07:12' );
```

The following statements return the `TIMESTAMP` value 1999-07-14 00:00:00.000. It is recommended that you use the second example (`DATEADD`).

```
SELECT DAYS( CAST('1998-07-13' AS DATE ), 366 );  
SELECT DATEADD( day, 366, '1998-07-13' );
```

DB_PROPERTY function [System]

Returns the value of the given property. Not supported by UltraLite Java edition databases.

Syntax

```
DB_PROPERTY( property-name )
```

Parameters

- ***property-name*** The database property name.

Returns

VARCHAR, LONG VARCHAR

Remarks

Returns a string.

To set an option in UltraLite, use the SET OPTION statement or your component's API-specific Set Database Option method.



Cloud note: Because of tenant database isolation rules, when this function is called in the cloud it returns only information about the current tenant database.

See also

- [“SET OPTION statement \[UltraLite\]” on page 446](#)
- UltraLite C++: [“ULConnection.SetDatabaseOption method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“ULDatabaseSchema.SetDatabaseOption method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#)

Example

The following statement returns the page size of the current database, in bytes.

```
SELECT DB_PROPERTY( 'page_size');
```

DEGREES function [Numeric]

Converts a number from radians to degrees.

Syntax

```
DEGREES( numeric-expression )
```

Parameters

- ***numeric-expression*** An angle in radians.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns the degrees of the angle given by *numeric-expression*. If the parameter is NULL, the result is NULL.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 29.79380534680281.

```
SELECT DEGREES( 0.52 );
```

DIFFERENCE function [String]

Returns the difference in the SOUNDEX values between the two string expressions.

Syntax

DIFFERENCE (*string-expression-1*, *string-expression-2*)

Parameters

- ***string-expression-1*** The first SOUNDEX argument.
- ***string-expression-2*** The second SOUNDEX argument.

Returns

SMALLINT

Remarks

The DIFFERENCE function compares the SOUNDEX values of two strings and evaluates the similarity between them, returning a value from 0 through 4, where 4 is the best match.

This function always returns some value. The result is NULL only if one of the arguments are NULL.

See also

- [“SOUNDEX function \[String\]” on page 383](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns similarity between the words test and chest:

```
SELECT DIFFERENCE( 'test', 'chest' );
```

DOW function [Date and time]

Returns a number from 1 to 7 representing the day of the week of a date, where Sunday=1, Monday=2, and so on.

Syntax

DOW(*date-expression*)

Parameters

- ***date-expression*** The value (of type DATE) to be evaluated.

Returns

SMALLINT

Remarks

The DOW function is not affected by the value specified for the `first_day_of_week` database option. For example, even if `first_day_of_week` is set to Monday, the DOW function returns a 2 for Monday.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 5.

```
SELECT DOW( '1998-07-09' );
```

The following statement returns the value 1.

```
SELECT DOW( CAST( '2010/05/30 11:33:00.000000+04:00' as TIMESTAMP WITH TIME ZONE ) );
```

The following statement queries the `Employees` table and returns the employees `StartDate`, expressed as the number of the day of the week:

```
SELECT DOW( StartDate ) FROM Employees;
```

EXP function [Numeric]

Returns the result of the base of natural logarithms e raised to the power of the given argument. Not supported by UltraLite Java edition databases.

Syntax

```
EXP( numeric-expression )
```

Parameters

- ***numeric-expression*** The exponent.

Returns

DOUBLE

Remarks

The EXP function returns the result of raising the base of natural logarithms e by the value specified by *numeric-expression*.

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

Standards and compatibility

- **SQL/2008** The EXP function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The statement returns the value 3269017.3724721107.

```
SELECT EXP( 15 );
```

EXPLANATION function [Miscellaneous]

Returns the optimization strategy of a SQL statement as a plain text string.

Syntax

```
EXPLANATION( string-expression )
```

Parameters

- **string-expression** The SQL statement, which is commonly a SELECT statement, but can also be an UPDATE, MERGE, or DELETE statement.

Returns

LONG VARCHAR

Remarks

The statement's access plan is returned as a string. For information about interpreting the result, see [“Advanced: Query execution plans” \[SQL Anywhere Server - SQL Usage\]](#).

The GRAPHICAL_PLAN function offers significantly greater information about access plans, including system properties that may have affected how the statement was optimized.

This information can help you decide which indexes to add or how to structure your database for better performance.

See also

- [“Execution plans in UltraLite” on page 462](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement passes a SELECT statement as a string parameter and returns the plan for executing the query.

```
SELECT EXPLANATION( 'SELECT * FROM Departments WHERE DepartmentID > 100' );
```

FLOOR function [Numeric]

Returns the largest integer not greater than the given number.

Syntax

FLOOR(*numeric-expression*)

Parameters

- ***numeric-expression*** The value to be truncated, typically a fixed numeric type with non-zero scale or an approximate numeric type (DOUBLE, REAL, or FLOAT).

Returns

DOUBLE

Remarks

This function converts its arguments to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

See also

- [“CEILING function \[Numeric\]” on page 315](#)

Standards and compatibility

- **SQL/2008** The FLOOR function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The following statement returns a Floor value of 123:

```
SELECT FLOOR (123);
```

The following statement returns a value of 123:

```
SELECT FLOOR (123.45);
```

The following statement returns a value of -124:

```
SELECT FLOOR (-123.45);
```

GETDATE function [Date and time]

Returns the current year, month, day, hour, minute, second and fraction of a second.

Syntax

GETDATE()

Returns

TIMESTAMP

Remarks

The accuracy is limited by the accuracy of the system clock.

The information the GETDATE function returns is equivalent to the information returned by the NOW function and the CURRENT_TIMESTAMP special value.

See also

- [“NOW function \[Date and time\]” on page 364](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the system date and time.

```
SELECT GETDATE( );
```

GREATER function [Miscellaneous]

Returns the greater of two parameter values.

Syntax

```
GREATER( expression-1, expression-2 )
```

Parameters

- ***expression-1*** The first parameter value to be compared.
- ***expression-2*** The second parameter value to be compared.

Returns

The return type for this function depends on the expressions specified. That is, when the database server evaluates the function, it first searches for a data type in which all the expressions can be compared. When found, the database server compares the expressions and then returns the result in the type used for the comparison. If the database server cannot find a common comparison type, an error is returned.

Remarks

If the parameters are equal, the first is returned.

See also

- [“LESSER function \[Miscellaneous\]” on page 349](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 10.

```
SELECT GREATER( 10, 5 ) FROM dummy;
```

HEXTOINT function [Data type conversion]

Returns the decimal integer equivalent of a hexadecimal string.

The CAST, CONVERT, HEXTOINT, and INTTOHEX functions can be used to convert to and from hexadecimal values. For more information about using these functions, see [“Converting to and from hexadecimal values” \[SQL Anywhere Server - SQL Reference\]](#).

Syntax

```
HEXTOINT( hexadecimal-string )
```

Parameters

- ***hexadecimal-string*** The string to be converted to an integer.

Returns

The HEXTOINT function returns as INT the platform-independent SQL INTEGER equivalent of the hexadecimal string. The hexadecimal value represents a negative integer if the 8th digit from the right is one of the digits 8-9 and the uppercase or lowercase letters A-F and the previous leading digits are all uppercase or lowercase letter F. The following is not a valid use of HEXTOINT since the argument represents a positive integer value that cannot be represented as a signed 32-bit integer:

```
SELECT HEXTOINT( '0x0080000001' );
```

Remarks

The HEXTOINT function accepts string literals or variables consisting only of digits and the uppercase or lowercase letters A-F, with or without a 0x prefix. The following are all valid uses of HEXTOINT:

```
SELECT HEXTOINT( '0xFFFFFFFF' );  
SELECT HEXTOINT( '0x00000100' );  
SELECT HEXTOINT( '100' );  
SELECT HEXTOINT( '0xffffffff80000001' );
```

The HEXTOINT function removes the 0x prefix, if present. If the data exceeds 8 digits, it must represent a value that can be represented as a signed 32-bit integer value.

See also

- [“INTTOHEX function \[Data type conversion\]” on page 345](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 420.

```
SELECT HEXTOINT( '1A4' );
```

HOURLY function [Date and time]

Returns the hour component of a `TIMESTAMP` value.

Syntax

```
HOURLY( timestamp-expression )
```

Parameters

- ***timestamp-expression*** A `TIMESTAMP` value.

Returns

`SMALLINT`

Remarks

The value returned is the hour portion of the `TIMESTAMP` expression, a `SMALLINT` value between 0 and 23.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 21:

```
SELECT HOURLY( '1998-07-09 21:12:13' );
```

HOURS function [Date and time]

The `HOURS` function manipulates a `TIMESTAMP`, or returns the number of hours between two `TIMESTAMP` values. For specific details, see this function's usage.

Syntax 1

```
HOURS ( timestamp-expression )
```

Syntax 2

```
HOURS ( timestamp-expression, timestamp-expression )
```

Syntax 3

```
HOURS ( time-or-timestamp-expression, integer-expression )
```

Parameters

- ***time-or-timestamp-expression*** A value of type TIME or TIMESTAMP.
- ***timestamp-expression*** A value of type TIMESTAMP.
- ***integer-expression*** The number of hours to be added to *time-or-timestamp-expression*. If *integer-expression* is negative, the appropriate number of hours is subtracted from *time-or-timestamp-expression*.

For information about casting data types, see “[CAST function \[Data type conversion\]](#)” on page 314.

Returns

INTEGER with Syntax 1 or Syntax 2.

TIME or TIMESTAMP with Syntax 3.

Remarks

The result of the HOURS function depends on its arguments.

- **Syntax 1** If you pass a single *timestamp-expression* to the HOURS function, it will return the number of hours between midnight 0000-02-29 and *timestamp-expression* as an INTEGER.

Note

0000-02-29 is not meant to imply an actual date; it is the default TIMESTAMP value used by the HOURS function.

- **Syntax 2** If you pass two TIMESTAMP values to the HOURS function, the function returns the integer number of hours between them.
- **Syntax 3** If you pass a TIMESTAMP value and an INTEGER value to the HOURS function, the function returns the TIMESTAMP result of adding the integer number of hours to *time-or-timestamp-expression* argument. Similarly, if you pass a TIME value as the first argument, a TIME value is returned as the result. Syntax 3 does not support implicit conversion of the first argument. It may be necessary to explicitly cast the first argument to a DATE, TIME or TIMESTAMP value. If the first argument is a DATE, midnight is assumed for the time portion.

Instead of Syntax 2, use the DATEDIFF function. Instead of Syntax 3, use the DATEADD function.

See also

- “[DATEDIFF function \[Date and time\]](#)” on page 327
- “[DATEADD function \[Date and time\]](#)” on page 326

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements return the value 4, signifying that the second TIMESTAMP value is four hours after the first. It is recommended that you use the second example (DATEDIFF).

```
SELECT HOURS( '1999-07-13 06:07:12', '1999-07-13 10:07:12' );  
SELECT DATEDIFF( hour, '1999-07-13 06:07:12', '1999-07-13 10:07:12' );
```

The following statement returns the value 17517342.

```
SELECT HOURS( '1998-07-13 06:07:12' );
```

The following statements return the datetime 1999-05-13 02:05:07.000. It is recommended that you use the second example (DATEADD).

```
SELECT HOURS( CAST( '1999-05-12 21:05:07' AS DATETIME ), 5 );  
SELECT DATEADD( hour, 5, '1999-05-12 21:05:07' );
```

IFNULL function [Miscellaneous]

If the first expression is the NULL value, then the value of the second expression is returned. If the first expression is not NULL, the value of the third expression is returned. If the first expression is not NULL and there is no third expression, NULL is returned.

Syntax

```
IFNULL( expression-1, expression-2 [ , expression-3 ] )
```

Parameters

- ***expression-1*** The expression to be evaluated. Its value determines whether *expression-2* or *expression-3* is returned.
- ***expression-2*** The return value if *expression-1* is NULL.
- ***expression-3*** The return value if *expression-1* is not NULL.

Returns

The data type returned depends on the data type of *expression-2* and *expression-3*.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value -66.

```
SELECT IFNULL( NULL, -66 );
```

The following statement returns NULL, because the first expression is not NULL and there is no third expression.

```
SELECT IFNULL( -66, -66 );
```


INSERTSTR function [String]

Inserts a string into another string at a specified position.

Syntax

```
INSERTSTR( integer-expression, string-expression-1, string-expression-2 )
```

Parameters

- ***integer-expression*** The position after which the string is to be inserted. Use zero to insert a string at the beginning.
- ***string-expression-1*** The string into which the other string is to be inserted.
- ***string-expression-2*** The string to be inserted.

Returns

LONG VARCHAR

See also

- [“STUFF function \[String\]” on page 395](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value backoffice.

```
SELECT INSERTSTR( 0, 'office ', 'back' );
```

INTTOHEX function [Data type conversion]

Returns a string containing the hexadecimal equivalent of an integer.

Syntax

```
INTTOHEX( integer-expression )
```

Parameters

- ***integer-expression*** The integer to be converted to hexadecimal.

Returns

VARCHAR

Remarks

The CAST, CONVERT, HEXTOINT, and INTTOHEX functions can be used to convert to and from hexadecimal values. For more information, see “[Converting to and from hexadecimal values](#)” [*SQL Anywhere Server - SQL Reference*].

See also

- “[HEXTOINT function \[Data type conversion\]](#)” on page 341

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 0000009c.

```
SELECT INTTOHEX( 156 );
```

ISDATE function [Data type conversion]

Tests if a string argument can be converted to a date.

Syntax

```
ISDATE( string )
```

Parameters

- **string** The string to be analyzed to determine if the string represents a valid date.

Returns

INT

Remarks

If a conversion is possible, the function returns 1; otherwise, 0 is returned. If the argument is NULL, 0 is returned.

Standards and compatibility

- **SQL/2008** Vendor extension.

ISNULL function [Miscellaneous]

Returns the first non-NULL expression from a list. This function is identical to the COALESCE function.

Syntax

```
ISNULL( expression, expression [, ...] )
```

Parameters

- **expression** An expression to be tested against NULL.

At least two expressions must be passed into the function, and all expressions must be comparable.

Returns

The return type for this function depends on the expressions specified. That is, when the database server evaluates the function, it first searches for a data type in which all the expressions can be compared. When found, the database server compares the expressions and then returns the result in the type used for the comparison. If the database server cannot find a common comparison type, an error is returned.

See also

- [“COALESCE function \[Miscellaneous\]” on page 319](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value -66.

```
SELECT ISNULL( NULL , -66, 55, 45, NULL, 16 );
```

LCASE function [String]

Converts all characters in a string to lowercase.

Syntax

```
LCASE( string-expression )
```

Parameters

- **string-expression** The string to be converted to lowercase.

Returns

- CHAR
- LONG VARCHAR
- VARCHAR

Remarks

The LCASE function is identical to the LOWER function.

See also

- [“LOWER function \[String\]” on page 353](#)
- [“UCASE function \[String\]” on page 403](#)
- [“UPPER function \[String\]” on page 404](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension. The equivalent function LOWER is a core feature of the SQL/2008 standard.

Example

The following statement returns the value chocolate.

```
SELECT LCASE( 'ChoCOLatE' );
```

LEFT function [String]

Returns multiple characters from the beginning of a string.

Syntax

```
LEFT( string-expression, integer-expression )
```

Parameters

- ***string-expression*** The string.
- ***integer-expression*** The number of characters to return.

Returns

LONG VARCHAR

Remarks

If the string contains multibyte characters, and the proper collation is being used, the number of bytes returned may be greater than the specified number of characters.

You can specify an *integer-expression* that is larger than the value in the argument string expression. In this case, the entire value is returned.

Whenever possible, if the input string uses character-length semantics, the return value is described in character-length semantics.

See also

- [“RIGHT function \[String\]” on page 375](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the first 5 characters of each Surname value in the Customers table.

```
SELECT LEFT( Surname, 5) FROM Customers;
```

LENGTH function [String]

Returns the number of characters in the specified string.

Syntax

```
LENGTH( string-expression )
```

Parameters

- ***string-expression*** The string.

Returns

INT

Remarks

Use this function to determine the length of a string. For example, specify a column name for *string-expression* to determine the length of values in the column.

If the string contains multibyte characters, and the proper collation is being used, LENGTH returns the number of characters, not the number of bytes. If the string is of data type BINARY, the LENGTH function behaves as the BYTE_LENGTH function.

Note

You can use the LENGTH function and the CHAR_LENGTH function interchangeably for CHAR, VARCHAR, and LONG VARCHAR data types. However, you must use the LENGTH function for BINARY and bit array data types.

See also

- [“BYTE_LENGTH function \[String\]” on page 313](#)
- [“International languages and character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** The LENGTH function is a vendor extension; however, its semantics are identical to those of the CHAR_LENGTH function in the SQL/2008 standard.

Example

The following statement returns the value 9.

```
SELECT LENGTH( 'chocolate' );
```

LESSER function [Miscellaneous]

Returns the lesser of two parameter values.

Syntax

LESSER(*expression-1*, *expression-2*)

Parameters

- ***expression-1*** The first parameter value to be compared.
- ***expression-2*** The second parameter value to be compared.

Returns

The return type for this function depends on the expressions specified. That is, when the database server evaluates the function, it first searches for a data type in which all the expressions can be compared. When found, the database server compares the expressions and then returns the result in the type used for the comparison. If the database server cannot find a common comparison type, an error is returned.

Remarks

If the parameters are equal, the first value is returned.

See also

- [“GREATER function \[Miscellaneous\]” on page 340](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 5.

```
SELECT LESSER( 10, 5 ) FROM dummy;
```

LIST function [Aggregate]

Returns a delimited list of values for every row in a group.

Syntax

```
LIST(  
[ DISTINCT ] string-expression  
[ , delimiter-string ] )
```

Parameters

- ***string-expression*** A string expression, usually a column name. For each row in the group, the value of *string-expression* is added to the result string, with values separated by *delimiter-string*. When DISTINCT is specified, only unique *string-expression* values are added.
- ***delimiter-string*** A delimiter string for the list items. The default setting is a comma. There is no delimiter if a value of NULL or an empty string is supplied. The *delimiter-string* must be a constant.

Returns

LONG VARCHAR

Remarks

The LIST function returns the concatenation (with delimiters) of all the non-NULL values of X for each row in the group. If there does not exist at least one row in the group with a definite X-value, then LIST(X) returns the empty string.

NULL values and empty strings are ignored by the LIST function.

A LIST function cannot be used as a window function, but it can be used as input to a window function.

Standards and compatibility

- **SQL/2008** Vendor extension.

Examples

The following statement returns all street addresses from the Employees table.

```
SELECT LIST( Street ) FROM Employees;
```

LOCATE function [String]

Returns the position of one string within another.

Syntax

```
LOCATE( string-expression-1, string-expression-2 [, integer-expression ] )
```

Parameters

- ***string-expression-1*** The string to be searched.
- ***string-expression-2*** The string to be searched for.
- ***integer-expression*** The character position in the string to begin the search. The first character is position 1. If the starting offset is negative, the locate function returns the last matching string offset rather than the first. A negative offset indicates how much of the end of the string is to be excluded from the search. The number of bytes excluded is calculated as $(-1 * \text{offset}) - 1$.

Returns

INT

Remarks

If *integer-expression* is specified, the search starts at that offset into the string.

The first string can be a long string (longer than 255 bytes), but the second is limited to 255 bytes. If a long string is given as the second argument, the function returns a NULL value. If the string is not found, 0 is returned. Searching for a zero-length string will return 1. If any of the arguments are NULL, the result is NULL.

If multibyte characters are used, with the appropriate collation, then the starting position and the return value may be different from the *byte* positions.

See also

- [“UltraLite string functions” on page 305](#)
- [“CHARINDEX function \[String\]” on page 318](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 8.

```
SELECT LOCATE(  
    'office party this week - rsvp as soon as possible',  
    'party',  
    2 );
```

LOG function [Numeric]

Returns the natural logarithm of a number. Not supported by UltraLite Java edition databases.

Syntax

LOG(*numeric-expression*)

Parameters

- ***numeric-expression*** The number.

Returns

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

Remarks

The argument is an expression that returns the value of any built-in numeric data type.

See also

- [“LOG10 function \[Numeric\]” on page 353](#)

Standards and compatibility

- **SQL/2008** The SQL/2008 standard defines the natural logarithm function using the keyword LN. The natural logarithm function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The following statement returns the natural logarithm of 50.


```
SELECT LOG( 50 );
```

LOG10 function [Numeric]

Returns the base 10 logarithm of a number. Not supported by UltraLite Java edition databases.

Syntax

```
LOG10( numeric-expression )
```

Parameters

- ***numeric-expression*** The number.

Returns

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic. If the parameter is NULL, the result is NULL.

Remarks

The argument is an expression that returns the value of any built-in numeric data type.

See also

- [“LOG function \[Numeric\]” on page 352](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the base 10 logarithm for 50.

```
SELECT LOG10( 50 );
```

LOWER function [String]

Converts all characters in a string to lowercase. This function is identical to the LCASE function.

Syntax

```
LOWER( string-expression )
```

Parameters

- ***string-expression*** The string to be converted to lowercase.

Returns

CHAR, VARCHAR, or LONG VARCHAR corresponding to the data type of the argument.

Remarks

The LCASE function is identical to the LOWER function.

See also

- [“LCASE function \[String\]” on page 347](#)
- [“UCASE function \[String\]” on page 403](#)
- [“UPPER function \[String\]” on page 404](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** The LOWER function is a core feature of the SQL/2008 standard.

Example

The following statement returns the value chocolate.

```
SELECT LOWER( 'chOCOLate' );
```

LTRIM function [String]

Removes leading blanks from the string.

Syntax

```
LTRIM( string-expression )
```

Parameters

- ***string-expression*** The string to be trimmed.

Returns

- VARCHAR
- LONG VARCHAR

Remarks

The actual length of the result is the length of the expression minus the number of characters removed. If all the characters are removed, the result is an empty string.

If the parameter can be null, the result can be null.

If the parameter is null, the result is the null value.

See also

- [“RTRIM function \[String\]” on page 377](#)
- [“TRIM function \[String\]” on page 401](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

The TRIM specifications defined by the SQL/2008 standard (LEADING and TRAILING) are supplied by the SQL Anywhere LTRIM and RTRIM functions respectively.

Example

The following statement returns the value Test Message with all leading blanks removed.

```
SELECT LTRIM( '      Test Message' );
```

MAX function [Aggregate]

Returns the maximum *expression* value found in each group of rows.

Syntax 1

```
MAX( [ DISTINCT ] expression )
```

Parameters

- ***expression*** The expression for which the maximum value is to be calculated. This is commonly a column name.
- **DISTINCT *expression*** Returns the same as MAX(*expression*), and is included for completeness.

Returns

The same data type as the argument.

Remarks

Rows where *expression* is NULL are ignored. Returns NULL for a group containing no rows.

For simple comparisons of two expressions, you can also use the GREATER function.

See also

- [“MIN function \[Aggregate\]” on page 356](#)

Standards and compatibility

- **SQL/2008** Core feature.

See also

- [“Aggregate functions and outer references” \[SQL Anywhere Server - SQL Usage\]](#)

Example

The following statement returns the value 138948.000, representing the maximum salary in the Employees table.

```
SELECT MAX( Salary )
FROM Employees;
```

MIN function [Aggregate]

Returns the minimum expression value found in each group of rows.

Syntax 1

```
MIN( [ DISTINCT ] expression )
```

Parameters

- **expression** The expression for which the minimum value is to be calculated. This is commonly a column name.
- **DISTINCT expression** Returns the same as `MIN(expression)`, and is included for completeness.

Returns

The same data type as the argument.

Remarks

Rows where *expression* is NULL are ignored. Returns NULL for a group containing no rows.

For simple comparisons of two expressions, you can also use the LESSER function. See [“LESSER function \[Miscellaneous\]” \[SQL Anywhere Server - SQL Reference\]](#).

See also

- [“MAX function \[Aggregate\]” on page 355](#)

Standards and compatibility

- **SQL/2008** Core feature.

Example

The following statement returns the value 24903.000, representing the minimum salary in the Employees table.

```
SELECT MIN( Salary )
FROM Employees;
```

MINUTE function [Date and time]

Returns the minute component of a TIMESTAMP value.

Syntax

```
MINUTE( timestamp-expression )
```

Parameters

- ***timestamp-expression*** The `TIMESTAMP` value.

Returns

`SMALLINT`

Remarks

The value returned is the minute portion of the `TIMESTAMP` expression, a `SMALLINT` value between 0 and 59.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 22.

```
SELECT MINUTE( '1998-07-13 12:22:34' );
```

MINUTES function [Date and time]

The `MINUTES` function manipulates a `TIMESTAMP`, or returns the number of minutes between two `TIMESTAMP` values. See the `Remarks` section below.

Syntax 1

```
MINUTES( timestamp-expression )
```

Syntax 2

```
MINUTES( timestamp-expression, timestamp-expression )
```

Syntax 3

```
MINUTES( timestamp-or-time-expression, integer-expression )
```

Parameters

- ***timestamp-expression*** An expression of type `TIMESTAMP`.
- ***timestamp-or-time-expression*** An expression of type `TIME` or `TIMESTAMP`.
- ***integer-expression*** The number of minutes to be added to *timestamp-or-time-expression*. If *integer-expression* is negative, the appropriate number of minutes is subtracted from *timestamp-or-time-expression*.

Returns

`INTEGER` with Syntax 1 or Syntax 2.

`TIME` or `TIMESTAMP` with Syntax 3.

Remarks

The result of the MINUTES function depends on its arguments.

- **Syntax 1** If you pass a single *timestamp-expression* to the MINUTES function, it will return the number of minutes between midnight 0000-02-29 and *timestamp-expression* as an INTEGER.

Note

0000-02-29 is not meant to imply an actual date; it is the default date used by the MINUTES function.

- **Syntax 2** If you pass two **TIMESTAMP** values to the MINUTES function, the function returns the integer number of minutes between them.
- **Syntax 3** If you pass a **TIMESTAMP** value and an **INTEGER** value to the MINUTES function, the function returns the **TIMESTAMP** result of adding the integer number of minutes to *timestamp-expression* argument. Similarly, if the first argument to MINUTES is a **TIME** value, then the result is also a **TIME** value. Syntax 3 does not support implicit conversion of the first argument. It may be necessary to explicitly cast the first argument to a **DATE**, **TIME** or **TIMESTAMP** value. If the first argument is of type **DATE**, midnight is assumed for the time portion.

Since MINUTES returns an integer, overflow can occur when Syntax 1 is used with **TIMESTAMP** values greater than or equal to 4083-03-23 02:08:00.

Instead of Syntax 2, use the DATEDIFF function. Instead of Syntax 3, use the DATEADD function.

See also

- [“DATEDIFF function \[Date and time\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“DATEADD function \[Date and time\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“CAST function \[Data type conversion\]” on page 314](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements return the value 240, signifying that the second **TIMESTAMP** value is 240 minutes after the first. It is recommended that you use the second example (DATEDIFF).

```
SELECT MINUTES( '1999-07-13 06:07:12',  
               '1999-07-13 10:07:12' );  
  
SELECT DATEDIFF( minute,  
               '1999-07-13 06:07:12',  
               '1999-07-13 10:07:12' );
```

The following statement returns the value 1051040527.

```
SELECT MINUTES( '1998-07-13 06:07:12' );
```

The following statements return the **TIMESTAMP** value 1999-05-12 21:10:07.000. Note that the first statement requires an explicit cast of the literal string parameter. It is recommended that you use the second example (DATEADD).

```
SELECT MINUTES( CAST( '1999-05-12 21:05:07' AS TIMESTAMP ), 5);  
SELECT DATEADD( minute, 5, '1999-05-12 21:05:07' );
```

ML_GET_SERVER_NOTIFICATION function [System]

This function allows UltraLite users to use light weight polling to query a notifier on a MobiLink server for server-initiated sync requests. Not supported by UltraLite Java edition databases.

Syntax

```
ML_GET_SERVER_NOTIFICATION( notifier, address, key )
```

Parameters

- **notifier** The name of the notifier on the MobiLink server to poll.
- **address** The stream parameters, specified as:

```
tcpip{host=pc1;port=1234}
```
- **key** Optional. The notification request key.

Returns

Returns the subject and content of a notification request for the given request key.

Remarks

If there are no requests for the given request key, or if the notifier name does not exist on the MobiLink server, the result is NULL. If NULL is provided for the request key, then the remote ID of the user is used as the request key. If a request does exist, the resulting message is returned in the form: *[subject]content* (for example, *[sync]profile1*).

This function communicates over the network as it retrieves responses from the MobiLink server. As a result, this function may require a long execution time resulting from network latency. During execution, there may be periods when the function can execute in the background, allowing work to be performed in the runtime on other connections. These periods are not guaranteed however, and depend on the complexity of the SQL. The recommended method for users to retrieve a MobiLink address to use in this function is to use the `sync_profile_option_value` function with an existing synchronization profile to get the value for the **Stream** profile option. The value returned by this function call can be used directly as the MobiLink address parameter.

See also

- [“SYNC_PROFILE_OPTION_VALUE function \[System\]” on page 398](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

```
SELECT ML_GET_SERVER_NOTIFICATION('Notifier1',  
    'tcpip{host=sybase;port=1234}', 'MyKey');
```

MOD function [Numeric]

Returns the remainder when one whole number is divided by another.

Syntax

MOD(*dividend*, *divisor*)

Parameters

- **dividend** The dividend, or numerator of the division.
- **divisor** The divisor, or denominator of the division.

Returns

- SMALLINT
- INT
- NUMERIC

Remarks

Division involving a negative dividend gives a negative or zero result. The sign of the divisor has no effect.

See also

- [“REMAINDER function \[Numeric\]” on page 372](#)

Standards and compatibility

- **SQL/2008** The MOD function is part of optional SQL/2008 language feature T441.

Example

The following statement returns the value 2.

```
SELECT MOD( 5, 3 );
```

MONTH function [Date and time]

Returns the month of the given date.

Syntax

MONTH(*date-expression*)

Parameters

- ***date-expression*** A value of type DATE.

Returns

SMALLINT

Remarks

The value returned is a number between 1 and 12, corresponding to the month of the given date.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 7.

```
SELECT MONTH( '1998-07-13' );
```

MONTHNAME function [Date and time]

Returns the name of the month from a date.

Syntax

```
MONTHNAME( date-expression )
```

Parameters

- ***timestamp-expression*** A TIMESTAMP value.

Returns

VARCHAR

Remarks

The MONTHNAME function returns a string, even if the result is numeric, such as 2 for the month of February.

See also

- [“DATEPART function \[Date and time\]” on page 330](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value September.

```
SELECT MONTHNAME( '1998-09-05' );
```

MONTHS function [Date and time]

The MONTHS function manipulates a `TIMESTAMP`, or returns the number of months between two `TIMESTAMP` values. See the Remarks section below.

Syntax 1

`MONTHS(timestamp-expression)`

Syntax 2

`MONTHS(timestamp-expression, timestamp-expression)`

Syntax 3

`MONTHS(timestamp-expression, integer-expression)`

Parameters

- ***timestamp-expression*** A date and time of type `TIMESTAMP`.
- ***integer-expression*** The integer number of months (of type `SMALLINT`) to be added to the *timestamp-expression*. If *integer-expression* is negative, the appropriate number of months is subtracted from *timestamp-expression*. If you supply an *integer-expression*, the *timestamp-expression* must be explicitly cast as a `TIME`, `DATE` or `TIMESTAMP` data type. If *timestamp-expression* is a `TIME` value, the current month is assumed.

For information about casting data types, see “[CAST function \[Data type conversion\]](#)” on page 314.

Returns

`INTEGER` with Syntax 1 or Syntax 2.

`TIMESTAMP` with Syntax 3.

Remarks

The result of the MONTHS function depends on its arguments. The MONTHS function ignores hours, minutes, and seconds in its arguments.

- **Syntax 1** If you pass a single *timestamp-expression* to the MONTHS function, it will return the number of months between 0000-02 and *timestamp-expression* as an `INTEGER`.

Note

0000-02 is not meant to imply an actual date; it is the default date used by the MONTHS function.

- **Syntax 2** If you pass two `TIMESTAMP` values to the MONTHS function, the function returns the integer number of months between them.
- **Syntax 3** If you pass a `TIMESTAMP` value and a `SMALLINT` value to the MONTHS function, the function returns the `TIMESTAMP` result of adding the integer number of months to *timestamp-expression*.

Instead of Syntax 2, use the `DATEDIFF` function. Instead of Syntax 3, use the `DATEADD` function.

The value of MONTHS is calculated from the number of first days of the month between the two dates.

See also

- [“DATEDIFF function \[Date and time\]” on page 327](#)
- [“DATEADD function \[Date and time\]” on page 326](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements return the value 2, signifying that the second date is two months after the first. It is recommended that you use the second example (DATEDIFF).

```
SELECT MONTHS( '1999-07-13 06:07:12', '1999-09-13 10:07:12' );  
  
SELECT DATEDIFF( month,  
    '1999-07-13 06:07:12',  
    '1999-09-13 10:07:12' );
```

The following statement returns the value 23981.

```
SELECT MONTHS( '1998-07-13 06:07:12' );
```

The following statements return the TIMESTAMP value 1999-10-12 21:05:07.000. It is recommended that you use the second example (DATEADD).

```
SELECT MONTHS( CAST( '1999-05-12 21:05:07' AS DATETIME ), 5 );  
  
SELECT DATEADD( month, 5, '1999-05-12 21:05:07' );
```

NEWID function [Miscellaneous]

Generates a UUID (Universally Unique Identifier) value. A UUID is the same as a GUID (Globally Unique Identifier).

Syntax

```
NEWID( )
```

Parameters

There are no parameters associated with the NEWID function.

Returns

```
UNIQUEIDENTIFIER
```

Remarks

The NEWID function can be used in a DEFAULT clause for a column.

UUIDs can be used to uniquely identify rows in a table. A value produced on one computer does not match a value produced on another computer, so they can be used as keys in synchronization and replication environments.

See also

- [“The NEWID default” \[SQL Anywhere Server - SQL Usage\]](#)
- [“STRTOUUID function \[String\]” on page 394](#)
- [“UIDTOSTR function \[String\]” on page 404](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement creates a table named mytab with two columns. Column pk has a unique identifier data type, and assigns the NEWID function as the default value. Column c1 has an integer data type.

```
CREATE TABLE mytab(  
  pk UNIQUEIDENTIFIER PRIMARY KEY DEFAULT NEWID(),  
  c1 INT );
```

The following statement returns a unique identifier as a string:

```
SELECT UIDTOSTR( NEWID() );
```

For example, the value returned might be 96603324-6FF6-49DE-BF7D-F44C1C7E6856.

NOW function [Date and time]

Returns the current date and time as a **TIMESTAMP** value. The accuracy is limited by the accuracy of the system clock.

Syntax

```
NOW( [ * ] )
```

Returns

TIMESTAMP

Remarks

NOW is equivalent to the GETDATE function and the CURRENT_TIMESTAMP special value. NOW(*) and NOW() are equivalent constructions.

Each instance of the NOW function in a request is evaluated at most once. Multiple instances of NOW in the same request may or may not share the identical **TIMESTAMP** value.

See also

- [“GETDATE function \[Date and time\]” on page 339](#)
- [“CURRENT TIMESTAMP special value” on page 263](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the current date and time.

```
SELECT NOW( * );
```

NULLIF function [Miscellaneous]

Provides an abbreviated CASE expression by comparing expressions.

Syntax

```
NULLIF( expression-1, expression-2 )
```

Parameters

- ***expression-1*** An expression to be compared.
- ***expression-2*** An expression to be compared.

Returns

Data type of the first argument.

Remarks

NULLIF compares the values of the two expressions.

If the first expression equals the second expression, NULLIF returns NULL.

If the first expression does not equal the second expression, or if the second expression is NULL, NULLIF returns the first expression.

The NULLIF function provides a short way to write some CASE expressions.

See also

- [“CASE expressions” on page 283](#)

Standards and compatibility

- **SQL/2008** Core feature.

Example

The following statement returns the value a:

```
SELECT NULLIF( 'a', 'b' );
```

The following statement returns NULL.

```
SELECT NULLIF( 'a', 'a' );
```

PATINDEX function [String]

Returns an integer representing the starting position of the first occurrence of a pattern in a string.

Syntax

```
PATINDEX( '%pattern%', string-expression )
```

Parameters

- **pattern** The pattern to be searched for. If the leading percent wildcard is omitted, the PATINDEX function returns one (1) if the pattern occurs at the beginning of the string, and zero if it does not.

The pattern for UltraLite uses the following wildcards:

Wildcard	Matches
_ (underscore)	Any one character
% (percent)	Any string of zero or more characters
[]	Any single character in the specified range or set
[^]	Any single character not in the specified range or set

- **string-expression** The string to be searched for the pattern.

Returns

INT

Remarks

The PATINDEX function returns the starting position of the first occurrence of the pattern. If the pattern is not found, it returns zero (0).

See also

- [“LOCATE function \[String\]” on page 351](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 2.

```
SELECT PATINDEX( '%hoco%', 'chocolate' );
```

The following statement returns the value 11.

```
SELECT PATINDEX( '%4_5_', '0a1A 2a3A 4a5A' );
```

The following statement returns 14 which is the first non-alphanumeric character in the string expression. The pattern '%[^a-z0-9]%' can be used instead of '%[^a-zA-Z0-9]%' if the database is case insensitive.

```
SELECT PATINDEX( '%[^a-zA-Z0-9]%', 'SQLAnywhere12 has many new features' );
```

To get everything up to and including the first non-alphanumeric character in a string, you can use something like the following:

```
SELECT LEFT( @string, PATINDEX( '%[^a-zA-Z0-9]%', @string ) );
```

The following statements create a table, myTable, and populate it with various strings containing alphanumeric characters, spaces (blanks), and non-alphanumeric characters. Then, the SELECT statement and subsequent results show how you can use PATINDEX to find the starting position of spaces and non-alphanumeric characters in the strings:

```
CREATE TABLE myTable( col1 LONG VARCHAR );

INSERT INTO myTable (col1) VALUES( 'the quick brown fox jumped over the lazy dog' ),
( 'the quick brown fox $$$$ jumped over the lazy dog' ),
( 'the quick brown fox 0999 jumped over the lazy dog' ),
( 'the quick brown fox ** jumped over the lazy dog' ),
( 'thequickbrownfoxjumpedoverthelazydog' ),
( 'thequickbrownfoxjum999pedoverthelazydog' ),
( 'thequick$$$brownfox' ),
( 'the quick brown fox$$ jumped over the lazy dog' );

SELECT col1,
//position of first non-alphanumeric character or space:
PATINDEX( '%[^a-z0-9]%', col1) AS blank_posn,
//position of first non-alphanumeric char that isn't a space:
PATINDEX( '%[^ a-z0-9]%', col1) AS non_alpha_char,
//everything up to and including first non-alphanumeric char that isn't a
space:
LEFT ( col1, PATINDEX( '%[^ a-zA-Z0-9]%', col1) ) AS left_str,
//first non-alphanumeric char that isn't a space, and everything to the
right:
SUBSTRING ( col1, PATINDEX( '%[^ a-zA-Z0-9]%', col1) ) AS sub_str
FROM myTable;
```

col1	blank_posn	non_alpha_char	left_str	sub_str
the quick brown fox jumped over the lazy dog	4	0		the quick brown fox jumped over the lazy dog
the quick brown fox \$\$\$\$ jumped over the lazy dog	4	21	the quick brown fox \$	\$\$\$\$ jumped over the lazy dog

col1	blank_posn	non_alpha_char	left_str	sub_str
the quick brown fox 0999 jumped over the lazy dog	4	0		the quick brown fox 0999 jumped over the lazy dog
the quick brown fox ** jumped over the lazy dog	4	21	the quick brown fox *	** jumped over the lazy dog
thequickbrownfoxjumpe- doverthelazydog	0	0		thequickbrownfoxjum- pedoverthelazydog
thequickbrownfox- jum999pedoverthelazydog	0	0		thequickbrownfox- jum999pedoverthelazy- dog
thequick\$\$\$\$brownfox	9	9	thequick\$	\$\$\$\$brownfox
the quick brown fox\$\$ jumped over the lazy dog	4	20	the quick brown fox \$	\$\$ jumped over the lazy dog

PI function [Numeric]

Returns the numeric value PI.

Syntax

PI([*])

Returns

DOUBLE

Standards and compatibility

- **SQL/2008** Vendor extension.

Remarks

This function returns a DOUBLE value.

PI(*) and PI() are semantically equivalent.

Example

The following statement returns the value 3.141592653...

```
SELECT PI( * );
```


POWER function [Numeric]

Calculates one number raised to the power of another.

Syntax

```
POWER( numeric-expression-1, numeric-expression-2 )
```

Parameters

- ***numeric-expression-1*** The base.
- ***numeric-expression-2*** The exponent.

Returns

DOUBLE

Remarks

This function converts its arguments to DOUBLE, and performs the computation in double-precision floating-point arithmetic. If any argument is NULL, the result is a NULL value.

Standards and compatibility

- **SQL/2008** The POWER function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The following statement returns the value 64.

```
SELECT POWER( 2, 6 );
```

QUARTER function [Date and time]

Returns a number indicating the quarter of the year from the supplied TIMESTAMP expression.

Syntax

```
QUARTER( timestamp-expression )
```

Parameters

- ***timestamp-expression*** The date you want the quarter for.

Returns

INTEGER

Remarks

The quarters are as follows:

Quarter	Period (inclusive)
1	January 1 to March 31
2	April 1 to June 30
3	July 1 to September 30
4	October 1 to December 31

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 2.

```
SELECT QUARTER( '1987/05/02' );
```

RADIANS function [Numeric]

Converts a number from degrees to radians.

Syntax

```
RADIANS( numeric-expression )
```

Parameters

- ***numeric-expression*** A number, in degrees. This angle is converted to radians.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, and performs the computation in double-precision floating-point arithmetic.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns a value of approximately 0.5236.

```
SELECT RADIANS( 30 );
```

RAND function [Numeric]

Returns a random number in the interval 0 to 1, with an optional seed. Not supported by UltraLite Java edition databases.

Syntax

RAND([*integer-expression*])

Parameters

- ***integer-expression*** An optional seed used to create a random number. This argument allows you to create repeatable random number sequences.

Returns

DOUBLE

Remarks

The RAND function is a multiplicative linear congruential random number generator. See Park and Miller (1988), CACM 31(10), pp. 1192-1201 and Press et al. (1992), Numerical Recipes in C (2nd edition, Chapter 7, pp. 279). The result of calling the RAND function is a pseudo-random number n where $0 < n < 1$ (neither 0.0 nor 1.0 can be the result).

When a connection is made to the server, the random number generator seeds an initial value. Each connection is uniquely seeded so that it sees a different random sequence from other connections. You can also specify a seed value (*integer-expression*) as an argument. Normally, you should only do this once before requesting a sequence of random numbers through successive calls to the RAND function. If you initialize the seed value more than once, the sequence is restarted. If you specify the same seed value, the same sequence is generated. Seed values that are close in value generate similar initial sequences, with divergence further out in the sequence.

Never combine the sequence generated from one seed value with the sequence generated from a second seed value, in an attempt to obtain statistically random results. In other words, do not reset the seed value at any time during the generation of a sequence of random values.

The RAND function is treated as a non-deterministic function. The query optimizer does not cache the results of the RAND function.

For more information about non-deterministic functions, see [“Function caching” \[SQL Anywhere Server - SQL Usage\]](#).

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements produce eleven random results. Each subsequent call to the RAND function where a seed is not specified continues to produce different results:

```
SELECT RAND( 1 );  
SELECT RAND( ), RAND( ), RAND( ), RAND( ), RAND( );  
SELECT RAND( ), RAND( ), RAND( ), RAND( ), RAND( );
```

The following example produces two sets of results with identical sequences, since the seed value is specified twice:

```
SELECT RAND( 1 ), RAND( ), RAND( ), RAND( ), RAND( );  
SELECT RAND( 1 ), RAND( ), RAND( ), RAND( ), RAND( );
```

The following example produces five results that are near each other in value, and do not have a random distribution. For this reason, calling the RAND function more than once with similar seed values is not recommended:

```
SELECT RAND( 1 ), RAND( 2 ), RAND( 3 ), RAND( 4 ), RAND( 5 );
```

The following example produces five identical results, and should be avoided:

```
SELECT RAND( 1 ), RAND( 1 ), RAND( 1 ), RAND( 1 ), RAND( 1 );
```

REMAINDER function [Numeric]

Returns the remainder when one whole number is divided by another.

Syntax

REMAINDER(*dividend*, *divisor*)

Parameters

- **dividend** The dividend, or numerator of the division.
- **divisor** The divisor, or denominator of the division.

Returns

- INTEGER
- NUMERIC

Remarks

You can also use the MOD function to return the remainder.

See also

- [“MOD function \[Numeric\]” on page 360](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 2.

```
SELECT REMAINDER( 5, 3 );
```

REPEAT function [String]

Concatenates a string a specified number of times.

Syntax

```
REPEAT( string-expression, integer-expression )
```

Parameters

- ***string-expression*** The string to be repeated.
- ***integer-expression*** The number of times the string is to be repeated. If *integer-expression* is negative, an empty string is returned.

Returns

LONG VARCHAR

Remarks

If the actual length of the result string exceeds the maximum for the return type, an error occurs. The result is truncated to the maximum string size allowed.

The behavior of this function is identical to that of the REPLICATE function.

See also

- [“REPLICATE function \[String\]” on page 374](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value repeatrepeatrepeat.

```
SELECT REPEAT( 'repeat', 3 );
```

REPLACE function [String]

Replaces a string with another string, and returns the new results.

Syntax

```
REPLACE( original-string, search-string, replace-string )
```

Parameters

If any argument is NULL, the function returns NULL.

- ***original-string*** The string to be searched. This can be any length.

- ***search-string*** The string to be searched for and replaced with *replace-string*. This string is limited to 255 bytes. If *search-string* is an empty string, the original string is returned unchanged.
- ***replace-string*** The replacement string, which replaces *search-string*. This can be any length. If *replacement-string* is an empty string, all occurrences of *search-string* are deleted.

Returns

LONG VARCHAR

Remarks

This function replaces all occurrences.

Comparisons are case-sensitive on case-sensitive databases.

See also

- [“SUBSTRING function \[String\]” on page 395](#)
- [“CHARINDEX function \[String\]” on page 318](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value xx.def.xx.ghi.

```
SELECT REPLACE( 'abc.def.abc.ghi', 'abc', 'xx' );
```

The following statement generates a result set containing ALTER PROCEDURE statements which, when executed, would repair stored procedures that reference a table that has been renamed. (To be useful, the table name must be unique.)

```
SELECT REPLACE(
    REPLACE( proc_defn, 'OldTableName', 'NewTableName' ),
    'CREATE PROCEDURE',
    'ALTER PROCEDURE' )
FROM SYS.SYSPROCEDURE
WHERE proc_defn LIKE '%OldTableName%';
```

REPLICATE function [String]

Concatenates a string a specified number of times.

Syntax

```
REPLICATE( string-expression, integer-expression )
```

Parameters

- ***string-expression*** The string to be repeated.

- ***integer-expression*** The number of times the string is to be repeated.

Returns

LONG VARCHAR

Remarks

If the actual length of the result string exceeds the maximum for the return type, an error occurs. The result is truncated to the maximum string size allowed.

The behavior of this function is identical to that of the REPEAT function.

See also

- [“REPEAT function \[String\]” on page 373](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value repeatrepeatrepeat.

```
SELECT REPLICATE( 'repeat', 3 );
```

RIGHT function [String]

Returns the rightmost characters of a string.

Syntax

```
RIGHT( string-expression, integer-expression )
```

Parameters

- ***string-expression*** The string to return the rightmost characters for.
- ***integer-expression*** The number of characters at the end of the string to return.

Returns

LONG VARCHAR

Remarks

If the string contains multibyte characters, and the proper collation is being used, the number of bytes returned may be greater than the specified number of characters.

You can specify an *integer-expression* that is larger than the value in the column. In this case, the entire value is returned.

Whenever possible, if the input string uses character-length semantics, the return value is described in character-length semantics.

See also

- [“LEFT function \[String\]” on page 348](#)
- [“International languages and character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the last 5 characters of each Surname value in the Customers table.

```
SELECT RIGHT( Surname, 5) FROM Customers;
```

ROUND function [Numeric]

Rounds the *numeric-expression* to the specified *integer-expression* amount of places after the decimal point.

Syntax

ROUND(*numeric-expression*, *integer-expression*)

Parameters

- ***numeric-expression*** The number, passed into the function, to be rounded.
- ***integer-expression*** A positive integer specifies the number of significant digits to the right of the decimal point at which to round. A negative expression specifies the number of significant digits to the left of the decimal point at which to round.

Returns

NUMERIC

Remarks

The result of this function is either numeric or double. When there is a numeric result and the integer *integer-expression* is a negative value, the precision is increased by one.

See also

- [“TRUNCNUM function \[Numeric\]” on page 402](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 123.200.

```
SELECT ROUND( 123.234, 1 );
```

RTRIM function [String]

Removes trailing blanks from the string.

Syntax

```
RTRIM( string-expression )
```

Parameters

- ***string-expression*** The string to be trimmed.

Returns

- VARCHAR
- LONG VARCHAR

Remarks

The actual length of the result is the length of the expression minus the number of characters removed. If all the characters are removed, the result is an empty string.

If the argument is null, the result is the NULL value.

See also

- [“TRIM function \[String\]” on page 401](#)
- [“LTRIM function \[String\]” on page 354](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

The TRIM specifications defined by the SQL/2008 standard (LEADING and TRAILING) are supplied by the SQL Anywhere LTRIM and RTRIM functions respectively.

Example

The following statement returns the string Test Message, with all trailing blanks removed.

```
SELECT RTRIM( 'Test Message      ' );
```

SECOND function [Date and time]

Returns the seconds value of the TIMESTAMP argument.

Syntax**SECOND**(*timestamp-expression*)**Parameters**

- ***timestamp-expression*** The **TIMESTAMP** value.

Returns

SMALLINT

Remarks

Returns a number from 0 to 59 corresponding to the seconds component of the given **TIMESTAMP** argument value.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 25.

```
SELECT SECOND( '1998-07-13 21:21:25' );
```

SECONDS function [Date and time]

The **SECONDS** function manipulates a **TIMESTAMP**, or returns the number of seconds between two **TIMESTAMP** values. See the **Remarks** section below.

Syntax 1**SECONDS**(*timestamp-expression*)**Syntax 2****SECONDS**(*timestamp-expression*, *timestamp-expression*)**Syntax 3****SECONDS**(*time-or-timestamp-expression*, *integer-expression*)**Parameters**

- ***timestamp-expression*** A **TIMESTAMP** value.
- ***time-or-timestamp-expression*** A value of type **TIME** or **TIMESTAMP**.
- ***integer-expression*** The number of seconds to be added to the *time-or-timestamp-expression*. If *integer-expression* is negative, the appropriate number of seconds is subtracted from *time-or-timestamp-expression*. If you supply an integer expression, the *time-or-timestamp-expression* must be explicitly cast as a **TIME**, **DATE**, or **TIMESTAMP** data type. If *time-or-timestamp-expression* is a **DATE** type, its time portion is assumed to be midnight.

Returns

UNSIGNED BIGINT with Syntax 1.

SIGNED BIGINT with Syntax 2.

TIME or TIMESTAMP with Syntax 3.

Remarks

The result of the SECONDS function depends on its arguments.

- **Syntax 1** If you pass a single *timestamp-expression* to the SECONDS function, it will return the number of seconds between midnight 0000-02-29 and *timestamp-expression* as an UNSIGNED BIGINT.

Note

0000-02 is not meant to imply an actual date; it is the default date used by the SECONDS function.

- **Syntax 2** If you pass two **TIMESTAMP** values to the SECONDS function, the function returns the integer number of seconds between them as a SIGNED BIGINT value.
- **Syntax 3** If you pass a **TIMESTAMP** value and a **INTEGER** value to the SECONDS function, the function returns the **TIMESTAMP** result of adding the integer number of seconds to *time-or-timestamp-expression*. Similarly, if you pass a **TIME** value to the SECONDS function, the function returns a value of type **TIME**.

Instead of Syntax 2, use the DATEDIFF function. Instead of Syntax 3, use the DATEADD function.

See also

- [“CAST function \[Data type conversion\]” on page 314](#)
- [“DATEADD function \[Date and time\]” on page 326](#)
- [“DATEDIFF function \[Date and time\]” on page 327](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements return the value 14400, signifying that the second **TIMESTAMP** value is 14400 seconds after the first.

```
SELECT SECONDS( '1999-07-13 06:07:12',
               '1999-07-13 10:07:12' );
SELECT DATEDIFF( second,
               '1999-07-13 06:07:12',
               '1999-07-13 10:07:12' );
```

The following statement returns the value 63062431632.

```
SELECT SECONDS( '1998-07-13 06:07:12' );
```

The following statements return the **TIMESTAMP** value 1999-05-12 21:05:12.000.

```
SELECT SECONDS( CAST( '1999-05-12 21:05:07' AS TIMESTAMP ), 5);
SELECT DATEADD( second, 5, '1999-05-12 21:05:07' );
```

SHORT_PLAN function [Miscellaneous]

Returns a short description of the UltraLite plan optimization strategy of a SQL statement, as a string. The description is the same as that returned by the EXPLANATION function.

Syntax

SHORT_PLAN(*string-expression*)

Remarks

For some queries, the execution plan for UltraLite may differ from the plan selected for SQL Anywhere.

Parameters

- ***string-expression*** The SQL statement, which is commonly a SELECT statement, but can also be an UPDATE or DELETE statement.

Returns

LONG VARCHAR

See also

- [“EXPLANATION function \[Miscellaneous\]” on page 338](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement passes a SELECT statement as a string parameter and returns the plan for executing the query.

```
SELECT SHORT_PLAN(
  'SELECT * FROM Departments WHERE DepartmentID > 100' );
```

SIGN function [Numeric]

Returns the sign (positive or negative) of the given number.

Syntax

SIGN(*numeric-expression*)

Parameters

- ***numeric-expression*** The number for which the sign is to be returned. *numeric-expression* may be of type INTEGER, DOUBLE, or NUMERIC.

Returns

SMALLINT

Remarks

For negative numbers, the SIGN function returns -1.

For zero, the SIGN function returns 0.

For positive numbers, the SIGN function returns 1.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value -1

```
SELECT SIGN( -550 );
```

SIMILAR function [String]

Returns a number indicating the similarity between two strings.

Syntax

```
SIMILAR( string-expression-1, string-expression-2 )
```

Parameters

- ***string-expression-1*** The first string to be compared.
- ***string-expression-2*** The second string to be compared.

Returns

SMALL INT

Remarks

The function returns an integer between 0 and 100 representing the similarity between the two strings. The result can be interpreted as the percentage of characters matched between the two strings. A value of 100 indicates that the two strings are identical.

This function can be used to correct a list of names (such as customers). Some customers may have been added to the list more than once with slightly different names. You can use the SIMILAR function to find similar customer names by joining the customer table to itself, producing a report of all similarities greater than 90 percent, but less than 100 percent.

The calculation performed for the SIMILAR function is more complex than just the number of characters that match.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 75, indicating that the two values are 75% similar.

```
SELECT SIMILAR( 'toast', 'coast' );
```

SIN function [Numeric]

Returns the sine of a number.

Syntax

```
SIN( numeric-expression )
```

Parameters

- ***numeric-expression*** The angle, in radians.

Returns

DOUBLE

Remarks

The SIN function returns the sine of the argument, where the argument is an angle expressed in radians. The SIN and ASIN functions are inverse operations.

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result.

See also

- [“ASIN function \[Numeric\]” on page 309](#)
- [“COS function \[Numeric\]” on page 322](#)
- [“COT function \[Numeric\]” on page 323](#)
- [“TAN function \[Numeric\]” on page 399](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the SIN value of 0.52.

```
SELECT SIN( 0.52 );
```

SOUNDEX function [String]

Returns a number representing the sound of a string.

Syntax

SOUNDEX(*string-expression*)

Parameters

- ***string-expression*** The string to be evaluated.

Returns

SMALLINT

Remarks

The SOUNDEX function value for a string is based on the first letter and the next three consonants other than H, Y, and W. Vowels in *string-expression* are ignored unless they are the first letter of the string. Doubled letters are counted as one letter. For example, the word "apples" is based on the letters A, P, L, and S.

Multibyte characters are ignored by the SOUNDEX function.

Although it is not perfect, the SOUNDEX function normally returns the same number for words that sound similar and that start with the same letter.

The SOUNDEX function works best with English words. It is less useful for other languages.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns two identical numbers, 3827, representing the sound of each name.

```
SELECT SOUNDEX( 'Smith' ), SOUNDEX( 'Smythe' );
```

SPACE function [String]

Returns a specified number of spaces.

Syntax

SPACE(*integer-expression*)

Parameters

- ***integer-expression*** The number of spaces to return.

Returns

LONG VARCHAR

Remarks

If *integer-expression* is negative, a null string is returned.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns a string containing 10 spaces.

```
SELECT SPACE( 10 );
```

SQRT function [Numeric]

Returns the square root of a number.

Syntax

```
SQRT( numeric-expression )
```

Parameters

- ***numeric-expression*** The number for which the square root is to be calculated.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result.

Standards and compatibility

- **SQL/2008** The SQRT function comprises part of optional SQL/2008 language feature T621, "Enhanced numeric functions".

Example

The following statement returns the value 3.

```
SELECT SQRT( 9 );
```


ST_AsBinary function [Spatial]

The ST_AsBinary function returns a binary string representing the geometry. The output format is WKB as defined by OGC SFS 1.1. This format does not contain Z and M values.

Syntax

ST_AsBinary(*geometry-expression*)

Returns

- **BINARY** Returns the WKB representation of the *geometry-expression*.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.37

Example

The following returns the result 0x010100000000000000000000f03f0000000000000040.

```
select ST_AsBinary(ST_Point(1.0, 2.0, 4326))
```

The following returns the result 0x010100000000000000000000f03f0000000000000040. The server implicitly invokes the ST_AsBinary() function when converting geometries to BINARY.

```
select cast(ST_Point(1.0, 2.0, 4326) as binary(50))
```

ST_AsExtText function [Spatial]

The ST_AsExtText function returns a binary string representing the geometry. The output format is EWKT.

Syntax

ST_AsExtText(*geometry-expression*)

Returns

- **VARCHAR** Returns the EWKT representation of the *geometry-expression*.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.37

Example

The following returns the result SRID=4326;Point (1 2). The SRID is included in the result as a prefix.

```
select ST_AsExtText(ST_Point(1.0, 2.0, 4326))
```

The following returns the result SRID=4326;Point (1 2). The ST_AsExtText() function is implicitly invoked when converting geometries to VARCHAR types.

```
select cast(ST_Point(1.0, 2.0, 4326) as varchar(25))
```

ST_AsText function [Spatial]

The ST_AsText function returns a binary string representing the geometry. The output format is WKT as defined by OGC SFS 1.1.

Syntax

```
ST_AsText(geometry-expression)
```

Returns

- **VARCHAR** Returns the WKT representation of the *geometry-expression*.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.37

Example

The following returns the result `Point (1 2)`.

```
select ST_AsText(ST_Point(1.0, 2.0, 4326))
```

ST_Distance function [Spatial]

Returns the smallest distance between two specified geometry values. If the points are in SRID 4326, the units are in meters.

Syntax

```
ST_Distance(geo1,geo2 )
```

Parameters

Name	Type	Description
<i>geo1</i>	ST_Geometry	The first geometry value to be used to calculate the distance between two geometry values.
<i>geo2</i>	ST_Geometry	The second geometry value to be used to calculate the distance between two geometry values.

Returns

- **DOUBLE** Returns the smallest distance between the specified geometry values.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.23

Example

```
Select ST_Distance(ST_Point(-79.38,43.65,4326),ST_Point(-123.1,49.28,4326))
```

Returns 3367142.4632130372.

ST_Equals function [Spatial]

Tests whether an ST_Geometry value is spatially equal to another ST_Geometry value. Two geometry values can be considered equal if they have the same x and y coordinates and are in the same reference system.

The test may be limited by the resolution of the spatial reference system or the accuracy of the data.

The ST_Equals function defines the semantics used for comparison predicates (= and <>), IN list predicates, DISTINCT, and GROUP BY.

Syntax

```
ST_Equals( geo1,geo2 )
```

Parameters

Name	Type	Description
<i>geo1</i>	ST_Geometry	The first geometry value to be compared.
<i>geo2</i>	ST_Geometry	The second geometry value to be compared.

Returns

- **BIT** Returns 1 if the two geometry values are spatially equal, otherwise 0.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.24

Example

```
Select ST_Equals(ST_Point(1,1,4326),ST_Point(1,1,4326))
```

Returns 1.

ST_IntersectsRect function [Spatial]

The ST_IntersectsRect function tests if point is located within the box defined by the two points, min and max.

Syntax

```
ST_IntersectsRect(location,min,max)
```

Parameters

Name	Type	Description
<i>location</i>	ST_Geometry	The point to be tested.
<i>min</i>	ST_Geometry	The minimum point value used to define the box.
<i>max</i>	ST_Geometry	The maximum point value used to define the box.

Returns

- **BIT** Returns 1 if *location* intersects with the specified box, otherwise 0.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** Vendor extension

Example

```
Select ST_IntersectsRect(ST_Point(1,1,4326),ST_Point(0,0,4326),  
ST_Point(3,3,4326))
```

Returns 1.

ST_Point function [Spatial]

Constructs a point based on x and y coordinates.

Syntax

ST_Point(*x,y,SRID*)

Parameters

Name	Type	Description
<i>x</i>	DOUBLE	The x coordinate to use to construct the point.
<i>y</i>	DOUBLE	The y coordinate to use to construct the point.
<i>SRID</i>	INTEGER	The SRID value associated with the point.

Returns

- **ST_Point** Returns an ST_Geometry value created from the input string.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.8

Example

```
ST_Point(10.0,20.0,2163)
```

Creates a point (10.0,20.0) in the 2163 reference system.

ST_PointFromExtText function [Spatial]

Returns an ST_Geometry value, which is transformed from a VARCHAR value containing the EWKT representation of an ST_Geometry.

Syntax

```
ST_PointFromText(ewkt, srid)
```

Parameters

Name	Type	Description
<i>ewkt</i>	VARCHAR	The EWKT representation.
<i>srid</i>	INT	The spatial reference system identifier of the result is indicated by the srid parameter.

Returns

- **ST_Geometry** Returns an ST_Geometry value created from the input string.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.8

Example

```
ST_PointFromExtText('SRID=4326;Point(10 20)')
```

Creates a point (10,20) in the 4326 reference system.

ST_PointFromText function [Spatial]

Returns an ST_Geometry value, which is transformed from a VARCHAR value containing the WKT representation of an ST_Geometry.

Syntax

```
ST_PointFromText(wkt)
```

Parameters

Name	Type	Description
<i>wkt</i>	VARCHAR	The WKT representation.

Returns

- **ST_Geometry** Returns an ST_Geometry value created from the input string.

The spatial reference system identifier of the result is the given by parameter *srid*.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.8

Example

```
ST_PointFromText('Point(10 20)', 4326)
```

Creates a point (10, 20) in the 4326 reference system.

ST_PointFromWKB function [Spatial]

Returns an ST_Geometry value, which is transformed from a BINARY value containing the WKB representation of an ST_Geometry.

Syntax

```
ST_PointFromWKB(wkb, SRID)
```

Parameters

Name	Type	Description
<i>wkb</i>	BINARY	The WKB representation.
<i>SRID</i>	INTEGER	The SRID value associated with the point.

Returns

- **ST_Geometry** Returns an ST_Geometry value created from the input string.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.9

Example

The following returns point (1.0, 2.0, 4326).

```
ST_PointFromWKB(0x010100000000000000000000f03f0000000000000040, 4326)
```

ST_SRID function [Spatial]

Retrieves the spatial reference system (SRID) associated with the geometry value.

Syntax

```
ST_SRID( geo1, srid )
```

Parameters

Name	Type	Description
<i>geo1</i>	ST_Geometry	The point value.
<i>srid</i>	INTEGER	The SRID value associated with the point.

Returns

- **INT** Returns the SRID of the geometry.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 5.1.5

Example

The following returns 4326.

```
Select ST_SRID(ST_Point(10,20,4326))
```

ST_X function [Spatial]

Returns the x coordinate of the ST_Geometry value.

Syntax

```
ST_X(geo1)
```

Parameters

Name	Type	Description
<i>geo1</i>	ST_Geometry	The ST_Geometry value to determine the x value of.

Returns

- **DOUBLE** Returns the x coordinate of the ST_Geometry value.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.3

Example

The following example returns the result 10 . 0.

```
SELECT ST_X(ST_Point(10.0,20.0,4326))
```

ST_Y function [Spatial]

Returns the y coordinate of the ST_Geometry value.

Syntax

```
.ST_Y(geo1)
```

Parameters

Name	Type	Description
<i>geo1</i>	ST_Geometry	The ST_Geometry value to determine the y value of.

Returns

- **DOUBLE** Returns the y coordinate of the ST_Geometry value.

Standards and compatibility

- **SQL/MM (ISO/IEC 13249-3: 2006)** 6.1.3

Example

The following example returns the result 20 . 0.

```
SELECT ST_Y(ST_Point(10.0,20.0,4326))
```

STR function [String]

Returns the string equivalent of a number.

Syntax

```
STR( numeric-expression [, length [, decimal ] ] )
```

Parameters

- ***numeric-expression*** Any approximate numeric (float, real, or double precision) expression between -1E126 and 1E127.
- ***length*** The number of characters to be returned (including the decimal point, all digits to the right and left of the decimal point, and blanks). The default is 10.
- ***decimal*** The number of decimal digits to be returned. The default is 0.

Returns

VARCHAR

Remarks

If the integer portion of the number cannot fit in the length specified, then the result is a string of the specified length containing all asterisks. For example, the following statement returns ***.

```
SELECT STR( 1234.56, 3 );
```

Note

The maximum length that is supported is 128. Any length that is not between 1 and 128 yields a result of NULL.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns a string of six spaces followed by 1235, for a total of ten characters.

```
SELECT STR( 1234.56 );
```

The following statement returns the result 1234.6.

```
SELECT STR( 1234.56, 6, 1 );
```

STRING function [String]

Concatenates one or more strings into one large string.

Syntax

```
STRING( string-expression [, ... ] )
```

Parameters

- ***string-expression*** The string to be evaluated.

If only one argument is supplied, it is converted into a single expression. If multiple arguments are supplied, they are concatenated into a single string.

Returns

- LONG VARCHAR
- LONG BINARY

Remarks

Numeric or date parameters are converted to strings before concatenation. The `STRING` function can also be used to convert any single expression to a string by supplying that expression as the only parameter.

If all parameters are `NULL`, `STRING` returns `NULL`. If any parameters are non-`NULL`, then any `NULL` parameters are treated as empty strings.

See also

- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value `testing123`.

```
SELECT STRING( 'testing', NULL, 123 );
```

STRTOUUID function [String]

Converts a string value to a unique identifier (UUID or GUID) value.

Note

In databases created before version 9.0.2, the `STRTOUUID` and `UUIDTOSTR` functions were needed to convert between binary and string representations of UUID values.

In databases created using version 9.0.2 or later, the `UNIQUEIDENTIFIER` data type was changed to a native data type. You do not need to use `STRTOUUID` and `UUIDTOSTR` functions with these versions.

For more information, see [“Data types in UltraLite” on page 265](#).

Syntax

```
STRTOUUID( string-expression )
```

Parameters

- ***string-expression*** A string in the format `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`.

Returns

`UNIQUEIDENTIFIER`

Remarks

Converts a string in the format `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`, where `x` is a hexadecimal digit, to a unique identifier value.

This function is useful for inserting UUID values into a database.

See also

- [“UIDTOSTR function \[String\]” on page 404](#)
- [“NEWID function \[Miscellaneous\]” on page 363](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

STUFF function [String]

Deletes multiple characters from one string and replaces them with another string.

Syntax

STUFF(*string-expression-1*, *start*, *length*, *string-expression-2*)

Parameters

- ***string-expression-1*** The string to be modified by the STUFF function.
- ***start*** The character position at which to begin deleting characters. The first character in the string is position 1.
- ***length*** The number of characters to delete.
- ***string-expression-2*** The string to be inserted. To delete a portion of a string using the STUFF function, use a replacement string of NULL.

Returns

VARCHAR

See also

- [“INSERTSTR function \[String\]” on page 345](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value chocolate pie.

```
SELECT STUFF( 'chocolate cake', 11, 4, 'pie' );
```

SUBSTRING function [String]

Returns a substring of a string.

Syntax

```
{ SUBSTRING | SUBSTR } ( string-expression, start  
[, length ] )
```

Parameters

- ***string-expression*** The string from which a substring is to be returned.
- ***start*** The start position of the substring to return, in characters.
- ***length*** The length of the substring to return, in characters. If *length* is specified, the substring is restricted to that length.

Returns

- LONG BINARY
- LONG VARCHAR

Remarks

In UltraLite, the database does not have an `ansi_substring` option, but the SUBSTR function behaves as if `ansi_substring` is set to on by default. The function's behavior corresponds to ANSI/ISO SQL/2008 behavior:

- **Start value** The first character in the string is at position 1. A negative or zero start offset is treated as if the string were padded on the left with non-characters.
- **Length value** A positive *length* specifies that the substring ends *length* characters to the right of the starting position.

A negative *length* returns an error.

A *length* of zero returns an empty string.

If *string-expression* is of binary data type, the SUBSTRING function behaves as BYTE_SUBSTR.

To obtain characters at the end of a string, use the RIGHT function.

Whenever possible, if the input string uses character-length semantics, the return value is described in character-length semantics.

See also

- [“BYTE_SUBSTR function \[String\]” on page 313](#)
- [“LEFT function \[String\]” on page 348](#)
- [“RIGHT function \[String\]” on page 375](#)
- [“CHARINDEX function \[String\]” on page 318](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** SUBSTRING is a core feature of the SQL/2008 standard. The standard's implementation differs slightly from the UltraLite implementation: in the standard, SUBSTRING is

defined with three parameters using the keywords FROM and FOR, neither of which are required by UltraLite.

Example

The following table shows the values returned by the SUBSTRING function.

Example	Result
SUBSTRING('front yard', 1, 4)	fron
SUBSTRING('back yard', 6, 4)	yard
SUBSTR('abcdefgh', 0, -2)	Returns an error
SUBSTR('abcdefgh', -2, 2)	Returns an empty string

SUM function [Aggregate]

Returns the total of the specified expression for each group of rows.

Syntax 1

SUM([DISTINCT] *expression*)

Parameters

- ***expression*** The name of the expression to be summed. This is commonly a column name.
- **DISTINCT *expression*** Computes the sum of the unique values of *expression* within each group.

Returns

- INTEGER
- DOUBLE
- NUMERIC

Remarks

Rows where the specified expression is NULL are not included.

Returns NULL for a group containing no rows.

This function can generate an overflow error, resulting in an error being returned. You can use the CAST function on *numeric-expression* to avoid the overflow error. See [“CAST function \[Data type conversion\]” on page 314](#).

See also

- [“COUNT function \[Aggregate\]” on page 323](#)
- [“AVG function \[Aggregate\]” on page 312](#)

Standards and compatibility

- **SQL/2008** Core feature.

Example

The following statement returns the value 3749146.740.

```
SELECT SUM( Salary )
FROM Employees;
```

SWITCHOFFSET function [Date and time]

Returns a **TIMESTAMP WITH TIME ZONE** value that is converted from its original time zone offset to the specified time zone offset.

Syntax

```
SWITCHOFFSET( tmz-expression, time-zone-offset )
```

Parameters

- ***tmz-expression*** The **TIMESTAMP WITH TIME ZONE** value to be converted.
- ***time-zone-offset*** The time zone offset of the result. The value can be an integer representing the minutes before or after Coordinated Universal Time (UTC), a string in the form { + | - } hh:nn, or Z for the Zulu Time Zone. The Zulu Time Zone is the same time zone as UTC.

Returns

TIMESTAMP WITH TIME ZONE

See also

- [“TIMESTAMP WITH TIME ZONE data type” \[SQL Anywhere Server - SQL Reference\]](#)
- [“SYSDATETIMEOFFSET function \[Date and time\]” \[SQL Anywhere Server - SQL Reference\]](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following example changes a time zone offset value from -04:00 hours to -07:00 hours. The value returned is 2009-04-03 11:45:12.123-07:00.

```
SELECT CAST ( '2009-04-03 14:45:12.123-04:00' AS datetimeoffset ) AS EDT,
SWITCHOFFSET( EDT, '-07:00' ) AS PDT;
```

SYNC_PROFILE_OPTION_VALUE function [System]

Returns the value of the option corresponding to the given option name. Not supported by UltraLite Java edition databases.

Syntax

SYNC_PROFILE_OPTION_VALUE(*profile_name*, *option_name*)

Parameters

- **profile_name** The name of the sync profile to inspect.
- **option_name** The name of the option to retrieve the corresponding value for.

Returns

Returns the value of the option corresponding to the given option name.

Remarks

Option names with periods will retrieve values from a sublist with the given base option name before the period, and the given sublist option name after the period.

See also

- [“ML_GET_SERVER_NOTIFICATION function \[System\]” on page 359](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

Consider the profile:

```
MobiLinkId=joe;Stream=tcip{host=sybase;port=1234};Ping=1
```

- **MobiLinkId** joe
- **Stream** tcip{host=sybase;port=1234}
- **Stream.host** sybase
- **Stream.port** 1234
- **Ping** 1

TAN function [Numeric]

Returns the tangent of a number.

Syntax

TAN(*numeric-expression*)

Parameters

- **numeric-expression** An angle, in radians.

Returns

DOUBLE

Remarks

The ATAN and TAN functions are inverse operations.

This function converts its argument to DOUBLE, performs the computation in double-precision floating-point arithmetic, and returns a DOUBLE as the result.

See also

- [“COS function \[Numeric\]” on page 322](#)
- [“SIN function \[Numeric\]” on page 382](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value of the tan of 0.52.

```
SELECT TAN( 0.52 );
```

TODATETIMEOFFSET function [Date and time]

Converts a **TIMESTAMP** value to a **TIME STAMP WITH TIME ZONE** value using the specified time zone offset.

Syntax

```
TODATETIMEOFFSET( timestamp-expression, time-zone-offset )
```

Parameters

- ***timestamp-expression*** The **TIMESTAMP** expression to be converted.
- ***time-zone-offset*** The time zone offset. The value can be an **INTEGER** representing minutes before or after UTC, a **VARCHAR** string in the form of { + | - }hh:nn, or the string "Z" for the Zulu Time Zone. The Zulu Time Zone is the same time zone as UTC.

Returns

TIMESTAMP WITH TIME ZONE

See also

- [“TIMESTAMP WITH TIME ZONE data type” \[SQL Anywhere Server - SQL Reference\]](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following example converts a `TIMESTAMP` value to a `TIMESTAMP WITH TIME ZONE` value.

```
SELECT CAST('2009-04-03 14:45:12.123' AS TIMESTAMP) AS orig,  
       TODATETIMEOFFSET (orig, '+11:00');
```

TODAY function [Date and time]

Returns the current date as a `DATE` value.

Syntax

```
TODAY([ * ])
```

Returns

`DATE`

Remarks

`TODAY(*)` and `TODAY()` are semantically equivalent. `TODAY` is equivalent to the `CURRENT DATE` special value.

Each instance of the `TODAY` function in a request is evaluated at most once. Multiple instances of `TODAY` in the same request may or may not share the identical `DATE` value.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements return the current day according to the system clock.

```
SELECT TODAY( * );  
SELECT CURRENT DATE;
```

TRIM function [String]

Removes leading and trailing blanks from a string.

Syntax

```
TRIM( string-expression )
```

Parameters

- ***string-expression*** The string to be trimmed.

Returns

- `VARCHAR`
- `LONG VARCHAR`

See also

- [“LTRIM function \[String\]” on page 354](#)
- [“RTRIM function \[String\]” on page 377](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** The TRIM function is a SQL/2008 core feature.

SQL Anywhere does not support the additional parameters *trim specification* and *trim character*, as defined in SQL/2008. The SQL Anywhere implementation of TRIM corresponds to a TRIM specification of BOTH.

For the other TRIM specifications defined by the SQL/2008 standard (LEADING and TRAILING), SQL Anywhere supplies the LTRIM and RTRIM functions respectively.

Example

The following statement returns the value chocolate with no leading or trailing blanks.

```
SELECT TRIM( '  chocolate  ' );
```

TRUNCNUM function [Numeric]

Truncates a number at a specified number of places after the decimal point.

Syntax

```
{ TRUNCNUM | TRUNCATE }( numeric-expression, integer-expression )
```

Parameters

- ***numeric-expression*** The number to be truncated. This argument may be of type NUMERIC or DOUBLE.
- ***integer-expression*** A positive integer specifies the number of significant digits to the right of the decimal point at which to round. A negative value specifies the number of significant digits to the left of the decimal point at which to round.

Returns

NUMERIC or DOUBLE

Remarks

If any parameter is NULL, the result is NULL.

You should use the TRUNCNUM function, not the TRUNCATE function, when truncating numbers.

Use of the TRUNCATE function is not recommended because the word truncate is a keyword, and therefore requires you to either set the `quoted_identifier` option to OFF, or put quotes around the word TRUNCATE.

See also

- [“ROUND function \[Numeric\]” on page 376](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 600.

```
SELECT TRUNCNUM( 655, -2 );
```

The following statement: returns the value 655.340.

```
SELECT TRUNCNUM( 655.348, 2 );
```

UCASE function [String]

Converts all characters in a string to uppercase.

Syntax

```
UCASE( string-expression )
```

Parameters

- ***string-expression*** The string to be converted to uppercase.

Returns

CHAR, VARCHAR, or LONG VARCHAR corresponding to the data type of the argument.

Remarks

This function is identical to the UPPER function.

See also

- [“UPPER function \[String\]” on page 404](#)
- [“LCASE function \[String\]” on page 347](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension. The UPPER function is SQL/2008 compliant.

Example

The following statement returns the value CHOCOLATE.

```
SELECT UCASE( 'ChocoLate' );
```

UPPER function [String]

Converts all characters in a string to uppercase.

Syntax

UPPER(*string-expression*)

Parameters

- ***string-expression*** The string to be converted to uppercase.

Returns

CHAR, VARCHAR, or LONG VARCHAR corresponding to the data type of the argument.

Remarks

This function is identical to the UCASE function.

See also

- [“UCASE function \[String\]” on page 403](#)
- [“LCASE function \[String\]” on page 347](#)
- [“LOWER function \[String\]” on page 353](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** The UPPER function is a core feature of the SQL/2008 standard.

Example

The following statement returns the value CHOCOLATE.

```
SELECT UPPER( 'ChocoLate' );
```

UUIDTOSTR function [String]

Converts a unique identifier value (UUID, also known as GUID) to a string value.

Note

In databases created before version 9.0.2, the STRTOUUID and UUIDTOSTR functions were needed to convert between binary and string representations of UUID values.

In databases created using version 9.0.2 or later, the UNIQUEIDENTIFIER data type was changed to a native data type. You do not need to use STRTOUUID and UUIDTOSTR functions with these versions.

For more information, see [“Data types in UltraLite” on page 265](#).

Syntax

UUIDTOSTR(*uuid-expression*)

Parameters

- ***uuid-expression*** A unique identifier value.

Returns

VARCHAR

Remarks

Converts a unique identifier to a string value in the format *xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx*, where *x* is a hexadecimal digit. If the binary value is not a valid uniqueidentifier, NULL is returned.

This function is useful if you want to view a UUID value.

See also

- [“NEWID function \[Miscellaneous\]” on page 363](#)
- [“STRTOUUID function \[String\]” on page 394](#)
- [“UltraLite string functions” on page 305](#)

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement creates a table `mytab` with two columns. Column `pk` has a unique identifier data type, and column `c1` has an integer data type. It then inserts two rows with the values 1 and 2 respectively into column `c1`.

```
CREATE TABLE mytab(
    pk UNIQUEIDENTIFIER PRIMARY KEY DEFAULT NEWID(),
    c1 INT );
INSERT INTO mytab( c1 ) values ( 1 );
INSERT INTO mytab( c1 ) values ( 2 );
```

Executing the following SELECT statement returns all the data in the newly created table.

```
SELECT * FROM mytab;
```

You will see a two-column, two-row table. The value displayed for column `pk` will be binary values.

To convert the unique identifier values into a readable format, execute the following command:

```
SELECT UUIDTOSTR(pk), c1 FROM mytab;
```

The `UUIDTOSTR` function is not needed for databases created with version 9.0.2 or later.

WEEKS function [Date and time]

The `WEEKS` function manipulates a `TIMESTAMP`, or returns the number of weeks between two `TIMESTAMP` values. See the Remarks section below.

Syntax 1

WEEKS(*timestamp-expression*)

Syntax 2

WEEKS(*timestamp-expression*, *timestamp-expression*)

Syntax 3

WEEKS(*timestamp-expression*, *integer-expression*)

Parameters

- ***timestamp-expression*** A date and time value of type **TIMESTAMP**.
- ***integer-expression*** The number of weeks to be added to *timestamp-expression*. If *integer-expression* is negative, the appropriate number of weeks is subtracted from *timestamp-expression*. If you supply an *integer-expression*, *timestamp-expression* must be explicitly cast as a **DATE** or **TIMESTAMP**.

Returns

INTEGER with Syntax 1 or Syntax 2.

TIMESTAMP with Syntax 3.

Remarks

Given a single date (Syntax 1), the **WEEKS** function returns the number of weeks since 0000-02-29.

Given two dates (Syntax 2), the **WEEKS** function returns the number of weeks between them. The **WEEKS** function is similar to the **DATEDIFF** function, however the method used to calculate the number of weeks between two dates is not the same and can return a different result. The return value for **WEEKS** is determined by dividing the number of days between the two dates by seven, and then rounding down. However, **DATEDIFF** uses number of week boundaries in its computation. This can cause the values returned from the two functions to be different. For example, if the first date is a Friday and the second date is the following Monday, the **WEEKS** function returns a difference of 0, but the **DATEDIFF** function returns a difference of 1. While neither method is better than the other, you should consider the difference when choosing between **WEEKS** and **DATEDIFF**.

For more information about the **DATEDIFF** function, see [“DATEDIFF function \[Date and time\]” on page 327](#).

Given a date and an integer (Syntax 3), the **WEEKS** function adds the integer number of weeks to *timestamp-expression*. With Syntax 3, you must explicitly cast *timestamp-expression* as a **TIME**, **DATE**, or **TIMESTAMP** data type. If *timestamp-expression* is a **TIME** value, the current date is assumed. Instead of Syntax 3, use the **DATEADD** function.

For more information about the **DATEADD** function, see [“DATEADD function \[Date and time\]” on page 326](#).

See also

For information about casting data types, see [“CAST function \[Data type conversion\]” on page 314](#).

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 8, signifying that 2008-09-13 10:07:12 is eight weeks after 2008-07-13 06:07:12.

```
SELECT WEEKS( '2008-07-13 06:07:12', '2008-09-13 10:07:12' );
```

The following statement returns the value 104792, signifying that the date is 104792 weeks after 0000-02-29.

```
SELECT WEEKS( '2008-07-13 06:07:12' );
```

The following statement returns the **TIMESTAMP** value 2008-06-16 21:05:07.0, indicating the date and time five weeks after 2008-05-12 21:05:07.

```
SELECT WEEKS( CAST( '2008-05-12 21:05:07' AS TIMESTAMP ), 5 );
```

YEAR function [Date and time]

Returns the year component of the **TIMESTAMP** argument.

Syntax

```
YEAR( timestamp-expression )
```

Parameters

- ***timestamp-expression*** A **TIMESTAMP** value.

Returns

SMALLINT

Remarks

The value returned is the years component of the given **TIMESTAMP** value, returned as a **SMALLINT**.

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following example returns the value 2001.

```
SELECT YEAR( '2001-09-12' );
```

YEARS function [Date and time]

The YEARS function manipulates a TIMESTAMP, or returns the number of years between two TIMESTAMP values. See the Remarks section below.

Syntax 1

YEARS(*timestamp-expression*)

Syntax 2

YEARS(*timestamp-expression*, *timestamp-expression*)

Syntax 2

YEARS(*timestamp-expression*, *integer-expression*)

Parameters

- ***timestamp-expression*** A date and time value of type TIMESTAMP.
- ***integer-expression*** The number of years (as a SMALLINT value) to be added to *timestamp-expression*. If *integer-expression* is negative, the appropriate number of years is subtracted from *timestamp-expression*. If you supply an *integer-expression*, the *timestamp-expression* must be explicitly cast as a DATE, TIME, or TIMESTAMP value. If *timestamp-expression* is a TIME, the current year is assumed.

For information about casting data types, see [“CAST function \[Data type conversion\]”](#) on page 314.

Returns

SMALLINT with Syntax 1 or Syntax 2.

TIMESTAMP with Syntax 3.

Remarks

The value of YEARS is computed by counting the number of first days of the year between the two dates.

See also

- [“DATEDIFF function \[Date and time\]”](#) on page 327
- [“DATEADD function \[Date and time\]”](#) on page 326

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statements both return -4.

```
SELECT YEARS( '1998-07-13 06:07:12',  
             '1994-03-13 08:07:13' );
```



```
SELECT DATEDIFF( year,
  '1998-07-13 06:07:12',
  '1994-03-13 08:07:13' );
```

The following statements return 1998.

```
SELECT YEARS( '1998-07-13 06:07:12' )
SELECT DATEPART( year, '1998-07-13 06:07:12' );
```

The following statements return the given date advanced 300 years.

```
SELECT YEARS( CAST( '1998-07-13 06:07:12' AS TIMESTAMP ), 300 )
SELECT DATEADD( year, 300, '1998-07-13 06:07:12' );
```

YMD function [Date and time]

Returns a date value corresponding to the given year, month, and day of the month. Arguments are SMALLINT values from -32768 to 32767.

Syntax

```
YMD( smallint-expression1, smallint-expression2, smallint-expression3 )
```

Parameters

- ***smallint-expression1*** The year.
- ***smallint-expression2*** The number of the month. The year is adjusted if the month is outside the range 1-12.
- ***smallint-expression3*** The day number. The day can be any integer; the date is adjusted accordingly.

Returns

DATE

Standards and compatibility

- **SQL/2008** Vendor extension.

Example

The following statement returns the value 1998-06-12.

```
SELECT YMD( 1998, 06, 12 );
```

If the values are outside their normal range, the date is adjusted accordingly. For example, the following statement returns the DATE value 2000-03-01.

```
SELECT YMD( 1999, 15, 1 );
```

UltraLite SQL statements

The SQL statements supported by UltraLite SQL are a subset of the statements supported by SQL Anywhere databases.

Before you begin

- Tables in UltraLite do not support the concept of an owner. As a convenience for existing SQL and for SQL that is programmatically generated, UltraLite still allows the syntax *owner.table-name*. However, the owner is not checked because table owners are not supported in UltraLite.
- UltraLite SQL statement documentation follows the same syntax conventions used by SQL Anywhere statements. Ensure you understand these conventions and how they are used to represent SQL syntax.
- Using UltraLite SQL creates a transaction. A transaction consists of all changes (insert, update, and delete statements) since the last rollback or commit statement.

These changes can be made permanent by executing a COMMIT. A ROLLBACK statement causes the changes to be removed.
- Statements used by Interactive SQL have **[Interactive SQL]** after the statement name.

See also

- “SQL statements” [[SQL Anywhere Server - SQL Reference](#)]
- “Syntax conventions” [[SQL Anywhere Server - SQL Reference](#)]
- “UltraLite transaction processing” on page 470
- “COMMIT statement [UltraLite]” on page 420
- “ROLLBACK statement [UltraLite]” on page 443

UltraLite statement categories

SQL statements are organized and identified by the initial word in a statement, which is almost always a verb. This action-oriented syntax makes the nature of the language into a set of imperative statements (commands) to the database. In UltraLite, supported SQL statements can be classified as follows:

- **Data retrieval statements** Also known as queries. These statements allow select rows of data expressions from tables. Data retrieval is achieved with the SELECT statement.
- **Data manipulation statements** Allow you to change content in the database. Data manipulation is achieved with the following statements:
 - INSERT
 - UPDATE
 - DELETE

- **Data definition statements** Allow you to define the structure or schema of a database. The schema can be changed with the following statements:
 - ALTER DATABASE SCHEMA FROM FILE
 - CREATE INDEX
 - CREATE TABLE
 - DROP INDEX
 - DROP TABLE
 - ALTER TABLE
 - TRUNCATE TABLE
- **Transaction control statements** Allow you to control transactions within your UltraLite application. Transaction control is achieved with the following statements:
 - CHECKPOINT
 - COMMIT
 - ROLLBACK
- **Synchronization management** Allow you to temporarily control synchronization with a MobiLink server. Synchronization management is achieved with:
 - START SYNCHRONIZATION DELETE
 - STOP SYNCHRONIZATION DELETE
 - CREATE PUBLICATION
 - ALTER PUBLICATION
 - DROP PUBLICATION

See also

- [“Expressions in UltraLite” on page 280](#)
- [“Operators in UltraLite” on page 296](#)
- [“SELECT statement \[UltraLite\]” on page 444](#)
- [“INSERT statement \[UltraLite\]”](#)
- [“UPDATE statement \[UltraLite\]”](#)
- [“DELETE statement \[UltraLite\]”](#)
- [“ALTER DATABASE SCHEMA FROM FILE statement \[UltraLite\]”](#)
- [“CREATE INDEX statement \[UltraLite\]”](#)
- [“CREATE TABLE statement \[UltraLite\]”](#)
- [“DROP INDEX statement \[UltraLite\]”](#)
- [“DROP TABLE statement \[UltraLite\]”](#)
- [“ALTER TABLE statement \[UltraLite\]”](#)
- [“TRUNCATE TABLE statement \[UltraLite\]”](#)
- [“CHECKPOINT statement \[UltraLite\]”](#)
- [“COMMIT statement \[UltraLite\]”](#)
- [“ROLLBACK statement \[UltraLite\]”](#)
- [“START SYNCHRONIZATION DELETE statement \[UltraLite\]”](#)
- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]”](#)
- [“CREATE PUBLICATION statement \[UltraLite\]”](#)
- [“ALTER PUBLICATION statement \[UltraLite\]”](#)
- [“DROP PUBLICATION statement \[UltraLite\]”](#)

ALTER DATABASE SCHEMA FROM FILE statement [UltraLite]

Use this statement to modify the schema definition of an existing UltraLite database using a SQL script.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

```
ALTER DATABASE SCHEMA FROM FILE filename
```

Parameters

filename Defines the name and path to the SQL script used to upgrade the schema of an existing UltraLite database.

Remarks

Use either `ulinit` or `ulunload` to extract the DDL statements required for your script. By using these utilities, you ensure that the DDL statements are syntactically correct. Use `ulinit (-l logfile option)` or `ulunload (using the -n -s output-file options)`.

Backup the database before executing this statement.

The character set of the SQL script file must match the character set of the database you want to upgrade.

Ensure that your device is not reset while this statement is executing. If you reset the device during a schema upgrade, the UltraLite database becomes unusable.

Any rows that do not fit into the schema will be dropped (for instance if a uniqueness constraint is added and multiple rows contain the same values, all but one row will be dropped). In this case, the `SQLE_ROW_DROPPED_DURING_SCHEMA_UPGRADE` warning is generated. You can use this warning to detect the error and restore the database from the backup version.

See also

- [“UltraLite database schemas” on page 45](#)
- [“Deploying UltraLite database schema upgrades” on page 123](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 210](#)
- [“UltraLite Database Unload utility \(ulunload\)” on page 227](#)

Example

The following statement modifies the schema of the database using a SQL script, *MySchema.sql*:

```
ALTER DATABASE SCHEMA FROM FILE 'MySchema.sql';
```

ALTER PUBLICATION statement [UltraLite]

Use this statement to alter a publication. A publication identifies data in a remote database that is to be synchronized.

Syntax

```
ALTER PUBLICATION publication-name alterpub-clause
```

alterpub-clause :

```
  ADD TABLE table-name [ WHERE search-condition ]
| ALTER TABLE table-name [ WHERE search-condition ]
| { DROP | DELETE } TABLE table-name
| RENAME publication-name
```

Side effects

Automatic commit.

See also

- [“Search conditions in UltraLite” on page 287](#)
- [“UltraLite client synchronization design” on page 72](#)
- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)
- [“DROP PUBLICATION statement \[UltraLite\]” on page 434](#)
- [“START SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 446](#)
- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 447](#)

Example

The following ALTER PUBLICATION statement adds the Customers table to the pub_contact publication.

```
ALTER PUBLICATION pub_contact
  ADD TABLE Customers;
```

ALTER SYNCHRONIZATION PROFILE statement [UltraLite]

Use this statement to alter an UltraLite synchronization profile. Synchronization profiles define how an UltraLite database synchronizes with the MobiLink server.

Syntax

```
ALTER SYNCHRONIZATION PROFILE sync-profile-name
MERGE sync-option [; ... ]
```

sync-option :

```
sync-option-name = sync-option-value
```

sync-option-name : *string*

sync-option-value : *string*

Parameters

- **sync-profile-name** The name of the synchronization profile.
- **MERGE clause** Use this clause to change existing, or add new, options to a synchronization profile.
- **sync-option** A string of one or more synchronization option value pairs, separated by semicolons. For example, 'option1=value1;option2=value2'.
- **sync-option-name** The name of the synchronization profile option.
- **sync-option-value** The value for the synchronization profile option.

Remarks

You can use the MERGE clause to make changes to an existing synchronization profile. When using this clause, only the synchronization options that are specified in the MERGE clause are changed. To remove a synchronization option from a synchronization profile, the sync-option string should look like 'option1=;' (to set the option to an empty value).

The STREAM synchronization profile option is different from the other options because its value contains a sub-list. For example: 'STREAM=TCPIP{host=192.168.1.1;port=1234}'. In this case 'host=192.168.1.1;port=1234' is the sub-list. To add or remove a sub-list value, use a period between the STREAM sync-option-name and the sub-option-name. For example, MERGE 'stream.port=5678;stream.host=;compression=zlib' results in a synchronization profile of: stream=TCPIP{port=5678;compression=zlib}. Attempting to set the stream to a new value will replace the entire stream value. For example: MERGE 'stream=HTTPS' results in a synchronization profile of: stream=HTTPS{}

Side effects

None.

See also

- [“Synchronization profile options” on page 224](#)
- [“DROP SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 434](#)
- [“SYNCHRONIZE statement \[UltraLite\]” on page 448](#)

Example

The following is an example of the ALTER SYNCHRONIZATION PROFILE...REPLACE statement:

```
CREATE SYNCHRONIZATION PROFILE myProfile1;  
ALTER SYNCHRONIZATION PROFILE myProfile1  
  REPLACE 'publications=p1;uploadonly=on';
```

The following is an example of the ALTER SYNCHRONIZATION PROFILE...MERGE statement.

```
CREATE SYNCHRONIZATION PROFILE myProfile2 'publications=p1';  
ALTER SYNCHRONIZATION PROFILE myProfile2  
  MERGE 'publications=p2;uploadonly=on';
```

The following examples illustrate the changes that occur after executing a sequence of ALTER SYNCHRONIZATION PROFILE commands with different options.

Suppose myProfile1= 'MobiLinkUID=mary;ScriptVersion=default'.

After executing ALTER SYNCHRONIZATION PROFILE myProfile1 REPLACE 'MobiLinkPwd=sql;ScriptVersion=1', myProfile1 is 'MobiLinkPwd=sql;ScriptVersion=1'.

After executing ALTER SYNCHRONIZATION PROFILE myProfile1 MERGE 'MobiLinkUID=mary;STREAM=tcPIP', myProfile1 is 'MobiLinkPwd=sql;ScriptVersion=1;MobiLinkUID=mary;STREAM=tcPIP'.

After executing ALTER SYNCHRONIZATION PROFILE myProfile1 MERGE 'MobiLinkUID=;STREAM.host=192.168.1.1;STREAM.port=1234;ScriptVersion=;', myProfile1 is 'MobiLinkPwd=sql;STREAM=tcPIP{192.168.1.1;port=1234}'.

After executing ALTER SYNCHRONIZATION PROFILE myProfile1 MERGE 'MobiLinkPwd=;Ping=yes;STREAM =HTTP', myProfile1 is 'Ping=yes;STREAM=HTTP'.

After executing ALTER SYNCHRONIZATION PROFILE myProfile1 MERGE 'STREAM=HTTP{host=192.168.1.1}', myProfile1 is 'Ping=yes;STREAM=HTTP{host=192.168.1.1}'.

ALTER TABLE statement [UltraLite]

Use this statement to modify a table definition.

Syntax

```
ALTER TABLE table-name {
  add-clause
  | modify-clause
  | drop-clause
  | rename-clause
}
```

add-clause :
ADD { *column-definition* | *table-constraint* }

modify-clause :
ALTER *column-definition* | *sync-constraint*

drop-clause :
DROP { *column-name* | **CONSTRAINT** *constraint-name* }

rename-clause :
RENAME {
 new-table-name
 | [*old-column-name* **TO**] *new-column-name*
 | **CONSTRAINT** *old-constraint-name* **TO** *new-constraint-name* }

column-definition :
column-name data-type
[[**NOT**] **NULL**]
[**DEFAULT** *column-default*]
[**UNIQUE**]

column-default :
GLOBAL AUTOINCREMENT [(*number*)]
AUTOINCREMENT
CURRENT DATE
CURRENT TIME
CURRENT TIMESTAMP
NULL
NEWID()
constant-value

table-constraint :
[**CONSTRAINT** *constraint-name*]
{ *fkey-constraint* | *unique-key-constraint* }
[**WITH MAX HASH SIZE** *integer*]

fkey-constraint :
[**NOT NULL**] **FOREIGN KEY** [*role-name*] (*ordered-column-list*)
REFERENCES *table-name* (*column-name*, ...)
[**CHECK ON COMMIT**]

unique-key-constraint :
UNIQUE (*ordered-column-list*)

ordered-column-list :
(*column-name* [**ASC** | **DESC**], ...)

sync-constraint : **SYNCHRONIZE** { **ON** | **OFF** | **ALL** | **DOWNLOAD** }

Parameters

add-clause Adds a new column or table constraint to the table:

Adds a new column or table constraint to the table:

- **ADD *column-definition clause*** Adds a new column to the table. If the column has a default value, all rows in the new column are populated with that default value. For descriptions of the keywords and subclauses for this clause, see [“CREATE TABLE statement \[UltraLite\]” on page 425](#).
- **ADD *table-constraint clause*** Adds a constraint to the table. The optional constraint name allows you to modify or drop individual constraints at a later time, rather than having to modify the entire table constraint. For descriptions of the keywords and subclauses for this clause, see [“CREATE TABLE statement \[UltraLite\]” on page 425](#).

When adding a new unique constraint, all constraint columns must be non nullable. To add a unique constraint, alter the column to be NOT NULL.

Note

You cannot add a primary key in UltraLite.

modify-clause Change a single column definition. Note that you cannot use primary keys in the *column-definition* when part of an ALTER statement. If necessary, the data in the modified column is converted to the new data type. If a conversion error occurs, the operation will fail and the table is left unchanged. For a full explanation of *column-definition*, see [“CREATE TABLE statement \[UltraLite\]” on page 425](#).

drop-clause Delete a column or a table constraint:

- **DROP *column-name*** Delete the column from the table. If the column is contained in any index, uniqueness constraint, foreign key, or primary key, then the object must be deleted *before* UltraLite can delete the column.
- **DROP CONSTRAINT *table-constraint*** Delete the named constraint from the table definition. For a full explanation of *table-constraint*, see [“CREATE TABLE statement \[UltraLite\]” on page 425](#).

Note

You cannot drop a primary key in UltraLite.

rename-clause Change the name of a table, column, or constraint:

- **RENAME *new-table-name*** Change the name of the table to *new-table-name*. Note that any applications using the old table name must be modified. Foreign keys that were automatically assigned the old table name will not change names.
- **RENAME *old-column-name* TO *new-column-name*** Change the name of the column to the *new-column-name*. Note that any applications using the old column name will need to be modified.
- **RENAME *old-constraint-name* TO *new-constraint-name*** Change the name of the constraint to the *new-constraint-name*. Note that any applications using the old constraint name need to be modified.

Note

You cannot rename a primary key in UltraLite.

column-constraint A column constraint restricts the values the column can hold to ensure the integrity of data in the database. A column constraint can only be UNIQUE.

UNIQUE Identifies one or more columns that uniquely identify each row in the table. No two rows in the table can have the same values in all the named column(s). A table may have more than one unique constraint.

sync-constraint clause Specify a sync constraint to determine whether a table can be synchronized or not and whether all rows are uploaded, just changes to the table are uploaded, or no changes to the table are uploaded.

- **SYNCHRONIZE ON** Default setting - the table can be synchronized and only changes to the table are sent in the upload.
- **SYNCHRONIZE OFF** The table cannot be synchronized and it is an error to include the table in a publication.

- **SYNCHRONIZE ALL** The table can be synchronized and all rows in the table are sent in the upload. This constraint is not supported by UltraLite Java edition databases.
- **SYNCHRONIZE DOWNLOAD** The table can be synchronized with changes to the consolidated database but no local changes are uploaded.

Remarks

Only one *table-constraint* or *column-constraint* can be added, modified, or deleted in one ALTER TABLE statement.

The role name is the name of the foreign key. The main function of the *role-name* is to distinguish two foreign keys to the same table. Alternatively, you can name the foreign key with CONSTRAINT *constraint-name*. However, do not use both methods to name a foreign key.

You cannot MODIFY a table or column constraint. To change a constraint, you must DELETE the old constraint and ADD the new constraint.

ALTER TABLE cannot execute if a statement that affects the table is already being referenced by another request or query. Similarly, UltraLite does not process requests referencing the table while that table is being altered. Furthermore, you cannot execute ALTER TABLE when the database includes active queries or uncommitted transactions.

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time.

See also

- [“UltraLite database schemas” on page 45](#)
- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)
- [“DROP TABLE statement \[UltraLite\]” on page 435](#)
- [“Data types in UltraLite” on page 265](#)
- [“Table alteration” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Table and column constraints” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Overriding partition sizes for autoincremented columns” on page 71](#)
- [“Determine the most recently assigned GLOBAL AUTOINCREMENT value” on page 70](#)

Examples

The following statement drops the Street column from a fictitious table called MyEmployees.

```
ALTER TABLE MyEmployees
DROP Street;
```

The following example changes the Street column of the fictitious table, MyCustomers, to hold approximately 50 characters.

```
ALTER TABLE MyCustomers
ALTER Street CHAR(50);
```

ALTER USER statement [UltraLite]

Alters user settings.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax 1

```
ALTER USER user-name [ IDENTIFIED BY password ]
```

Parameters

user-name The name of the user.

IDENTIFIED BY clause The password for the user.

Remarks

- User IDs cannot:
 - begin with white space, single quotes, or double quotes
 - end with white space
 - contain semicolons
- Passwords are case-sensitive and they cannot:
 - begin with white space, single quotes, or double quotes
 - end with white space
 - contain semicolons
 - be longer than 255 bytes in length

Permissions

Any user can change their own password. All other changes require DBA authority.

Side effects

None.

See also

- [“CREATE USER statement \[UltraLite\]” on page 431](#)
- [“DROP USER statement \[UltraLite\]” on page 436](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following alters a user named SQLTester. The password is set to "welcome".

```
ALTER USER SQLTester IDENTIFIED BY welcome;
```

CHECKPOINT statement [UltraLite]

Use this statement to checkpoint the database.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

CHECKPOINT

Remarks

You can use the CHECKPOINT statement as a trigger for a commit flush. A commit flush writes uncommitted transactions to storage.

If you are using the embedded SQL API, you can also use the ULCheckpoint method. If you are writing a C++ component application, you can also use the Checkpoint method on a connection object. All other APIs must use this statement.

Side effects

While this statement flushes any pending committed transactions to storage, it does not commit or flush current transactions.

See also

- [“Flush single or grouped transactions” on page 471](#)
- [“COMMIT statement \[UltraLite\]” on page 420](#)
- [“UltraLite COMMIT_FLUSH connection parameter” on page 172](#)
- [“ULCheckpoint method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.Checkpoint method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)

Example

The following statement performs a checkpoint of the database:

```
CHECKPOINT ;
```

COMMIT statement [UltraLite]

Use this statement to make changes to the database permanent.

Syntax

COMMIT [WORK]

Remarks

Using UltraLite SQL creates a transaction. A transaction consists of all changes (INSERTs, UPDATEs, and DELETEs) since the last ROLLBACK or COMMIT. The COMMIT statement ends the current transaction and makes all changes made during the transaction permanent in the database.

Changes to the database objects using the ALTER, CREATE, and DROP statements are committed automatically.

See also

- [“CHECKPOINT statement \[UltraLite\]” on page 420](#)
- [“ROLLBACK statement \[UltraLite\]” on page 443](#)

Example

The following statement makes the changes in the current transaction permanent in the database:

```
COMMIT;
```

CREATE INDEX statement [UltraLite]

Use this statement to create an index on a specified table.

Syntax

```
CREATE [ UNIQUE ] INDEX [ IF NOT EXISTS ] [ index-name ]
ON table-name ( ordered-column-list )
[ WITH MAX HASH SIZE integer ]
```

```
ordered-column-list :
( column-name [ ASC | DESC ], ... )
```

Parameters

UNIQUE The UNIQUE attribute ensures that there will not be two rows in the table with identical values in all the columns in the index. Each index key must be unique or contain a NULL in at least one column.

There is a difference between a unique constraint on a table and a unique index. Columns of a unique index are allowed to be NULL, while columns in a unique constraint are not. A foreign key can reference either a primary key or a unique constraint, but not a unique index, because it can include multiple instances of NULL.

If the columns in a unique constraint are changed during an update, and a foreign key references that unique constraint, any rows no longer referencing rows in the unique constraint are deleted from the remote.

IF NOT EXISTS clause When the IF NOT EXISTS attribute is specified and the named index already exists, no changes are made and an error is not returned.

ordered-column-list An ordered list of columns. Column values in the index can be sorted in ascending or descending order.

WITH MAX HASH SIZE Sets the hash size (in bytes) for this index. This value overrides the default MaxHashSize property in effect for the database. This is not supported by UltraLite Java edition databases. To learn the default size, see [“Reading database properties” on page 33](#).

Remarks

UltraLite automatically creates indexes for primary keys and for unique constraints.

Indexes can improve query performance by providing quick ways for UltraLite to look up specific rows. Conversely, because they have to be maintained, indexes may slow down synchronization and INSERT, DELETE, and UPDATE statements.

Indexes are automatically used to improve the performance of queries issued to the database, and to sort queries with an ORDER BY clause. Once an index is created, it is never referenced in a SQL statement again except to remove it with DROP INDEX.

Indexes use space in the database. Also, the additional work required to maintain indexes can affect the performance of data modification operations. For these reasons, you should avoid creating indexes that do not improve query performance.

UltraLite does not process requests or queries referencing the index while the CREATE INDEX statement is being processed. Furthermore, you cannot execute CREATE INDEX when the database includes active queries or uncommitted transactions.

UltraLite can provide developers with execution plans to optimize queries. See [“Execution plans in UltraLite” on page 462](#).

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“UltraLite database schemas” on page 45](#).

Side effects

- Automatic commit.

See also

- [“UltraLite performance tips” on page 455](#)
- [“DROP INDEX statement \[UltraLite\]” on page 433](#)
- [“UltraLite max_hash_size creation parameter” on page 149](#)
- [“UltraLite indexes” on page 54](#)

Example

The following statement creates a two-column index on the Employees table.

```
CREATE INDEX employee_name_index
ON Employees ( Surname, GivenName );
```

The following statement creates an index on the SalesOrderItems table for the ProductID column.

```
CREATE INDEX item_prod
ON SalesOrderItems ( ProductID );
```

CREATE PUBLICATION statement [UltraLite]

Use this statement to create a publication. A publication identifies synchronized data in an UltraLite remote database.

Syntax

```
CREATE PUBLICATION [ IF NOT EXISTS ] publication-name
( TABLE table-name [ WHERE search-condition ], ... )
```

Parameters

- **IF NOT EXISTS clause** When the IF NOT EXISTS clause is specified and the named publication already exists, no changes are made and an error is not returned.
- **TABLE clause** Use the table to include a TABLE in the publication. There is no limit to the number of TABLE clauses.
- **WHERE clause** If a WHERE clause is specified, only rows satisfying *search-condition* are considered for upload from the associated table during synchronization.

If you do not specify a WHERE clause, every row in the table that has changed in UltraLite since the last synchronization is considered for upload.

Remarks

A publication establishes tables that are synchronized during a single synchronization operation, and determines which data is uploaded to the MobiLink server. The MobiLink server may send back rows for these (and only these) tables during its download session; however, rows that are downloaded do not have to satisfy the WHERE clause for a table.

Only entire tables can be published. You cannot publish specific columns of a table in UltraLite.

Side effects

- Automatic commit.

See also

- [“Search conditions in UltraLite” on page 287](#)
- [“UltraLite clients” on page 67](#)
- [“DROP PUBLICATION statement \[UltraLite\]” on page 434](#)
- [“ALTER PUBLICATION statement \[UltraLite\]” on page 413](#)
- [“Search conditions in UltraLite” on page 287](#)

Example

The following statement publishes all the columns and rows of two tables.

```
CREATE PUBLICATION pub_contact (
  TABLE Contacts,
  TABLE Customers
);
```

The following statement publishes only the rows of the Customers table where the State column contains MN.

```
CREATE PUBLICATION pub_customer (  
  TABLE Customers  
  WHERE State = 'MN'  
);
```

CREATE SYNCHRONIZATION PROFILE statement [UltraLite]

Use this statement to create or replace an UltraLite synchronization profile. Synchronization profiles define how an UltraLite database synchronizes with the MobiLink server.

Syntax

```
CREATE [OR REPLACE] SYNCHRONIZATION PROFILE sync-profile-name sync-option [;...]
```

sync-option :

sync-option-name = *sync-option-value*

sync-option-name : *string*

sync-option-value : *string*

Parameters

- **OR REPLACE clause** If the named synchronization profile already exists, then it will be replaced. If the profile does not exist, it will be created.
- **sync-profile-name** The name of the synchronization profile.
- **sync-option** A string of one or more synchronization option value pairs, separated by semicolons. For example, 'option1=value1;option2=value2'.
- **sync-option-name** The name of the synchronization profile option.
- **sync-option-value** The value for the synchronization profile option.

Remarks

You can use the REPLACE clause to make changes to an existing synchronization profile. This clause replaces the contents of the synchronization profile with whatever is contained in the new sync-option string. This approach is the same as dropping the synchronization profile and then creating one with the same name but using the new string. Therefore, a synchronization profile does not need to contain a full synchronization definition because parameters can be merged in or overridden at synchronization time.

The STREAM synchronization profile option is different from the other options because its value contains a sub-list. For example: 'STREAM=TCPIP{host=192.168.1.1;port=1234}'. In this case 'host=192.168.1.1;port=1234' is the sub-list. To add or remove a sub-list value, use a period between the STREAM sync-option-name and the sub-option-name. For example, MERGE

'stream.port=5678;stream.host=;compression=zlib' results in a synchronization profile of: stream=TCPIP{port=5678;compression=zlib}. Attempting to set the stream to a new value will replace the entire stream value. For example: MERGE 'stream=HTTPS' results in a synchronization profile of: stream=HTTPS{}

Side effects

None.

See also

- “Synchronization profile options” on page 224
- “ALTER SYNCHRONIZATION PROFILE statement [UltraLite]” on page 413
- “DROP SYNCHRONIZATION PROFILE statement [UltraLite]” on page 434
- “SYNCHRONIZE statement [UltraLite]” on page 448

Example

The following creates a synchronization profile called Test1.

```
CREATE SYNCHRONIZATION PROFILE Test1
'MobiLinkId=mary;Stream=TCPIP{host=192.168.1.1;port=1234}'
```

CREATE TABLE statement [UltraLite]

Use this statement to create a table.

Syntax

```
CREATE TABLE [ IF NOT EXISTS ] table-name (
  { column-definition | table-constraint | sync-constraint }, ...
)
```

column-definition :

```
column-name data-type
[ [ NOT ] NULL ]
[ DEFAULT column-default ]
[ STORE AS FILE (file-name-column) [ CASCADE DELETE ]
[ column-constraint ]
```

column-default :

```
AUTOFILENAME(prefix,extension)
| GLOBAL AUTOINCREMENT [ ( number ) ]
| AUTOINCREMENT
| CURRENT DATE
| CURRENT TIME
| CURRENT TIMESTAMP
| CURRENT UTC TIMESTAMP
| NULL
| NEWID( )
| constant-value
```

file-name
"filename"

column-constraint :

PRIMARY KEY
| UNIQUE

table-constraint :

{ [**CONSTRAINT** *constraint-name*]
 pkey-constraint
 | *fkey-constraint*
 | *unique-key-constraint* }
[**WITH MAX HASH SIZE** *integer*]

pkey-constraint :

PRIMARY KEY [*ordered-column-list*]

fkey-constraint :

[**NOT NULL**] **FOREIGN KEY** [*role-name*] (*ordered-column-list*)
REFERENCES *table-name* (*column-name*, ...)
[**CHECK ON COMMIT**]

unique-key-constraint :

UNIQUE (*ordered-column-list*)

ordered-column-list :

(*column-name* [**ASC** | **DESC**], ...)

sync-constraint : **SYNCHRONIZE** { **ON** | **OFF** | **ALL** | **DOWNLOAD** }

Parameters

IF NOT EXISTS clause Use this clause to create a table. No changes are made if the named table already exists, and an error is not returned.

column-definition Defines a column in a table. Available parameters for this clause include:

- **column-name** The column name is an identifier. Two columns in the same table cannot have the same name. See [“Identifiers in UltraLite” on page 259](#).

UltraLite Java edition databases support the partitioning of database files such that external files may now be used to store large BLOB values, with the files referenced by two columns: the first stores the file name and should have a data type of CHAR(size) ... AUTOFILENAME(...), and the second is used to access the file contents and is declared as a LONG BINARY STORE AS ... data type. The file contents column is read-only.

- **data-type** The data type of the column. See [“Data types in UltraLite” on page 265](#).
- **[NOT] NULL** If NOT NULL is specified, or if the column is in a PRIMARY KEY or UNIQUE constraint, the column cannot contain NULL in any row. Otherwise, NULL is allowed.
- **column-default** Sets the default value for the column. If a DEFAULT value is specified, it is used as the value for the column in any INSERT statement that does not specify a value for the column. If no DEFAULT is specified, it is equivalent to DEFAULT NULL. Default options include:
 - **AUTOFILENAME** This clause supports the storing of external BLOB files in a partitioned UltraLite Java edition database.

When partitioning the database, the column designated to store the file names requires the `AUTOFILENAME(prefix,extension)` clause. This clause specifies how new filenames are to be generated for downloaded BLOB values. The prefix and extension values are string literal constants. See [“Connection.OPTION_BLOB_FILE_BASE_DIR variable \[BlackBerry\] \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#).

- **AUTOINCREMENT** When using `AUTOINCREMENT`, the column must be one of the integer data types, or an exact numeric type. On inserts into the table, if a value is not specified for the `AUTOINCREMENT` column, a unique value larger than any other value in the column is generated. If an `INSERT` specifies a value for the column that is larger than the current maximum value for the column, that value is used as a starting point for subsequent inserts.

Tip

In UltraLite, the autoincrement value is not set to 0 when the table is created, and `AUTOINCREMENT` generates negative numbers when a signed data type is used for the column. Therefore, declare `AUTOINCREMENT` columns as unsigned integers to prevent negative values from being used.

- **GLOBAL AUTOINCREMENT** Similar to `AUTOINCREMENT`, except that the domain is partitioned. Each partition contains the same number of values. You assign each copy of the database a unique global database identification number. UltraLite supplies default values in a database only from the partition uniquely identified by that database's number.

Tip

If the column is of type `BIGINT` or `UNSIGNED BIGINT`, the default partition size is $2^{32} = 4294967296$; for columns of all other types, the default partition size is $2^{16} = 65536$. Since these defaults may be inappropriate, especially if your column is not of type `INT` or `BIGINT`, it is best to specify the partition size explicitly.

See [“Using GLOBAL AUTOINCREMENT in UltraLite” on page 69](#), and [“UltraLite global_database_id option” on page 195](#).

- **[NOT] NULL** Controls whether the column can contain NULLs.
- **NEWID()** A function that generates a unique identifier value. See [“NEWID function \[Miscellaneous\]” on page 363](#).
- **CURRENT TIMESTAMP** Combines `CURRENT DATE` and `CURRENT TIME` to form a `TIMESTAMP` value containing the year, month, day, hour, minute, second, and fraction of a second. The fraction of a second is stored to 3 decimal places. The accuracy is limited by the accuracy of the system clock. See [“CURRENT TIMESTAMP special value” on page 263](#).
- **| CURRENT UTC TIMESTAMP** A `TIMESTAMP WITH TIMEZONE` value containing the Coordinated Universal Time (UTC) comprised of the year, month, day, hour, minute, second, fraction of a second, and time zone. The fraction of a second is stored to 3 decimal places. The accuracy is limited by the accuracy of the system clock. See [“CURRENT UTC TIMESTAMP special value” on page 263](#).

- **CURRENT DATE** Stores the current year, month, and day. See “[CURRENT DATE special value](#)” on page 262.
- **CURRENT TIME** Stores the current hour, minute, second and fraction of a second. See “[CURRENT TIME special value](#)” on page 262.
- **constant-value** A constant for the data type of the column. Typically the constant is a number or a string.
- **STORE AS FILE (*file-name-column*) [CASCADE DELETE]** Supported by UltraLite Java edition databases only.

Specify that a LONG BINARY column is to be stored externally (partitioning the database) and specify the **file-name-column** to name the column that will be used to store the file names of the externally stored BLOB values. A column with this clause must be of type LONG BINARY and behave as a read-only column.

- **column-constraint clause** Specify a column constraint to restrict the values allowed in a column. A column constraint can be one of:
 - **PRIMARY KEY** When set as part of a *column-constraint*, the PRIMARY KEY clause sets the column as the primary key for the table. Primary keys uniquely identify each row in a table. By default, columns included in primary keys do not allow NULL.
 - **UNIQUE** Identifies one or more columns that uniquely identify each row in the table. No two rows in the table can have the same values in all the named column(s). A table may have more than one unique constraint. NULL values are not allowed.

table-constraint clause Specify a table constraint to restrict the values that one or more columns in the table can hold. Use the CONSTRAINT clause to specify an identifier for the table constraint. Table constraints can be in the form of a primary key constraint, a foreign key constraint, or a unique constraint, as defined below:

- **pkey-constraint clause** Sets the specified column(s) as the primary key for the table. Primary keys uniquely identify each row in a table. Columns included in primary keys cannot allow NULLs.
- **fkey-constraint clause** Specify a foreign key constraint to restrict values for one or more columns that must match the values in a primary key (or a unique constraint) of another table.
 - **NOT NULL clause** Specify NOT NULL to disallow NULLs in the foreign key columns. A NULL in a foreign key means that no row in the primary table corresponds to this row in the foreign table. If at least one value in a multi-column foreign key is NULL, there is no restriction on the values that can be held in other columns of the key.
 - **role-name clause** Specify a *role-name* to name the foreign key. *role-name* is used to distinguish foreign keys within the same table. Alternatively, you can name the foreign key using CONSTRAINT *constraint-name*. However, do not use both methods to name a foreign key.
 - **REFERENCES clause** Specify the REFERENCES clause to define one or more columns in the primary table to use as the foreign key constraint. Any *column-name* you specify in a

REFERENCES column constraint must be a column in the primary table, and must be subject to a unique constraint or primary key constraint.

- **CHECK ON COMMIT** not supported by UltraLite Java edition databases. Specify CHECK ON COMMIT to cause the database server to wait for a COMMIT before enforcing foreign key constraints. By default, foreign key constraints are enforced immediately during insert, update, or delete operations. However, when CHECK ON COMMIT is set, database changes can be made in any order, even if they violate foreign key constraints, if inconsistent data is resolved before the next COMMIT.
 - **unique-key-constraint clause** Specify a unique constraint to identify one or more columns that uniquely identify each row in the table. No two rows in the table can have the same values in all the named column(s). A table may have more than one unique constraint.
 - **WITH MAX HASH SIZE** Sets the hash size (in bytes) for this index. This value overrides the default MaxHashSize property in effect for the database. This clause is not supported for UltraLite Java edition databases. To learn the default size, see [“Reading database properties” on page 33](#). Also see [“UltraLite max_hash_size creation parameter” on page 149](#).
- sync-constraint clause** Specify a sync constraint to determine whether a table can be synchronized or not and whether all rows are uploaded, just changes to the table are uploaded, or no changes to the table are uploaded.
- **SYNCHRONIZE ON** Default setting - the table can be synchronized and only changes to the table are sent in the upload.
 - **SYNCHRONIZE OFF** The table cannot be synchronized and it is an error to include the table in a publication.
 - **SYNCHRONIZE ALL** The table can be synchronized and all rows in the table are sent in the upload. This constraint is not supported by UltraLite Java edition databases.
 - **SYNCHRONIZE DOWNLOAD** The table can be synchronized with changes to the consolidated database but no local changes are uploaded.

Remarks

Column constraints are normally used unless the constraint references more than one column in the table. In these cases, a table constraint must be used. If a statement causes a violation of a constraint, execution of the statement does not complete. Any changes made by the statement before error detection are undone, and an error is reported.

Each row in the table has a unique primary key value.

If no role name is specified, the role name is assigned as follows:

1. If there is no foreign key with a role name the same as the table name, the table name is assigned as the role name.
2. If the table name is already taken, the role name is the table name concatenated with a zero-padded, three-digit number unique to the table.

Schema changes Statements are not released if schema changes are initiated at the same time. See [“UltraLite database schemas” on page 45](#).

UltraLite does not process requests or queries referencing the table while the CREATE TABLE statement is being processed. Furthermore, you cannot execute CREATE TABLE when the database includes active queries or uncommitted transactions.

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Synchronization of external BLOB columns (UltraLiteJ only) On the consolidated database, the filename column is stored as a regular CHAR column and the BLOB file column is stored as a regular BLOB (LONG BINARY) column. On a download, the filename column is ignored and a new filename is generated using the database option (Connection.OPTION_BLOB_FILE_BASE_DIR) and the prefix and extension strings specified on the DEFAULT AUTOFILENAME clause. For J2SE value stored in the filename column will match the syntax <database_option_blob_file_base_dir><prefix><auto generated integer value>.<extension> and for the BlackBerry the value stored will match the syntax <prefix><auto generated integer value>.<extension>. Therefore, for the BlackBerry generated filenames are always relative.

Accessing external BLOB columns (UltraLiteJ only) Files containing external BLOB values are only opened when the client application attempts to read the column value. At that time, the filename stored in the column designated by the STORE AS FILE clause must be a valid filename. For the BlackBerry, relative filenames are resolved against the database option OPTION_BLOB_FILE_BASE_DIR. If UltraLite determines that a filename does not begin with the prefix "file://", it will prepend the filename with the value of the OPTION_BLOB_FILE_BASE_DIR option before attempting to open it.

A BLOB file is inserted into the database by specifying a filename for the **file-name-column** column. The inserted filename must be a valid filename and, for the BlackBerry, it must follow the format for a fully qualified, absolute path file name as described in the file URL format in IETF RFCs 1738 & 2396 (see the Package javax.microedition.io.file Description of BlackBerry JDE API documentation). On insert, a value for the STORE AS FILE column may not be specified.

Once a file is inserted into the database, the database assumes full control over the file and assumes that no outside modification will be made.

Side effects

Automatic commit.

See also

- [“Expressions in UltraLite” on page 280](#)
- [“DROP TABLE statement \[UltraLite\]” on page 435](#)
- [“CREATE TABLE statement” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Data types in UltraLite” on page 265](#)
- [“Overriding partition sizes for autoincremented columns” on page 71](#)

Example

The following statement creates a table for a library database to hold book information.

```
CREATE TABLE library_books (
  isbn CHAR(20) PRIMARY KEY,
  copyright_date DATE,
  title CHAR(100),
  author CHAR(50),
  location CHAR(50),
  FOREIGN KEY location REFERENCES room
);
```

The following statement creates a table for a library database to hold information on borrowed books. The default value for date_borrowed indicates that the book is borrowed on the day the entry is made. The date_returned column is NULL until the book is returned.

```
CREATE TABLE borrowed_book (
  loaner_name CHAR(100) PRIMARY KEY,
  date_borrowed DATE NOT NULL DEFAULT CURRENT DATE,
  date_returned DATE,
  book CHAR(20),
  FOREIGN KEY (book) REFERENCES library_books (isbn)
);
```

The following statement creates tables for a sales database to hold order and order item information.

```
CREATE TABLE Orders (
  order_num INTEGER NOT NULL PRIMARY KEY,
  date_ordered DATE,
  name CHAR(80)
);
CREATE TABLE Order_item (
  order_num INTEGER NOT NULL,
  item_num SMALLINT NOT NULL,
  PRIMARY KEY (order_num, item_num),
  FOREIGN KEY (order_num)
  REFERENCES Orders (order_num)
);
```

CREATE USER statement [UltraLite]

Creates a database user or group.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

```
CREATE USER user-name IDENTIFIED BY password
```

Parameters

user-name The name of the user you are creating.

password The password for the user you are creating.

Remarks

- User IDs cannot:
 - begin with white space, single quotes, or double quotes
 - end with white space
 - contain semicolons
- Passwords are case-sensitive and they cannot:
 - begin with white space, single quotes, or double quotes
 - end with white space
 - contain semicolons
 - be longer than 255 bytes in length

Permissions

DBA authority

Side effects

None.

See also

- [“ALTER USER statement \[UltraLite\]” on page 419](#)
- [“DROP USER statement \[UltraLite\]” on page 436](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following example creates a user named SQLTester with the password welcome.

```
CREATE USER SQLTester IDENTIFIED BY welcome;
```

DELETE statement [UltraLite]

Use this statement to delete rows from a table in the database.

Syntax

```
DELETE [ FROM ] table-name[[AS] correlation-name]  
[ WHERE search-condition ]
```

Parameters

correlation-name An identifier to use when referencing the table from elsewhere in the statement.

WHERE clause If a WHERE clause is specified, only rows satisfying *search-condition* are deleted.

The WHERE clause does not support non-deterministic functions (like RAND) or variables. Nor does this clause restrict columns; columns may need to reference another table when used in a subquery.

Remarks

The way in which UltraLite traces row states is unique. Be sure you understand the implication of deletes and row states.

See also

- [“Search conditions in UltraLite” on page 287](#)
- [“UltraLite database row state management” on page 469](#)
- [“START SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 446](#)
- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 447](#)

Example

The following statement removes employee 105 from the Employees table.

```
DELETE
FROM Employees
WHERE EmployeeID = 105;
```

The following statement removes all data before the year 2000 from the FinancialData table.

```
DELETE
FROM FinancialData
WHERE Year < 2000;
```

DROP INDEX statement [UltraLite]

Use this statement to delete an index.

Syntax

```
DROP INDEX[ IF EXISTS ] [ table-name.]index-name
```

Remarks

You cannot drop the primary index of a table.

UltraLite does not process requests or queries referencing the index while the DROP INDEX statement is being processed. Furthermore, you cannot execute DROP INDEX when the database includes active queries or uncommitted transactions.

Use the IF EXISTS clause if you do not want an error returned when the DROP INDEX statement attempts to remove an index that does not exist.

When you specify the IF EXISTS clause and the named table cannot be located, an error is returned.

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“UltraLite database schemas” on page 45](#).

See also

- [“CREATE INDEX statement \[UltraLite\]” on page 421](#)
- [“UltraLite indexes” on page 54](#)

Example

The following statement deletes a fictitious index, `fin_codes_idx`, on the `FinancialData` table:

```
DROP INDEX FinancialData.fin_codes_idx;
```

DROP PUBLICATION statement [UltraLite]

Use this statement to delete publications.

Syntax

```
DROP PUBLICATION[ IF EXISTS ] publication-name, ...
```

Remarks

Use the `IF EXISTS` clause if you do not want an error returned when the `DROP PUBLICATION` statement attempts to remove a publication that does not exist.

See also

- [“UltraLite client synchronization design” on page 72](#)
- [“ALTER PUBLICATION statement \[UltraLite\]” on page 413](#)
- [“CREATE PUBLICATION statement \[UltraLite\]” on page 423](#)

Example

The following statement drops the `pub_contact` publication.

```
DROP PUBLICATION pub_contact;
```

DROP SYNCHRONIZATION PROFILE statement [UltraLite]

Use this statement to delete an UltraLite or UltraLiteJ synchronization profile. Synchronization profiles define how an UltraLite or UltraLite Java edition database synchronizes with the MobiLink server.

Syntax

```
DROP SYNCHRONIZATION PROFILE [ IF EXISTS ] sync-profile-name
```

Parameters

- **sync-profile-name** The name of the synchronization profile.

Remarks

Use the IF EXISTS clause if you do not want an error returned when the DROP SYNCHRONIZATION PROFILE statement attempts to remove a synchronization profile that does not exist.

Side effects

None.

See also

- [“CREATE SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 424](#)
- [“ALTER SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 413](#)
- [“SYNCHRONIZE statement \[UltraLite\]” on page 448](#)

Example

The following example shows the syntax for dropping a synchronization profile called Test1.

```
DROP SYNCHRONIZATION PROFILE Test1;
```

DROP TABLE statement [UltraLite]

Use this statement to remove a table, and all its data, from a database.

Syntax

```
DROP TABLE [ IF EXISTS ] table-name
```

Remarks

The DROP TABLE statement drops the specified table from the database. All data in the table and any indexes and keys are also removed.

UltraLite does not process requests or queries referencing the table, or its indexes, while the DROP TABLE statement is being processed. Furthermore, you cannot execute DROP TABLE when there are active queries or uncommitted transactions.

Use the IF EXISTS clause if you do not want an error returned when the DROP TABLE statement attempts to remove a table that does not exist.

For UltraLite.NET, you cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“UltraLite database schemas” on page 45](#).

See also

- [“ALTER TABLE statement \[UltraLite\]” on page 415](#)
- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)

Example

The following statement deletes a fictitious table, EmployeeBenefits, from the database:

```
DROP TABLE EmployeeBenefits;
```

DROP USER statement [UltraLite]

Drops a user.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

```
DROP USER userid IDENTIFIED BY password
```

Parameters

- **userid** The userid of the user you are dropping.
- **password** The password for the user.

Permissions

DBA authority.

Remarks

None.

Side effects

None.

See also

- [“ALTER USER statement \[UltraLite\]” on page 419](#)
- [“CREATE USER statement \[UltraLite\]” on page 431](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following example drops the user SQLTester from a database.

```
DROP USER SQLTester;
```

FROM clause [UltraLite]

Use this clause to specify the tables or views involved in a SELECT statement.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

FROM *table-expression*, ...

table-expression :

table-name [[**AS**] *correlation-name*]

| (*select-list*) [**AS**] *derived-table-name* (*column-name*, ...)

| (*table-expression*)

| *table-expression* *join-operator* *table-expression* [**ON** *search-condition*] ...

join-operator :

| **INNER JOIN**

| **CROSS JOIN**

| **LEFT OUTER JOIN**

| **JOIN**

Parameters

table-name A base table or temporary table. Tables cannot be owned by different users in UltraLite. If you qualify tables with user ID, the ID is ignored.

correlation-name An identifier to use when referencing the table from elsewhere in the statement. For example, in the following statement, a is defined as the correlation name for the Contacts table, and b is the correlation name for the Customers table.

```
SELECT *
FROM Contacts a, Customers b
WHERE a.CustomerID=b.ID;
```

derived-table-name A derived table is a nested SELECT statement in the FROM clause.

Items from the SELECT list of the derived table are referenced by the (optional) derived table name followed by a period (.) and the column name. You can use the column name by itself if it is unambiguous.

You cannot reference derived tables from within the SELECT statement.

join-operator Specify the type of join. If you specify a comma (,), or CROSS JOIN, you cannot specify an ON subclause. If you specify JOIN, you must specify an ON subclause. For INNER JOIN and LEFT OUTER JOIN, the ON clause is optional.

Remarks

When there is no FROM clause, the expressions in the SELECT statement must be a constant expression.

Note

Although this description refers to tables, it also applies to derived tables unless otherwise noted.

The FROM clause creates a result set consisting of all the columns from all the tables specified. Initially, all combinations of rows in the specified tables are in the result set, and the number of combinations is usually reduced by JOIN conditions and/or WHERE conditions.

If you do not specify the type of join, and instead list the tables as a comma-separated list, a CROSS JOIN is used, by default.

For INNER joins, restricting results of the join using an ON clause or WHERE clause returns equivalent results. For OUTER joins, they are not equivalent.

Note

UltraLite does not support KEY JOINS nor NATURAL joins.

See also

- “Subqueries in expressions” on page 285
- “Joins: Retrieving data from several tables” [*SQL Anywhere Server - SQL Usage*]
- “DELETE statement [UltraLite]” on page 432
- “SELECT statement [UltraLite]” on page 444
- “UPDATE statement [UltraLite]” on page 452

Example

The following are valid FROM clauses:

```
...
FROM Employees
...

...
FROM Customers
CROSS JOIN SalesOrders
CROSS JOIN SalesOrderItems
CROSS JOIN Products
...
```

The following query uses a derived table to return the names of the customers in the Customers table who have more than three orders in the SalesOrders table:

```
SELECT Surname, GivenName, number_of_orders
FROM Customers JOIN
    ( SELECT CustomerID, COUNT(*)
      FROM SalesOrders
      GROUP BY CustomerID )
  AS sales_order_counts( CustomerID, number_of_orders )
ON ( Customers.id = sales_order_counts.CustomerID )
WHERE number_of_orders > 3;
```

INSERT statement [UltraLite]

Use this statement to insert rows into a table.

Syntax

```
INSERT [ INTO ]
 [ ( column-name, ... ) ]
{ VALUES ( expression, ... ) | select-statement }
```

Remarks

The INSERT statement can be used to insert a single row, or to insert multiple rows from a query result set.

If columns are specified, values are inserted one for one into the specified columns. If the list of column names is not specified, values are inserted into the table columns in the order in which they appear in the table (the same order as retrieved with SELECT *). Rows are inserted into the table in an arbitrary order.

Character strings inserted into tables are always stored in the same case as they are entered, regardless of whether the database is case sensitive.

See also

- [“SELECT statement \[UltraLite\]” on page 444](#)

Example

The following statement adds an Eastern Sales department to the database.

```
INSERT
INTO Departments ( DepartmentID, DepartmentName )
VALUES ( 230, 'Eastern Sales' );
```

The following statement adds the values of a and b into mytable from othertable where the value of c in othertable is greater than 10.

```
INSERT INTO mytable( col1, col2 ) SELECT a, b FROM othertable WHERE c > 10
```

LOAD TABLE statement [UltraLite]

Use this statement to import bulk data into a database table from an external file. This statement also provides support for handling the output of the SQL Anywhere dbunload utility (the reload.sql file). LOAD TABLE is only available using DBISQL on Windows.

Note

This statement is not supported by UltraLite Java edition databases.

The recommended method for unloading and reloading UltraLite databases is to use the ulunload and ulload utilities. Note also that the ulunit utility can load schema and data directly from a SQL Anywhere database.

```
LOAD [ INTO ] TABLE [ owner.]table-name
( column-name, ... )
FROM stringfilename
[ load-option ... ]
```

load-option :

```
CHECK CONSTRAINTS { ON | OFF }  
| COMPUTES { ON | OFF }  
| DEFAULTS { ON | OFF }  
| DELIMITED BY string  
| ENCODING encoding  
| ESCAPES { ON }  
| FORMAT { ASCII | TEXT }  
| ORDER { ON | OFF }  
| QUOTES { ON | OFF }  
| SKIP integer  
| STRIP { ON | OFF | BOTH }  
| WITH CHECKPOINT { ON | OFF }
```

comment-prefix : string

encoding : string

Parameters

- **column-name** Use this clause to specify one or more columns to load data into. Any columns not present in the column list become NULL if DEFAULTS is OFF. If DEFAULTS is ON and the column has a default value, that value is used. If DEFAULTS is OFF and a non-nullable column is omitted from the column list, the database server attempts to convert the empty string to the column's type.

When a column list is specified, it lists the columns that are expected to exist in the file and the order in which they are expected to appear. Column names cannot be repeated.

- **FROM string-filename** Use this to specify a file from which to load the data. The *string-filename* is passed to the database server as a string. The string is therefore subject to the same database formatting requirements as other SQL strings. In particular:
 - To indicate directory paths, the backslash character (\) must be represented by two backslashes. The statement to load data from the file *c:\temp\input.dat* into the Employees table is:

```
LOAD TABLE Employees  
FROM 'c:\\temp\\input.dat' ...
```
 - The path name is relative to the database server, not to the client application.
 - You can use UNC path names to load data from files on computers other than the database server.
- **load-option clause** There are several load options you can specify to control how data is loaded. The following list gives the supported load options:
 - **CHECK CONSTRAINTS clause** This clause controls whether constraints are checked during loading. CHECK CONSTRAINTS is ON by default, but the Unload utility (ulunload) writes out LOAD TABLE statements with CHECK CONSTRAINTS set to OFF. Setting CHECK CONSTRAINTS to OFF disables check constraints, which can be useful, for example, during database rebuilding.
 - **COMPUTES clause** This option is processed but ignored by UltraLite.

- **DEFAULTS clause** By default, DEFAULTS is set to OFF. If DEFAULTS is OFF, any column not present in the list of columns is assigned NULL. If DEFAULTS is set to OFF and a non-nullable column is omitted from the list, the database server attempts to convert the empty string to the column's type. If DEFAULTS is set to ON and the column has a default value, that value is used.
- **DELIMITED BY clause** Use this clause to specify the column delimiter string. The default column delimiter string is a comma; however, it can be any string up to 255 bytes in length (for example, . . . DELIMITED BY '###' . . .). The formatting requirements of other SQL strings apply. If you want to specify tab-delimited values, you could specify the hexadecimal escape sequence for the tab character (9), . . . DELIMITED BY '\x09'
- **ENCODING clause** This clause specifies the character encoding used for the data being loaded into the database. Note that UltraLite does not perform character set translation: the encoding of the data file must match the database.
- **ESCAPES clause** ESCAPES is always ON, therefore characters following the backslash character are recognized and interpreted as special characters by the database server. Newline characters can be included as the combination \n, and other characters can be included in data as hexadecimal ASCII codes, such as \x09 for the tab character. A sequence of two backslash characters (\\) is interpreted as a single backslash. A backslash followed by any character other than n, x, X, or \ is interpreted as two separate characters. For example, \q inserts a backslash and the letter q.
- **FORMAT clause** This clause specifies the format of the data source you are loading data from. With TEXT, input lines are assumed to be characters (as defined by the ENCODING option), one row per line, with values separated by the column delimiter string. ASCII is also supported.
- **QUOTES clause** This clause specifies whether strings are enclosed in quotes. UltraLite only supports ON, therefore the LOAD TABLE statement expects strings to be enclosed in quote characters. The quote character is an apostrophe (single quote). The first such character encountered in a string is treated as the quote character for the string. Strings must be terminated by a matching quote.

Column delimiter strings can be included in column values. Also, quote characters are assumed not to be part of the value. Therefore, the following line is treated as two values, not three, despite the presence of the comma in the address. Also, the quotes surrounding the address are not inserted into the database.

```
'123 High Street, Anytown', (715)398-2354
```

To include a quote character in a value, you must use two quotes. The following line includes a value in the third column that is a single quote character:

```
'123 High Street, Anytown', '(715)398-2354', ''''
```

- **SKIP clause** Use this clause to specify whether to ignore lines at the beginning of a file. The *integer* argument specifies the number of lines to skip. You can use this clause to skip over a line containing column headings, for example.

- **STRIP clause** This clause is processed but ignored. This clause specifies whether unquoted values should have leading or trailing blanks stripped off before they are inserted. The STRIP option accepts the following options:
 - **STRIP ON** Strip leading blanks.
 - **STRIP OFF** Do not strip off leading or trailing blanks.
 - **STRIP BOTH** Strip both leading and trailing blanks.
- **WITH CHECKPOINT clause** Use this clause to specify whether to perform a checkpoint. The default setting is OFF. If this clause is set to ON, a checkpoint is issued after successfully completing the statement.

Remarks

LOAD TABLE allows efficient mass insertion into a database table from a file. It is provided primarily as a means of supporting the output of the SQL Anywhere dbunload utility (the reload.sql file).

With FORMAT TEXT, a NULL value is indicated by specifying no value. For example, if three values are expected and the file contains 1 , , ' Fred ' , , then the values inserted are 1, NULL, and Fred. If the file contains 1 , 2 , , then the values 1, 2, and NULL are inserted. Values that consist only of spaces are also considered NULL values. For example, if the file contains 1 , , ' Fred ' , , then values 1, NULL, and Fred are inserted. All other values are considered not NULL. For example, " (a single quote followed by single quote) is an empty string. 'NULL' is a string containing four letters.

If a column being loaded by LOAD TABLE does not allow NULL values and the file value is NULL, then numeric columns are given the value 0 (zero), character columns are given an empty string ("). If a column being loaded by LOAD TABLE allows NULL values and the file value is NULL, then the column value is NULL (for all types).

If the table contains columns a, b, and c, and the input data contains a, b, and c, but the LOAD statement only specifies a and b as columns to load data into, the following values are inserted into column c:

- if DEFAULTS ON is specified, and column c has a default value, the default value is used.
- if column c does not have a default value, and NULLs are allowed, a NULL is used.
- if column c has no default value and does not allow NULLs, either a zero (0) or an empty string ("), is used, or an error is returned, depending on the data type of the column.

Side effects

Automatic commit.

See also

- [“INSERT statement \[UltraLite\]” on page 438](#)
- [“UltraLite Database Unload utility \(ulunload\)” on page 227](#)
- [“Unload utility \(dbunload\)” \[SQL Anywhere Server - Database Administration\]](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

Following is an example of LOAD TABLE. First, you create a table, and then load data into it using a file called *input.txt*.

```
CREATE TABLE t( a CHAR(100) primary key, let_me_default INT DEFAULT 1, c
CHAR(100) );
```

Following is the content of a file called *input.txt*:

```
'this_is_for_column_c', 'this_is_for_column_a', ignore_me
```

The following LOAD statement loads the file called *input.txt*:

```
LOAD TABLE T ( c, a ) FROM 'input.txt' FORMAT TEXT DEFAULTS ON;
```

The command `SELECT * FROM t` yields the result set:

a	let_me_default	c
this_is_for_column_a	1	this_is_for_column_c

ROLLBACK statement [UltraLite]

Use this statement to end a transaction and revert any changes made to data since the last COMMIT or ROLLBACK statement was executed.

Syntax

```
ROLLBACK [ WORK ]
```

Remarks

Using UltraLite SQL creates a transaction. A transaction consists of all changes (INSERTs, UPDATEs, and DELETEs) since the last ROLLBACK or COMMIT. The ROLLBACK statement ends the current transaction and undoes all changes made to the database since the previous COMMIT or ROLLBACK.

See also

- [“COMMIT statement \[UltraLite\]” on page 420](#)

Example

The following statement rolls the database back to the state it was in at the previous commit:

```
ROLLBACK;
```

SELECT statement [UltraLite]

Use this statement to retrieve information from the database.

Syntax

```
SELECT [ DISTINCT ] [ row-limitation ]  
select-list  
[ FROM table-expression, ... ]  
[ WHERE search-condition ]  
[ GROUP BY group-by-expression, ... ]  
[ ORDER BY order-by-expression, ... ]  
[ FOR { UPDATE | READ ONLY } ]  
[ OPTION ( FORCE ORDER ) ]
```

row-limitation :

```
FIRST  
| TOP n [ START AT m ]
```

select-list :

```
expression [ [ AS ] alias-name ], ...
```

order-by-expression :

```
{ integer | expression } [ ASC | DESC ]
```

Parameters

DISTINCT clause Specify DISTINCT to eliminate duplicate rows from the results. If you do not specify DISTINCT, all rows that satisfy the clauses of the SELECT statement are returned, including duplicate rows. Many statements take significantly longer to execute when DISTINCT is specified, so you should reserve DISTINCT for cases where it is necessary.

row-limitation clause Use row limitations to return a subset of the results. For example, specify FIRST to retrieve the first row of a result set. Use TOP*n* to return the first *n* rows of the results. Specify START AT*m* to control the location of the starting row when retrieving the TOP*n* rows. To order the rows so that these clauses return meaningful results, specify an ORDER BY clause for the SELECT statement.

select-list A list of expressions specifying what to retrieve from the database. Usually, the expressions in a SELECT list are column names. However, they can be other types of expressions, such as functions. Use an asterisk (*) to select all columns of all tables listed in the FROM clause. Optionally, you can define an alias for each expression in the *select-list*. Using an alias allows you to reference the *select-list* expressions from elsewhere in the query, such as from within the WHERE and ORDER BY clauses.

FROM clause Rows are retrieved from the tables and views specified in the *table-expression*. See [“FROM clause \[UltraLite\]” on page 436](#).

WHERE clause If a WHERE clause is specified, only rows satisfying *search-condition* are selected. See [“Search conditions in UltraLite” on page 287](#).

GROUP BY clause The result of the query that has a GROUP BY clause contains one row for each distinct set of values in the GROUP BY expression. The resulting rows are often referred to as groups

since there is one row in the result for each group of rows from the table list. Aggregate functions can be applied to the rows in these groups. NULL is considered to be a unique value if it occurs.

ORDER BY clause This clause sorts the results of a query according to the expression specified in the clause. Each expression in the ORDER BY clause can be sorted in ascending (ASC) or descending (DESC) order (the default). If the expression is an integer n , then the query results are sorted by the n th expression in the SELECT list.

The only way to ensure that rows are returned in a particular order is to use ORDER BY. In the absence of an ORDER BY clause, UltraLite returns rows in whatever order is most efficient.

UltraLite does not support the ordering of LONG VARCHAR or LONG BINARY values.

FOR clause This clause has two variations that control the query's behavior:

- **FOR READ ONLY** This clause indicates the query is not being used for updates. You should specify this clause whenever possible, since UltraLite can sometimes achieve better performance when it is known that a query is not going to be used for updates. For example, UltraLite could perform a direct table scan when it learns that read-only access is required. FOR READ ONLY is the default behavior. See [“Direct page scans” on page 464](#).
- **FOR UPDATE** This clause allows the query to be used for updates. This clause must be explicitly specified otherwise updates are not permitted (FOR READ ONLY is the default behavior).

OPTION (FORCE ORDER) clause This clause is not recommended for general use. It overrides the UltraLite choice of the order in which to access tables, and requires that UltraLite access the tables in the order they appear in the query. Only use this clause when the query order is determined to be more efficient than the UltraLite order.

UltraLite can also use execution plans to optimize queries. See [“Execution plans in UltraLite” on page 462](#).

Remarks

Always remember to close the query. Otherwise memory cannot be freed and the number of temporary tables that remain can proliferate unnecessarily.

See also

- [“UltraLite performance tips” on page 455](#)
- [“SELECT statement” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Queries” \[SQL Anywhere Server - SQL Usage\]](#)

Example

The following statement selects the number of employees from the Employees table.

```
SELECT COUNT(*)  
FROM Employees;
```

The following statement selects 10 rows from the Employees table starting from the 40th row and ending at the 49th row.

```
SELECT TOP 10 START AT 40 * FROM Employees;
```

SET OPTION statement [UltraLite]

Use this statement to change the values of database options.

Syntax

```
SET OPTION option-name=[option-value]
```

option-name: identifier

option-value: string, identifier, or number

Remarks

This statement allows you to set options on an UltraLite database. Note that most UltraLite options are set when the database is initially created and cannot be modified afterward.

You cannot specify whether an option is persistent or not. UltraLite determines whether it is a persistent or temporary option. Persistent options are stored in the database. Temporary options are used only until the connection or database stops running.

UltraLite performs a commit when persistent options are set: **global_database_id** and **ml_remote_id**. UltraLite does not perform a commit on temporary or connection-based options.

The only database option that can be unset is **ml_remote_id**. For example:

```
SET OPTION ml_remote_id=;
```

The result is that the ID is set to NULL. When this occurs, UltraLite will automatically generate a new value at the next synchronization.

See also

- [“UltraLite global_database_id option” on page 195](#)
- [“UltraLite ml_remote_id option” on page 196](#)
- [“UltraLite database options” on page 192](#)
- [“DB_PROPERTY function \[System\]” on page 334](#)

Example

The following statement sets the **global_database_id** option to 100:

```
SET OPTION global_database_id=100;
```

START SYNCHRONIZATION DELETE statement [UltraLite]

Use this statement to restart the logging of deleted rows for MobiLink synchronization.

Syntax**START SYNCHRONIZATION DELETE****Remarks**

UltraLite databases automatically log changes made to rows that need to be synchronized. UltraLite uploads these changes to the consolidated database during the next synchronization. This statement allows you to restart logging of deleted rows, previously stopped by a `STOP SYNCHRONIZATION DELETE` statement.

When a `STOP SYNCHRONIZATION DELETE` statement is executed, none of the delete operations executed on that connection are synchronized. The effect continues until a `START SYNCHRONIZATION DELETE` statement is executed.

Do not use `START SYNCHRONIZATION DELETE` if your application does not synchronize data.

The way in which UltraLite traces row states is unique. Be sure you understand the implication of deletes and row states.

See also

- [“UltraLite database row state management” on page 469](#)
- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 447](#)

Example

The following sequence of SQL statements illustrates how to use `START SYNCHRONIZATION DELETE` and `STOP SYNCHRONIZATION DELETE`.

```
STOP SYNCHRONIZATION DELETE;
DELETE FROM PROPOSAL
  WHERE last_modified < months( CURRENT_TIMESTAMP, -1 );
START SYNCHRONIZATION DELETE;
COMMIT;
```

STOP SYNCHRONIZATION DELETE statement [UltraLite]

Use this statement to stop the logging of deleted rows for MobiLink synchronization.

Syntax**STOP SYNCHRONIZATION DELETE****Remarks**

UltraLite databases automatically log changes made to rows that need to be synchronized. UltraLite uploads these changes to the consolidated database during the next synchronization. This statement allows you to stop the logging of deleted rows, previously started using a `STOP SYNCHRONIZATION DELETE` statement. This command can be useful when making corrections to a remote database, but should be used with caution as it effectively disables MobiLink synchronization. You should only stop deletion logging temporarily.

When a `STOP SYNCHRONIZATION DELETE` statement is executed, no further delete operations executed on that connection are synchronized. The effect continues until a `START SYNCHRONIZATION DELETE` statement is executed.

Do not use `STOP SYNCHRONIZATION DELETE` if your application does not synchronize data.

The way in which UltraLite traces row states is unique. Be sure you understand the implication of deletes and row states.

See also

- [“UltraLite database row state management” on page 469](#)
- [“START SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 446](#)

Example

The following sequence of SQL statements illustrates how to use `START SYNCHRONIZATION DELETE` and `STOP SYNCHRONIZATION DELETE`.

```
STOP SYNCHRONIZATION DELETE;
DELETE FROM PROPOSAL
WHERE last_modified < months( CURRENT_TIMESTAMP, -1 );
START SYNCHRONIZATION DELETE;
COMMIT;
```

SYNCHRONIZE statement [UltraLite]

Use this statement to synchronize an UltraLite or UltraLite Java edition database via the MobiLink server. The synchronization is configured according to the parameters in the synchronization profile, or the parameters can be specified in the statement itself.

Syntax

```
SYNCHRONIZE {
PROFILE sync-profile-name [ MERGE sync-option [ ;... ] ]
| USING sync-option [ ;... ]
}
```

sync-option :
sync-option-name = *sync-option-value*

sync-option-name : *string*

sync-option-value : *string*

Parameters

- **sync-profile-name** The name of the synchronization profile.
- **MERGE clause** Use this clause when you want to add or override options that are provided in the synchronization profile.
- **USING clause** Use this clause when you want to specify the synchronization options without referencing a synchronization profile.

- **sync-option** A string of one or more synchronization option value pairs, separated by semicolons. For example, 'option1=value1;option2=value2'.
- **sync-option-name** The name of the synchronization option.
- **sync-option-value** The value for the synchronization option.

Remarks

By allowing sync options to be merged in, developers can choose to omit storing some options in the database (like the MobiLinkPwd for instance).

If a synchronization callback function is defined and registered with UltraLite, whenever a SYNCHRONIZE statement is executed, progress information for that synchronization is passed to the callback function. If no callback is registered, progress information is suppressed.

Side effects

None.

See also

- [“CREATE SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 424](#)
- [“ALTER SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 413](#)
- [“DROP SYNCHRONIZATION PROFILE statement \[UltraLite\]” on page 434](#)
- [“ULSetSynchronizationCallback method \[UltraLite Embedded SQL\]” \[*UltraLite - C and C++ Programming*\]](#)

Example

The following example shows the syntax for synchronizing a synchronization profile called Test1 where the MobiLinkPwd has not been stored as part of the profile:

```
SYNCHRONIZE PROFILE Test1 MERGE 'MobiLinkPwd=sql'
```

The following example shows the syntax for adding the publication and uploadonly options to a synchronization profile called Test1.

```
SYNCHRONIZE PROFILE Test1
MERGE 'publications=p2;uploadonly=on';
```

The following example illustrates how to use USING.

```
SYNCHRONIZE USING
'MobiLinkUid=joe;MobiLinkPwd=sql;ScriptVersion=1;Stream=TCPIP{host=localhost}
```

The following example shows the syntax for synchronizing the publication and uploadonly options.

```
SYNCHRONIZE
USING 'publications=p2;uploadonly=on';
```

TRUNCATE TABLE statement [UltraLite]

Use this statement to delete all rows from a table, without deleting the table.

Syntax

```
TRUNCATE TABLE table-name
```

Remarks

The TRUNCATE TABLE statement deletes all rows from a table and the MobiLink server is not informed of their removal upon subsequent synchronization. It is equivalent to executing the following statements:

```
STOP SYNCHRONIZATION DELETE;  
DELETE FROM TABLE;  
START SYNCHRONIZATION DELETE;
```

Note

This statement should be used with great care on a database involved in synchronization or replication. Because the MobiLink server is not notified, this deletion can lead to inconsistencies that can cause synchronization or replication to fail.

After a TRUNCATE TABLE statement, the table structure, all the indexes, and the constraints and column definitions continue to exist; only data is deleted.

TRUNCATE TABLE cannot execute if a statement that affects the table is already being referenced by another request or query. Similarly, UltraLite does not process requests referencing the table while that table is being altered. Furthermore, you cannot execute TRUNCATE TABLE when the database includes active queries or uncommitted transactions.

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“ULBulkCopy.Dispose method \[UltraLite.NET\]” \[UltraLite - .NET Programming\]](#).

Schema changes Statements are not released if schema changes are initiated at the same time.

Side effects

If the table contains a column defined as DEFAULT AUTOINCREMENT or DEFAULT GLOBAL AUTOINCREMENT, TRUNCATE TABLE resets the next available value for the column.

Once rows are marked as deleted with TRUNCATE TABLE, they are no longer accessible to the user who performed this action, unless the user issues a ROLLBACK statement. However, they do remain accessible from other connections. Use COMMIT to make the deletion permanent, thereby making the data inaccessible from all connections.

If you synchronize the truncated table, all INSERT statements applied to the table take precedence over a TRUNCATE TABLE statement.

See also

- [“UltraLite database schemas” on page 45](#)
- [“DELETE statement \[UltraLite\]” on page 432](#)
- [“START SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 446](#)
- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 447](#)

Example

The following statement deletes all rows from the Departments table.

```
TRUNCATE TABLE Departments;
```

If you execute this example, be sure to execute a ROLLBACK statement to revert your change.

UNION statement [UltraLite]

Use this statement to combine the results of two or more select statements.

Note

This statement is not supported by UltraLite Java edition databases.

Syntax

```
select-statement-without-ordering
[ UNION [ ALL | DISTINCT ] select-statement-without-ordering ]...
[ ORDER BY [ number [ ASC | DESC ] , ... ]
```

Remarks

The results of several SELECT statements can be combined into a larger result using UNION. Each SELECT statement must have the same number of expressions in their respective SELECT list, and cannot contain an ORDER BY clause.

The results of UNION ALL are the combined results of the unioned SELECT statements. Specify UNION or UNION DISTINCT to get results without duplicate rows; however, note that removing duplicate rows adds to the total execution time for the statement. Specify UNION ALL to allow duplicate rows.

When attempting to combine corresponding expressions that are of different data types, UltraLite attempts find a data type in which to represent the combined values. If this is not possible, the union operation fails and an error is returned (for example "Cannot convert 'Surname' to a numeric").

The column names displayed in the results are column names (or aliases) used for the first SELECT statement.

ORDER BY for UNION is restricted to the integer format. The ORDER BY clause uses integers to establish the ordering, where the integer indicates the query expression(s) on which to sort the results.

See also

- [“SELECT statement \[UltraLite\]” on page 444](#)

Example

The following example lists all distinct surnames found in the Employees and Customers tables, combined.

```
SELECT Surname FROM Employees
UNION
SELECT Surname FROM Customers;
```

UPDATE statement [UltraLite]

Use this statement to modify rows in a table.

Syntax

```
UPDATE table-name[[AS] correlation-name]
SET column-name = expression, ...
[ WHERE search-condition ]
```

Parameters

table-name The *table-name* specifies the name of the table to update. Only a single table is allowed.

correlation-name An identifier to use when referencing the table from elsewhere in the statement.

SET clause Each named column is set to the value of the expression on the right side of the equal sign. There are no restrictions on the expression. If the expression is a *column-name*, the old value is used.

Only columns specified in the SET clause have their values changed. In particular, you cannot use UPDATE to set a column's value to its default.

WHERE clause If a WHERE clause is specified, only rows satisfying *search-condition* are updated.

Remarks

The UPDATE statement modifies values in a table.

Character strings inserted into tables are always stored in the same case as they are entered, regardless of whether the database is case sensitive.

See also

- [“INSERT statement \[UltraLite\]” on page 438](#)
- [“DELETE statement \[UltraLite\]” on page 432](#)
- [“Search conditions in UltraLite” on page 287](#)

Example

The following statement transfers employee Philip Chin (employee 129) from the sales department to the marketing department (department 400).

```
UPDATE Employees
SET DepartmentID = 400
WHERE EmployeeID = 129;
```

An example using *correlation-name*.

```
UPDATE Employee E
SET salary = salary * 1.05
WHERE EXISTS( SELECT 1 FROM Sales S HAVING E.Sales > Avg( S.sales)
GROUP BY S.dept_no )
```

UltraLite performance tips

This section describes ways you can tune various UltraLite features to improve the performance of queries, inserts or updates.

Adjusting the cache size for an UltraLite database

UltraLite database cache sizes increase dynamically in response to data operations and as available device memory allows within the parameters you specify. Normally you do not need to specify any parameters. If your database is large (400 MB for instance), you may want to specify the `CACHE_MAX_SIZE` parameter to raise the maximum limit beyond the default. UltraLite allocates some data structures based on the maximum cache size, so the default is not extremely large: you must explicitly request a large maximum to incur this extra memory overhead. There is no benefit to specifying a maximum cache size that is much larger than your maximum actual database file size.

Note

Dynamic cache sizing is not supported by UltraLite Java edition databases. For more information about UltraLite Java edition database cache sizes, see [“UltraLite Java edition database cache sizes” on page 456](#).

Although explicitly adjusting the cache sizing is not required, you typically want to adjust the cache size when your UltraLite database application is requested to reduce its memory usage by the operating system on mobile devices.

UltraLite does not shrink the cache automatically. The database cache size can only be controlled explicitly in your application with the `cache_allocation` database option.

Adjust the UltraLite database cache size in an UltraLite C++ application

This procedure illustrates the dynamic cache sizing functionality using the UltraLite C++ API as an example. It assumes that you want to update an existing C++ application to explicitly control the cache size.

1. Update the connection string used to connect to your UltraLite database to set your cache size limits.

In this example, you set the maximum cache size to 100 MB:

```
static ul_char const * ConnectionParms =  
    "UID=DBA;PWD=sql;DBF=sample.udb;CACHE_MAX_SIZE=100m";
```

2. In response to a low memory event raised by the operating system, adjust the `cache_allocation` database option after connecting to the database.

In this example, you reduce the cache allocation in half to resize the cache using the following code:

```
ULConnection * conn = ULDatabaseManager::OpenConnection(ConnectionParms);  
ul_u_long percent;
```

```
percent = conn->GetDatabasePropertyInt( "cache_allocation" );
conn->SetDatabaseOptionInt( "cache_allocation", percent / 2 );
```

See also

- [“UltraLite cache_allocation option” on page 192](#)
- [“UltraLite CACHE_SIZE connection parameter” on page 169](#)
- [“UltraLite CACHE_MIN_SIZE connection parameter” on page 168](#)
- [“UltraLite CACHE_MAX_SIZE connection parameter” on page 167](#)

UltraLite Java edition database cache sizes

UltraLite databases are stored on a set of pages on all platforms. A working set of pages is maintained in cache and is managed using a first-in, first-out (FIFO) scheme. Pages that are currently in use are locked in cache to avoid being swapped out.

For larger databases, the database can be configured each time it is opened to limit the number of rows and index pages that resident in memory at one time.

You can graph the UltraLite Java edition database performance with benchmark tests to determine the optimal cache size for your business solution.

You can test different cache sizes and watch for performance to change abruptly. Your cache should be large enough to have a good working set of pages. Consider the following ideas to help you stress the cache:

- Create multiple indexes on the table and add foreign keys.
- Insert rows randomly (something other than the index order).
- Create large rows, at least 25% of the database page size.
- Set the index hash to something other than 0. This increased size also increases the page accesses needed.
- Start graphing performance based on the smallest cache size. For example, 256 KB on Windows (the smallest allowed cache for this platform) or 64 KB on all other platforms.
- Encrypt or obfuscate the database. Obfuscation uses less code compared to strong encryption and performs fewer computations. Simple encryption performance should only be marginally slower than no encryption at all. However, your security requirements ultimately dictate whether to use strong encryption.

See also

- [“Use index hashing” on page 458](#)
- [“Encrypting data in an UltraLite Java edition database” \[*UltraLite - Java Programming*\]](#)

Query performance tips

This section illustrates methods for improving the performance of executing queries in UltraLite databases.

Use index scans

You can create one or more indexes to improve the performance of your queries, or, depending on the type of index you create, to ensure that row values remain unique.

An index provides an ordering of a table's rows based on the values in some or all of the columns. When creating indexes, the order in which you select columns to be indexed becomes the order in which the columns actually appear in the index. Indexes can greatly improve the performance of searches on the indexed column(s) when used strategically.

Use the following recommended practices for improving query performance:

- Create an index on any column:
 - for values that you search for on a regular basis
 - that the query uses to join tables
 - that are commonly used in ORDER BY, GROUP BY, or WHERE clauses
- Create a composite index and ensure that the first column of the index is used most often by the predicate in your query when creating it.
- Ensure that the update maintenance overhead an index introduces is not too high for the memory of your device.
- Do not create or maintain unnecessary indexes: indexes must be updated when the data in a column is modified, so all insert, update, and delete operations are performed on the indexes as well.
- Create an index on large tables.
- Do not create redundant indexes. For example, if you create an index on table T with columns (x, y), you can create a redundancy if there is another existing index on T with columns (x, y, z).

See also

- [“UltraLite indexes” on page 54](#)
- [“Manage temporary tables” on page 463](#)
- [“Direct page scans” on page 464](#)
- [“Viewing an UltraLite execution plan” on page 465](#)
- [“About composite indexes” on page 55](#)
- [“EXPLANATION function \[Miscellaneous\]” on page 338](#)
- [“ULPreparedStatement.GetPlan method \[UltraLite C++\]” \[*UltraLite - C and C++ Programming*\]](#)
- [“ULCommand.Plan property \[UltraLite.NET\]” \[*UltraLite - .NET Programming*\]](#)
- [“PreparedStatement.getPlan method \[UltraLite for M-Business Anywhere\]” \[*UltraLite - M-Business Anywhere Programming \(deprecated\)*\]](#)
- [“PreparedStatement.getPlan method \[UltraLiteJ\]” \[*UltraLite - Java Programming*\]](#)

Use index hashing

You can tune the performance of your queries by choosing a specific size for the maximum **hash**. A hash key represents the actual values of the indexed column. An index hash key aims to avoid the expensive operation of finding, loading, and then unpacking the rows to determine the indexed value. It prevents these operations by including enough of the actual row data with a row ID.

A row ID allows UltraLite to locate the actual row data in the database file. If you set the hash size to 0 (which disables index hashing), then the index entry only contains this row ID. If you set the hash size to anything other than 0, then a hash key is also used. A hash key can contain all or part of the transformed data in that row, and is stored with the row ID in the index page.

How much row data the hash key includes is determined by:

- The maximum hash size property you configure.
- How much is actually needed for the data type of the column.

A hash example

The value of an index hash maintains the order of the actual row data of indexed columns. For example, if you have indexed a LastName column for a table called Employees, you may see four names ordered as follows:

- Anders
- Anderseck
- Andersen
- Anderson

If you hashed the first six letters, your hash keys for these row values would appear as follows:

- Anders

- Anders
- Anders
- Anders

While these entries look the same, note that the first Anders in the list is used to represent the actual row value of **Anders**. The last Anders in the list, however, is used to represent the actual row value **Anderson**.

Consider the following statement:

```
SELECT *  
FROM Employees  
WHERE LastName = 'Andersen';
```

If the Employees table only contained a very high proportion of names similar to Andersen, then the hash key may not offer enough uniqueness to gain any performance benefits. In this case, UltraLite cannot determine if any of the hash keys actually meets the conditions of this statement. When duplicate index hash keys exist, UltraLite still needs to:

1. Find the table row that matches the row ID in question.
2. Load and then unpack the data so the value can be evaluated.

Performance benefits only occur when UltraLite can discern a proportionate number of unique hash so that the query condition evaluation is immediate to the index itself. For example, if the Employees table had thousands of names, there is still enough benefit to be gained by a hash of six letters. However, if the Employees table only contained an inordinate number of names that begin with Anders*, then you should hash at least seven letters so the degree of unique keys increases. Therefore, the original four names at the start of this example how are now represented with these hash keys:

- Anders
- Anderse
- Anderse
- Anderso

In this example, only two of the four row values would need to be unpacked and evaluated, rather than all four.

See also

- [“Choose and set the maximum hash size” on page 459](#)
- [“Adding an UltraLite index” on page 57](#)
- [“UltraLite max_hash_size creation parameter” on page 149](#)

Choose and set the maximum hash size

The UltraLite default maximum hash size of 4 bytes was chosen to suit most deployments. You can increase the size to include more data with the row ID. However, this change could increase the size of

the index and fragment it among multiple pages. This change can possibly increase the size of the database as a result. The impact of an increased maximum hash size depends on the number of rows in the table: for example, if you only have a few rows, a large index hash key would still fit on the index page. No index fragmentation occurs in this case.

When choosing an optimal hash size, consider the data type, the row data, and the database size (especially if a table contains many rows).

The only way to determine if you have chosen an optimal hash size is to run benchmark tests against your UltraLite client application on the target device. You need to observe how various hash sizes affect the application and query performance, in addition to the changes in database size itself.

The data type

If you want to hash the entire value in a column, note the size required by each data type in the table that follows. UltraLite only uses the maximum hash size if it really needs to, and it never exceeds the maximum hash size you specify. UltraLite always use a smaller hash size if the column type does not use the full byte limit.

Data type	Bytes used to hash the entire value
FLOAT, DOUBLE, and REAL	Not hashed.
BIT and TINYINT	1
SMALL INT and SHORT	2
INTEGER, LONG, and DATE	4
DATETIME, TIME, TIME-STAMP, and BIG	8
CHAR and VARCHAR	<p>To hash the entire string, the maximum hash size in bytes must match the declared size of the column. In a UTF-8 encoded database, always multiply the declared size by a factor of 2, but only to the allowed maximum of 32 bytes.</p> <p>For example, if you declare a column VARCHAR(10) in a non-UTF-8 encoded database, the required size is 10 bytes. However, if you declare the same column in a UTF-8 encoded database, the size used to hash the entire string is 20 bytes.</p>
BINARY	<p>The maximum hash size in bytes must match the declared size of the column.</p> <p>For example, if you declare a column BINARY(30), the required size is 30 bytes.</p>

Data type	Bytes used to hash the entire value
UUID	16

For example, if you set a maximum hash size of 6 bytes for a two-column composite index that you declared as INTEGER and BINARY (20) respectively, then based on the data type size requirements, the following occurs:

- The entire value of the row in the INTEGER column is hashed and stored in the index because only 4 bytes are required to hash integer data types.
- Only the first 2 bytes of the BINARY column are hashed and stored in the index because the first 4 bytes are used by the INTEGER column. If these remaining 2 bytes do not hash an appropriate amount of the BINARY column, increase the maximum hash size.

The row data

The row values of the data being stored in the database also influence the effectiveness of a hashed index.

For example, if you have a common prefix shared among entries of a given column, you may render the hash ineffective if you choose a size that only hashes prefixes. In this case, you need to choose a size that ensures more than just the common prefix is hashed. If the common prefix is long, you should consider not hashing the values at all.

When a non-unique index stores many duplicate values, and UltraLite cannot hash the entire value, the hash likely cannot improve performance.

The database size

Each index page has some fixed overhead, but the majority of the page space is used by the actual index entries. A larger hash size means each index entry is bigger, which means that fewer entries can fit on a page. For large tables, indexes with large hashes use more pages than indexes with small or no hashes. The more pages required increases the database size and degrades performance. The latter typically occurs because the cache can only hold a fixed number of pages thereby causing UltraLite to swap pages.

The following table gives you an approximation of how the hash size can affect the number of pages required to store data in an index:

Table	Page size	Hash size	Number of entries	Pages required
Table A	4 KB	0	1200	3 pages
Table B	4 KB	32 bytes	116	3 pages
Table C	4 KB	32 bytes	1200 entries	11 pages

Set the hash size

You can set the maximum hash size in two ways:

- By storing a database *default* for the maximum size, you set the `max_hash_size` creation parameter when you create your database. If you do not want to hash indexes by default, set this value to 0. Otherwise, you can change it to any value up to 32 bytes, or keep the UltraLite default of 4 bytes.
- If you want to *override* the default, you can set a specific hash size when you create a new index. Do one of the following methods:
 - In Sybase Central, set the Maximum Hash Size property when creating a new index.
 - With SQL, use the `WITH MAX HASH SIZE` clause in either the `CREATE TABLE` or `CREATE INDEX` statement.

See also

- [“Adding an UltraLite index” on page 57](#)
- [“Data types in UltraLite” on page 265](#)
- [“UltraLite max_hash_size creation parameter” on page 149](#)
- [“CREATE INDEX statement \[UltraLite\]” on page 421](#)
- [“CREATE TABLE statement \[UltraLite\]” on page 425](#)

Execution plans in UltraLite

UltraLite execution plans show how tables and indexes are accessed when a query is executed. UltraLite includes a **query optimizer**. The optimizer is an internal component of the UltraLite runtime that attempts to produce an efficient plan for the query. It tries to avoid the use of temporary tables to store intermediate results and attempts to ensure that only the pertinent subset of a table is accessed when a query joins two tables.

Overriding the optimizer

The optimizer always aims identify the most efficient access plan possible, but this goal is not guaranteed—especially with a complicated query where a great number of possibilities may exist. In extreme cases, you can override the table order it selects by adding the `OPTION (FORCE ORDER)` clause to a query, which forces UltraLite to access the tables in the order they appear in the query. *This option is not recommended for general use.* If performance is slow, a better approach is usually to create appropriate indexes to speed up execution.

Tip

If you are not going to update data with the query, you should specify the `FOR READ ONLY` clause in your query. This clause may yield better performance. See [“SELECT statement \[UltraLite\]” on page 444](#).

Determining the access method used by the optimizer

The UltraLite optimizer uses sophisticated optimization strategies when choosing an index for query optimization. However, with simple queries you cannot easily predetermine which index the optimizer

uses to optimize the query performance, or if an index is used at all. As the complexity increases, the index selected depends on the clauses required by your query. Usually, the presence of a FOR READ ONLY clause may cause the optimizer to choose a direct table scan instead of an index to yield better query performance.

When optimizing a query, the optimizer looks at the requirements of the query and checks if there are any indexes that it can use to improve performance. If performance cannot be improved with any index, then the optimizer does not scan one: either a temporary table or a direct page scan is used instead. Therefore, you may need to experiment with your indexes and frequently check the generated execution plans to ensure that:

- You are not maintaining indexes that are not being used by the optimizer.
- You are minimizing the number of temporary tables being created.

For complex queries, knowing which index is used is even less predictable. For example, when a query contains a WHERE predicate and a GROUP BY clause in addition to an ORDER BY clause, one index alone might not satisfy the search conditions of this query. So, if you have created an index to meet the selectivity requirements of the WHERE predicate, you may find that the optimizer does not actually use it. Instead, the optimizer may use an index that offers better performance for the ORDER BY conditions because this clause could require the most processing.

Checking the execution plan

You can check the execution plan either programmatically with the appropriate API call or in the Plan Viewer in Interactive SQL:

- **If no index is used** the execution plan appears as follows:

```
scan(T)
```
- **If a temporary table is used** the execution plan appears as follows:

```
temp [scan(T)]
```
- **If an index is used** the index name is included the execution plan:

```
scan (T, index_name)
```

See also

- [“Manage temporary tables” on page 463](#)

Manage temporary tables

In general, the optimizer always tries to avoid creating temporary tables to return query results because the entire temporary table must be populated before the first row can be returned. If an index exists, the optimizer tries to use the index first and only creates a temporary table as a last resort.

A temporary table is used by an access plan to store data during its execution in a transient or temporary work table. This table only exists while the access plan is being executed. Generally, temporary tables are used when intermediate results do not fit in the available memory, such as:

- When subqueries need to be evaluated early in the access plan.
- When data in a temporary table is held for a single connection only.
- When a query contains an ORDER BY on a column other than an index.
- When a query contains a GROUP BY on a column other than an index.

It is difficult to anticipate whether an index you have created avoids the necessity for a temporary table. Therefore, you should always check the plans for a query to ensure the indexes you have created are actually being used by the UltraLite query optimizer.

You can avoid using temporary tables by using an index for the columns used in the ORDER BY or GROUP BY clauses.

See also

- [“UltraLite TEMP_DIR connection parameter” on page 186](#)
- [“Determining the access method used by the optimizer” on page 462](#)
- [“Reading UltraLite execution plans” on page 466](#)
- [“When to view an execution plan” on page 465](#)

Direct page scans

UltraLite uses direct page scans as an alternative to index scans when it is more efficient to access information directly from the database page. A direct page scan is only used after the optimizer confirms that:

- No preexisting index can return results more efficiently.
- You are not using the query to perform updates. For example, you have declared the statement to be FOR READ ONLY (the default setting if no FOR clause has been specified), or have written the query in such a way that it is obvious that data is not being updated.

Because UltraLite reads the rows directly from the pages on which the rows are stored, query results are returned without order. The order of subsequent query results is unpredictable. If you need the order of rows to be predictable and deterministic, use an ORDER BY clause to get results in a consistent order. However, if order is not important, you can omit the ORDER BY clause to improve query performance.

Note

You cannot use direct page scans if you are using a ULTable class in an UltraLite API to program your application.

You can check to see when UltraLite scans a page directly or which index was used to return results.

See also

- [“Determining the access method used by the optimizer” on page 462](#)
- [“Use index scans” on page 457](#)

When to view an execution plan

View an execution plan in Interactive SQL when you need to know:

- What index will be used to return the results. An index scan object contains the name of the table and the index on that table that is being used.
- Whether a temporary table will be used to return the results. Temporary tables are written to the UltraLite temporary file.
- Which order tables are joined. This information allows you to determine how performance is affected.
- Why a query is running slowly or to ensure that a query does not run slowly.

See also

- [“UltraLite TEMP_DIR connection parameter” on page 186](#)

Viewing an UltraLite execution plan

As a development aid, you can use Interactive SQL to display an UltraLite plan that summarizes how a prepared statement is to be executed. The text plan is displayed in the Interactive SQL Plan Viewer.

In UltraLite, an execution plan is strictly a short textual summary of the plan. No other plan types are supported. However, being a short plan, it allows you to compare plans quickly, because information is summarized on a single line.

View an execution plan in the Plan Viewer

1. Click **Tools » Plan Viewer**.
2. In the **SQL** pane, type a query.
3. Click **Get Plan** to generate a plan for the specified SQL statements.

Example

The text plan appears in the lower pane of the Plan Viewer.

Consider the following statement:

```
SELECT I.inv_no, I.name, T.quantity, T.prod_no
FROM Invoice I, Transactions T
WHERE I.inv_no = T.inv_no;
```

This statement might produce the following plan:

```
join[scan(Invoice,primary),index-scan(Transactions,secondary)]
```

The plan indicates that the join operation is completed by reading all rows from the Invoice table (following index named primary). It then uses the index named secondary from the Transactions table to read only the row whose `inv_no` column matches.

See also

- [“Interactive SQL utility \(dbisql\)” \[SQL Anywhere Server - Database Administration\]](#)
- [“Reading UltraLite execution plans” on page 466](#)

Reading UltraLite execution plans

Because UltraLite short plans are textual summaries of how a query is accessed, you need to understand how the operations of either a join or a scan of a table are implemented.

- **For scan operations** Represented with a single operand, which applies to a single table only and uses an index. The table name and index name are displayed as round brackets ((,)) following the operation name.
- **For other operations** Represented with one or more operands, which can also be plans in and of themselves. In UltraLite, these operands are comma-separated lists contained by square brackets ([]).

Operation list

Operations supported by UltraLite are listed in the table that follows.

Operation	Description
count (*)	Counts the number of rows in a table.
distinct [<i>plan</i>]	Implements the DISTINCT aspect of a query to compare and eliminate duplicate rows. It is used when the underlying plan sorts rows in such a way that duplicate contiguous rows are eliminated. If two contiguous rows match, only the first row is added to the result set.
dummy	No operation performed. It only occurs in two cases: <ul style="list-style-type: none">• When you specify DUMMY in a FROM clause.• When the FROM clause is missing from the query.
filter [<i>plan</i>]	Executes a search condition for each row supplied by the underlying plan. Only the rows that evaluate to true are forwarded as part of the result set.

Operation	Description
group-by [<i>plan</i>]	Creates an aggregate of GROUP BY results, to sort multiple rows of grouped data. Rows are listed in the order they occur and are grouped by comparing contiguous rows.
group-single [<i>plan</i>]	Creates an aggregate of GROUP BY results, but only when it is known that a single row will be returned.
keyset [<i>plan</i>]	Records which rows were used to create rows in a temporary table so UltraLite can update the original rows. If you do not want those rows to be updated, then use the FOR READ ONLY clause in the query to eliminate this operation.
index-scan (<i>table-name, index-name</i>)	Reads only part of the table; the index is used to find the starting row.
join [<i>plan, plan</i>]	Performs an inner join between two plans.
lojoin [<i>plan, plan</i>]	Performs a left outer join between two plans.
like-scan (<i>table-name, index-name</i>)	Reads only part of a table; the index is used to find the starting row by pattern matching.
rowlimit [<i>plan</i>]	Performs the row limiting operation on propagated rows. Row limits are set by the TOP n or FIRST clause of the SELECT statement.
scan (<i>table-name, index-name</i>)	Reads an entire table following the order indicated by the index.
sub-query [<i>plan</i>]	Marks the start of a subquery.
temp [<i>plan</i>]	<p>Creates a temporary table from the rows in the underlying plan. UltraLite uses a temporary table when underlying rows must be ordered and no index was found to do this ordering.</p> <p>You can add an index to eliminate the need for a temporary table. However, each additional index used increases the duration needed to insert or synchronize rows in the table for which the index applies.</p>
union-all [<i>plan, ..., plan</i>]	Performs a UNION ALL operation on the rows generated in the underlying plan.

Insert and Update performance tips

This section illustrates methods for improving the performance of executing data inserts and updates in UltraLite databases.

Manage transactions and row states

UltraLite maintains state information along with the data in the database. UltraLite tracks and stores state information so it can manage:

- Concurrent connections, so UltraLite can share resources as required.
- Synchronization progress state, to ensure that synchronization occurs successfully.
- Row state, to maintain data integrity by tracking how data has changed between synchronizations.
- Transactions, to determine when and how data gets committed. In UltraLite, a transaction is processed in its entirety or not at all.
- Recovery and backup information, to protect data against operating system crashes, and end-user actions such as removing storage cards, or device resets while UltraLite is running.

See also

- [“UltraLite concurrency” on page 468](#)
- [“UltraLite synchronization client features” on page 2](#)
- [“UltraLite database row state management” on page 469](#)
- [“UltraLite transaction processing” on page 470](#)
- [“UltraLite and UltraLite Java edition database back up and recovery” on page 43](#)

UltraLite concurrency

UltraLite automatically uses the following methods to manage concurrent database accesses:

- **Multiple UltraLite database access** A single application can open connections to multiple databases. UltraLite Java edition does not support simultaneous access.
- **Multiple applications** An UltraLite or UltraLite Java edition database can only be opened by one process at a time.
- **Multiple threads** UltraLite supports multi-threaded applications. A single application can be written to use multiple threads, each of which can connect to the same or different databases.
- **Multiple transactions/requests** Each connection can have a single transaction in progress at one time. Transactions can consist of a single request or multiple requests. Data modifications made during a transaction are not made permanent in the database until the transaction is committed. Either all data modifications made in a transaction are committed, or all are rolled back.

- **Synchronization** During upload and download, read-write access to the database is permitted. However, if an application changes a row that the download then attempts to change, the download fails and rolls back. Use the Disable Concurrency synchronization parameter to disable access to data during synchronization.

If synchronization fails, UltraLite supports resumable downloads on all platforms.

See also

- [“Failed downloads” \[MobiLink - Server Administration\]](#)
- [“UltraLite and UltraLite Java edition database limitations” on page 8](#)
- [“UltraLite transaction processing” on page 470](#)
- [“UltraLite clients” on page 67](#)
- [“Additional Parameters synchronization parameter” on page 88](#)

UltraLite database row state management

Maintaining row state information is a powerful part of the UltraLite database feature set. Tracking the state of tables and rows is particularly important for data synchronization.

Note

UltraLite Java edition does not support row state management. UltraLite Java edition databases use transaction logs to track changes that need to be synchronized.

An internal marker is used to keep track of the row state in an UltraLite database. Row states control transaction processing, recovery, and synchronization. When an application inserts, updates, or deletes a row, UltraLite modifies the state of the row to reflect the operation and the connection that performed the operation. When a transaction is committed, the states of all rows affected by the transaction are modified to reflect the commit. If an unexpected failure occurs during a commit, the entire transaction is rolled back. The following list summarizes these behaviors:

- **When a delete is issued** The state of each affected row is changed to reflect the fact that it was deleted. When a delete is undone through a rollback, the original state of the row is restored.
- **When a delete is committed** The affected rows are not always removed from memory. If the row has never been synchronized, then it is removed. If the row has been synchronized, then it is not removed, because the delete operation needs to be synchronized to the consolidated database first. After the next synchronization, the row is removed from memory.
- **When a row is updated** A new version of the row is created. The states of the old and new rows are set so the old row is no longer visible and the new row is visible.
- **When a row update is committed** When a transaction is committed, the states of all rows affected by the transaction are modified to reflect the commit. When an update is synchronized, both the old and new versions of the row are needed to allow conflict detection and resolution. The old row is then deleted from the database and the new row simply becomes a normal row.
- **When a row is added** The row is added to the database and is marked as not committed.

- **When an added row is committed** The row is marked as committed and is also flagged as requiring synchronization with the consolidated database.

See also

- [“UltraLite and UltraLite Java edition database back up and recovery” on page 43](#)
- [“Flush single or grouped transactions” on page 471](#)
- [“UltraLite transaction processing” on page 470](#)

UltraLite transaction processing

A transaction is a logical set of operations that are executed atomically: either all operations in the transaction are stored in the database or none are. An UltraLite application's access to the UltraLite runtime is serialized. While it is possible for multiple transactions to be open simultaneously, UltraLite only processes one transaction at a time. This behavior means that an application cannot:

- Have blocked transactions (also known as deadlocks). UltraLite never blocks a request based on an existing row lock. In this case, UltraLite immediately returns an error.
- Overwrite outstanding changes. A transaction cannot overwrite another transaction's outstanding changes. When a transaction changes a row, UltraLite locks that row until the transaction is **committed** or **rolled back**. The lock prevents other transactions from changing the row, although they can still read the row.

Tip

All UltraLite APIs—except the UltraLiteJ and C++ APIs—can operate in **autocommit** mode.

In autocommit mode, UltraLite executes a commit after each operation. Some APIs use autocommit by default. If you are using one of these interfaces, you must set autocommit to off to exploit multi-operation transactions. The way of turning autocommit off depends on the programming interface you are using. In most interfaces it is a property of the connection object.

See:

- UltraLite C++: [“Managing transactions”](#) [*UltraLite - C and C++ Programming*]
- UltraLite.NET: [“Transaction management”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“Transaction management”](#) [*UltraLite - M-Business Anywhere Programming (deprecated)*]
- UltraLiteJ: [“Transaction management”](#) [*UltraLite - Java Programming*]

For example, two applications, A and B, are reading the same row from the database and they both calculate new values for one of its columns based on the data they read. If A updates the row with its new value and B then tries to modify the same row, B gets an error. An attempt to change a locked row sets the error `SQLCODE SQLE_LOCKED`, while an attempt to change a deleted row sets the error `SQLCODE SQLE_NOTFOUND`. Therefore, you should program your application so it checks the `SQLCODE` value after attempting to modify data.

See also

- [UltraLite C++: “Error handling” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite.NET: “Error handling” \[UltraLite - .NET Programming\]](#)
- [UltraLite for M-Business Anywhere: “Error handling” \[UltraLite - M-Business Anywhere Programming \(deprecated\)\]](#)
- [UltraLiteJ: “Error handling” \[UltraLite - Java Programming\]](#)

Flush single or grouped transactions

You can choose your recovery point in UltraLite by delaying committed transaction flushes. When UltraLite releases the commit to storage, the recovery point helps control when a subset of SQL statements in a transaction triggers additional operational overhead.

By default, UltraLite uses an operational-based default that flushes individual transactions to storage immediately upon a commit. For some deployments, these frequent operations can be excessive and limit the amount of transaction throughput. To reduce the performance expense caused by this default, you may choose a state-based approach. Especially for applications that rely on autocommit operations, this approach delays the additional overhead required to flush the committed transactions to storage:

- **On checkpoint** You can set your own checkpoint, and then use it to release the work performed over the course of time. You can use as many checkpoints as you require, either within a single transaction or over multiple transactions.
- **Grouped** You can choose a transaction count threshold and/or a timeout threshold to release the work performed.

Delaying commit flushes based on state yields better performance and a cleaner application design because applications are not required to wait for a response from UltraLite. By delaying commit flushes you also minimize the exposure to transactions by giving more granular control over data for which work has not been fully completed. For example, in a sales application, an order may be available to a second application before all items have been added or even approved.

However, it is important for you to take into account the recoverability of a transaction for which commit flushes have been delayed. Transactions that have not been released cannot be recovered. Therefore, you need to evaluate the trade-off between the data integrity of your application and its performance.

See also

- [“UltraLite COMMIT_FLUSH connection parameter” on page 172](#)
- [“UltraLite commit_flush_count option \[temporary\]” on page 193](#)
- [“UltraLite commit_flush_timeout option \[temporary\]” on page 194](#)
- [“CHECKPOINT statement \[UltraLite\]” on page 420](#)
- [“ULCheckpoint method \[UltraLite Embedded SQL\]” \[UltraLite - C and C++ Programming\]](#)
- [“ULConnection.Checkpoint method \[UltraLite C++\]” \[UltraLite - C and C++ Programming\]](#)
- [“Connection.checkpoint method \[UltraLiteJ\]” \[UltraLite - Java Programming\]](#)

UltraLite benchmark tips

Benchmark testing activity is generally performed before reaching the production stage of the application development cycle. This phase requires that you test the UltraLite database with the application and ensure both components inter-operate as efficiently as possible. If you discover through your tests that performance is not as efficient as it could be, you can then tune your database and/or optimize your application to improve benchmark results.

Note

If your UltraLite deployment is part of a MobiLink synchronization environment, remember to test synchronization performance as well. See “[MobiLink tuning for performance](#)” [[MobiLink - Server Administration](#)].

Types of benchmark tests

You can run benchmark tests to check the performance of:

- SQL statements

The UltraLite database has been optimized to handle SQL queries efficiently and return results quickly. Nonetheless, you should see how well important queries perform to improve database performance.

- Synchronization

The key to achieving optimal MobiLink synchronization throughput is to have multiple synchronizations occurring simultaneously and executing efficiently.

- Indexes

- Table design

- Application code

- Device configuration

For example, compare using external flash memory and a device's internal memory as deployment locations for UltraLite.

- Database configuration

For example, try different cache sizes, page sizes, reserve sizes, indexes, hash sizes, and so on.

- Data throughput (based on transactions per second)

While UltraLite is not typically a database intended for mass data entry processing, depending on your business requirements, this benchmark is one that you may want to run tests for.

- Software changes

You should test the effect of software changes between two different versions of UltraLite or different versions of an application.

See also

- [“Test SQL queries” on page 473](#)
- [“MobiLink tuning for performance” \[MobiLink - Server Administration\]](#)

Test SQL queries

SQL queries can range from simple to complex. Depending on the nature and importance of your queries, there are two types of query benchmark testing available to you.

Representational SQL benchmarks

This type of testing requires that you test a selection of statements that are representative of typical transactions that the application performs during day-to-day operations. Different applications require different benchmark tests because the fundamental business use for each application varies. For example, a meter reading application might simply test a basic INSERT statement. However a mobile sales force application might test multiple INSERT statements, in addition to multiple SELECT statements and perhaps an UPDATE statement.

The volume of queries in your application can limit what you can realistically test. If you have excessive query processing, you may be limited to performing the targeted SQL benchmark testing.

Targeted SQL benchmarks

If you have a lot of statements used by your applications, you may want to narrow the scope of your test to include some or all of the following:

- The most frequently-used statements.
- The statements that process high volumes of data.
- The statements that have time-sensitive requirements.
- The statements that are most important to the business case of your application.
- The most complex statements: for example, those that have the largest number of table joins or that use many subqueries. These types of statements can use a large amount of device resources. Even if the statements are only used infrequently, you may want to check that they don't exceed the capacity of the device.
- The statements that are not supported by an index.
- The statements that use a large amount of memory resources.

Methodology

1. The preparation phase

Allows you to finalize your database design and reach a stable point in your application development before starting your benchmark testing.

2. The creation phase

Allows you to build a custom program that replicates the end-user behavior you predict for your UltraLite deployment.

3. The execution and analysis phase

Allows you to fine tune different elements of your database and record the results of those changes so you can analyze them. Tests are repeated until the maximum benefit of all modifications has been reached.

Understanding the preparation phase

The preparation phase allows you to get your database and application in a state where they can be successful benchmark candidates, and to determine what goals you hope to achieve from your tests.

Development cycle milestones checklist

1. Complete the logical design of the database.

Ensure you have:

- Created and populated tables with representative data.
- Created indexes to retrieve data in those tables more effectively.

2. Prepare the physical deployment environment of both the database and application. The deployment environment must accurately represent the final production environment: that is, the lab and production environments should share the same memory and disk configurations on the same platform/device type.

3. Ensure you have reached a stable point in your application programming phase. Remember you are looking for performance optimizations not defects; however, the latter may also be revealed as a result of testing.

All database queries should access and return as much data as required. If the production environment requires sorting of data, ensure queries include this data as well. Otherwise, the application cannot accurately test representative memory requirements.

4. Deploy a copy of the database and application to the final disk location.

5. Decide what element of the database performance you want to check and potentially tune.

See also

- [“Types of benchmark tests” on page 472](#)
- [“Methodology” on page 474](#)

The creation phase

You must create tests that yield reliable results. Otherwise, you cannot legitimately compare results over time.

The following characteristics make a benchmark test effective and reliable:

- **Goal** Are you looking to capture a performance ratio or are you trying to see the duration required to process a command against a database? For the former, if you are testing SQL performance, you may want to run one or more statements repeatedly until a set time interval has expired. This testing gives you the throughput ratio, which can be summarized as follows:

statement-number ÷ time-interval = throughput ratio

- **Environment** Establish a test environment as your baseline and record the design and scope of it. If you cannot run the same test under the same conditions, you cannot legitimately compare results of that test. Additionally, the hardware and software you use in the lab as part of your benchmark test should match that of your production environment.
- **State** Reliable benchmark tests always start each iteration with the same action. Decide whether or not third-party applications should operate concurrently with UltraLite. If they affect performance, you should add them to the benchmark test design. For third-party applications that should not be running, always exit these applications completely—even minimized or idle applications/processes could skew results because memory is still being used.
- **Results** Results of benchmarks must be captured in a consistent way after each iteration of the test. Over time the results can indicate a trend and help you determine what changes can yield an improvement in UltraLite performance—either in the database or in the application (or both).
- **Timing mechanism** Benchmark tests simulate user actions; therefore, you typically track the elapsed execution times of these actions. Ensure your timing mechanism is systematic so execution times are accurately reflected in the results of your tests.

See also

- [“Methodology” on page 474](#)
- [“Test SQL queries” on page 473](#)

Understanding the execution and analysis phase

Benchmark execution is the phase during which you tune your database by iteratively running the test, modifying something about the database (for example, the value of one or more database properties or connection parameters), and then running the test again to see the outcome of any changes.

The following procedure assumes you are testing different database properties and/or connection parameters to find the maximum benefit. Repeat this procedure until all parameters that require testing have been tested.

Tip

Choose only those properties or parameters that are significant to the workload and the objectives of your UltraLite deployment.

Execute your benchmark tests

1. Create a baseline by running the first iteration of the test. In this case, because you are testing different database properties and/or connection parameters, you would use UltraLite defaults wherever possible.
2. Begin your normal test runs by tuning only one database property or connection parameter at a time. This limitation ensures that the results you collect are systematic in their approach and helps you more readily determine when you have reached the maximum benefit of your tuning activities.
3. Output from the benchmark program should include:
 - an identifier or label for each test
 - the iteration of the program execution
 - the name of the element being checked and what you did to change it
 - the recorded elapsed time

For example, even though you could test other database parameters, if you limited your test to just varying page sizes, cache sizes, and reserve sizes, your output might be saved to a table that looks similar to the example that follows:

PROP/PARM	VALUES		
TEST_NUMBER	001	002	003
page_size	1	2	8
CACHE_SIZE	128	256	512
RESERVE_SIZE	128	256	512

STMT ID	EXECUTION (seconds)		
01	01.55	01.50	01.49
02	02.01	02.20	01.59
03	00.33	00.55	00.44

4. When you have completed an iteration, return the database to the baseline state to ensure you do not inadvertently contaminate results of subsequent runs.
5. Depending on the results of the benchmark test, do one of the following:
 - If performance improves, change the value of the same property or parameter and rerun the test. Keep tuning this value until you cannot improve performance any further.
 - If the performance worsens, return the value of the property or parameter to the previous value.
6. Test a new property or parameter.

See also

- [“Methodology” on page 474](#)

UltraLite troubleshooting

Unable to start the UltraLite engine

Symptom

You have to use the START connection parameter to start the UltraLite engine with the following definition; however, the client returns `SQL_UNABLE_TO_CONNECT_OR_START`.

```
START="\Program Files\ulengl2.exe"
```

Explanation

The location of the quotes is incorrect.

Recommendation

For this parameter to work, the first quotation mark must follow the `\` character. For example, you can delimit spaces in this path as follows:

```
START=\"Program Files\ulengl2.exe"
```

or

```
START='\"Program Files\ulengl2.exe''
```

Unable to connect to databases after upgrade

Symptom

You have upgraded UltraLite. You discover that you are able to create an empty database using the administration tools. However, when you try to connect to this or any other database (including *CustDB.udb*) with Sybase Central, you receive an error. Connecting to SQL Anywhere databases works without incident, however.

Explanation

You did not close all SQL Anywhere applications and processes. Therefore, your UltraLite plug-ins were not installed correctly.

Recommendation

Remove and reinstall SQL Anywhere.

1. Close Sybase Central, Interactive SQL, and any running database engines.
2. Run the following commands:

```
dbisql -terminate
```

```
scjview -terminate
```

3. Open the Windows **Task Manager**, and end any *scjview.exe* and *dbisql.exe* processes.
4. Reinstall the latest version of UltraLite.

See also

- [“UltraLite upgrades” \[SQL Anywhere 12 - Changes and Upgrading\]](#)

UltraLite database corruption

Symptom

Your UltraLite database may be corrupt if it:

- Generates the following errors:
 - SQLE_DEVICE_ERROR
 - SQLE_DATABASE_ERROR (can also be a symptom of other issues)
 - SQLE_MEMORY_ERROR (can also be a symptom of other issues)
- Crashes or returns invalid query results.

Explanation

There are two more typical causes corruption:

- The more frequent cause occurs if the device has problems storing the file, thereby spuriously changing the contents of it. This issue usually stops the UltraLite database from functioning fairly quickly.
- The less frequent cause occurs if an error in the UltraLite code fails to maintain an index correctly. These issues can go undetected for much longer because the change to the results of a query are more subtle.

Recommendation

Checksums are used to detect offline corruption in an UltraLite database, which can help reduce the chances of other data being corrupted as the result of a bad critical page. If a checksum validation fails when the UltraLite database loads a page, UltraLite immediately stops the database and reports a fatal error. This error cannot be corrected. Instead you must:

1. Report the error. It is helpful if you know the sequence of events that caused the corruption to occur, and if the error is reproducible.
2. If you need the data, unload the contents of the UltraLite database to a file.
3. Create a new UltraLite database.
4. Repopulate the data either by synchronizing or by loading the unloaded data.

See also

- [“UltraLite checksum_level creation parameter” on page 142](#)

Database size not stabilizing

Symptom

Your application collects a lot of large binary objects among multiple client devices, synchronizes this information to a consolidated database, and then the synchronized data is deleted from each client device. However, the database size remains large despite the data being removed from the database. This is a concern because file size needs to be managed carefully due to limited resources of the device.

Explanation

Database size should only increase if your data grows in the database. However, once grown, the database file keeps that size, and does not shrink on its own. Free space is maintained internally to the file.

Recommendation

Ensure you are not using the STOP SYNCHRONIZATION DELETE or TRUNCATE statements for tables that do not get synchronized. Instead use the DELETE statement with a FROM *table-name* clause for tables that do not get synchronized.

Recreate the database post-synchronization:

1. Create the database that is deployed to the devices.
2. Creating a SQL script of DDL statements that define the schema required by the client devices. See [“Deploying UltraLite database schema upgrades” on page 123](#).
3. Synchronize the data.
4. Drop the database.
5. Create a new, empty database and use standard database schema with the ALTER DATABASE SCHEMA FROM FILE statement.

See also

- [“STOP SYNCHRONIZATION DELETE statement \[UltraLite\]” on page 447](#)
- [“TRUNCATE TABLE statement \[UltraLite\]” on page 450](#)
- [“DELETE statement \[UltraLite\]” on page 432](#)
- [“ALTER DATABASE SCHEMA FROM FILE statement \[UltraLite\]” on page 412](#)

Importing ASCII data into a new UltraLite database

Symptom

You have created a new UltraLite database, but have a .csv ASCII data file that you cannot import.

Explanation

The .csv format is not supported by any of the UltraLite administration tools.

Recommendation

You can try one of the following techniques:

- Use Interactive SQL (dbisql) to import the data. You can connect to the UltraLite database and then click **Data » Import Data**. Alternatively, you can connect to the UltraLite database and run the INPUT statement (this statement cannot be used in an UltraLite PreparedStatement object).

Note

UltraLite requires primary keys. Although Interactive SQL can create the table for you, it does not automatically create the primary keys for them. Always connect to an empty UltraLite database you have created for this purpose.

- If you incorporate this functionality as part of a batch process, you must write your own code.

See also

- [“INPUT statement \[Interactive SQL\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Interactive SQL for UltraLite utility \(dbisql\)” on page 198](#)

Utilities still running as the previous version

Symptom

You have just installed UltraLite 12. However, when you try to run any of the UltraLite utilities, the previous version starts.

Explanation

If you have multiple versions of UltraLite on your computer, you must pay attention to your system path when using the administration. Since the installation adds the most recently installed version executable directory to the end of your system path, it is possible to install a new version of the software, and still inadvertently be running the previously installed version.

Recommendation

There are various workarounds to this problem.

See also

- [“Utility versions and the upgrade process” \[SQL Anywhere 12 - Changes and Upgrading\]](#)

Result set changes unpredictably

Symptom

You run a query and the result set you expect changes each time you run it.

Explanation

Carefully review the result set you are getting. Are the results in the set truly different? Or are they simply being returned in the most efficient order each time. The order selected can change each time you execute the query, depending on when you last accessed the row and other factors.

Recommendation

If your result set must be returned in a predictable or consistent order, ensure that the `SELECT` statement includes an `ORDER BY` clause. If the result set is still returning results incorrectly, your database may be corrupt.

See also

- [“SELECT statement \[UltraLite\]” on page 444](#)
- [“UltraLite database corruption” on page 480](#)

UltraLite engine client fails with error -764

Applies to

Windows Mobile

Symptom

You are running the UltraLite engine on Windows Mobile device, and the client returns a -764 error.

Explanation

An error of -764 means that the engine could not be found and was unable to start.

Recommendation

Consider one of the following actions:

- Consider redeploying the engine to the recommended deployment location, the `\Windows` directory. UltraLite automatically looks for the engine files in this location.
- If you have install the engine to any other location, ensure your connection code uses the `START` connection parameter.
- If you have used the `START` connection parameter, and you are sure the path to the engine is correct, ensure you have used the correct escape sequences for special characters in the path name.

For example, you may need to change this code:

```
ULConnection conn = new ULConnection(@"dbf=\Program Files\HelloEngine
\HelloEngine.udb;
START=\Windows\uleng12.exe")
```

To something similar to:

```
ULConnection conn = new ULConnection(@"dbf=\\\"Program Files \"\
\HelloEngine\HelloEngine.udb;
START=\\Windows\uleng12.exe");
```

See also

- [“Start a deployed UltraLite engine” on page 125](#)
- [“UltraLite START connection parameter” on page 185](#)

Index

Symbols

- % operator
 - modulo function, UltraLite, 360
- &
 - UltraLite bitwise operator, 298
- comment indicator
 - UltraLite about, 260
- case option
 - UltraLite initialize database utility (ulinit), 210
- checksum_level option
 - UltraLite initialize database utility (ulinit), 210
- collation option
 - UltraLite initialize database utility (ulinit), 210
- databaseid option
 - UltraLite initialize database utility (ulinit), 210
- datacopy option
 - UltraLite initialize database utility (ulinit), 210
- date_format option
 - UltraLite initialize database utility (ulinit), 210
- date_order option
 - UltraLite initialize database utility (ulinit), 210
- dba option
 - UltraLite initialize database utility (ulinit), 210
- exactschema option
 - UltraLite initialize database utility (ulinit), 210
- exclude option
 - UltraLite initialize database utility (ulinit), 210
- fips option
 - UltraLite initialize database utility (ulinit), 210
- help option
 - UltraLite initialize database utility (ulinit), 210
- identity--password option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
- identity-file option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
- insertforupload option
 - UltraLite initialize database utility (ulinit), 210
- key option
 - UltraLite initialize database utility (ulinit), 210
- list_collation option
 - UltraLite initialize database utility (ulinit), 210
- UltraLite initialize database utility (ulinit), 210
- log option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite synchronization utility (ulsync), 222
- max_hash_size option
 - UltraLite initialize database utility (ulinit), 210
- mirror_file option
 - UltraLite initialize database utility (ulinit), 210
- no_warnings option
 - UltraLite initialize database utility (ulinit), 210
- obfuscate option
 - UltraLite initialize database utility (ulinit), 210
- overwrite option
 - UltraLite initialize database utility (ulinit), 210
- page_size option
 - UltraLite initialize database utility (ulinit), 210
- precision option
 - UltraLite initialize database utility (ulinit), 210
- prompt option
 - UltraLite initialize database utility (ulinit), 210
- publication option
 - UltraLite initialize database utility (ulinit), 210
- quiet option
 - UltraLite initialize database utility (ulinit), 210
- reserve_size option
 - UltraLite initialize database utility (ulinit), 210
- rootcert option
 - UltraLite initialize database utility (ulinit), 210
- SAconnect option
 - UltraLite initialize database utility (ulinit), 210
- scale option
 - UltraLite initialize database utility (ulinit), 210
- sql option
 - UltraLite initialize database utility (ulinit), 210
- sync_publication option
 - UltraLite initialize database utility (ulinit), 210
- time_format option
 - UltraLite initialize database utility (ulinit), 210
- timestamp_increment option
 - UltraLite initialize database utility (ulinit), 210
- timestamp_with_time_zone_format option
 - UltraLite initialize database utility (ulinit), 210
- utf8_encoding option
 - UltraLite initialize database utility (ulinit), 210
- /? option
 - UltraLite initialize database utility (ulinit), 210
- a option
 - ulinit creation parameters, 22

- UltraLite initialize database utility (ulinit), 210
- UltraLite load XML to database utility (ulload), 218
- UltraLiteJ database load utility (ULjLoad), 244
- b option
 - UltraLite database unload utility (ulunload), 227
 - UltraLiteJ database unload utility (ULjUnload), 246
- c option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite information utility (ulinfo), 209
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
 - UltraLiteJ database information utility (ULjInfo), 243
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
- d option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
- d1 option
 - UltraLite Interactive SQL utility (dbisql), 199
- E option
 - UltraLite load XML to database utility (ulload), 218
- e option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLite validate database utility (ulvalid) utility, 230
 - UltraLiteJ database unload utility (ULjUnload), 246
- f option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
- g option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
- h option
 - UltraLite SQL preprocessor utility (sqlpp), 204
- i option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLiteJ database load utility (ULjLoad), 244
- K option
 - UltraLite initialize database utility (ulinit), 210
- k option
 - UltraLite erase database utility (ulerase), 208
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
- l option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
- m option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
- n option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
- nearest_century option
 - UltraLite initialize database utility (ulinit), 210
- nogui option
 - UltraLite Interactive SQL utility (dbisql), 199

-
- o option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - onerror option
 - UltraLite Interactive SQL utility (dbisql), 199
 - p option
 - UltraLite erase database utility (ulerase), 208
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite synchronization utility (ulsync), 222
 - UltraLiteJ database information utility (ULjInfo), 243
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
 - q option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite erase database utility (ulerase), 208
 - UltraLite information utility (ulinfo), 209
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
 - UltraLiteJ database information utility (ULjInfo), 243
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
 - r option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLite synchronization utility (ulsync), 222
 - S option
 - UltraLite initialize database utility (ulinit), 210
 - s option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - SQL command
 - UltraLite Interactive SQL utility (dbisql), 199
 - t option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLiteJ database unload utility (ULjUnload), 246
 - u option
 - UltraLite erase database utility (ulerase), 208
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - ud option
 - UltraLite Engine utility (uleng12), 206
 - ul option
 - UltraLite Interactive SQL utility (dbisql), 199
 - v option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
 - UltraLiteJ database information utility (ULjInfo), 243
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
 - version
 - UltraLite Interactive SQL utility (dbisql), 199
 - w option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - x option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - y option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite load XML to database utility (ulload), 218
 - UltraLiteJ database load utility (ULjLoad), 244
 - UltraLiteJ database unload utility (ULjUnload), 246
 - Z option
 - UltraLite initialize database utility (ulinit), 210

- z option
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite SQL preprocessor utility (sqlpp), 204
- /* comment indicator
 - UltraLite about, 260
- // comment indicator
 - UltraLite about, 260
- 10054
 - UltraLite synchronization stream system errors, 104
- 130 error
 - SQL code for UltraLite schema upgrade, 123
- 256-bit strong encryption
 - UltraLite connection parameter for, 176
 - UltraLite usage of, 26
- ?
 - UltraLite input parameter, 286
- @data option
 - UltraLite database unload utility (ulunload), 227
 - UltraLite erase database utility (ulerase), 208
 - UltraLite information utility (ulinfo), 209
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
- ^
 - UltraLite bitwise operator, 298
- |
 - UltraLite bitwise operator, 298
- ~
 - UltraLite bitwise operator, 298
- A**
- ABS function
 - UltraLite syntax, 307
- ACOS function
 - UltraLite syntax, 307
- ActiveSync
 - deploy ActiveSync provider for UltraLite , 126
 - deploying MobiLink UltraLite applications, 86
 - registering applications for UltraLite clients, 127
 - UltraLite deploying provider files, 126
- ActiveSync provider installation utility (mlasinst)
 - registering applications for UltraLite clients, 127
- adding
 - UltraLite column methods, 48
 - UltraLite columns, 415
 - UltraLite indexes, 58
 - UltraLite users with connection parameters, 61
 - UltraLite users with SQL statements, 63
 - UltraLite users with Sybase Central, 64
- adding synchronization
 - UltraLite applications, 78
- adding UltraLite indexes
 - about, 57
- additional parameters
 - UltraLite synchronization parameter, 88
- administration tools
 - UltraLite troubleshooting, 479
- administration utilities
 - UltraLite utilities reference, 197
- AES encryption algorithm
 - UltraLite deployment steps, 13
 - UltraLite fips creation parameter, 148
 - UltraLite fips usage, 26
 - UltraLite usage, 26
- aggregate expressions
 - UltraLite SQL syntax, 284
- aggregate functions
 - UltraLite alphabetical list, 300
- aliases
 - UltraLite columns, 444
 - UltraLite equivalents, 267
- ALL search condition
 - UltraLite SQL, 290
- AllowDownloadDupRows
 - UltraLite synchronization parameter, 88
- allsync tables
 - UltraLite synchronizing tables, 75
- ALTER DATABASE SCHEMA FROM FILE statement
 - UltraLite schema changes impact, 45
 - UltraLite syntax, 412
 - usage, 123
- ALTER PUBLICATION statement
 - UltraLite syntax, 413
- ALTER SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 413
- ALTER TABLE statement
 - UltraLite Interactive SQL example, 49
 - UltraLite syntax, 415
- ALTER USER statement

-
- UltraLite syntax, 419
 - altering
 - UltraLite ALTER PUBLICATION statement, 413
 - UltraLite ALTER TABLE statement, 415
 - UltraLite ALTER USER statement, 419
 - UltraLite column methods, 49
 - UltraLite columns methods, 49
 - UltraLite table methods, 49
 - altering UltraLite column definitions
 - about, 49
 - ambiguous string to date conversions
 - UltraLite, 151
 - AND
 - UltraLite bitwise operators, 298
 - UltraLite logical operators, 289
 - ANY search condition
 - UltraLite SQL, 291
 - application profiling
 - UltraLite benchmark tests, 472
 - applications
 - (*see also* UltraLite applications)
 - ApplyFile method
 - UltraLite replacement for schema upgrade , 123
 - arc-cosine function
 - UltraLite ACOS function, 307
 - arc-sine function
 - UltraLite ASIN function, 309
 - arc-tangent function
 - UltraLite ATAN function, 310
 - architecture diagram
 - UltraLite, 1
 - ARGN function
 - UltraLite syntax, 308
 - arithmetic operators
 - UltraLite operators, 297
 - UltraLite SQL syntax, 297
 - articles
 - UltraLite copying method, 52
 - UltraLite databases, 76
 - UltraLite restrictions, 76
 - ASCII
 - UltraLite sorting, 23
 - UltraLite syntax, 309
 - ASCII files
 - UltraLite importing, 481
 - ASIN function
 - UltraLite syntax, 309
 - assembling parameters into connection strings
 - UltraLite about, 29
 - ATAN function
 - UltraLite syntax, 310
 - ATAN2 function
 - UltraLite syntax, 311
 - authenticating UltraLite users
 - about, 59
 - authentication
 - UltraLite users, 59
 - authentication parameters
 - UltraLite synchronization parameter, 90
 - authentication status
 - UltraLite synchronization parameter, 90
 - authentication value
 - UltraLite synchronization parameter, 92
 - autocommit
 - UltraLite transaction overview , 470
 - AUTOINCREMENT
 - UltraLite syntax, 426
 - UltraLiteJ syntax, 426
 - AvantGo (*see* M-Business Anywhere)
 - AvantGo M-Business Server (*see* M-Business Anywhere)
 - average function
 - UltraLite AVG function, 312
 - AVG function
 - UltraLite syntax, 312
- B**
- backing up (*see* backups)
 - backup
 - UltraLite about, 43
 - backups
 - UltraLite databases on Windows Mobile, 32
 - UltraLite internal mechanism, 43
 - UltraLite transaction overview , 470
 - base 10 logarithm
 - UltraLite LOG10 function, 353
 - benchmark testing
 - UltraLite characteristics of reliable tests, 475
 - UltraLite databases, 472
 - UltraLite methodology overview, 474
 - UltraLite SQL tests for, 473
 - UltraLite test creation, 475
 - UltraLite test execution and analysis, 475
 - UltraLite test preparation, 474
 - UltraLite types, 472

- BETWEEN search condition
 - UltraLite SQL, 292
- BIGINT data type
 - UltraLite, 265
- binary
 - UltraLite sorting, 23
- binary data types
 - UltraLite, 265
 - UltraLite maximum size, 8
- bitwise operators
 - UltraLite SQL syntax, 298
- BlackBerry
 - object store limitations, 12
 - ULjDbT utility, 248
 - ULjInfo utility, 243
 - ULjLoad utility, 244
 - ULjUnload utility, 246
 - UltraLite Java edition system tables, 251
 - utilities, 243
- blob_file_base_dir database option
 - UltraLite Java edition description, 238
- blob_file_base_dir property
 - UltraLite Java edition description, 235
- browsing
 - UltraLite table information, 51
 - UltraLite table methods, 51
- building
 - UltraLite CustDB application, 131
- bulk loading
 - UltraLite LOAD TABLE statement, 439
- BYTE_LENGTH function
 - UltraLite syntax, 313
- BYTE_SUBSTR function
 - UltraLite syntax, 313
- C**
- C++ APIs
 - (*see also* UltraLite C/C++ API)
- C++ applications
 - (*see also* UltraLite C/C++)
- cache
 - UltraLite cache_allocation property, 188
- cache allocation
 - UltraLite configuration, 192
- cache size
 - UltraLite performance, 455
 - UltraLite usage, 153
- cache_allocation database option
 - UltraLite, 192
- CACHE_MAX_SIZE connection parameter
 - UltraLite syntax, 167
- CACHE_MIN_SIZE connection parameter
 - UltraLite syntax, 168
- CACHE_SIZE connection parameter
 - UltraLite syntax, 169
- callback
 - UltraLite schema upgrade errors, 124
- cascading deletes
 - UltraLite limitations, 3
- cascading updates
 - UltraLite limitations, 3
- case creation parameter
 - UltraLite description, 141
- CASE expression
 - UltraLite NULLIF function, 365
 - UltraLite SQL syntax, 283
- case property
 - UltraLite description, 188
- case sensitivity
 - LIKE search condition, 296
 - UltraLite case creation parameter, 141
 - UltraLite case property, 188
 - UltraLite comparison operators, 288
 - UltraLite strings, 260
- case sensitivity considerations
 - UltraLite about, 141
- CAST function
 - UltraLite syntax, 314
- casting
 - UltraLite data types list, 268
- catalog
 - UltraLite system tables, 231
- CE_FILE connection parameter
 - UltraLite syntax, 171
- CEILING function
 - UltraLite syntax, 315
- central administration of remote databases
 - UltraLite databases, 18
- Certicom
 - UltraLite cryptographic module, 26
 - UltraLite TLS-enabled synchronization, 115
- certificates
 - UltraLite application access to encryption information, 223
- CHAR data type

- UltraLite, 265
- CHAR function
 - UltraLite syntax, 316
- CHAR_LENGTH function
 - UltraLite syntax, 317
- char_set property
 - UltraLite description, 188
- character functions
 - UltraLite alphabetical list, 305
- character set conversion
 - passwords, 432
- character sets
 - UltraLite char_set property, 188
 - UltraLite collation creation parameter, 143
 - UltraLite databases, 24
 - UltraLite on Windows, 24
 - UltraLite on Windows Mobile, 24
 - UltraLite strings, 260
- character strings
 - UltraLite embedded SQL, 204
- CHARINDEX function
 - UltraLite syntax, 318
- CHECK constraints
 - UltraLite limitations, 3
- CHECK CONSTRAINTS clause
 - UltraLite LOAD TABLE statement, 440
- CHECKPOINT statement
 - UltraLite syntax, 420
- checkpointing
 - UltraLite CHECKPOINT syntax, 420
- checkpoints
 - UltraLite performance optimization, 471
- CheckpointStore
 - UltraLite synchronization parameter, 88
- checksum_level creation parameter
 - UltraLite description, 142
- checksum_level property
 - UltraLite description, 188
- checksums
 - UltraLite checksum_level creation parameter, 142
 - UltraLite checksum_level property, 188
- choosing an index type
 - UltraLite about, 57
- client databases
 - UltraLite options, 87
- clients
 - UltraLite MobiLink clients, 67
- COALESCE function
 - UltraLite syntax, 319
- code points
 - UltraLite, 23
- CodeWarrior
 - UltraLite building CustDB application, 131
- collation creation parameter
 - UltraLite description, 143
- collation property
 - UltraLite description, 188
- collation sequences
 - UltraLite about, 23
 - UltraLite changing, 23
 - UltraLite LIKE search condition, 295
- collations
 - UltraLite collation creation parameter, 143
 - UltraLite CollationName property, 188
 - UltraLite unsupported, 217
- column compression
 - UltraLite SQL ALTER TABLE statement, 415
- column names
 - UltraLite SQL syntax, 282
- column names in expressions
 - UltraLite about, 282
- columns
 - syscolumn, UltraLite Java edition system table, 252
 - UltraLite adding methods, 48
 - UltraLite aliases, 444
 - UltraLite ALTER TABLE statement, 415
 - UltraLite altering methods, 49
 - UltraLite altering usage, 49
 - UltraLite copying method, 52
 - UltraLite limitations, 8
- comma-separated lists
 - UltraLite LIST function syntax, 350
- command line utilities
 - UltraLite database unload (ulunload) syntax, 227
 - UltraLite engine start (uleng12) syntax, 206
 - UltraLite engine stop (ulstop) syntax, 207
 - UltraLite erase database (ulerase) syntax, 208
 - UltraLite information (ulinfo) syntax, 209
 - UltraLite initialize database (ulinit) syntax, 210
 - UltraLite Interactive SQL (dbisql) syntax, 198
 - UltraLite load XML to database (ulload) syntax, 217
 - UltraLite SQL preprocessor (sqlpp) syntax, 204
 - UltraLite synchronization (ulsync) syntax, 221
 - UltraLite validate database (ulvalid) syntax, 230

- command prompts
 - Interactive SQL mode, 198
- comments
 - UltraLite syntax, 260
- commit flush
 - UltraLite configuration, 193
- COMMIT statement
 - UltraLite syntax, 420
- COMMIT_FLUSH connection parameter
 - UltraLite syntax, 172
- commit_flush_count database option
 - UltraLite, 193
- commit_flush_count property
 - UltraLite description, 188
- commit_flush_timeout database option
 - UltraLite, 194
- commit_flush_timeout property
 - UltraLite description, 188
- committing
 - UltraLite COMMIT syntax, 420
 - UltraLite database rows, 469
 - UltraLite transaction overview , 470
- comparing
 - UltraLite and SQL Anywhere databases, 3
- comparing UltraLite and SQL Anywhere
 - about, 3
- comparison operators
 - UltraLite dynamic SQL syntax, 288
 - UltraLite SQL, 288
- comparisons
 - UltraLite search conditions, 287
- compatibility
 - UltraLite SQL, 288
- compressed columns
 - UltraLite ALTER TABLE statement, 415
- computed columns
 - UltraLite limitations, 3
- COMPUTES clause
 - UltraLite LOAD TABLE statement, 440
- CON connection parameter
 - UltraLite syntax, 173
- concatenating strings
 - UltraLite string operators, 297
- concurrency
 - UltraLite issues, 468
 - UltraLite synchronization, 468
- conditions
 - UltraLite ALL conditions, 290
 - UltraLite ANY, 291
 - UltraLite BETWEEN, 292
 - UltraLite EXISTS, 292
 - UltraLite IN, 293
 - UltraLite search, 287
- connCount property
 - UltraLite description, 188
- connecting
 - MobiLink UltraLite Stream Type synchronization parameter, 105
 - UltraLite database troubleshooting, 30
- connecting to an UltraLite database
 - about, 29
- connection methods
 - UltraLite about, 29
- connection parameters
 - alphabetical list (UltraLite), 165
 - CE, 179
 - DBN for UltraLite, 176
 - iPhone, 179
 - Mac, 177
 - NT, 177
 - UltraLite, 29
 - UltraLite CACHE_MAX_SIZE, 167
 - UltraLite CACHE_MIN_SIZE, 168
 - UltraLite CACHE_SIZE, 169
 - UltraLite CE_FILE, 171
 - UltraLite choosing between, 31
 - UltraLite COMMIT_FLUSH, 172
 - UltraLite CON, 173
 - UltraLite connection summary, 29
 - UltraLite DBF, 174
 - UltraLite DBKEY, 176
 - UltraLite desktop, 177
 - UltraLite device, 179
 - UltraLite file_name, 174
 - UltraLite key, 176
 - UltraLite MIRROR_FILE, 180
 - UltraLite NT_FILE, 181
 - UltraLite overview, 29
 - UltraLite password, 183
 - UltraLite precedence of, 30
 - UltraLite PWD, 183
 - UltraLite RESERVE_SIZE , 184
 - UltraLite START , 185
 - UltraLite TEMP_DIR , 186
 - UltraLite UID , 187
 - UltraLite userid , 187

- connection strings
 - UltraLite connection parameters, 165
 - UltraLite parameters overview, 29
 - UltraLite setting , 29
- connections
 - UltraLite concurrency, 468
 - UltraLite connCount property, 188
 - UltraLite overview, 29
 - UltraLite troubleshooting, 479
- consolidated databases
 - UltraLite choosing, 2
 - UltraLite compatibility, 67
 - UltraLite sample, 136
- constants
 - UltraLite SQL syntax, 282
- constants in expressions
 - UltraLite about, 282
- constraints
 - UltraLite ALTER TABLE statement, 415
 - UltraLite referential integrity, 77
 - UltraLite renaming, 415
- controlling synchronization
 - UltraLite publications, 76
- conversion
 - UltraLite CAST, 314
- conversion functions
 - UltraLite alphabetical list, 300
- CONVERT function
 - UltraLite syntax, 319
- converting
 - UltraLite ambiguous dates, 151
 - UltraLite data types list, 268
- converting data types
 - UltraLite about, 268
- converting strings
 - UltraLite about, 305
- coordinated universal timestamp
 - UltraLite CURRENT UTC TIMESTAMP, 263
- copying
 - UltraLite table method, 52
- copying data
 - UltraLite databases, 52
- COS function
 - UltraLite syntax, 322
- cosine function
 - UltraLite COS function, 322
- COT function
 - UltraLite syntax, 323
- cotangent function
 - UltraLite COT function, 323
- COUNT function
 - UltraLite syntax, 323
- count operation
 - UltraLite execution plans, 466
- COUNT_UPLOAD_ROWS function
 - syntax, 324
- CPU
 - UltraLite limits, 8
- create database wizard
 - UltraLite usage, 17
- CREATE INDEX statement
 - UltraLite example, 58
 - UltraLite syntax, 421
 - UNIQUE parameter, 421
- create index wizard
 - UltraLite using, 57
- CREATE PUBLICATION statement
 - UltraLite Interactive SQL example, 83
 - UltraLite Interactive SQL subset example, 84
 - UltraLite Interactive SQL whole table example, 83
 - UltraLite syntax, 423
 - UltraLite usage, 77
- create publication wizard
 - UltraLite rows publishing, 84
 - UltraLite usage, 83
- CREATE SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 424
- CREATE TABLE statement
 - UltraLite syntax , 425
 - UltraLiteJ syntax , 425
- create table wizard
 - UltraLite usage, 47
- CREATE USER statement
 - UltraLite dynamic SQL syntax, 431
- create user wizard
 - UltraLite usage, 64
- creating
 - reference databases for UltraLite, 19
 - UltraLite CREATE PUBLICATION statement, 423
 - UltraLite CREATE TABLE statement, 425
 - UltraLite databases, 17
 - UltraLite databases at a command prompt, 17
 - UltraLite databases from a MobiLink sync model, 18

- UltraLite databases from XML, 20
 - UltraLite databases using central administration of remote databases, 18
 - UltraLite databases with ulinit, 210
 - UltraLite indexes, 58
 - UltraLite publications, 76
 - UltraLite publications tables, 83
 - UltraLite remote databases, 68
 - UltraLite table methods, 47
 - UltraLite users with connection parameters, 61
 - UltraLite users with SQL statements, 63
 - UltraLite users with Sybase Central, 64
 - creating databases
 - Sybase Central UltraLite plug-in, 17
 - UltraLite creation parameters, 139
 - UltraLite remote database , 68
 - creating publications
 - UltraLite databases with MobiLink applications, 76
 - creating remote databases
 - UltraLite clients, 68
 - creating UltraLite databases
 - about, 17
 - creating UltraLite tables
 - about, 47
 - creation parameters (UltraLite)
 - about, 139
 - case, 141
 - checksum_level, 142
 - collation, 143
 - date_format, 144
 - date_order, 147
 - fips creation parameter, 148
 - max_hash_size, 149
 - nearest_century, 151
 - obfuscate, 152
 - page_size, 153
 - precision, 155
 - scale , 156
 - time_format, 157
 - timestamp_format, 159
 - timestamp_increment, 162
 - utf8_encoding, 164
 - CROSS JOIN clause
 - UltraLite syntax, 436
 - cryptography
 - UltraLite Certicom module, 26
 - CSV files
 - UltraLite importing, 481
 - CURRENT DATE function
 - UltraLite TODAY function, 401
 - CURRENT DATE special value
 - UltraLite syntax, 262
 - current row
 - UltraLite concurrency, 468
 - CURRENT TIME special value
 - UltraLite syntax, 262
 - CURRENT TIMESTAMP special value
 - UltraLite feature comparison, 3
 - UltraLite syntax, 263
 - CURRENT UTC TIMESTAMP special value
 - UltraLite syntax, 263
 - CURRENT_TIMESTAMP special value
 - UltraLite syntax, 263
 - cursors
 - UltraLite current row, 468
 - UltraLite dirty reads, 41
 - CustDB
 - UltraLite application readme files, 130
 - UltraLite building application, 131
 - UltraLite limitations on instances of, 129
 - UltraLite running application, 131
 - UltraLite running multiple instances, 129
 - UltraLite sample , 129
 - UltraLite sample file locations, 130
 - UltraLite starting, 132
 - custdb.db
 - location of, 130
 - custdb.sql
 - calling synchronization scripts, 136
 - MobiLink project, 136
 - custdb.udb
 - UltraLite location of, 130
 - cycles
 - UltraLite foreign key issues, 77
 - UltraLite foreign keys, 77
- ## D
- daemon
 - running the UltraLite engine as a, 206
 - data
 - UltraLite loading, 217
 - UltraLite selecting rows, 444
 - UltraLite unloading, 227
 - UltraLite viewing methods, 51
 - UltraLiteJ selecting rows, 444

- data consistency
 - UltraLite dirty reads, 41
- data management
 - UltraLite components, 13
 - UltraLite description, 1
 - UltraLite state information, 468
- Data Manager
 - UltraLite database storage, 31
- data sources
 - UltraLite CustDB tutorial example, 134
- data type conversion functions
 - UltraLite about, 300
- data type conversions
 - UltraLite CAST, 314
- data types
 - spatial (UltraLite), 38
 - UltraLite about, 265
 - UltraLite alias equivalents, 267
 - UltraLite casting of values, 268
 - UltraLite limitations, 8
 - UltraLite special value, 261
 - UltraLite SQL conversion functions, 300
- data types in UltraLite
 - about, 265
- database files
 - (*see also* UltraLite databases)
 - UltraLite connection parameters, 31
 - UltraLite encrypting, 176
 - UltraLite maximum size, 8
- database management
 - UltraLite layers of, 1
- database objects
 - UltraLite copying method, 52
- database options
 - (*see also* database options (UltraLite))
 - UltraLite, 192
 - UltraLite Java edition, 237
- database options (UltraLite)
 - browsing, 34
 - cache_allocation, 192
 - commit_flush_count, 193
 - commit_flush_timeout, 194
 - DB_PROPERTY function, 334
 - global_database_id, 195
 - isolation_level, 196
 - ml_remote_id, 196
 - SET OPTION statement, 446
- database options window
 - accessing, 34
- database page size considerations
 - UltraLite about, 153
- database properties
 - (*see also* database properties (UltraLite))
 - UltraLite, 188
 - UltraLite Java edition, 235
- database properties (UltraLite)
 - browsing, 33
 - creation parameters, 22
 - fips usage, 26
 - obfuscate usage, 26
 - utf8_encoding usage, 24
- database schemas
 - UltraLite system tables, 231
- database sizes
 - UltraLite limit, 8
- database utilities
 - UltraLite database connections, 29
- database validation
 - UltraLite, 42
- database_name property
 - UltraLite Java edition description, 235
- databases
 - (*see also* UltraLite databases)
 - comparing UltraLite and SQL Anywhere, 3
 - creating with UltraLite plug-in, 17
 - UltraLite loading bulk data into, 439
- DATALENGTH function
 - UltraLite syntax, 325
- date considerations
 - UltraLite about, 144
- DATE data type
 - UltraLite, 265
- DATE function
 - UltraLite syntax, 326
- date functions
 - UltraLite alphabetical list, 300
- date parts
 - about, 301
 - available in UltraLite, 144
- date_format creation parameter
 - UltraLite description, 144
- date_format database option
 - UltraLite Java edition description, 238
- date_format property
 - UltraLite description, 188
 - UltraLite Java edition description, 235

- date_order creation parameter
 - UltraLite description, 147
- date_order database option
 - UltraLite Java edition description, 239
- date_order property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- DATEADD function
 - UltraLite syntax, 326
- DATEDIFF function
 - UltraLite syntax, 327
- DATEFORMAT function
 - UltraLite syntax, 329
- DATENAME function
 - UltraLite syntax, 329
- DATEPART function
 - UltraLite syntax, 330
- dates
 - query the current system date, 340
 - UltraLite ambiguous string conversions , 151
 - UltraLite conversion functions, 300
 - UltraLite date_format property, 188
 - UltraLite date_order property, 188
 - UltraLite formatting, 144
 - UltraLite Java edition date_format property, 235
 - UltraLite Java edition date_order property, 235
 - UltraLite ordering, 147
 - UltraLite rollover point, 151
- datetime
 - UltraLite conversion functions, 300
- DATETIME function
 - UltraLite syntax, 331
- DAY function
 - UltraLite syntax, 332
- day of week
 - UltraLite DOW function, 336
- DAYNAME function
 - UltraLite syntax, 332
- DAYS function
 - UltraLite syntax, 333
- DB_PROPERTY function
 - UltraLite syntax, 334
- DBF connection parameter
 - UltraLite syntax, 174
- dbisql utility
 - UltraLite exit codes, 203
 - UltraLite syntax, 198
 - UltraLite troubleshooting data imports, 481
- dbisql.com
 - about, 202
- dbisql.exe
 - about, 202
 - shutting down before installing, 479
- DBKEY connection parameter
 - UltraLite syntax, 176
- dblgen12.dll
 - ActiveSync conduit deployment in UltraLite, 126
- DBN connection parameter
 - UltraLite syntax, 176
- DDL
 - UltraLite schema changes with, 45
- deadlocks
 - UltraLite prevention of, 470
- DECIMAL data type
 - UltraLite, 265
- decimal point position considerations
 - UltraLite precision, 155
 - UltraLite scale creation parameter, 156
- decimal precision
 - UltraLite precision, 155
- decimals
 - UltraLite, 261
- declaring default global autoincrement columns
 - UltraLite clients in MobiLink systems, 71
- DEFAULT TIMESTAMP columns
 - UltraLite syntax, 426
 - UltraLiteJ syntax, 426
- default values
 - UltraLite CURRENT DATE, 262
 - UltraLite CURRENT TIME, 262
 - UltraLite CURRENT TIMESTAMP, 263
 - UltraLite CURRENT UTC TIMESTAMP, 263
 - UltraLite SQLCODE, 264
- defaults
 - UltraLite autoincrement, 426
 - UltraLiteJ autoincrement, 426
- DEFAULTS clause
 - UltraLite LOAD TABLE statement, 441
- DEGREES function
 - syntax, 335
- delaying commits
 - performance enhancements, 471
- DELETE statement
 - UltraLite dynamic SQL syntax, 432
 - UltraLiteJ dynamic SQL syntax, 432
- deleting

- UltraLite columns, 415
- UltraLite databases, 469
- UltraLite indexes, 58
- UltraLite publications, 84
- UltraLite START SYNCHRONIZATION
- DELETE statement, 446
- UltraLite table methods, 50
- UltraLite TRUNCATE TABLE statement, 450
- UltraLite users with connection parameters, 61
- UltraLite users with SQL statements, 63
- UltraLite users with Sybase Central, 64
- UltraLiteJ START SYNCHRONIZATION
- DELETE statement, 446
- UltraLiteJ TRUNCATE TABLE statement, 450
- deleting data
 - UltraLite file size impact of, 481
- DELIMITED BY clause
 - UltraLite LOAD TABLE statement, 441
- demos
 - (*see also* tutorials)
- deploying
 - applications that use ActiveSync for UltraLite clients, 86
 - changes to UltraLite database files, 115
 - UltraLite ActiveSync provider files, 126
 - UltraLite databases, 122
 - UltraLite engine troubleshooting, 483
 - UltraLite FIPS-enabled applications, 149
 - UltraLite schema upgrades, 123
 - UltraLite to devices, 115
 - UltraLite upgrades to devices, 115
- derived tables
 - UltraLite FROM clause, 436
 - UltraLite SQL, 285
- Designing synchronization in UltraLite
 - about, 72
- desktop connection parameter
 - UltraLite syntax, 177
- desktop creation
 - UltraLite about, 17
- device connection parameter
 - UltraLite syntax, 179
- devices
 - deploying UltraLite to, 115
 - UltraLite deployment techniques, 115
 - UltraLite multiple connection parameters for, 31
- diagnostics
 - UltraLite benchmark tests, 472
- DIFFERENCE function
 - UltraLite syntax, 336
- digits
 - UltraLite maximum number, 155
- direct page scans
 - UltraLite about, 464
- dirty reads
 - UltraLite isolation levels, 41
- disable concurrency
 - UltraLite synchronization parameter overview, 468
- DisableConcurrency
 - UltraLite synchronization parameter, 88
- DISTINCT keyword
 - UltraLite SQL, 444
- distinct operation
 - UltraLite execution plans, 466
- DOUBLE data type
 - UltraLite, 265
- double hyphen
 - UltraLite comment indicator, 260
- double slash
 - UltraLite comment indicator, 260
- DOW function
 - UltraLite syntax, 336
- download acknowledgements
 - UltraLite send_download_ack synchronization parameter, 102
- download only
 - UltraLite synchronization parameter, 92
- download only synchronization
 - UltraLite download_only synchronization parameter , 92
- download only tables
 - UltraLite non-synchronizing tables , 74
- download-only synchronization
 - UltraLite defining overview, 78
 - UltraLite synchronization parameter, 92
- download_only synchronization parameter
 - UltraLite reference, 92
- DROP INDEX statement
 - UltraLite Interactive SQL example, 59
 - UltraLite syntax, 433
 - UltraLiteJ syntax, 433
- DROP PUBLICATION statement
 - UltraLite Interactive SQL example, 85
 - UltraLite syntax, 434
- DROP SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 434

- UltraLiteJ syntax, 434
- DROP TABLE statement
 - UltraLite Interactive SQL example, 51
 - UltraLite syntax, 435
 - UltraLiteJ syntax, 435
- DROP USER statement
 - UltraLite syntax, 436
- dropping
 - UltraLite columns, 415
 - UltraLite DROP SYNCHRONIZATION PROFILE statement, 434
 - UltraLite DROP USER statement, 436
 - UltraLite indexes, 58
 - UltraLite SQL CREATE INDEX statement, 421
 - UltraLite SQL DROP INDEX statement, 433
 - UltraLite SQL DROP PUBLICATION statement, 434
 - UltraLite SQL DROP TABLE statement, 435
 - UltraLite table methods, 50
 - UltraLiteJ DROP SYNCHRONIZATION PROFILE statement, 434
 - UltraLiteJ SQL DROP INDEX statement, 433
 - UltraLiteJ SQL DROP TABLE statement, 435
- dropping an index
 - UltraLite about, 58
- dropping publications
 - UltraLite clients, 84
- dropping UltraLite tables
 - about, 50
- dummy operation
 - UltraLite execution plans, 466
- dynamic SQL
 - UltraLite arithmetic operators, 297
 - UltraLite bitwise operators, 298
 - UltraLite logical operators, 289
 - UltraLite operator precedence, 298
 - UltraLite string operators, 297
- E**
- editing
 - UltraLite table methods, 51
- ELSE
 - UltraLite CASE expression, 283
 - UltraLite IF expressions, 282
- embedded SQL
 - (*see also* UltraLite embedded SQL)
 - UltraLite NULL values, 261
 - embedded SQL library functions (*see* UltraLite embedded SQL library functions)
 - empty databases
 - UltraLite populating after running ulinit, 216
 - encoding
 - UltraLite utf8_encoding creation parameter, 164
 - UltraLite utf8_encoding usage, 24
 - ENCODING clause
 - UltraLite LOAD TABLE statement, 441
 - encryption
 - UltraLite changing key, 148
 - UltraLite data deployment considerations, 13
 - UltraLite deployment steps, 13
 - UltraLite encryption keys, 176
 - UltraLite encryption property, 188
 - UltraLite fips creation parameter, 148
 - UltraLite fips property usage, 26
 - UltraLite fips usage, 26
 - UltraLite obfuscate creation parameter, 152
 - UltraLite obfuscate property usage, 26
 - UltraLite synchronization deployment considerations, 115
 - UltraLite TLS synchronization configuration, 115
 - encryption keys
 - UltraLite changing, 148
 - encryption property
 - UltraLite description, 188
 - END
 - UltraLite CASE expression, 283
 - ENDIF
 - UltraLite IF expressions, 282
 - engines
 - (*see also* UltraLite engine)
 - entity-relationship diagrams
 - UltraLite about, 53
 - entity-relationship tab
 - UltraLite using, 53
 - environment variables
 - ERRORLEVEL for UltraLite, 203
 - UltraLite ULSQLCONNECT, 31
 - UltraLite usage, 29
 - ER diagram tab
 - UltraLite about, 53
 - erase database utility
 - UltraLite syntax, 208
 - error callback
 - UltraLite schema upgrade errors, 124
 - error codes

-
- UltraLite database synchronization utility (ulsync), 198
 - UltraLite load XML to database utility (uload), 198
 - UltraLite unload data to XML utility (ulunload), 198
 - UltraLite utilities, 198
 - ERRORLEVEL environment variable
 - Interactive SQL return code for UltraLite, 203
 - errors
 - UltraLite client error -764, 483
 - UltraLite `SQL_DATABASE_ERROR`, 480
 - UltraLite `SQL_DEVICE_ERROR`, 480
 - UltraLite `SQL_MEMORY_ERROR`, 480
 - escape sequences
 - UltraLite engine paths, 483
 - ESCAPES clause
 - UltraLite `LOAD TABLE` statement, 441
 - ESQL (*see* embedded SQL) (*see* UltraLite SQL)
 - event notifications
 - UltraLite working with, 34
 - examples
 - (*see also* samples)
 - (*see also* tutorials)
 - exclusive OR
 - UltraLite bitwise operator, 298
 - execution plans
 - UltraLite checking for index usage, 462
 - UltraLite how to read, 466
 - UltraLite operations, 466
 - UltraLite overriding, 462
 - UltraLite text of, 465
 - UltraLite working with, 462
 - EXISTS search condition
 - UltraLite search conditions, 287
 - UltraLite SQL, 292
 - exit codes
 - Interactive SQL utility (dbisql) for UltraLite, 203
 - UltraLite database synchronization utility (ulsync), 198
 - UltraLite load XML to database utility (uload), 198
 - UltraLite unload data to XML utility (ulunload), 198
 - EXP function
 - UltraLite syntax, 337
 - EXPLANATION function
 - UltraLite syntax, 338
 - exponential function
 - UltraLite EXP function, 337
 - exponents
 - UltraLite, 261
 - export tools
 - UltraLite ulunload utility, 227
 - exporting
 - UltraLite databases with ulunload, 227
 - expressions
 - UltraLite aggregate, 284
 - UltraLite CASE expressions, 283
 - UltraLite column names, 282
 - UltraLite constants, 282
 - UltraLite IF expressions, 282
 - UltraLite input parameters, 286
 - UltraLite SQL, 280
 - UltraLite SQL operator precedence, 298
 - UltraLite subqueries, 285
 - expressions in UltraLite
 - about, 280
- ## F
- failures
 - UltraLite preventing out-of-memory errors, 184
 - UltraLite schema upgrade errors, 124
 - features
 - UltraLite comparison list, 3
 - Federal Information Processing Standards (*see* FIPS)
 - fetches
 - UltraLite, 41
 - fetching rows
 - UltraLite concurrency, 41
 - file names
 - UltraLite connection parameters, 31
 - file property
 - UltraLite description, 188
 - file size
 - UltraLite database troubleshooting, 481
 - file systems
 - (*see also* VFS)
 - files
 - transferring with MLFileTransfer, 80
 - UltraLite ActiveSync provider, 126
 - UltraLite CustDB sample location, 130
 - filter operation
 - UltraLite execution plans, 466
 - FIPS
-

- UltraLite encrypted database deployment, 13
- UltraLite fips creation parameter, 148
- UltraLite fips property usage, 26
- UltraLite setup and deployment, 149
- UltraLite TLS-enabled synchronization, 115
- FIPS creation parameter
 - UltraLite description, 148
- fips database property
 - UltraLite deployment steps, 13
 - UltraLite usage, 26
- fips network protocol option
 - UltraLite TLS-enabled synchronization, 115
- FIRST clause
 - UltraLite SELECT statement, 444
 - UltraLiteJ SELECT statement, 444
- FLOAT data type
 - UltraLite, 265
- FLOOR function
 - UltraLite syntax, 339
- flush count
 - UltraLite formatting, 193
- flush timeout
 - UltraLite formatting, 194
- flushing
 - UltraLite databases, 193
- footprint
 - UltraLite databases, 3
- FOR clause
 - UltraLite SELECT statement, 445
- FOR READ ONLY clause
 - UltraLite direct page scans with, 464
- FORCE ORDER clause
 - UltraLite SELECT statement, 445
- foreign key columns
 - UltraLite Java edition system table for, 256
- foreign key cycles
 - UltraLite about, 77
 - UltraLite issues, 77
- foreign keys
 - UltraLite characteristics, 55
 - UltraLite copying method, 52
 - UltraLite foreign keys, 426
 - UltraLite Java edition system table for, 256
 - UltraLite of unnamed, 426
 - UltraLiteJ foreign keys, 426
 - UltraLiteJ of unnamed, 426
- FORMAT clause
 - UltraLite LOAD TABLE statement, 441
- format options (UltraLite)
 - utf8_encoding usage, 24
- FROM clause
 - UltraLite LOAD TABLE statement, 440
 - UltraLite SELECT statement, 444
 - UltraLite syntax, 436
- functions
 - return NULL if you specify NULL argument, 299
 - types of function for UltraLite, 299
 - UltraLite aggregate, 300
 - UltraLite data type conversion SQL, 300
 - UltraLite date and time, 300
 - UltraLite introduction, 299
 - UltraLite miscellaneous, 303
 - UltraLite numeric, 303
 - UltraLite string, 305
 - UltraLite system SQL, 306
- functions, aggregate
 - about, 300
 - UltraLite AVG, 312
 - UltraLite COUNT, 323
 - UltraLite LIST, 350
 - UltraLite MAX, 355
 - UltraLite MIN, 356
 - UltraLite SUM, 397
- functions, data type conversion
 - UltraLite about, 300
 - UltraLite CAST, 314
 - UltraLite CONVERT, 319
 - UltraLite HEXTOINT, 341
 - UltraLite INTTOHEX, 345
 - UltraLite ISDATE, 346
 - UltraLite ISNULL, 346
- functions, date and time
 - SWITCHOFFSET, 398
 - TODATETIMEOFFSET, 400
 - UltraLite about, 300
 - UltraLite DATE, 326
 - UltraLite DATEADD, 326
 - UltraLite DATEDIFF, 327
 - UltraLite DATEFORMAT, 329
 - UltraLite DATENAME, 329
 - UltraLite DATEPART, 330
 - UltraLite DATETIME, 331
 - UltraLite DAY, 332
 - UltraLite DAYNAME, 332
 - UltraLite DAYS, 333
 - UltraLite DOW, 336

- UltraLite GETDATE, 339
- UltraLite HOUR, 342
- UltraLite HOURS, 342
- UltraLite MINUTE, 356
- UltraLite MINUTES, 357
- UltraLite MONTH, 360
- UltraLite MONTHNAME, 361
- UltraLite MONTHS, 362
- UltraLite NOW, 364
- UltraLite QUARTER, 369
- UltraLite SECOND, 377
- UltraLite SECONDS, 378
- UltraLite TODAY, 401
- UltraLite WEEKS, 405
- UltraLite YEAR, 407
- UltraLite YEARS, 408
- UltraLite YMD, 409
- functions, miscellaneous
 - UltraLite about, 303
 - UltraLite ARGN, 308
 - UltraLite COALESCE, 319
 - UltraLite EXPLANATION, 338
 - UltraLite GREATER, 340
 - UltraLite IFNULL, 344
 - UltraLite LESSER, 349
 - UltraLite NEWID, 363
 - UltraLite NULLIF, 365
 - UltraLite SHORT_PLAN, 380
- functions, numeric
 - COUNT_UPLOAD_ROWS, 324
 - DEGREES, 335
 - RAND, 371
 - UltraLite about, 303
 - UltraLite ABS, 307
 - UltraLite ACOS, 307
 - UltraLite ASIN, 309
 - UltraLite ATAN, 310
 - UltraLite ATAN2, 311
 - UltraLite CEILING, 315
 - UltraLite COS, 322
 - UltraLite COT, 323
 - UltraLite EXP, 337
 - UltraLite FLOOR, 339
 - UltraLite LOG, 352
 - UltraLite LOG10, 353
 - UltraLite MOD, 360
 - UltraLite PI, 368
 - UltraLite POWER, 369
 - UltraLite RADIANS, 370
 - UltraLite REMAINDER, 372
 - UltraLite ROUND, 376
 - UltraLite SIGN, 380
 - UltraLite SIN, 382
 - UltraLite SQRT, 384
 - UltraLite TAN, 399
 - UltraLite TRUNCATE, 402
 - UltraLite TRUNCNUM, 402
 - UltraLiteJ RAND, 371
- functions, string
 - UltraLite about, 305
 - UltraLite ASCII, 309
 - UltraLite BYTE_LENGTH, 313
 - UltraLite BYTE_SUBSTR, 313
 - UltraLite CHAR, 316
 - UltraLite CHAR_LENGTH, 317
 - UltraLite CHARINDEX, 318
 - UltraLite DIFFERENCE, 336
 - UltraLite INSERTSTR, 345
 - UltraLite LCASE, 347
 - UltraLite LEFT, 348
 - UltraLite LENGTH, 349
 - UltraLite LOCATE, 351
 - UltraLite LOWER, 353
 - UltraLite LTRIM, 354
 - UltraLite PATINDEX, 366
 - UltraLite REPEAT, 373
 - UltraLite REPLACE, 373
 - UltraLite REPLICATE, 374
 - UltraLite RIGHT, 375
 - UltraLite RTRIM, 377
 - UltraLite SIMILAR, 381
 - UltraLite SOUNDEX, 383
 - UltraLite SPACE, 383
 - UltraLite STR, 392
 - UltraLite STRING, 393
 - UltraLite STRTOUUID, 394
 - UltraLite STUFF, 395
 - UltraLite SUBSTRING, 395
 - UltraLite TRIM, 401
 - UltraLite UCASE, 403
 - UltraLite UPPER, 404
 - UltraLite UUIDTOSTR, 404
- functions, system
 - DB_PROPERTY, 334
 - ML_GET_SERVER_NOTIFICATION, 359
 - SYNC_PROFILE_OPTION_VALUE, 398

- UltraLite DATALENGTH, 325
- fundamentals
 - UltraLite database management, 468
- G**
- GETDATE function
 - query the current system date, 340
 - UltraLite syntax, 339
- GetLastIdentity method
 - UltraLite synchronization, 70
- getScriptVersion method
 - UltraLite example, 111
- getStream method
 - UltraLite example, 105
- getUploadOK method
 - UltraLite example, 108
- global autoincrement
 - exhausted range in UltraLite, 69
 - UltraLite clients in MobiLink systems , 68
 - UltraLite global_database_id, 195
 - UltraLite, setting, 69
 - UltraLite, setting defaults, 71
- global database ID considerations
 - UltraLite about, 195
- global database identifier
 - UltraLite global_database_id, 195
 - UltraLite, setting, 69
- global_database_id option
 - UltraLite, 195
 - UltraLite CREATE TABLE statement, 426
 - UltraLite, setting, 69
 - UltraLiteJ CREATE TABLE statement, 426
- global_database_id property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- globally unique identifiers
 - UltraLite clients in MobiLink systems, 68
 - UltraLite SQL syntax for NEWID function, 363
- graphical plans
 - UltraLite not supported, 465
- GREATER function
 - UltraLite syntax, 340
- GROUP BY clause
 - UltraLite SELECT statement, 444
- group-by operation
 - UltraLite execution plans, 466
- GUIDs

- UltraLite clients in MobiLink systems, 68
- UltraLite SQL syntax for NEWID function, 363
- UltraLite SQL syntax for STRTOUUID function, 394
- UltraLite SQL syntax for UUIDTOSTR function, 404

H

- hash
 - UltraLite configuring size for, 461
 - UltraLite max_hash_size property, 188
 - UltraLite optimal size for, 459
 - UltraLite size considerations, 149
- hashing
 - UltraLite indexes, 149
- HAVING clause
 - UltraLite search conditions, 287
- hexadecimal strings
 - UltraLite about, 341
- HEXTOINT function
 - UltraLite syntax, 341
- host
 - UltraLite ULSynchronize arguments, 106
- HOUR function
 - UltraLite syntax, 342
- HOURS function
 - UltraLite syntax, 342

I

- IDENTIFIED BY password clause
 - ALTER USER statement, 419
- identifiers
 - UltraLite SQL, 259
- IDs
 - ml_remote_id option, 196
 - UltraLite global database, 195
- IF expressions
 - UltraLite SQL syntax, 282
- IF NOT EXISTS clause
 - CREATE PUBLICATION statement [UltraLite], 423
 - DROP INDEX statement [UltraLite] [UltraLiteJ], 433
 - DROP PUBLICATION statement [UltraLite] [UltraLiteJ], 433
 - DROP SYNCHRONIZATION PROFILE statement [UltraLite] [UltraLiteJ], 433

- DROP TABLE statement [UltraLite] [UltraLiteJ], 433
- UltraLite, CREATE TEXT INDEX statement, 421
- IFNULL function
 - UltraLite syntax, 344
- ignored rows
 - UltraLite synchronization parameter, 94
- ignored_rows synchronization parameter
 - UltraLite reference, 94
- importing data
 - UltraLite troubleshooting, 481
- importing data into databases
 - UltraLite ulload utility, 217
- IN search condition
 - UltraLite SQL, 293
- in-process runtime
 - UltraLite about, 13
- index performance considerations
 - UltraLite about, 149
- index scans
 - UltraLite, 457
- index-based UltraLite optimizations
 - UltraLite about, 457
- index-scan operation
 - UltraLite execution plans, 466
- indexes
 - sysindex, UltraLite Java edition system table, 253
 - sysindexcolumn, UltraLite Java edition system table, 254
 - UltraLite bypassing use of, 464
 - UltraLite copying method, 52
 - UltraLite creating, 58
 - UltraLite deleting, 58
 - UltraLite determining which index is scanned, 462
 - UltraLite hash considerations, 149
 - UltraLite hash value, 149
 - UltraLite introduction of, 457
 - UltraLite limitations, 8
 - UltraLite non-unique index characteristics, 57
 - UltraLite page_size usage, 153
 - UltraLite performance enhancements, 458
 - UltraLite primary keys, 19
 - UltraLite sysindex system table, 233
 - UltraLite sysixcol system table, 234
 - UltraLite types, 57
 - UltraLite unique index characteristics, 57
 - UltraLite unique key characteristics, 57
 - UltraLite UNIQUE SQL parameter, 421
 - UltraLite when to create, 57
 - UltraLite when to use, 56
 - UltraLite working with, 54
- indicators
 - UltraLite comments in SQL block, 260
- initializing
 - UltraLite databases with ulinit, 210
- initializing databases
 - Sybase Central UltraLite plug-in, 17
- INNER JOIN clause
 - UltraLite syntax, 436
- inner references
 - UltraLite subqueries, 285
- input parameters
 - UltraLite about, 286
- INPUT statement
 - UltraLite troubleshooting, 481
- INSERT statement
 - UltraLite Interactive SQL example, 53
 - UltraLite syntax, 438
 - UltraLiteJ syntax, 438
- inserting
 - UltraLite data using LOAD TABLE statement, 439
 - UltraLite INSERT statement, 438
 - UltraLite rows in bulk, 439
 - UltraLiteJ INSERT statement, 438
- INSERTSTR function
 - UltraLite syntax, 345
- installation
 - UltraLite troubleshooting, 479
- installing
 - UltraLite on devices, 115
- INTEGER data type
 - UltraLite, 265
- integrity
 - UltraLite CREATE TABLE statement, 426
 - UltraLiteJ CREATE TABLE statement, 426
- integrity constraints
 - UltraLite usage, 426
 - UltraLiteJ usage, 426
- Interactive SQL
 - command line, 202
 - UltraLite command line, 198
 - UltraLite displaying plans, 462
 - UltraLite plan interpretation, 466
 - UltraLite plan operations, 466
 - UltraLite text plans, 465
 - UltraLite troubleshooting data imports, 481

- Interactive SQL utility (dbisql)
 - UltraLite exit codes, 203
 - UltraLite syntax, 198
- INTTOHEX function
 - UltraLite syntax, 345
- iPad (*see* iOS)
- iPhone (*see* iOS)
- IS
 - UltraLite logical operators, 289
- ISDATE function
 - UltraLite syntax, 346
- ISNULL function
 - UltraLite syntax, 346
- isolation levels
 - UltraLite about, 40
 - UltraLite dirty reads, 41
 - UltraLite isolation levels, 40
- isolation_level option
 - UltraLite, 196
- J**
- join operation
 - UltraLite execution plans, 466
- joins
 - UltraLite FROM clause syntax, 436
- K**
- keep partial download synchronization parameter
 - UltraLite reference , 94
- key connection parameter
 - UltraLite syntax, 176
- KEY JOIN clause
 - UltraLite syntax, 436
- keys
 - UltraLite index creation from, 57
 - UltraLite index hash, 149
 - UltraLite primary, 47
- keyset operation
 - UltraLite execution plans, 466
- keywords
 - UltraLite SQL, 259
- L**
- LCASE function
 - UltraLite syntax, 347
- LEFT function
 - UltraLite syntax, 348
- LEFT OUTER JOIN clause
 - UltraLite syntax, 436
- LENGTH function
 - UltraLite syntax, 349
- LESSER function
 - UltraLite syntax, 349
- libraries
 - UltraLite choices, 13
 - UltraLite deploying uleng to Windows Mobile, 13
 - UltraLite FIPS-enabled applications, 149
- libulbase.lib
 - UltraLite C++ development, 115
- libulecc.lib
 - UltraLite C++ development, 115
- libulrsa.lib
 - UltraLite C++ development, 115
- libulrt.lib
 - UltraLite C++ development, 115
- LIKE search condition
 - UltraLite case-sensitivity, 295
 - UltraLite collations, 295
 - UltraLite pattern length, 294
 - UltraLite search conditions, 287
 - UltraLite SQL, 294
- like-scan operation
 - UltraLite execution plans, 466
- limitations
 - UltraLite, 8
 - UltraLite data types, 265
- limits
 - UltraLite, 8
- line length
 - UltraLite sqlpp utility output, 204
- linking
 - UltraLite runtime libraries, 14
- Linux
 - UltraLite runtime libraries for, 115
- LIST function
 - UltraLite syntax, 350
- lists
 - UltraLite LIST function syntax, 350
- literals
 - UltraLite constants, 282
- LOAD TABLE statement
 - UltraLite syntax, 439
- loading
 - UltraLite bulk inserts, 439
 - UltraLite LOAD TABLE statement, 439

-
- loading data
 - UltraLite LOAD TABLE statement, 439
 - loading databases
 - UltraLite databases with ulload, 217
 - LOCATE function
 - UltraLite syntax, 351
 - locking
 - UltraLite concurrency, 470
 - LOG function
 - UltraLite syntax, 352
 - LOG10 function
 - UltraLite syntax, 353
 - logging
 - UltraLite internal mechanism, 43
 - logic
 - UltraLite capturing for synchronization design, 72
 - logical operators
 - UltraLite SQL syntax, 289
 - lojoin operation
 - UltraLite execution plans, 466
 - LONG BINARY data type
 - UltraLite, 265
 - LONG VARCHAR data type
 - UltraLite, 265
 - LOWER function
 - UltraLite syntax, 353
 - lowercase strings
 - UltraLite LCASE function, 347
 - UltraLite LOWER function, 353
 - LTRIM function
 - UltraLite syntax, 354
- M**
- M-Business Anywhere
 - (*see also* UltraLite for M-Business Anywhere)
 - maintaining
 - UltraLite on devices, 115
 - maintaining primary key uniqueness
 - UltraLite clients in MobiLink systems, 68
 - management tools
 - UltraLite utilities reference, 197
 - managing databases
 - UltraLite data and state, 468
 - managing temporary tables
 - UltraLite about, 463
 - managing transactions
 - UltraLite increased throughput, 471
 - mathematical expressions
 - UltraLite arithmetic operators, 297
 - MAX function
 - UltraLite syntax, 355
 - max_hash_size creation parameter
 - UltraLite description, 149
 - max_hash_size property
 - UltraLite description, 188
 - maximum
 - UltraLite date ranges, 265
 - UltraLite physical limits, 8
 - maximum length
 - UltraLite user ID, 187
 - mechanism
 - UltraLite application and database deployment, 115
 - media failures
 - UltraLite transaction overview , 470
 - memory
 - UltraLite limits, 8
 - memory failures
 - UltraLite preventing, 184
 - memory usage
 - UltraLite database storage, 31
 - UltraLite indexes, 457
 - UltraLite row states, 469
 - UltraLite SQL performance, 473
 - metadata
 - UltraLite considering for reserve size, 184
 - MIN function
 - UltraLite syntax, 356
 - minimum
 - UltraLite date ranges, 265
 - MINUTE function
 - UltraLite syntax, 356
 - MINUTES function
 - UltraLite syntax, 357
 - MIRROR_FILE connection parameter
 - UltraLite syntax , 180
 - ml_add_connection_script system procedure
 - adding, 136
 - ml_add_table_script system procedure
 - adding, 136
 - ML_GET_SERVER_NOTIFICATION
 - syntax, 359
 - ml_remote_id option
 - UltraLite, 196
 - UltraLite property configuration, 209

- ml_remote_id property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- mlasdesk.dll
 - deploying UltraLite applications, 126
- mlasdev.dll
 - deploying UltraLite applications, 126
- mlasinst utility
 - registering applications for UltraLite clients, 127
 - UltraLite deploying with DLL files, 126
- MobiLink
 - UltraLite clients, 67
 - UltraLite CREATE PUBLICATION statement, 423
 - UltraLite SQL statements, 410
 - UltraLite user ID uniqueness, 196
- MobiLink ActiveSync provider installation utility (mlasinst)
 - registering applications for UltraLite clients, 127
- MobiLink clients
 - UltraLite progress counter, 2
- MobiLink file transfers
 - UltraLite client overview of, 80
- MobiLink project
 - UltraLite sample, 136
- MobiLink synchronization
 - setting timestamp_increment creation parameter, 162
 - UltraLite clients, 67
- MOD function
 - UltraLite syntax, 360
- modeling
 - UltraLite databases from MobiLink, 18
- modifying
 - UltraLite columns, 415
 - UltraLite tables, 415
- MONEY data type
 - UltraLite equivalent, 267
- monitoring synchronization
 - UltraLite observer synchronization parameter, 96
 - UltraLite setObserver method, 96
- MONTH function
 - UltraLite syntax, 360
- MONTHNAME function
 - UltraLite syntax, 361
- MONTHS function
 - UltraLite syntax, 362
- multi-table joins
 - UltraLite databases, 3
- multi-threaded
 - UltraLite applications, 468
- multiple databases
 - UltraLite maximum of, 468
- multiple devices
 - UltraLite connection parameters for, 31
- N**
- name property
 - UltraLite description, 188
- names
 - UltraLite column names, 282
- NATURAL JOIN clause
 - UltraLite syntax, 436
- nearest century conversion considerations
 - UltraLite about, 151
- nearest_century creation parameter
 - UltraLite description, 151
- nearest_century database option
 - UltraLite Java edition description, 239
- nearest_century property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- network protocols
 - UltraLite supported, 3
 - UltraLite Sync Result synchronization parameter, 107
 - UltraLite synchronization using HTTP, 105
 - UltraLite synchronization using HTTPS, 105
 - UltraLite synchronization using TCP/IP, 105
- new password
 - UltraLite synchronization parameter, 95
- NEWID function
 - UltraLite syntax, 363
- newmobilinkpwd synchronization parameter
 - UltraLite reference, 95
- Non-repeatable reads
 - UltraLite about, 41
- non-unique indexes
 - UltraLite characteristics, 55
 - UltraLite index creation from, 57
- nosync tables
 - UltraLite non-synchronizing tables , 73
- NOT
 - UltraLite bitwise operator, 298
 - UltraLite logical operators, 289

-
- NOW function
 - UltraLite syntax, 364
 - NT_FILE connection parameter
 - UltraLite syntax , 181
 - NULL
 - returned by functions if a NULL argument is specified, 299
 - UltraLite ISNULL function, 346
 - NULL values
 - UltraLite SQL, 261
 - NULLIF function
 - UltraLite about, 365
 - UltraLite using with CASE expressions, 284
 - number of authentication parameters
 - UltraLite synchronization parameter, 96
 - numbers
 - UltraLite SQL, 261
 - NUMERIC data type
 - UltraLite, 265
 - numeric functions
 - UltraLite alphabetical list, 303
 - numeric precision
 - UltraLite precision, 155
- O**
- obfuscate database creation parameter
 - UltraLite description, 152
 - obfuscate database property
 - UltraLite usage, 26
 - obfuscation
 - UltraLite usage, 26
 - observer synchronization parameter
 - UltraLite description, 96
 - ON phrase
 - UltraLite search conditions, 287
 - on-device creation
 - about, 21
 - operator precedence
 - UltraLite SQL syntax, 298
 - operators
 - UltraLite arithmetic operators, 297
 - UltraLite bitwise operators, 298
 - UltraLite comparison operators, 288
 - UltraLite logical operators, 289
 - UltraLite precedence of operators, 298
 - UltraLite SQL syntax, 296
 - UltraLite string operator, 297
 - optimization
 - UltraLite benchmark tests, 472
 - UltraLite checkpoints, 471
 - UltraLite execution plan access options, 462
 - UltraLite indexes, 459
 - UltraLite queries, 445
 - UltraLite SQL, 462
 - optimizer
 - (*see also* query optimizer)
 - UltraLite execution plan access options, 462
 - UltraLite impact of, 462
 - UltraLite overriding, 462
 - UltraLite plan interpretation, 466
 - UltraLite plan operations, 466
 - UltraLite using, 462
 - optimizing UltraLite
 - query performance, 455
 - options
 - UltraLite browsing, 34
 - UltraLite database unload utility (ulunload), 227
 - UltraLite erase database utility (ulerase), 208
 - UltraLite information utility (ulinfo), 209
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite Java edition, 237
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLite synchronization profiles, 224
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
 - options (UltraLite)
 - cache_allocation, 192
 - commit_flush_count, 193
 - commit_flush_timeout, 194
 - DB_PROPERTY function, 334
 - global_database_id, 195
 - isolation_level, 196
 - ml_remote_id, 196
 - SET OPTION statement, 446
 - OR
 - UltraLite bitwise operators, 298
 - UltraLite logical operators, 289
 - Oracle consolidated databases
 - UltraLite issues with, 67
 - ORDER BY clause
 - UltraLite SELECT statement, 445

- UltraLite troubleshooting with, 483
- order of operations
 - UltraLite SQL operator precedence, 298
- ordering result sets
 - UltraLite troubleshooting, 483
- outer references
 - UltraLite subqueries, 285
- overflow errors
 - AVG function, 312
 - SUM function, 397
- overhead
 - UltraLite considering for reserve size, 184

P

- packed rows
 - UltraLite about, 46
- packing rows
 - UltraLite effect of, 153
- page sizes
 - UltraLite Java edition page_size property, 235
 - UltraLite page_size property, 188
- page_size creation parameter
 - UltraLite description, 153
- page_size property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- pages
 - UltraLite size considerations, 153
- parameters
 - UltraLite connection list, 165
 - UltraLite connection overview, 29
 - UltraLite database creation, 139
 - UltraLite SQL input, 286
- partial download retained synchronization parameter
 - UltraLite reference, 97
- partition sizes
 - UltraLite defaults, choosing, 69
 - UltraLite defaults, overriding, 71
 - UltraLite exhausting defaults, 69
- partitioning
 - UltraLite primary keys, 69
 - UltraLite rows publishing, 83
- password
 - UltraLite synchronization parameter, 98
- Password connection parameter
 - UltraLite syntax, 183
- passwords

- character set conversion, 432
- maximum length, 432
- PWD UltraLite connection parameter, 183
- salt value for UltraLite connection parameter, 183
- UltraLite databases, 59
- UltraLite new MobiLink password parameter, 95
- UltraLite password synchronization parameter, 98

- paths
 - UltraLite connection parameters, 31
 - UltraLite engine, 479
- PATINDEX function
 - UltraLite syntax, 366
- pattern length
 - UltraLite LIKE search condition, 294
- pattern matching
 - UltraLite case-sensitivity, 295
 - UltraLite collations, 295
 - UltraLite LIKE search condition, 294
 - UltraLite PATINDEX function, 366
 - UltraLite pattern length, 294
- performance
 - separating commits from checkpoints, 471
 - UltraLite CACHE_MAX_SIZE connection parameter, 167
 - UltraLite CACHE_MIN_SIZE connection parameter, 168
 - UltraLite CACHE_SIZE connection parameter, 169
 - UltraLite choosing optimal hash size, 459
 - UltraLite download-only synchronization, 92
 - UltraLite index hashing, 149
 - UltraLite page sizes, 153
 - UltraLite preventing memory failures, 184
 - UltraLite query optimization, 445
 - UltraLite query optimization techniques, 455
 - UltraLite query tuning with index hashing, 458
 - UltraLite specifying FOR READ ONLY clause, 462
 - UltraLite upload only synchronization, 109
 - UltraLite using index, 56
 - UltraLite using index for applications, 19
 - UltraLite using index for improving query performance, 422
 - UltraLite using index for large tables, 54
 - UltraLite viewing an execution plan, 465
- persistent memory
 - UltraLite database storage, 31
- phantom rows

- UltraLite, 41
- physical limitations
 - UltraLite, 8
- PI function
 - UltraLite syntax, 368
- ping
 - UltraLite synchronization parameter, 99
- placeholder
 - UltraLite SQL input parameter, 286
- planning
 - UltraLite synchronization design overview, 72
- plans
 - UltraLite cursors, 338
 - UltraLite operations for, 466
 - UltraLite plan interpretation, 466
 - UltraLite plan operations, 466
 - UltraLite queries, overriding, 462
 - UltraLite queries, reading, 466
 - UltraLite queries, working with, 462
 - UltraLite syntax, 338
 - UltraLite text plans, 465
- platforms
 - UltraLite file storage, 31
 - UltraLite multiple connection parameters for, 31
- plug-ins
 - UltraLite troubleshooting, 479
- pools
 - UltraLite unused Global IDs, 69
- populating
 - UltraLite databases created with ulinit, 216
- port number
 - UltraLite ULSynchronize arguments, 106
- POWER function
 - UltraLite syntax, 369
- precedence
 - UltraLite SQL operator precedence, 298
- precision creation parameter
 - UltraLite description, 155
- precision database option
 - UltraLite Java edition description, 240
- precision property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- predicates
 - UltraLite ALL, 290
 - UltraLite ANY, 291
 - UltraLite BETWEEN, 292
 - UltraLite comparison operators, 288
 - UltraLite EXISTS, 292
 - UltraLite IN, 293
 - UltraLite LIKE search condition, 294
 - UltraLite SQL, 287
- prefix
 - UltraLite for Windows Mobile databases , 32
- prepared statements
 - UltraLite input parameters, 286
- primary key indexes
 - UltraLite bypassing use of, 464
- primary keys
 - UltraLite characteristics, 55
 - UltraLite generating unique values, 363
 - UltraLite generating unique values using UUIDs, 363
 - UltraLite indexing, 19
 - UltraLite integrity constraints, 426
 - UltraLite order of columns, 426
 - UltraLite table order, 77
 - UltraLite tables, 47
 - UltraLite troubleshooting data imports, 481
 - UltraLite UUIDs and GUIDs, 363
 - UltraLiteJ integrity constraints, 426
 - UltraLiteJ order of columns, 426
- procedures
 - UltraLite limitations, 3
- progress counter
 - UltraLite offset mismatches with , 2
- properties
 - DB_PROPERTY function, 334
 - UltraLite browsing, 33
 - UltraLite database creation parameters, 22
 - UltraLite databases, 188
 - UltraLite Java edition databases, 235
 - UltraLite Java edition system tables, 257
- properties (UltraLite)
 - DB_PROPERTY function, 334
- provider files
 - UltraLite deploying ActiveSync, 126
- providers
 - UltraLite ActiveSync files, 126
- public certificate
 - UltraLite application access to, 223
- publication creation wizard
 - UltraLite creating publications, 76
- publications
 - sysarticles, UltraLite Java edition system table, 255

- syspublications, UltraLite Java edition system table, 255
- UltraLite ALTER PUBLICATION statement, 413
- UltraLite copying method, 52
- UltraLite CREATE PUBLICATION statement, 423
- UltraLite design plans with, 72
- UltraLite dropping , 84
- UltraLite Java edition database schema description for, 255
- UltraLite Java edition table listing in schema , 255
- UltraLite limit, 210
- UltraLite publication synchronization parameter, 100
- UltraLite publishing overview, 76
- UltraLite publishing tables, 83
- UltraLite rows publishing, 83
- UltraLite schema description for, 234
- UltraLite SQL CREATE INDEX statement, 421
- UltraLite SQL DROP INDEX statement, 433
- UltraLite SQL DROP PUBLICATION statement, 434
- UltraLite SQL DROP TABLE statement, 435
- UltraLite synchronization, 100
- UltraLite synchronization parameter for, 100
- UltraLite sysarticle system table, 235
- UltraLite syspublication system table, 234
- UltraLite table listing in schema , 235
- UltraLite WHERE clause usage, 83
- UltraLite working with, 82
- UltraLiteJ SQL DROP INDEX statement, 433
- UltraLiteJ SQL DROP TABLE statement, 435
- publishing
 - UltraLite rows, 83
 - UltraLite tables, 83
 - UltraLite whole table, 76
- publishing whole tables
 - UltraLite, 83
- PWD connection parameter
 - UltraLite syntax, 183
- Q**
- QUARTER function
 - UltraLite syntax, 369
- queries
 - UltraLite optimization, 445
 - UltraLite troubleshooting unpredictable result sets, 483
- query optimization
 - (*see also* optimizer)
 - UltraLite SQL, 462
- query optimizer
 - (*see also* optimizer)
 - UltraLite, 462
- QUOTES clause
 - UltraLite LOAD TABLE statement, 441
- R**
- RADIANS function
 - UltraLite syntax, 370
- RAND function
 - syntax, 371
 - UltraLiteJ syntax, 371
- random numbers
 - RAND function, 371
 - UltraLiteJ RAND, 371
- range
 - UltraLite date type, 265
- read committed
 - UltraLite isolation levels, 40
- read uncommitted
 - UltraLite isolation levels, 40
- read-only tables
 - UltraLite databases, 78
 - UltraLite synchronizing, 78
- reading
 - UltraLite table rows, 41
- reading UltraLite access plans
 - about, 466
- readme files
 - UltraLite CustDB applications, 130
- REAL data type
 - UltraLite, 265
- recovery
 - UltraLite, 43
 - UltraLite about, 43
 - UltraLite introduction, 469
 - UltraLite transaction overview , 470
- recovery from system failure
 - UltraLite internal mechanism, 43
- reference databases
 - about, 19
 - creating for UltraLite, 19

- options for UltraLite, 19
- referential integrity
 - UltraLite databases, 3
 - UltraLite indexes, 56
 - UltraLite table order, 77
- registering
 - MobiLink UltraLite applications with ActiveSync, 127
- reload.sql
 - UltraLite loading, 439
- REMAINDER function
 - UltraLite syntax, 372
- remote databases
 - creating UltraLite clients, 68
 - UltraLite synchronization count, 2
- remote IDs
 - setting in UltraLite databases, 196
- remote servers
 - UltraLite CREATE TABLE statement, 425
- removing
 - UltraLite users with connection parameters, 61
 - UltraLite users with SQL statements, 63
 - UltraLite users with Sybase Central, 64
- removing data
 - UltraLite file size impact of, 481
- renaming
 - UltraLite database objects during upgrade, 123
 - UltraLite tables, 415
- REPEAT function
 - UltraLite syntax, 373
- REPLACE function
 - UltraLite syntax, 373
- REPLICATE function
 - UltraLite syntax, 374
- requests
 - UltraLite concurrency, 468
 - UltraLite management of, 468
- RESERVE_SIZE connection parameter
 - UltraLite syntax, 184
- reserved words
 - UltraLite SQL, 259
- restartable downloads
 - UltraLite keep partial download, 94
 - UltraLite partial download retained, 97
 - UltraLite resume partial download, 101
- restoring
 - UltraLite transaction overview , 470
- result sets
 - UltraLite troubleshooting unpredictable changes, 483
- resume partial download synchronization parameter
 - UltraLite reference, 101
- return codes
 - Interactive SQL utility (dbisql) for UltraLite, 203
- RIGHT function
 - UltraLite syntax, 375
- RIGHT OUTER JOIN clause
 - UltraLite syntax, 436
- role names
 - UltraLite foreign keys, 426
 - UltraLite role names, 426
 - UltraLiteJ foreign keys, 426
 - UltraLiteJ role names, 426
- ROLLBACK statement
 - UltraLite syntax, 443
 - UltraLiteJ syntax, 443
- rollbacks
 - UltraLite databases, 469
- rolling back
 - UltraLite transaction overview , 470
 - UltraLite transactions, 443
 - UltraLiteJ transactions, 443
- ROUND function
 - UltraLite syntax, 376
- rounding
 - UltraLite scale, 156
 - UltraLite scale creation parameter , 156
- row packing
 - UltraLite about, 46
 - UltraLite effect of, 153
 - UltraLite observing, 184
- rowlimit operation
 - UltraLite execution plans, 466
- rows
 - UltraLite deleting all rows from a table, 450
 - UltraLite fetching, 41
 - UltraLite INSERT statement, 438
 - UltraLite inserting in bulk, 439
 - UltraLite locking of, 470
 - UltraLite publishing, 83
 - UltraLiteJ deleting all rows from a table, 450
 - UltraLiteJ INSERT statement, 438
- RSA encryption algorithm
 - UltraLite TLS-enabled synchronization, 115
- RTRIM function
 - UltraLite syntax, 377

- running
 - UltraLite CustDB application, 131
- runtime libraries
 - UltraLite, 13
- runtimes
 - (*see also* UltraLite runtime)
- S**
- salt value for password
 - UltraLite syntax, 183
- sample applications
 - starting CustDB in UltraLite, 132
- samples
 - (*see also* examples)
 - (*see also* tutorials)
- scale creation parameter
 - UltraLite description, 156
- scale database option
 - UltraLite Java edition description, 240
- scale property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- scan operation
 - UltraLite execution plans, 466
- scanning
 - UltraLite database pages in queries, 464
 - UltraLite indexes in queries, 457
- schema
 - UltraLite catalog of tables for, 45
 - UltraLite changes, precautions, 45
- schemas
 - UltraLite SQL ALTER DATABASE SCHEMA FROM FILE syntax, 412
 - UltraLite system tables, 231
 - working with UltraLite, 45
- scjview
 - terminating before installing, 479
- script versions
 - UltraLite getScriptVersion, 111
 - UltraLite setScriptVersion method , 111
 - UltraLite version synchronization parameter , 111
- scripted upload
 - UltraLite CREATE PUBLICATION syntax, 423
- search conditions
 - UltraLite ALL, 290
 - UltraLite ANY, 291
 - UltraLite BETWEEN, 292
 - UltraLite EXISTS, 292
 - UltraLite IN, 293
 - UltraLite LIKE, 294
 - UltraLite SQL, 287
- SECOND function
 - UltraLite syntax, 377
- SECONDS function
 - UltraLite syntax, 378
- security
 - UltraLite, 26
 - UltraLite AES_FIPS database deployment, 13
 - UltraLite FIPS encryption, 148
 - UltraLite overview of, 26
 - UltraLite TLS-enabled synchronization, 115
- SELECT statement
 - UltraLite browsing data example, 51
 - UltraLite copying result sets with, 53
 - UltraLite syntax, 444
 - UltraLite troubleshooting result sets, 483
 - UltraLiteJ syntax, 444
- selecting
 - UltraLite SELECT statement, 444
 - UltraLiteJ SELECT statement, 444
- send column names
 - UltraLite synchronization parameter, 102
- send download acknowledgement
 - UltraLite synchronization parameter, 102
- send_column_names synchronization parameter
 - UltraLite reference, 102
- send_download_ack synchronization parameter
 - UltraLite reference, 102
- sensitivity
 - UltraLite case creation parameter, 141
- serialized transaction
 - UltraLite processing of, 470
- SET OPTION statement
 - UltraLite syntax, 446
- setObserver method
 - UltraLite example, 96
- setScriptVersion method
 - UltraLite example, 111
- setStream method
 - UltraLite example, 105
- setStreamParms method
 - UltraLite example, 106
- setting the global database identifier
 - UltraLite clients in MobiLink systems, 69
- setting the hash size

- about, 461
- setUserData method
 - UltraLite example, 109
- SHORT_PLAN function
 - UltraLite syntax, 380
- SIGN function
 - UltraLite syntax, 380
- SIMILAR function
 - UltraLite syntax, 381
- SIN function
 - UltraLite syntax, 382
- SKIP clause
 - UltraLite LOAD TABLE statement, 441
- slash-asterisk
 - UltraLite comment indicator, 260
- SMALLINT data type
 - UltraLite, 265
- SMALLMONEY data type
 - UltraLite equivalent, 267
- sort order
 - UltraLite collations, 24
- SOUNDEX function
 - UltraLite syntax, 383
- SPACE function
 - UltraLite syntax, 383
- spaces
 - UltraLite filepath definitions, 479
- spatial
 - list of supported functions (UltraLite), 39
- spatial data
 - about (UltraLite), 37
 - introduction (UltraLite), 37
 - standards compliance (UltraLite), 37
- spatial data types
 - ST_Geometry (UltraLite), 38
- spatial functions
 - ST_AsBinary, UltraLite, 385
 - ST_AsExtText, UltraLite, 385
 - ST_AsText, UltraLite, 386
 - ST_Distance, UltraLite, 386
 - ST_Equals, UltraLite, 387
 - ST_IntersectsRect, UltraLite, 387
 - ST_Point, UltraLite, 388
 - ST_PointFromExtText, UltraLite, 389
 - ST_PointFromText UltraLite, 389
 - ST_PointFromWKB, UltraLite, 390
 - ST_SRID, UltraLite, 391
 - ST_X, UltraLite, 391
 - ST_Y, UltraLite, 392
- Spatial SQL API
 - ST_Geometry type for UltraLite, 38
- special values
 - UltraLite CURRENT DATE, 262
 - UltraLite CURRENT TIME, 262
 - UltraLite CURRENT TIMESTAMP, 263
 - UltraLite CURRENT UTC TIMESTAMP, 263
 - UltraLite SQL, 261
 - UltraLite SQLCODE, 264
- SQL
 - (*see also* UltraLite SQL)
 - data types in UltraLite, 265
 - UltraLite comparison operators, 288
 - UltraLite expressions, 280
 - UltraLite identifiers, 259
 - UltraLite keywords, 259
 - UltraLite numbers, 261
 - UltraLite operators, 296
 - UltraLite reserved words, 259
 - UltraLite schema changes with, 45
 - UltraLite search conditions, 287
 - UltraLite statement types, 410
 - UltraLite strings, 260
 - UltraLite variables, 299
- SQL Anywhere databases
 - database comparison with UltraLite, 3
- SQL code
 - UltraLite schema upgrade errors, 124
- SQL Flagger
 - UltraLite usage, 204
- SQL functions
 - return NULL if you specify NULL argument, 299
 - types of function for UltraLite, 299
 - UltraLite aggregate, 300
 - UltraLite data type conversion, 300
 - UltraLite date and time, 300
 - UltraLite introduction, 299
 - UltraLite miscellaneous, 303
 - UltraLite numeric, 303
 - UltraLite string, 305
 - UltraLite system, 306
- SQL preprocessor utility (sqlpp)
 - UltraLite syntax , 204
- SQL standards
 - spatial data (UltraLite), 37
- SQL statements
 - alphabetical list of UltraLite statements, 410

- UltraLite, 410
 - SQL syntax
 - UltraLite CASE expression, 283
 - UltraLite column names, 282
 - UltraLite comments, 260
 - UltraLite constants, 282
 - UltraLite functions, 299
 - UltraLite IF expressions, 282
 - UltraLite input parameters, 286
 - UltraLite predicates, 287
 - UltraLite special values, 261
 - UltraLite SQLCODE special value, 264
 - SQL/MM standard
 - about (UltraLite), 37
 - SQLCODE
 - UltraLite concurrency checks, 470
 - UltraLite special value, 264
 - SQLCODE SQLE_LOCKED
 - UltraLite concurrency error, 470
 - SQLE_CONVERSION_ERROR
 - UltraLite upgrade warning, 124
 - SQLE_DATABASE_ERROR
 - UltraLite data corruption, 480
 - SQLE_DEVICE_ERROR
 - UltraLite data corruption, 480
 - SQLE_DOWNLOAD_CONFLICT error
 - UltraLite synchronization, 78
 - SQLE_MAX_ROW_SIZE_EXCEEDED
 - UltraLite error, 46
 - SQLE_MEMORY_ERROR
 - UltraLite data corruption, 480
 - SQLE_NOTFOUND
 - UltraLite concurrency error, 470
 - SQLE_ROW_DROPPED_DURING_SCHEMA_UPGRADE
 - UltraLite upgrade warning, 123
 - SQLE_UNABLE_TO_CONNECT_OR_START
 - troubleshooting, 479
 - sqlpp utility
 - UltraLite syntax, 204
 - SQRT function
 - UltraLite syntax, 384
 - square root function
 - UltraLite SQRT function, 384
 - ST_AsBinary
 - function, UltraLite, 385
 - ST_AsExtText
 - function, UltraLite, 385
 - ST_AsText
 - function, UltraLite, 386
 - ST_Distance
 - function, UltraLite, 386
 - ST_Equals
 - function, UltraLite, 387
 - ST_Geometry
 - type for UltraLite, 38
 - ST_IntersectsRect
 - function, UltraLite, 387
 - ST_Point
 - function, UltraLite, 388
 - ST_PointFromExtText
 - function, UltraLite, 389
 - ST_PointFromText
 - function, UltraLite, 389
 - ST_PointFromWKB
 - function, UltraLite, 390
 - ST_SRID
 - function, UltraLite, 391
 - ST_X
 - function, UltraLite, 391
 - ST_Y
 - function, UltraLite, 392
- START connection parameter
 - UltraLite for uleng12 on Windows Mobile, 13
 - UltraLite syntax, 185
 - START SYNCHRONIZATION DELETE statement
 - UltraLite syntax, 446
 - UltraLiteJ syntax, 446
 - state bytes
 - UltraLite databases, 469
 - state management
 - UltraLite overview, 468
 - statements
 - UltraLite, 410
 - UltraLite prepared, input parameters for, 286
 - UltraLite statement performance tuning, 473
 - UltraLite types, 410
 - STOP SYNCHRONIZATION DELETE statement
 - UltraLite syntax, 447
 - UltraLiteJ syntax, 447
 - storage
 - UltraLite limits, 8
 - UltraLite reserve size, 184
 - stored procedures
 - UltraLite limitations, 3
 - STR function

- UltraLite syntax, 392
- stream error
 - UltraLite synchronization parameter, 103
- stream parameters
 - UltraLite synchronization parameter, 106
- stream synchronization parameters
 - UltraLite synchronization parameter, 106
- stream type
 - UltraLite synchronization parameter, 105
- stream_error synchronization parameter
 - UltraLite reference, 103
 - UltraLite ul_stream_error structure, 103
- stream_parms synchronization parameter
 - UltraLite reference, 106
- STRING function
 - UltraLite syntax, 393
- string functions
 - UltraLite alphabetical list, 305
- string length
 - UltraLite LENGTH function, 349
- string literals
 - UltraLite constants, 282
- string operators
 - UltraLite dynamic SQL syntax, 297
- string position
 - UltraLite LOCATION function, 351
- strings
 - UltraLite case sensitivity, 260
 - UltraLite maximum size, 8
 - UltraLite nearest_century conversions to dates with , 151
 - UltraLite removing trailing blanks , 377
 - UltraLite replacing, 373
 - UltraLite SQL, 260
 - UltraLite SQL functions, 305
- STRIP clause
 - UltraLite LOAD TABLE statement, 442
- strong encryption
 - UltraLite deployment steps, 13
 - UltraLite fips creation parameter, 148
 - UltraLite usage, 26
- STRTOUUID function
 - UltraLite syntax, 394
- STUFF function
 - UltraLite syntax, 395
- subqueries
 - UltraLite SQL, 285
- subquery operation
 - UltraLite execution plans, 466
- SUBSCRIBE BY clause
 - UltraLite synchronization limitations, 76
- SUBSTR function
 - UltraLite syntax, 395
- SUBSTRING function
 - UltraLite syntax, 395
- substrings
 - UltraLite about, 395
 - UltraLite replacing, 373
- SUM function
 - UltraLite syntax, 397
- supplying UltraLite connection parameters
 - about, 29
- switches
 - UltraLite database unload utility (ulunload), 227
 - UltraLite erase database utility (ulerase), 208
 - UltraLite information utility (ulinfo), 209
 - UltraLite initialize database utility (ulinit), 210
 - UltraLite Interactive SQL utility (dbisql), 199
 - UltraLite load XML to database utility (ulload), 218
 - UltraLite SQL preprocessor utility (sqlpp), 204
 - UltraLite synchronization utility (ulsync), 222
 - UltraLite validate database utility (ulvalid) utility, 230
- SWITCHOFFSET function
 - syntax, 398
- Sybase Central
 - troubleshooting UltraLite connections, 479
- sync result
 - UltraLite synchronization parameter, 107
- SYNC_PROFILE_OPTION_VALUE function
 - syntax, 398
- synchronization
 - character sets in UltraLite, 23
 - setting timestamp_increment creation parameter, 162
 - ulsync utility for UltraLite databases, 221
 - UltraLite ALTER TABLE statement, 415
 - UltraLite application implementation , 78
 - UltraLite client-specific data, 75
 - UltraLite clients, 67
 - UltraLite CREATE TABLE statement, 426
 - UltraLite CustDB tutorial, 134
 - UltraLite design overview, 72
 - UltraLite download_only parameter, 92
 - UltraLite excluding tables with publications, 76

- UltraLite foreign keys, 77
- UltraLite ignored rows, 94
- UltraLite introduction , 67
- UltraLite M-Business Anywhere channel, 80
- UltraLite monitoring, 96
- UltraLite progress counting mechanism, 2
- UltraLite read-only tables , 78
- UltraLite referential integrity, 77
- UltraLite schema changes during, 45
- UltraLite SQLE_DOWNLOAD_CONFLICT error, 78
- UltraLite stopping, 96
- UltraLite system table, 231
- UltraLite task overview , 67
- UltraLite upload only parameter, 109
- UltraLiteJ CREATE TABLE statement, 426
- synchronization logic
 - browsing Sybase Central in UltraLite, 136
- synchronization models
 - UltraLite databases, 18
- synchronization parameters
 - UltraLite, 87
 - UltraLite additionalparms, 88
 - UltraLite Authentication Value, 92
 - UltraLite Disable Concurrency overview, 468
 - UltraLite download_only, 92
 - UltraLite getScriptVersion , 111
 - UltraLite getStream method , 105
 - UltraLite getUploadOK method, 108
 - UltraLite keep partial download, 94
 - UltraLite newmobilinkpwd, 95
 - UltraLite observer , 96
 - UltraLite partial download retained, 97
 - UltraLite password, 98
 - UltraLite ping, 99
 - UltraLite publication, 100
 - UltraLite required, 87
 - UltraLite resume partial download, 101
 - UltraLite send_column_names, 102
 - UltraLite send_download_ack , 102
 - UltraLite setObserver method, 96
 - UltraLite setScriptVersion method, 111
 - UltraLite setStream method , 105
 - UltraLite setStreamParms method, 106
 - UltraLite setUserData method, 109
 - UltraLite stream type , 105
 - UltraLite stream_error, 103
 - UltraLite stream_parms, 106
 - UltraLite Sync Result , 107
 - UltraLite upload_ok , 108
 - UltraLite upload_only, 109
 - UltraLite user_data, 109
 - UltraLite user_name, 110
 - UltraLite version , 111
- synchronization profile options
 - about, UltraLite clients, 224
- synchronization profiles
 - UltraLite ALTER SYNCHRONIZATION PROFILE statement, 413
 - UltraLite CREATE SYNCHRONIZATION PROFILE statement, 424
 - UltraLite DROP SYNCHRONIZATION PROFILE statement, 434
 - UltraLite SYNCHRONIZE statement, 448
 - UltraLiteJ DROP SYNCHRONIZATION PROFILE statement, 434
 - UltraLiteJ SYNCHRONIZE statement, 448
- synchronization scripts
 - browsing the UltraLite sample, 136
- synchronization stream parameters
 - UltraLite stream type , 105
- synchronization streams
 - UltraLite getStream method, 105
 - UltraLite setStream method, 105
 - UltraLite setStreamParms method, 106
 - UltraLite setting , 105
 - UltraLite stream synchronization parameter , 105
 - UltraLite stream_error synchronization parameter, 103
 - UltraLite stream_parms synchronization parameter, 106
 - UltraLite ULHTTPSSStream, 105
 - UltraLite ULHTTPStream, 105
- SYNCHRONIZE statement
 - UltraLite syntax, 448
 - UltraLiteJ syntax, 448
- synchronizing
 - UltraLite databases on Windows Mobile, 85
- synchronizing data
 - UltraLite file size impact of, 481
- syntax
 - UltraLite arithmetic operators, 297
 - UltraLite bitwise operators, 298
 - UltraLite CASE expression, 283
 - UltraLite column names, 282
 - UltraLite comparison operators, 288

- UltraLite constants, 282
- UltraLite CURRENT DATE special value, 262
- UltraLite CURRENT TIMESTAMP special value, 263
- UltraLite CURRENT UTC TIMESTAMP special value, 263
- UltraLite IF expressions, 282
- UltraLite logical operators, 289
- UltraLite predicates, 287
- UltraLite search conditions, 287
- UltraLite special values, 261
- UltraLite SQL comments, 260
- UltraLite SQL CURRENT TIME special value, 262
- UltraLite SQL functions, 299
- UltraLite SQL input parameters, 286
- UltraLite SQL operator precedence, 298
- UltraLite SQL operators, 296
- UltraLite SQLCODE special value, 264
- UltraLite string operators, 297
- SYS**
 - UltraLite system tables, 231
- sysarticle system table [UltraLite]
 - about, 235
- sysarticles system table [UltraLite Java edition databases]
 - about, 255
- syscolumn system table [UltraLite Java edition databases]
 - about, 252
- syscolumn system table [UltraLite]
 - about, 232
- sysfkcol system table [UltraLite Java edition databases]
 - about, 256
- sysforeignkey system table [UltraLite Java edition databases]
 - about, 256
- sysindex system table [UltraLite Java edition databases]
 - about, 253
- sysindex system table [UltraLite]
 - about, 233
- sysindexcolumn system table [UltraLite Java edition databases]
 - about, 254
- sysixcol system table [UltraLite]
 - about, 234
- syspublication system table [UltraLite]
 - about, 234
- syspublications system table [UltraLite Java edition databases]
 - about, 255
- systable system table [UltraLite Java edition databases]
 - about, 251
- systable system table [UltraLite]
 - about, 231
- system failure
 - UltraLite recover from, 43
- system failures
 - UltraLite transaction overview , 470
- system functions
 - UltraLite about, 306
 - UltraLite limitations, 3
- system tables
 - UltraLite about, 231
 - UltraLite browsing methods, 51
 - UltraLite Java edition sysarticles, 255
 - UltraLite Java edition syscolumn, 252
 - UltraLite Java edition sysfkcol, 256
 - UltraLite Java edition sysforeignkey, 256
 - UltraLite Java edition sysindex, 253
 - UltraLite Java edition sysindexcolumn, 254
 - UltraLite Java edition syspublications, 255
 - UltraLite Java edition systable, 251
 - UltraLite Java edition sysuldata, 257
 - UltraLite sysarticle, 235
 - UltraLite syscolumn, 232
 - UltraLite sysindex, 233
 - UltraLite sysixcol, 234
 - UltraLite syspublication, 234
 - UltraLite systable, 231
- system_error_code values
 - UltraLite synchronization stream errors, 104
- sysuldata system table [UltraLite Java edition databases]
 - about, 257

T

- table constraints
 - UltraLite adding, deleting, or modifying, 415
 - UltraLite CREATE TABLE statement, 426
 - UltraLiteJ CREATE TABLE statement, 426
- table expressions

- UltraLite subqueries, 285
- table owners
 - UltraLite, 259
- table size
 - number of rows, 8
 - UltraLite limit, 8
- TableOrder
 - UltraLite synchronization parameter, 88
- tables
 - UltraLite allsync, 75
 - UltraLite ALTER TABLE statement, 415
 - UltraLite altering methods, 49
 - UltraLite browsing methods, 51
 - UltraLite bulk loading, 439
 - UltraLite controlling synchronization with publications, 76
 - UltraLite copying method, 52
 - UltraLite copying methods, 52
 - UltraLite CREATE TABLE statement, 425
 - UltraLite creating methods, 47
 - UltraLite deleting methods, 50
 - UltraLite download only, 74
 - UltraLite editing methods, 51
 - UltraLite INSERT statement, 438
 - UltraLite limitations, 8
 - UltraLite nosync, 73
 - UltraLite order of, 77
 - UltraLite publications, 76
 - UltraLite rows publishing, 83
 - UltraLite size limit, 46
 - UltraLite temporary table usage, 465
 - UltraLite TRUNCATE TABLE statement, 450
 - UltraLite working with, 46
 - UltraLiteJ CREATE TABLE statement, 425
 - UltraLiteJ INSERT statement, 438
 - UltraLiteJ TRUNCATE TABLE statement, 450
- TAN function
 - UltraLite syntax, 399
- TCP/IP
 - (*see also* TCP/IP synchronization)
- temp operation
 - UltraLite execution plans, 466
- TEMP_DIR connection parameter
 - UltraLite syntax, 186
- temporary files
 - UltraLite limitation, 186
- temporary tables
 - UltraLite about, 463
 - UltraLite limitations, 8
 - UltraLite queries, 465
 - UltraLite synchronization using, 75
- terminating processes
 - UltraLite troubleshooting upgrades, 479
- TEXT data type
 - UltraLite equivalent, 267
- text plans
 - UltraLite viewing, 465
- THEN
 - UltraLite IF expressions, 282
- threads
 - UltraLite concurrency, 468
- time considerations
 - UltraLite about, 157
- TIME data type
 - UltraLite, 265
- time functions
 - UltraLite alphabetical list, 300
- time_format creation parameter
 - UltraLite description, 157
- time_format database option
 - UltraLite Java edition description, 241
- time_format property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- times
 - UltraLite conversion functions, 300
- TIMESTAMP data type
 - UltraLite, 265
 - UltraLite column limitations, 3
- TIMESTAMP special value
 - UltraLite column limitations, 3
 - UltraLite TIMESTAMP columns, 426
 - UltraLiteJ TIMESTAMP columns, 426
- timestamp_format creation parameter
 - UltraLite description, 159
- timestamp_format database option
 - UltraLite Java edition description, 241
- timestamp_format property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- timestamp_increment creation parameter
 - UltraLite description, 162
 - using in MobiLink synchronization, 162
- timestamp_increment database option
 - UltraLite Java edition description, 242
- timestamp_increment property

- UltraLite description, 188
- UltraLite Java edition description, 235
- timestamp_with_time_zone property
 - UltraLite description, 188
 - UltraLite Java edition description, 235
- timestamp_with_time_zone_format creation parameter
 - UltraLite, 163
- timestamp_with_time_zone_format database option
 - UltraLite Java edition description, 242
- timestamps
 - UltraLite timestamp_format creation parameter, 159
 - UltraLite timestamp_increment creation parameter, 162
- TINYINT data type
 - UltraLite, 265
- TLS
 - UltraLite client configuration, 115
 - UltraLite synchronization deployment considerations, 115
- TLS_TYPE network protocol option
 - UltraLite TLS-enabled synchronization, 115
- TODATETIMEOFFSET function
 - syntax, 400
- TODAY function
 - UltraLite syntax, 401
- TOP clause
 - UltraLite SELECT statement, 444
 - UltraLiteJ SELECT statement, 444
- tracking
 - UltraLite synchronization, 2
- transaction logging
 - UltraLite internal mechanism, 43
- transaction management
 - UltraLite, 468
- transaction processing
 - checkpoints and commits, 471
 - UltraLite COMMIT_FLUSH connection parameter, 172
 - UltraLite commit_flush_count option, 193
 - UltraLite commit_flush_timeout option, 194
- transactions
 - UltraLite checkpoints for, 420
 - UltraLite COMMIT statement, 420
 - UltraLite concurrency, 468
 - UltraLite databases, 469
 - UltraLite flushing, 193
 - UltraLite isolation level, 41
 - UltraLite rolling back, 443
 - UltraLite schema changes impact, 45
 - UltraLiteJ rolling back, 443
- transferring files
 - UltraLite files with MLFileTransfer, 80
- transport-layer security
 - (*see also* TLS)
- triggers
 - UltraLite limitations, 3
- TRIM function
 - UltraLite syntax, 401
- troubleshooting
 - avoiding synchronization issues with foreign key cycles, 77
 - UltraLite backing up application, 67
 - UltraLite checksum failures, 142
 - UltraLite connection parameter precedence, 30
 - UltraLite getUploadOK method , 108
 - UltraLite global ID numbers, 69
 - UltraLite implementing resumable downloads, 79
 - UltraLite maintaining timestamps and timestamp increments, 160
 - UltraLite ping synchronization parameter, 99
 - UltraLite retrieving GLOBAL AUTOINCREMENT value, 70
 - UltraLite Stream Error synchronization parameter, 103
 - UltraLite Sync Result synchronization parameter , 107
 - UltraLite upload_ok synchronization parameter, 108
- TRUNCATE function
 - UltraLite syntax, 402
- TRUNCATE TABLE statement
 - UltraLite syntax, 450
 - UltraLiteJ syntax, 450
- truncating
 - UltraLite tables, 450
 - UltraLiteJ tables, 450
- TRUNCNUM function
 - UltraLite syntax, 402
- trusted certificates
 - UltraLite application access to encryption information, 223
- tuning performance
 - UltraLite max_hash_size, 149
 - UltraLite with index hashing, 458

tutorials

- UltraLite CustDB, 129
- UltraLite CustDB database synchronization, 134

U

UCASE function

- UltraLite syntax, 403

UDB

- (*see also* UltraLite databases)

UID connection parameter

- UltraLite syntax, 187

ul_stream_error structure

- UltraLite example, 103

UL_SYNC_ALL macro

- UltraLite publications list, 100

UL_SYNC_ALL_PUBS macro

- UltraLite publications list, 100

uleng12 utility

- ud option, 206
- deploying to Windows Mobile, 13
- in-process database support, 206
- syntax, 206

ulerase utility

- syntax, 208

ULHTTPSSStream function [UL ESQL]

- UltraLite synchronization stream, 105

ULHTTPStream function [UL ESQL]

- UltraLite synchronization stream, 105

ulinfo utility

- syntax, 209

ulinit utility

- syntax, 210
- unsupported collation workaround, 217
- using, 17

ULjDbT utility

- syntax, 248

ULjInfo utility

- syntax, 243

ULjLoad utility

- syntax, 244

ULjUnload utility

- syntax, 246

ulload utility

- syntax, 217
- using, 20

ULSocketStream function

- UltraLite synchronization stream, 105

ULSQLCONNECT environment variable

- description, 31

ulstop utility

- syntax, 207

ulsync utility

- synchronization profile options, 224
- syntax, 221

UltraLite

- (*see also* UltraLite APIs)

- (*see also* UltraLite applications)

- (*see also* UltraLite databases)

- (*see also* UltraLite embedded SQL)

- (*see also* UltraLite SQL)

- (*see also* UltraLite utilities)

- about, 1

- architecture, 1

- collations unsupported, 217

- concurrency, 468

- CustDB sample, 136

- data conversion, 300

- database creation parameters, 139

- deploying applications and databases, 115

- deploying databases, 122

- engine and runtime for, 13

- engine starting, 206

- engine stopping, 207

- environment variables, 29

- error codes, 198

- installing applications and databases, 115

- spatial function, ST_AsBinary, 385

- spatial function, ST_AsExtText, 385

- spatial function, ST_AsText, 386

- spatial function, ST_Distance, 386

- spatial function, ST_Equals, 387

- spatial function, ST_IntersectsRect, 387

- spatial function, ST_Point, 388

- spatial function, ST_PointFromExtText, 389

- spatial function, ST_PointFromText, 389

- spatial function, ST_PointFromWKB, 390

- spatial function, ST_SRID, 391

- spatial function, ST_X, 391

- spatial function, ST_Y, 392

- SQL functions, aggregate, 300

- SQL functions, data type conversion, 300

- SQL functions, types of, 299

- SQL statement reference, 410

- SQLE_MAX_ROW_SIZE_EXCEEDED error, 46

- supported network protocols, 3

- system functions, 306
- table owners, 259
- timestamp_with_time_zone_format creation parameter, 163
- troubleshooting, 479
- users, 59
- UTF8BIN collation for UNICODE characters, 24
- utilities reference , 197
- Windows Mobile considerations, 12
- UltraLite administration tools
 - troubleshooting, 482
- UltraLite APIs (*see* UltraLite C/C++ API) (*see* UltraLite for AppForge API) (*see* UltraLite for M-Business Anywhere API) (*see* UltraLite.NET API)
- UltraLite applications
 - (*see also* UltraLite C/C++ API)
 - (*see also* UltraLite for AppForge API)
 - (*see also* UltraLite for M-Business Anywhere API)
 - (*see also* UltraLite.NET API)
 - bench marking, 472
 - building CustDB, 131
 - concurrency, 468
 - CustDB applications and readme files, 130
 - defining a temporary file location at connection time, 186
 - defining engine location at connection time, 185
 - deploying ActiveSync provider files, 126
 - deploying FIPS-enabled, 149
 - deploying to devices, 13
 - error -764, 483
 - managing multiple requests, 468
 - public certificate access, 223
 - synchronization , 67
 - TLS-enabled synchronization, 115
 - transferring files with MobiLink, 80
 - troubleshooting
 - SQLE_UNABLE_TO_CONNECT_OR_START, 479
 - Windows Mobile APIs, 13
- UltraLite benchmark testing
 - about, 472
- UltraLite C/C++
 - building CustDB application, 131
 - CustDB sample and readme files, 131
 - engine and runtime support, 13
- UltraLite client synchronization
 - parameters and options about, 87
- UltraLite clients
 - about MobiLink, 67
- UltraLite columns
 - renaming during schema upgrade, 123
- UltraLite connection parameters
 - about, 29
 - list, 165
- UltraLite connections
 - troubleshooting, 479
- UltraLite creation parameters
 - about, 139
- UltraLite database management system
 - about, 1
- UltraLite database properties
 - about, 188
- UltraLite database synchronization utility (ulsync)
 - synchronization profile options, 224
- UltraLite databases
 - articles in, 235
 - backing up on Windows Mobile, 32
 - benchmark testing, 472
 - checkpoint usage, 471
 - collation sequences, 23
 - columns adding, 48
 - columns altering, 49
 - columns of tables in, 234
 - connection overview, 29
 - connection parameter list, 165
 - connection parameters overview, 29
 - connections overview, 468
 - counting synchronizations, 2
 - create with the create database wizard, 17
 - creating, 17
 - creating at a command prompt, 17
 - creating from SQL Anywhere reference database, 19
 - creating from XML, 20
 - creation parameters, 22
 - data and state management, 468
 - database comparison with SQL Anywhere, 3
 - database integrity, 469
 - deadlocks, 470
 - deploying for data encryption, 13
 - deploying for synchronization encryption, 115
 - deploying to devices, 115
 - encrypting with the fips creation parameter, 148
 - entity-relationship diagrams, 53
 - erasing, 208
 - fetching rows, 41

- file path definition, 31
- indexes creating, 421
- indexes hashing, 149
- indexes types of, 57
- indexes working with, 54
- indexes, when to create, 57
- indexing primary keys, 19
- initializing at a command prompt, 19
- initializing from Sybase Central, 17
- isolation levels, 41
- maintaining row state, 469
- management fundamentals, 468
- managing multiple, 468
- managing users, 61, 63, 64
- memory usage, 469
- modeling from MobiLink, 18
- multi-table joins, 3
- objects copy method, 52
- Oracle as reference database for, 67
- page_size creation parameter, 153
- populating after running ulinit, 216
- properties browsing, 33
- publications about, 82
- publications dropping, 84
- publications in, 234
- publishing rows, 83
- publishing tables, 83
- recovery, 469
- renaming objects during schema upgrade, 123
- requests overview, 468
- rollbacks, 469
- row fetching, 41
- row locking, 470
- row packing, 153
- row packing introduction, 46
- row size limit, 46
- rows deleting, 469
- schema, 231
- schema changes, 45
- schema overview, 45
- security overview, 26
- size reduction, 481
- state bytes used, 469
- storing indexes in, 233
- supported index types, 55
- synchronization introduction, 468
- synchronization profile options, 224
- synchronization utility (ulsync) syntax, 221
- system failure, backups, and recovery, 43
- tables copying, 52
- tables creating, 47
- tables dropping, 50
- tables, browsing, 51
- temporary files, 186
- temporary table management, 463
- threads overview, 468
- transactions overview, 468
- troubleshooting connections, 479
- UltraLite file property, 188
- UltraLite name property, 188
- UltraLite schema upgrades reference, 412
- UltraLite size limit, 8
- unique keys, 57
- upgrade previous versions, 17
- upgrading schema on devices, 123
- user IDs and passwords, 59
- utilities reference , 197
- validating, 42
- viewing option settings, 34
- Windows Mobile file paths, 32
- UltraLite embedded SQL
 - (*see also* UltraLite embedded SQL library functions)
 - character strings, 204
 - line numbers, 204
 - preprocessor, 204
- UltraLite engine
 - about, 13
 - concurrency in, 468
 - defining a temporary file location at connection time, 186
 - defining executable location at connection time, 185
 - erase database utility, 208
 - error codes, 198
 - running as a daemon, 206
 - start utility, 206
 - stop utility, 207
 - troubleshooting, 479
 - troubleshooting error -764, 483
 - Windows Mobile deployment, 13
- UltraLite engine start utility
 - syntax, 206
- UltraLite engine stop utility (ulstop)
 - syntax, 207
- UltraLite erase database utility

-
- syntax, 208
 - UltraLite for M-Business Anywhere
 - building CustDB application, 132
 - UltraLite CustDB sample and readme files, 132
 - UltraLite information utility (ulinfo)
 - syntax, 209
 - UltraLite initialize database utility (ulinit)
 - syntax, 210
 - UltraLite Java edition
 - about, 1
 - UltraLite Java edition database options
 - about, 237
 - UltraLite Java edition database properties
 - about, 235
 - UltraLite Java edition databases
 - (*see also* UltraLite Java edition management system)
 - (*see also* UltraLiteJ)
 - columns, 252
 - describing columns of tables in, 254
 - describing foreign key columns in, 256
 - describing foreign keys in, 256
 - describing publications in, 255
 - describing tables in, 251
 - limitations, 8
 - storing indexes in, 253
 - storing properties of, 257
 - storing publications in, 255
 - ULjDbT utility, 248
 - ULjInfo utility, 243
 - ULjLoad utility, 244
 - ULjUnload utility, 246
 - utilities, 243
 - UltraLite Java edition management system
 - (*see also* UltraLite Java edition databases)
 - (*see also* UltraLiteJ)
 - UltraLite load XML to database utility
 - syntax, 217
 - UltraLite M-Business Anywhere (*see* UltraLite for M-Business Anywhere)
 - UltraLite optimizer
 - execution plan access options, 462
 - UltraLite passwords
 - about, 59
 - UltraLite plug-in
 - troubleshooting, 479
 - UltraLite publications
 - renaming during schema upgrade, 123
 - UltraLite runtime
 - about, 13
 - concurrency in, 468
 - UltraLite runtime libraries
 - Linux, 115
 - Windows, 115
 - Windows Mobile, 115
 - UltraLite schema
 - upgrading on device, 123
 - UltraLite SQL
 - (*see also* UltraLite embedded SQL)
 - comma-separated lists, 350
 - comments, 260
 - data types, 265
 - dates, 264
 - execution plans for, 462
 - expressions, 280
 - identifiers, 259
 - index-based optimizations, 457
 - keywords, 259
 - NULL values, 261
 - numbers, 261
 - operators, 296
 - page-based optimizations, 464
 - special values, 261
 - SQL functions date and time, 300
 - SQL functions, miscellaneous, 303
 - SQL functions, numeric, 303
 - SQL functions, string, 305
 - strings, 260
 - times, 264
 - troubleshooting queries, 483
 - tuning statement performance, 473
 - variables, 299
 - UltraLite SQL statements
 - about, 410
 - ALTER DATABASE SCHEMA FROM FILE statement syntax, 412
 - ALTER PUBLICATION statement syntax, 413
 - ALTER SYNCHRONIZATION PROFILE statement syntax, 413
 - ALTER TABLE statement syntax, 415
 - ALTER USER statement syntax, 419
 - categories of, 410
 - CHECKPOINT statement syntax, 420
 - COMMIT statement syntax, 420
 - CREATE INDEX statement syntax, 421
 - CREATE PUBLICATION statement syntax, 423

- CREATE SYNCHRONIZATION PROFILE statement syntax, 424
- CREATE TABLE statement syntax, 425
- CREATE USER statement syntax, 431
- DELETE statement syntax, 432
- DROP INDEX statement syntax, 433
- FROM clause, 436
- INSERT statement, 438
- LOAD TABLE statement syntax, 439
- ROLLBACK statement syntax, 443
- SELECT statement syntax, 444
- SET OPTION statement syntax, 446
- START SYNCHRONIZATION DELETE statement syntax, 446
- STOP SYNCHRONIZATION DELETE statement syntax, 447
- SYNCHRONIZE statement syntax, 448
- TRUNCATE TABLE syntax, 450
- UltraLite DROP TABLE statement syntax, 435
- UltraLite DROP USER statement syntax, 436
- UNION operation syntax, 451
- UNIQUE parameter, 421
- UPDATE statement syntax, 452
- UltraLite synchronization
 - about, 67
 - remote IDs and user IDs, 196
- UltraLite synchronization utility (ulsync) syntax, 221
- UltraLite system table reference
 - about, 231
- UltraLite system tables
 - about, 231
- UltraLite tables
 - excluding from creation process, 210
 - renaming during schema upgrade, 123
- UltraLite temporary files
 - about, 186
- UltraLite temporary tables
 - managing, 463
- UltraLite unload database to XML utility (ulunload) syntax, 227
- UltraLite user IDs
 - about, 59
 - MobiLink uniqueness, 196
- UltraLite utilities
 - about, 197
 - UltraLite database unload (ulunload) syntax, 227
 - UltraLite engine (uleng12) syntax, 206
 - UltraLite engine start (uleng12) syntax, 206
 - UltraLite engine stop (ulstop) syntax, 207
 - UltraLite erase database (ulerase) syntax, 208
 - UltraLite information (ulinfo) syntax, 209
 - UltraLite initialize database (ulinit) syntax, 210
 - UltraLite Interactive SQL (dbisql) syntax, 198
 - UltraLite load XML to database (ulload) syntax, 217
 - UltraLite SQL preprocessor (sqlpp) syntax, 204
 - UltraLite synchronization (ulsync) syntax, 221
 - UltraLite validate database (ulvalid) syntax, 230
- UltraLite validate database utility (ulvalid)
 - about, 230
- UltraLite.NET
 - building CustDB application, 132
 - CustDB sample and readme files, 132
 - UltraLite engine and runtime support, 13
- UltraLiteJ
 - (*see also* UltraLite Java edition databases)
 - (*see also* UltraLite Java edition management system)
 - UltraLiteJ database information utility (ULjInfo) syntax, 243
 - UltraLiteJ database load utility (ULjLoad) syntax, 244
 - UltraLiteJ database transfer utility (ULjDbT) syntax, 248
 - UltraLiteJ database unload utility (ULjUnload) syntax, 246
 - UltraLiteJ SQL statements
 - DELETE statement syntax, 432
 - DROP INDEX statement syntax, 433
 - INSERT statement, 438
 - ROLLBACK statement syntax, 443
 - SELECT statement syntax, 444
 - START SYNCHRONIZATION DELETE statement syntax, 446
 - STOP SYNCHRONIZATION DELETE statement syntax, 447
 - SYNCHRONIZE statement syntax, 448
 - TRUNCATE TABLE syntax, 450
 - UltraLite DROP TABLE statement syntax, 435
 - ulunload utility
 - syntax, 227
 - ulvalid utility
 - about, 230
 - uncommitted transactions
 - UltraLite isolation level, 41

- UltraLite overview, 470
- undoing
 - UltraLite transactions, 443
 - UltraLiteJ transactions, 443
- UNICODE characters
 - UltraLite collation for, 24
- UNION operation
 - UltraLite syntax, 451
- UNION statement
 - UltraLite syntax, 451
- union-all operation
 - UltraLite execution plans, 466
- unions
 - UltraLite multiple select statements, 451
- UNIQUE
 - UltraLite CREATE INDEX parameter, 421
- unique constraints
 - UltraLite characteristics, 55
 - UltraLite copying method, 52
 - UltraLite CREATE TABLE statement, 426
 - UltraLiteJ CREATE TABLE statement, 426
- unique indexes
 - UltraLite characteristics, 55
 - UltraLite index creation from, 57
 - UltraLite UNIQUE SQL parameter, 421
- unique keys
 - UltraLite characteristics, 55
 - UltraLite index creation from, 57
- UNIQUEIDENTIFIER data type
 - UltraLite, 265
- universally unique identifiers
 - (*see also* UUIDs)
 - UltraLite SQL syntax for NEWID function, 363
- unloading
 - UltraLite databases, 227
 - UltraLite databases with unload, 227
- unnamed foreign keys
 - UltraLite usage, 426
 - UltraLiteJ usage, 426
- unpacked rows
 - UltraLite about, 46
- UPDATE statement
 - UltraLite syntax, 452
- updates
 - UltraLite databases, 469
- updating
 - UltraLite updating rows, 452
- UpgradeSchemaFromFile method
 - UltraLite replacement for schema upgrade , 123
- upgrading
 - UltraLite schema error callback, 124
 - UltraLite schema process, 123
 - UltraLite SQL ALTER DATABASE SCHEMA FROM FILE syntax, 412
- upgrading UltraLite
 - troubleshooting connections, 479
- upload ok
 - UltraLite synchronization parameter, 108
- upload only
 - UltraLite synchronization parameter, 109
- upload only synchronization
 - UltraLite databases, 109
 - UltraLite upload_only synchronization parameter, 109
- upload_ok synchronization parameter
 - UltraLite reference, 108
- upload_only synchronization parameter
 - UltraLite reference, 109
- UPPER function
 - UltraLite syntax, 404
- uppercase characters
 - UltraLite UPPER function, 404
- uppercase strings
 - UltraLite UCASE function, 403
 - UltraLite UPPER function, 404
- user authentication
 - custom MobiLink, 90
 - PWD UltraLite connection parameter, 183
 - UltraLite, 59
 - UltraLite Authentication Value synchronization parameter, 92
 - UltraLite custom for synchronization, 96
 - UltraLite getUsername method, 110
 - UltraLite password synchronization parameter, 98
 - UltraLite synchronization status reports, 90
 - UltraLite user_name synchronization , 110
- user data
 - UltraLite synchronization parameter, 109
- user IDs
 - maximum length for UltraLite, 187
 - UltraLite databases, 59
- user name
 - UltraLite synchronization parameter, 110
- user-defined data types
 - UltraLite equivalents, 267
 - UltraLite unsupported in, 265

- user_data
 - UltraLite synchronization parameter, 109
- user_name
 - UltraLite synchronization parameter, 110
- users
 - managing in UltraLite databases, 61, 63, 64
 - UltraLite ALTER USER statement, 419
 - UltraLite databases, 59
 - UltraLite DROP USER statement, 436
- utf8_encoding creation parameter
 - UltraLite description, 164
- utf8_encoding database property
 - UltraLite usage, 24
- UTF8BIN collation
 - UltraLite considerations, 24
- utilities
 - UltraLite, 197
 - UltraLite database unload (ulunload) syntax, 227
 - UltraLite engine start (uleng12) syntax, 206
 - UltraLite engine stop (ulstop) syntax, 207
 - UltraLite erase database (ulerase) syntax, 208
 - UltraLite error codes, 198
 - UltraLite information (ulinfo) syntax, 209
 - UltraLite initialize database (ulinit) syntax, 210
 - UltraLite Interactive SQL (dbisql) syntax, 198
 - UltraLite load XML to database (ulload) syntax, 217
 - UltraLite SQL preprocessor (sqlpp) syntax, 204
 - UltraLite synchronization (ulsync) syntax, 221
 - UltraLite troubleshooting, 482
 - UltraLite validate database (ulvalid) syntax, 230
 - UltraLiteJ database information (ULjInfo) syntax, 243
 - UltraLiteJ database load (ULjLoad) syntax, 244
 - UltraLiteJ database transfer utility (ULjDbT) syntax, 248
 - UltraLiteJ database unload (ULjUnload) syntax, 246
 - Windows Mobile database administration on device, 32
- UUIDs
 - UltraLite SQL syntax for NEWID function, 363
 - UltraLite SQL syntax for STRTOUUID function, 394
 - UltraLite SQL syntax for UUIDTOSTR function, 404
- UUIDTOSTR function
 - UltraLite syntax, 404
- V**
 - validate database utility (ulvalid)
 - UltraLite about, 230
 - validating
 - (*see also* validating databases)
 - UltraLite checksum_level creation parameter, 142
 - UltraLite databases, 42
 - UltraLite databases with ulvalid, 230
 - validating databases
 - UltraLite validate database utility (ulvalid), 230
 - values
 - UltraLite index hash, 149
 - VARBINARY data type
 - UltraLite, 265
 - VARCHAR data type
 - UltraLite, 265
 - variables
 - UltraLite SQL, 299
 - version
 - UltraLite synchronization parameter, 111
 - version synchronization parameter
 - UltraLite reference, 111
 - versions
 - UltraLite troubleshooting utilities, 482
 - viewing
 - UltraLite execution plant, 465
 - UltraLite table methods, 51
 - viewing UltraLite database settings
 - about, 33
 - virtual file system (*see* VFS)
 - Visual Studio
 - UltraLite building CustDB application, 131
- W**
 - WEEKS function
 - UltraLite syntax, 405
 - WHEN
 - UltraLite CASE expression, 283
 - WHERE clause
 - CREATE PUBLICATION statement [MobiLink] [SQL Remote], 423
 - UltraLite CREATE PUBLICATION statement, 423
 - UltraLite DELETE statement, 432
 - UltraLite publication usage, 83
 - UltraLite search conditions, 287
 - UltraLite SELECT statement, 444

- UltraLite synchronization limitations, 76
- UltraLite UPDATE statement, 452
- whole tables
 - UltraLite publishing, 76
- wildcards
 - UltraLite LIKE search condition, 294
 - UltraLite PATINDEX function, 366
- Windows
 - (*see also* Windows ME)
 - (*see also* Windows Mobile 5)
 - (*see also* Windows Vista)
 - (*see also* Windows XP/200x)
 - UltraLite character sets, 24
 - UltraLite runtime libraries for, 115
- Windows desktop
 - UltraLite databases, 32
- Windows Mobile
 - choosing an API , 13
 - troubleshooting error -764, 483
 - UltraLite ActiveSync deployment, 126
 - UltraLite building CustDB application using .NET, 132
 - UltraLite character sets, 24
 - UltraLite considerations, 12
 - UltraLite engine and runtime support, 13
 - UltraLite engine deployment , 13
 - UltraLite file path prefix, 32
 - UltraLite FIPS enablement, 149
 - UltraLite MobiLink clients, 85
 - UltraLite runtime libraries for, 115
 - UltraLite uleng12 deployment, 13
- WITH CHECKPOINT clause
 - UltraLite LOAD TABLE statement, 442
- words
 - UltraLite keywords, 259
 - UltraLite reserved words, 259
- working with indexes
 - UltraLite about, 54

X

XML

- loading to database, 217
- sourcing UltraLite databases from , 20
- UltraLite unloading databases, 221

XML data type

- UltraLite equivalent, 267

Y

- YEAR function
 - UltraLite syntax, 407
- YEARS function
 - UltraLite syntax, 408
- YMD function
 - UltraLite syntax, 409

Z

- zero-padding
 - controlling with
 - timestamp_with_time_zone_format option, 164
 - UltraLite date_format creation parameter, 146
 - UltraLite timestamp_format creation parameter, 161
