



UltraLite® Database Management and Reference

February 2009

Version 11.0.1

Copyright and trademarks

Copyright © 2009 iAnywhere Solutions, Inc. Portions copyright © 2009 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	xi
About the SQL Anywhere documentation	xii
Introducing UltraLite	1
UltraLite feature comparison	2
UltraLite limitations	8
UltraLite data architecture	10
UltraLite storage and file name conventions	11
UltraLite transaction and state management	13
UltraLite isolation levels	18
Implementing an UltraLite solution	21
UltraLite supported platforms and protocols	22
Choosing an UltraLite data management component	23
Choosing an UltraLite programming interface	24
Using UltraLite Databases	27
Creating and configuring UltraLite databases	29
Creating an UltraLite database	30
Choosing database creation parameters for UltraLite	34
UltraLite character sets	36
Securing UltraLite databases	40
Connecting to an UltraLite database	43
UltraLite database connection parameters	44
UltraLite user authentication	51
Storing UltraLite parameters with the ULSQLCONNECT environment variable	53
Deploying UltraLite to devices	55
Deploy Multiple UltraLite applications with the UltraLite engine	57
Deploy UltraLite with AES_FIPS database encryption	59
Deploy UltraLite with TLS-enabled synchronization	60

Deploy the UltraLite HotSync conduit	62
Deploy the ActiveSync provider for UltraLite	64
Register applications with the ActiveSync Manager	66
Deploying UltraLite schema upgrades	67
Working with UltraLite databases	69
Working with UltraLite tables and columns	70
Working with UltraLite indexes	79
Working with UltraLite publications	83
Working with UltraLite users	87
Working with event notifications	89
UltraLite support for SQL passthrough	92
UltraLite CustDB samples	93
CustDB sample file locations	95
Lesson 1: Build and run the CustDB application	97
Lesson 2: Log in and populate the UltraLite remote database	99
Lesson 3: Use the CustDB client application	100
Lesson 4: Synchronize with the CustDB consolidated database	102
Lesson 5: Browse MobiLink synchronization scripts	104
Build your own application	106
UltraLite performance and optimization	107
Using index scans	108
Determining the access method used by the optimizer	110
Tuning query performance with index hashing	111
Choosing an optimal hash size	113
Setting the maximum hash size	116
Managing temporary tables	117
Flushing single or grouped transactions	119
How database encryption and obfuscation affect performance	120
UltraLite optimization strategies	121

UltraLite as a MobiLink Client 123

UltraLite clients	125
Built-in UltraLite synchronization features	126
Customizing UltraLite client synchronization behavior	127
Primary key uniqueness in UltraLite	129

Designing synchronization in UltraLite	133
Using MobiLink file transfers	142
Using ActiveSync and HotSync with UltraLite	145
HotSync on Palm OS	146
ActiveSync on Windows Mobile	150
UltraLite synchronization parameters and network protocol options	153
Synchronization parameters for UltraLite	154
Network protocol options for UltraLite synchronization streams	178
UltraLite Database Reference	179
UltraLite creation parameters	181
UltraLite case creation parameter	183
UltraLite checksum_level creation parameter	184
UltraLite collation creation parameter	186
UltraLite date_format creation parameter	187
UltraLite date_order creation parameter	190
UltraLite fips creation parameter	192
UltraLite max_hash_size creation parameter	194
UltraLite nearest_century creation parameter	196
UltraLite obfuscate creation parameter	198
UltraLite page_size creation parameter	199
UltraLite precision creation parameter	201
UltraLite scale creation parameter	203
UltraLite time_format creation parameter	205
UltraLite timestamp_format creation parameter	207
UltraLite timestamp_increment creation parameter	210
UltraLite utf8_encoding creation parameter	211
UltraLite database properties	213
Accessing UltraLite database properties	218
UltraLite database options	219
UltraLite commit_flush_count option [temporary]	220
UltraLite commit_flush_timeout option [temporary]	221
UltraLite global_database_id option	222
UltraLite ml_remote_id option	223
Change UltraLite persistent database option settings	224
UltraLite connection parameters	225

UltraLite CACHE_SIZE connection parameter	226
UltraLite CE_FILE connection parameter	227
UltraLite COMMIT_FLUSH connection parameter	229
UltraLite CON connection parameter	230
UltraLite DBF connection parameter	231
UltraLite DBKEY connection parameter	233
UltraLite DBN connection parameter	234
UltraLite MIRROR_FILE connection parameter	235
UltraLite NT_FILE connection parameter	237
UltraLite ORDERED_TABLE_SCAN connection parameter [deprecated]	239
UltraLite PALM_ALLOW_BACKUP connection parameter	240
UltraLite PALM_FILE connection parameter	241
UltraLite PWD connection parameter	243
UltraLite RESERVE_SIZE connection parameter	244
UltraLite START connection parameter	245
UltraLite UID connection parameter	246
UltraLite utilities	247
Supported exit codes	248
Interactive SQL utility for UltraLite (dbisql)	249
SQL Preprocessor for UltraLite utility (sqlpp)	252
UltraLite Create Database utility (ulcreate)	255
UltraLite Data Management utility for Palm OS (ULDBUtil)	258
UltraLite Engine utility (uleng11)	260
UltraLite Engine Stop utility (ulstop)	261
UltraLite Erase database (ulerase)	262
UltraLite HotSync Conduit Installation utility for Palm OS (ulcond11)	263
UltraLite Information utility (ulinfo)	265
UltraLite Initialize Database utility (ulinit)	269
UltraLite Load XML to Database utility (ulload)	272
UltraLite Synchronization utility (ulsync)	275
Synchronization profile options	278
UltraLite Unload Database utility (ulunload)	281
UltraLite Unload Old Database utility (ulunloadold)	284
UltraLite Validate Database utility (ulvalid)	286
UltraLite system tables	289
View or hide UltraLite system tables	290

systable system table	291
syscolumn system table	292
sysindex system table	293
sysixcol system table	295
syspublication system table	296
sysarticle system table	297
sysuldata system table	298
UltraLite SQL Reference	299
UltraLite SQL elements	301
Keywords in UltraLite	302
Identifiers in UltraLite	303
Strings in UltraLite	304
Comments in UltraLite	305
Numbers in UltraLite	306
The NULL value in UltraLite	307
Special values in UltraLite	308
Dates and times in UltraLite	311
Data types in UltraLite	312
Expressions in UltraLite	325
Operators in UltraLite	338
Variables in UltraLite	341
Execution plans in UltraLite	342
UltraLite SQL functions	347
Function types	348
SQL functions (A-D)	354
SQL functions (E-O)	383
SQL functions (P-Z)	411
UltraLite SQL statements	447
UltraLite statement categories	449
UltraLite ALTER DATABASE SCHEMA FROM FILE statement	450
UltraLite ALTER PUBLICATION statement	451
UltraLite ALTER SYNCHRONIZATION PROFILE statement	452
UltraLite ALTER TABLE statement	454
UltraLite CHECKPOINT statement	458

UltraLite COMMIT statement	459
UltraLite CREATE INDEX statement	460
UltraLite CREATE PUBLICATION statement	462
UltraLite CREATE SYNCHRONIZATION PROFILE statement	464
UltraLite CREATE TABLE statement	468
UltraLite DELETE statement	473
UltraLite DROP INDEX statement	474
UltraLite DROP PUBLICATION statement	475
UltraLite DROP SYNCHRONIZATION PROFILE statement	476
UltraLite DROP TABLE statement	477
UltraLite FROM clause	478
UltraLite INSERT statement	480
UltraLite LOAD TABLE statement	481
UltraLite ROLLBACK statement	485
UltraLite SELECT statement	486
UltraLite SET OPTION statement	488
UltraLite START SYNCHRONIZATION DELETE statement	489
UltraLite STOP SYNCHRONIZATION DELETE statement	490
UltraLite SYNCHRONIZE statement	491
UltraLite TRUNCATE TABLE statement	493
UltraLite UNION statement	495
UltraLite UPDATE statement	496

Troubleshooting UltraLite 497

Unable to start the UltraLite engine	498
Unable to connect to databases after upgrade	499
Database corruption	500
Database size not stabilizing	501
Importing ASCII data into a new database	502
Utilities still running as the previous version	503
Result set changes unpredictably	504
UltraLite engine client fails with error -764	505

Glossary 507

Glossary 509

Index 539

About this book

Subject

This book introduces the UltraLite database system for small devices.

Audience

This book is intended for all developers who want to take advantage of the performance, resource efficiency, robustness, and security of an UltraLite relational database for data storage and synchronization for embedded or mobile devices.

About the SQL Anywhere documentation

The complete SQL Anywhere documentation is available in four formats that contain identical information.

- **HTML Help** The online Help contains the complete SQL Anywhere documentation, including the books and the context-sensitive help for SQL Anywhere tools.

If you are using a Microsoft Windows operating system, the online Help is provided in HTML Help (CHM) format. To access the documentation, choose **Start » Programs » SQL Anywhere 11 » Documentation » Online Books**.

The administration tools use the same online documentation for their Help features.

- **Eclipse** On Unix platforms, the complete online Help is provided in Eclipse format. To access the documentation, run *sadoc* from the *bin32* or *bin64* directory of your SQL Anywhere 11 installation.
- **DocCommentXchange** DocCommentXchange is a community for accessing and discussing SQL Anywhere documentation.

Use DocCommentXchange to:

- View documentation
- Check for clarifications users have made to sections of documentation
- Provide suggestions and corrections to improve documentation for all users in future releases

Visit <http://dcx.sybase.com>.

- **PDF** The complete set of SQL Anywhere books is provided as a set of Portable Document Format (PDF) files. You must have a PDF reader to view information. To download Adobe Reader, visit <http://get.adobe.com/reader/>.

To access the PDF documentation on Microsoft Windows operating systems, choose **Start » Programs » SQL Anywhere 11 » Documentation » Online Books - PDF Format**.

To access the PDF documentation on Unix operating systems, use a web browser to open *install-dir/documentation/en/pdf/index.html*.

About the books in the documentation set

The SQL Anywhere documentation consists of the following books:

- **SQL Anywhere 11 - Introduction** This book introduces SQL Anywhere 11, a comprehensive package that provides data management and data exchange, enabling the rapid development of database-powered applications for server, desktop, mobile, and remote office environments.
- **SQL Anywhere 11 - Changes and Upgrading** This book describes new features in SQL Anywhere 11 and in previous versions of the software.
- **SQL Anywhere Server - Database Administration** This book describes how to run, manage, and configure SQL Anywhere databases. It describes database connections, the database server, database

files, backup procedures, security, high availability, replication with the Replication Server, and administration utilities and options.

- **SQL Anywhere Server - Programming** This book describes how to build and deploy database applications using the C, C++, Java, PHP, Perl, Python, and .NET programming languages such as Visual Basic and Visual C#. A variety of programming interfaces such as ADO.NET and ODBC are described.
- **SQL Anywhere Server - SQL Reference** This book provides reference information for system procedures, and the catalog (system tables and views). It also provides an explanation of the SQL Anywhere implementation of the SQL language (search conditions, syntax, data types, and functions).
- **SQL Anywhere Server - SQL Usage** This book describes how to design and create databases; how to import, export, and modify data; how to retrieve data; and how to build stored procedures and triggers.
- **MobiLink - Getting Started** This book introduces MobiLink, a session-based relational-database synchronization system. MobiLink technology allows two-way replication and is well suited to mobile computing environments.
- **MobiLink - Client Administration** This book describes how to set up, configure, and synchronize MobiLink clients. MobiLink clients can be SQL Anywhere or UltraLite databases. This book also describes the Dbmlsync API, which allows you to integrate synchronization seamlessly into your C++ or .NET client applications.
- **MobiLink - Server Administration** This book describes how to set up and administer MobiLink applications.
- **MobiLink - Server-Initiated Synchronization** This book describes MobiLink server-initiated synchronization, a feature that allows the MobiLink server to initiate synchronization or perform actions on remote devices.
- **QAnywhere** This book describes QAnywhere, which is a messaging platform for mobile, wireless, desktop, and laptop clients.
- **SQL Remote** This book describes the SQL Remote data replication system for mobile computing, which enables sharing of data between a SQL Anywhere consolidated database and many SQL Anywhere remote databases using an indirect link such as email or file transfer.
- **UltraLite - Database Management and Reference** This book introduces the UltraLite database system for small devices.
- **UltraLite - C and C++ Programming** This book describes UltraLite C and C++ programming interfaces. With UltraLite, you can develop and deploy database applications to handheld, mobile, or embedded devices.
- **UltraLite - M-Business Anywhere Programming** This book describes UltraLite for M-Business Anywhere. With UltraLite for M-Business Anywhere you can develop and deploy web-based database applications to handheld, mobile, or embedded devices, running Palm OS, Windows Mobile, or Windows.
- **UltraLite - .NET Programming** This book describes UltraLite.NET. With UltraLite.NET you can develop and deploy database applications to computers, or handheld, mobile, or embedded devices.
- **UltraLiteJ** This book describes UltraLiteJ. With UltraLiteJ, you can develop and deploy database applications in environments that support Java. UltraLiteJ supports BlackBerry smartphones and Java SE environments. UltraLiteJ is based on the iAnywhere UltraLite database product.

- **Error Messages** This book provides a complete listing of SQL Anywhere error messages together with diagnostic information.

Documentation conventions

This section lists the conventions used in this documentation.

Operating systems

SQL Anywhere runs on a variety of platforms. In most cases, the software behaves the same on all platforms, but there are variations or limitations. These are commonly based on the underlying operating system (Windows, Unix), and seldom on the particular variant (AIX, Windows Mobile) or version.

To simplify references to operating systems, the documentation groups the supported operating systems as follows:

- **Windows** The Microsoft Windows family includes Windows Vista and Windows XP, used primarily on server, desktop, and laptop computers, and Windows Mobile used on mobile devices.

Unless otherwise specified, when the documentation refers to Windows, it refers to all Windows-based platforms, including Windows Mobile.

- **Unix** Unless otherwise specified, when the documentation refers to Unix, it refers to all Unix-based platforms, including Linux and Mac OS X.

Directory and file names

In most cases, references to directory and file names are similar on all supported platforms, with simple transformations between the various forms. In these cases, Windows conventions are used. Where the details are more complex, the documentation shows all relevant forms.

These are the conventions used to simplify the documentation of directory and file names:

- **Uppercase and lowercase directory names** On Windows and Unix, directory and file names may contain uppercase and lowercase letters. When directories and files are created, the file system preserves letter case.

On Windows, references to directories and files are *not* case sensitive. Mixed case directory and file names are common, but it is common to refer to them using all lowercase letters. The SQL Anywhere installation contains directories such as *Bin32* and *Documentation*.

On Unix, references to directories and files *are* case sensitive. Mixed case directory and file names are not common. Most use all lowercase letters. The SQL Anywhere installation contains directories such as *bin32* and *documentation*.

The documentation uses the Windows forms of directory names. In most cases, you can convert a mixed case directory name to lowercase for the equivalent directory name on Unix.

- **Slashes separating directory and file names** The documentation uses backslashes as the directory separator. For example, the PDF form of the documentation is found in *install-dir\Documentation\en\PDF* (Windows form).

On Unix, replace the backslash with the forward slash. The PDF documentation is found in *install-dir/documentation/en/pdf*.

- **Executable files** The documentation shows executable file names using Windows conventions, with a suffix such as *.exe* or *.bat*. On Unix, executable file names have no suffix.

For example, on Windows, the network database server is *dbsrv11.exe*. On Unix, it is *dbsrv11*.

- **install-dir** During the installation process, you choose where to install SQL Anywhere. The environment variable `SQLANY11` is created and refers to this location. The documentation refers to this location as *install-dir*.

For example, the documentation may refer to the file *install-dir\readme.txt*. On Windows, this is equivalent to `%SQLANY11%\readme.txt`. On Unix, this is equivalent to `$(SQLANY11)/readme.txt` or `{SQLANY11}/readme.txt`.

For more information about the default location of *install-dir*, see [“SQLANY11 environment variable” \[SQL Anywhere Server - Database Administration\]](#).

- **samples-dir** During the installation process, you choose where to install the samples included with SQL Anywhere. The environment variable `SQLANY11SAMP` is created and refers to this location. The documentation refers to this location as *samples-dir*.

To open a Windows Explorer window in *samples-dir*, from the **Start** menu, choose **Programs » SQL Anywhere 11 » Sample Applications And Projects**.

For more information about the default location of *samples-dir*, see [“SQLANY11SAMP environment variable” \[SQL Anywhere Server - Database Administration\]](#).

Command prompts and command shell syntax

Most operating systems provide one or more methods of entering commands and parameters using a command shell or command prompt. Windows command prompts include Command Prompt (DOS prompt) and 4NT. Unix command shells include Korn shell and bash. Each shell has features that extend its capabilities beyond simple commands. These features are driven by special characters. The special characters and features vary from one shell to another. Incorrect use of these special characters often results in syntax errors or unexpected behavior.

The documentation provides command line examples in a generic form. If these examples contain characters that the shell considers special, the command may require modification for the specific shell. The modifications are beyond the scope of this documentation, but generally, use quotes around the parameters containing those characters or use an escape character before the special characters.

These are some examples of command line syntax that may vary between platforms:

- **Parentheses and curly braces** Some command line options require a parameter that accepts detailed value specifications in a list. The list is usually enclosed with parentheses or curly braces. The documentation uses parentheses. For example:

```
-x tcpip(host=127.0.0.1)
```

Where parentheses cause syntax problems, substitute curly braces:

```
-x tcpip{host=127.0.0.1}
```

If both forms result in syntax problems, the entire parameter should be enclosed in quotes as required by the shell:

```
-x "tcPIP(host=127.0.0.1)"
```

- **Quotes** If you must specify quotes in a parameter value, the quotes may conflict with the traditional use of quotes to enclose the parameter. For example, to specify an encryption key whose value contains double-quotes, you might have to enclose the key in quotes and then escape the embedded quote:

```
-ek "my \"secret\" key"
```

In many shells, the value of the key would be my "secret" key.

- **Environment variables** The documentation refers to setting environment variables. In Windows shells, environment variables are specified using the syntax `%ENVVAR%`. In Unix shells, environment variables are specified using the syntax `$ENVVAR` or `${ENVVAR}`.


Graphic icons

The following icons are used in this documentation.


- A client application.



- A database server, such as Sybase SQL Anywhere.



- A database. In some high-level diagrams, the icon may be used to represent both the database and the database server that manages it.



- Replication or synchronization middleware. These assist in sharing data among databases. Examples are the MobiLink server and the SQL Remote Message Agent.



- A programming interface.



Contacting the documentation team

We would like to receive your opinions, suggestions, and feedback on this Help.

To submit your comments and suggestions, send an email to the SQL Anywhere documentation team at iasdoc@sybase.com. Although we do not reply to emails, your feedback helps us to improve our documentation, so your input is welcome.

DocCommentXchange

You can also leave comments directly on help topics using DocCommentXchange. DocCommentXchange (DCX) is a community for accessing and discussing SQL Anywhere documentation. Use DocCommentXchange to:

- View documentation
- Check for clarifications users have made to sections of documentation
- Provide suggestions and corrections to improve documentation for all users in future releases

Visit <http://dcx.sybase.com>.

Finding out more and requesting technical support

Additional information and resources are available at the Sybase iAnywhere Developer Community at <http://www.sybase.com/developer/library/sql-anywhere-techcorner>.

If you have questions or need help, you can post messages to the Sybase iAnywhere newsgroups listed below.

When you write to one of these newsgroups, always provide details about your problem, including the build number of your version of SQL Anywhere. You can find this information by running the following command:
dbeng11 -v.

The newsgroups are located on the *forums.sybase.com* news server.

The newsgroups include the following:

- [sybase.public.sqlanywhere.general](#)
- [sybase.public.sqlanywhere.linux](#)
- [sybase.public.sqlanywhere.mobilink](#)
- [sybase.public.sqlanywhere.product_futures_discussion](#)
- [sybase.public.sqlanywhere.replication](#)
- [sybase.public.sqlanywhere.ultralite](#)
- [ianywhere.public.sqlanywhere.qanywhere](#)

For web development issues, see <http://groups.google.com/group/sql-anywhere-web-development>.

Newsgroup disclaimer

iAnywhere Solutions has no obligation to provide solutions, information, or ideas on its newsgroups, nor is iAnywhere Solutions obliged to provide anything other than a systems operator to monitor the service and ensure its operation and availability.

iAnywhere Technical Advisors, and other staff, assist on the newsgroup service when they have time. They offer their help on a volunteer basis and may not be available regularly to provide solutions and information. Their ability to help is based on their workload.

Introducing UltraLite

Contents

UltraLite feature comparison	2
UltraLite limitations	8
UltraLite data architecture	10
UltraLite storage and file name conventions	11
UltraLite transaction and state management	13
UltraLite isolation levels	18

UltraLite is a compact relational database with many of the features and functionality of SQL Anywhere. UltraLite can be installed as part of a SQL Anywhere solution for enterprise-wide mobile data management, or as part of a standalone embedded solution.

UltraLite provides built-in functionality to mobilize corporate data. When deployed as a MobiLink client, you can implement a synchronization solution that delivers mission-critical information in a timely and reliable way. Users can record and access the data they require—even when they cannot directly access the corporate network.

An application development interface allows you to write applications and create added value for your users.

See also

- [“Understanding MobiLink synchronization” \[MobiLink - Getting Started\]](#)
- [“MobiLink clients” \[MobiLink - Client Administration\]](#)
- [MobiLink - Getting Started](#)
- [MobiLink - Client Administration](#)
- [MobiLink - Server Administration](#)

UltraLite feature comparison

In the C/C++ version, the UltraLite database and management system adds 400-500 KB to the size of your application. The SQL Anywhere database, database server, and synchronization client add approximately 6 MB.

Feature	SQL Anywhere	Ultra-Lite	Considerations
Transaction processing, referential integrity, 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	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. See “Avoiding synchronization issues with foreign key cycles” on page 138.
Cascading updates and deletes	X	Limited	
Dynamic, multiple database support	X	X	With the UltraLite engine only.

Feature	SQL Anywhere	Ultra-Lite	Considerations
Multi-threaded application support	X	X	
Row-level locking	X	X	
XML unload and load utilities		X	UltraLite uses separate administration tools to complete XML load and unloads. It is not built into the runtime. See “UltraLite Load XML to Database utility (ulload)” on page 272 and “UltraLite Unload Database utility (ulunload)” on page 281 .
XML export and import utilities	X		SQL Anywhere uses SQL statements to export/import data to XML. You can also use dbunload to export your data. See “Importing and exporting data” [SQL Anywhere Server - SQL Usage] .
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 a <code>Feature not available in UltraLite</code> error. See “UltraLite SQL functions” on page 347 .
SQL statements	X	Limited	The scope of SQL statements are limited in UltraLite compared to SQL Anywhere. See “UltraLite SQL statements” on page 447 .
Integrated HTTP server	X		
Strong encryption for database files and network communications	X	X	

Feature	SQL Anywhere	Ultra-Lite	Considerations
Event scheduling and handling	X		
High-performance, self-tuning, cost-based query optimizer	X		UltraLite has a query optimizer, but it is not as extensive as that of SQL Anywhere. Therefore, the UltraLite optimizer may not provide as high performance as the SQL Anywhere optimizer on complex queries.
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 “Choosing an UltraLite programming interface” on page 24.
Cursor support	X	X	See “UltraLite limitations” on page 8.
Dynamic cache sizing with an advanced cache management system	X		Cache sizing is static in UltraLite. Nonetheless, UltraLite allows you to set the cache size when the database is started, which gives you the ability to scale cache size. See “UltraLite CACHE_SIZE connection parameter” on page 226.
Database recovery after system or application failure	X	X	
Binary Large Object (BLOB) support	X	X	UltraLite cannot index or compare BLOBs.

Feature	SQL Anywhere	Ultra-Lite	Considerations
Windows Performance Monitor integration	X		
Online table and index defragmentation	X		
Online backup	X		
Small footprint, which can be as small as 500 KB		X	Small footprint devices tend to have relatively slow processors. UltraLite employs algorithms and data structures that are targeted for these devices, so UltraLite continues to provide high performance and low memory use.
Supports Palm OS		X	
Direct device connections to a Windows Mobile device from the desktop.		X	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 simply need to prefix the connection string with WCE:\ . See “Windows Mobile” on page 47 .
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 194 .

Feature	SQL Anywhere	Ultra-Lite	Considerations
Synchronizing to Oracle, DB2, Sybase Adaptive Server Enterprise, 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 125 .
In-process execution		X	
Computed columns	X		
Declared temporary tables/global temporary tables	X		
System functions	X		UltraLite does not support SQL Anywhere system functions, including property functions. You cannot include them as part of your UltraLite application.
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.

Feature	SQL Anywhere	Ultra-Lite	Considerations
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 user authentication” on page 51 .

UltraLite limitations

To compare UltraLite limitations with SQL Anywhere limitations, see [“SQL Anywhere size and number limitations” \[SQL Anywhere Server - Database Administration\]](#).

Statistic	Maximum for UltraLite
Number of connections per database	Up to 14 for single threaded applications.
Number of concurrent open connections	Up to 8 for Palm OS. Up to 32 for all other OS.
Total number of concurrent connections per application	Up to 16 on Palm OS. Up to 64 for all other OS.
SQL communication areas	Up to 63.
File-based persistent store (database size)	2 GB file or OS limit on file size.
Palm Computing Platform database size	128 MB (primary storage). 2 GB (expansion card file system).
Rows per table	Up to 16 million. ¹
Rows per database	Limited by persistent store.
Table size	Limited only by database size.
Tables per database	Limited only by database size.
Columns per table	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.
Indexes per table	Limited only by database size.
Tables referenced per transaction	No limit.

Statistic	Maximum for UltraLite
Stored procedure length	Not applicable.
Stored procedures per database	Not applicable.
Triggers per database	Not applicable.
Nesting	Not applicable.
Number of publications	32 publications.
Page size	16 KB.
Row size	The length of each packed row must not exceed the page size. See “Row packing and table definitions” on page 70. 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.
Cursors per connection	The number of maximum allowable cursors on a given connection to an UltraLite database is 64 (all platforms).
Long binary/long varchar size	Limited only by database size.

¹ In some cases, 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 16.

UltraLite data architecture

UltraLite is a mobile database with true platform independence. UltraLite is designed to create custom solutions for small-footprint devices such as cell phones, handheld computers, and embedded devices.

UltraLite provides you with a complete database management system including:

- **Development layer** UltraLite supports several programming interfaces that keep you from getting locked into one deployment platform, development tool, or set of IT infrastructure products.

For information about which API you should choose, see [“Choosing an UltraLite programming interface” on page 24](#).


To help you maintain your UltraLite project, UltraLite completes its development support with a comprehensive set of administration tools. You can run these tools either as command line utilities or as wizards in the UltraLite plug-in for Sybase Central.

- **Data management layer and synchronization client** You can connect to an UltraLite database with an in-process library called a runtime, or a separate process called an engine. Both processes control connection and data access requests. They also include a built-in bi-directional synchronization client that links UltraLite databases with enterprise back-end systems.

For information about which process you should choose, see [“Choosing an UltraLite data management component” on page 23](#).

- **Data layer** This layer is the local data repository either stored as a file or a series of data records (for Palm OS). See [“UltraLite storage and file name conventions” on page 11](#).

The data, management, and development layers are represented in the following figure.



UltraLite storage and file name conventions

The file management requirements of a device dictate how an UltraLite database is stored and what database name conventions must be used. While most platforms use traditional file-based storage, others, such as the Palm OS record-based store, require that the database be saved differently. However, the UltraLite database is typically referred to as a file. Depending on platform restrictions, you may be able to create a database on the development desktop, and then deploy it to one or more platforms.

See also

- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)

UltraLite database schema

The logical framework of the database is known as a **schema**. In UltraLite, the schema is maintained as a catalog of system tables that hold the metadata for the UltraLite database. Metadata stored in the system tables includes:

- Index definitions. See [“sysindex system table” on page 293](#) and [“sysixcol system table” on page 295](#).
- Table definitions. See [“systable system table” on page 291](#).
- Column definitions. See [“syscolumn system table” on page 292](#).
- Publication definitions. See [“syspublication system table” on page 296](#) and [“sysarticle system table” on page 297](#).
- User names and passwords. See [“sysuldata system table” on page 298](#).

Schema changes with DDL statements

You can change the schema of a database with the appropriate Data Definition Language (DDL) statements. If you require more than just a few changes, you can use 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 schema (for example, synchronization, prepared but unreleased statements, or executing database operations).

If either of these conditions is true, the DDL statement fails. When the DDL statement is executing, any other attempt to use the database is blocked until the DDL statements completes the schema change.

See also

- [“UltraLite ALTER DATABASE SCHEMA FROM FILE statement” on page 450](#)
- [“UltraLite SQL statements” on page 447](#)

UltraLite temporary files

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.

Temporary files are only used by UltraLite and are maintained in the same file path (if one exists) as the UltraLite database itself. The temporary file has the same file name as the database, but with the following difference:

- **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 this file.
- **For record-based platforms** The tilde is appended to the end of the name of the file. For example, if *CustDB.udb* existed as a record-based file for Palm OS, the temporary file for it would be maintained as *CustDB.udb~*.

Tip

You can safely delete the temporary file without loss of data—as long as UltraLite is not running. It does not contain information that is required across sessions.

UltraLite temporary tables

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 subqueries need to be evaluated early in the access plan, or when intermediate results do not fit in the available memory.

Data in a temporary table is held for a single connection only. Temporary 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.

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 temporary files” on page 12](#)
- [“UltraLite performance and optimization” on page 107](#)
- [“When to view an execution plan” on page 342](#)

UltraLite transaction and state management

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. See [“Concurrency in UltraLite” on page 13](#).
- Synchronization progress counts, to ensure that synchronization occurs successfully. See [“Built-in UltraLite synchronization features” on page 126](#).
- Row state, to maintain data integrity by tracking how data has changed between synchronizations. See [“UltraLite row states” on page 14](#).
- Transactions, to determine when and how data gets committed. In UltraLite, a transaction is processed in its entirety or not at all. See [“UltraLite transaction processing” on page 16](#).
- 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 [“Backing up and recovering data in UltraLite” on page 15](#).

Concurrency in UltraLite

Concurrency is the manner in which UltraLite shares resources by allowing multiple connections at the same time. UltraLite uses the following methods to manage concurrency:

- **Multiple databases** A single UltraLite application can open connections to multiple databases. UltraLite can open up to 8 databases on Palm OS and up to 32 databases for all other platforms.
- **Multiple applications** The UltraLite database can only be opened by one process at a time. If you plan to support concurrency among multiple applications, choose the UltraLite engine as your data management component. See [“Choosing an UltraLite data management component” on page 23](#).
- **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.

If you are managing your database with the runtime, ensure that you do not exceed the concurrent connection limits supported by UltraLite. These are:

- 16 concurrent connections for Palm OS.
- 64 concurrent connections for all other platforms.

If you are managing your database connections with the UltraLite engine, the number of SQLCAs you can use is typically restricted to 128. However, the implementation of UltraLite.NET API effectively reduces this limit to 128 less the number of running UltraLite.NET clients.

- **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. See [“UltraLite transaction processing” on page 16](#).

- **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. See [“Additional Parameters synchronization parameter” on page 155](#).

If synchronization fails, UltraLite supports resumable downloads on all platforms. See [“Handling failed downloads” \[MobiLink - Server Administration\]](#).

See also

- [“UltraLite clients” on page 125](#)
- [“Additional Parameters synchronization parameter” on page 155](#)

UltraLite row states

Maintaining row state information is a powerful part of the UltraLite feature set. Tracking the state of tables and rows is particularly important for data synchronization.

Changes to data

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 a row is inserted, updated, or deleted, the state of the row is modified 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. Restore the original state of the row to undo the delete.
- **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

- [“Backing up and recovering data in UltraLite” on page 15](#)
- [“Flushing single or grouped transactions” on page 119](#)
- [“UltraLite transaction processing” on page 16](#)

Validate an UltraLite database

It is recommended that you validate when there is no other activity on the database.

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

- The **Validate Database Wizard** in Sybase Central.
- Calling the `ValidateDatabase` function.
- The `ulvalid` utility.
- The `ValidateDatabase` method.

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.

To validate a database (Sybase Central)

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

Tip

You can also access the **Validate Database Wizard** from within Sybase Central using any of the following methods:

- Right-clicking the UltraLite database, and choosing **Validate Database**.
- Selecting the UltraLite database, and choosing **Tools » UltraLite 11 » Validate Database**.

See also

- [“ValidateDatabase method” \[UltraLite - .NET Programming\]](#)
- [“ValidateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- [“UltraLite Validate Database utility \(ulvalid\)” on page 286](#)

Backing up and recovering data in UltraLite

If an application using an UltraLite database stops unexpectedly, the UltraLite database automatically recovers to a consistent state when the application is restarted. All committed transactions flushed to memory

prior to the unexpected failure are present in the UltraLite database. All transactions not flushed at the time of the failure are rolled back.

UltraLite 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. See [“UltraLite row states” on page 14](#).

Backups

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

In smaller UltraLite deployments, you can copy the database file to a desktop computer to provide a manual backup.

Note

On Palm OS, you can also set the `PALM_ALLOW_BACKUP` connection parameter to true. This setting allows you to back up the database over HotSync. See [“UltraLite PALM_ALLOW_BACKUP connection parameter” on page 240](#).

See also

- [“Flushing single or grouped transactions” on page 119](#)

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.

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.

For more information about how to handle errors, see:

- UltraLite for C/C++: “[Handling errors](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[Handling errors](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[Handling errors](#)” [*UltraLite - M-Business Anywhere Programming*]

Programming tip

All UltraLite APIs—except the C++ API—can operate in *autocommit* mode. Some APIs use autocommit by default.

This default means that each transaction is automatically committed after each operation. 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.NET: “[Managing transactions](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[Managing transactions](#)” [*UltraLite - M-Business Anywhere Programming*]

UltraLite isolation levels

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. For .NET, read-committed is the default isolation level for new transactions created by calling `ULConnection.BeginTransaction` without parameters. The default UltraLite isolation level provides the best performance while ensuring data consistency.

With the `ReadCommitted` isolation level:

- Dirty reads are prevented
- No read locks are applied
- Uncommitted deletes are visible
- Non-repeatable reads and phantom rows are allowed
- No guarantee that data will not change during transaction

With the `ReadUncommitted` isolation level:

- Dirty reads are allowed
- No read locks are applied
- Non-repeatable reads and phantom rows are allowed
- No guarantee that concurrent transaction will not modify row or roll back changes to row

You can change the isolation level from `ReadCommitted` to `ReadUncommitted`. For UltraLite C++, use the `SetDatabaseOption` method to change the isolation level. For UltraLite.NET 2.0, call the `ULConnection.BeginTransaction` to create a transaction with the `ReadUncommitted` isolation level.

Note

For .NET, executing `SetDatabaseOption` while a transaction is active is not recommended. It changes the isolation level of the connection, but does not update the `ULTransaction.IsolationLevel`.

Do not use `SetDatabaseOption` to change the isolation while a transaction is in progress; unpredictable results might occur.

Isolation level side effects

As UltraLite operates by default at an isolation level of 0, the following side effects are possible:

- No locking operations are required when executing a `SELECT` statement.
- Applications can access uncommitted data (also known as dirty reads). In this scenario, an application 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: Dirty reads” \[SQL Anywhere Server - SQL Usage\]](#). For a tutorial that demonstrates a phantom row, see [“Tutorial: 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: Non-repeatable reads” \[SQL Anywhere Server - SQL Usage\]](#).

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.

See also

- [“BeginTransaction\(\) method” \[UltraLite - .NET Programming\]](#)
- [“BeginTransaction\(IsolationLevel\) method” \[UltraLite - .NET Programming\]](#)
- [“SetDatabaseOption method” \[UltraLite - .NET Programming\]](#)

Implementing an UltraLite solution

Contents

UltraLite supported platforms and protocols	22
Choosing an UltraLite data management component	23
Choosing an UltraLite programming interface	24

When implementing an UltraLite solution, consider the following:

- How many applications need to connect to the UltraLite database? The number of concurrent connections affects whether you need the UltraLite in-process runtime or the UltraLite engine. To understand how they differ, see [“Choosing an UltraLite data management component”](#) on page 23.
- What platform(s) do you want to support? This choice can affect which APIs are available to program your application. See [“Choosing an UltraLite programming interface”](#) on page 24.
- Which platforms will the database run on? Because file formats have been consolidated, you may be able to create a database that runs on multiple platforms. See [“Creating and configuring UltraLite databases”](#) on page 29.

Tip

If you need to create a file format that suits multiple platforms, use the **Create Database Wizard** in Sybase Central to help you determine whether a single database is possible.

UltraLite supported platforms and protocols

You can synchronize the data in UltraLite databases with a central consolidated database over the TCP/IP, HTTP, or HTTPS network protocols.

For more information about which platforms different network protocols (also known as streams) are supported on, see <http://www.sybase.com/detail?id=1061806>.

See also

- “UltraLite clients” on page 125
- “Network protocol options for UltraLite synchronization streams” on page 178
- “UltraLite Synchronization utility (ulsync)” on page 275

Choosing an UltraLite data management component

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 library types.

UltraLite in-process runtime library

In UltraLite, the runtime and the application are part of the same process, which makes the database specific to the application. For all platforms, the runtime manages UltraLite databases and built-in synchronization operations. The UltraLite runtime can manage a maximum of 14 databases at any one time. Note that only one application can access a database at any given time.

Linking to the runtime requires that you use a specific import library/DLL pair than that of the engine. For details, see [“Compiling and linking your application” \[UltraLite - C and C++ Programming\]](#). If you require TLS-enabled synchronization, there are additional libraries you also require. See [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#).

UltraLite database engine client library

The UltraLite engine is only available for Windows desktop and Windows Mobile platforms. The engine is a separate executable, which provides concurrent database/application connections. Each application must use a client library when using the UltraLite engine. This client library allows each application to communicate with the UltraLite engine. The UltraLite engine requires more system resources than the UltraLite runtime and may yield lower performance.

Linking to the engine requires that you a different import library/DLL pair than that of the runtime. If you require TLS-enabled synchronization or AES FIPS database encryption, there are additional libraries you also require. See [“Deploy UltraLite with AES_FIPS database encryption” on page 59](#) and [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#).

Note

Ensure that you do not exceed the concurrent connection limits supported by the UltraLite engine.

See also

- [“UltraLite Engine utility \(uleng11\)” on page 260](#)
- [“Concurrency in UltraLite” on page 13](#)
- [“UltraLite feature comparison” on page 2](#)
- [“UltraLite limitations” on page 8](#)

Choosing an UltraLite programming interface

All UltraLite APIs expose core database functionality. Some of the following APIs integrate with the development environment to simplify programming tasks:

- C/C++ interface
- Embedded SQL for C/C++
- UltraLite.NET (C# or VB.NET)
- M-Business Anywhere (JavaScript)

UltraLite APIs offer different 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 that allows for the creation of applications for varied deployment environments.

Note

Pocket Builder is not supported in this version of UltraLite. Sybase PocketBuilder is not included with SQL Anywhere. Contact Sybase for details (<http://www.sybase.com/products/developmentintegration/pocketbuilder>).

See also

- UltraLite.NET: [UltraLite - .NET Programming](#)
- UltraLite C/C++ and embedded SQL: [UltraLite - C and C++ Programming](#)
- UltraLite for M-Business Anywhere: [UltraLite - M-Business Anywhere Programming](#)

To choose your programming interface

1. Choose your target platform(s). UltraLite supports Palm OS, Windows Mobile, Windows XP/embedded Windows XP, and Java OS.
2. For each platform you need to support, determine if the API supports that platform. Different APIs support different platforms. If you are doing cross-platform development, choose an API that supports all of your intended targets.

Use this support matrix to help you quickly identify your development options.

Deployment targets	UltraLite for C/C++ and embedded SQL	UltraLite.NET ¹	UltraLite for M-Business Anywhere ²
Palm OS	version 4 +	N/A	version 5.0 +
Windows Mobile	CE 5.0 +	CE 5.0 with .NET compact framework 2.0	version 3.0 +

Deployment targets	UltraLite for C/C++ and embedded SQL	UltraLite.NET ¹	UltraLite for M-Business Anywhere ²
Embedded Windows XP	Supported	.NET framework 2.0	version 5.0 +
Java	Java SE (1.5 and later) RIM BlackBerry OS (4.1 and later)	N/A	N/A

¹ Development as an extension to Microsoft Visual Studio.NET. The driver supports ADO.NET versions 2.0.

² For browser-based deployments of UltraLite programming in JavaScript.

3. Consider the effects of the following requirements, and then finalize your selection:

SQL Anywhere compatibility If database compatibility with SQL Anywhere is a concern, consider the following:

- SQL Anywhere embedded SQL support provides a common programming interface for UltraLite and SQL Anywhere databases.
- ADO.NET provides common programming models that are shared between UltraLite components and SQL Anywhere.

Maintaining a common interface may be particularly useful on platforms such as Windows Mobile, where both databases are available. If you need to move from UltraLite to a SQL Anywhere database, you should use embedded SQL or ADO.NET to make application migration easier.

Tip

Even though a common interface exists, good development practice dictates that you create an abstract data-access layer when writing your application.

Simplified deployments If simplifying your UltraLite deployment is an issue, consider programming with the M-Business Anywhere API. Your end-users can download both the UltraLite application and the database concurrently.

Application size If creating the smallest application footprint is a priority, you should program your application with the C/C++ API. These applications typically yield the best performance and still maintain a small application file size.

Application performance Each API has yields a different performance result. For example, the performance of UltraLite.NET might not be as good as the performance of UltraLite C++, and UltraLite embedded SQL might perform better than UltraLite C++.

Using UltraLite Databases

This section describes how to create and use UltraLite databases.

Creating and configuring UltraLite databases	29
Connecting to an UltraLite database	43
Deploying UltraLite to devices	55
Working with UltraLite databases	69
UltraLite CustDB samples	93
UltraLite performance and optimization	107

Creating and configuring UltraLite databases

Contents

Creating an UltraLite database	30
Choosing database creation parameters for UltraLite	34
UltraLite character sets	36
Securing UltraLite databases	40

There are two 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.

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

If you do not want to use default values, you must set the creation parameter values when creating a database. The values cannot be changed after the database is created. If you need to change any database property, you must re-create the database, specifying a different value for the creation parameter.

Sharing a database among multiple platforms

Within the configuration differences imposed by different operating systems, you might be able to copy the database from one device to another. If you are unsure of property compatibility among multiple platforms, create a database in Sybase Central with the **UltraLite Create Database Wizard**. This wizard handles the file compatibility logic for you. In so doing, it prevents you from creating a file that is not supported on your combination of deployment devices.

See also

- [“UltraLite optimization strategies” on page 121](#)
- [“Choosing database creation parameters for UltraLite” on page 34](#)
- [“Change UltraLite persistent database option settings” on page 224](#)
- [“Working with UltraLite databases” on page 69](#)

Creating an UltraLite database

The UltraLite 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 `ulcreate` utility. See [“Create a database with the Create Database Wizard” on page 30](#).

The MobiLink Create Synchronization Model Wizard Choose this method if you are creating a synchronization system with remote UltraLite databases and a centralized consolidated database.

See [“Create an UltraLite database from a MobiLink synchronization model” on page 31](#).

The command line Choose any of the following utilities:

- Use the `ulcreate` utility if you are familiar with database creation parameters and want a fast alternative to creating a new, empty UltraLite database. This option is particularly useful for creating databases in batch operations. See [“Create an UltraLite database from the command prompt” on page 30](#).
- Use the `ulinit` utility if you want to create a new, empty UltraLite database sourced from a SQL Anywhere reference database schema. See [“Create an UltraLite database from a SQL Anywhere reference database” on page 31](#).
- Use the `ulload` 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 [“Create an UltraLite database from XML” on page 32](#).

Create a database with the Create Database Wizard

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

To create a new UltraLite database (Sybase Central)

1. Choose **Start » Programs » SQL Anywhere 11 » Sybase Central**.
2. Choose **Tools » UltraLite 11 » Create Database**.
3. Follow the instructions in the **Create Database Wizard**.

Create an UltraLite database from the command prompt

Use the `ulcreate` utility to create a database from a command prompt. With this utility, you can include utility options to configure the database.

To create a new UltraLite database (command prompt)

Run the `ulcreate` utility, including any necessary options. For example, to create a case-sensitive UTF-8 database called `test.udb`, overwriting the database if it already exists, run the following command:

```
ulcreate -c "DBF=test.udb" -o "case=respect;utf_encoding=1" -y
```


Supplying a database file in a connection string is an alternative to specifying the database file after the other options. For example:

```
ulcreate -o "case=respect;utf_encoding=1" -y test.udb
```

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.

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 “[MobiLink models](#)” [*MobiLink - Getting Started*].

Create 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.

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.

To initialize/extract a new UltraLite database from a reference database (command prompt)

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 “[Creating a database](#)” [*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
- Date order
- Nearest century
- Precision
- Scale

- Time format
 - Timestamp format
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 [“Working with UltraLite tables and columns” on page 70](#).
 - **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 79](#).
 - **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 [“Publications in UltraLite” on page 137](#).

Performance 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 [“Using index scans” on page 108](#).

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

For example, to initialize an UltraLite database called `customer.udb` with tables contained in two distinct publications, run the following command. `Pub1` may contain a small subset of tables for priority synchronization, while `Pub2` could contain the bulk of the data.

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

See also

- [“Create a database with the Create Database Wizard” on page 30](#)
- [“Create an UltraLite database from the command prompt” on page 30](#)
- [“Create an UltraLite database from XML” on page 32](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“Choosing database creation parameters for UltraLite” on page 34](#)
- [“Upgrading UltraLite” \[SQL Anywhere 11 - Changes and Upgrading\]](#)

Create an UltraLite database from XML

You can use XML as an intermediate format for managing your UltraLite database, provided that 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 that you can check into a version control system.

UltraLite cannot use an arbitrary XML file. The `install-dir\UltraLite` directory contains a `usm.xsd` file. Use this file to review the XML format.

To create an UltraLite database from an XML file

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. See [“Exporting relational data as XML” \[SQL Anywhere Server - SQL Usage\]](#).
 - 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

- [“Create a database with the Create Database Wizard” on page 30](#)
- [“Create an UltraLite database from the command prompt” on page 30](#)
- [“Create an UltraLite database from a SQL Anywhere reference database” on page 31](#)
- [“Upgrading UltraLite” \[SQL Anywhere 11 - Changes and Upgrading\]](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“Choosing database creation parameters for UltraLite” on page 34](#)

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

- [“Creating an UltraLite database” on page 30](#)
- UltraLite for C/C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)

Choosing database creation parameters for UltraLite

UltraLite database creation parameters are recorded as *name=value* pairs in the system tables that make up the schema of the database. By storing these creation parameters in system tables, users and/or applications can access them in the same way as other data. See [“sysuldata system table” on page 298](#).

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. See [“Accessing UltraLite database properties” on page 218](#).

You can also access the database properties programmatically from the UltraLite application by calling the `GetDatabaseProperty` function appropriate to the API.

For API-specific details, see:

- UltraLite C/C++: [“GetDatabaseProperty function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“GetDatabaseProperty method” \[UltraLite - .NET Programming\]](#)
- M-Business: [“getDatabaseProperty method” \[UltraLite - M-Business Anywhere Programming\]](#)

In addition to these database creation parameters, you can further configure other aspects of your database with database options and connection parameters. See:

- [“UltraLite database options” on page 219](#)
- [“Connecting to an UltraLite database” on page 43](#)
- [“UltraLite connection parameters” on page 225](#)

Property	Description
case	Sets the case-sensitivity of string comparisons in the UltraLite database. See “UltraLite case creation parameter” on page 183 .
checksum_level	Sets the level of checksum validation in the database. See “UltraLite checksum_level creation parameter” on page 184 .
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 36 and “UltraLite collation creation parameter” on page 186 and “UltraLite utf8_encoding creation parameter” on page 211 .
date_format	Sets the default string format in which dates are retrieved from the database. See “UltraLite date_format creation parameter” on page 187 .
date_order	Controls the interpretation of date ordering of months, days, and years. See “UltraLite date_order creation parameter” on page 190 .

Property	Description
fips	Controls AES FIPS compliant data encryption, by using a Certicom certified cryptographic algorithm. FIPs encoding is a form of strong encryption. See “Securing UltraLite databases” on page 40 and “UltraLite fips creation parameter” on page 192 .
max_hash_size	Sets the default index hash size in bytes. See “UltraLite max_hash_size creation parameter” on page 194 .
nearest_century	Controls the interpretation of two-digit years in string-to-date conversions. See “UltraLite nearest_century creation parameter” on page 196 .
obfuscate	Controls whether data in the database is obfuscated. Obfuscation is a form of simple encryption. See “Securing UltraLite databases” on page 40 and “UltraLite obfuscate creation parameter” on page 198 .
page_size	Defines the database page size. See “UltraLite page_size creation parameter” on page 199 .
precision	Specifies the maximum number of digits in decimal point arithmetic results. See “UltraLite precision creation parameter” on page 201 .
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 203 .
time_format	Sets the format for times retrieved from the database. See “UltraLite time_format creation parameter” on page 205 .
timestamp_format	Sets the format for timestamps retrieved from the database. See “UltraLite timestamp_format creation parameter” on page 207 .
timestamp_increment	Determines how the timestamp is truncated in UltraLite. See “UltraLite timestamp_increment creation parameter” on page 210 .
utf8_encoding	Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode. See “UltraLite character sets” on page 36 and “UltraLite utf8_encoding creation parameter” on page 211 .

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. Generally speaking, 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.

See also

- [“UltraLite platform requirements for character set encoding” on page 37](#)
- [“UltraLite collation creation parameter” on page 186](#)
- [“UltraLite utf8_encoding creation parameter” on page 211](#)
- [“Overview of character sets, encodings, and collations” \[*SQL Anywhere Server - Database Administration*\]](#)
- [“UltraLite connection parameters” on page 225](#)
- [“Character set considerations” \[*MobiLink - Server Administration*\]](#)
- [“UltraLite case creation parameter” on page 183](#)
- [“Securing UltraLite databases” on page 40](#)

Change the character set and collation

You cannot change the character set, collation, or encoding properties of the database once UltraLite writes the file. You must recreate the database.

To recreate an UltraLite database with a new collation

1. Unload the database.

You can use either the **Unload Database Wizard** in Sybase Central or the `ulunload` utility. See [“UltraLite Unload Database utility \(ulunload\)” on page 281](#).

2. Create a new database with the appropriate collation.
3. Optional. Convert the data to match the new collation if required.
4. Reload the database.

You can use either the **Load Database Wizard** in Sybase Central or the ulload utility. See [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#).

UltraLite platform requirements for character set encoding

Each platform has specific character set and encoding requirements.

Palm OS

Never use UTF-8 encoding; Palm does not support Unicode characters. Always choose a collation that matches the code page of your intended device.

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 211](#)
- [“Overview of character sets, encodings, and collations” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite connection parameters” on page 225](#)
- [“Character set considerations” \[MobiLink - Server Administration\]](#)
- [“Securing UltraLite databases” on page 40](#)

UltraLite supported collations

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

```
ulcreate -l
```

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

Collation label	Description
1251CYR	Code Page 1251, Windows Cyrillic
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

Collation label	Description
ISO_BINENG	Binary ordering, English ISO/ASCII 7-bit letter case mappings
UTF8BIN	UTF-8, 8-bit multibyte encoding for Unicode, binary ordering

Securing UltraLite databases

By default, UltraLite databases are unencrypted on disk. When using a viewing tool such as a hex editor, text and binary columns can be read. To encrypt data for greater security, consider the following options:

- **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 198](#).
- **AES 128-bit strong encryption** UltraLite databases can be strongly encrypted using the AES 128-bit algorithm, which is the same algorithm used to encrypt SQL Anywhere databases. 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 wizards by selecting the **Encrypt Database** option and then selecting **AES Strong Encryption**. Using a creation utility, you set the key with the `key` connection parameter. This same parameter is used by end users when connecting to the database after it has been created. You do not need any special configuration to use AES encryption on your device. See [“UltraLite fips creation parameter” on page 192](#).
- **AES FIPS 140-2 compliant encryption** UltraLite provides encryption libraries compliant with the FIPS 140-2 US and Canadian government standard (using a Certicom certified cryptographic module). You set FIPS compliant encryption with the `fips` creation parameter. The user must supply the required key in their connection string. AES FIPS encryption requires that you configure your device appropriately. See [“Deploy UltraLite with AES_FIPS database encryption” on page 59](#), and [“UltraLite fips creation parameter” on page 192](#).

Tip

The MobiLink server's synchronization streams can use public/private keys to encrypt streamed data. For ease of deployment, you can embed these certificates in the UltraLite database when you create it. See [“Configuring MobiLink clients to use transport-layer security” \[SQL Anywhere Server - Database Administration\]](#).

Notes

Both the FIPS and AES database encryption types use 128-bit AES. Therefore, if you use the same encryption key, the database is encrypted the same way irrespective of the standard you choose.

Caution

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

- UltraLite for C++: “[ChangeEncryptionKey function](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[ChangeEncryptionKey method](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[changeEncryptionKey method](#)” [*UltraLite - M-Business Anywhere Programming*]

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

For strongly encrypted databases, be sure to 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.

See also

- “[UltraLite fips creation parameter](#)” on page 192
- “[UltraLite obfuscate creation parameter](#)” on page 198
- “[UltraLite DBKEY connection parameter](#)” on page 233
- “[Deploy UltraLite with TLS-enabled synchronization](#)” on page 60
- UltraLite.NET: “[Encryption and obfuscation](#)” [*UltraLite - .NET Programming*]
- UltraLite for C++: “[Encrypting data](#)” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “[Database encryption and obfuscation](#)” [*UltraLite - M-Business Anywhere Programming*]

Connecting to an UltraLite database

Contents

UltraLite database connection parameters	44
UltraLite user authentication	51
Storing UltraLite parameters with the ULSQLCONNECT environment variable	53

Any application that uses a database must establish a connection to that database before any transactions can occur. An application can be an UltraLite command line utility, a connection window from either Sybase Central tool or Interactive SQL, 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.

UltraLite database connection parameters

Some of the connection parameters you configure overlap with properties you define at creation time. At creation time, you typically are defining the property required. At connection time, you must supply the value that was configured at creation time. It is important that you understand the differences and commonalities of these shared settings so you know what is expected and when.

Parameter name	Description
CACHE_SIZE	Defines the size of the database cache. See “UltraLite CACHE_SIZE connection parameter” on page 226 .
CON	Specifies a name of the current connection. See “UltraLite CON connection parameter” on page 230 .
DBF, and CE_FILE, PALM_FILE, or NT_FILE	<p>At creation time these parameters set the location of the database. 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 231 • “UltraLite CE_FILE connection parameter” on page 227 • “UltraLite PALM_FILE connection parameter” on page 241 • “UltraLite NT_FILE connection parameter” on page 237
DBN	Identifies a running database by name rather than file name. See “UltraLite DBN connection parameter” on page 234 .
DBKEY	At creation-time, this parameter sets the encryption key to use. For subsequent connections, names and then passes the same encryption key for the database. If the incorrect key is named, the connection fails. See “UltraLite DBKEY connection parameter” on page 233 .
MIRROR_FILE	Specifies the name of a database mirror file. See “UltraLite MIRROR_FILE connection parameter” on page 235 .
PALM_ALLOW_BACKUP	Controls backup behavior over HotSync on Palm devices. See “UltraLite PALM_ALLOW_BACKUP connection parameter” on page 240 .
PWD	At creation-time, sets the initial password for a user. For subsequent connections, supplies the password for the user ID. See “UltraLite PWD connection parameter” on page 243 .

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 244.
START	Specifies the location of the UltraLite engine executable and then starts it. See “UltraLite START connection parameter” on page 245.
UID	At creation time, sets the initial user ID. For subsequent connections, 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 246.

See also

- [“Interpreting user ID and password combinations”](#) on page 51
- [“UltraLite page_size creation parameter”](#) on page 199

Supplying UltraLite connection parameters

How connection information is collected depends on how systematic or automated you require the input to be. The more systematic the input, the more reliable the connection information is.

Connection details can be collected via different methods, depending on whether you are connecting from a custom UltraLite application or from one of the SQL Anywhere administration tools for UltraLite.

Method	Administration tools	Custom applications
<p>Prompt the end user at connection time when you require a user to authenticate as one of the four supported database users. The UltraLite graphical administration tools use a connection object.</p> <p>Where possible, use either the <code>ULConnectionParms</code> or <code>ConnectionParms</code> object. It provides easier checking and a more systematic interface than using a connection string that is an argument for the <code>Open</code> method. See:</p> <ul style="list-style-type: none"> • UltraLite.NET: “Authenticating users” [<i>UltraLite - .NET Programming</i>] • UltraLite C/C++ : “Authenticating users” [<i>UltraLite - C and C++ Programming</i>] • UltraLite for M-Business Anywhere: “Authenticating users” [<i>UltraLite - M-Business Anywhere Programming</i>] • UltraLite for embedded SQL: “Authenticating users” [<i>UltraLite - C and C++ Programming</i>] 	X	X

Method	Administration tools	Custom applications
<p>Use a connection string if user authentication is not required. Common reasons for not authenticating a user 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. The UltraLite command line utilities typically use a connection string if a connection to a database is required. You can also program your UltraLite application to read the values from a stored file, or hard code it into your application. See:</p> <ul style="list-style-type: none"> ● UltraLite.NET: “Connecting to a database” [UltraLite - .NET Programming] ● UltraLite C/C++ : “Connecting to a database” [UltraLite - C and C++ Programming] ● UltraLite for M-Business Anywhere: “Connecting to an UltraLite database” [UltraLite - M-Business Anywhere Programming] ● UltraLite for embedded SQL: “Connecting to a database” [UltraLite - C and C++ Programming] 	X ¹	X ³
<p>Use the ULSQLCONNECT environment variable if you want to store connection parameters you use repeatedly. By storing parameters, you don't need to provide them repeatedly during the development phase. Values you supply as a parameter in ULSQLCONNECT become defaults for the UltraLite desktop administration tools.</p> <p>All UltraLite desktop administration tools check the ULSQLCONNECT values for any missing parameters not supplied in a connection string, following parameter precedence rules. To override these values, supply the alternate value in the connection string. See “Storing UltraLite parameters with the ULSQLCONNECT environment variable” on page 53.</p>	X ²	N/A

¹ Typically user-supplied.

² For desktop administration tools only.

³ Typically hard coded or stored in a file.

See also

- “[Precedence of connection parameters for UltraLite administration tools](#)” on page 49

Specifying file paths in an UltraLite connection parameter

The physical storage of your device determines:

- Whether the database is saved as a file.
- What naming conventions you must follow when identifying your database.

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
```

Windows Mobile tip

You can use the UltraLite administration tools to administer databases already deployed to an attached device. See [“Windows Mobile” on page 47](#).

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

```
Connection = DatabaseMgr.OpenConnection( "UID=JDoe;PWD=ULdb;  
CE_FILE=\database\MyCEDB.udb;PALM_FILE=MyPalmDB" )
```

Windows desktop

Desktops allow either absolute or relative paths.

Windows Mobile

Windows Mobile devices require that all paths be absolute.

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.

Palm OS

Palm OS does not necessarily use the concept of file paths. Therefore, how you define it depends on the store type (that is, record-based, or VFS).

File-based stores (VFS) For databases on a VFS volume, define the file with the following syntax:

```
vfs: [ volume-label: | volume-ordinal: ] filename
```

You can set the *volume-label* as **INTERNAL** for the built-in drive, or **CARD** for either an expansion card or the label name of the volume. There is no default string for the *volume-label*.

Alternatively, you can set the *volume-ordinal* to identify the volume. Since the enumeration of mounted volumes can vary, ensure that you set the correct ordinal volume for your chosen internal or external volume. The default value is 0 (which is the first volume enumerated by the platform).

For the *filename*, always specify the absolute file path, following the file and path naming convention of Palm OS. If directories specified in the path do not already exist, they are created.

Record-based data stores For record-based data stores, database names must follow all conventions for Palm OS database names. For example, database names cannot exceed the 32 character limit and cannot contain a path.

Also ensure that you use the appropriate value for DBF or PALM_FILE according to the database's location.

- Use the *.pdb* extension with DBF when you store a Palm OS database anywhere other than on the device itself (for example, with ulload).
- Once you move the file to the device, the *.pdb* extension is dropped by the HotSync conduit. For example, if the database you created on the desktop is named *CustDB.pdb*, then when you deploy it to the device, the filename changes to *CustDB*.

Note

You cannot deploy databases to VFS volumes using the Palm install tool. Instead, you need to copy the database directly onto the media with a card reader, or some other tool.

See also

- [“UltraLite DBF connection parameter” on page 231](#)
- [“UltraLite NT_FILE connection parameter” on page 237](#)
- [“UltraLite CE_FILE connection parameter” on page 227](#)
- [“UltraLite PALM_FILE connection parameter” on page 241](#)

Opening UltraLite connections with connection strings

A connection string is a set of parameters that is passed from an application to the runtime so that a connection can be defined and established.

There are three steps that take place before a connection to a database is opened:

1. The parameter definition phase

You must define the connection via a combination of supported parameters. Some connection parameters are always required to open a connection. Others are used to adjust database features for a single connection.

How these parameters are supplied can vary depending on whether you are connecting from an UltraLite administration tool or an UltraLite application. See [“Supplying UltraLite connection parameters” on page 45](#).

2. The string assembly phase

Either you or the application assembles the supplied parameters into a string. Connection strings contain a list of parameters defined as *keyword=value* pairs in a semicolon delimited list. See [“Assembling parameters into UltraLite connection strings” on page 49](#).

For example, a connection string fragment that supplies a file name, user ID, and password is written as follows:

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

3. The transmittal phase

When the connection string has been assembled, it is passed to the database via an UltraLite API to the UltraLite runtime for processing. If the connection attempt is validated, the connection is granted. Connection failures can occur if:

- The database file does not exist.
- Authentication was unsuccessful.

Assembling parameters into UltraLite connection strings

An assembly of connection parameters supplied in any application's connection code (be it an administration tool or a custom UltraLite application) is called a **connection string**. In some cases, applications parse the fields of a ConnectionParms object into a string. In others, you 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 named database file until you supply this string:

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

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

Precedence of connection parameters for UltraLite administration tools

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

- If specified, the CE_FILE, NT_FILE, and PALM_FILE parameters always take precedence over DBF.
- If you specify two DBF parameters, the last one specified takes precedence.
- 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.

- The ULSQLCONNECT environment variable is checked for connection parameters that are not supplied in the connection string.
- 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

- [“Storing UltraLite parameters with the ULSQLCONNECT environment variable” on page 53](#)

UltraLite user authentication

You cannot disable UltraLite user authentication. A successful connection requires that a user be authenticated. Unlike SQL Anywhere, UltraLite database users are created and managed solely for the purposes of authentication—not for the purposes of object ownership. Once a user authenticates and connects to the database, the user has unrestricted access to everything in that database, including schema data.

You can only add or modify UltraLite users from an existing connection. Therefore, any changes to your UltraLite user base can only occur after you have connected with a valid user ID and password.

If this is your first time connecting, the UID and PWD required are the same values set when you first created the database. If you did not set an initial user, then you must authenticate with the defaults of **UID=DBA** and **PWD=sql**.

Bypass authentication

Although you cannot disable authentication, you can bypass it by using UltraLite defaults when you create and connect to the database.

If you do not supply the UID *and* the PWD parameters, UltraLite assumes the defaults of **UID=DBA** and **PWD=sql**.

To bypass authentication in UltraLite

1. Do not set the UID and PWD connection parameters when you create a database.
2. Do not delete or modify the default user in your UltraLite database.
3. Do not set the UID and PWD connection parameters when you connect to the database you have created.

See also

- [“Limitations” on page 87](#)
- [“Working with UltraLite users” on page 87](#)
- [“Interpreting user ID and password combinations” on page 51](#)

Interpreting user ID and password combinations

UltraLite allows you to set one, none, or both of the UID and PWD parameters—except when a partial definition prevents a user from being identified by UltraLite. The table below tells you how UltraLite interprets incomplete user definitions.

If you create a database with...	It has this impact...
No user ID <i>and</i> password.	UltraLite creates a default user with a UID of DBA and PWD of sql . You do not need to supply these connection parameters upon future connection attempts.

If you create a database with...	It has this impact...
<p>The user ID parameter only.</p> <p>Examples:</p> <ul style="list-style-type: none">• UID=JaneD• UID=JaneD;PWD=• UID=JaneD;PWD=""	<p>UltraLite creates a default user with a UID of JaneD and an empty PWD. When connecting, you must always supply the UID parameter. The PWD parameter is not required.</p>
<p>The password parameter only.</p> <p>Examples:</p> <ul style="list-style-type: none">• PWD=3saBys• UID=;PWD=3saBys• UID="";PWD=3saBys	<p>UltraLite generates an error. UltraLite cannot set a password without a user ID.</p>

See also

- [“UltraLite user authentication” on page 51](#)
- [“Limitations” on page 87](#)
- [“Working with UltraLite users” on page 87](#)

Storing UltraLite parameters with the ULSQLCONNECT environment variable

The ULSQLCONNECT environment variable is optional, and is not set by the installation program. ULSQLCONNECT contains a list of parameters defined as *keyword=value* pairs in a semicolon delimited list.

Use ULSQLCONNECT to avoid having to supply the same connection parameters repeatedly to UltraLite administration tools during development. You cannot use ULSQLCONNECT for custom applications.

Caution

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

To set ULSQLCONNECT for UltraLite desktop tools

1. Run the following command:

```
set ULSQLCONNECT="parameter=value; ..."
```

2. If an administration tool 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.

See also

- [“Precedence of connection parameters for UltraLite administration tools” on page 49](#)
- [“Supplying UltraLite connection parameters” on page 45](#)

Example

To use ULSQLCONNECT 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 *-c* connection option for these defaults values—unless you need to override these values.

For example, if you were using *ulload* to add additional information to your database from an *extra.xml* file, you would run the following command:

```
ulload -a extra.xml
```

Deploying UltraLite to devices

Contents

Deploy Multiple UltraLite applications with the UltraLite engine	57
Deploy UltraLite with AES_FIPS database encryption	59
Deploy UltraLite with TLS-enabled synchronization	60
Deploy the UltraLite HotSync conduit	62
Deploy the ActiveSync provider for UltraLite	64
Register applications with the ActiveSync Manager	66
Deploying UltraLite schema upgrades	67

In the majority of cases, development occurs on a Windows desktop 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.

Initial installation

Initially installing your UltraLite solution on a device is a required step so that the device can be continuously maintained.

Ongoing maintenance

The majority of the benefits and costs are associated with the application being deployed and used in the field. To that end, UltraLite supports includes various mechanisms to deliver applications and synchronize data, both of which are key features in any UltraLite deployment that make UltraLite so flexible.

Technique	Windows Mobile	Palm OS	Java
File transfers	yes	record based and VFS	
UltraLite as a MobiLink synchronization client	yes	record based and VFS	

Technique	Windows Mobile	Palm OS	Java
On device file system browsing tool or card reader		VFS only (true) ¹	
ActiveSync	yes	n/a	
HotSync	n/a	record based only	
ALTER DATABASE SCHEMA	yes	yes	

¹ You cannot deploy UltraLite databases with the Palm install tool, if the target is a VFS volume. Instead, you need to copy the database directly onto the media using a card reader, or some other tool.

See also

- [“Using MobiLink file transfers” on page 142](#)
- [“UltraLite clients” on page 125](#)
- [“HotSync on Palm OS” on page 146](#)
- [“ActiveSync on Windows Mobile” on page 150](#)
- [“Deploying UltraLite schema upgrades” on page 67](#)
- [“UltraLite ALTER DATABASE SCHEMA FROM FILE statement” on page 450](#)

Deploy Multiple UltraLite applications with the UltraLite engine

The UltraLite engine is a data management module that manages concurrent UltraLite database connections from applications on 32-bit Windows desktops and Windows Mobile only. The engine is automatically installed to the desktop with the SQL Anywhere installer. Therefore, you only need to deploy the engine to the Windows Mobile device.

To deploy uleng to a Windows Mobile device

1. Copy the *uleng11.exe* file and the appropriate **.dll* files. The **.dll* files you copy should include any database encryption, synchronization encryption, or compression **.dll* files required.

Filename	Basic	ECC TLS	RSA TLS	FIPS RSA TLS	HTT PS	Compression	FIPS AES data encryption
<i>uleng11.exe</i>	X	X	X	X	X	X	X
<i>mlcecc11.dll</i> ¹		X			X		
<i>mlcrsa11.dll</i> ¹			X		X		
<i>mlcrsafips11.dll</i>				X	X		
<i>mlczlib11.dll</i>						X	
<i>sbgse2.dll</i>				X	X		X
<i>ulfips11.dll</i>							X
<i>ulrt11.dll</i> ²	X	X	X	X	X	X	X

¹ File not needed if an application links directly against *ulecc.lib* and *ulrsa.lib* respectively.

² File only needed if an application links against *ulimp.lib*.

2. Save the files to an appropriate directory. In most cases you would use one of the following destination directories:
 - The `\windows` directory. This location is the recommended location, as the client automatically looks for the engine in this location.
 - The directory for other UltraLite application files.
3. If you use any location other than the `\windows` directory, include the `START` connection parameter. This parameter starts the UltraLite engine when the application connects the UltraLite database.

For example, a connection string to the database or connection code for a Windows Mobile client application, might use this `START` parameter:

```
"START=\Program Files\MyApp\ulengl1.exe"
```

See also

- [“Deploy UltraLite with AES_FIPS database encryption” on page 59](#)
- [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#)

Deploy UltraLite with AES_FIPS database encryption

Strong database encryption technology makes a database inoperable and inaccessible without a key (a type of a password). An algorithm encodes the information contained in your database and transaction log files so they cannot be deciphered. However, database encryption requires that you deploy the appropriate number of files with your database.

When you connect to UltraLite with the `-fips` option, you can run databases encrypted with AES or AES_FIPS strong encryption. To ensure you are running with AES_FIPS, use `-fips=1`.

If you are encrypting your database with AES FIPS encryption, you must configure and deploy your device for each platform.

To set up your application and device for an AES FIPS encrypted UltraLite database

1. Create an UltraLite database with the property `fips=1`. See [“UltraLite fips creation parameter” on page 192](#).
2. Use the following connection parameter in your application's connection string: `DBKEY=key`. See [“UltraLite DBKEY connection parameter” on page 233](#).
3. On Palm OS, enable database encryption by calling `ULEnableRsaFipsStrongEncryption`. See [“ULEnableFIPSStrongEncryption function” \[UltraLite - C and C++ Programming\]](#).
4. On Palm OS, ensure that you link to the following files in addition to `ulrt.lib`:
 - `ulfips.lib`
 - `gse1st.lib`
5. Ensure that you deploy the appropriate files to your device.
 - Windows desktop, Windows Mobile, require `ulfips11.dll` and `sbgse2.dll`. The Windows Mobile component also requires the component DLL files.
 - Palm OS requires `libsgse_4i.prc`.

See also

- [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#)
- UltraLite.NET: [“Encryption and obfuscation” \[UltraLite - .NET Programming\]](#)
- UltraLite for C++: [“Encrypting data” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Database encryption and obfuscation” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“UltraLite fips creation parameter” on page 192](#)
- [“Securing UltraLite databases” on page 40](#)

Deploy UltraLite with TLS-enabled synchronization

Separately licensed component required

ECC encryption and FIPS-certified encryption require a separate license. All strong encryption technologies are subject to export regulations.

See “Separately licensed components” [[SQL Anywhere 11 - Introduction](#)].

UltraLite client applications of MobiLink must be configured to use enable TLS synchronization. Transport-layer security enables encryption, tamper detection, and certificate-based authentication. See “Introduction to transport-layer security” [[SQL Anywhere Server - Database Administration](#)].

Platform support

RSA, ECC, and FIPS encryption are not available on all platforms. For information about which platforms support which encryption method, see <http://www.sybase.com/detail?id=1061806>.

To set up TLS synchronization on an UltraLite client application and device

1. Enable encrypted synchronization by calling one of the following in your application code:
 - To enable RSA encryption, call `ULEnableRsaSyncEncryption`. See “[ULEnableRsaSyncEncryption function](#)” [[UltraLite - C and C++ Programming](#)].
 - To enable ECC encryption, call `ULEnableEccSyncEncryption`. See “[ULEnableEccSyncEncryption function](#)” [[UltraLite - C and C++ Programming](#)].
 - To enable FIPS RSA encryption, call `ULEnableRsaFipsSyncEncryption`. This function is only required for Palm OS clients. See “[ULEnableRsaFipsSyncEncryption function](#)” [[UltraLite - C and C++ Programming](#)].
2. Set the synchronization information stream to either TLS or HTTPS.
3. If you are enabling ECC or FIPS encryption, you also need to:
 - **ECC** Set the `tls_type` network protocol option to **ECC**. See “[tls_type](#)” [[MobiLink - Client Administration](#)].
 - **FIPS** Set the `fips` network protocol option to **Yes**. See “[fips](#)” [[MobiLink - Client Administration](#)].
4. Ensure that you have linked to the appropriate libraries:

Platform	Linking	RSA encryption	ECC encryption	FIPS encryption
Windows desk-top	static ¹	<i>ulrsa.lib</i>	<i>ulecc.lib</i>	none
Windows desk-top	dynamic ²	none	none	none

Platform	Linking	RSA encryption	ECC encryption	FIPS encryption
Windows Mobile	static ¹	<i>ulrsa.lib</i>	<i>ulecc.lib</i>	none
Windows Mobile	dynamic ²	none	none	none
Palm OS	static ¹	<i>ulrsa.lib</i>	<i>ulecc.lib</i>	<i>ulfips.lib, gse1st.lib</i>

¹ You must also link to *ulrt.lib*.

² You must also link to *ulimp.lib*.

5. Ensure that the appropriate files are copied to the device:

Platform	Linking	RSA encryption	ECC encryption	FIPS encryption
Windows desktop	static	none	none	<i>mlcrsafips11.dll</i> <i>sbgse2.dll</i>
Windows desktop	dynamic ¹	<i>mlcrsa11.dll</i>	<i>mlcecc11.dll</i>	<i>mlcrsafips11.dll</i> <i>sbgse2.dll</i>
Windows Mobile	static	none	none	<i>mlcrsafips11.dll</i> <i>sbgse2.dll</i>
Windows Mobile	dynamic ¹	<i>mlcrsa11.dll</i>	<i>mlcecc11.dll</i>	<i>mlcrsafips11.dll</i> <i>sbgse2.dll</i>
Palm OS	static	none	none	<i>libsbgse_4i.prc</i>
Windows Mobile components and UltraLite engine	static ²	<i>mlcrsa11.dll</i>	<i>mlcecc11.dll</i>	<i>mlcrsafips11.dll</i> <i>sbgse2.dll</i>

¹ You must also deploy *ulrt11.dll*.

² You must also deploy your component *.dll* file and/or *uleng11.exe*.

See also

- “Configuring UltraLite clients to use transport-layer security” [[SQL Anywhere Server - Database Administration](#)]
- UltraLite.NET: “Encryption and obfuscation” [[UltraLite - .NET Programming](#)]
- UltraLite for C++: “Encrypting data” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Database encryption and obfuscation” [[UltraLite - M-Business Anywhere Programming](#)]

Deploy the UltraLite HotSync conduit

The UltraLite HotSync conduit 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 either Windows Mobile ActiveSync and Palm OS HotSync software.

During development you install UltraLite onto your desktop with the SQL Anywhere installer. However, you then need to deploy the required HotSync conduit files onto the end user's computer. HotSync Manager must be installed on your computer for the installer to detect the software and run.

UltraLite HotSync conduit files

- ***install-dir\win32\Condmgr\condmgr.dll*** The utility DLL that locates the HotSync installation path and registers the conduit with HotSync.
- ***install-dir\Bin32\ulcond11.exe*** The UltraLite HotSync Conduit Installation utility that installs or removes the UltraLite HotSync conduit on the desktop computer. See “[UltraLite HotSync Conduit Installation utility for Palm OS \(ulcond11\)](#)” on page 263.
- ***install-dir\Bin32\dbhsync11.dll*** The conduit DLL that is called by HotSync.
- ***install-dir\Bin32\dblgen11.dll*** The language resource library. For languages other than English, the letters *en* in the file name are replaced by a two-letter abbreviation for the language, such as *dblge11.dll* for German, or *dblja11.dll* for Japanese.
- **Stream DLLs** Optional. The stream DLL required for encrypted network communication between the UltraLite HotSync conduit and the MobiLink server.
 - For RSA encryption with TLS and HTTPS, *install-dir\Bin32\mlcrsa11.dll*.
 - For ECC encryption with TLS and HTTPS, *install-dir\Bin32\mlcecc11.dll*.
 - For RSA FIPS encryption with TLS and HTTPS, *install-dir\Bin32\mlcrsafips11.dll*.

Separately licensed component required

ECC encryption and FIPS-approved encryption stream DLLs require a separate license. All strong encryption technologies are subject to export regulations.

To order a separately licensed component, visit http://www.iAnywhere.com/products/separately_licensed_components.html.

For information about component and platform support, see <http://www.sybase.com/detail?id=1002288>.

To deploy and register the UltraLite HotSync conduit

1. On the end-user's desktop, create the following directories:
 - *MyDir\win32*
 - *MyDir\win32\condmgr*
2. Deploy a copy of the following files to the *MyDir\win32* directory:
 - *ulcond11.exe*
 - *dbhsync11.dll*

- *dblgen11.dll*
3. Deploy a copy of the *Condmgr.dll* file to the *MyDir\win32\condmgr* directory.
 4. Locate the following registry key:

```
HKEY_CURRENT_USER\Software\Sybase\SQL Anywhere\11.0\
```
 5. Create a value named **Location** in this key and set this value data as the root deployment folder for the conduit. For example, *MyDir*.
 6. If the end user requires a certificate to encrypt the communication stream, install the root certificate on the desktop computer so it can be accessed by the conduit.
 7. Run **ulcond11**, ensuring that you have set connection string for each the UltraLite database with either the **-c** option, and possibly the **-a** option. You must also set the correct creator ID.

This utility deploys and correctly configures the UltraLite HotSync conduit.

Tip

If you are using an encryption key, avoid setting the key in the connection string. This location can pose a security risk. Instead, allow the conduit to prompt the user for the key.

For example, the following command installs a conduit for the application with creator ID Syb2, named CustDB.

```
ulcond11 -c "DBF=custdb.udb;UID=DBA;PWD=sql" -n CustDB Syb2
```

8. If you did not include synchronization parameters in your UltraLite application's *ul_synch_info* structure, configure this information either in HotSync or use *ulcond11*. See [“Setting protocol options for MobiLink synchronization” on page 148](#).

To check that the HotSync conduit is properly deployed

- In the computer's system tray, right-click **HotSync Manager** and choose **Custom**.

A list of conduits appears for each HotSync user. Verify that the conduit is listed before proceeding to start the MobiLink server and test your synchronization operation.

Deploy 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 either Windows Mobile ActiveSync and Palm OS HotSync software.

During development you install UltraLite onto your desktop with the SQL Anywhere installer. However, when you deploy UltraLite to 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.
- **dblggen11.dll** The language resource library.

For a list of supported provider platforms, see <http://www.sybase.com/detail?id=1002288>.

To 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 "install-dir\UltraLite\ce\arm.50"  
"install-dir\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 “Register applications with the ActiveSync Manager” on page 66.

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**, choose **MobiLink Clients**, and click **Settings**.

See also

- [“Register applications with the ActiveSync Manager” on page 66](#)
- [“ActiveSync provider installation utility \(mlasinst\)” \[*MobiLink - Client Administration*\]](#)

Register applications with the ActiveSync Manager

You can register your application for use with ActiveSync either by using the ActiveSync provider install utility or using the ActiveSync Manager itself. This section describes how to use the ActiveSync Manager.

To register an application for use with the ActiveSync Manager

1. Launch ActiveSync.
2. From the ActiveSync window, choose **Options**.
3. From the list of information types, choose **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. See [“Assigning class names for applications”](#) [*UltraLite - C and C++ Programming*].
 - **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

- [“ActiveSync provider installation utility \(mlasinst\)”](#) [*MobiLink - Client Administration*]

Deploying UltraLite schema upgrades

To make schema upgrades, use the SQL statement ALTER DATABASE SCHEMA FROM FILE.

The upgrade process

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."

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 value for a row is **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.

To upgrade the UltraLite schema

1. Create a SQL script of DDL statements that define the new schema. The character set of the SQL script file must match the character set of the database you want to upgrade.

You should use either ulinit or ulunload 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.

- If you are using ulunload, use the -n and -s [*file*] options.
- If you are using ulinit, use the -l [*file*] option.

If you do not use either ulunload or ulinit, review the script and ensure that:

- You do not rename tables, columns, or publications. RENAME operations are not supported. If you rename a table, it is processed as a DROP TABLE and CREATE TABLE operation.
- You have not included non-DDL statements. Including non-DDL statements may not have the effect you expect.
- Words in the SQL statement are separated by spaces.
- Only one SQL statement can appear in each line.
- Comments are prepended with double hyphens (--), and only occur at the start of a line.

2. Backup the database against which the upgrade will be performed.
3. Run the new statement. 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 ALTER DATABASE SCHEMA FROM FILE statement” on page 450](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Unload Database utility \(ulunload\)” on page 281](#)
- [“UltraLite SQL statements” on page 447](#)
- [“Comments in UltraLite” on page 305](#)

Working with UltraLite databases

Contents

Working with UltraLite tables and columns 70

Working with UltraLite indexes 79

Working with UltraLite publications 83

Working with UltraLite users 87

Working with event notifications 89

UltraLite support for SQL passthrough 92

Working with 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 199](#)
- [“Database tables” \[SQL Anywhere 11 - Introduction\]](#)
- [“Creating databases in SQL Anywhere” \[SQL Anywhere Server - SQL Usage\]](#)
- [“UltraLite system tables” on page 289](#)

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.

To create an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. In the left pane, right-click **Tables** and choose **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, choose **Save Table**.

Interactive SQL

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

To 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

- [“UltraLite CREATE TABLE statement” on page 468](#)
- [“Adding a column to an UltraLite table” on page 72](#)

Using allsync and nosync suffixes

You can append either **_allsync** or **_nosync** to a table name to control data restriction for synchronization. You can use these suffixes as an alternative to using publications to control data restrictions. To control data priority, define one or more publications.

- If you create a table with a name ending in **_allsync**, all rows of that table are synchronized at each synchronization—even if they have not changed since the last synchronization.

Tip

You can store user specific or client specific data in allsync tables. You can then upload the data in the UltraLite table to a temporary table in the consolidated database on synchronization. Synchronization scripts can control the data and save you from having to maintain that data in the consolidated database.

- If you create a table with a name ending in **_nosync**, all rows of that table are excluded from synchronization. You can use these tables for persistent data that is not required in the consolidated database's table.

See also

- [“Working with UltraLite publications” on page 83](#)
- [“Designing synchronization in UltraLite” on page 133](#)
- [“Nosync tables in UltraLite” on page 136](#)
- [“Allsync tables in UltraLite” on page 136](#)
- [“UltraLite CustDB samples” on page 93](#)

Example

In the *CustDB.udb* sample database, you can see that one table was declared a nosync table because the table name is named `ULIdentifyEmployee_nosync`. Therefore, no matter how data changes in this table, it is never synchronized with MobiLink and information will not appear in the *CustDB.db* consolidated database.

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.

To 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 choose **New » Column**.
5. Set the attributes for the new column.
6. From the **File** menu, choose **Save Table**.

Interactive SQL

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

To 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

- “Choosing object names” [[SQL Anywhere Server - SQL Usage](#)]
- “Data types in UltraLite” on page 312
- “Choosing column data types” [[SQL Anywhere Server - SQL Usage](#)]
- “UltraLite CREATE TABLE statement” on page 468
- “UltraLite ALTER TABLE statement” on page 454

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.

To 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, choose **Save Table**.

Interactive SQL

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

To 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

- “Choosing object names” [[SQL Anywhere Server - SQL Usage](#)]
- “Data types in UltraLite” on page 312
- “Choosing column data types” [[SQL Anywhere Server - SQL Usage](#)]
- “UltraLite ALTER TABLE statement” on page 454

Deleting UltraLite tables

You can drop any table provided that 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.

To 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 choose **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.

To 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

- [“UltraLite DROP TABLE statement” on page 477](#)

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.

To browse UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. If system tables are hidden and you want to browse the data in one or more tables, right-click the white space of the **Contents** pane and choose **Show System Objects**.
3. To view a list of tables, double-click **Tables**.
4. To view table data, double-click a table and click the **Data** tab in the right pane.

To filter UltraLite system tables (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the database you are connected to and choose either **Hide System Objects** or **Show System Objects**.

Interactive SQL

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

To browse UltraLite user tables (Interactive SQL)

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

To 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 289](#)

Copying and pasting data 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.

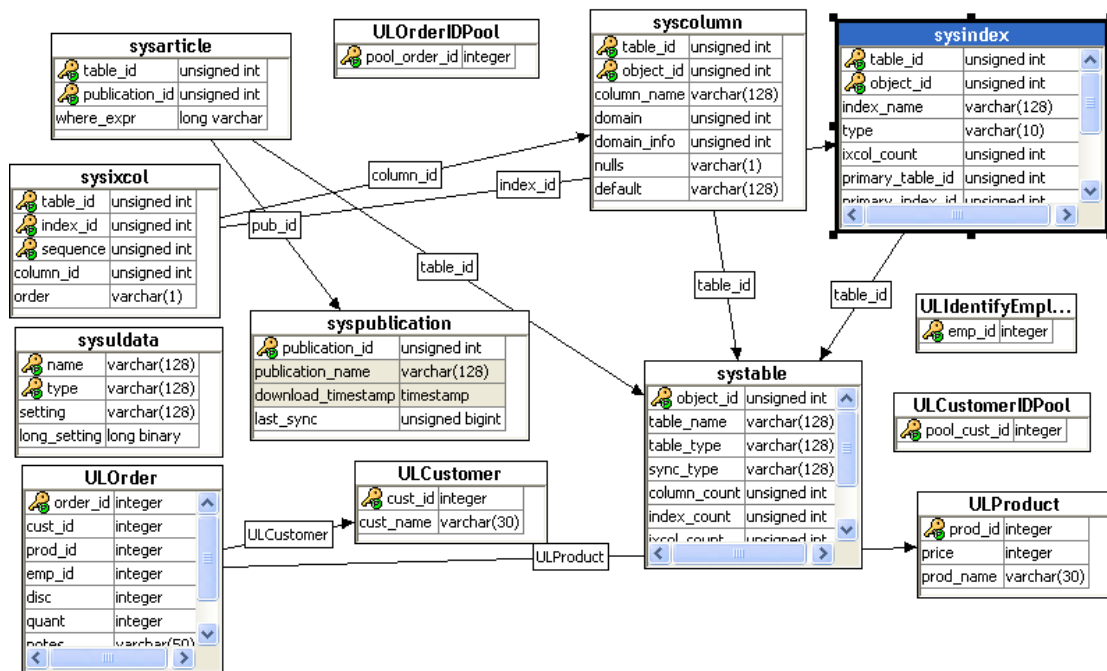
- Use the SELECT statement results into a named object.
- Use the INSERT statement to insert a row or selection of rows from elsewhere in the database into a table.

See also

- “Copying database objects in the SQL Anywhere plug-in” [[SQL Anywhere Server - Database Administration](#)]
- “UltraLite INSERT statement” on page 480
- “UltraLite SELECT statement” on page 486

Viewing 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

- “Creating databases in SQL Anywhere” [[SQL Anywhere Server - SQL Usage](#)]

Working with UltraLite indexes

An index provides an ordering (either ascending or descending) of a table's rows based on the values in one or more columns. 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.

Performance 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.

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

- [“Using index scans” on page 108](#)
- [“Execution plans in UltraLite” on page 342](#)
- [“Composite indexes” \[SQL Anywhere Server - SQL Usage\]](#)
- UltraLite.NET: [“Accessing and manipulating data with the Table API” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Prepare method” \[UltraLite - .NET Programming\]](#)
- UltraLite for C++: [“Accessing data with the table API” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“UltraLite_PreparedStatement class” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Working with data using the Table API” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“PreparedStatement class” \[UltraLite - M-Business Anywhere Programming\]](#)

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.

Choosing an index type

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

Notes

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 UltraLite indexes” on page 81](#)

Adding UltraLite indexes

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.

To create a new index for a given UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click **Indexes**, and choose **New » Index**.
3. Follow the instructions in the wizard.

Interactive SQL

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

To 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. See [“UltraLite CREATE INDEX statement” on page 460](#).

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

- [“UltraLite CREATE INDEX statement” on page 460](#)

Dropping an 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.

To 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 choose **Delete**.
4. Click **Yes**.

Interactive SQL

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

To 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

- [“UltraLite DROP INDEX statement” on page 474](#)

Working with 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 (except with HotSync on Palm OS).

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.
- The publication determines which columns are selected, but it does not determine the order in which they are sent. 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 138](#)
- [“Publishing data” \[MobiLink - Client Administration\]](#)
- [“Designing synchronization in UltraLite” on page 133](#)
- [“Introduction to synchronization scripts” \[MobiLink - Server Administration\]](#)

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.

To publish one or more whole UltraLite tables (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the **Publications** folder, and choose **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, select tables from the **Available Tables** list. Click **Add**.
5. Click **Finish**.

Interactive SQL

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

To 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

- [“UltraLite CREATE PUBLICATION statement” on page 462](#)
- [“UltraLite clients” on page 125](#)

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.

Palm OS

You cannot use a CREATE PUBLICATION statement with a WHERE clause on this platform.

Sybase Central

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

To publish only some rows in an UltraLite table (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the **Publications** folder, and choose **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, select a table and click **Add**.
6. Click the **WHERE Clauses** tab, and select 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.

To 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

- [“UltraLite CREATE PUBLICATION statement” on page 462](#)
- [“UltraLite clients” on page 125](#)

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.

To 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 choose **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.

To 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

- [“UltraLite DROP PUBLICATION statement” on page 475](#)
- [“UltraLite clients” on page 125](#)

Working with UltraLite users

User IDs and passwords are encrypted in the UltraLite database, so you can only view the list of defined users in Sybase Central.

UltraLite user IDs are not the same as MobiLink user names or SQL Anywhere user IDs.

Limitations

When creating unique user IDs, bear the following limitations in mind:

- UltraLite supports up to four unique users per database.
- Both the user ID and password values have a limit of 31 characters.
- Passwords are always case sensitive and user IDs are always case insensitive. You can change a password anytime from Sybase Central.
- Any leading or trailing spaces the user ID are ignored. The user ID cannot include leading single quotes('), leading double quotes ("), or semicolons(;).
- You cannot change a user ID once it is created. Instead, you must delete the user ID and then add a new one.
- Passwords can be changed using Sybase Central.

Adding a new UltraLite user

UltraLite does not support the creation of users with Interactive SQL. However, you can add users by:

- Using Sybase Central to add users to the User folder.
- Using the GrantConnectTo function on the Connection object to add new users from an UltraLite application.

To create a new UltraLite user (Sybase Central)

1. Connect to the UltraLite database.
2. Right-click the **Users** folder, and choose **New » User**.
3. Follow the instructions in the wizard. Ensure you understand how UltraLite interprets different user ID and password combinations. See [“Interpreting user ID and password combinations” on page 51](#).

See also

- UltraLite.NET: [“Authenticating users” \[UltraLite - .NET Programming\]](#)
- UltraLite C/C++ : [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Authenticating users” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for embedded SQL: [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)

Deleting an existing UltraLite user

UltraLite does not support the deletion of users using a SQL statement. However, you can delete users by using:

- Sybase Central to delete users from the User folder.
- The `RevokeConnectFrom` function on the Connection object to remove users from an UltraLite application.

To delete an existing UltraLite user (Sybase Central)

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

See also

- UltraLite.NET: “Authenticating users” [[UltraLite - .NET Programming](#)]
- UltraLite C/C++ : “Authenticating users” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Authenticating users” [[UltraLite - M-Business Anywhere Programming](#)]
- UltraLite for embedded SQL: “Authenticating users” [[UltraLite - C and C++ Programming](#)]

Working with 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	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. The <i>object_name</i> parameter specifies the table to monitor. A value of "*" means all tables in the database. The <i>table_name</i> notification parameter is the name of the modified table.

```
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 SQL passthrough

The SQL passthrough feature allows the publisher of the consolidated database to send SQL statement scripts to UltraLite for execution. SQL statement scripts can be executed automatically or manually by users with DBA authority.

SQL statement scripts are downloaded to UltraLite automatically during synchronization and stored in the `syssql` table. SQL statement scripts are not downloaded when a file-based download or a Ping is performed or when a download is restarted.

SQL statement scripts are run automatically the next time the database is started unless:

- The **flags** parameter is set to **manual** in the execution script.
- The connection parameter `dont_run_scripts` is set.
- The upload failed.

If the **flags** parameter is set to **manual** in the execution script you can use the following methods to manually apply scripts sent from the consolidated database:

- `ExecuteSQLPassthroughScripts`
- `ExecuteNextSQLPassthroughScript`
- `GetSQLPassthroughScriptCount`

During synchronization, UltraLite uploads the status of any scripts executed since the last synchronization. If UltraLite encounters a problem executing a script, no further scripts are executed until the consolidated database is notified and instructions for proceeding are sent. The status of executed scripts is stored in the “`ml_passthrough_status`” [*MobiLink - Server Administration*] table in the consolidated database. You can review the table to determine the success of distributed passthrough scripts.

You can use an observer callback function to monitor the progress of automatically or manually executed scripts. See “[Callback function for `ULRegisterSQLPassthroughCallback`](#)” [*UltraLite - C and C++ Programming*].

See also

- “[Introduction to SQL passthrough](#)” [*MobiLink - Client Administration*]
- “[ExecuteSQLPassthroughScripts method](#)” [*UltraLite - .NET Programming*]
- “[ExecuteNextSQLPassthroughScript method](#)” [*UltraLite - .NET Programming*]
- “[GetSQLPassthroughScriptCount method](#)” [*UltraLite - .NET Programming*]

UltraLite CustDB samples

Contents

CustDB sample file locations	95
Lesson 1: Build and run the CustDB application	97
Lesson 2: Log in and populate the UltraLite remote database	99
Lesson 3: Use the CustDB client application	100
Lesson 4: Synchronize with the CustDB consolidated database	102
Lesson 5: Browse MobiLink synchronization scripts	104
Build your own application	106

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. This data is pre-populated with sales status data.
- A **remote** UltraLite database. This database is initially empty.
- An UltraLite client application.
- A MobiLink server synchronization sample with 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 implements UltraLite features with some variation 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 file locations” on page 95](#)
- [“Scenario” \[*MobiLink - Getting Started*\]](#)
- [“Users in the CustDB sample” \[*MobiLink - Getting Started*\]](#)
- [“Tables in the CustDB databases” \[*MobiLink - Getting Started*\]](#)

CustDB sample file locations

SQL Anywhere CustDB database

This is the consolidated database. During installation, an ODBC data source called SQL Anywhere 11 CustDB is created for this database.

The CustDB installation depends on whether you want to use the existing sample or recreate a new file:

- For the existing sample: *samples-dir\UltraLite\CustDB\custdb.db*
- Erases changes that were synchronized into the consolidated *CustDB.db* file, so you have a clean version to work with: *samples-dir\UltraLite\CustDB\newdb.bat*

For more information about the schema of this file, see [“Exploring the CustDB sample for MobiLink” \[MobiLink - Getting Started\]](#).

For more information about the default location of *samples-dir*, by operating system, see [“Samples directory” \[SQL Anywhere Server - Database Administration\]](#).

The UltraLite CustDB database

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: *samples-dir\UltraLite.NET\CustDB\Common*
- For all other platforms and APIs: *samples-dir\UltraLite\CustDB\custdb.udb*

For more information about the default location of *samples-dir*, by operating system, see [“Samples directory” \[SQL Anywhere Server - Database Administration\]](#).

RDBMS-specific build scripts

The SQL scripts that rebuild a CustDB consolidated database for any one of the supported RDBMSs.

In the *samples-dir\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 more information about setting up a consolidated database, see [“Setting up the CustDB consolidated database” \[MobiLink - Getting Started\]](#).

For more information about the default location of *samples-dir*, by operating system, see [“Samples directory” \[SQL Anywhere Server - Database Administration\]](#).

UltraLite CustDB client applications and ReadMe files

These are the end-user tools that provide a user-friendly interface to the UltraLite remote database. There is a sample client installed for each supported API.

Each client application also contains a *ReadMe.html* or *ReadMe.txt* file. The contents of these files vary. However in some cases they outline the steps required to build and run that sample.

The location the applications and its ReadMe depends on your development environment. See [“Lesson 1: Build and run the CustDB application” on page 97](#).

SQL synchronization logic

These are the SQL statements needed to query and modify information from the UltraLite database and the calls required to start synchronization with the consolidated database.

samples-dir\UltraLite\CustDB\custdb.sqc

For more information about the default location of *samples-dir*, by operating system, see [“Samples directory” \[SQL Anywhere Server - Database Administration\]](#).

See also

- [“MobiLink consolidated databases” \[MobiLink - Server Administration\]](#)

Lesson 1: Build and run the CustDB application

The CustDB application is built for many development environments. For a general procedure that applies to all environments, see the following section.

To 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 *install-dir\UltraLite\win32\386* 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 *samples-dir\UltraLite\Custdb* directory.

For information about all versions of C/C++ CustDB applications, see *samples-dir\UltraLite\Custdb\readme.txt*.

- **CodeWarrior for Palm OS** You can find the project files in the *samples-dir\UltraLite\CustDB\cwcommon* and *samples-dir\UltraLite\CustDB\cw* directories.

For more information about building the C/C++ CustDB application, see [“Building the CustDB sample application in CodeWarrior” \[UltraLite - C and C++ Programming\]](#).

- **Visual Studio** You can find project files in the *samples-dir\UltraLite\CustDB\vs7* and *samples-dir\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.

UltraLite for embedded SQL

You can find project files specific to eMbedded Visual C++ in the *samples-dir\UltraLite\CustDB\EVC* and *samples-dir\UltraLite\CustDB\EVC40* directories depending on your version of embedded Visual C++.

For more information about building the embedded SQL CustDB application for Windows Mobile using eMbedded Visual C++, see [“Building the CustDB sample application” \[UltraLite - C and C++ Programming\]](#).

UltraLite.NET

You can find project files specific to Visual Studio.NET in the *samples-dir\UltraLite.NET\CustDB* directory.

For instructions on building the CustDB application for Windows Mobile using Visual Studio.NET, see *samples-dir\UltraLite.NET\CustDB\ce\ReadMe.html*.

To obtain deployment directory information for Microsoft Windows desktop, and information on where to download additional UltraLite.NET samples, see *samples-dir\UltraLite.NET\CustDB\Desktop\ReadMe.html*.

UltraLite for M-Business Anywhere

You can find project files specific to M-Business Anywhere in the *samples-dir\UltraLiteForMBusinessAnywhere\CustDB* directory.

For more information about building the CustDB application using M-Business Anywhere, see [“UltraLite for M-Business Anywhere quick start” \[UltraLite - M-Business Anywhere Programming\]](#). The instructions are applicable to Windows Mobile, Windows, and Palm OS.

Lesson 2: Log in and populate 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.

To start and synchronize the sample application

1. Choose **Start » Programs » SQL Anywhere 11 » MobiLink » Synchronization Server Sample**. Or, execute the following command:

```
mksrv11 -c "DSN=SQL Anywhere 11 CustDB" -zu+ -vcrs
```

A command prompt appears displaying messages about the MobiLink server's status.

2. Choose **Start » Programs » SQL Anywhere 11 » UltraLite » Windows Sample Application**.
3. In the **Employee ID** field, type **50**. Press Enter.

The application synchronizes and the MobiLink server messages window displays messages showing the synchronization taking place.

The default 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: Use 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 only displays a subset of columns 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.

To browse orders

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.

1. To scroll down the list of customers, click **Next**.
2. To scroll up through the list of customers, click **Previous**.

To add an order

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.

1. Choose **Order » New**.
2. In the **Customer** list, choose **Basements R Us**.
3. In the **Product** list, choose **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 press Enter.

To approve, deny, and delete orders

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.

1. 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, choose **Good**.
 - d. Press Enter.

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, choose **Discount Is Too High**.
 - d. Press Enter.

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 until you synchronize changes to the consolidated database.

See also

- [“Tables in the CustDB databases” \[MobiLink - Getting Started\]](#)

Lesson 4: Synchronize 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: Log in and populate the UltraLite remote database” on page 99](#).

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.

To synchronize the UltraLite remote

1. To synchronize your data, from the **File** menu choose **Synchronize**.
2. Confirm that synchronization took place.
 - At the remote database, you can confirm that all required transactions occurred by checking that the approved order for **Apple Street Builders** 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.

To confirm the synchronization (Sybase Central)

1. Choose **Start » Programs » SQL Anywhere 11 » Sybase Central**.
2. Choose **Connections » Connect With SQL Anywhere 11**.
3. Click **ODBC Data Source Name**.
4. Click **Browse** and choose **SQL Anywhere 11 CustDB**.
5. Click **OK**.
6. Click **OK**.
7. Double-click **Tables**.
8. Double-click **ULOrder (DBA)**.
9. Click the **Data** tab and verify that order 5100 is approved, order 5101 is denied, and order 5102 is deleted.

To confirm the synchronization (Interactive SQL)

1. Connect to the consolidated database from Interactive SQL.
 - a. From the **Start** menu, choose **Programs » SQL Anywhere 11 » Interactive SQL**.
 - b. Click **ODBC Data Source Name** and choose **SQL Anywhere 11 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: Browse 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” \[MobiLink - Getting Started\]](#).

To browse the synchronization scripts

1. From the **Start** menu, choose **Programs » SQL Anywhere 11 » Sybase Central**.
2. From the **Connections** menu, choose **Connect With MobiLink 11**.
3. Click **ODBC Data Source Name**.
4. Click **Browse** and choose **SQL Anywhere 11 CustDB**.
5. Click **OK**.
6. Click **OK**.
7. 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.

8. Click **Synchronized Tables**.
9. 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

- “Writing synchronization scripts” [*MobiLink - Server Administration*]
- “UltraLite clients” on page 125
- “Connection scripts” [*MobiLink - Server Administration*]
- “Table scripts” [*MobiLink - Server Administration*]

Build your own application

Use one of the supported interfaces to build your own application. For more information, see:

- UltraLite C++: [“Tutorial: Build an application using the C++ API”](#) [*UltraLite - C and C++ Programming*]
- UltraLite.NET: [“Tutorial: Build an UltraLite.NET application”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“Tutorial: A sample application for M-Business Anywhere”](#) [*UltraLite - M-Business Anywhere Programming*]

UltraLite performance and optimization

Contents

Using index scans	108
Determining the access method used by the optimizer	110
Tuning query performance with index hashing	111
Choosing an optimal hash size	113
Setting the maximum hash size	116
Managing temporary tables	117
Flushing single or grouped transactions	119
How database encryption and obfuscation affect performance	120
UltraLite optimization strategies	121

UltraLite provides excellent SQL query performance. Index scans, direct page scans, and temporary tables are internal optimization techniques that help you to achieve the most from the product. However, you can further tune these features, depending on the results of any query performance tests you run.

Using index scans

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. An index is a database object. The index is maintained automatically by UltraLite after it has been created. 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. So, when you use them strategically, indexes can greatly improve the performance of searches on the indexed column(s).

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.

Performance tips

- 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
- When creating a composite index, the first column of the index should be the one that is used most often by the predicate in your query.
- Ensure 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

- [“Managing temporary tables” on page 117](#)
- [“Using direct page scans” on page 117](#)
- [“View an UltraLite execution plan” on page 342](#)
- [“About composite indexes” on page 79](#)
- [“EXPLANATION function \[Miscellaneous\]” on page 383](#)
- UltraLite C++: [“GetPlan function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“getPlan method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“UltraLite page_size creation parameter” on page 199](#)

Determining the access method used by the optimizer

The UltraLite optimizer uses sophisticated optimization strategies when choosing an index for query optimization. Except in the case of 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. In some cases, 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. See [“Managing temporary tables” on page 117](#).

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)
```


Tuning query performance with 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:

- Partly by the maximum hash size property you configure. See [“Choosing an optimal hash size” on page 113](#).
- Partly by 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**.

Now, 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

Now, only two of the four row values would need to be unpacked and evaluated, rather than all four.

See also

- [“UltraLite max_hash_size creation parameter” on page 194](#)
- [“Choosing an optimal hash size” on page 113](#)
- [“UltraLite performance and optimization” on page 107](#)
- [“Adding UltraLite indexes” on page 81](#)

Choosing an optimal 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>

Data type	Bytes used to hash the entire value
BINARY	The maximum hash size in bytes must match the declared size of the column. For example, if you declare a column BINARY(30), the required size is 30 bytes.
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.

In cases where 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

See also

- [“UltraLite max_hash_size creation parameter” on page 194](#)
- [“UltraLite performance and optimization” on page 107](#)
- [“Adding UltraLite indexes” on page 81](#)
- [“Data types in UltraLite” on page 312](#)

Setting the maximum hash size

You can set the maximum hash size in two ways:

- To store a database *default* for the maximum size, you can 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:
 - 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

- [“UltraLite CREATE INDEX statement” on page 460](#)
- [“UltraLite CREATE TABLE statement” on page 468](#)

Managing 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.

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.

See also

- [“UltraLite temporary tables” on page 12](#)
- [“Determining the access method used by the optimizer” on page 110](#)
- [“Reading UltraLite execution plans” on page 343](#)

Using 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 pre-existing 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. On the other hand, 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 the Table 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 [“Determining the access method used by the optimizer” on page 110](#).

Reverting to primary key index order

In version 10.0.0 and earlier of UltraLite, the primary key was used to return results when no other index was used by the UltraLite optimizer. As a result, rows were ordered according to the order of the primary key index.

If your results must be ordered by primary key, you should re-write your queries to include the ORDER BY clause. If it is impractical to sort your rows with this clause, you can consider using the ORDERED_TABLE_SCAN connection parameter.

Tip

It is recommended that you use the ORDER BY clause whenever possible.

See also

- [“UltraLite ORDERED_TABLE_SCAN connection parameter \[deprecated\]” on page 239](#)
- [“Using index scans” on page 108](#)
- [“UltraLite SELECT statement” on page 486](#)

Flushing 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 tradeoff between the data integrity of your application and its performance.

See also

- [“UltraLite COMMIT_FLUSH connection parameter” on page 229](#)
- [“UltraLite commit_flush_count option \[temporary\]” on page 220](#)
- [“UltraLite commit_flush_timeout option \[temporary\]” on page 221](#)
- [“UltraLite CHECKPOINT statement” on page 458](#)
- [UltraLite for embedded SQL: “ULCheckpoint function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C++: “Checkpoint function” \[*UltraLite - C and C++ Programming*\]](#)

How database encryption and obfuscation affect performance

You can encrypt the database to increase the security of the information stored in UltraLite. However, you should note that 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. To determine what the optimal cache size for your scenario is, you can graph the database performance with benchmark tests.

Stressing the cache

You can benchmark 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 NT (the smallest allowed cache for this platform) or 64 KB on all other platforms.

If you find that increasing the cache does not improve the performance of an encrypted database, consider obfuscating the data rather than encrypting it. Obfuscation can yield better performance while still offering some security benefits; the obfuscation algorithm 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 must ultimately dictate whether you choose to use strong encryption or not.

See also

- [“UltraLite performance and optimization” on page 107](#)
- [“UltraLite page_size creation parameter” on page 199](#)
- [“UltraLite fips creation parameter” on page 192](#)
- [“UltraLite CACHE_SIZE connection parameter” on page 226](#)
- [“UltraLite CREATE INDEX statement” on page 460](#)

UltraLite optimization strategies

Different UltraLite deployments of UltraLite may require very specific settings to tune the database for your needs.

Early versions of Palm OS

Early versions (for example, version 4.x) of Palm devices only have approximately 200 KB of RAM. This limitation can cause problems due to the dynamic memory requirements of UltraLite 11.

Platform strategies There are two alternatives:

- Upgrade to Palm OS 5.x. This suggestion is the preferred option.
- Downgrade UltraLite to version 9.0.x.

Database optimization strategies You should ensure that you are doing the following:

- **For the `CACHE_SIZE` creation parameter** Do not establish a connection using this parameter, and allow UltraLite to set the default value. If you set your own value, you may inadvertently cause problems with the dynamic memory required for synchronization.
- **For the `page_size` creation database property** Create a new database with the smallest possible page size for your application. While 1 KB is the smallest page size, a 2 KB page size is the size typically used for Palm OS.

UltraLite as a MobiLink Client

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

UltraLite clients	125
Using ActiveSync and HotSync with UltraLite	145
UltraLite synchronization parameters and network protocol options	153

UltraLite clients

Contents

Built-in UltraLite synchronization features	126
Customizing UltraLite client synchronization behavior	127
Primary key uniqueness in UltraLite	129
Designing synchronization in UltraLite	133
Using MobiLink file transfers	142

Because the UltraLite runtime includes a built-in bi-directional synchronization framework that links field and mobile workers with enterprise back-end systems, synchronizing UltraLite data is less involved than other remote clients. This built-in framework means that all data in an UltraLite database is synchronized automatically 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.

For more information about UltraLite, see [“Introducing UltraLite” on page 1](#). For information about how to use SQL Anywhere databases as MobiLink clients, see [“SQL Anywhere clients” \[MobiLink - Client Administration\]](#).

Tip

The best way to back up an UltraLite application is to synchronize with a consolidated database. To restore an UltraLite database, take an empty database and populate it from the consolidated database through synchronization.

Tip

If you need to deploy multiple files, or if you need to target versions of the file to specific user IDs, consider using the MobiLink server to transfer files. See [“Using MobiLink file transfers” on page 142](#).

Built-in UltraLite synchronization features

UltraLite contains MobiLink synchronization technology in the data management layer for UltraLite. Unlike SQL Anywhere remote, 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 the progress counter.

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 a one-byte marker to keep track of the state of the row. In addition to synchronization, UltraLite also uses the row states to control transaction processing and data recovery. See [“UltraLite row states” on page 14](#).

The progress counter

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.

See also

- [“ml_subscription”](#) [*MobiLink - Server Administration*]

Customizing UltraLite client synchronization behavior

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.

See “[Maintaining unique primary keys](#)” [*MobiLink - Server Administration*] and “[Primary key uniqueness in UltraLite](#)” on page 129.

- **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. See “[Oracle consolidated database](#)” [*MobiLink - Server Administration*] and “[UltraLite precision creation parameter](#)” on page 201.
- **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. See “[Designing synchronization in UltraLite](#)” on page 133.

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 high usage costs associated with them, you would 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 initializing 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 assigning the MobiLink user. However, there are other parameters you can configure as well: for example, use the `upload_only` and `download_only` parameters to change the default bi-directional synchronization to one-way only. See “[Adding synchronization to your UltraLite application](#)” on page 139.

All other important synchronization behaviors are controlled with MobiLink synchronization scripts. You need multiple scripts because each MobiLink remote database can contain a different subset of the data in the consolidated database.

These include:

- What data is downloaded as updates 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

- [“MobiLink consolidated databases”](#) [*MobiLink - Server Administration*]
- [“Introduction to synchronization scripts”](#) [*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. See [“Maintaining unique primary keys” \[MobiLink - Server Administration\]](#).

One of these methods is to use a global ID. The global ID affects how the GLOBAL AUTOINCREMENT column values are generated. You need to use columns declared as GLOBAL AUTOINCREMENT when a MobiLink server must manage synchronizations for multiple clients.

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 ID.

See also

- [“UltraLite global_database_id option” on page 222](#)

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 ID for the UltraLite remote.

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.

To 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. See [“UltraLite global_database_id option” on page 222](#).

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 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 ID number, you allow UltraLite to set appropriate values from another partition.

One method of choosing a new global ID is to maintain a pool of unused global ID values. This pool is maintained in the same manner as a pool of primary keys. See [“Using primary key pools” \[SQL Remote\]](#).

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

- [“Overriding partition sizes for autoincremented columns” on page 131](#)
- [“UltraLite global_database_id option” on page 222](#)
- UltraLite.NET: [“Connection property” \[UltraLite - .NET Programming\]](#)
- UltraLite C/C++: [“Synchronize function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“setDatabaseID method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for embedded SQL: [“ULSetDatabaseID function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“GlobalAutoIncrementUsage property” \[UltraLite - .NET Programming\]](#)
- UltraLite C/C++: [“GetSynchResult function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“getGlobalAutoIncrementUsage method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for embedded SQL: [“ULGlobalAutoincUsage function” \[UltraLite - C and C++ Programming\]](#)

Determining 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. See [“GetLastIdentity function” \[UltraLite - C and C++ Programming\]](#).
- **UltraLite.NET** Use the `LastIdentity` property on the `ULConnection` class. See [“LastIdentity property” \[UltraLite - .NET Programming\]](#).
- **UltraLite for M-Business Anywhere** Use the `GetLastIdentity` method on the `Connection` class. See [“getLastIdentity method” \[UltraLite - M-Business Anywhere Programming\]](#).

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.

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.

To override UltraLite partition values (Sybase Central)

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

To declare autoincrement columns in UltraLite (Interactive 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. See [“UltraLite CREATE TABLE statement” on page 468](#) and [“UltraLite ALTER TABLE statement” on page 454](#).

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

- UltraLite.NET: [“GetColumnPartitionSize method” \[UltraLite - .NET Programming\]](#)
- UltraLite C/C++: [“GetGlobalAutoincPartitionSize function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“getColumnPartitionSize method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for embedded SQL: [“ULGlobalAutoincUsage function” \[UltraLite - C and C++ Programming\]](#)

Designing synchronization in UltraLite

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

Once you become comfortable with the 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 yourself 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 exclude tables from synchronization?	The nosync table name suffix allows you to identify any tables that you do not want to synchronize. See “Nosync tables in UltraLite” on page 136 .
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. See “Allsync tables in UltraLite” on page 136 .

Design question	If you answer yes, use the following
<p>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>	<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. See “Publications in UltraLite” on page 137.</p>
<p>Do you require a table order for synchronization because you have cycles of foreign keys?</p>	<p>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. See “Table order in UltraLite” on page 138.</p>

Design question	If you answer yes, use the following
<p>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>	<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. <p>See “UltraLite synchronization parameters and network protocol options” on page 153.</p>
<p>Do you want synchronization triggers to be time-based (scheduled), cradle-triggered, or user-initiated? Or do you require a combination of the above?</p>	<p>Different behavior can be achieved programmatically via an appropriate interface. In some cases, HotSync or ActiveSync may manage the synchronization process. See “Adding synchronization to your UltraLite application” on page 139.</p>

Design question	If you answer yes, use the following
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 “Deploy UltraLite with TLS-enabled synchronization” on page 60.

See also

- [“The synchronization process”](#) [*MobiLink - Getting Started*]
- [“The upload and the download”](#) [*MobiLink - Getting Started*]

Nosync tables in UltraLite

By adding the `_nosync` suffix to the table name, 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.

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. See [“Publications in UltraLite”](#) on page 137.

Allsync tables in UltraLite

By adding the `_allsync` suffix to the table name, 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.

Some UltraLite applications require user/client-specific data that you can store in allsync tables. Therefore, if you upload the data in the table to a temporary table in the consolidated database, you can 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.

Publications in UltraLite

Publications define a set of articles that describe the data to be synchronized. In general, 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 the `_nosync` table suffix approach and allow you to have more granular control. 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.

Adding publications

You can add publications to the UltraLite database with Sybase Central, or from Interactive SQL. For UltraLite synchronization, each article in a publication may include either a complete table, or may include a WHERE clause (except with HotSync on Palm OS).

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.

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.

To 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 choose **Add Publication**.
3. Enter a name for the new publication. Click **Next**.
4. On the **Tables** tab, select a table from 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. If you are using HotSync synchronization, do not specify a WHERE clause.
7. Click **Finish**.

To publish data from an UltraLite database (Interactive 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

- [“Download Only synchronization parameter” on page 159](#)
- [“Upload Only synchronization parameter” on page 174](#)
- [“Additional Parameters synchronization parameter” on page 155](#)
- [“Publishing data” \[MobiLink - Client Administration\]](#)
- [“UltraLite CREATE PUBLICATION statement” on page 462](#)
- [“Working with UltraLite publications” on page 83](#)
- [“Introduction to synchronization scripts” \[MobiLink - Server Administration\]](#)

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 Table Order parameter programmatically or as part of the `ulsync` utility during testing. The Table Order parameter specifies the order of tables that are to be uploaded. See [“Additional Parameters synchronization parameter” on page 155](#).

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 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 check the referential integrity before committing operations. 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

- [“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, Palm OS devices also support HotSync synchronization, and Windows Mobile devices also support ActiveSync synchronization.

Defining the connection

Each UltraLite remote 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 with the `user_name` and `version` parameters.

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 `disable_concurrency` synchronization parameter.

To 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_synch_info` structure:

```
ul_synch_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 TCP/IP-based 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 HotSync synchronization, you must use the `ULSetSynchInfo` function, supplying the `ul_synch_info` structure as an argument. This attaches the `ul_synch_info` structure to the current database for use on a subsequent synchronization.

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 [“Handling failed downloads” \[MobiLink - Server Administration\]](#) and [“Resuming failed downloads” \[MobiLink - Server Administration\]](#).

See also

- [“Using ActiveSync and HotSync with UltraLite” on page 145](#)
- [“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 153](#)
- [UltraLite.NET: “Synchronization in UltraLite applications” \[UltraLite - .NET Programming\]](#)
- [UltraLite C/C++: “Synchronizing data” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite C/C++: “Adding HotSync synchronization to Palm applications” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite C/C++: “Adding ActiveSync synchronization to your application” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Synchronizing data” \[UltraLite - M-Business Anywhere Programming\]](#)
- [UltraLite for embedded SQL: “Adding synchronization to your application” \[UltraLite - C and C++ Programming\]](#)

Using MobiLink file transfers

All UltraLite libraries support the ability to transfer files with the MobiLink server, except M-Business Anywhere. 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 Palm OS, if your destination is an external storage media, you must prefix the destination of the local path with **vfs:**.

If the destination is NULL, mlfiletransfer assumes that the file you need to download is a Palm record database (a *.pdb file) to the device's record store.

The file name must follow file name conventions for Palm OS. See [“Palm OS” on page 47](#).

- 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 47](#).

- 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 [“Windows desktop” on page 47](#).

2. Set the 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. See [“UltraLite synchronization parameters and network protocol options” on page 153](#).

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.

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 download 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.

See also

- [“-ftr option” \[MobiLink - Server Administration\]](#)
- [“MobiLink file transfer utility \(mlfiletransfer\)” \[MobiLink - Client Administration\]](#)
- UltraLite for C/C++: [“MLFileTransfer function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“ULFileTransfer class” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business: Not supported

Using ActiveSync and HotSync with UltraLite

Contents

HotSync on Palm OS 146
ActiveSync on Windows Mobile 150

HotSync on Palm OS

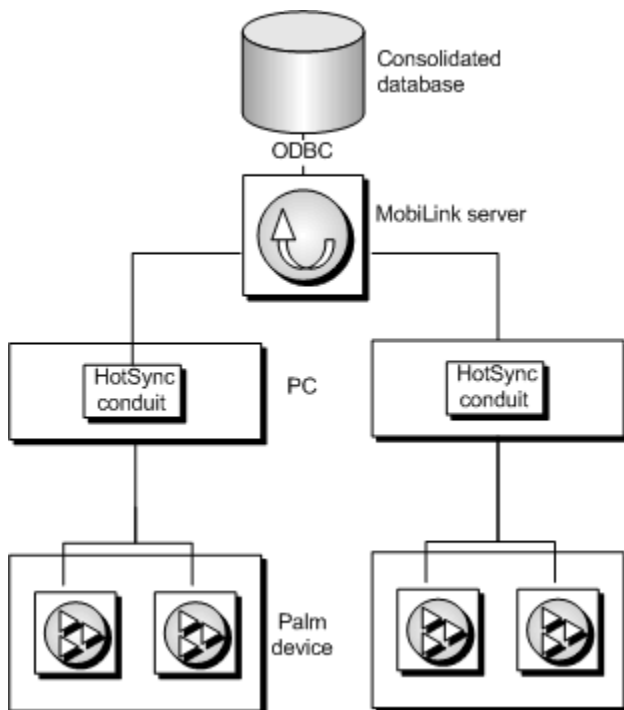
While you can synchronize data from a Palm OS device over an ethernet, Wi-Fi, or RAS connection, this section describes how to configure your desktop and device for HotSync synchronization.

HotSync synchronization allows you to manage synchronization for all UltraLite databases on the Palm device simultaneously from the user's desktop. HotSync synchronization is initiated externally by HotSync when the device is connected to the desktop. This requires that you program your application for HotSync synchronization. See [“Adding HotSync synchronization to Palm applications” \[UltraLite - C and C++ Programming\]](#).

If you were to synchronize UltraLite from an application directly (that is, without Hotsync), your end-users need to open each application separately and synchronize each database in turn. You implement this programmatically by initializing synchronization with an API specific synchronization function. See [“Adding TCP/IP, HTTP, or HTTPS synchronization to Palm applications” \[UltraLite - C and C++ Programming\]](#).

The HotSync architecture

The HotSync architecture requires a consolidated database to which data to and from remote databases are synchronized. The MobiLink server manages synchronization events among these databases. The UltraLite HotSync conduit manages synchronization events locally from the end user's desktop.



You must use the HotSync Conduit Installation utility (ulcond11) to install a conduit and register each UltraLite database. By registering each database with the conduit, you are configuring the UltraLite HotSync conduit to:

- Manage HotSync synchronization for all databases associated with an application (and therefore its registered creator ID). The conduit must be installed for each applicable creator ID.
- Connect to each database with the appropriate connection string. If you need to register multiple databases, use the -a option with the ulcond11 utility. This appends each additional connection string to the one you define with the -c option.

This allows you to synchronize all databases simultaneously by initializing synchronization with HotSync from the desktop.

See also

- [“UltraLite HotSync Conduit Installation utility for Palm OS \(ulcond11\)” on page 263](#)
- [UltraLite C/C++ : “Developing UltraLite applications for the Palm OS” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “Adding synchronization to your application” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for M-Business Anywhere: “Synchronizing data” \[UltraLite - M-Business Anywhere Programming\]](#)

HotSync synchronization overview

1. Ensure that you have specified the required synchronization parameters in the application with the ul_synch_info structure. These synchronization parameters are then set for HotSync synchronization by calling the SetSynchInfo function. The synchronization parameters also include the network protocol options you need to communicate with the MobiLink server. See [“Setting protocol options for MobiLink synchronization” on page 148](#).
2. When your UltraLite application is closed, the state of your UltraLite application is stored in a temporary file, separately from the database. See [“Maintaining state in UltraLite Palm applications \(deprecated\)” \[UltraLite - C and C++ Programming\]](#).
3. When you synchronize your Palm device, HotSync calls the UltraLite HotSync conduit on the desktop for a specific creator ID.
4. The UltraLite HotSync conduit connects to all registered databases using the appropriate connection string and then synchronizes them. Databases that do not call SetSynchInfo properly will fail to synchronize.
5. When your application is launched, it loads the previously saved state of your UltraLite application. See [“Restoring state in UltraLite Palm applications \(deprecated\)” \[UltraLite - C and C++ Programming\]](#).

See also

- [“ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- [“SetSynchInfo function” \[UltraLite - C and C++ Programming\]](#)

Setting protocol options for MobiLink synchronization

Protocol options describe the connection from the UltraLite HotSync conduit to the MobiLink server. Typically you would add this information as part of the application's synchronization code. However, you can also enter the required parameters from HotSync, or even `ulcond11`.

If you are using the `ul_synch_info` structure, the argument has the following form:

```
stream=name;parameters
```

stream=name indicates the type of network protocol. The default value for the *stream=name* is TCPIP.

parameters is a set of network protocol options for use by the conduit, and has the same form as the **stream_parms** argument for the protocol in use. For the given stream, the *parameters* adopts the same defaults as the **stream_parms** argument for the protocol.

Note

If you do not supply protocol options, the conduit searches in the registry for the protocol name and protocol options that you may have also supplied when you ran `ulcond11`.

If HotSync finds no valid network protocol, the default protocol and protocol options are used. This default stream parameter setting is:

```
stream=tcPIP;host=localhost
```

See also

- [“Network protocol options for UltraLite synchronization streams” on page 178](#)
- UltraLite.NET: [“ULStreamType enumeration” \[UltraLite - .NET Programming\]](#)
- UltraLite for embedded SQL: [“ULSetSynchInfo function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“SetSynchInfo function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business: [“setStream method” \[UltraLite - M-Business Anywhere Programming\]](#)

Debugging HotSync operations

A HotSync log file is maintained in the `Palm-install\User-dir` directory. By default, this file contains information on:

- The time of the synchronization event.
- Status of each registered conduit.

You can obtain additional debugging information in this file by setting the `UL_DEBUG_CONDUIT_LOG` environment variable. There are two levels you can choose from, depending on the amount of information you need to capture:

UL_DEBUG_CONDUIT_LOG = 1 Record basic information. For example, synchronization parameters, and registry locations.

UL_DEBUG_CONDUIT_LOG = 2 Record additional UltraLite details, including input/output operations.

Restart HotSync so the new setting takes effect.

See also

- [“Introduction to SQL Anywhere environment variables” \[SQL Anywhere Server - Database Administration\]](#)

ActiveSync 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. 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.

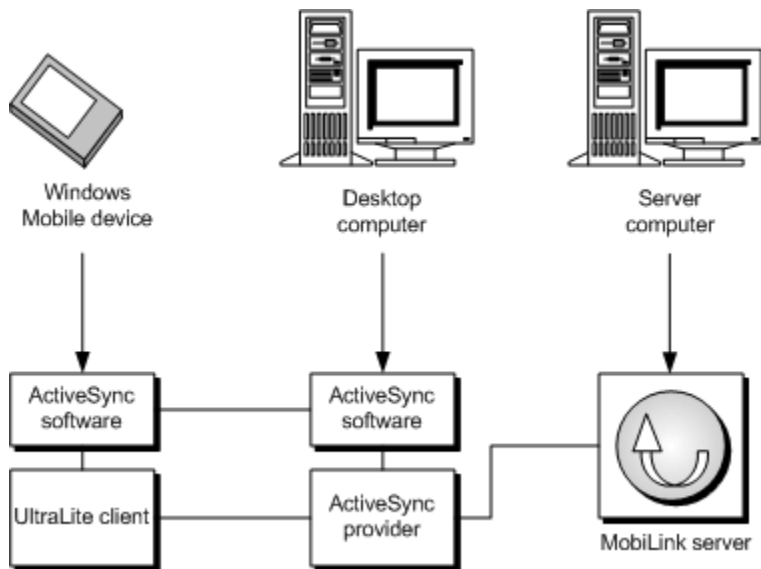
To use ActiveSync initiated synchronization requires that you:

- 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++ : “Adding 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: “Synchronizing data” [[UltraLite - M-Business Anywhere Programming](#)]

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

Contents

Synchronization parameters for UltraLite 154
Network protocol options for UltraLite synchronization streams 178

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 chapter describes the effects of the parameters, and provides links to other locations for information about how to set them.

Note

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

For API-specific details, see:

- UltraLite for Embedded SQL: “[ULSynchronize function](#)” [[UltraLite - C and C++ Programming](#)]
- UltraLite.NET: “[ULSyncParms class](#)” [[UltraLite - .NET Programming](#)]
- UltraLite C/C++: “[Synchronize function](#)” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “[SyncParms class](#)” [[UltraLite - M-Business Anywhere Programming](#)]

Required parameters

The following parameters are required:

- **Stream Type** See “[Stream Type synchronization parameter](#)” on page 171.
- **User Name** See “[User Name synchronization parameter](#)” on page 175.
- **Version** See “[Version synchronization parameter](#)” on page 176.

If you do not set these parameters, the synchronization function throws an exception (for example, `SQLCode.SQLLE_SYNC_INFO_INVALID` or its equivalent).

Conflicting parameters

You can only specify one of these parameters:

- **Download Only** See “[Download Only synchronization parameter](#)” on page 159.
- **Ping** See “[Ping synchronization parameter](#)” on page 165.
- **Upload Only** See “[Upload Only synchronization parameter](#)” on page 174.

If you set more than one of these parameters to true, the synchronization function throws an exception (for example, `SQLCode.SQLLE_SYNC_INFO_INVALID` or its equivalent).

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.

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.
DisableConcurrency	Disallows database access from other threads during synchronization during the upload phase. Set this property to 0 to allow concurrent database access; otherwise, set to 1 . By default, this property is set to 0 .

Property name	Description
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 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

- UltraLite for C/C++: “ul_synch_info_a struct” [[UltraLite - C and C++ Programming](#)]
- UltraLite.NET: “AdditionalParams property” [[UltraLite - .NET Programming](#)]
- UltraLite for M-Business Anywhere: “setAdditionalParams method” [[UltraLite - M-Business Anywhere Programming](#)]

Example

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

```

ul_synch_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. See [“Number of Authentication Parameters parameter” on page 162](#).

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 162](#)
- [“Authentication parameters” \[MobiLink - Server Administration\]](#)
- [“authenticate_parameters connection event” \[MobiLink - Server Administration\]](#)
- UltraLite for C/C++: [“ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“AuthenticationParms property” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“setAuthenticationParms method” \[UltraLite - M-Business Anywhere Programming\]](#)

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. See [“authenticate_user connection event” \[MobiLink - Server Administration\]](#).

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

- [“MobiLink users” \[MobiLink - Client Administration\]](#)
- [UltraLite for C/C++: “ul_synch_info_a_struct” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite.NET: “AuthStatus property” \[UltraLite - .NET Programming\]](#)
- [UltraLite for M-Business Anywhere: “SyncResult class” \[UltraLite - M-Business Anywhere Programming\]](#)

Example

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

```
ul_synch_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.

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 157
- UltraLite for C/C++: “`ul_synch_info_a` struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “AuthValue property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncResult class” [*UltraLite - M-Business Anywhere Programming*]

Example

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

```
ul_synch_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 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 174](#)
- [“The progress counter” on page 126](#)
- [“UltraLite Synchronization utility \(`ulsync`\)” on page 275](#)
- UltraLite for C/C++: [“`ul_synch_info_a` struct” \[*UltraLite - C and C++ Programming*\]](#)
- UltraLite.NET: [“`DownloadOnly` property” \[*UltraLite - .NET Programming*\]](#)
- UltraLite for M-Business Anywhere: [“`SyncParms` class” \[*UltraLite - M-Business Anywhere Programming*\]](#)

Examples

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

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

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

```
ul_synch_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

- UltraLite for C/C++: “[ul_synch_info_a struct](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[IgnoredRows property](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[SyncResult class](#)” [*UltraLite - M-Business Anywhere Programming*]

Example

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

```
ul_synch_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.

Default

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

Allowed values

Boolean

See also

- “[Resuming failed downloads](#)” [*MobiLink - Server Administration*]
- “[Resume Partial Download synchronization parameter](#)” on page 167
- UltraLite for C/C++: “[ul_synch_info_a struct](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[KeepPartialDownload property](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[SyncParms class](#)” [*UltraLite - M-Business Anywhere Programming*]

Example

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

```
ul_synch_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 275
- UltraLite for C/C++: [“ul_synch_info_a struct”](#) [*UltraLite - C and C++ Programming*]
- UltraLite.NET: [“NewPassword property”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“SyncParms class”](#) [*UltraLite - M-Business Anywhere Programming*]

Example

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

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

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

```
ul_synch_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.

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 156](#)
- [“authenticate_parameters connection event” \[MobiLink - Server Administration\]](#)
- [“Authentication parameters” \[MobiLink - Server Administration\]](#)
- UltraLite for C/C++: [“ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“AuthenticationParms property” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“SyncParms class” \[UltraLite - M-Business Anywhere Programming\]](#)

Example

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

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

Observer synchronization parameter

Specifies a pointer to a callback function or event handler that monitors synchronization.

Syntax

The syntax varies depending on the API you use.

See also

- [“User Data synchronization parameter” on page 175](#)
- UltraLite for C/C++: [“ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“ULSyncProgressListener members” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“synchronizeWithParm method” \[UltraLite - M-Business Anywhere Programming\]](#)

Example

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

```
ul_synch_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.

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 [“Keep Partial Download synchronization parameter”](#) on page 161.

See also

- [“Resuming failed downloads”](#) [*MobiLink - Server Administration*]
- [“Resume Partial Download synchronization parameter”](#) on page 167
- UltraLite for C/C++: [“ul_synch_info_a struct”](#) [*UltraLite - C and C++ Programming*]
- UltraLite.NET: [“PartialDownloadRetained property”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“SyncResult class”](#) [*UltraLite - M-Business Anywhere Programming*]

Example

Access the parameter as follows:

```
ul_synch_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. See [“User Name synchronization parameter”](#) on page 175.

If the MobiLink client already has a password, use the New Password parameter to change it. See [“New Password synchronization parameter”](#) on page 162.

See also

- “MobiLink users” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 275
- UltraLite for C/C++: “ul_synch_info_a struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “Password property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Examples

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

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

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

```
ul_synch_info info;
// ...
info.password = UL_TEXT( "myspassword" );
```

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 option” [*MobiLink - Client Administration*]
- “UltraLite Synchronization utility (ulsync)” on page 275
- UltraLite for C/C++: “ul_synch_info_a_struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “PingOnly property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Examples

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

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

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

```
ul_synch_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

- [“Publications in UltraLite” on page 137](#)
- [“Working with UltraLite publications” on page 83](#)
- [“UltraLite Synchronization utility \(ulsync\)” on page 275](#)
- UltraLite for C/C++: “ul_synch_info_a struct” [[UltraLite - C and C++ Programming](#)]
- UltraLite.NET: “ULPublicationSchema class” [[UltraLite - .NET Programming](#)]
- UltraLite for M-Business Anywhere: “PublicationSchema class” [[UltraLite - M-Business Anywhere Programming](#)]

Example

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

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

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

```
ul_synch_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.

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

- “Resuming failed downloads” [*MobiLink - Server Administration*]
- “Keep Partial Download synchronization parameter” on page 161
- “Partial Download Retained synchronization parameter” on page 163
- “UltraLite Synchronization utility (ulsync)” on page 275
- UltraLite for C/C++: “ul_synch_info_a struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “ResumePartialDownload property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Example

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

```
ul_synch_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. See “[Direct row handling](#)” [*MobiLink - Server Administration*].

See also

- “UltraLite Synchronization utility (ulsync)” on page 275
- UltraLite for C/C++: “ul_synch_info_a struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “SendColumnNames property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Examples

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

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

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

```
ul_synch_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 275
- UltraLite for C/C++: “`ul_synch_info` a struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “SendDownloadAck property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Examples

`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_synch_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** Not required. The value is always 0.
For a listing of error numbers, see “[MobiLink communication error messages](#)” [*Error Messages*]. For the error code suffixes, see `install-dir\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 [the Microsoft Developer Network web site](#).
- **error_string** An application-provided error message. The string may or may not be empty. A non-empty `error_string` provides information in addition to the `stream_error_code`. For instance, for a write error (error code 9) the error string is a number showing how many bytes it was trying to write.
- **error_string_length** Deprecated. The size of the error string buffer.

Remarks

UltraLite applications other than the UltraLite C++ Component receive communications error information as part of the Sync Result parameter. See “[Sync Result synchronization parameter](#)” on page 173.

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 `install-dir\SDK\Include\sserror.h`.

Check for `SQLE_COMMUNICATIONS_ERROR`:

```

Connection conn;
ul_synch_info info;
...
conn.InitSynchInfo( &info );
info.stream_error.error_string = error_buff;
info.stream_error.error_string_length =
    sizeof( error_buff );
if( !conn.Synchronize( &synch_info ) ){
    if( SQLCODE == SQLE_COMMUNICATIONS_ERROR ){
        printf( error_buff );
        // more error handline here
    }
}

```

See also

- “ul_synch_info_a struct” [*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. See “[Stream Parameters synchronization parameter](#)” on page 172.

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.
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

- “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 (ulsynch)” on page 275
- “Network protocol options for UltraLite synchronization streams” on page 178
- UltraLite for C/C++: “ul_synch_info_a struct” [[UltraLite - C and C++ Programming](#)]
- UltraLite.NET: “Stream property” [[UltraLite - .NET Programming](#)]
- UltraLite for M-Business Anywhere: “SyncParms class” [[UltraLite - M-Business Anywhere Programming](#)]

Example

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

```
Connection conn;
ul_synch_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 `ulsynch`.

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 (ulsynch)” on page 275
- “Network protocol options for UltraLite synchronization streams” on page 178
- UltraLite for C/C++: “ul_synch_info_a struct” [[UltraLite - C and C++ Programming](#)]
- UltraLite.NET: “StreamParms property” [[UltraLite - .NET Programming](#)]
- UltraLite for M-Business Anywhere: “SyncParms class” [[UltraLite - M-Business Anywhere Programming](#)]

Example

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

```
ul_synch_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_synch_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. See [“Authentication Status synchronization parameter” on page 157](#).
- **Ignored Rows** Reports the number of ignored rows. See [“Ignored Rows synchronization parameter” on page 160](#).
- **Stream Error information** The Stream Error information includes a Stream Error Code, Stream Error Context, Stream Error ID, and Stream Error System. See [“Stream Error synchronization parameter” on page 169](#).
- **Upload OK** Reports the success or failure of the upload phase. See [“Upload OK synchronization parameter” on page 173](#).

See also

- UltraLite.NET: [“ULSyncParms class” \[UltraLite - .NET Programming\]](#)
- UltraLite C/C++: [“ul_synch_result struct” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“SyncParms class” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for embedded SQL: [“ULGetSynchResult function” \[UltraLite - C and C++ Programming\]](#)

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

- UltraLite C/C++: “[ul_synch_info_a struct](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[UploadOK property](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[SyncResult class](#)” [*UltraLite - M-Business Anywhere Programming*]

Example

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

```
ul_synch_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

- [“Designing synchronization in UltraLite” on page 133](#)
- [“Download Only synchronization parameter” on page 159](#)
- [“UltraLite Synchronization utility \(ulsync\)” on page 275](#)
- UltraLite C/C++: “`ul_synch_info_a struct`” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “`UploadOnly property`” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “`SyncParms class`” [*UltraLite - M-Business Anywhere Programming*]

Examples

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_synch_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 163](#)
- UltraLite C/C++: “`ul_synch_info_a struct`” [*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. See [“Remote IDs” \[MobiLink - Client Administration\]](#).

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. See [“Password synchronization parameter” on page 164](#).

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

- [“MobiLink users” \[MobiLink - Client Administration\]](#)
- [“UltraLite user authentication” \[MobiLink - Client Administration\]](#)
- [“UltraLite Synchronization utility \(ulsync\)” on page 275](#)
- [UltraLite C/C++: “ul_synch_info_a struct” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite.NET: “UserName property” \[UltraLite - .NET Programming\]](#)
- [UltraLite for M-Business Anywhere: “SyncParms class” \[UltraLite - M-Business Anywhere Programming\]](#)

Examples

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_synch_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 275
- UltraLite C/C++: “`ul_synch_info_a` struct” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “Version property” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “SyncParms class” [*UltraLite - M-Business Anywhere Programming*]

Examples

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_synch_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 HotSync synchronization on the Palm OS and 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.

For a list of protocol options, see [“MobiLink client network protocol option summary”](#) [*MobiLink - Client Administration*].

See also

- [“Configuring UltraLite clients to use transport-layer security”](#) [*SQL Anywhere Server - Database Administration*]
- [“Deploy UltraLite with TLS-enabled synchronization”](#) on page 60
- [“MobiLink client network protocol options”](#) [*MobiLink - Client Administration*]
- [“Stream Parameters synchronization parameter”](#) on page 172
- [-x option in “UltraLite Synchronization utility \(ulsync\)”](#) on page 275

Setting 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 [“Stream Parameters synchronization parameter”](#) on page 172.

However, if you are using HotSync, and if you did not specify a Stream Parameter value, or even if you specified the value as NULL, you can enter the required parameters from the HotSync Manager. See [“Setting protocol options for MobiLink synchronization”](#) on page 148.

For information about including stream parameters in your UltraLite synchronization call, see:

- UltraLite for C/C++: [“Adding HotSync synchronization to Palm applications”](#) [*UltraLite - C and C++ Programming*]
- UltraLite for UltraLite.NET: [“StreamParms property”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“setStreamParms method”](#) [*UltraLite - M-Business Anywhere Programming*]

zlib compression note

zlib compression is not supported on Palm OS.

UltraLite Database Reference

This section provides a reference for UltraLite database properties, options, connection parameters, and utilities.

UltraLite creation parameters	181
UltraLite database properties	213
UltraLite database options	219
UltraLite connection parameters	225
UltraLite utilities	247
UltraLite system tables	289

UltraLite creation parameters

Contents

UltraLite case creation parameter	183
UltraLite checksum_level creation parameter	184
UltraLite collation creation parameter	186
UltraLite date_format creation parameter	187
UltraLite date_order creation parameter	190
UltraLite fips creation parameter	192
UltraLite max_hash_size creation parameter	194
UltraLite nearest_century creation parameter	196
UltraLite obfuscate creation parameter	198
UltraLite page_size creation parameter	199
UltraLite precision creation parameter	201
UltraLite scale creation parameter	203
UltraLite time_format creation parameter	205
UltraLite timestamp_format creation parameter	207
UltraLite timestamp_increment creation parameter	210
UltraLite utf8_encoding creation parameter	211

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

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

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. For example:

```
ulcreate -o "case=respect;utf_encoding=1" -y test.udb
```

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

See also

- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite case creation parameter

Sets the case sensitivity of string comparisons in the UltraLite database.

Syntax

```
{ ulcreate | ulinit | ulload } -o case=value;
```

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. See [“Strings in UltraLite” on page 304](#).

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, select 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

- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite checksum_level creation parameter

Sets the level of checksum validation for the database.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 213
- “UltraLite Create Database utility (ulcreate)” on page 255
- “UltraLite Initialize Database utility (ulinit)” on page 269
- “UltraLite Load XML to Database utility (ulload)” on page 272
- UltraLite for embedded SQL: “ULCreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite for C++: “CreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “CreateDatabase method” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “createDatabase method” [*UltraLite - M-Business Anywhere Programming*]
- “UltraLite optimization strategies” on page 121
- “UltraLite performance and optimization” on page 107
- “UltraLite page_size creation parameter” on page 199
- “Connecting to an UltraLite database” on page 43

UltraLite collation creation parameter

Sets the database collation.

Syntax

```
{ ulcreate | ulinit | ulload } -o collation=value;
```

Allowed values

String

Default

1252Latin1

Remarks

For a list of supported collations in UltraLite, see [“UltraLite supported collations” on page 37](#).

You can also view a list of supported collations in UltraLite by executing the following command:

```
ulcreate -l
```

From Sybase Central, you can set the collation in any wizard that creates a database. On the **New Database Collation And Character Set** page, choose 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 Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“UltraLite character sets” on page 36](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite date_format creation parameter

Sets the format for dates retrieved from the database.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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[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.

Symbol	Description
<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 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, select 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 190
- “UltraLite Create Database utility (ulcreate)” on page 255
- “UltraLite Initialize Database utility (ulinit)” on page 269
- “UltraLite Load XML to Database utility (ulload)” on page 272
- UltraLite for embedded SQL: “ULCreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite for C++: “CreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “CreateDatabase method” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “createDatabase method” [*UltraLite - M-Business Anywhere Programming*]
- “Accessing creation parameter values” on page 34

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
<i>yyyy/mm/ddddd</i>	2001/05/21/thu
<i>jjj</i>	141
<i>mmm yyyy</i>	may 2001
<i>mm-yyyy</i>	05-2001

UltraLite date_order creation parameter

Controls the interpretation of date formats.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 187](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

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.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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 strong encryption in any wizard that creates a database. On the **New Database Storage Settings** page, select the **AES FIPS Algorithm** option. You must also set and confirm the encryption key.

On Palm OS, use the `ULEnableFipsStrongEncryption` method.

On all other platforms, set this parameter as one of the creation parameters for the create database method on the database manager class.

To deploy a FIPS-enabled database, copy all appropriate libraries for your platform. See [“Deploy UltraLite with AES_FIPS database encryption” on page 59](#).

See also

- [“Strong encryption” \[SQL Anywhere Server - Database Administration\]](#)
- [“Securing UltraLite databases” on page 40](#)
- [“UltraLite obfuscate creation parameter” on page 198](#)
- [“UltraLite DBKEY connection parameter” on page 233](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“Saving, retrieving, and clearing encryption keys on Palm OS” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite C/C++: “ULChangeEncryptionKey function” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for C++: “CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite.NET: “CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- [UltraLite.NET: “ChangeEncryptionKey method” \[UltraLite - .NET Programming\]](#)
- [UltraLite for M-Business Anywhere: “createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [UltraLite for M-Business Anywhere: “changeEncryptionKey method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite max_hash_size creation parameter

Sets the maximum default index hash size in bytes.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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 tradeoff 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, select 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 and optimization” on page 107](#)
- [“Working with UltraLite indexes” on page 79](#)
- [“Choosing an optimal hash size” on page 113](#)
- [“UltraLite CREATE INDEX statement” on page 460](#)
- [“UltraLite CREATE TABLE statement” on page 468](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [UltraLite for embedded SQL: “ULCreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C++: “CreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite.NET: “CreateDatabase method” \[*UltraLite - .NET Programming*\]](#)
- [UltraLite for M-Business Anywhere: “createDatabase method” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite nearest_century creation parameter

Controls the interpretation of two-digit years in string-to-date conversions.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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. In the case of 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. See [“Using unambiguous dates and times” \[SQL Anywhere Server - SQL Reference\]](#).

- **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, select 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

- “Ambiguous string to date conversions” [*SQL Anywhere Server - SQL Reference*]
- “UltraLite Create Database utility (ulcreate)” on page 255
- “UltraLite Initialize Database utility (ulinit)” on page 269
- “UltraLite Load XML to Database utility (ulload)” on page 272
- UltraLite for embedded SQL: “ULCreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite for C++: “CreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “CreateDatabase method” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “createDatabase method” [*UltraLite - M-Business Anywhere Programming*]
- “Accessing creation parameter values” on page 34

UltraLite obfuscate creation parameter

Controls obfuscation of data in the database. Obfuscation is a form of simple encryption.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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. See [“UltraLite fips creation parameter” on page 192](#).

From Sybase Central, you can set obfuscation in any wizard that creates a database. On the **New Database Storage Settings** page, select 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

- [“Simple encryption” \[SQL Anywhere Server - Database Administration\]](#)
- [“Securing UltraLite databases” on page 40](#)
- [“UltraLite fips creation parameter” on page 192](#)
- [“UltraLite DBKEY connection parameter” on page 233](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++ : [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite page_size creation parameter

Defines the database page size.

Syntax

```
{ ulcreate | ulinit | ulload } -o page_size=size[ k ];...
```

Allowed values

1k, 2k, 4k, 8k, 16k

Default

4k

Remarks

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. In some cases, 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. See [“Row packing and table definitions” on page 70](#).
- **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. When your cache cannot hold enough pages, performance suffers as UltraLite begins swapping frequently-used pages to disk. See [“UltraLite CACHE_SIZE connection parameter” on page 226](#).
- **Index entries** Page size also affects indexes. The larger the database page, the more index entries it can hold. See [“Working with UltraLite indexes” on page 79](#).
- **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, select 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

- [“UltraLite optimization strategies” on page 121](#)
- [“Row packing and table definitions” on page 70](#)
- [“UltraLite case creation parameter” on page 183](#)
- [“UltraLite CACHE_SIZE connection parameter” on page 226](#)
- [“UltraLite RESERVE_SIZE connection parameter” on page 244](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

Example

To set the page size of the database to 8 KB, specify **page_size=8k** or **page_size=8192**:

```
ulcreate test.udb -o page_size=8k
```

UltraLite precision creation parameter

Specifies the maximum number of digits in decimal point arithmetic results.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 scale creation parameter” on page 203](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [UltraLite for embedded SQL: “ULCreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C++: “CreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite.NET: “CreateDatabase method” \[*UltraLite - .NET Programming*\]](#)
- [UltraLite for M-Business Anywhere: “createDatabase method” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [“Accessing creation parameter values” on page 34](#)

UltraLite scale creation parameter

Specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum precision.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 201
- [“UltraLite Create Database utility \(ulcreate\)”](#) on page 255
- [“UltraLite Initialize Database utility \(ulinit\)”](#) on page 269
- [“UltraLite Load XML to Database utility \(ulload\)”](#) on page 272
- UltraLite for embedded SQL: [“ULCreateDatabase function”](#) [*UltraLite - C and C++ Programming*]
- UltraLite for C++: [“CreateDatabase function”](#) [*UltraLite - C and C++ Programming*]
- UltraLite.NET: [“CreateDatabase method”](#) [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: [“createDatabase method”](#) [*UltraLite - M-Business Anywhere Programming*]
- [“Accessing creation parameter values”](#) on page 34

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.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 207](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)

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.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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 maintain timestamps and timestamp increments to the same resolution. By setting the 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.

Symbol	Description
<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 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, select 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 210](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [UltraLite for embedded SQL: “ULCreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C++: “CreateDatabase function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite.NET: “CreateDatabase method” \[*UltraLite - .NET Programming*\]](#)
- [UltraLite for M-Business Anywhere: “createDatabase method” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [“Accessing creation parameter values” on page 34](#)
- [“Timestamp-based downloads” \[*MobiLink - Server Administration*\]](#)
- [“Concurrency in UltraLite” on page 13](#)

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.

Syntax

```
{ ulcreate | ulinit | ulload } -o 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, select 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 207](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: [“ULCreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C++: [“CreateDatabase function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“CreateDatabase method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“createDatabase method” \[UltraLite - M-Business Anywhere Programming\]](#)
- [“Accessing creation parameter values” on page 34](#)
- [“Timestamp-based downloads” \[MobiLink - Server Administration\]](#)
- [“Concurrency in UltraLite” on page 13](#)

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 utf8_encoding creation parameter

Encodes data using the UTF-8 format, 8-bit multibyte encoding for Unicode.

Syntax

```
{ ulcreate | ulinit | ulload } -o utf8_encoding=value;...
```

Values

Boolean.

Default

0 (databases are not 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.

By encoding your database in UTF-8, UltraLite uses the UTF8BIN collation to sort characters. The UTF8BIN character set is not specific to any particular native language; no specific code page is associated with this character set. So, you can synchronize data from multiple native languages to the same consolidated database. 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, select 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 37](#)
- [“UltraLite character sets” on page 36](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- UltraLite for embedded SQL: “ULCreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite for C++: “CreateDatabase function” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “CreateDatabase method” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “createDatabase method” [*UltraLite - M-Business Anywhere Programming*]
- [“Accessing creation parameter values” on page 34](#)

UltraLite database properties

Contents

Accessing UltraLite database properties 218

Property	Description
case	Returns the status of the case sensitivity feature. Returns On if the database is case sensitive. Otherwise, it returns Off. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite case creation parameter” on page 183 .
char_set	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. See also: <ul style="list-style-type: none">• “UltraLite utf8_encoding creation parameter” on page 211• “UltraLite collation creation parameter” on page 186 The value of this property is set when the database is created, and can only be changed by creating a new database.
checksum_level	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). The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite checksum_level creation parameter” on page 184 .
collation	Returns the name of the database's collation sequence. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite collation creation parameter” on page 186 .

Property	Description
commit_flush_count	Returns the value of the <code>commit_flush_count</code> option that sets a commit count threshold. See “UltraLite commit_flush_count option [temporary]” on page 220 .
commit_flush_timeout	Returns the value of the <code>commit_flush_timeout</code> option that sets a time interval threshold. See “UltraLite commit_flush_timeout option [temporary]” on page 221 .
conn_count	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.
date_format	Returns the date format the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite date_format creation parameter” on page 187 .
date_order	Returns the date order the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite date_order creation parameter” on page 190 .

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"> ● UltraLite C/C++: “ULChangeEncryptionKey function” [UltraLite - C and C++ Programming] ● UltraLite.NET: “ChangeEncryptionKey method” [UltraLite - .NET Programming] ● UltraLite for M-Business Anywhere: “changeEncryptionKey method” [UltraLite - M-Business Anywhere Programming] ● “Securing UltraLite databases” on page 40 ● “UltraLite fips creation parameter” on page 192 ● “UltraLite obfuscate creation parameter” on page 198 ● “UltraLite DBKEY connection parameter” on page 233
file	<p>Returns the name of the database root file for the current connection, the including path. This is the value specified in the DBF connection parameter value. See:</p> <ul style="list-style-type: none"> ● UltraLite C/C++: “GetDatabaseProperty function” [UltraLite - C and C++ Programming] ● UltraLite.NET: “GetDatabaseProperty method” [UltraLite - .NET Programming] ● UltraLite for M-Business Anywhere: “getDatabaseProperty method” [UltraLite - M-Business Anywhere Programming] ● “UltraLite DBF connection parameter” on page 231
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 222.</p>
max_hash_size	<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 194.</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 223 .
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 234 ● “UltraLite DBF connection parameter” on page 231
nearest_century	Returns the nearest century the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite nearest_century creation parameter” on page 196 .
page_size	Returns the page size of the database, in bytes. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite page_size creation parameter” on page 199 .
precision	Returns the floating-point precision the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite precision creation parameter” on page 201 .
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. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite scale creation parameter” on page 203 .
time_format	Returns the time format the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite time_format creation parameter” on page 205 .
timestamp_format	Returns the timestamp format the database uses for string conversions. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite timestamp_format creation parameter” on page 207 .

Property	Description
timestamp_increment	Returns the minimum difference between two unique timestamps, in microseconds. The value of this property is set when the database is created, and can only be changed by creating a new database. See “UltraLite timestamp_increment creation parameter” on page 210.

Accessing UltraLite database properties

UltraLite provides a set of 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.

To browse UltraLite database properties (Sybase Central)

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

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

3. If the database properties have changed since you started browsing them, click **Refresh**.

To get the value of a database property (C/C++)

- In C/C++, call the `GetDatabaseProperty` function.

For example, to get the value of the `conn_count` property, call:

```
GetDatabaseProperty( ul_database_property_id conn_count )
```

To get the value of the `char_set` property, call:

```
GetDatabaseProperty( ul_database_property_id char_set )
```

See also

- UltraLite C/C++: “[GetDatabaseProperty function](#)” [*UltraLite - C and C++ Programming*]
- UltraLite.NET: “[GetDatabaseProperty method](#)” [*UltraLite - .NET Programming*]
- UltraLite for M-Business Anywhere: “[getDatabaseProperty method](#)” [*UltraLite - M-Business Anywhere Programming*]

UltraLite database options

Contents

UltraLite commit_flush_count option [temporary]	220
UltraLite commit_flush_timeout option [temporary]	221
UltraLite global_database_id option	222
UltraLite ml_remote_id option	223
Change UltraLite persistent database option settings	224

Options are used to configure database behavior. Database options can be set or modified at any time. In UltraLite, options can be persistent or temporary. Persistent options are stored in the database in the sysuldata system table. Temporary option settings only persist while the database is running.

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

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

You can obtain the current setting of a database option by querying the value of the corresponding database property or by using the appropriate get database property method. For example, to obtain the current setting of the commit_flush_timeout database option, execute the following SQL statement:

```
SELECT DB_PROPERTY ( 'commit_flush_timeout' );
```

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.

You must set this option each time you start the database if it is required.

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.

See also

- [“Flushing single or grouped transactions” on page 119](#)
- [“UltraLite `commit_flush_timeout` option \[temporary\]” on page 221](#)
- [“UltraLite `COMMIT_FLUSH` connection parameter” on page 229](#)
- [“UltraLite `SET OPTION` statement” on page 488](#)
- UltraLite for C/C++: [“`SetDatabaseOption` function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“`ULSetDatabaseOptionString` function” \[UltraLite - C and C++ Programming\]](#)

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.

You must set this option each time you start database, if it is required.

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.

See also

- [“UltraLite commit_flush_count option \[temporary\]” on page 220](#)
- [“UltraLite COMMIT_FLUSH connection parameter” on page 229](#)
- [“UltraLite SET OPTION statement” on page 488](#)
- UltraLite for C/C++: [“SetDatabaseOption function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“ULSetDatabaseOptionString function” \[UltraLite - C and C++ Programming\]](#)

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. See [“Using global autoincrement” \[MobiLink - Server Administration\]](#).

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.

You can also set this option using the ulinfo utility:

```
ulinfo -g ID ...
```

To set the global_database_id option from a client application, use the set database ID function for the programming interface you are using.

See also

- [“Change UltraLite persistent database option settings” on page 224](#)
- [“UltraLite Information utility \(ulinfo\)” on page 265](#)
- [“UltraLite SET OPTION statement” on page 488](#)
- [“Using GLOBAL AUTOINCREMENT in UltraLite” on page 129](#)
- UltraLite for C/C++: [“SetDatabaseID function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“ULSetDatabaseID function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“DatabaseID property” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“setDatabaseID method” \[UltraLite - M-Business Anywhere Programming\]](#)

Example

To autoincrement UltraLite database columns from 3001 to 4000, set the global ID to 3.

```
SET OPTION global_database_id="3";
```

UltraLite ml_remote_id option

A unique identifier in UltraLite that is used for MobiLink synchronization.

Allowed values

Any value that uniquely identifies the database for MobiLink synchronization.

Default

Null

Remarks

The **remote ID** is a unique identifier for an UltraLite remote that is used for MobiLink synchronization. The MobiLink remote ID is initially set to NULL. During the first synchronization, the MobiLink server sets the remote ID to a GUID. However, the remote ID can be any string that has meaning to you, provided that the string remains unique among all remote MobiLink clients. This uniqueness requirement is always enforced.

The remote ID stores the synchronization progress for the MobiLink user name. By including a unique remote ID, 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 client database. In this case, your synchronization scripts should reference the remote ID and not just the user name.

You can also set this option using the ulinfo utility:

```
ulinfo -r ID ...
```

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.

See also

- [“Change UltraLite persistent database option settings” on page 224](#)
- [“Remote IDs” \[MobiLink - Client Administration\]](#)
- [“Grant REMOTE permission” \[SQL Remote\]](#)
- [“UltraLite Information utility \(ulinfo\)” on page 265](#)
- [“UltraLite SET OPTION statement” on page 488](#)
- [UltraLite for C/C++: “SetDatabaseOption function” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite for embedded SQL: “ULSetDatabaseOptionString function” \[UltraLite - C and C++ Programming\]](#)
- [UltraLite.NET: “SetDatabaseOption method” \[UltraLite - .NET Programming\]](#)
- [“UltraLite clients” on page 125](#)
- [“UltraLite user authentication” \[MobiLink - Client Administration\]](#)
- [“User Name synchronization parameter” on page 175](#)
- [and “Password synchronization parameter” on page 164](#)

Change UltraLite persistent database option settings

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.

To browse or modify persistent UltraLite database options (Sybase Central)

1. Connect to the database.
2. Right-click the database and choose **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 SET OPTION statement” on page 488](#)
- [“Accessing UltraLite database properties” on page 218](#)
- [“DB_PROPERTY function \[System\]” on page 380](#)
- UltraLite C/C++: [“GetDatabaseProperty function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“GetDatabaseProperty method” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business Anywhere: [“getDatabaseProperty method” \[UltraLite - M-Business Anywhere Programming\]](#)

UltraLite connection parameters

Contents

UltraLite CACHE_SIZE connection parameter	226
UltraLite CE_FILE connection parameter	227
UltraLite COMMIT_FLUSH connection parameter	229
UltraLite CON connection parameter	230
UltraLite DBF connection parameter	231
UltraLite DBKEY connection parameter	233
UltraLite DBN connection parameter	234
UltraLite MIRROR_FILE connection parameter	235
UltraLite NT_FILE connection parameter	237
UltraLite ORDERED_TABLE_SCAN connection parameter [deprecated]	239
UltraLite PALM_ALLOW_BACKUP connection parameter	240
UltraLite PALM_FILE connection parameter	241
UltraLite PWD connection parameter	243
UltraLite RESERVE_SIZE connection parameter	244
UltraLite START connection parameter	245
UltraLite UID connection parameter	246

UltraLite CACHE_SIZE connection parameter

Defines the size of the database cache.

Syntax

`CACHE_SIZE=number{ k | m | g }`

Default

The default cache size is determined by the amount of memory available on your system and the size of the database.

Remarks

If the cache size is not specified, or if you set the size to 0, the default size is used. If your testing shows the need for better performance, you should increase the cache size.

By default, the size is in bytes. Use k, m, or g to specify units of kilobytes, megabytes, or gigabytes, respectively.

If you exceed the maximum cache size, it is automatically replaced with your platform's upper cache size limit. Increasing the cache size beyond the size of the database does not provide any performance improvement, and a large cache size can interfere with the number of other applications you can use.

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 optimization strategies” on page 121](#)
- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite page_size creation parameter” on page 199](#)
- [“UltraLite RESERVE_SIZE connection parameter” on page 244](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following connection string fragment sets the cache size to 2 MB.

```
"CACHE_SIZE=2m"
```

UltraLite CE_FILE connection parameter

Names the new database file when creating a database. When opening a connection to an existing database, this connection parameter identifies the database.

Syntax

```
CE_FILE=path\ce-db
```

Default

- DBF connection parameter.
- If you do not set a value for the DBF connection parameter or this connection parameter, then the default value is `\UltraLiteDB\ulstore.udb`.

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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 49](#)
- [“UltraLite DBF connection parameter” on page 231](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following example creates a new connection and identifies different database files for the Windows desktop and Windows Mobile platforms:

UltraLite connection parameters

```
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

- [“Flushing single or grouped transactions” on page 119](#)
- [“UltraLite commit_flush_count option \[temporary\]” on page 220](#)
- [“UltraLite commit_flush_timeout option \[temporary\]” on page 221](#)
- [“UltraLite CHECKPOINT statement” on page 458](#)
- UltraLite embedded SQL: [“ULCheckpoint function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite C++: [“Checkpoint function” \[UltraLite - C and C++ Programming\]](#)

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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite DBN connection parameter” on page 234](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following connection string fragment sets the first connection name to MyFirstCon.

```
"CON=MyFirstCon"
```


UltraLite DBF connection parameter

Names the new database file when creating a database. When opening a connection, this connection parameter indicates which database file you want to load and connect to.

Syntax

DBF=*ul-db*

Default

- **On desktop platforms** If you do not specify an NT_FILE or DBF value, UltraLite sets the file name to `\UltraLiteDB\ulstore.udb`.
- **On Windows Mobile** If you do not specify a CE_FILE or DBF value, UltraLite sets the file name to `\UltraLiteDB\ulstore.udb`.
- **On Palm OS** If you do not specify a PALM_FILE or DBF value, UltraLite sets the file name to `ulstore.udb`.

Remarks

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
- PALM_FILE
- 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.

Palm OS You cannot deploy databases to VFS volumes using the Palm install tool. Instead, you need to copy the database directly onto the media with a card reader, or some other tool.

If you are creating a file with the DBF parameter, you must name the database with a `.pdb` extension. This database name is used while administering the Palm OS database on the desktop. Administration tools and utilities can connect to the database with the DBF connection parameter and the database with the `.pdb` extension. Once you deploy the file to the device however, the `.pdb` extension is dropped. Your application must then connect to the database using the Palm_FILE parameter and no extension.

For example, if you name a file `CustDB.pdb` use **DBF=CustDB.pdb** with your administration tools desktop. Then, when the database is deployed to the device the application connects to the same database with **PALM_FILE=CustDB**.

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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 49](#)
- [“UltraLite DBN connection parameter” on page 234](#)
- [“UltraLite CE_FILE connection parameter” on page 227](#)
- [“UltraLite PALM_FILE connection parameter” on page 241](#)
- [“UltraLite NT_FILE connection parameter” on page 237](#)
- [UltraLite for C/C++: “Connecting to a database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C/C++: “OpenConnection function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connecting to a database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connecting to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [UltraLite for M-Business Anywhere: “openConnection method” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [UltraLite.NET: “Connecting to a database” \[*UltraLite - .NET Programming*\]](#)
- [UltraLite.NET: “Open method” \[*UltraLite - .NET Programming*\]](#)

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

Provides an encryption key for the database when creating a new database. When opening a connection to an existing database, this connection parameter 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 128-bit or AES FIPS 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.

On Palm OS, applications are automatically shut down by the system whenever a user switches to a different application. However, you can program your UltraLite client to circumvent the need to re-enter the key each time a user switches back to the application again.

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 for C/C++: “Saving, retrieving, and clearing encryption keys on Palm OS” [[UltraLite - C and C++ Programming](#)]
- “Opening UltraLite connections with connection strings” on page 48
- “Securing UltraLite databases” on page 40
- “UltraLite obfuscate creation parameter” on page 198
- UltraLite for C/C++: “Connecting to a database” [[UltraLite - C and C++ Programming](#)]
- UltraLite for C/C++: “OpenConnection function” [[UltraLite - C and C++ Programming](#)]
- UltraLite for embedded SQL: “Connecting to a database” [[UltraLite - C and C++ Programming](#)]
- UltraLite for M-Business Anywhere: “Connecting to an UltraLite database” [[UltraLite - M-Business Anywhere Programming](#)]
- UltraLite for M-Business Anywhere: “openConnection method” [[UltraLite - M-Business Anywhere Programming](#)]
- UltraLite.NET: “Connecting to a database” [[UltraLite - .NET Programming](#)]
- UltraLite.NET: “Open method” [[UltraLite - .NET Programming](#)]

UltraLite DBN connection parameter

Differentiates databases by name when applications connect to more than one database.

Syntax

DBN=*db-name*

Default

- **Palm OS** The creator ID.
- **All other platforms** The file name without the path and extension, if they exist. Cannot exceed 16 characters in length.

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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite DBF connection parameter” on page 231](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

Use the following parameters to connect to the running UltraLite database named Kitchener:

```
DBN=Kitchener;DBF=cities.udb
```

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, including Palm VFS, but not on Palm record databases.

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. See [“UltraLite checksum_level creation parameter” on page 184](#).

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 [“UltraLite Validate Database utility \(ulvalid\)” on page 286](#).

See also

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 49](#)
- [UltraLite for C/C++: “Connecting to a database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for C/C++: “OpenConnection function” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for embedded SQL: “Connecting to a database” \[*UltraLite - C and C++ Programming*\]](#)
- [UltraLite for M-Business Anywhere: “Connecting to an UltraLite database” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [UltraLite for M-Business Anywhere: “openConnection method” \[*UltraLite - M-Business Anywhere Programming*\]](#)
- [UltraLite.NET: “Connecting to a database” \[*UltraLite - .NET Programming*\]](#)
- [UltraLite.NET: “Open method” \[*UltraLite - .NET Programming*\]](#)

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

Names the new database file when creating a database. When opening a connection to an existing database, the parameter identifies the database.

Syntax

`NT_FILE=path\nt-db`

Default

- DBF connection parameter.
- If you do not set a value for the DBF connection parameter or this connection parameter, then the default value is `\UltraLiteDB\ulstore.udb`.

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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)
- [“Precedence of connection parameters for UltraLite administration tools” on page 49](#)
- [“UltraLite DBF connection parameter” on page 231](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following example creates a new connection and identifies different database files for the desktop, Palm OS, and Windows Mobile platforms:

UltraLite connection parameters

```
Connection = DatabaseMgr.OpenConnection( "UID=JDoe;PWD=ULdb;  
NT_FILE=c:\test\MyTestDB.udb;CE_FILE=\database  
\MyCEDB.udb;PALM_FILE=MyPalmDB_MyCreatorID" )
```


UltraLite ORDERED_TABLE_SCAN connection parameter [deprecated]

Controls whether queries access rows via a primary key index or from the database page directly.

Syntax

```
ORDERED_TABLE_SCAN={ yes | no }
```

Default

no (the database page is scanned directly)

Remarks

If you want results returned with the row order of the primary key index, rewrite your queries to include the ORDER BY clause first. Otherwise, you cannot take advantage of the performance benefits introduced by allowing direct page scans for this connection. Only use this connection parameter if it is impractical to rewrite your queries to use the ORDER BY clause.

In versions 10.0.0 and earlier of UltraLite, UltraLite used the primary key index to return results when no other index was selected by the UltraLite optimizer.

See also

- [“Using direct page scans” on page 117](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

UltraLite PALM_ALLOW_BACKUP connection parameter

Controls backup behavior over HotSync, which is disabled in UltraLite by default.

Syntax

PALM_ALLOW_BACKUP={ yes | no }

Default

no

Remarks

This connection parameter is supported for desktop backup using the UltraLite HotSync conduit for Palm OS devices.

On Palm, you can back up the database to desktop using HotSync. In most UltraLite client applications, data is backed up by synchronization. There is no need to use informal backups to the desktop. Therefore, the UltraLite runtime disables Palm's backup behavior. However, if your deployment explicitly requires that HotSync back up the UltraLite database to the desktop while it is also being synchronized, use this connection parameter to override the UltraLite default.

Once you have enabled backups for HotSync, you do not need to configure backups in ULDBUtil. Backups occur with each synchronization attempt until you disable them.

See also

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite Data Management utility for Palm OS \(ULDBUtil\)” on page 258](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

UltraLite PALM_FILE connection parameter

Names the new database file when creating a database. When opening a connection to an existing database, this connection parameter identifies the database.

Syntax 1: record based stores

PALM_FILE=*name*

Syntax 2: file-based stores

PALM_FILE=*vfs*:*[volume-label:| volume-ordinal:]filename*

Default

- DBF connection parameter.
- If you do not set a value for the DBF connection parameter or this connection parameter, then the default value is *ulstore.udb*.

Remarks

You should use the PALM_FILE connection parameter for UltraLite client applications that use the same connection string to connect to a Palm device, and other platforms.

This 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 Palm OS database, use DBF.

The value of PALM_FILE must meet the file name requirements for Palm OS platforms.

You cannot deploy databases to VFS volumes using the Palm install tool. Instead, you need to copy the database directly onto the media with a card reader, or some other tool.

If you are creating a file, you must use the DBF connection parameter and name the database with a *.pdb* extension. This database name is used while administering the Palm OS database on the desktop.

Administration tools and utilities can connect to the database with the DBF connection parameter and the database with the *.pdb* extension. Once you deploy the file to the device, the *.pdb* extension is dropped. Your application must then connect to the database using the PALM_FILE connection parameter and no extension.

For example, if you name a file *CustDB.pdb* use **DBF=CustDB.pdb** with your administration tools desktop. When the database is deployed to the device, the application connects to the same database with **PALM_FILE=CustDB**.

On VFS stores, always specify the absolute file path. If directories are specified in the full path name, they are created if they do not already exist.

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

- “Opening UltraLite connections with connection strings” on page 48
- “Specifying file paths in an UltraLite connection parameter” on page 47
- “Precedence of connection parameters for UltraLite administration tools” on page 49
- “Registering the Palm creator ID” [*UltraLite - C and C++ Programming*]
- “UltraLite DBF connection parameter” on page 231
- UltraLite for C/C++: “Connecting to a database” [*UltraLite - C and C++ Programming*]
- UltraLite for C/C++: “OpenConnection function” [*UltraLite - C and C++ Programming*]
- UltraLite for embedded SQL: “Connecting to a database” [*UltraLite - C and C++ Programming*]
- UltraLite for M-Business Anywhere: “Connecting to an UltraLite database” [*UltraLite - M-Business Anywhere Programming*]
- UltraLite for M-Business Anywhere: “openConnection method” [*UltraLite - M-Business Anywhere Programming*]
- UltraLite.NET: “Connecting to a database” [*UltraLite - .NET Programming*]
- UltraLite.NET: “Open method” [*UltraLite - .NET Programming*]

Example

The following example creates a new connection and identifies different database files for the desktop, Palm OS, and Windows Mobile platforms:

```
Connection = DatabaseMgr.OpenConnection("UID=JDoe;PWD=ULdb;  
NT_FILE=c:\test\MyTestDB.udb;CE_FILE=\database  
\MyCEDB.udb;PALM_FILE=MyPalmDB_MyCreatorID")
```

UltraLite PWD connection parameter

Defines the password for a user ID that is used for authentication.

Syntax

PWD=*password*

Default

If you do not set both the UID and PWD, UltraLite opens connections with **UID=DBA** and **PWD=sql**.

Remarks

You can set passwords to NULL or an empty string, but they cannot exceed the maximum length of 31 characters.

Every user of a database has a password. UltraLite supports up to four user ID/password combinations.

This connection parameter is not encrypted. However, UltraLite hashes the password before saving it, so you can only modify a password from Sybase Central.

See also

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite user authentication” on page 51](#)
- [“Interpreting user ID and password combinations” on page 51](#)
- [“UltraLite UID connection parameter” on page 246](#)
- UltraLite for C/C++: [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Authenticating users” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Authenticating users” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

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. By reserving the file system space means that the space cannot be 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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite CACHE_SIZE connection parameter” on page 226](#)
- [“UltraLite page_size creation parameter” on page 199](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

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.

Syntax

START=*path*\uleng11.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

- [“UltraLite Engine utility \(uleng11\)” on page 260](#)
- [“Choosing an UltraLite data management component” on page 23](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following command starts the UltraLite engine that is located in the *Program Files* directory:

```
Start="Program Files\uleng11.exe"
```

An alternative way to define this path is to put the entire string in single quotes:

```
Start='\"Program Files\uleng11.exe\"'
```

UltraLite UID connection parameter

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, 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

- [“Opening UltraLite connections with connection strings” on page 48](#)
- [“UltraLite user authentication” on page 51](#)
- [“Interpreting user ID and password combinations” on page 51](#)
- UltraLite for C/C++: [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for C/C++: [“OpenConnection function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Authenticating users” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for embedded SQL: [“Connecting to a database” \[UltraLite - C and C++ Programming\]](#)
- UltraLite for M-Business Anywhere: [“Authenticating users” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“Connecting to an UltraLite database” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite for M-Business Anywhere: [“openConnection method” \[UltraLite - M-Business Anywhere Programming\]](#)
- UltraLite.NET: [“Authenticating users” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Connecting to a database” \[UltraLite - .NET Programming\]](#)
- UltraLite.NET: [“Open method” \[UltraLite - .NET Programming\]](#)

Example

The following connection string fragment supplies the user ID DBA and password sql for a database:

```
"UID=DBA;PWD=sql"
```

UltraLite utilities

Contents

Supported exit codes	248
Interactive SQL utility for UltraLite (dbisql)	249
SQL Preprocessor for UltraLite utility (sqlpp)	252
UltraLite Create Database utility (ulcreate)	255
UltraLite Data Management utility for Palm OS (ULDBUtil)	258
UltraLite Engine utility (uleng11)	260
UltraLite Engine Stop utility (ulstop)	261
UltraLite Erase database (ulerase)	262
UltraLite HotSync Conduit Installation utility for Palm OS (ulcond11)	263
UltraLite Information utility (ulinfo)	265
UltraLite Initialize Database utility (ulinit)	269
UltraLite Load XML to Database utility (ulload)	272
UltraLite Synchronization utility (ulsync)	275
Synchronization profile options	278
UltraLite Unload Database utility (ulunload)	281
UltraLite Unload Old Database utility (ulunloadold)	284
UltraLite Validate Database utility (ulvalid)	286

UltraLite includes a set of utilities that are designed to perform basic database administration activities from 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 chapter are case sensitive, unless otherwise noted. Type options *exactly* as they are displayed.

Supported exit codes

The ulcreate, 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 utility for UltraLite (dbisql)

Executes SQL commands or runs command files.

Syntax

dbisql -c "*connection-string*" **-ul** [*options*] [*dbisql-command* | *command-file*]

Option	Description
@data	<p>Read options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.</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 “Using configuration files” [SQL Anywhere Server - Database Administration].</p>
-c " <i>connection-string</i> "	<p>Required. Connect to the database as identified in the DBF or file_name connection 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.</p>
-d <i>delimiter</i>	<p>Specify 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>Command delimiters are used for all connections in that Interactive SQL session, regardless of the setting stored in the database.</p>
-dl	<p>Echo all statements explicitly executed by the user to the Command window (STDOUT). This echo can provide useful feedback for debugging SQL scripts, or when Interactive SQL is processing a long SQL script.</p>
-f <i>file-name</i>	<p>Open but do not run the file called <i>filename</i>. Enclose the file name in quotation marks if the file name contains a space.</p> <p>If the -f option is given, the -c option is ignored; that is, no connection is made to the database.</p> <p>If you do not include a full path to the file, a path relative to the current directory is assumed. If the file does not exist, or if it is a directory instead of a file, Interactive SQL prints an error message and then shuts down.</p> <p>This option is only supported when Interactive SQL is run as a windowed application.</p>
-nogui	<p>Run in command-prompt mode. If you specify either <i>dbisql-command</i> or <i>command-file</i>, then -nogui is assumed.</p>

Option	Description
-onerror <i>behavior</i>	Control what happens if an error is encountered while reading data from the specified command file. Define one of the following supported <i>behavior</i> values: <ul style="list-style-type: none"> • Stop Interactive SQL stops executing statements. • Prompt Interactive SQL prompts the user to see if the user wants to continue. • Continue The error is ignored and Interactive SQL continues executing statements. • Exit Interactive SQL terminates. • Notify_Continue The error is reporting, and the user is prompted to press Enter or click OK to continue. • Notify_Stop The error is reported, and the user is prompted to press Enter or click OK to stop executing statements. • Notify_Exit The error is reported and the user is prompted to press Enter or click OK to terminate Interactive SQL.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-ul	Connect to UltraLite databases by default. By default, Interactive SQL assumes that you are connecting to SQL Anywhere databases. When you specify the -ul option, the default becomes UltraLite connections. Regardless of the type of database set as the default, you can connect to either SQL Anywhere or UltraLite databases. Perform this action by choosing the database type from the dropdown list on the Connect window.
-version	Displays the version number of Interactive SQL. You can also view the version number from within Interactive SQL; from the Help menu, choose About Interactive SQL .
-x	Scan commands, but do not execute them. This option is useful for checking long command files for syntax errors.
<i>SQL-command</i> <i>command-file</i>	Execute the SQL statement or the specified <i>command-file</i> . If you do not specify a <i>SQL-statement</i> or <i>command-file</i> , Interactive SQL enters interactive mode, where you can type a command into a command window.

Remarks

Interactive SQL allows you to execute SQL commands or run command files. It also provides feedback about:

- the number of rows affected
- the time required for each command
- the execution plan of queries
- any error messages

When you connect to an UltraLite database, menu items that are SQL Anywhere-specific are not displayed in the interface. For example, **Tools » Lookup Procedure Name**, or **Tools » Index Consultant**.

In UltraLite, collations include a code page plus a sort order. Therefore, code page numbers correspond to the number displayed as part of the UltraLite collation name. To see a list of supported collations (and their corresponding code pages), run `ulcreate -l` at a command prompt.

You can specify a code page to use when reading or writing files using the `ENCODING` clause of the `INPUT`, `OUTPUT`, or `READ` statement. For example, on an English 32-bit Windows desktop computer, windowed programs use the 1252 (ANSI) code page. If you want Interactive SQL to read a file named `status.txt` created using the 297 (IBM France) code page, use the following statement:

```
READ
ENCODING 297
status.txt;
```

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).

In command line 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 also

- “Using 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 225
- “UltraLite Create Database utility (ulcreate)” on page 255
- “Supported exit codes” on page 248

Example

The following command runs the command 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. If there is an error in the command 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. Note that 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 <i>flag</i>	<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 <i>width</i>	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
-n	<p>Generate line number information in the C file by using <code>#line</code> directives in the appropriate places in the generated code.</p> <p>Use this option to the report source errors and to debug source on line numbers in the <i>sql-filename</i> file, rather than in the <i>output-filename</i> file.</p>
-o <i>O/S spec</i>	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 <i>string-length</i>	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 <i>level</i>	<p>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.</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>
-x	Change multibyte strings to escape sequences, so that they can be passed through a compiler.

Option	Description
<code>-z collation-sequence</code>	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 *sql*. 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 its corresponding codepage), run `ulcreate -l` at a command prompt.

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 [“Testing SQL compliance using the SQL Flagger” \[SQL Anywhere Server - SQL Usage\]](#).

See also

- [“SQL Anywhere embedded SQL” \[SQL Anywhere Server - Programming\]](#)
- [“UltraLite character sets” on page 36](#)

Example

The following command preprocesses the *srcfile.sql* embedded SQL file in quiet mode for an UltraLite application.

```
sqlpp -u -q MyEsqFile.sql
```


UltraLite Create Database utility (ulcreate)

Creates an UltraLite database with the properties you define.

Syntax

ulcreate [*options*][*new-database-file*]

Option	Description
-c " <i>connection-string</i> "	Create 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. If you do specify a userid and password, those will be required to access the database. If you do not provide a file name as a parameter in the connection string, ulcreate checks the end of the command for the file you specified as <i>new-database-file</i> .
-g <i>global-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 222 .
-l	List the available collation sequences and exit.
-o [<i>extended-options</i>]	Specify a semicolon separated list of UltraLite database creation parameters. See “Choosing database creation parameters for UltraLite” on page 34 .
-ol	List the available database creation parameters and exit. See “Choosing database creation parameters for UltraLite” on page 34 .
-p <i>creator-ID</i>	Required for Palm OS when the database is being installed to the record store. Create the database with the specified four character <i>creator-ID</i> of the UltraLite client application. If you are deploying the database to a VFS store, then do not use this option.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-t <i>file</i>	Loads the specified file containing the public trusted root certificate into the database, and eliminates the requirement to download the trusted_certificates sync parameter. This certificate is required for server authentication.
-v	Print verbose messages.
-y	Overwrite the database file if it exists.
-z <i>collation-sequence</i>	Specify the label of the collation to be used.

Option	Description
<i>new-database-file</i>	Create a file with the specified name. Only use this standalone file name if you are not using a connection string to set initial database parameters like a user ID (UID) or a password (PWD). Ensure the standalone file name you set is appropriate for your platform.

Remarks

If you do not set any database properties, `ulcreate` creates a case insensitive database with a collation sequence that depends on the current locale.

Database passwords are always case sensitive, regardless of the case-sensitivity of the database. Database case sensitivity depends on whether you use the `case=respect` creation parameter.

The collation sequence is used for all string comparisons in the database. In UltraLite, collations include a code page plus a sort order. If you do not specify `-z`, `ulcreate` attempts to determine a reasonable collation to use based on the current locale of the desktop.

To see a list of supported collations (and its corresponding codepage), run `ulcreate -l` at a command prompt.

The operating system of your device determines if UTF-8 encoding should be used.

Palm OS databases written to the desktop must be identified with the `.pdb` extension. However, once you deploy the database to the device, the extension is dropped. For more information about file name formats, see [“Palm OS” on page 47](#).

You cannot deploy UltraLite databases with the Palm install tool if the target is a VFS volume. Instead, you need to copy the database directly onto the media using a card reader, or some other tool.

This utility returns error codes. Any value other than 0 means that the operation failed.

See also

- [“Creating and configuring UltraLite databases” on page 29](#)
- [“Specifying file paths in an UltraLite connection parameter” on page 47](#)
- [“UltraLite connection parameters” on page 225](#)
- [“Supported exit codes” on page 248](#)
- [“Supported and alternate collations” \[SQL Anywhere Server - Database Administration\]](#)
- [“Registering the Palm creator ID” \[UltraLite - C and C++ Programming\]](#)
- [“UltraLite case creation parameter” on page 183](#)
- [“UltraLite platform requirements for character set encoding” on page 37](#)
- [“Introduction to MobiLink models” \[MobiLink - Getting Started\]](#)

Example

Create an UltraLite database called `test.udb` as a case-insensitive, non-Unicode database with a collation sequence that depends on the current locale:

```
ulcreate test.udb
```

Create a case sensitive database called `test.udb`, so that the database is created with ISO-compatible date formatting and ordering:

```
ulcreate -c DBF=test.udb -o case=respect;date_format=YYYY-MM-DD;date_order=YMD
```

Create an encrypted database called *test.udb* with the **afvc_1835** encryption key:

```
ulcreate -c "DBF=test.udb;DBKEY=afvc_1835"
```

UltraLite Data Management utility for Palm OS (ULDBUtil)

Deletes the database from the device, or backs up the database at the next synchronization.

Syntax

N/A.

Remarks

You can use this utility to perform the following tasks:

- Delete the database from the device when the device is shared among different users. Deleting the file allows you to save space or to maintain privacy. You can then reinstall the database or even have the database create a new unpopulated database.
- Back up the database upon next synchronization. Use this feature to perform an initial synchronization, and then back up the database. This allows you to deploy the database to other devices, so they do not need to perform an initial synchronization.

This utility is installed as the following file:

```
install-dir\UltraLite\Palm\68k\ULDBUtil.prc
```

Once installed, ULDBUtil uses a file name filter. This filter only detects UltraLite database files that meet UltraLite extension requirements. The file name filter list excludes all files with the following extensions: *.bak*, *.dat*, *.dba*, *.htm*, *.html*, *.in*, *.jpeg*, *.jpg*, *.prc*, *.pqa*, and *.tda*. Ensure you do not name an UltraLite database with any extension used by this filter.

To delete UltraLite application data from a Palm OS device

1. Switch to ULDBUtil.
2. If your device has expansion cards, pick the media (internal/external volume or record-based) from which the application database file is to be deleted.
3. From the list of UltraLite version 11 databases, select a database.
4. Tap Delete on the Palm device to remove the data.

The following procedure only applies to databases deployed with the HotSync conduit.

To back up the database from a Palm OS device to a desktop

1. Switch to ULDBUtil.
2. Select the Backup option so HotSync knows to back up the database on the next synchronization attempt. You need to select this option for each subsequent back up you need. This option is disabled for databases stored on VFS volumes such as cards.

Tip

If you want to enable automatic backups and thereby avoid the need to select this option each time you synchronize, use the PALM_ALLOW_BACKUP parameter. See [“UltraLite PALM_ALLOW_BACKUP connection parameter”](#) on page 240.

See also

- [“UltraLite HotSync Conduit Installation utility for Palm OS \(ulcond11\)”](#) on page 263

Example

N/A.

UltraLite Engine utility (uleng11)

Manages concurrent UltraLite database connections from applications on 32-bit Windows desktops and Windows Mobile.

Syntax

uleng11

Remarks

The UltraLite engine does not display a messages window on startup.

Linking to the UltraLite engine requires a different set of libraries than those used by the UltraLite runtime. If you also require TLS-enabled synchronization or AES FIPS database encryption, further libraries are required.

See:

- [“Compiling and linking your application” \[UltraLite - C and C++ Programming\]](#)
- [“Deploy UltraLite with AES_FIPS database encryption” on page 59](#)
- [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#)

See also

- [“Working with UltraLite databases” on page 69](#)
- [“Choosing an UltraLite data management component” on page 23](#)
- [“UltraLite Engine Stop utility \(ulstop\)” on page 261](#)
- [“UltraLite START connection parameter” on page 245](#)

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.

There are no options for this utility.

See also

- [“Choosing an UltraLite data management component” on page 23](#)
- [“UltraLite Engine utility \(uleng11\)” on page 260](#)

UltraLite Erase database (ulerase)

Erases an UltraLite database.

Syntax

ulerase -c "*connection-string*" [*options*] [*dbname*]

Option	Description
-c " <i>connection-string</i> " <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.
-ek <i>key</i>	Specify the encryption key for an encrypted database.
-ep	Specify that you want to be prompted for the encryption key.
-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.

Remarks

The database must be accessible. The User ID and password combination must allow a connection, otherwise the database will not be erased.

Encrypted databases require a key provided in the connection string, or using one of **-ek <key>** or **-ep**.

UltraLite HotSync Conduit Installation utility for Palm OS (ulcond11)

Installs and registers each Palm OS database with the HotSync conduit so the conduit can manage HotSync synchronization operations for each database. You can also use this utility to uninstall the conduit.

Syntax

ulcond11 -c "connection-string" [options] creator-ID

Option	Description
<i>creator-ID</i>	Required. Set the creator ID of the application that uses the conduit. If a conduit already exists for the specified creator ID, it is replaced by the new conduit.
-a	Append additional database connection strings to the connection string configured with the -c option. Use this option to register more than one database with the conduit.
-c "connection-string sync_profile=profile_name{ }	<p>Required. Connect to the Palm database on the device as identified by the DBF parameter of your <i>connection-string</i>. The connection string you define registers the deployed Palm database with the conduit. The connection parameters are stored as part of the conduit's configuration information.</p> <p>sync_profile specifies the sync profile to use, along with any merge parameters. The default profile is ul_palm_conduit.</p> <p>If you do not specify both a user ID and a password, the default UID of DBA and PWD of sql are assumed.</p>
-d filename	<p>Set the name of the plug-in <i>.dll</i> file.</p> <p>The utility creates a registry key (<i>Software\Sybase\SQL Anywhere\version\Conduit\creator-ID</i>) with a value named <i>PluginDLL</i>. The value for <i>PluginDLL</i> is the <i>filename</i> you set.</p> <p>Supply an empty string (-d"") to clear an existing file name from the registry.</p>
-n name	<p>Set the name displayed by the HotSync manager. The default value is Conduit.</p> <p>Do not use this option with the -u option.</p>
-q	<p>Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.</p> <p>If you are running ulcond11 on Windows Vista with the correct privileges, the window is shut down immediately when no errors are reported. Otherwise, a five second delay applies.</p>

Option	Description
-QQ	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however. If you are running ulcond11 on Windows Vista with the correct privileges, the window is shut down immediately.
-u	Uninstall the conduit for the creator ID. If you do not specify -u, a conduit is installed. If you specify -u, do not specify -n.

Remarks

ulcond11 is a command line utility. On Windows Vista, ulcond11 runs elevated. The executable window does not close for five seconds, allowing you to view the output.

HotSync records when each synchronization took place and whether each installed conduit worked as expected. The HotSync log file is in the subdirectory *User* of your Palm desktop installation directory.

See also

- [“UltraLite connection parameters” on page 225](#)
- [“UltraLite Data Management utility for Palm OS \(ULDBUtil\)” on page 258](#)
- [“Deploy the UltraLite HotSync conduit” on page 62](#)
- [“Network protocol options for UltraLite synchronization streams” on page 178](#)
- [“HotSync on Palm OS” on page 146](#)
- [“Registering the Palm creator ID” \[*UltraLite - C and C++ Programming*\]](#)

Examples

The following command installs a conduit for the application with creator ID **Syb2**, named **CustDB**. These are the settings for the CustDB sample application:

```
ulcond11 -c "DBF=custdb.udb;UID=DBA;PWD=sql" -n CustDB Syb2
```

The following command uninstalls the conduit for the same CustDB sample application:

```
ulcond11 -u Syb2
```

The following command installs the conduit for palmdb using the default sync profile for Palm:

```
ulcond11 ... -c  
"dbf=palmdb;sync_profile=ul_palm_conduit{stream.host=override_host}" ...
```

UltraLite Information utility (ulinfo)

Displays information about an UltraLite database and changes or clears the `global_id` or `ml_remote_id` database options.

Syntax

ulinfo -c "*connection-string*" [*options*]

Option	Description
-c " <i>connection-string</i> "	Required. Connect to the database as identified in the <code>DBF</code> or <code>file_name</code> 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.
-g <i>ID</i>	Set the initial global database ID to the value you assign. This value is used by the database for all new rows that have global autoincrement columns. The database uses this base value to autoincrement IDs associated with each additional row and/or column. When deploying an application, you must assign a different identification number to each database for synchronization with the MobiLink server.
-oa	Cancel the process if it identifies that the database was created by a previous version of UltraLite (which would cause the database to be upgraded).
-or	Open the database in read-only mode. UltraLite makes a copy of the original file, which you can then use to test your scripts without altering the database. Changes to the copied file are discarded upon completion. If you are connecting directly from the desktop to a database already deployed to a Windows Mobile device, this option is not supported.
-ou	Upgrade the database if it was created with an older UltraLite release.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-r <i>ID</i>	Set the initial <code>ml_remote_id</code> to the value you assign. By default, new UltraLite databases set the MobiLink remote ID to NULL. You can keep this default if you choose: both UltraLite and <code>dbmlsync</code> automatically set the MobiLink remote ID to a unique user ID (UUID) at the start of synchronization.
-rc	Set the MobiLink remote ID to NULL.
-v	Print verbose messages. Display current database properties in addition to the database internals for the named database.

Remarks

Warning messages generated when opening an UltraLite database are always displayed unless you use the `-q` option.

The `ulinfo` utility displays information about existing synchronization profiles and the number of pending SQL pass-through scripts unless the quiet (`-q`) flag is used. For example:

```
Sync profile cv_prosync:
"TableOrder=cv_prosync;MobiLinkId=cv_prosync;ScriptVersion=test;"
Number of available SQL pass-through scripts: 0
```

See also

- [“UltraLite connection parameters” on page 225](#)
- [“UltraLite global_database_id option” on page 222](#)
- [“UltraLite ml_remote_id option” on page 223](#)

Example

Show basic database internals for a file named *sample.udb* that has already been synchronized:

```
ulinfo -c DBF=cv_dbattr.udb

ulinfo -c DBF=cv_dbattr.udb
Utility Version 11.*suppressed*
Collation: 1252LATIN1
Number of Users: 1
  1. User: 'DBA'
Page size: 4096
Default index maximum hash size: 4
Checksum level: 0
MobiLink Remote ID: not set
Global database ID: 1000
Global autoincrement usage: 0%
Number of tables: 3
Number of columns: 7
Number of publications: 3
Number of tables that will always be uploaded: 0
Number of tables that are never synchronized: 0
Number of primary Keys: 3
Number of foreign Keys: 0
Number of indexes: 0
This database has not yet been synchronized.
Synchronization by publication number:
  1. Publication cv_sync
    Number of rows in next upload: 0
    Last download: 1900-01-01 00:00:00.000
  2. Publication cv_syncPub2
    Number of rows in next upload: 0
    Last download: 1900-01-01 00:00:00.000
  3. Publication cv_syncPub3
    Number of rows in next upload: 0
    Last download: 1900-01-01 00:00:00.000
Number of available SQL pass-through scripts: 0
ulsync ...
Utility Version 11.*suppressed*
Results of this synchronization:
Succeeded
Download timestamp: 20XX-XX-XX XX:XX:XX.XXXXXX
Upload OK
No ignored rows
```

```

No part download remaining
Authentication value: 1000 (0x3e8)
ulinfo -c dbf=cv_dbattr.udb
Utility Version 11.*suppressed*
Collation: 1252LATIN1
Number of Users: 1
  1. User: 'DBA'
Page size: 4096
Default index maximum hash size: 4
Checksum level: 0
MobiLink Remote ID: XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX
Global database ID: 1000
Global autoincrement usage: 0%
Number of tables: 3
Number of columns: 7
Number of publications: 3
Number of tables that will always be uploaded: 0
Number of tables that are never synchronized: 0
Number of primary Keys: 3
Number of foreign Keys: 0
Number of indexes: 0
Last synchronization completed successfully
Download occurred: *suppressed*
Upload OK
Upload rows not ignored
No partial downloads
Actual MobiLink Authentication value: 1000
Authentication valid
Synchronization by publication number:
  1. Publication cv_sync
     Number of rows in next upload: 0
     Last download: 20XX-XX-XX XX:XX:XX.XXXXXX
  2. Publication cv_syncPub2
     Number of rows in next upload: 0
     Last download: 1900-01-01 00:00:00.000
  3. Publication cv_syncPub3
     Number of rows in next upload: 0
     Last download: 1900-01-01 00:00:00.000
Number of available SQL pass-through scripts: 0

```

Show database internals for a file named *CustDB.udb* and display database properties by enabling verbose messaging:

```

ulinfo -c DBF=CustDB.udb -v

Utility Version 11.*suppressed*
Database information
Database name: custdb
Disk file 'C:\Documents and Settings\All Users\Shared Documents\SQL Anywhere
11\Samples\UltraLite\CustDB\custdb.udb'
Collation: 1252LATIN1
Number of Users: 1
  1. User: 'DBA'
Page size: 4096
Default index maximum hash size: 8
Checksum level: 0
MobiLink Remote ID: not set
Global database ID: not set
Encryption: None
Character encoding set: 1252LATIN1
Case sensitive: OFF
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: 6
Number of columns: 16
Number of publications: 1
Number of tables that will always be uploaded: 0
Number of tables that are never synchronized: 1
Number of primary Keys: 6
Number of foreign Keys: 2
Number of indexes: 3
This database has not yet been synchronized.
Synchronization by publication number:
  1. Publication test
      Number of rows in next upload: 0
      Last download: 1900-01-01 00:00:00.000
Done.
```

Set the `ml_remote_id` to NULL for a file called `sample.udb`:

```
ulinfo -c DBF=sample.udb -rc
```

UltraLite Initialize Database utility (ulinit)

Creates an UltraLite database from an existing SQL Anywhere database.

Syntax

```
ulinit -a "SAconnection-string" -c "ULconnection-string" -n pubname [ options ]
```

Option	Description
-a "SAconnection-string"	Required. Connect to the SQL Anywhere reference database specified in the <i>SAconnection-string</i> .
-c "ULconnection-string"	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	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 .
-e table, ...	Exclude the named <i>table</i> . The named table(s) are not created in the UltraLite database. You can name multiple tables in a comma-separated list. For example: <i>-e mydbtable1,mydbtable5</i>
-I	Use with -d . Include inserted rows in the next upload synchronization. By default, rows inserted by this utility are not uploaded during synchronization.
-l logfile	Log DDL database schema creation SQL statements, as executed, to <i>logfile</i> .
-n pubname	Required. Add tables to the UltraLite database schema. <i>pubname</i> specifies a publication in the reference database. Tables in the publication are added to the UltraLite database. 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* .
-o [extended-options]	Specify a semicolon separated list of UltraLite database creation parameters. See “Choosing database creation parameters for UltraLite” on page 34.
-p creator-ID	Required for Palm OS when the database is being installed to the record store. Create the database with the specified four character <i>creator-ID</i> of the UltraLite client application. If you are deploying the database to a VFS store, then do not use this option.

Option	Description
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-s <i>pubname</i>	<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 has no named publications.</p> <p>For more information about how to create publications for MobiLink synchronization, see “Publications in UltraLite” on page 137.</p>
-t <i>file</i>	Specify the file containing the trusted root certificate. This certificate is required for server authentication.
-w	Do not display warnings.

Remarks

The SQL Anywhere reference database 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.

If you want to create an UltraLite database without using a SQL Anywhere reference database, try one of the following methods:

- 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.
- If you want to create an empty UltraLite database that you can configure independent of any kind of reference database, use the `ulcreate` utility or the **Create Database Wizard** for UltraLite.

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 as part of your *extended-options* list.

To see a list of supported collations (and corresponding codepages), run `ulcreate -l` 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.

Palm databases written to the desktop must be identified with the *.pdb* extension. However, once you deploy the database to the device, the extension is dropped. For more information about file name formats, see [“Palm OS” on page 47](#).

You cannot deploy UltraLite databases with the Palm install tool if the target is a VFS volume. Instead, you must copy the database directly onto the media using a card reader, or some other tool.

See also

- [“Create an UltraLite database from a SQL Anywhere reference database” on page 31](#)
- [“Introduction to MobiLink models” \[*MobiLink - Getting Started*\]](#)
- [“UltraLite Create Database utility \(ulcreate\)” on page 255](#)
- [“UltraLite connection parameters” on page 225](#)

Examples

Create a file called *customer.udb* that contains the tables defined in TestPublication:

```
ulinit -a "DSN=dbdsn;UID=JimmyB;PWD=secret" -c DBF=customer.udb -n  
TestPublication
```

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=dbdsn;UID=JimmyB;PWD=secret" -c DBF=customer.udb -n Pub1 -n  
Pub2 -s Pub1 -s Pub2
```

Create a file called *customer.udb* for Palm OS using a registered creator ID:

```
ulinit -a "DSN=dbdsn;UID=JimmyB;PWD=secret" -c DBF=customer.udb.pdb -n  
TutCustomersPub -p creator-id
```

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
-a	Add data and schema definitions into an existing database. If you are adding data into a pre-existing record-based database for Palm OS (database with a <i>.pdb</i> extension), do not use the -p option.
-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.
-d	Load data only, ignoring any schema meta-data in the XML file input.
-f <i>directory</i>	Set the directory that contains files with additional data to load. See the -f option in ulunload (“ UltraLite Unload Database utility (ulunload) ” on page 281).
-g <i>ID</i>	Set the initial database ID to the INTEGER value you assign. This value is used by the database for all new rows that have global autoincrement columns. The database uses this base value to autoincrement IDs associated with each additional row and/or column. When deploying an application, you must assign a different identification number to each database for synchronization with the MobiLink server.
-I	Include inserted rows in the next upload synchronization. By default, rows inserted by this utility are not uploaded during synchronization.
-n	Load schema meta-data only, ignoring any data in the XML input file.
-o [<i>extended-options</i>]	Specify a semicolon separated list of UltraLite database creation parameters. See “ Choosing database creation parameters for UltraLite ” on page 34.
-oa	Cancel the process if it identifies that the database was created by a previous version of UltraLite (which would cause the database to be upgraded).
-ol	List the available database creation parameters and exit. See “ Choosing database creation parameters for UltraLite ” on page 34.

Option	Description
-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.
-or	Open database in read-only mode. UltraLite makes a copy of the original file, and then uses the copy to test your scripts without altering the database. Changes to the copied file are discarded upon completion. If you are connecting directly from the desktop to a database already deployed to a Windows Mobile device, this option is not supported.
-ou	Upgrade the database if it was created with an older UltraLite release.
-p creator-ID	Required for Palm OS when the database is being installed to the record store. Create the database with the specified four character <i>creator-ID</i> of the UltraLite client application. If you are deploying the database to a VFS store, then do not use this option. If you are appending data into a pre-existing record-based database for Palm OS (database with a <i>.pdb</i> extension), do not use this option with the <i>-a</i> option.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-s file	Log the SQL statements used to load the database into the specified <i>file</i> .
-t file	Specify the file containing the trusted root certificate. This certificate is required for server authentication.
-v	Print verbose messages.
-y	Overwrite the database file without confirmation. This only applies when you use ulload to create a new database.
<i>xml-file</i>	Specify the name of the XML file from which data is loaded.

Remarks

The ulload utility takes an input XML file generated by ulunload, ulunloadold, 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 meta-data for the schema and/or meta-data for the database data. **-d** ignores the schema meta-data, only adding data to the *.udb* file. **-n** ignores the data and the meta-data, 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.

Palm databases written to the desktop must be identified with the *.pdb* extension. However, once you deploy the database to the device, the extension is dropped. For more information about file name formats, see [“Palm OS” on page 47](#).

You cannot deploy UltraLite databases with the Palm install tool if the target is a VFS volume. Instead, you need to copy the database directly onto the media using a card reader, or some other tool.

See also

- [“Registering the Palm creator ID” \[UltraLite - C and C++ Programming\]](#)
- [“UltraLite connection parameters” on page 225](#)
- [“UltraLite Unload Database utility \(ulunload\)” on page 281](#)
- [“Supported exit codes” on page 248](#)
- [“UltraLite global_database_id option” on page 222](#)

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
```

Load XML from a file named *test_data.xml* into the copy UltraLite makes of the database named *sample.udb*. Discard those changes upon completion. This file allows you to check for errors in the XML data and correct them. When data loads successfully, you can run the command without the *-or* option to keep the XML updates:

```
ulload -or -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 "*connection-string*" [*options*] [*synchronization parameters* { *REPLACE* | *MERGE* } *profile-string*]

Option	Description
-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.
-oa	Cancel the process if it identifies that the database was created by a previous version of UltraLite (which would cause the database to be upgraded).
-or	Synchronize the database in read-only mode. UltraLite makes a copy of the original file, and then uses the copy to test your scripts without altering the database. Changes to the copied file are discarded upon completion. If you are connecting directly from the desktop to a database already deployed to a Windows Mobile device, the parameter is not supported.
-ou	Upgrade the database if it was created with an older UltraLite release.
-p	Synchronize using the named sync profile, equivalent to: <code>synchronize profileName merge 'syncOptions'</code> where sync options are taken from the trailing ulsync options. For example: <code>ulsync -p profileName "MobiLinkUid=ml;ScriptVersion=Version001...syncOptions"</code> See “Synchronization profile options” on page 278 .
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-r	Display last synchronization results and exit.
-s	Run available SQL passthrough scripts after syncing.
-v	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 “UltraLite CREATE SYNCHRONIZATION PROFILE statement” on page 464 .

Remarks

The following options that were valid for versions 10 and preceding 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** have been folded into 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`, `ulload`, or `ulcreate`.
- Referencing an external certificate file at synchronization time with the **trusted_certificates=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 278](#)
- [“UltraLite connection parameters” on page 225](#)
- [“UltraLite clients” on page 125](#)
- [“Supported exit codes” on page 248](#)
- [“MobiLink file transfer utility \(mlfiletransfer\)” \[MobiLink - Client Administration\]](#)
- [“Deploy UltraLite with TLS-enabled synchronization” on page 60](#)

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_certificates=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_certificates=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
11\Samples\UltraLite\SyncEncrypt\custdb.udb"
MobiLinkUid=50;ScriptVersion=custdb 11.0;Stream=tcpip{port=2439}
```

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
Allow-Download-DupRows	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 synchronisation to fail. The default value is "no." See “Additional Parameters synchronization parameter” on page 155
Auth-Parms	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 156 .
Check-pointStore	Boolean	Adds additional checkpoints of the database during synchronization to limit database growth during the synchronization process. See “Additional Parameters synchronization parameter” on page 155 .
Continue-Download	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 “Resuming failed downloads” [MobiLink - Server Administration] .
Disable-Concurrency	Boolean	Disallow database access from other threads during synchronization. See “Additional Parameters synchronization parameter” on page 155 .
DownloadOnly	Boolean	Performs a download-only synchronization. See “Download Only synchronization parameter” on page 159 .
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 161 .
MobiLinkPwd	String	Specifies the existing MobiLink password associated with the user name. See “MobiLinkPwd (mp) extended option” [MobiLink - Client Administration] .
MobiLinkUid	String	Specifies the MobiLink user name. See “-u option” [MobiLink - Client Administration] . See “-mn option” [MobiLink - Client Administration] .

Synchronization profile option	Valid values	Description
NewMobi-LinkPwd	String	Supplies a new password for the MobiLink user. Use this option when you want to change an existing password. See “-mn option” [MobiLink - Client Administration] .
Ping	Boolean	Confirms communications with the server only; no synchronization is performed. See “Ping synchronization parameter” on page 165 .
Publications	String (comma separated)	Specifies the publications(s) to synchronize. The publications determine the tables on the remote 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 “Publications in UltraLite” on page 137 .
Script-Version	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” [MobiLink - Client Administration] .
SendColumn-Names	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 168 .
Send-DownloadACK	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 169 .
Stream	String (with sublist)	Specifies the MobiLink network synchronization protocol. See “Stream Type synchronization parameter” on page 171 .
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 155 .
UploadOnly	String	Specifies that synchronization will only include an upload, and no download will occur. See “Upload Only synchronization parameter” on page 174 .

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

- [“UltraLite ALTER SYNCHRONIZATION PROFILE statement” on page 452](#)
- [“UltraLite DROP SYNCHRONIZATION PROFILE statement” on page 476](#)
- [“UltraLite SYNCHRONIZE statement” on page 491](#)
- [“UltraLite creation parameters” on page 181](#)

UltraLite Unload Database 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
-b <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>connec-tion-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.
-d	Only unload the data from the database to the output file. Do not unload any schema information.
-e <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>
-f <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.
-n	Unload schema only, ignoring any data in the database.
-oa	Cancel the process if it identifies that the database was created by a previous version of UltraLite (which would cause the database to be upgraded).
-or	Open the database in read-only mode. UltraLite makes a copy of the original file and uses that to unload. This ensures that databases created with earlier versions of the software are not upgraded. If you are connecting directly from the desktop to a database already deployed to a Windows Mobile device, the parameter is not supported.
-ou	Upgrade the database if it was created with an older UltraLite release.
-q	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
-s	Unload as SQL Anywhere-compatible SQL statements. SQL file output can be read by UltraLite or SQL Anywhere using DBISQL.

Option	Description
-t <i>table</i> ,...	Unload data in the named <i>table</i> only. You can name multiple tables in a comma separated list. For example: <code>-t mydbtable2,mydbtable6</code>
-v	Print verbose messages.
-x <i>owner</i>	Output tables so they are owned by a specific user ID. You can use this option with the <code>-s</code> option.
-y	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 <code>-s</code> 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

- [“UltraLite connection parameters” on page 225](#)
- [“Supported exit codes” on page 248](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“UltraLite Information utility \(ulinfo\)” on page 265](#)

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 sample.sql
```

UltraLite Unload Old Database utility (ulunloadold)

Unloads UltraLite version 8.0.2 to 9.0.x databases and/or schema files (*.usm) into an XML file.

Syntax

ulunloadold -c "*connection-string*" [*options*] *xml-file*

Option	Description
-b <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>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.
-f <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 XML file.
-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.
-y	Overwrite <i>xml-file</i> without confirmation.
<i>xml-file</i>	Set the name of the XML file that data will be unloaded into.

Remarks

UltraLite version 11 cannot directly upgrade version 8.x or version 9.x databases. Use this tool to generate an XML file that can then be used by ulload to create a version 11 database. Do not unload UltraLite version 11 databases with this utility. Use the ulunload utility instead.

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 name follows this naming convention:

tablename-columnname-rownumber.bin

Any error code value other than 0 means that the operation failed.

This utility cannot be used to unload databases directly on Windows Mobile devices. You must first copy them to a desktop computer.

See also

- [“UltraLite connection parameters” on page 225](#)
- [“UltraLite Load XML to Database utility \(ulload\)” on page 272](#)
- [“UltraLite Unload Database utility \(ulunload\)” on page 281](#)
- [“UltraLite Information utility \(ulinfo\)” on page 265](#)

Example

Upgrading an UltraLite 8.0.x schema file named *dbschema8.usm* into an UltraLite version 11 database named *db.udb* requires these two commands:

```
ulunloadold -c SCHEMA_FILE=dbschema8.usm dbschema.xml
ulload -c DBF=db.udb dbschema.xml
```

Upgrading an UltraLite version 9.0.x database for Palm OS named *palm9db.pdb* to an UltraLite version 11 database named *palm11db.pdb* requires these two commands:

```
ulunloadold -c DBF=palm9db.pdb dbdata.xml
ulload -c DBF=palm11db.pdb -p Sybl dbdata.xml
```

UltraLite Validate Database utility (ulvalid)

Performs a full ("normal") validation of an UltraLite database, including:

- database pages - validate all database pages, using checksums when enabled (see the `checksum_level` database creation parameter). Note that certain critical pages always have checksums and even pages without checksums undergo a basic validity check.
- table - validate table(s) by checking that the table row count matches the count in each index
- index - validate indexes by checking that entries refer to valid rows. `ulvalid -e` performs an express check, which includes only table validation.

Syntax

`ulvalid -c "connection-string" [options]`

Option	Description
<code>-c "connection-string"</code>	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.
<code>-e</code>	Express validation. Only perform table validation. This option provides a faster validation than normal validation
<code>-oa</code>	Cancel the process if it identifies that the database was created by a previous version of UltraLite (which would cause the database to be upgraded).
<code>-or</code>	Open the database in read-only mode. UltraLite makes a copy of the original file, which you can then use to test your scripts without altering the database. Changes to the copied file are discarded upon completion. If you are connecting directly from the desktop to a database already deployed to a Windows Mobile device, the parameter is not supported.
<code>-ou</code>	Upgrade the database if it was created with an older UltraLite release.
<code>-q</code>	Set the utility to run in quiet mode. Suppress informational banners, version numbers, and status messages. Error messages are still displayed, however.
<code>-v</code>	Print verbose messages.

Remarks

Validating a database verifies the accuracy of the table meta-data and ensures the file has not been corrupted.

See also

- [“Validate an UltraLite database” on page 15](#)

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

Contents

View or hide UltraLite system tables	290
systable system table	291
syscolumn system table	292
sysindex system table	293
sysixcol system table	295
syspublication system table	296
sysarticle system table	297
sysuldata system table	298

The schema of every UltraLite database is described in several system tables. Because UltraLite does not support table ownership, these system tables are accessible to any UltraLite user.

The contents of system tables can be changed only by UltraLite itself. UPDATE, DELETE, and INSERT commands cannot be used to modify the contents of these tables. Furthermore, the structure of these tables cannot be changed using the ALTER TABLE and DROP statements.

View or hide UltraLite system tables

To hide or show system objects (Sybase Central)

1. Connect to the database.
2. Browse the objects of the database.
3. Right-click the contents pane and choose **Hide/Show System Objects**.
4. Click **Tables** and browse the available tables.

systable system table

Each row in the systable system table describes one table in the database.

Column name	Column type	Description
column_count	UNSIGNED INT	The number of columns in the table.
index_count	UNSIGNED INT	The number of indexes in the table.
indexcol_count	UNSIGNED INT	The total number of columns in all indexes in the table.
map_handle	UNSIGNED INT	Internal use only.
table_name	VARCHAR(128)	The name of the table.
object_id	UNSIGNED INT	A unique identifier for that table.
sync_type	VARCHAR(128)	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_name	VARCHAR(128)	The name of the table.
table_type	TINYINT	One of either sys for system tables or user for regular tables.
tpd_handle	UNSIGNED INT	Internal user only.

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)	A unique identifier 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	SMALLINT	Used with a variable sized domain.
nulls	CHAR(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	UNSIGNED INT	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	UNSIGNED INT	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	SMALLINT (10)	The type of index. Can be one of: <ul style="list-style-type: none"> ● primary ● foreign ● key ● unique ● index
hash_size	SHORT	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 295](#)

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	CHAR(1)	Indicates whether the column in the index is kept in ascending (A) or descending (D) order.
sequence	SMALLINT	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 293](#)

syspublication system table

Each row in the syspublication system table describes a publication.

Column name	Column type	Description
download_timestamp	UNSIGNED INT	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	CHAR(128)	The name of the publication.

Constraints

PRIMARY KEY (publication_id)

See also

- [“sysarticle system table” on page 297](#)

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	TINY INT	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 296](#)

sysuldata system table

Each row in the sysuldata system table names value pairs of options and properties.

Column name	Column type	Description
long_setting	LONGBINAR Y	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 sys for internals, opt for options, or prop for properties

Constraints

PRIMARY KEY (name, type)

See also

- [“UltraLite database properties” on page 213](#)

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	301
UltraLite SQL functions	347
UltraLite SQL statements	447

UltraLite SQL elements

Contents

Keywords in UltraLite	302
Identifiers in UltraLite	303
Strings in UltraLite	304
Comments in UltraLite	305
Numbers in UltraLite	306
The NULL value in UltraLite	307
Special values in UltraLite	308
Dates and times in UltraLite	311
Data types in UltraLite	312
Expressions in UltraLite	325
Operators in UltraLite	338
Variables in UltraLite	341
Execution plans in UltraLite	342

Keywords in UltraLite

Each SQL statement contains one or more keywords. SQL keywords are case insensitive, but throughout these manuals, keywords are indicated in uppercase. Some keywords cannot be used as identifiers without surrounding them in double quotes. These are called **reserved words**. See [“Reserved words” \[SQL Anywhere Server - SQL Reference\]](#).

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.

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. See [“Reserved words” \[SQL Anywhere Server - SQL Reference\]](#).
- 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.

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 36](#)

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 “NULL 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 325](#)
- [“GETDATE function \[Date and time\]” on page 385](#)
- [“NOW function \[Date and time\]” on page 408](#)

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 325](#)
- [“GETDATE function \[Date and time\]” on page 385](#)
- [“NOW function \[Date and time\]” on page 408](#)

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 308](#)
- [“Expressions in UltraLite” on page 325](#)
- [“NOW function \[Date and time\]” on page 408](#)
- [“GETDATE function \[Date and time\]” on page 385](#)
- [“NOW function \[Date and time\]” on page 408](#)

SQLCODE special value

Current SQLCODE value at the time the special value was evaluated.

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 325](#)
- [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 [“Date parts” on page 349](#).

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 339.

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].
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 <i>select-list</i> of a query. See “ LONG VARCHAR data type ” [SQL Anywhere Server - SQL Reference]. You can cast strings to/from LONGVARCHAR data.

Data type	Description
[UNSIGNED] BIGINT	An integer requiring 8 bytes of storage. See “BIGINT data type” [SQL Anywhere Server - SQL Reference].
{ DECIMAL DEC NUMERIC } (precision , scale)	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 201, and “UltraLite scale creation parameter” on page 203.
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 [(precision)]	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].
DATETIME	Identical to TIMESTAMP. See “DATETIME data type” [SQL Anywhere Server - SQL Reference].
TIMESTAMP	A point in time, containing year, month, day, hour, minute, second, and fraction of a second. See “TIMESTAMP data type” [SQL Anywhere Server - SQL Reference].
VARBINARY (max-length)	Identical to BINARY. See “VARBINARY data type” [SQL Anywhere Server - SQL Reference].

Data type	Description
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 ” [SQL Anywhere Server - SQL Reference].
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 <i>select-list</i> of a query. See “ LONG BINARY data type ” [SQL Anywhere Server - SQL Reference]. 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 ” [SQL Anywhere Server - SQL Reference].

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
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

In most cases, self-casting 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.

In some cases, however, 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 **SQL_CONVERT_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 **SQL_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 LONG INT SIGNED LONG 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
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	
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	

From	Always	Never	Value-dependent
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
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

From	Always	Never	Value-dependent
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
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

From	Always	Never	Value-dependent
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
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	
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	
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 326
- “Special values in UltraLite” on page 308
- “Column names in expressions” on page 326
- “UltraLite SQL functions ” on page 347
- “Subqueries in expressions” on page 329
- “Search conditions in UltraLite” on page 331
- “Data types in UltraLite” on page 312
- “CASE expressions” on page 327
- “Input parameters” on page 330

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

- [“UltraLite FROM clause” on page 478](#)

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 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 with references to 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

- [“UltraLite SELECT statement” on page 486](#)
- [“Using subqueries” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Search conditions in UltraLite” on page 331](#)

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

- [“Using host variables” \[UltraLite - C and C++ Programming\]](#)
- [“Preparing statements” \[SQL Anywhere Server - Programming\]](#)
- UltraLite C/C++: [“Data manipulation: Insert, Delete, and Update” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“Data manipulation: INSERT, UPDATE, and DELETE” \[UltraLite - .NET Programming\]](#)
- UltraLite for M-Business: [“Data manipulation: INSERT, UPDATE, and DELETE” \[UltraLite - M-Business Anywhere Programming\]](#)

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

Specifies a search condition 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 compare expression
| expression IS [ NOT ] { NULL | TRUE | FALSE | UNKNOWN }
| expression [ NOT ] BETWEEN expression AND expression
| expression [ NOT ] IN ( expression, ... )
| expression [ NOT ] IN ( subquery )
| expression [ NOT ] { ANY | ALL } ( subquery )
| expression [ NOT ] EXISTS ( subquery )
| NOT search-condition
| search-condition AND search-condition
| search-condition OR search-condition
| ( search-condition )
```

Parameters

```
compare:
= | > | < | >= | <= | <> | != | !< | !>
```

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.

The different types of search conditions supported by UltraLite include:

- ALL conditions
- ANY conditions
- BETWEEN conditions
- EXISTS conditions
- IN conditions

These conditions are discussed in separate sections that follow.

Note

Subqueries form an important class of expression that is used in many search conditions.

See also

- [“Comparison operators” on page 332](#)
- [“Three-valued logic” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Subqueries in expressions” on page 329](#)

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
[NOT] LIKE	a text comparison, possibly using regular expressions
>	greater than
<	less than
>=	greater than or equal to
<=	less than or equal to
!=	not equal to
<>	not equal to
!>	not greater than

Operator	Interpretation
!<	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 333](#)
- [“Search conditions in UltraLite” on page 331](#)

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 **OR** *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 332](#)
- [“Search conditions in UltraLite” on page 331](#)

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 conditions

Use the ALL condition in conjunction with a 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 332

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 conditions

Use the ANY condition in conjunction with a 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 332

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 conditions

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 cheaper than \$10 or more expensive than \$15.

```
SELECT Name, UnitPrice
FROM Products
WHERE UnitPrice NOT BETWEEN 10 AND 15;
```

EXISTS conditions

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 conditions

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');
```

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 340](#)
- [“Comparison operators” on page 332](#)
- [“Arithmetic operators” on page 338](#)
- [“String operators” on page 339](#)
- [“Bitwise operators” on page 339](#)
- [“Logical operators” on page 333](#)

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.

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.

Performance 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 [“UltraLite SELECT statement” on page 486](#).

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. See [“UltraLite temporary files” on page 12](#).
- 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.

View 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.

To view an execution plan in the Plan Viewer

1. Choose **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 343

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.

Operation	Description
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.
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.

Operation	Description
temp [<i>plan</i>]	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. 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.
union-all [<i>plan</i> , ..., <i>plan</i>]	Performs a UNION ALL operation on the rows generated in the underlying plan.

UltraLite SQL functions

Contents

Function types	348
SQL functions (A-D)	354
SQL functions (E-O)	383
SQL functions (P-Z)	411

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. For a complete list of syntax conventions, see [“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.

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 \[Aggregate\]” on page 359](#)
- [“COUNT function \[Aggregate\]” on page 370](#)
- [“LIST function \[Aggregate\]” on page 395](#)
- [“MAX function \[Aggregate\]” on page 400](#)
- [“MIN function \[Aggregate\]” on page 401](#)
- [“SUM function \[Aggregate\]” on page 432](#)

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 \[Data type conversion\]” on page 361](#)
- [“CONVERT function \[Data type conversion\]” on page 366](#)
- [“HEXTOINT function \[Data type conversion\]” on page 386](#)
- [“INTTOHEX function \[Data type conversion\]” on page 390](#)
- [“ISDATE function \[Data type conversion\]” on page 391](#)

UltraLite date and time functions

Date and time functions perform operations on date and time data types or return date or time information.

In this chapter, the term **datetime** is used to mean date or time or timestamp. The specific data type DATETIME is indicated as DATETIME.

For more information about datetime data types, see [“Data types in UltraLite” on page 312](#).

List of functions

The following date and time functions are available:

- “DATE function [Date and time]” on page 372
- “DATEADD function [Date and time]” on page 372
- “DATEDIFF function [Date and time]” on page 373
- “DATEFORMAT function [Date and time]” on page 375
- “DATENAME function [Date and time]” on page 375
- “DATEPART function [Date and time]” on page 376
- “DATETIME function [Date and time]” on page 377
- “DAY function [Date and time]” on page 377
- “DAYNAME function [Date and time]” on page 378
- “DAYS function [Date and time]” on page 379
- “DOW function [Date and time]” on page 382
- “GETDATE function [Date and time]” on page 385
- “HOUR function [Date and time]” on page 387
- “HOURS function [Date and time]” on page 388
- “MINUTE function [Date and time]” on page 401
- “MINUTES function [Date and time]” on page 402
- “MONTH function [Date and time]” on page 405
- “MONTHNAME function [Date and time]” on page 406
- “MONTHS function [Date and time]” on page 406
- “NOW function [Date and time]” on page 408
- “QUARTER function [Date and time]” on page 413
- “SECOND function [Date and time]” on page 420
- “SECONDS function [Date and time]” on page 421
- “TODAY function [Date and time]” on page 435
- “WEEKS function [Date and time]” on page 442
- “YEAR function [Date and time]” on page 444
- “YEARS function [Date and time]” on page 444
- “YMD function [Date and time]” on page 445

Date parts

Many of the date functions use dates built from **date parts**. The following table displays allowed values of date parts.

Date part	Abbreviation	Values
Year	yy	1-9999
Quarter	qq	1-4

Date part	Abbreviation	Values
Month	mm	1-12
Week	wk	1-54. Weeks begin on Sunday.
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
Calyearofweek	cyr	Integer. 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.
Calweekofyear	cwk	1-54. The week number within the year that contains the specified date.
Caldayofweek	cdw	1-7. (Monday = 1, ..., Sunday = 7)

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 \[Miscellaneous\]” on page 355](#)
- [“COALESCE function \[Miscellaneous\]” on page 366](#)
- [“EXPLANATION function \[Miscellaneous\]” on page 383](#)
- [“GREATER function \[Miscellaneous\]” on page 385](#)
- [“IFNULL function \[Miscellaneous\]” on page 389](#)
- [“ISNULL function \[Miscellaneous\]” on page 392](#)
- [“LESSER function \[Miscellaneous\]” on page 395](#)
- [“NEWID function \[Miscellaneous\]” on page 407](#)
- [“NULLIF function \[Miscellaneous\]” on page 409](#)

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 \[Numeric\]” on page 354](#)
- [“ACOS function \[Numeric\]” on page 354](#)
- [“ASIN function \[Numeric\]” on page 357](#)
- [“ATAN function \[Numeric\]” on page 357](#)
- [“ATAN2 function \[Numeric\]” on page 358](#)
- [“CEILING function \[Numeric\]” on page 362](#)
- [“COS function \[Numeric\]” on page 369](#)
- [“COT function \[Numeric\]” on page 369](#)
- [“DEGREES function \[Numeric\]” on page 381](#)
- [“EXP function \[Numeric\]” on page 383](#)
- [“FLOOR function \[Numeric\]” on page 384](#)
- [“LOG function \[Numeric\]” on page 397](#)
- [“LOG10 function \[Numeric\]” on page 398](#)
- [“MOD function \[Numeric\]” on page 404](#)
- [“PI function \[Numeric\]” on page 412](#)
- [“POWER function \[Numeric\]” on page 412](#)
- [“RADIANS function \[Numeric\]” on page 414](#)
- [“REMAINDER function \[Numeric\]” on page 414](#)
- [“ROUND function \[Numeric\]” on page 419](#)
- [“SIGN function \[Numeric\]” on page 423](#)
- [“SIN function \[Numeric\]” on page 424](#)
- [“SQRT function \[Numeric\]” on page 427](#)
- [“TAN function \[Numeric\]” on page 435](#)
- [“TRUNCNUM function \[Numeric\]” on page 437](#)

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]” on page 356
- “BYTE_LENGTH function [String]” on page 360
- “BYTE_SUBSTR function [String]” on page 360
- “CHAR function [String]” on page 363
- “CHARINDEX function [String]” on page 364
- “CHAR_LENGTH function [String]” on page 365
- “DIFFERENCE function [String]” on page 381
- “INSERTSTR function [String]” on page 390
- “LCASE function [String]” on page 392
- “LEFT function [String]” on page 393
- “LENGTH function [String]” on page 394
- “LOCATE function [String]” on page 396
- “LOWER function [String]” on page 398
- “LTRIM function [String]” on page 399
- “PATINDEX function [String]” on page 411
- “REPEAT function [String]” on page 415
- “REPLACE function [String]” on page 416
- “REPLICATE function [String]” on page 417
- “RIGHT function [String]” on page 418
- “RTRIM function [String]” on page 419
- “SIMILAR function [String]” on page 424
- “SOUNDEX function [String]” on page 425
- “SPACE function [String]” on page 426
- “STR function [String]” on page 427
- “STRING function [String]” on page 428
- “STRTOUUID function [String]” on page 429
- “STUFF function [String]” on page 430
- “SUBSTRING function [String]” on page 430
- “TRIM function [String]” on page 436
- “UCASE function [String]” on page 438
- “UPPER function [String]” on page 439
- “UUIDTOSTR function [String]” on page 441

UltraLite system functions

System functions return system information.

List of functions

The following system functions are available in UltraLite:

- [“DB_PROPERTY function \[System\]” on page 380](#)
- [“ML_GET_SERVER_NOTIFICATION \[System\]” on page 403](#)
- [“SYNC_PROFILE_OPTION_VALUE function \[System\]” on page 434](#)

SQL functions (A-D)

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 348](#).

See also

- [“SQL functions \(E-O\)” \[SQL Anywhere Server - SQL Reference\]](#)
- [“SQL functions \(P-Z\)” \[SQL Anywhere Server - SQL Reference\]](#)

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/2003** SQL foundation feature outside core SQL.

Example

The following statement returns the value 66.

```
SELECT ABS( -66 );
```

ACOS function [Numeric]

Returns the arc-cosine, in radians, of a numeric expression.

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.

See also

- [“ASIN function \[Numeric\]” on page 357](#)
- [“ATAN function \[Numeric\]” on page 357](#)
- [“ATAN2 function \[Numeric\]” on page 358](#)
- [“COS function \[Numeric\]” on page 369](#)

Standards and compatibility

- **SQL/2003** 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/2003** 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

- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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.

See also

- [“ACOS function \[Numeric\]” on page 354](#)
- [“ATAN function \[Numeric\]” on page 357](#)
- [“ATAN2 function \[Numeric\]” on page 358](#)
- [“SIN function \[Numeric\]” on page 424](#)

Standards and compatibility

- **SQL/2003** 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.

Syntax

ATAN(*numeric-expression*)

Remarks

The ATAN and TAN functions are inverse operations.

Parameters

- **numeric-expression** The tangent of the angle.

Returns

DOUBLE

Remarks

This function converts its argument to DOUBLE, and performs the computation in double-precision floating point.

See also

- [“ACOS function \[Numeric\]” on page 354](#)
- [“ASIN function \[Numeric\]” on page 357](#)
- [“ATAN2 function \[Numeric\]” on page 358](#)
- [“TAN function \[Numeric\]” on page 435](#)

Standards and compatibility

- **SQL/2003** 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.

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.

See also

- [“ACOS function \[Numeric\]” on page 354](#)
- [“ASIN function \[Numeric\]” on page 357](#)
- [“ATAN function \[Numeric\]” on page 357](#)
- [“TAN function \[Numeric\]” on page 435](#)

Standards and compatibility

- **SQL/2003** 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( numeric-expression | DISTINCT numeric-expression )
```

Parameters

- **numeric-expression** The expression whose average is calculated over a set of rows.
- **DISTINCT clause** Computes the average of the unique numeric values in the input.

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.

See also

- [“SUM function \[Aggregate\]” on page 432](#)
- [“COUNT function \[Aggregate\]” on page 370](#)

Standards and compatibility

- **SQL/2003** Core feature.

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 365](#)
- [“DATALENGTH function \[System\]” on page 371](#)
- [“LENGTH function \[String\]” on page 394](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

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

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.

BINARY

VARCHAR

NVARCHAR

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 430](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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 on 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.
- **data type** The target data type.

Returns

The specified data type.

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

- [“CONVERT function \[Data type conversion\]” on page 366](#)
- [“LEFT function \[String\]” on page 393](#)

Standards and compatibility

- **SQL/2003** Core feature.

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(*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.

See also

- [“FLOOR function \[Numeric\]” on page 384](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

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 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value Y.

```
SELECT CHAR( 89 );
```

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.

See also

- [“SUBSTRING function \[String\]” on page 430](#)
- [“REPLACE function \[String\]” on page 416](#)
- [“LOCATE function \[String\]” on page 396](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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
Kelly	Moira

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 360](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Core feature.

Example

The following statement returns the value 8.

```
SELECT CHAR_LENGTH( 'Chemical' );
```

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

ANY

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 392](#).

See also

- [“ISNULL function \[Miscellaneous\]” on page 392](#)

Standards and compatibility

- **SQL/2003** 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 on 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 outputted 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]yymmdd
-	13 or 113	dd Mmm yyyy hh:nn:ss:sss (24 hour clock, Europe default + milliseconds, 4-digit year)
-	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)

Without century (yy) style codes	With century (yyyy) style codes	Output format
-	21 or 121	yyyy-mm-dd hh:nn:ss.sss (24 hour clock, ODBC canonical with milliseconds, 4-digit year)

Returns

The data type specified.

Remarks

If no *format-style* argument is provided, style code 0 is used.

For a description of the styles produced by each output symbol (such as Mmm), see [“UltraLite date_format creation parameter” on page 187](#).

See also

- [“CAST function \[Data type conversion\]” on page 361](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

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
25.03.2000
...

```
SELECT CONVERT( CHAR( 20 ), OrderDate, 7 ) FROM SalesOrders;
```

OrderDate
Mar 16, 00
Mar 20, 00

OrderDate
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, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

See also

- [“ACOS function \[Numeric\]” on page 354](#)
- [“COT function \[Numeric\]” on page 369](#)
- [“SIN function \[Numeric\]” on page 424](#)
- [“TAN function \[Numeric\]” on page 435](#)

Standards and compatibility

- **SQL/2003** 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, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

See also

- [“COS function \[Numeric\]” on page 369](#)
- [“SIN function \[Numeric\]” on page 424](#)
- [“TAN function \[Numeric\]” on page 435](#)

Standards and compatibility

- **SQL/2003** 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(  
*  
| expression  
| DISTINCT expression  
)
```

Parameters

- * Return the number of rows in each group.
- **expression** The expression for which to return the number of rows.
- **DISTINCT expression** The expression for which to return the number of distinct rows.

Returns

INT

Remarks

Rows where the value is the NULL value are not included in the count.

See also

- “AVG function [Aggregate]” on page 359
- “SUM function [Aggregate]” on page 432

Standards and compatibility

- **SQL/2003** Core feature. Syntax 2 is feature T611.

Example

The following statement returns each unique city, and the number of rows with that city value.

```
SELECT City, COUNT( * ) FROM Employees GROUP BY City;
```

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/2003** 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 190](#).

Syntax

```
DATE( expression )
```

Returns

DATE

Parameters

- **expression** The value to be converted to date format, typically a string.

Standards and compatibility

- **SQL/2003** 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 the date produced by adding multiple date parts to a date.

Syntax

```
DATEADD( date-part, numeric-expression, date-expression )
```

```
date-part :
year
| quarter
| month
| week
| day
```

```

| dayofyear
| hour
| minute
| second
| millisecond

```

Parameters

- **date-part** The date part to be added to the date. For more information about date parts, see “[Date parts](#)” on page 349.
- **numeric-expression** The number of date parts to be added to the date. The *numeric-expression* can be any numeric type, but the value is truncated to an integer.
- **date-expression** The date to be modified.

Returns

TIMESTAMP

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value: 1995-11-02 00:00:00.000.

```
SELECT DATEADD( month, 102, '1987/05/02' );
```

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

```

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 [“Date parts” on page 349](#).

- **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

Remarks

This function calculates the number of date parts between two specified dates. The result is a signed integer value equal to (date2 - date1), 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.

For the smaller time units there are overflow values:

- **milliseconds** 24 days
- **seconds** 68 years
- **minutes** 4083 years
- **others** No overflow limit

The function returns an overflow error if you exceed these limits.

Standards and compatibility

- **SQL/2003** Transact-SQL 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 187](#).

Returns

VARCHAR

Remarks

Any allowable date format can be used for the string-expression.

Standards and compatibility

- **SQL/2003** 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 datetime value, as a character string.

Syntax

DATENAME(*date-part*, *date-expression*)

Parameters

- **date-part** The date part to be named.
For a complete listing of allowed date parts, see [“Date parts” on page 349](#).
- **date-expression** The date for which the date part name is to be returned. The date must 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.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value May.

```
SELECT DATENAME( month, '1987/05/02' );
```

DATEPART function [Date and time]

Returns the value of part of a datetime value.

Syntax

DATEPART(*date-part*, *date-expression*)

Parameters

- **date-part** The date part to be returned.
For a complete listing of allowed date parts, see [“Date parts” on page 349](#).
- **date-expression** The date for which the part is to be returned.

Returns

INT

Remarks

The date must contain the *date-part* field.

Standards and compatibility

- **SQL/2003** 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.

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.

Standards and compatibility

- **SQL/2003** 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 an integer from 1 to 31.

Syntax

DAY(*date-expression*)

Parameters

- **date-expression** The date.

Returns

SMALLINT

Remarks

The integers 1 to 31 correspond to the day of the month in a date.

Standards and compatibility

- **SQL/2003** 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/2003** Vendor extension.

Example

The following statement returns the value Saturday.

```
SELECT DAYNAME ( '1987/05/02' );
```

DAYS function [Date and time]

A function that evaluates days. For specific details, see this function's usage.

Syntax 1: integer

DAYS([*datetime-expression*,] *datetime-expression*)

Syntax 2: timestamp

DAYS(*datetime-expression*, *integer-expression*)

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of days to be added to the *datetime-expression*. If the *integer-expression* is negative, the appropriate number of days is subtracted from the timestamp. If you supply an integer expression, the *datetime-expression* must be explicitly cast as a date or timestamp.

For information about casting data types, see [“CAST function \[Data type conversion\]” on page 361](#).

Returns

INT when you specify two datetime expressions.

TIMESTAMP when the second argument you specify is an integer.

Remarks

The behavior of this function varies depending on the information you specify:

- If you give a single date, this function returns the number of days since 0000-02-29.

Note

0000-02-29 is not meant to imply an actual date; it is the date used by the date algorithm.

- If you give two dates, this function returns the integer number of days between them. Instead, use the DATEDIFF function.
- If you give a date and an integer, this function adds the integer number of days to the specified date. Instead, use the DATEADD function.

This function ignores hours, minutes, and seconds.

See also

- [“DATEDIFF function \[Date and time\]” on page 373](#)
- [“DATEADD function \[Date and time\]” on page 372](#)

Standards and compatibility

- **SQL/2003** 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 is 366 days prior to 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 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.

Syntax

```
DB_PROPERTY( property-name )
```

Parameters

- **property-name** The database property name.

Returns

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.

See also

- [“UltraLite SET OPTION statement” on page 488](#)
- UltraLite C++: [“SetDatabaseOption function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite.NET: [“SetDatabaseOption method” \[UltraLite - .NET Programming\]](#)

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

Returns the degrees of the angle given by *numeric-expression*.

DOUBLE

Remarks

This function converts its argument to DOUBLE, AND performs the computation in double-precision floating point. If the parameter is NULL, the result is NULL.

Standards and compatibility

- **SQL/2003** 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 425](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 3.

```
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 date to evaluate.

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/2003** Vendor extension.

Example

The following statement returns the value 5.

```
SELECT DOW( '1998-07-09' );
```

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;
```

SQL functions (E-O)

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 348](#).

See also

- [“SQL functions \(A-D\)” \[SQL Anywhere Server - SQL Reference\]](#)
- [“SQL functions \(P-Z\)” \[SQL Anywhere Server - SQL Reference\]](#)

EXP function [Numeric]

Returns the exponential function, e to the power of a number.

Syntax

EXP(*numeric-expression*)

Parameters

- **numeric-expression** The exponent.

Returns

DOUBLE

Remarks

The EXP function returns the exponential of the value specified by *numeric-expression*.

This function converts its argument to DOUBLE, performs the computation in double-precision floating point, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The statement returns the value 3269017.3724721107.

```
SELECT EXP( 15 );
```

EXPLANATION function [Miscellaneous]

Returns the plan optimization strategy of a SQL statement.

Syntax

EXPLANATION(*string-expression*)

Parameters

- **string-expression** The SQL statement, which is commonly a SELECT statement, but can also be an UPDATE or DELETE statement.

Returns

LONG VARCHAR

Remarks

The optimization is returned as a string.

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 342](#)

Standards and compatibility

- **SQL/2003** 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 a number.

Syntax

FLOOR(*numeric-expression*)

Parameters

- **numeric-expression** The value, usually a FLOAT.

Returns

DOUBLE

Remarks

This function converts its arguments to DOUBLE, and performs the computation in double-precision floating point.

See also

- [“CEILING function \[Numeric\]” on page 362](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statements returns a Floor value of 123

```
SELECT FLOOR (123);
```

The following statements returns a Floor value of 123

```
SELECT FLOOR (123.45);
```

The following statements returns a Floor 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 408](#)

Standards and compatibility

- **SQL/2003** 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

ANY

Remarks

If the parameters are equal, the first is returned.

See also

- [“LESSER function \[Miscellaneous\]” on page 395](#)

Standards and compatibility

- **SQL/2003** 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 on 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 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( '0x008000001' );
```

INT

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 390

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 420.

```
SELECT HEXTOINT( '1A4' );
```

HOURL function [Date and time]

Returns the hour component of a datetime.

Syntax

```
HOURL( datetime-expression )
```

Parameters

- **datetime-expression** The datetime.

Returns

SMALLINT

Remarks

The value returned is a number from 0 to 23 corresponding to the datetime hour.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 21:

```
SELECT HOUR( '1998-07-09 21:12:13' );
```

HOURS function [Date and time]

A function that evaluates hours. For specific details, see this function's usage.

Syntax 1: integer

```
HOURS ( [ datetime-expression, ] datetime-expression )
```

Syntax 2: timestamp

```
HOURS ( datetime-expression, integer-expression )
```

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of hours to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of hours is subtracted from the datetime. If you supply an integer expression, the *datetime-expression* must be explicitly cast as a DATETIME data type.

For information about casting data types, see “[CAST function \[Data type conversion\]](#)” on page 361.

Returns

INT

TIMESTAMP

Remarks

The information you specify changes the behavior of this function:

- If you give a single date, this function returns the number of hours since 0000-02-29.

Note

0000-02-29 is not meant to imply an actual date; it is the date used by the date algorithm.

- If you give two time stamps, this function returns the integer number of hours between them. Instead, use the DATEDIFF function.
- If you give a date and an integer, this function adds the integer number of hours to the specified timestamp. Instead, use the DATEADD function.

See also

- “[DATEDIFF function \[Date and time\]](#)” on page 373
- “[DATEADD function \[Date and time\]](#)” on page 372

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statements return the value 4, signifying that the second timestamp 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/2003** Transact-SQL 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 430](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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 386](#)

Standards and compatibility

- **SQL/2003** Transact-SQL 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/2003** 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 366](#)

Standards and compatibility

- **SQL/2003** 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. This function is identical the LOWER function.

Syntax

```
LCASE( string-expression )
```

Parameters

- **string-expression** The string to be converted to lowercase.

Returns

CHAR

NCHAR

LONG VARCHAR

VARCHAR

NVARCHAR

Remarks

The LCASE function is similar to the LOWER function.

See also

- [“LOWER function \[String\]” on page 398](#)
- [“UCASE function \[String\]” on page 438](#)
- [“UPPER function \[String\]” on page 439](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

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

LONG NVARCHAR

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 will be described in terms of character length semantics.

See also

- [“RIGHT function \[String\]” on page 418](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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 360](#)
- [“International languages and character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

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 385](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 5.

```
SELECT LESSER( 10, 5 ) FROM dummy;
```

LIST function [Aggregate]

Returns a comma-delimited list of values.

Syntax

```
LIST(  
[ DISTINCT ] string-expression  
[ , delimiter-string ] )
```

Parameters

- **string-expression** A string expression, usually a column name. For each row in the column, the value is added to a comma-separated list. When DISTINCT is specified, only unique 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

LONG NVARCHAR

Remarks

NULL values are not added to the list. LIST (X) 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.

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/2003** 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 * 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 352](#)
- [“CHARINDEX function \[String\]” on page 364](#)

Standards and compatibility

- **SQL/2003** 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.

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, 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 398](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the natural logarithm of 50.

```
SELECT LOG( 50 );
```

LOG10 function [Numeric]

Returns the base 10 logarithm of a number.

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. 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 397](#)

Standards and compatibility

- **SQL/2003** 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 the LCASE function.

Syntax

```
LOWER( string-expression )
```


Parameters

- **string-expression** The string to be converted.

Returns

CHAR

NCHAR

LONG VARCHAR

VARCHAR

NVARCHAR

Remarks

The LCASE function is identical to the LOWER function.

See also

- [“LCASE function \[String\]” on page 392](#)
- [“UCASE function \[String\]” on page 438](#)
- [“UPPER function \[String\]” on page 439](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Core feature.

Example

The following statement returns the value chocolate.

```
SELECT LOWER( 'chOCOLate' );
```

LTRIM function [String]

Removes leading and trailing blanks from a string.

Syntax

```
LTRIM( string-expression )
```

Parameters

- **string-expression** The string to be trimmed.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

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 419](#)
- [“TRIM function \[String\]” on page 436](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

The TRIM specifications defined by the SQL/2003 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( expression | 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.

See also

- [“MIN function \[Aggregate\]” on page 401](#)

Standards and compatibility

- **SQL/2003** Core feature.

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( expression | 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.

See also

- [“MAX function \[Aggregate\]” on page 400](#)

Standards and compatibility

- **SQL/2003** 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 a minute component of a datetime value.

Syntax

MINUTE(*datetime-expression*)

Parameters

- **datetime-expression** The datetime value.

Returns

SMALLINT

Remarks

The value returned is a number from number from 0 to 59 corresponding to the datetime minute.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 22.

```
SELECT MINUTE( '1998-07-13 12:22:34' );
```

MINUTES function [Date and time]

The information you specify changes the behavior of this function:

- If you give a single date, this function returns the number of minutes since 0000-02-29.

Note

0000-02-29 is not meant to imply an actual date; it is the date used by the date algorithm.

- If you give two time stamps, this function returns the integer number of minutes between them. Instead, use the DATEDIFF function.
- If you give a date and an integer, this function adds the integer number of minutes to the specified timestamp. Instead, use the DATEADD function.

Syntax 1: integer

MINUTES([*datetime-expression*,] *datetime-expression*)

Syntax 2: timestamp

MINUTES(*datetime-expression*, *integer-expression*)

Parameters

- **datetime-expression** A date and time.

- **integer-expression** The number of minutes to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of minutes is subtracted from the datetime value. If you supply an integer expression, the *datetime-expression* must be explicitly cast as a DATETIME data type.

Returns

INT
TIMESTAMP

Remarks

Since this function returns an integer, overflow can occur when syntax 1 is used with time stamps greater than or equal to 4083-03-23 02:08:00.

See also

- [“CAST function \[Data type conversion\]” on page 361](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statements return the value 240, signifying that the second timestamp 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 1999-05-12 21:10:07.000. It is recommended that you use the second example (DATEADD).

```
SELECT MINUTES( CAST( '1999-05-12 21:05:07'
                    AS DATETIME ), 5);

SELECT DATEADD( minute, 5, '1999-05-12 21:05:07' );
```

ML_GET_SERVER_NOTIFICATION [System]

This function allows Ultralite users to use light weight polling to query a notifier on a MobiLink server for server-initiated sync requests.

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}
```

for example.

- **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 434](#)

Standards and compatibility**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 414](#)

Standards and compatibility

- **SQL/2003** SQL foundation feature outside core SQL.

Example

The following statement returns the value 2.

```
SELECT MOD( 5, 3 );
```

MONTH function [Date and time]

Returns a month of the given date.

Syntax

```
MONTH( date-expression )
```

Parameters

- **date-expression** A datetime value.

Returns

SMALLINT

Remarks

The value returned is a number from number from 1 to 12 corresponding to the datetime month.

Standards and compatibility

- **SQL/2003** 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

- **date-expression** The datetime value.

Returns

VARCHAR

Remarks

The MONTHNAME function returns a string, even if the result is numeric, such as 2 for the month of February.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value September.

```
SELECT MONTHNAME( '1998-09-05' );
```

MONTHS function [Date and time]

The information you specify changes the behavior of this function:

- If you give a single date, this function returns the number of months since 0000-02.

Note

0000-02 is not meant to imply an actual date; it is the date used by the date algorithm.

- If you give two time stamps, this function returns the integer number of months between them. Instead, use the DATEDIFF function.
- If you give a date and an integer, this function adds the integer number of minutes to the specified timestamp. Instead, use the DATEADD function.

Syntax 1: integer

MONTHS([*datetime-expression*,] *datetime-expression*)

Syntax 2: timestamp

MONTHS(*datetime-expression*, *integer-expression*)

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of months to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of months is subtracted from the datetime value. If you supply an *integer-expression*, the *datetime-expression* must be explicitly cast as a datetime data type.

For information about casting data types, see [“CAST function \[Data type conversion\]”](#) on page 361.

Returns

INT

TIMESTAMP

Remarks

The value of MONTHS is calculated from the number of first days of the month between the two dates.

Standards and compatibility

- **SQL/2003** 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 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 429
- “UIDTOSTR function [String]” on page 441

Standards and compatibility

- **SQL/2003** 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 NEWID();
```

For example, the value returned might be 96603324-6FF6-49DE-BF7D-F44C1C7E6856.

NOW function [Date and time]

Returns the current year, month, day, hour, minute, second, and fraction of a second. The accuracy is limited by the accuracy of the system clock.

Syntax

```
NOW( * )
```

Returns

TIMESTAMP

Remarks

The information the NOW function returns is equivalent to the information returned by the GETDATE function and the CURRENT_TIMESTAMP special value.

See also

- [“GETDATE function \[Date and time\]” on page 385](#)
- [“CURRENT_TIMESTAMP special value” on page 309](#)

Standards and compatibility

- **SQL/2003** 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 327](#)

Standards and compatibility

- **SQL/2003** Core feature.

Example

The following statement returns the value a:

```
SELECT NULLIF( 'a', 'b' );
```

The following statement returns NULL.

```
SELECT NULLIF( 'a', 'a' );
```

SQL functions (P-Z)

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 348](#).

See also

- [“SQL functions \(A-D\)” \[SQL Anywhere Server - SQL Reference\]](#)
- [“SQL functions \(E-O\)” \[SQL Anywhere Server - SQL Reference\]](#)

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 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 396](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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. Note that 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]%', 'SQLAnywhere11 has many new features' );
```

To get the first alphanumeric word in a string, you can use something like the following:

```
SELECT LEFT( @string, PATINDEX( '%[^a-zA-Z0-9]%', @string ) );
```

PI function [Numeric]

Returns the numeric value PI.

Syntax

PI(*)

Returns

DOUBLE

Standards and compatibility

- **SQL/2003** Vendor extension.

Remarks

This function returns a DOUBLE value.

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. If any argument is NULL, the result is a NULL value.

Standards and compatibility

- **SQL/2003** Vendor extension.

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 date expression.

Syntax

QUARTER(*date-expression*)

Parameters

- **date-expression** The date.

Returns

INT

Remarks

The quarters are as follows:

Quarter	Period (inclusive)
1	January 1 to March 31
2	April 1 to June 30

Quarter	Period (inclusive)
3	July 1 to September 30
4	October 1 to December 31

Standards and compatibility

- **SQL/2003** 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.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns a value of approximately 0.5236.

```
SELECT RADIANS( 30 );
```

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

Alternatively, try using the MOD function.

See also

- [“MOD function \[Numeric\]” on page 404](#)

Standards and compatibility

- **SQL/2003** 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

LONG NVARCHAR

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.

Alternatively, try using the REPLICATE function.

See also

- [“REPLICATE function \[String\]” on page 417](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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

LONG NVARCHAR

Remarks

This function replaces all occurrences.

Comparisons are case-sensitive on case-sensitive databases.

See also

- “SUBSTRING function [String]” on page 430
- “CHARINDEX function [String]” on page 364
- “UltraLite string functions” on page 352

Standards and compatibility

- **SQL/2003** 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

LONG NVARCHAR

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.

Alternatively, try using the REPEAT function.

See also

- [“REPEAT function \[String\]” on page 415](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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 be left-truncated.
- **integer-expression** The number of characters at the end of the string to return.

Returns

LONG VARCHAR

LONG NVARCHAR

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 will be described in terms of character length semantics.

See also

- [“LEFT function \[String\]” on page 393](#)
- [“International languages and character sets” \[SQL Anywhere Server - Database Administration\]](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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 437](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 123.200.

```
SELECT ROUND( 123.234, 1 );
```

RTRIM function [String]

Removes leading and trailing blanks from a string.

Syntax

```
RTRIM( string-expression )
```

Parameters

- **string-expression** The string to be trimmed.

Returns

VARCHAR
NVARCHAR
LONG VARCHAR
LONG NVARCHAR

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 436](#)
- [“LTRIM function \[String\]” on page 399](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

The TRIM specifications defined by the SQL/2003 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 a second of the given date.

Syntax

```
SECOND( datetime-expression )
```

Parameters

- **datetime-expression** The datetime value.

Returns

SMALLINT

Remarks

Returns a number from 0 to 59 corresponding to the second component of the given datetime value.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 25.

```
SELECT SECOND( '1998-07-13 21:21:25' );
```

SECONDS function [Date and time]

The behavior of this function can vary depending on what you supply:

- If you give a single date, this function returns the number of seconds since 0000-02-29.

Note

0000-02-29 is not meant to imply an actual date; it is the date used by the date algorithm.

- If you give two time stamps, this function returns the integer number of seconds between them. Instead, use the DATEDIFF function.
- If you give a date and an integer, this function adds the integer number of seconds to the specified timestamp. Instead, use the DATEADD function.

Syntax 1: integer

```
SECONDS( [ datetime-expression, ] datetime-expression )
```

Syntax 2: timestamp

```
SECONDS( datetime-expression, integer-expression )
```

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of seconds to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of minutes is subtracted from the datetime value. If you supply an integer expression, the *datetime-expression* must be explicitly cast as a datetime data type.

Returns

INTEGER

TIMESTAMP

See also

- “CAST function [Data type conversion]” on page 361
- “DATEADD function [Date and time]” on page 372
- “DATEDIFF function [Date and time]” on page 373

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statements return the value 14400, signifying that the second timestamp 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 datetime 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 383](#)

Standards and compatibility

- **SQL/2003** 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' );
```

This information can help with decisions about indexes to add or how to structure your database for better performance.

SIGN function [Numeric]

Returns the sign of a number.

Syntax

```
SIGN( numeric-expression )
```

Parameters

- **numeric-expression** The number for which the sign is to be returned.

Returns

SMALL INT

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/2003** 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. Join the table to itself and produce 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 352](#)

Standards and compatibility

- **SQL/2003** 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, and returns a DOUBLE as the result.

See also

- [“ASIN function \[Numeric\]” on page 357](#)
- [“COS function \[Numeric\]” on page 369](#)
- [“COT function \[Numeric\]” on page 369](#)
- [“TAN function \[Numeric\]” on page 435](#)

Standards and compatibility

- **SQL/2003** 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 352](#)

Standards and compatibility

- **SQL/2003** 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 352](#)

Standards and compatibility

- **SQL/2003** 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, and returns a DOUBLE as the result.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value 3.

```
SELECT SQRT( 9 );
```

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 352](#)

Standards and compatibility

- **SQL/2003** 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 more than one argument is supplied, they are concatenated into a single string.

Returns

LONG VARCHAR

LONG NVARCHAR

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 352](#)

Standards and compatibility

- **SQL/2003** 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.

Not needed in newer databases

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 312](#).

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

- [“UUIDTOSTR function \[String\]” on page 441](#)
- [“NEWID function \[Miscellaneous\]” on page 407](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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

LONG NVARCHAR

See also

- [“INSERTSTR function \[String\]” on page 390](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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

LONG NVARCHAR

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/2003 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 terms of character length semantics.

See also

- [“BYTE_SUBSTR function \[String\]” on page 360](#)
- [“LEFT function \[String\]” on page 393](#)
- [“RIGHT function \[String\]” on page 418](#)
- [“CHARINDEX function \[String\]” on page 364](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Core feature.

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(*expression* | **DISTINCT** *expression*)

Parameters

- **expression** The object to be summed. This is commonly a column name.
- **DISTINCT expression** Computes the sum of the unique values of *expression* in the input.

Returns

INTEGER

DOUBLE

NUMERIC

Remarks

Rows where the specified expression is NULL are not included.

Returns NULL for a group containing no rows.

See also

- [“COUNT function \[Aggregate\]” on page 370](#)
- [“AVG function \[Aggregate\]” on page 359](#)

Standards and compatibility

- **SQL/2003** Core feature.

Example

The following statement returns the value 3749146.740.

```
SELECT SUM( Salary )  
FROM Employees;
```

SUSER_ID function [System]

Returns the numeric user ID for the specified user name.

Syntax

```
SUSER_ID( [ user-name ] )
```

Parameters

- **user-name** The user name for the user ID you are searching for.

Returns

INT

Remarks

If you do not specify *user-name*, the ID of the current user is returned.

See also

- “SUSER_NAME function [System]” [*SQL Anywhere Server - SQL Reference*]
- “USER_ID function [System]” [*SQL Anywhere Server - SQL Reference*]

Standards and compatibility

- **SQL/2003** Transact-SQL extension.

Example

The following statement returns the ID for the GROUPO user.

```
SELECT SUSER_ID( 'GROUPO' );
```

SUSER_NAME function [System]

Returns the user name for the specified user ID.

Syntax

```
SUSER_NAME( [ user-id ] )
```

Parameters

- **user-id** The user ID of the user you are searching for.

Returns

LONG VARCHAR

Remarks

If you do not specify *user-id*, the user name of the current user is returned.

See also

- [“SUSER_ID function \[System\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“USER_NAME function \[System\]” \[SQL Anywhere Server - SQL Reference\]](#)

Standards and compatibility

- **SQL/2003** Transact-SQL extension.

Example

The following statement returns the user name for a user with ID 101.

```
SELECT SUSER_NAME( 101 );
```

SYNC_PROFILE_OPTION_VALUE function [System]

Returns the value of the option corresponding to the given option name.

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 \[System\]” on page 403](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

Consider the profile:

```
MobilinkUid=joe;Stream=tcPIP{host=sybase;port=1234};Ping=1
```

- **MobilinkUid** joe
- **Stream** tcPIP{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, and returns a DOUBLE as the result.

See also

- [“COS function \[Numeric\]” on page 369](#)
- [“SIN function \[Numeric\]” on page 424](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value of the tan of 0.52.

```
SELECT TAN( 0.52 );
```

TODAY function [Date and time]

Returns the current date.

Syntax

TODAY(*)

Returns

DATE

Remarks

Use this syntax in place of the historical CURRENT DATE function.

Standards and compatibility

- **SQL/2003** 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

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

See also

- [“LTRIM function \[String\]” on page 399](#)
- [“RTRIM function \[String\]” on page 419](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** The TRIM function is a SQL/2003 core feature.

SQL Anywhere does not support the additional parameters *trim specification* and *trim character*, as defined in SQL/2003. The SQL Anywhere implementation of TRIM corresponds to a TRIM specification of BOTH.

For the other TRIM specifications defined by the SQL/2003 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.
- **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

If any parameter is NULL, the result is NULL.

See also

- [“ROUND function \[Numeric\]” on page 419](#)

Standards and compatibility

- **SQL/2003** 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 );
```

TSEQUAL function [System] (deprecated)

Compares two timestamp values and returns whether they are the same.

Syntax**TSEQUAL** (*timestamp1*, *timestamp2*)**Parameters**

- **timestamp1** Timestamp expression.
- **timestamp2** Timestamp expression.

Returns

BIT

Remarks

The TSEQUAL function can only be used in a WHERE clause and is most commonly used as part of an UPDATE statement.

If *timestamp1* is equal to *timestamp2*, a row has changed since it was fetched. If the row changed, its timestamp has been modified and the TSEQUAL function returns FALSE. When the TSEQUAL function returns FALSE, the application determines that no rows were updated and assumes that the row was modified by another user. The updated row is re-fetched.

You can use the TSEQUAL function to determine whether a row has been changed since it was fetched.

See also

- [“The data type of a timestamp column” \[SQL Anywhere Server - SQL Usage\]](#)
- [“TIMESTAMP special value” \[SQL Anywhere Server - SQL Reference\]](#)
- [“The special Transact-SQL timestamp column and data type” \[SQL Anywhere Server - SQL Usage\]](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

Suppose you create a TIMESTAMP column Products.LastUpdated to store the timestamp for the last time the row was updated. The following UPDATE statement uses the TSEQUAL function to determine whether to update the row. If the value of LastUpdated is '2010/12/25 11:08:34.173226', the row is updated.

```
UPDATE Products
SET Color = 'Yellow'
WHERE ID = '300'
AND TSEQUAL( LastUpdated, '2010/12/25 11:08:34.173226' );
```

UCASE function [String]

Converts all characters in a string to uppercase. This function is identical to the UPPER function.

Syntax**UCASE**(*string-expression*)

Parameters

- **string-expression** The string to be converted to uppercase.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

See also

- [“UPPER function \[String\]” on page 439](#)
- [“LCASE function \[String\]” on page 392](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value CHOCOLATE.

```
SELECT UCASE( 'ChocoLate' );
```

UPPER function [String]

Converts all characters in a string to uppercase. This function is identical the UCASE function.

Syntax

```
UPPER( string-expression )
```

Parameters

- **string-expression** The string to be converted to uppercase.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Remarks

The UCASE function is similar to the UPPER function.

See also

- [“UCASE function \[String\]” on page 438](#)
- [“LCASE function \[String\]” on page 392](#)
- [“LOWER function \[String\]” on page 398](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the value CHOCOLATE.

```
SELECT UPPER( 'ChocoLate' );
```

USER_ID function [System]

Returns the numeric user ID for the specified user name.

Syntax

```
USER_ID( [ user-name ] )
```

Parameters

- **user-name** The user name for the user ID you are searching for.

Returns

INT

Remarks

If you do not specify *user-name*, the ID of the current user is returned.

See also

- [“USER_NAME function \[System\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“SUSER_ID function \[System\]” \[SQL Anywhere Server - SQL Reference\]](#)

Standards and compatibility

- **SQL/2003** Transact-SQL extension.

Example

The following statement returns the GROUPO user ID.

```
SELECT USER_ID( 'GROUPO' );
```

USER_NAME function [System]

Returns the user name for the specified user ID.

Syntax

```
USER_NAME( [ user-id ] )
```

Parameters

- **user-id** The user ID of the user you are searching for.

Returns

LONG VARCHAR

Remarks

If you do not specify *user-id*, the user name of the current user is returned.

See also

- “USER_ID function [System]” [[SQL Anywhere Server - SQL Reference](#)]
- “SUSER_NAME function [System]” [[SQL Anywhere Server - SQL Reference](#)]

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following statement returns the user name for user ID 101.

```
SELECT USER_NAME( 101 );
```

UUIDTOSTR function [String]

Converts a unique identifier value (UUID, also known as GUID) to a string value.

Not needed in newer databases

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 312.

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-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 407](#)
- [“STRTOUUID function \[String\]” on page 429](#)
- [“UltraLite string functions” on page 352](#)

Standards and compatibility

- **SQL/2003** 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]

Returns the number of weeks between two dates.

Syntax 1

```
WEEKS( [ datetime-expression, ] datetime-expression )
```

Syntax 2

WEEKS(*datetime-expression*, *integer-expression*)

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of weeks to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of weeks is subtracted from the datetime value. If you supply an *integer-expression*, the *datetime-expression* must be explicitly cast as DATETIME.

Returns

Syntax 1 returns an INTEGER.

Syntax 2 returns a TIMESTAMP.

Remarks

Given a single date (Syntax 1), the WEEKS function returns the number of weeks since 0000-02-29.

Given two dates (Syntax 1), 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. This can cause the values returned 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, and 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 373](#).

Given a date and an integer (Syntax 2), the WEEKS function adds the integer number of weeks to the specified date. This function is similar to the DATEADD function.

For more information about the DATEADD function, see [“DATEADD function \[Date and time\]” on page 372](#).

See also

For information about casting data types, see [“CAST function \[Data type conversion\]” on page 361](#).

Standards and compatibility

- **SQL/2003** 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 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]

Takes a timestamp value as a parameter and returns the year specified by that timestamp.

Syntax

```
YEAR( datetime-expression )
```

Parameters

- **datetime-expression** A date, time, or timestamp.

Returns

SMALLINT

Remarks

The value is returned as a SMALL INT.

Standards and compatibility

- **SQL/2003** Vendor extension.

Example

The following example returns the value 2001.

```
SELECT YEAR( '2001-09-12' );
```

YEARS function [Date and time]

Given two dates, this function returns the integer number of years between them. It is recommended that you use the DATEDIFF function instead. See [“DATEDIFF function \[Date and time\]” on page 373](#).

Given one date, it returns the year. It is recommended that you use the DATEPART function instead. See [“DATEPART function \[Date and time\]” \[SQL Anywhere Server - SQL Reference\]](#).

Given one date and an integer, it adds the integer number of years to the specified date. It is recommended that you use the DATEADD function instead. See [“DATEADD function \[Date and time\]” \[SQL Anywhere Server - SQL Reference\]](#).

Syntax 1

```
YEARS( [ datetime-expression, ] datetime-expression )
```

Syntax 2

YEARS(*datetime-expression*, *integer-expression*)

Parameters

- **datetime-expression** A date and time.
- **integer-expression** The number of years to be added to the *datetime-expression*. If *integer-expression* is negative, the appropriate number of years is subtracted from the datetime value. If you supply an *integer-expression*, the *datetime-expression* must be explicitly cast as a datetime data type.

For information about casting data types, see [“CAST function \[Data type conversion\]” on page 361](#).

Returns

Syntax 1 returns an INTEGER. Syntax 2 returns a TIMESTAMP.

Remarks

The value of YEARS is calculated from the number of first days of the year between the two dates.

Standards and compatibility

- **SQL/2003** 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. Values are small integers from -32768 to 32767.

Syntax

YMD(
integer-expression1,

integer-expression2,
integer-expression3)

Parameters

- **integer-expression1** The year.
- **integer-expression2** The number of the month. The year is adjusted if the month is outside the range 1-12.
- **integer-expression3** The day number. The day can be any integer; the date is adjusted.

Returns

DATE

Standards and compatibility

- **SQL/2003** 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. For example, the following statement returns the value 2000-03-01.

```
SELECT YMD( 1999, 15, 1 );
```

UltraLite SQL statements

Contents

UltraLite statement categories	449
UltraLite ALTER DATABASE SCHEMA FROM FILE statement	450
UltraLite ALTER PUBLICATION statement	451
UltraLite ALTER SYNCHRONIZATION PROFILE statement	452
UltraLite ALTER TABLE statement	454
UltraLite CHECKPOINT statement	458
UltraLite COMMIT statement	459
UltraLite CREATE INDEX statement	460
UltraLite CREATE PUBLICATION statement	462
UltraLite CREATE SYNCHRONIZATION PROFILE statement	464
UltraLite CREATE TABLE statement	468
UltraLite DELETE statement	473
UltraLite DROP INDEX statement	474
UltraLite DROP PUBLICATION statement	475
UltraLite DROP SYNCHRONIZATION PROFILE statement	476
UltraLite DROP TABLE statement	477
UltraLite FROM clause	478
UltraLite INSERT statement	480
UltraLite LOAD TABLE statement	481
UltraLite ROLLBACK statement	485
UltraLite SELECT statement	486
UltraLite SET OPTION statement	488
UltraLite START SYNCHRONIZATION DELETE statement	489
UltraLite STOP SYNCHRONIZATION DELETE statement	490
UltraLite SYNCHRONIZE statement	491
UltraLite TRUNCATE TABLE statement	493
UltraLite UNION statement	495
UltraLite UPDATE statement	496

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. See “[Syntax conventions](#)” [*SQL Anywhere Server - SQL Reference*].
- Using UltraLite SQL creates a transaction. A transaction consists of all changes (INSERTs, UPDATEs, and DELETEs) since the last ROLLBACK or COMMIT. See “[UltraLite transaction processing](#)” on page 16.

These changes can be made permanent by executing a COMMIT. A ROLLBACK statement causes the changes to be removed. See “[UltraLite COMMIT statement](#)” on page 459 and “[UltraLite ROLLBACK statement](#)” on page 485.

- If you are looking for statements used by Interactive SQL, see “[SQL statements](#)” [*SQL Anywhere Server - SQL Reference*]. Statements used by Interactive SQL have **[Interactive SQL]** after the statement name. For example, “[CONFIGURE statement \[Interactive SQL\]](#)” [*SQL Anywhere Server - SQL Reference*].

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. See [“UltraLite SELECT statement” on page 486](#).
- **Data manipulation statements** Allow you to change content in the database. Data manipulation is achieved with:
 - [“UltraLite DELETE statement” on page 473](#)
 - [“UltraLite INSERT statement” on page 480](#)
 - [“UltraLite UPDATE statement” on page 496](#)
- **Data definition statements** Allow you to define the structure or schema of a database. The schema can be changed with:
 - [“UltraLite ALTER DATABASE SCHEMA FROM FILE statement” on page 450](#)
 - [“UltraLite CREATE INDEX statement” on page 460](#)
 - [“UltraLite CREATE TABLE statement” on page 468](#)
 - [“UltraLite DROP INDEX statement” on page 474](#)
 - [“UltraLite DROP TABLE statement” on page 477](#)
 - [“UltraLite ALTER TABLE statement” on page 454](#)
 - [“UltraLite TRUNCATE TABLE statement” on page 493](#)
- **Transaction control statements** Allow you to control transactions within your UltraLite application. Transaction control is achieved with:
 - [“UltraLite CHECKPOINT statement” on page 458](#)
 - [“UltraLite COMMIT statement” on page 459](#)
 - [“UltraLite ROLLBACK statement” on page 485](#)
- **Synchronization management** Allow you to temporarily control synchronization with a MobiLink server. Synchronization management is achieved with:
 - [“UltraLite START SYNCHRONIZATION DELETE statement” on page 489](#)
 - [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)
 - [“UltraLite CREATE PUBLICATION statement” on page 462](#)
 - [“UltraLite ALTER PUBLICATION statement” on page 451](#)
 - [“UltraLite DROP PUBLICATION statement” on page 475](#)

See also

- [“Expressions in UltraLite” on page 325](#)
- [“Operators in UltraLite” on page 338](#)

UltraLite ALTER DATABASE SCHEMA FROM FILE statement

Use this statement to modify the schema definition of an existing UltraLite database using a SQL script.

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). See [“UltraLite Initialize Database utility \(ulinit\)” on page 269](#), and [“UltraLite Unload Database utility \(ulunload\)” on page 281](#).

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

- [“Deploying UltraLite schema upgrades” on page 67](#)

Example

The following statement modifies the schema of the database using a SQL script, *MySchema.sql*:

```
ALTER DATABASE SCHEMA FROM FILE 'MySchema.sql';
```

UltraLite ALTER PUBLICATION statement

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 331](#)
- [“Designing synchronization in UltraLite” on page 133](#)
- [“UltraLite CREATE PUBLICATION statement” on page 462](#)
- [“UltraLite DROP PUBLICATION statement” on page 475](#)
- [“UltraLite START SYNCHRONIZATION DELETE statement” on page 489](#)
- [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)

Example

The following ALTER PUBLICATION statement adds the Customers table to the pub_contact publication.

```
ALTER PUBLICATION pub_contact  
  ADD TABLE Customers;
```

UltraLite ALTER SYNCHRONIZATION PROFILE statement

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  
{ REPLACE | 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.
- **REPLACE clause** Use this clause to drop the options currently defined for the profile, and add the specified options instead.
- **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

See [“UltraLite CREATE SYNCHRONIZATION PROFILE statement” on page 464](#) for a list of the synchronization profile options supported by UltraLite.

Side effects

None.

See also

- [“UltraLite DROP SYNCHRONIZATION PROFILE statement” on page 476](#)
- [“UltraLite SYNCHRONIZE statement” on page 491](#)

Example

The following is an example of the ALTER SYNCHRONIZATION PROFILE...REPLACE statement:

```
CREATE SYNCHRONIZATION PROFILE myProfile1;  
ALTER SYNCHRONIZATION PROFILE myProfile1  
  REPLACE 'publication=p1;uploadonly=on';
```

The following is an example of the ALTER SYNCHRONIZATION PROFILE...MERGE statement.

```
CREATE SYNCHRONIZATION PROFILE myProfile2 'publication=p1;verbosity=high';  
ALTER SYNCHRONIZATION PROFILE myProfile2  
    MERGE 'publication=p2;uploadonly=on';
```

UltraLite ALTER TABLE statement

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*

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**], ...)

Parameters

add-clause 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 [“UltraLite CREATE TABLE statement” on page 468](#).
- **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 [“UltraLite CREATE TABLE statement” on page 468](#).

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. For a full explanation of *column-definition*, see [“UltraLite CREATE TABLE statement” on page 468](#).

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 [“UltraLite CREATE TABLE statement” on page 468](#).

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.

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.

A table whose name ends with **nosync** can only be renamed to a table name that also ends with **nosync**. See [“Nosync tables in UltraLite” on page 136](#).

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 [“Dispose method” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“Schema changes with DDL statements” on page 11](#).

See also

- [“UltraLite CREATE TABLE statement” on page 468](#)
- [“UltraLite DROP TABLE statement” on page 477](#)
- [“Data types in UltraLite” on page 312](#)
- [“Altering tables” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Using table and column constraints” \[SQL Anywhere Server - SQL Usage\]](#)
- [“Overriding partition sizes for autoincremented columns” on page 131](#)
- [“Determining the most recently assigned GLOBAL AUTOINCREMENT value” on page 131](#)

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  
MODIFY Street CHAR(50);
```

UltraLite CHECKPOINT statement

Use this statement to checkpoint the database.

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

- [“Flushing single or grouped transactions” on page 119](#)
- [“UltraLite COMMIT statement” on page 459](#)
- [“UltraLite COMMIT_FLUSH connection parameter” on page 229](#)
- UltraLite embedded SQL: [“ULCheckpoint function” \[UltraLite - C and C++ Programming\]](#)
- UltraLite C++: [“Checkpoint function” \[UltraLite - C and C++ Programming\]](#)

Example

The following statement performs a checkpoint of the database:

```
CHECKPOINT ;
```

UltraLite COMMIT statement

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

- [“UltraLite CHECKPOINT statement” on page 458](#)
- [“UltraLite ROLLBACK statement” on page 485](#)

Example

The following statement makes the changes in the current transaction permanent in the database:

```
COMMIT;
```

UltraLite CREATE INDEX statement

Use this statement to create an index on a specified table.

Syntax

```
CREATE [ UNIQUE ] INDEX [ index-name ]  
ON table-name ( ordered-column-list )  
[ WITH MAX HASH SIZE integer ]
```

ordered-column-list :
(*column-name* [ASC | DESC], ...)

Parameters

UNIQUE The UNIQUE parameter ensures that there are not 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 in a unique index are allowed to be NULL, while columns in a unique constraint are not. Also, a foreign key can reference either a primary key or a column with a unique constraint. A foreign key cannot, however, reference a unique index.

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. To learn the default size, see [“Accessing UltraLite database properties” on page 218](#).

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 also use execution plans to optimize queries. See [“Execution plans in UltraLite” on page 342](#).

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“Dispose method” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“Schema changes with DDL statements” on page 11](#).

Side effects

- Automatic commit.

See also

- [“UltraLite performance and optimization” on page 107](#)
- [“UltraLite DROP INDEX statement” on page 474](#)
- [“Working with UltraLite indexes” on page 79](#)

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 );
```

UltraLite CREATE PUBLICATION statement

Use this statement to create a publication. A publication identifies synchronized data in an UltraLite remote database.

Syntax

```
CREATE PUBLICATION publication-name  
( TABLE table-name [ WHERE search-condition ], ... )
```

Parameters

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. See [“Search conditions in UltraLite” on page 331](#).

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

- [“UltraLite clients” on page 125](#)
- [“UltraLite DROP PUBLICATION statement” on page 475](#)
- [“UltraLite ALTER PUBLICATION statement” on page 451](#)
- [“Search conditions in UltraLite” on page 331](#)

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'  
);
```

UltraLite CREATE SYNCHRONIZATION PROFILE statement

Use this statement to create an UltraLite synchronization profile. Synchronization profiles define how an UltraLite database synchronizes with the MobiLink server.

Syntax

```
CREATE SYNCHRONIZATION PROFILE sync-profile-name 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.
- **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

There are two possible ways to make changes to an existing synchronization profile. The first is to use the REPLACE clause. This will replace the contents of the synchronization profile with whatever is contained in the new sync-option string. This is the same as dropping the synchronization profile and then creating one with the same name but using the new string. Note, therefore, that a synchronization profile does not need to contain a full synchronization definition because parameters can be merged in (or overridden) at synchronization time.

The second way to modify a synchronization profile is to use the MERGE clause. When using this clause, only the sync options that are specified in the MERGE clause are changed. To remove a sync 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{}

The following table lists the synchronization profile options supported by UltraLite.

Synchronization profile option	Valid values	Description
Allow-Download-DupRows	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 155.
Auth-Parms	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 156.
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 155.
Continue-Download	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 “Resuming failed downloads” [<i>MobiLink - Server Administration</i>].
Disable-Concurrency	Boolean	Disallow database access from other threads during synchronization. See “Additional Parameters synchronization parameter” on page 155.
DownloadOnly	Boolean	Performs a download-only synchronization. See “Download Only synchronization parameter” on page 159.
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 161.
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 option” [<i>MobiLink - Client Administration</i>]. See “-mn 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 option” [<i>MobiLink - Client Administration</i>].
Ping	Boolean	Confirms communications with the server only; no synchronization is performed. See “Ping synchronization parameter” on page 165.

Synchronization profile option	Valid values	Description
Publications	String (comma separated)	Specifies the publications(s) to synchronize. The publications determine the tables on the remote 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 “Publications in UltraLite” on page 137 .
Script-Version	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” [MobiLink - Client Administration] .
SendColumn-Names	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 168 .
Send-DownloadACK	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 169 .
Stream	String (with sub-list)	Specifies the MobiLink network synchronization protocol. See “Stream Type synchronization parameter” on page 171 .
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 155 .
UploadOnly	String	Specifies that synchronization will only include an upload, and no download will occur. See “Upload Only synchronization parameter” on page 174 .

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.

Side effects

None.

See also

- “UltraLite ALTER SYNCHRONIZATION PROFILE statement” on page 452
- “UltraLite DROP SYNCHRONIZATION PROFILE statement” on page 476
- “UltraLite SYNCHRONIZE statement” on page 491

Example

The following creates a synchronization profile called Test1.

```
CREATE SYNCHRONIZATION PROFILE Test1
'MobiLinkUid=mary;Stream=TCPIP{host=192.168.1.1;port=1234}'
```

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}'.

UltraLite CREATE TABLE statement

Use this statement to create a table.

Syntax

```
CREATE TABLE table-name (  
  { column-definition | table-constraint }, ...  
)
```

column-definition :
column-name *data-type*
[[**NOT**] **NULL**]
[**DEFAULT** *column-default*]
[*column-constraint*]

column-default :
GLOBAL AUTOINCREMENT [(*number*)]
| **AUTOINCREMENT**
| **CURRENT DATE**
| **CURRENT TIME**
| **CURRENT TIMESTAMP**
| **NULL**
| **NEWID()**
| *constant-value*

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**], ...)

Parameters

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 303](#).
- **data-type** The data type of the column. See [“Data types in UltraLite” on page 312](#).
- **[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:
 - **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. See [“Using GLOBAL AUTOINCREMENT in UltraLite” on page 129](#), and [“UltraLite global_database_id option” on page 222](#).
- **[NOT] NULL** Controls whether the column can contain NULLs.
- **NEWID()** A function that generates a unique identifier value. See [“NEWID function \[Miscellaneous\]” on page 407](#).
- **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 309](#).
- **CURRENT DATE** Stores the current year, month, and day. See [“CURRENT DATE special value” on page 308](#).
- **CURRENT TIME** Stores the current hour, minute, second and fraction of a second. See [“CURRENT TIME special value” on page 308](#).
- **constant-value** A constant for the data type of the column. Typically the constant is a number or a string.
- **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** 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, provided that inconsistent data is resolved prior to 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. To learn the default size, see [“Accessing UltraLite database properties” on page 218](#).

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 [“Schema changes with DDL statements” on page 11](#).

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 [“Dispose method” \[UltraLite - .NET Programming\]](#).

Side effects

Automatic commit.

See also

- [“Expressions in UltraLite” on page 325](#)
- [“UltraLite DROP TABLE statement” on page 477](#)
- [“CREATE TABLE statement” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Data types in UltraLite” on page 312](#)
- [“Overriding partition sizes for autoincremented columns” on page 131](#)

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)  
);
```

UltraLite DELETE statement

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. See [“Search conditions in UltraLite” on page 331](#).

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 [“UltraLite row states” on page 14](#).

See also

- [“UltraLite START SYNCHRONIZATION DELETE statement” on page 489](#)
- [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)

Example

The following statement removes employee 105 from the Employees table.

```
DELETE
FROM Employees
WHERE EmployeeID = 105;
```

The following statement removes all data prior to the year 2000 from the FinancialData table.

```
DELETE
FROM FinancialData
WHERE Year < 2000;
```

UltraLite DROP INDEX statement

Use this statement to delete an index.

Syntax

```
DROP INDEX [ 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.

For UltraLite.NET users: You cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“Dispose method” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“Schema changes with DDL statements” on page 11](#).

See also

- [“UltraLite CREATE INDEX statement” on page 460](#)
- [“Working with UltraLite indexes” on page 79](#)

Example

The following statement deletes a fictitious index, `fin_codes_idx`, on the `FinancialData` table:

```
DROP INDEX FinancialData.fin_codes_idx;
```

UltraLite DROP PUBLICATION statement

Use this statement to delete publications.

Syntax

```
DROP PUBLICATION publication-name, ...
```

See also

- [“Designing synchronization in UltraLite” on page 133](#)
- [“UltraLite ALTER PUBLICATION statement” on page 451](#)
- [“UltraLite CREATE PUBLICATION statement” on page 462](#)

Example

The following statement drops the `pub_contact` publication.

```
DROP PUBLICATION pub_contact;
```

UltraLite DROP SYNCHRONIZATION PROFILE statement

Use this statement to delete an UltraLite synchronization profile. Synchronization profiles define how an UltraLite database synchronizes with the MobiLink server.

Syntax

```
DROP SYNCHRONIZATION PROFILE sync-profile-name
```

Parameters

- **sync-profile-name** The name of the synchronization profile.

Remarks

None.

Side effects

None.

See also

- [“UltraLite CREATE SYNCHRONIZATION PROFILE statement” on page 464](#)
- [“UltraLite ALTER SYNCHRONIZATION PROFILE statement” on page 452](#)
- [“UltraLite SYNCHRONIZE statement” on page 491](#)

Example

The following example shows the syntax for dropping a synchronization profile called Test1.

```
DROP SYNCHRONIZATION PROFILE Test1;
```

UltraLite DROP TABLE statement

Use this statement to remove a table, and all its data, from a database.

Syntax

```
DROP TABLE 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.

For UltraLite.NET, you cannot execute this statement unless you also call the Dispose method for all data objects (for example, ULDataReader). See [“Dispose method” \[UltraLite - .NET Programming\]](#).

Statements are not released if schema changes are initiated at the same time. See [“Schema changes with DDL statements” on page 11](#).

See also

- [“UltraLite ALTER TABLE statement” on page 454](#)
- [“UltraLite CREATE TABLE statement” on page 468](#)

Example

The following statement deletes a fictitious table, EmployeeBenefits, from the database:

```
DROP TABLE EmployeeBenefits;
```

UltraLite FROM clause

Use this clause to specify the tables or views involved in a SELECT statement.

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. See [“Subqueries in expressions” on page 329](#).

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.

Derived tables

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

- “Joins: Retrieving data from several tables” [[SQL Anywhere Server - SQL Usage](#)]
- “UltraLite DELETE statement” on page 473
- “UltraLite SELECT statement” on page 486
- “UltraLite UPDATE statement” on page 496

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;
```

UltraLite INSERT statement

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

- [“UltraLite SELECT statement” on page 486](#)

Example

The following statement adds an Eastern Sales department to the database.

```
INSERT
INTO Departments ( DepartmentID, DepartmentName )
VALUES ( 230, 'Eastern Sales' );
```

UltraLite LOAD TABLE statement

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 [ 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.
- **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).

LOAD TABLE is only supported for Windows and Linux, not Palm OS or Windows Mobile.

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, " (single-quote 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

- “UltraLite INSERT statement” on page 480
- “UltraLite Unload Database utility (ulunload)” on page 281
- “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

UltraLite ROLLBACK statement

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

- [“UltraLite COMMIT statement” on page 459](#)

Example

The following statement rolls the database back to the state it was in at the previous commit:

```
ROLLBACK;
```

UltraLite SELECT statement

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 [“UltraLite FROM clause” on page 478](#).

WHERE clause If a WHERE clause is specified, only rows satisfying *search-condition* are selected. See [“Search conditions in UltraLite” on page 331](#).

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.

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 [“Using direct page scans” on page 117](#).
- **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 342](#).

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 and optimization” on page 107](#)
- [“SELECT statement” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Querying data” \[SQL Anywhere Server - SQL Usage\]](#)

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;
```

UltraLite SET OPTION statement

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

You can only set database options with this statement and properties cannot be modified after the database has been created. The exception to these rules is **ml_remote_id** (see below).

You cannot specify whether an option is persistent or not. The way an option has been implemented in UltraLite determines whether it is a persistent or temporary option. Persistent options are stored in the sysuldata table. Temporary options are used only until the database stops running.

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.

See also

- [“sysuldata system table” on page 298](#)
- [“UltraLite database options” on page 219](#)
- [“DB_PROPERTY function \[System\]” on page 380](#)
- [“UltraLite ml_remote_id option” on page 223](#)

Example

The following statement sets the global_database_id option to 100:

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

UltraLite START SYNCHRONIZATION DELETE statement

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 [“UltraLite row states” on page 14](#).

See also

- [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)

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;
```

UltraLite STOP SYNCHRONIZATION DELETE statement

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 [“UltraLite row states” on page 14](#).

See also

- [“UltraLite START SYNCHRONIZATION DELETE statement” on page 489](#)

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;
```

UltraLite SYNCHRONIZE statement

Use this statement to synchronize an UltraLite 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

See [“UltraLite CREATE SYNCHRONIZATION PROFILE statement” on page 464](#) for a list of the synchronization profile options supported by UltraLite.

See [“UltraLite ALTER SYNCHRONIZATION PROFILE statement” on page 452](#) to understand how sync options are merged with existing options in the synchronization profile.

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

- “UltraLite ALTER SYNCHRONIZATION PROFILE statement” on page 452
- “UltraLite DROP SYNCHRONIZATION PROFILE statement” on page 476
- “ULRegisterSynchronizationCallback” [*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 'publication=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 'publication=p2;uploadonly=on';
```

UltraLite TRUNCATE TABLE statement

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 [“Dispose method” \[UltraLite - .NET Programming\]](#).

Schema changes Statements are not released if schema changes are initiated at the same time. See [“Schema changes with DDL statements” on page 11](#).

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 DELETE statement” on page 473](#)
- [“UltraLite START SYNCHRONIZATION DELETE statement” on page 489](#)
- [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)

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. See [“UltraLite ROLLBACK statement” on page 485](#).

UltraLite UNION statement

Use this statement to combine the results of two or more select statements.

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.

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

- [“UltraLite SELECT statement” on page 486](#)

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;
```

UltraLite UPDATE statement

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-hand 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. See [“Search conditions in UltraLite” on page 331](#).

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

- [“UltraLite INSERT statement” on page 480](#)
- [“UltraLite DELETE statement” on page 473](#)
- [“Search conditions in UltraLite” on page 331](#)

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 )
```

Troubleshooting UltraLite

Contents

Unable to start the UltraLite engine	498
Unable to connect to databases after upgrade	499
Database corruption	500
Database size not stabilizing	501
Importing ASCII data into a new database	502
Utilities still running as the previous version	503
Result set changes unpredictably	504
UltraLite engine client fails with error -764	505

Unable to start the UltraLite engine

Symptom

You have use the START connection parameter to start the UltraLite engine with the following definition; however, the client returns SQLE_UNABLE_TO_CONNECT_OR_START.

```
START="\Program Files\uleng11.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\uleng11.exe"
```

or

```
START='"\Program Files\uleng11.exe"'
```

Unable to connect to databases after upgrade

Symptom

You have upgraded UltraLite. You discover that you are able to create an empty UltraLite database using the administration tools. However, when you try to connect to this or any other UltraLite 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

- [“Upgrading UltraLite” \[SQL Anywhere 11 - Changes and Upgrading\]](#)

Database corruption

Symptom

Your 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 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, 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 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 to iAnywhere. 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 database to a file.
3. Create a new database.
4. Repopulate the data either by synchronizing or by loading the unloaded data.

See also

- [“UltraLite checksum_level creation parameter” on page 184](#)
- [“ULSQLCode enumeration” \[*UltraLite - .NET Programming*\]](#)

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 your UltraLite 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 schema upgrades” on page 67](#).
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

- [“UltraLite STOP SYNCHRONIZATION DELETE statement” on page 490](#)
- [“UltraLite TRUNCATE TABLE statement” on page 493](#)
- [“UltraLite DELETE statement” on page 473](#)
- [“UltraLite ALTER DATABASE SCHEMA FROM FILE statement” on page 450](#)

Importing ASCII data into a new 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 choose **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 to 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 utility for UltraLite (dbisql)” on page 249

Utilities still running as the previous version

Symptom

You have just installed UltraLite 11. 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 [“Using the utilities” \[SQL Anywhere 11 - Changes and Upgrading\]](#).

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

- [“UltraLite SELECT statement” on page 486](#)
- [“Database corruption” on page 500](#)

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, in some cases you may need to change this code:

```
ULConnection conn = new ULConnection(@"dbf=\Program Files\HelloEngine
\HelloEngine.udb;
START=\Windows\uleng11.exe")
```

To something similar to:

```
ULConnection conn = new ULConnection(@"dbf=\\\\"Program Files \\"
\HelloEngine\HelloEngine.udb;
START=\\Windows\uleng11.exe");
```

See also

- [“Deploy Multiple UltraLite applications with the UltraLite engine” on page 57](#)
- [“UltraLite START connection parameter” on page 245](#)

Glossary

Glossary 509

Glossary

Adaptive Server Anywhere (ASA)

The relational database server component of SQL Anywhere Studio, intended for use in mobile and embedded environments or as a server for small and medium-sized businesses. In version 10.0.0, Adaptive Server Anywhere was renamed SQL Anywhere Server, and SQL Anywhere Studio was renamed SQL Anywhere.

See also: [“SQL Anywhere” on page 533](#).

agent ID

See also: [“client message store ID” on page 511](#).

article

In MobiLink or SQL Remote, an article is a database object that represents a whole table, or a subset of the columns and rows in a table. Articles are grouped together in a publication.

See also:

- [“replication” on page 531](#)
- [“publication” on page 528](#)

atomic transaction

A transaction that is guaranteed to complete successfully or not at all. If an error prevents part of an atomic transaction from completing, the transaction is rolled back to prevent the database from being left in an inconsistent state.

base table

Permanent tables for data. Tables are sometimes called **base tables** to distinguish them from temporary tables and views.

See also:

- [“temporary table” on page 535](#)
- [“view” on page 537](#)

bit array

A bit array is a type of array data structure that is used for efficient storage of a sequence of bits. A bit array is similar to a character string, except that the individual pieces are 0s (zeros) and 1s (ones) instead of characters. Bit arrays are typically used to hold a string of Boolean values.

business rule

A guideline based on real-world requirements. Business rules are typically implemented through check constraints, user-defined data types, and the appropriate use of transactions.

See also:

- [“constraint” on page 513](#)
- [“user-defined data type” on page 537](#)

carrier

A MobiLink object, stored in MobiLink system tables or a Notifier properties file, that contains information about a public carrier for use by server-initiated synchronization.

See also: [“server-initiated synchronization” on page 532](#).

character set

A character set is a set of symbols, including letters, digits, spaces, and other symbols. An example of a character set is ISO-8859-1, also known as Latin1.

See also:

- [“code page” on page 511](#)
- [“encoding” on page 517](#)
- [“collation” on page 511](#)

check constraint

A restriction that enforces specified conditions on a column or set of columns.

See also:

- [“constraint” on page 513](#)
- [“foreign key constraint” on page 518](#)
- [“primary key constraint” on page 528](#)
- [“unique constraint” on page 536](#)

checkpoint

The point at which all changes to the database are saved to the database file. At other times, committed changes are saved only to the transaction log.

checksum

The calculated number of bits of a database page that is recorded with the database page itself. The checksum allows the database management system to validate the integrity of the page by ensuring that the numbers match as the page is being written to disk. If the counts match, it's assumed that page was successfully written.

client message store

In QAnywhere, a SQL Anywhere database on the remote device that stores messages.

client message store ID

In QAnywhere, a MobiLink remote ID that uniquely identifies a client message store.

client/server

A software architecture where one application (the client) obtains information from and sends information to another application (the server). The two applications often reside on different computers connected by a network.

code page

A code page is an encoding that maps characters of a character set to numeric representations, typically an integer between 0 and 255. An example of a code page is Windows code page 1252. For the purposes of this documentation, code page and encoding are interchangeable terms.

See also:

- [“character set” on page 510](#)
- [“encoding” on page 517](#)
- [“collation” on page 511](#)

collation

A combination of a character set and a sort order that defines the properties of text in the database. For SQL Anywhere databases, the default collation is determined by the operating system and language on which the server is running; for example, the default collation on English Windows systems is 1252LATIN1. A collation, also called a collating sequence, is used for comparing and sorting strings.

See also:

- [“character set” on page 510](#)
- [“code page” on page 511](#)
- [“encoding” on page 517](#)

command file

A text file containing SQL statements. Command files can be built manually, or they can be built automatically by database utilities. The dbunload utility, for example, creates a command file consisting of the SQL statements necessary to recreate a given database.

communication stream

In MobiLink, the network protocol used for communication between the MobiLink client and the MobiLink server.

concurrency

The simultaneous execution of two or more independent, and possibly competing, processes. SQL Anywhere automatically uses locking to isolate transactions and ensure that each concurrent application sees a consistent set of data.

See also:

- [“transaction” on page 535](#)
- [“isolation level” on page 521](#)

conflict resolution

In MobiLink, conflict resolution is logic that specifies what to do when two users modify the same row on different remote databases.

connection ID

A unique number that identifies a given connection between a client application and the database. You can determine the current connection ID using the following SQL statement:

```
SELECT CONNECTION_PROPERTY( 'Number' );
```

connection-initiated synchronization

A form of MobiLink server-initiated synchronization in which synchronization is initiated when there are changes to connectivity.

See also: [“server-initiated synchronization” on page 532](#).

connection profile

A set of parameters that are required to connect to a database, such as user name, password, and server name, that is stored and used as a convenience.

consolidated database

In distributed database environments, a database that stores the master copy of the data. In case of conflict or discrepancy, the consolidated database is considered to have the primary copy of the data.

See also:

- [“synchronization” on page 535](#)
- [“replication” on page 531](#)

constraint

A restriction on the values contained in a particular database object, such as a table or column. For example, a column may have a uniqueness constraint, which requires that all values in the column be different. A table may have a foreign key constraint, which specifies how the information in the table relates to data in some other table.

See also:

- [“check constraint” on page 510](#)
- [“foreign key constraint” on page 518](#)
- [“primary key constraint” on page 528](#)
- [“unique constraint” on page 536](#)

contention

The act of competing for resources. For example, in database terms, two or more users trying to edit the same row of a database contend for the rights to edit that row.

correlation name

The name of a table or view that is used in the FROM clause of a query—either its original name, or an alternate name, that is defined in the FROM clause.

creator ID

In UltraLite Palm OS applications, an ID that is assigned when the application is created.

cursor

A named linkage to a result set, used to access and update rows from a programming interface. In SQL Anywhere, cursors support forward and backward movement through the query results. Cursors consist of two parts: the cursor result set, typically defined by a SELECT statement; and the cursor position.

See also:

- [“cursor result set” on page 513](#)
- [“cursor position” on page 513](#)

cursor position

A pointer to one row within the cursor result set.

See also:

- [“cursor” on page 513](#)
- [“cursor result set” on page 513](#)

cursor result set

The set of rows resulting from a query that is associated with a cursor.

See also:

- [“cursor” on page 513](#)
- [“cursor position” on page 513](#)

data cube

A multi-dimensional result set with each dimension reflecting a different way to group and sort the same results. Data cubes provide complex information about data that would otherwise require self-join queries and correlated subqueries. Data cubes are a part of OLAP functionality.

data definition language (DDL)

The subset of SQL statements for defining the structure of data in the database. DDL statements create, modify, and remove database objects, such as tables and users.

data manipulation language (DML)

The subset of SQL statements for manipulating data in the database. DML statements retrieve, insert, update, and delete data in the database.

data type

The format of data, such as CHAR or NUMERIC. In the ANSI SQL standard, data types can also include a restriction on size, character set, and collation.

See also: [“domain” on page 516](#).

database

A collection of tables that are related by primary and foreign keys. The tables hold the information in the database. The tables and keys together define the structure of the database. A database management system accesses this information.

See also:

- [“foreign key” on page 518](#)
- [“primary key” on page 528](#)
- [“database management system \(DBMS\)” on page 515](#)
- [“relational database management system \(RDBMS\)” on page 530](#)

database administrator (DBA)

The user with the permissions required to maintain the database. The DBA is generally responsible for all changes to a database schema, and for managing users and groups. The role of database administrator is automatically built into databases as user ID DBA with password sql.

database connection

A communication channel between a client application and the database. A valid user ID and password are required to establish a connection. The privileges granted to the user ID determine the actions that can be carried out during the connection.

database file

A database is held in one or more database files. There is an initial file, and subsequent files are called dbspaces. Each table, including its indexes, must be contained within a single database file.

See also: [“dbspace” on page 516](#).

database management system (DBMS)

A collection of programs that allow you to create and use databases.

See also: [“relational database management system \(RDBMS\)” on page 530](#).

database name

The name given to a database when it is loaded by a server. The default database name is the root of the initial database file.

See also: [“database file” on page 515](#).

database object

A component of a database that contains or receives information. Tables, indexes, views, procedures, and triggers are database objects.

database owner (dbo)

A special user that owns the system objects not owned by SYS.

See also:

- [“database administrator \(DBA\)” on page 514](#)
- [“SYS” on page 535](#)

database server

A computer program that regulates all access to information in a database. SQL Anywhere provides two types of servers: network servers and personal servers.

DBA authority

The level of permission that enables a user to do administrative activity in the database. The DBA user has DBA authority by default.

See also: [“database administrator \(DBA\)” on page 514](#).

dbspace

An additional database file that creates more space for data. A database can be held in up to 13 separate files (an initial file and 12 dbspaces). Each table, together with its indexes, must be contained in a single database file. The SQL command CREATE DBSPACE adds a new file to the database.

See also: [“database file” on page 515](#).

deadlock

A state where a set of transactions arrives at a place where none can proceed.

device tracking

In MobiLink server-initiated synchronization, functionality that allows you to address messages using the MobiLink user name that identifies a device.

See also: [“server-initiated synchronization” on page 532](#).

direct row handling

In MobiLink, a way to synchronize table data to sources other than the MobiLink-supported consolidated databases. You can implement both uploads and downloads with direct row handling.

See also:

- [“consolidated database” on page 512](#)
- [“SQL-based synchronization” on page 533](#)

domain

Aliases for built-in data types, including precision and scale values where applicable, and optionally including DEFAULT values and CHECK conditions. Some domains, such as the monetary data types, are pre-defined in SQL Anywhere. Also called user-defined data type.

See also: [“data type” on page 514](#).

download

The stage in synchronization where data is transferred from the consolidated database to a remote database.

dynamic SQL

SQL that is generated programmatically by your program before it is executed. UltraLite dynamic SQL is a variant designed for small-footprint devices.

EBF

Express Bug Fix. An express bug fix is a subset of the software with one or more bug fixes. The bug fixes are listed in the release notes for the update. Bug fix updates may only be applied to installed software with the same version number. Some testing has been performed on the software, but the software has not

undergone full testing. You should not distribute these files with your application unless you have verified the suitability of the software yourself.

embedded SQL

A programming interface for C programs. SQL Anywhere embedded SQL is an implementation of the ANSI and IBM standard.

encoding

Also known as character encoding, an encoding is a method by which each character in a character set is mapped onto one or more bytes of information, typically represented as a hexadecimal number. An example of an encoding is UTF-8.

See also:

- [“character set” on page 510](#)
- [“code page” on page 511](#)
- [“collation” on page 511](#)

event model

In MobiLink, the sequence of events that make up a synchronization, such as `begin_synchronization` and `download_cursor`. Events are invoked if a script is created for them.

external login

An alternate login name and password used when communicating with a remote server. By default, SQL Anywhere uses the names and passwords of its clients whenever it connects to a remote server on behalf of those clients. However, this default can be overridden by creating external logins. External logins are alternate login names and passwords used when communicating with a remote server.

extraction

In SQL Remote replication, the act of unloading the appropriate structure and data from the consolidated database. This information is used to initialize the remote database.

See also: [“replication” on page 531](#).

failover

Switching to a redundant or standby server, system, or network on failure or unplanned termination of the active server, system, or network. Failover happens automatically.

FILE

In SQL Remote replication, a message system that uses shared files for exchanging replication messages. This is useful for testing and for installations without an explicit message-transport system.

See also: [“replication” on page 531](#).

file-based download

In MobiLink, a way to synchronize data in which downloads are distributed as files, allowing offline distribution of synchronization changes.

file-definition database

In MobiLink, a SQL Anywhere database that is used for creating download files.

See also: [“file-based download” on page 518](#).

foreign key

One or more columns in a table that duplicate the primary key values in another table. Foreign keys establish relationships between tables.

See also:

- [“primary key” on page 528](#)
- [“foreign table” on page 518](#)

foreign key constraint

A restriction on a column or set of columns that specifies how the data in the table relates to the data in some other table. Imposing a foreign key constraint on a set of columns makes those columns the foreign key.

See also:

- [“constraint” on page 513](#)
- [“check constraint” on page 510](#)
- [“primary key constraint” on page 528](#)
- [“unique constraint” on page 536](#)

foreign table

The table containing the foreign key.

See also: [“foreign key” on page 518](#).

full backup

A backup of the entire database, and optionally, the transaction log. A full backup contains all the information in the database and provides protection in the event of a system or media failure.

See also: [“incremental backup” on page 520](#).

gateway

A MobiLink object, stored in MobiLink system tables or a Notifier properties file, that contains information about how to send messages for server-initiated synchronization.

See also: [“server-initiated synchronization” on page 532](#).

generated join condition

A restriction on join results that is automatically generated. There are two types: key and natural. Key joins are generated when you specify `KEY JOIN` or when you specify the keyword `JOIN` but do not use the keywords `CROSS`, `NATURAL`, or `ON`. For a key join, the generated join condition is based on foreign key relationships between tables. Natural joins are generated when you specify `NATURAL JOIN`; the generated join condition is based on common column names in the two tables.

See also:

- [“join” on page 522](#)
- [“join condition” on page 522](#)

generation number

In MobiLink, a mechanism for forcing remote databases to upload data before applying any more download files.

See also: [“file-based download” on page 518](#).

global temporary table

A type of temporary table for which data definitions are visible to all users until explicitly dropped. Global temporary tables let each user open their own identical instance of a table. By default, rows are deleted on commit, and rows are always deleted when the connection is ended.

See also:

- [“temporary table” on page 535](#)
- [“local temporary table” on page 522](#)

grant option

The level of permission that allows a user to grant permissions to other users.

hash

A hash is an index optimization that transforms index entries into keys. An index hash aims to avoid the expensive operation of finding, loading, and then unpacking the rows to determine the indexed value, by including enough of the actual row data with its row ID.

histogram

The most important component of column statistics, histograms are a representation of data distribution. SQL Anywhere maintains histograms to provide the optimizer with statistical information about the distribution of values in columns.

iAnywhere JDBC driver

The iAnywhere JDBC driver provides a JDBC driver that has some performance benefits and feature benefits compared to the pure Java jConnect JDBC driver, but which is not a pure-Java solution. The iAnywhere JDBC driver is recommended in most cases.

See also:

- [“JDBC” on page 521](#)
- [“jConnect” on page 521](#)

identifier

A string of characters used to reference a database object, such as a table or column. An identifier may contain any character from A through Z, a through z, 0 through 9, underscore (_), at sign (@), number sign (#), or dollar sign (\$).

incremental backup

A backup of the transaction log only, typically used between full backups.

See also: [“transaction log” on page 535](#).

index

A sorted set of keys and pointers associated with one or more columns in a base table. An index on one or more columns of a table can improve performance.

InfoMaker

A reporting and data maintenance tool that lets you create sophisticated forms, reports, graphs, cross-tabs, and tables, and applications that use these reports as building blocks.

inner join

A join in which rows appear in the result set only if both tables satisfy the join condition. Inner joins are the default.

See also:

- [“join” on page 522](#)
- [“outer join” on page 526](#)

integrated login

A login feature that allows the same single user ID and password to be used for operating system logins, network logins, and database connections.

integrity

Adherence to rules that ensure that data is correct and accurate, and that the relational structure of the database is intact.

See also: [“referential integrity” on page 530](#).

Interactive SQL

A SQL Anywhere application that allows you to query and alter data in your database, and modify the structure of your database. Interactive SQL provides a pane for you to enter SQL statements, and panes that return information about how the query was processed and the result set.

isolation level

The degree to which operations in one transaction are visible to operations in other concurrent transactions. There are four isolation levels, numbered 0 through 3. Level 3 provides the highest level of isolation. Level 0 is the default setting. SQL Anywhere also supports three snapshot isolation levels: snapshot, statement-snapshot, and readonly-statement-snapshot.

See also: [“snapshot isolation” on page 533](#).

JAR file

Java archive file. A compressed file format consisting of a collection of one or more packages used for Java applications. It includes all the resources necessary to install and run a Java program in a single compressed file.

Java class

The main structural unit of code in Java. It is a collection of procedures and variables grouped together because they all relate to a specific, identifiable category.

jConnect

A Java implementation of the JavaSoft JDBC standard. It provides Java developers with native database access in multi-tier and heterogeneous environments. However, the iAnywhere JDBC driver is the preferred JDBC driver for most cases.

See also:

- [“JDBC” on page 521](#)
- [“iAnywhere JDBC driver” on page 520](#)

JDBC

Java Database Connectivity. A SQL-language programming interface that allows Java applications to access relational data. The preferred JDBC driver is the iAnywhere JDBC driver.

See also:

- [“jConnect” on page 521](#)
- [“iAnywhere JDBC driver” on page 520](#)

join

A basic operation in a relational system that links the rows in two or more tables by comparing the values in specified columns.

join condition

A restriction that affects join results. You specify a join condition by inserting an ON clause or WHERE clause immediately after the join. In the case of natural and key joins, SQL Anywhere generates a join condition.

See also:

- [“join” on page 522](#)
- [“generated join condition” on page 519](#)

join type

SQL Anywhere provides four types of joins: cross join, key join, natural join, and joins using an ON clause.

See also: [“join” on page 522](#).

light weight poller

In MobiLink server-initiated synchronization, a device application that polls for push notifications from a MobiLink server.

See also: [“server-initiated synchronization” on page 532](#).

Listener

A program, dblsn, that is used for MobiLink server-initiated synchronization. Listeners are installed on remote devices and configured to initiate actions on the device when they receive push notifications.

See also: [“server-initiated synchronization” on page 532](#).

local temporary table

A type of temporary table that exists only for the duration of a compound statement or until the end of the connection. Local temporary tables are useful when you need to load a set of data only once. By default, rows are deleted on commit.

See also:

- [“temporary table” on page 535](#)
- [“global temporary table” on page 519](#)

lock

A concurrency control mechanism that protects the integrity of data during the simultaneous execution of multiple transactions. SQL Anywhere automatically applies locks to prevent two connections from changing the same data at the same time, and to prevent other connections from reading data that is in the process of being changed.

You control locking by setting the isolation level.

See also:

- [“isolation level” on page 521](#)
- [“concurrency” on page 512](#)
- [“integrity” on page 521](#)

log file

A log of transactions maintained by SQL Anywhere. The log file is used to ensure that the database is recoverable in the event of a system or media failure, to improve database performance, and to allow data replication using SQL Remote.

See also:

- [“transaction log” on page 535](#)
- [“transaction log mirror” on page 536](#)
- [“full backup” on page 518](#)

logical index

A reference (pointer) to a physical index. There is no indexing structure stored on disk for a logical index.

LTM

Log Transfer Manager (LTM) also called Replication Agent. Used with Replication Server, the LTM is the program that reads a database transaction log and sends committed changes to Sybase Replication Server.

See: [“Replication Server” on page 531](#).

maintenance release

A maintenance release is a complete set of software that upgrades installed software from an older version with the same major version number (version number format is *major.minor.patch.build*). Bug fixes and other changes are listed in the release notes for the upgrade.

materialized view

A materialized view is a view that has been computed and stored on disk. Materialized views have characteristics of both views (they are defined using a query specification), and of tables (they allow most table operations to be performed on them).

See also:

- [“base table” on page 509](#)
- [“view” on page 537](#)

message log

A log where messages from an application such as a database server or MobiLink server can be stored. This information can also appear in a messages window or be logged to a file. The message log includes informational messages, errors, warnings, and messages from the MESSAGE statement.

message store

In QAnywhere, databases on the client and server device that store messages.

See also:

- [“client message store” on page 511](#)
- [“server message store” on page 533](#)

message system

In SQL Remote replication, a protocol for exchanging messages between the consolidated database and a remote database. SQL Anywhere includes support for the following message systems: FILE, FTP, and SMTP.

See also:

- [“replication” on page 531](#)
- [“FILE” on page 517](#)

message type

In SQL Remote replication, a database object that specifies how remote users communicate with the publisher of a consolidated database. A consolidated database may have several message types defined for it; this allows different remote users to communicate with it using different message systems.

See also:

- [“replication” on page 531](#)
- [“consolidated database” on page 512](#)

metadata

Data about data. Metadata describes the nature and content of other data.

See also: [“schema” on page 532](#).

mirror log

See also: [“transaction log mirror” on page 536](#).

MobiLink

A session-based synchronization technology designed to synchronize UltraLite and SQL Anywhere remote databases with a consolidated database.

See also:

- [“consolidated database” on page 512](#)
- [“synchronization” on page 535](#)
- [“UltraLite” on page 536](#)

MobiLink client

There are two kinds of MobiLink clients. For SQL Anywhere remote databases, the MobiLink client is the dbmlsync command line utility. For UltraLite remote databases, the MobiLink client is built in to the UltraLite runtime library.

MobiLink Monitor

A graphical tool for monitoring MobiLink synchronizations.

MobiLink server

The computer program that runs MobiLink synchronization, mlsrv11.

MobiLink system table

System tables that are required by MobiLink synchronization. They are installed by MobiLink setup scripts into the MobiLink consolidated database.

MobiLink user

A MobiLink user is used to connect to the MobiLink server. You create the MobiLink user on the remote database and register it in the consolidated database. MobiLink user names are entirely independent of database user names.

network protocol

The type of communication, such as TCP/IP or HTTP.

network server

A database server that accepts connections from computers sharing a common network.

See also: [“personal server” on page 527](#).

normalization

The refinement of a database schema to eliminate redundancy and improve organization according to rules based on relational database theory.

Notifier

A program that is used by MobiLink server-initiated synchronization. Notifiers are integrated into the MobiLink server. They check the consolidated database for push requests, and send push notifications.

See also:

- [“server-initiated synchronization” on page 532](#)
- [“Listener” on page 522](#)

object tree

In Sybase Central, the hierarchy of database objects. The top level of the object tree shows all products that your version of Sybase Central supports. Each product expands to reveal its own sub-tree of objects.

See also: [“Sybase Central” on page 534](#).

ODBC

Open Database Connectivity. A standard Windows interface to database management systems. ODBC is one of several interfaces supported by SQL Anywhere.

ODBC Administrator

A Microsoft program included with Windows operating systems for setting up ODBC data sources.

ODBC data source

A specification of the data a user wants to access via ODBC, and the information needed to get to that data.

outer join

A join that preserves all the rows in a table. SQL Anywhere supports left, right, and full outer joins. A left outer join preserves the rows in the table to the left of the join operator, and returns a null when a row in the right table does not satisfy the join condition. A full outer join preserves all the rows from both tables.

See also:

- [“join” on page 522](#)
- [“inner join” on page 520](#)

package

In Java, a collection of related classes.

parse tree

An algebraic representation of a query.

PDB

A Palm database file.

performance statistic

A value reflecting the performance of the database system. The CURRREAD statistic, for example, represents the number of file reads issued by the database server that have not yet completed.

personal server

A database server that runs on the same computer as the client application. A personal database server is typically used by a single user on a single computer, but it can support several concurrent connections from that user.

physical index

The actual indexing structure of an index, as it is stored on disk.

plug-in module

In Sybase Central, a way to access and administer a product. Plug-ins are usually installed and registered automatically with Sybase Central when you install the respective product. Typically, a plug-in appears as a top-level container, in the Sybase Central main window, using the name of the product itself; for example, SQL Anywhere.

See also: [“Sybase Central” on page 534](#).

policy

In QAnywhere, the way you specify when message transmission should occur.

polling

In MobiLink server-initiated synchronization, the way a light weight poller, such as the MobiLink Listener, requests push notifications from a Notifier.

See also: [“server-initiated synchronization” on page 532](#).

PowerDesigner

A database modeling application. PowerDesigner provides a structured approach to designing a database or data warehouse. SQL Anywhere includes the Physical Data Model component of PowerDesigner.

PowerJ

A Sybase product for developing Java applications.

predicate

A conditional expression that is optionally combined with the logical operators AND and OR to make up the set of conditions in a WHERE or HAVING clause. In SQL, a predicate that evaluates to UNKNOWN is interpreted as FALSE.

primary key

A column or list of columns whose values uniquely identify every row in the table.

See also: [“foreign key” on page 518](#).

primary key constraint

A uniqueness constraint on the primary key columns. A table can have only one primary key constraint.

See also:

- [“constraint” on page 513](#)
- [“check constraint” on page 510](#)
- [“foreign key constraint” on page 518](#)
- [“unique constraint” on page 536](#)
- [“integrity” on page 521](#)

primary table

The table containing the primary key in a foreign key relationship.

proxy table

A local table containing metadata used to access a table on a remote database server as if it were a local table.

See also: [“metadata” on page 524](#).

publication

In MobiLink or SQL Remote, a database object that identifies data that is to be synchronized. In MobiLink, publications exist only on the clients. A publication consists of articles. SQL Remote users can receive a publication by subscribing to it. MobiLink users can synchronize a publication by creating a synchronization subscription to it.

See also:

- [“replication” on page 531](#)
- [“article” on page 509](#)
- [“publication update” on page 528](#)

publication update

In SQL Remote replication, a list of changes made to one or more publications in one database. A publication update is sent periodically as part of a replication message to the remote database(s).

See also:

- [“replication” on page 531](#)
- [“publication” on page 528](#)

publisher

In SQL Remote replication, the single user in a database who can exchange replication messages with other replicating databases.

See also: [“replication” on page 531](#).

push notification

In QAnywhere, a special message delivered from the server to a QAnywhere client that prompts the client to initiate a message transmission. In MobiLink server-initiated synchronization, a special message delivered from a Notifer to a device that contains push request data and internal information.

See also:

- [“QAnywhere” on page 529](#)
- [“server-initiated synchronization” on page 532](#)

push request

In MobiLink server-initiated synchronization, a row of values in a result set that a Notifier checks to determine if push notifications need to be sent to a device.

See also: [“server-initiated synchronization” on page 532](#).

QAnywhere

Application-to-application messaging, including mobile device to mobile device and mobile device to and from the enterprise, that permits communication between custom programs running on mobile or wireless devices and a centrally located server application.

QAnywhere agent

In QAnywhere, a process running on the client device that monitors the client message store and determines when message transmission should occur.

query

A SQL statement or group of SQL statements that access and/or manipulate data in a database.

See also: [“SQL” on page 533](#).

Redirector

A web server plug-in that routes requests and responses between a client and the MobiLink server. This plug-in also implements load-balancing and failover mechanisms.

reference database

In MobiLink, a SQL Anywhere database used in the development of UltraLite clients. You can use a single SQL Anywhere database as both reference and consolidated database during development. Databases made with other products cannot be used as reference databases.

referencing object

An object, such as a view, whose definition directly references another object in the database, such as a table.

See also: [“foreign key” on page 518](#).

referenced object

An object, such as a table, that is directly referenced in the definition of another object, such as a view.

See also: [“primary key” on page 528](#).

referential integrity

Adherence to rules governing data consistency, specifically the relationships between the primary and foreign key values in different tables. To have referential integrity, the values in each foreign key must correspond to the primary key values of a row in the referenced table.

See also:

- [“primary key” on page 528](#)
- [“foreign key” on page 518](#)

regular expression

A regular expression is a sequence of characters, wildcards, and operators that defines a pattern to search for within a string.

relational database management system (RDBMS)

A type of database management system that stores data in the form of related tables.

See also: [“database management system \(DBMS\)” on page 515](#).

remote database

In MobiLink or SQL Remote, a database that exchanges data with a consolidated database. Remote databases may share all or some of the data in the consolidated database.

See also:

- [“synchronization” on page 535](#)
- [“consolidated database” on page 512](#)

REMOTE DBA authority

In SQL Remote, a level of permission required by the Message Agent (dbremote). In MobiLink, a level of permission required by the SQL Anywhere synchronization client (dbmlsync). When the Message Agent (dbremote) or synchronization client connects as a user who has this authority, it has full DBA access. The user ID has no additional permissions when not connected through the Message Agent (dbremote) or synchronization client (dbmlsync).

See also: [“DBA authority” on page 515](#).

remote ID

A unique identifier in SQL Anywhere and UltraLite databases that is used by MobiLink. The remote ID is initially set to NULL and is set to a GUID during a database's first synchronization.

replication

The sharing of data among physically distinct databases. Sybase has three replication technologies: MobiLink, SQL Remote, and Replication Server.

Replication Agent

See: [“LTM” on page 523](#).

replication frequency

In SQL Remote replication, a setting for each remote user that determines how often the publisher's message agent should send replication messages to that remote user.

See also: [“replication” on page 531](#).

replication message

In SQL Remote or Replication Server, a communication sent between a publishing database and a subscribing database. Messages contain data, passthrough statements, and information required by the replication system.

See also:

- [“replication” on page 531](#)
- [“publication update” on page 528](#)

Replication Server

A Sybase connection-based replication technology that works with SQL Anywhere and Adaptive Server Enterprise. It is intended for near-real time replication between a few databases.

See also: [“LTM” on page 523](#).

role

In conceptual database modeling, a verb or phrase that describes a relationship from one point of view. You can describe each relationship with two roles. Examples of roles are "contains" and "is a member of."

role name

The name of a foreign key. This is called a role name because it names the relationship between the foreign table and primary table. By default, the role name is the table name, unless another foreign key is already using that name, in which case the default role name is the table name followed by a three-digit unique number. You can also create the role name yourself.

See also: [“foreign key” on page 518](#).

rollback log

A record of the changes made during each uncommitted transaction. In the event of a ROLLBACK request or a system failure, uncommitted transactions are reversed out of the database, returning the database to its former state. Each transaction has a separate rollback log, which is deleted when the transaction is complete.

See also: [“transaction” on page 535](#).

row-level trigger

A trigger that executes once for each row that is changed.

See also:

- [“trigger” on page 536](#)
- [“statement-level trigger” on page 534](#)

schema

The structure of a database, including tables, columns, and indexes, and the relationships between them.

script

In MobiLink, code written to handle MobiLink events. Scripts programmatically control data exchange to meet business needs.

See also: [“event model” on page 517](#).

script-based upload

In MobiLink, a way to customize the upload process as an alternative to using the log file.

script version

In MobiLink, a set of synchronization scripts that are applied together to create a synchronization.

secured feature

A feature specified by the -sf option when a database server is started, so it is not available for any database running on that database server.

server-initiated synchronization

A way to initiate MobiLink synchronization from the MobiLink server.

server management request

A QAnywhere message that is formatted as XML and sent to the QAnywhere system queue as a way to administer the server message store or monitor QAnywhere applications.

server message store

In QAnywhere, a relational database on the server that temporarily stores messages until they are transmitted to a client message store or JMS system. Messages are exchanged between clients via the server message store.

service

In Windows operating systems, a way of running applications when the user ID running the application is not logged on.

session-based synchronization

A type of synchronization where synchronization results in consistent data representation across both the consolidated and remote databases. MobiLink is session-based.

snapshot isolation

A type of isolation level that returns a committed version of the data for transactions that issue read requests. SQL Anywhere provides three snapshot isolation levels: snapshot, statement-snapshot, and readonly-statement-snapshot. When using snapshot isolation, read operations do not block write operations.

See also: [“isolation level” on page 521](#).

SQL

The language used to communicate with relational databases. ANSI has defined standards for SQL, the latest of which is SQL-2003. SQL stands, unofficially, for Structured Query Language.

SQL Anywhere

The relational database server component of SQL Anywhere that is intended for use in mobile and embedded environments or as a server for small and medium-sized businesses. SQL Anywhere is also the name of the package that contains the SQL Anywhere RDBMS, the UltraLite RDBMS, MobiLink synchronization software, and other components.

SQL-based synchronization

In MobiLink, a way to synchronize table data to MobiLink-supported consolidated databases using MobiLink events. For SQL-based synchronization, you can use SQL directly or you can return SQL using the MobiLink server APIs for Java and .NET.

SQL Remote

A message-based data replication technology for two-way replication between consolidated and remote databases. The consolidated and remote databases must be SQL Anywhere.

SQL statement

A string containing SQL keywords designed for passing instructions to a DBMS.

See also:

- [“schema” on page 532](#)
- [“SQL” on page 533](#)
- [“database management system \(DBMS\)” on page 515](#)

statement-level trigger

A trigger that executes after the entire triggering statement is completed.

See also:

- [“trigger” on page 536](#)
- [“row-level trigger” on page 532](#)

stored procedure

A stored procedure is a group of SQL instructions stored in the database and used to execute a set of operations or queries on a database server

string literal

A string literal is a sequence of characters enclosed in single quotes.

subquery

A SELECT statement that is nested inside another SELECT, INSERT, UPDATE, or DELETE statement, or another subquery.

There are two types of subquery: correlated and nested.

subscription

In MobiLink synchronization, a link in a client database between a publication and a MobiLink user, allowing the data described by the publication to be synchronized.

In SQL Remote replication, a link between a publication and a remote user, allowing the user to exchange updates on that publication with the consolidated database.

See also:

- [“publication” on page 528](#)
- [“MobiLink user” on page 525](#)

Sybase Central

A database management tool that provides SQL Anywhere database settings, properties, and utilities in a graphical user interface. Sybase Central can also be used for managing other Sybase products, including MobiLink.

synchronization

The process of replicating data between databases using MobiLink technology.

In SQL Remote, synchronization is used exclusively to denote the process of initializing a remote database with an initial set of data.

See also:

- [“MobiLink” on page 525](#)
- [“SQL Remote” on page 533](#)

SYS

A special user that owns most of the system objects. You cannot log in as SYS.

system object

Database objects owned by SYS or dbo.

system table

A table, owned by SYS or dbo, that holds metadata. System tables, also known as data dictionary tables, are created and maintained by the database server.

system view

A type of view, included in every database, that presents the information held in the system tables in an easily understood format.

temporary table

A table that is created for the temporary storage of data. There are two types: global and local.

See also:

- [“local temporary table” on page 522](#)
- [“global temporary table” on page 519](#)

transaction

A sequence of SQL statements that comprise a logical unit of work. A transaction is processed in its entirety or not at all. SQL Anywhere supports transaction processing, with locking features built in to allow concurrent transactions to access the database without corrupting the data. Transactions end either with a COMMIT statement, which makes the changes to the data permanent, or a ROLLBACK statement, which undoes all the changes made during the transaction.

transaction log

A file storing all changes made to a database, in the order in which they are made. It improves performance and allows data recovery in the event the database file is damaged.

transaction log mirror

An optional identical copy of the transaction log file, maintained simultaneously. Every time a database change is written to the transaction log file, it is also written to the transaction log mirror file.

A mirror file should be kept on a separate device from the transaction log, so that if either device fails, the other copy of the log keeps the data safe for recovery.

See also: [“transaction log” on page 535](#).

transactional integrity

In MobiLink, the guaranteed maintenance of transactions across the synchronization system. Either a complete transaction is synchronized, or no part of the transaction is synchronized.

transmission rule

In QAnywhere, logic that determines when message transmission is to occur, which messages to transmit, and when messages should be deleted.

trigger

A special form of stored procedure that is executed automatically when a user runs a query that modifies the data.

See also:

- [“row-level trigger” on page 532](#)
- [“statement-level trigger” on page 534](#)
- [“integrity” on page 521](#)

UltraLite

A database optimized for small, mobile, and embedded devices. Intended platforms include cell phones, pagers, and personal organizers.

UltraLite runtime

An in-process relational database management system that includes a built-in MobiLink synchronization client. The UltraLite runtime is included in the libraries used by each of the UltraLite programming interfaces, and in the UltraLite engine.

unique constraint

A restriction on a column or set of columns requiring that all non-null values are different. A table can have multiple unique constraints.

See also:

- [“foreign key constraint” on page 518](#)
- [“primary key constraint” on page 528](#)
- [“constraint” on page 513](#)

unload

Unloading a database exports the structure and/or data of the database to text files (SQL command files for the structure, and ASCII comma-separated files for the data). You unload a database with the Unload utility.

In addition, you can unload selected portions of your data using the UNLOAD statement.

upload

The stage in synchronization where data is transferred from a remote database to a consolidated database.

user-defined data type

See [“domain” on page 516](#).

validate

To test for particular types of file corruption of a database, table, or index.

view

A SELECT statement that is stored in the database as an object. It allows users to see a subset of rows or columns from one or more tables. Each time a user uses a view of a particular table, or combination of tables, it is recomputed from the information stored in those tables. Views are useful for security purposes, and to tailor the appearance of database information to make data access straightforward.

window

The group of rows over which an analytic function is performed. A window may contain one, many, or all rows of data that has been partitioned according to the grouping specifications provided in the window definition. The window moves to include the number or range of rows needed to perform the calculations for the current row in the input. The main benefit of the window construct is that it allows additional opportunities for grouping and analysis of results, without having to perform additional queries.

Windows

The Microsoft Windows family of operating systems, such as Windows Vista, Windows XP, and Windows 200x.

Windows CE

See [“Windows Mobile” on page 537](#).

Windows Mobile

A family of operating systems produced by Microsoft for mobile devices.

work table

An internal storage area for interim results during query optimization.

Index

Symbols

- % operator
 - modulo function, UltraLite, 404
- &
 - bitwise operator for UltraLite, 339
- comment indicator
 - UltraLite about, 305
- a option
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
- b option
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
- c option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite erase database [ulerase] utility, 262
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite information [ulinfo] utility, 265
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
 - UltraLite validate database utility (ulvalid) utility, 286
- creator-ID
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
- d option
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite validate database utility (ulvalid) utility, 286
- d1 option
 - UltraLite Interactive SQL [dbisql] utility, 249
- e option
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite validate database utility (ulvalid) utility, 286
- ek option
 - UltraLite erase database [ulerase] utility, 262
- ep option
 - UltraLite erase database [ulerase] utility, 262
- f option
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
- g option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite information [ulinfo] utility, 265
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
- h option
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
- I option
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
- k option
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
- l option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite initialize database [ulinit] utility, 269
- n option
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite initialize database [ulinit] utility, 269

- UltraLite load XML to database [ulload] utility, 272
- UltraLite SQL Preprocessor [sqlpp] utility, 252
- UltraLite unload data to XML [ulunload] utility, 281
- nogui option
 - UltraLite Interactive SQL [dbisql] utility, 249
- o option
 - ulcreate creation parameters, 34
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
- oa option
 - UltraLite information [ulinfo] utility, 265
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite validate database utility (ulvalid) utility, 286
- ol option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite load XML to database [ulload] utility, 272
- onerror option
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite load XML to database [ulload] utility, 272
- or option
 - UltraLite information [ulinfo] utility, 265
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite validate database utility (ulvalid) utility, 286
- ou option
 - UltraLite information [ulinfo] utility, 265
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
- UltraLite validate database utility (ulvalid) utility, 286
- output-file
 - UltraLite unload data to XML [ulunload] utility, 281
- p option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
- q option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite erase database [ulerase] utility, 262
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite information [ulinfo] utility, 265
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
 - UltraLite validate database utility (ulvalid) utility, 286
- qq option
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
- r option
 - UltraLite information [ulinfo] utility, 265
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite synchronization [ulsync] utility, 275
- rc option
 - UltraLite information [ulinfo] utility, 265
- s option
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
- SQL command
 - UltraLite Interactive SQL [dbisql] utility, 249

- t option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite unload data to XML [ulunload] utility, 281
 - u option
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - ul option
 - UltraLite Interactive SQL [dbisql] utility, 249
 - v option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite erase database [ulerase] utility, 262
 - UltraLite information [ulinfo] utility, 265
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite synchronization [ulsync] utility, 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
 - UltraLite validate database utility (ulvalid) utility, 286
 - version
 - Interactive SQL [dbisql] utility, 249
 - w option
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - x option
 - UltraLite Interactive SQL [dbisql] utility, 249
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite unload data to XML [ulunload] utility, 281
 - xml-file
 - UltraLite unload old database [ulunloadold] utility, 284
 - y option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
 - z option
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - .NET
 - UltraLite engine support, 23
 - .NET compatibility
 - UltraLite driver for ADO.NET, 24
 - /* comment indicator
 - UltraLite about, 305
 - // comment indicator
 - UltraLite about, 305
 - 10054
 - UltraLite synchronization stream system errors, 170
 - 128-bit strong encryption
 - UltraLite connection parameter for, 233
 - UltraLite usage of, 40
 - 130 error
 - SQL code for UltraLite schema upgrade, 67
 - ?
 - UltraLite input parameter, 330
 - @data option
 - UltraLite Interactive SQL [dbisql] utility, 249
 - ^
 - bitwise operator for UltraLite, 339
 - |
 - bitwise operator for UltraLite, 339
 - ~
 - bitwise operator for UltraLite, 339
- A**
- ABS function
 - UltraLite syntax, 354
 - ACOS function
 - UltraLite syntax, 354
 - ActiveSync
 - deploying MobiLink UltraLite applications, 150
 - registering applications for UltraLite clients, 66
 - UltraLite deploying provider files, 64
 - adding
 - UltraLite column methods, 72
 - UltraLite columns, 454
 - UltraLite indexes, 81
 - UltraLite users, 87
 - adding synchronization
 - UltraLite applications, 139
 - adding UltraLite indexes
 - about, 81

- adding UltraLite users
 - about, 87
- adding users
 - UltraLite, 87
- additional parameters
 - UltraLite synchronization parameter, 155
- administration tools
 - UltraLite troubleshooting, 499
- administration utilities
 - UltraLite utilities reference, 247
- ADO.NET
 - UltraLite drivers for, 24
- AES encryption algorithm
 - UltraLite deployment steps, 59
 - UltraLite fips creation parameter, 192
 - UltraLite fips usage, 40
 - UltraLite usage, 40
- agent IDs
 - glossary definition, 509
- aggregate expressions
 - UltraLite SQL syntax, 328
- aggregate functions
 - UltraLite alphabetical list, 348
- aliases
 - UltraLite columns, 486
 - UltraLite equivalents, 314
- ALL conditions
 - UltraLite SQL, 334
- AllowDownloadDupRows
 - UltraLite synchronization parameter, 155
- allsync tables
 - UltraLite overview, 72
 - UltraLite synchronizing tables, 136
- ALTER DATABASE SCHEMA FROM FILE statement
 - UltraLite schema changes impact, 11
 - UltraLite syntax, 450
 - usage, 67
- ALTER PUBLICATION statement
 - UltraLite syntax, 451
- ALTER SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 452
- ALTER TABLE statement
 - UltraLite Interactive SQL example, 74
 - UltraLite syntax, 454
- altering
 - UltraLite ALTER PUBLICATION statement, 451
 - UltraLite ALTER TABLE statement, 454
 - UltraLite column methods, 74
 - UltraLite columns methods, 73
 - UltraLite table methods, 74
- altering UltraLite column definitions
 - about, 73
- ambiguous string to date conversions
 - UltraLite, 196
- AND
 - bitwise operators for UltraLite, 339
 - logical operators for UltraLite, 333
- ANY conditions
 - UltraLite SQL, 335
- APIs
 - UltraLite choices, 24
- applications
 - (*see also* UltraLite applications)
- ApplyFile method
 - UltraLite replacement for schema upgrade , 67
- arc-cosine function
 - UltraLite ACOS function, 354
- arc-sine function
 - UltraLite ASIN function, 357
- arc-tangent function
 - UltraLite ATAN function, 357
- ARGN function
 - UltraLite syntax, 355
- arithmetic
 - operators and UltraLite SQL syntax, 338
- arithmetic operators
 - UltraLite SQL syntax, 338
- articles
 - glossary definition, 509
 - UltraLite copying method, 77
 - UltraLite databases, 137
 - UltraLite restrictions, 137
- ASCII
 - UltraLite sorting, 36
 - UltraLite syntax, 356
- ASCII files
 - UltraLite importing, 502
- ASIN function
 - UltraLite syntax, 357
- assembling parameters into connection strings
 - UltraLite about, 49
- ATAN function
 - UltraLite syntax, 357
- ATAN2 function
 - UltraLite syntax, 358

- atomic transactions
 - glossary definition, 509
- authentication
 - UltraLite bypassing, 51
 - UltraLite setup, 51
- authentication parameters
 - UltraLite synchronization parameter, 156
- authentication status
 - UltraLite synchronization parameter, 157
- authentication value
 - UltraLite synchronization parameter, 159
- autocommit
 - UltraLite transaction overview , 16
- AUTOINCREMENT
 - UltraLite syntax, 468
- average function
 - UltraLite AVG function, 359
- AVG function
 - UltraLite syntax, 359

B

- backing up (*see* backups)
- backup and data recovery
 - UltraLite about, 15
- backups
 - UltraLite databases on Palm, 258
 - UltraLite databases on Windows Mobile, 47
 - UltraLite internal mechanism, 15
 - UltraLite transaction overview , 16
- base 10 logarithm
 - UltraLite LOG10 function, 398
- base tables
 - glossary definition, 509
- BETWEEN conditions
 - UltraLite SQL, 336
- BIGINT data type
 - UltraLite, 312
- binary
 - UltraLite sorting, 36
- BINARY data types
 - UltraLite, 312
 - UltraLite maximum size, 8
- bit arrays
 - glossary definition, 510
- bitwise operators
 - UltraLite SQL syntax, 339
- browsing

- UltraLite table information, 76
- UltraLite table methods, 76
- bugs
 - providing feedback, xvii
- building
 - UltraLite CustDB application, 97
- bulk loading
 - UltraLite LOAD TABLE statement, 481
- business rules
 - glossary definition, 510
- BYTE_LENGTH function
 - UltraLite syntax, 360
- BYTE_SUBSTR function
 - UltraLite syntax, 360

C

- C-language programming
 - UltraLite support, 24
- cache
 - UltraLite maximum size, 8
 - UltraLite performance, 120
- cache size
 - UltraLite limit, 8
 - UltraLite usage, 199
- CACHE_SIZE connection parameter
 - UltraLite optimization for early Palm OS, 121
 - UltraLite syntax, 226
- callback
 - UltraLite schema upgrade errors, 68
- carriers
 - glossary definition, 510
- cascading deletes
 - UltraLite limitations, 2
- cascading updates
 - UltraLite limitations, 2
- case creation parameter
 - UltraLite description, 183
- case database property
 - UltraLite database creation [ulcreate] utility, 255
- CASE expression
 - UltraLite NULLIF function, 409
 - UltraLite SQL syntax, 327
- case property
 - UltraLite description, 213
- case sensitivity
 - comparison operators for UltraLite, 332
 - UltraLite case creation parameter, 183

- UltraLite case property, 213
- UltraLite strings, 304
- case sensitivity considerations
 - UltraLite about, 183
- CAST function
 - UltraLite syntax, 361
- casting
 - UltraLite data types list, 314
- catalog
 - UltraLite system tables, 289
- CE_FILE connection parameter
 - UltraLite syntax, 227
- CEILING function
 - UltraLite syntax, 362
- Certicom
 - UltraLite cryptographic module, 40
 - UltraLite modules for Palm and Windows Mobile, 59
 - UltraLite TLS-enabled synchronization, 60
- certificates
 - UltraLite application access to encryption information, 276
- CHAR data type
 - UltraLite, 312
- CHAR function
 - UltraLite syntax, 363
- CHAR_LENGTH function
 - UltraLite syntax, 365
- char_set property
 - UltraLite description, 213
- character functions
 - UltraLite alphabetical list, 352
- character set considerations
 - UltraLite, 36
- character sets
 - glossary definition, 510
 - synchronization for UltraLite, 36
 - UltraLite char_set property, 213
 - UltraLite collation creation parameter, 186
 - UltraLite databases, 37
 - UltraLite on Palm OS, 37
 - UltraLite on Windows, 37
 - UltraLite on Windows Mobile, 37
 - UltraLite strings, 304
- character strings
 - UltraLite embedded SQL, 252
- CHARINDEX function
 - UltraLite syntax, 364
- CHECK constraints
 - glossary definition, 510
 - UltraLite limitations, 2
- CHECK CONSTRAINTS clause
 - UltraLite LOAD TABLE statement, 482
- CHECKPOINT statement
 - UltraLite syntax, 458
- checkpointing
 - UltraLite CHECKPOINT syntax, 458
- checkpoints
 - glossary definition, 510
 - UltraLite performance optimization, 119
- CheckpointStore
 - UltraLite synchronization parameter, 155
- checksum_level creation parameter
 - UltraLite description, 184
- checksum_level property
 - UltraLite description, 213
- checksums
 - glossary definition, 511
 - UltraLite checksum_level creation parameter, 184
 - UltraLite checksum_level property, 213
- choosing a data management component
 - UltraLite about, 23
- choosing an index type
 - UltraLite about, 80
- choosing your programming interface
 - UltraLite about, 24
- client databases
 - UltraLite options, 154
- client message store IDs
 - glossary definition, 511
- client message stores
 - glossary definition, 511
- client/server
 - glossary definition, 511
- clients
 - UltraLite embedded engine, 23
 - UltraLite MobiLink clients, 125
- COALESCE function
 - UltraLite syntax, 366
- code pages
 - glossary definition, 511
- code points
 - UltraLite, 36
- CodeWarrior
 - UltraLite building CustDB application, 97
- COL_LENGTH function

- syntax, 353
- COL_NAME function
 - syntax, 353
- collation creation parameter
 - UltraLite description, 186
- collation property
 - UltraLite description, 213
- collation sequences
 - UltraLite about, 36
 - UltraLite changing, 36
- collations
 - glossary definition, 511
 - UltraLite collation creation parameter, 186
 - UltraLite CollationName property, 213
 - UltraLite unsupported, 270
- column compression
 - UltraLite SQL ALTER TABLE statement, 454
- column names
 - UltraLite SQL syntax, 326
- column names in expressions
 - UltraLite about, 326
- columns
 - UltraLite adding methods, 72
 - UltraLite aliases, 486
 - UltraLite ALTER TABLE statement, 454
 - UltraLite altering methods, 74
 - UltraLite altering usage, 73
 - UltraLite copying method, 77
 - UltraLite limitations, 8
- comma-separated lists
 - UltraLite LIST function syntax, 395
- command files
 - glossary definition, 511
- command line utilities
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite engine start [uleng11] utility, 260
 - UltraLite engine stop [ulstop] utility, 261
 - UltraLite erase database [ulerase] utility, 262
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite information [ulinfo] utility, 265
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite Interactive SQL [dbisql] syntax, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite Palm [ULDBUtl] utility, 258
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite synchronization [ulsync] , 275
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
- command prompts
 - conventions, xv
 - curly braces, xv
 - environment variables, xv
 - parentheses, xv
 - quotes, xv
- command shells
 - conventions, xv
 - curly braces, xv
 - environment variables, xv
 - parentheses, xv
 - quotes, xv
- comments
 - UltraLite syntax, 305
- commit flush
 - UltraLite configuration, 220
- COMMIT statement
 - UltraLite syntax, 459
- COMMIT_FLUSH connection parameter
 - UltraLite syntax, 229
- commit_flush_count database option
 - UltraLite, 220
- commit_flush_count property
 - UltraLite description, 213
- commit_flush_timeout database option
 - UltraLite, 221
- commit_flush_timeout property
 - UltraLite description, 213
- committing
 - UltraLite COMMIT syntax, 459
 - UltraLite database rows, 14
 - UltraLite transaction overview , 16
- communication streams
 - glossary definition, 512
- comparing
 - UltraLite and SQL Anywhere databases, 2
- comparing UltraLite and SQL Anywhere
 - about, 2
- comparison operators
 - dynamic SQL syntax for UltraLite, 332
 - UltraLite SQL, 332
- compatibility
 - UltraLite SQL, 332
- compressed columns

- UltraLite ALTER TABLE statement, 454
- computed columns
 - UltraLite limitations, 2
- COMPUTES clause
 - UltraLite LOAD TABLE statement, 482
- CON connection parameter
 - UltraLite syntax, 230
- concatenating strings
 - string operators for UltraLite, 339
- concurrency
 - glossary definition, 512
 - UltraLite issues, 13
 - UltraLite synchronization, 13
- concurrent access
 - UltraLite engine, 23
- conditions
 - ALL conditions for UltraLite SQL, 334
 - ANY for UltraLite SQL, 335
 - BETWEEN for UltraLite SQL, 336
 - EXISTS for UltraLite SQL, 336
 - IN for UltraLite SQL, 337
 - searching in UltraLite SQL, 331
- conduit
 - deploying UltraLite applications, 146
 - installing, 263
 - installing for CustDB, 263
 - UltraLite HotSync, 62
- conduit files
 - deploying HotSync, 146
- conduits
 - UltraLite HotSync conduit files, 62
- conflict resolution
 - glossary definition, 512
- conn_count property
 - UltraLite description, 213
- connecting
 - MobiLink UltraLite Stream Type synchronization parameter, 171
 - UltraLite database troubleshooting, 49
 - UltraLite databases, 51
- connecting to an UltraLite database
 - about, 43
- connection failures
 - UltraLite troubleshooting, 49
- connection IDs
 - glossary definition, 512
- connection methods
 - UltraLite about, 43
 - connection parameters
 - alphabetical list (UltraLite), 225
 - DBN for UltraLite, 234
 - UltraLite, 43
 - UltraLite CACHE_SIZE, 226
 - UltraLite CE_FILE, 227
 - UltraLite choosing between, 47
 - UltraLite COMMIT_FLUSH, 229
 - UltraLite CON, 230
 - UltraLite connection summary, 48
 - UltraLite DBF, 231
 - UltraLite DBKEY, 233
 - UltraLite file_name, 231
 - UltraLite key, 233
 - UltraLite list of, 44
 - UltraLite MIRROR_FILE, 235
 - UltraLite NT_FILE, 237
 - UltraLite overview, 48
 - UltraLite PALM_ALLOW_BACKUP, 240
 - UltraLite PALM_FILE, 241
 - UltraLite password, 243
 - UltraLite precedence of, 49
 - UltraLite PWD, 243
 - UltraLite RESERVE_SIZE , 244
 - UltraLite START , 245
 - UltraLite supplying, 45
 - UltraLite troubleshooting transmission of, 49
 - UltraLite UID , 246
 - UltraLite userid , 246
 - connection profiles
 - glossary definition, 512
 - connection strings
 - UltraLite connection parameters, 225
 - UltraLite parameters overview, 48
 - UltraLite setting , 49
 - connection-initiated synchronization
 - glossary definition, 512
 - connections
 - UltraLite concurrency, 13
 - UltraLite conn_count property, 213
 - UltraLite overview, 43
 - UltraLite troubleshooting, 499
 - consolidated databases
 - glossary definition, 512
 - UltraLite choosing, 126
 - UltraLite compatibility, 127
 - UltraLite sample, 104
 - constants

- UltraLite SQL syntax, 326
- constants in expressions
 - UltraLite about, 326
- constraints
 - glossary definition, 513
 - UltraLite ALTER TABLE statement, 454
 - UltraLite referential integrity, 138
 - UltraLite renaming, 454
- contention
 - glossary definition, 513
- controlling synchronization
 - UltraLite publications, 137
- conventions
 - command prompts, xv
 - command shells, xv
 - documentation, xiv
 - file names in documentation, xiv
- conversion functions
 - UltraLite alphabetical list, 348
- CONVERT function
 - UltraLite syntax, 366
- converting
 - UltraLite ambiguous dates, 196
 - UltraLite data types list, 314
- converting data types
 - UltraLite about, 314
- converting strings
 - UltraLite about, 352
- copying
 - UltraLite table method, 77
- copying data
 - UltraLite databases, 77
- correlation names
 - glossary definition, 513
- COS function
 - UltraLite syntax, 369
- cosine function
 - UltraLite COS function, 369
- COT function
 - UltraLite syntax, 369
- cotangent function
 - UltraLite COT function, 369
- COUNT function
 - UltraLite syntax, 370
- count operation
 - UltraLite execution plans, 343
- CPU
 - UltraLite limits, 8
- create database wizard
 - UltraLite usage, 30
- CREATE INDEX statement
 - UltraLite Interactive SQL example, 81
 - UltraLite syntax, 460
 - UNIQUE parameter, 460
- create index wizard
 - UltraLite using, 81
- CREATE PUBLICATION statement
 - UltraLite Interactive SQL example, 84
 - UltraLite Interactive SQL subset example, 85
 - UltraLite Interactive SQL whole table example, 84
 - UltraLite syntax, 462
 - UltraLite usage, 137
- create publication wizard
 - UltraLite rows publishing, 84
 - UltraLite usage, 83
- CREATE SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 464
- CREATE TABLE statement
 - UltraLite Interactive SQL example, 71
 - UltraLite syntax , 468
- create table wizard
 - UltraLite usage, 71
- create user wizard
 - UltraLite usage, 87
- creating
 - reference databases for UltraLite, 31
 - UltraLite CREATE PUBLICATION statement, 462
 - UltraLite CREATE TABLE statement, 468
 - UltraLite databases from a MobiLink sync model, 31
 - UltraLite databases from the command prompt, 30
 - UltraLite databases from XML, 32
 - UltraLite databases overview of methods, 29
 - UltraLite databases with ulcreate, 255
 - UltraLite databases with ulinit, 269
 - UltraLite indexes, 81
 - UltraLite publications, 137
 - UltraLite publications tables, 83
 - UltraLite remote databases, 128
 - UltraLite table methods, 71
 - UltraLite users, 87
- creating databases
 - Sybase Central UltraLite plug-in, 30
 - UltraLite creation parameters, 181
 - UltraLite remote , 128

- creating publications
 - UltraLite databases with MobiLink applications, 137
 - creating remote databases
 - UltraLite clients, 128
 - creating UltraLite databases
 - about, 29
 - creating UltraLite tables
 - about, 71
 - creation parameters (UltraLite)
 - about, 181
 - case, 183
 - checksum_level, 184
 - collation, 186
 - date_format, 187
 - date_order, 190
 - fips creation parameter, 192
 - max_hash_size, 194
 - nearest_century, 196
 - obfuscate, 198
 - page_size, 199
 - precision, 201
 - scale , 203
 - time_format, 205
 - timestamp_format, 207
 - timestamp_increment, 210
 - utf8_encoding, 211
 - creator ID
 - glossary definition, 513
 - CROSS JOIN clause
 - UltraLite syntax, 478
 - Crossfire
 - UltraLite support, 24
 - cryptography
 - UltraLite Certicom module, 40
 - CSV files
 - UltraLite importing, 502
 - CURRENT DATE
 - UltraLite special value, 308
 - current date function
 - UltraLite TODAY function, 435
 - current row
 - UltraLite concurrency, 13
 - CURRENT TIME
 - UltraLite special value, 308
 - CURRENT TIMESTAMP
 - UltraLite special value, 309
 - UltraLite special value for, 2
 - cursor positions
 - glossary definition, 513
 - cursor result sets
 - glossary definition, 513
 - cursors
 - glossary definition, 513
 - UltraLite current row, 13
 - UltraLite dirty reads, 18
 - CustDB
 - UltraLite application readme files, 96
 - UltraLite building application, 97
 - UltraLite database synchronization, 102
 - UltraLite limitations on instances of, 93
 - UltraLite running application, 97
 - UltraLite running multiple instances, 93
 - UltraLite sample , 93
 - UltraLite sample file locations, 95
 - UltraLite starting, 99
 - custdb.db
 - location of, 95
 - custdb.sql
 - calling synchronization scripts, 104
 - custdb.udb
 - UltraLite location of, 95
 - cycles
 - UltraLite foreign key issues, 138
 - UltraLite foreign keys, 138
- ## D
- data
 - UltraLite loading, 272
 - UltraLite selecting rows, 486
 - UltraLite unloading, 281
 - UltraLite viewing methods, 76
 - data consistency
 - UltraLite dirty reads, 18
 - data cube
 - glossary definition, 514
 - data management
 - UltraLite components available for, 23
 - UltraLite description, 10
 - UltraLite state information, 13
 - Data Manager
 - UltraLite database storage, 47
 - data manipulation language
 - glossary definition, 514
 - data sources

- UltraLite CustDB tutorial example, 102
- data type conversion functions
 - UltraLite about, 348
- data types
 - alphabetical list of UltraLite data types, 312
 - glossary definition, 514
 - UltraLite alias equivalents, 314
 - UltraLite casting of values, 314
 - UltraLite limitations, 8
 - UltraLite SQL conversion functions, 348
- data types in UltraLite
 - about, 312
- database administrator
 - glossary definition, 514
- database connections
 - glossary definition, 515
- database encryption
 - UltraLite performance impact of, 120
- database engines
 - UltraLite about, 23
- database files
 - (*see also* UltraLite databases)
 - glossary definition, 515
 - UltraLite connection parameters, 47
 - UltraLite encrypting, 233
 - UltraLite maximum size, 8
- database management
 - UltraLite layers of, 10
- database names
 - glossary definition, 515
- database objects
 - glossary definition, 515
 - UltraLite copying method, 77
- database options
 - (*see also* database options (UltraLite))
 - UltraLite, 219
- database options (UltraLite)
 - browsing, 224
 - commit_flush_count, 220
 - commit_flush_timeout, 221
 - DB_PROPERTY function , 380
 - global_database_id, 222
 - ml_remote_id, 223
 - SET OPTION statement, 488
- database options window
 - accessing, 224
- database owner
 - glossary definition, 515
- database page size considerations
 - UltraLite about, 199
- database properties
 - (*see also* database properties (UltraLite))
 - UltraLite, 213
 - UltraLite alphabetical list, 213
- database properties (UltraLite)
 - browsing, 218
 - creation parameters, 34
 - fiops usage, 40
 - obfuscate usage, 40
 - utf8_encoding usage, 37
- database schemas
 - UltraLite system tables, 289
- database servers
 - glossary definition, 515
- database sizes
 - UltraLite limit, 8
- database utilities
 - UltraLite database connections, 49
- database validation
 - UltraLite, 15
- databases
 - (*see also* UltraLite databases)
 - comparing UltraLite and SQL Anywhere, 2
 - creating with UltraLite plug-in, 30
 - glossary definition, 514
 - UltraLite loading bulk data into, 481
- DATALENGTH function
 - UltraLite syntax, 371
- date considerations
 - UltraLite about, 187
- DATE data type
 - UltraLite, 312
- DATE function
 - UltraLite syntax, 372
- date functions
 - UltraLite alphabetical list, 348
- date parts
 - about, 349
 - available in UltraLite, 187
- date_format creation parameter
 - UltraLite description, 187
- date_format property
 - UltraLite description, 213
- date_order creation parameter
 - UltraLite description, 190
- date_order property

- UltraLite description, 213
- DATEADD function
 - UltraLite syntax, 372
- DATEDIFF function
 - UltraLite syntax, 373
- DATEFORMAT function
 - UltraLite syntax, 375
- DATENAME function
 - UltraLite syntax, 375
- DATEPART function
 - UltraLite syntax, 376
- dates
 - UltraLite ambiguous string conversions , 196
 - UltraLite conversion functions, 348
 - UltraLite date_format property, 213
 - UltraLite date_order property, 213
 - UltraLite formatting, 187
 - UltraLite ordering, 190
 - UltraLite rollover point, 196
- datetime
 - UltraLite conversion functions, 348
- DATETIME function
 - UltraLite syntax, 377
- DAY function
 - UltraLite syntax, 377
- day of week
 - UltraLite DOW function, 382
- DAYNAME function
 - UltraLite syntax, 378
- DAYS function
 - UltraLite syntax, 379
- DB_PROPERTY function
 - UltraLite syntax, 380
- DBA authority
 - glossary definition, 515
- DBF connection parameter
 - UltraLite syntax, 231
- dbhsync11.dll
 - HotSync conduit in UltraLite, 146
- dbisql utility
 - UltraLite exit codes, 251
 - UltraLite syntax, 249
 - UltraLite troubleshooting data imports, 502
- dbisql.exe
 - shutting down before installing, 499
- DBKEY connection parameter
 - UltraLite syntax, 233
- dbngen11.dll
 - ActiveSync conduit deployment in UltraLite, 64
 - HotSync conduit deployment in UltraLite, 146
- dbmlhttp11.dll
 - deploying UltraLite applications, 146
- dbmlhttps11.dll
 - deploying UltraLite applications, 146
- dbmlsock11.dll
 - deploying UltraLite applications, 146
- dbmltls11.dll
 - deploying UltraLite applications, 146
- DBMS
 - glossary definition, 515
- DBN connection parameter
 - UltraLite syntax, 234
- dbser11.dll
 - deploying UltraLite applications, 146
- dbspaces
 - glossary definition, 516
- DCX
 - about, xii
- DDL
 - glossary definition, 514
- DDL statements
 - UltraLite schema changes with, 11
- deadlocks
 - glossary definition, 516
 - UltraLite prevention of, 16
- DECIMAL data type
 - UltraLite, 312
- decimal point position considerations
 - UltraLite precision, 201
 - UltraLite scale creation parameter, 203
- decimal precision
 - UltraLite precision, 201
- decimals
 - used in UltraLite, 306
- declaring default global autoincrement columns
 - UltraLite clients in MobiLink systems, 131
- DEFAULT TIMESTAMP columns
 - UltraLite syntax, 468
- default values
 - UltraLite CURRENT DATE, 308
 - UltraLite CURRENT TIME, 308
 - UltraLite CURRENT TIMESTAMP, 309
 - UltraLite SQLCODE, 309
- defaults
 - UltraLite autoincrement, 468
- DEFAULTS clause

- UltraLite LOAD TABLE statement, 482
- DEGREES function
 - syntax, 381
- delaying commits
 - performance enhancements, 119
- DELETE statement
 - UltraLite dynamic SQL syntax, 473
- deleting
 - UltraLite columns, 454
 - UltraLite databases, 14
 - UltraLite indexes, 82
 - UltraLite publications, 85
 - UltraLite START SYNCHRONIZATION DELETE statement, 489
 - UltraLite table methods, 75
 - UltraLite TRUNCATE TABLE statement, 493
 - UltraLite users, 88
 - UltraLite utility to delete databases, 258
- deleting data
 - UltraLite file size impact of, 501
- deleting UltraLite users
 - about, 88
- deleting users
 - UltraLite, 88
- DELIMITED BY clause
 - UltraLite LOAD TABLE statement, 482
- demos
 - (*see also* tutorials)
- deploying
 - applications that use ActiveSync for UltraLite clients, 150
 - HotSync conduit files, 146
 - UltraLite ActiveSync provider files, 64
 - UltraLite engine troubleshooting, 505
 - UltraLite FIPS-enabled applications, 192
 - UltraLite schema upgrades, 67
 - UltraLite to devices, 55
- derived tables
 - subqueries for UltraLite SQL, 329
 - UltraLite FROM clause, 478
 - UltraLite SQL, 329
- designing
 - UltraLite implementation , 21
- Designing synchronization in UltraLite
 - about, 133
- desktop creation
 - UltraLite about, 30
- developer community
 - newsgroups, xvii
- development platforms
 - UltraLite support for, 24
- device tracking
 - glossary definition, 516
- devices
 - UltraLite deployment techniques, 55
 - UltraLite multiple connection parameters for, 47
- DIFFERENCE function
 - UltraLite syntax, 381
- digits
 - UltraLite maximum number, 201
- direct page scans
 - UltraLite about, 117
- direct row handling
 - glossary definition, 516
- dirty reads
 - UltraLite isolation levels, 18
- disable concurrency
 - UltraLite synchronization parameter overview, 13
- DisableConcurrency
 - UltraLite synchronization parameter, 155
- DISTINCT keyword
 - UltraLite SQL, 486
- distinct operation
 - UltraLite execution plans, 343
- DML
 - glossary definition, 514
- DocCommentXchange (DCX)
 - about, xii
- documentation
 - conventions, xiv
 - SQL Anywhere, xii
- domains
 - glossary definition, 516
- DOUBLE data type
 - UltraLite, 312
- double hyphen
 - UltraLite comment indicator, 305
- double slash
 - UltraLite comment indicator, 305
- DOW function
 - UltraLite syntax, 382
- download acknowledgements
 - UltraLite send_download_ack synchronization parameter, 169
- download only
 - UltraLite synchronization parameter, 159

- download only synchronization
 - UltraLite download_only synchronization parameter , 159
- download-only synchronization
 - UltraLite defining overview, 139
 - UltraLite synchronization parameter, 159
- download_only synchronization parameter
 - UltraLite reference, 159
- downloads
 - glossary definition, 516
- DROP INDEX statement
 - UltraLite Interactive SQL example, 82
 - UltraLite syntax, 474
- DROP PUBLICATION statement
 - UltraLite Interactive SQL example, 86
 - UltraLite syntax, 475
- DROP SYNCHRONIZATION PROFILE statement
 - UltraLite syntax, 476
- DROP TABLE statement
 - UltraLite Interactive SQL example, 75
 - UltraLite syntax, 477
- dropping
 - UltraLite columns, 454
 - UltraLite DROP SYNCHRONIZATION PROFILE statement, 476
 - UltraLite indexes, 82
 - UltraLite SQL CREATE INDEX statement, 460
 - UltraLite SQL DROP INDEX statement, 474
 - UltraLite SQL DROP PUBLICATION statement, 475
 - UltraLite SQL DROP TABLE statement, 477
 - UltraLite table methods, 75
- dropping an index
 - UltraLite about, 82
- dropping publications
 - UltraLite clients, 85
- dropping UltraLite tables
 - about, 75
- dummy operation
 - UltraLite execution plans, 343
- dynamic SQL
 - arithmetic operators for UltraLite, 338
 - bitwise operators for UltraLite, 339
 - glossary definition, 516
 - logical operators for UltraLite, 333
 - operator precedence for UltraLite, 340
 - string operators for UltraLite, 339

E

- EBFs
 - glossary definition, 516
- editing
 - UltraLite table methods, 76
- ELSE
 - CASE expression for UltraLite, 327
 - UltraLite IF expressions, 327
- embedded SQL
 - glossary definition, 517
 - NULL values in UltraLite, 307
- eMbedded Visual C++
 - UltraLite CustDB sample and readme files, 97
- empty databases
 - UltraLite populating after running ulinit, 270
- encoding
 - glossary definition, 517
 - UltraLite utf8_encoding creation parameter, 211
 - UltraLite utf8_encoding usage, 37
- ENCODING clause
 - UltraLite LOAD TABLE statement, 482
- encryption
 - UltraLite changing key, 192
 - UltraLite data deployment considerations, 59
 - UltraLite deployment steps, 59
 - UltraLite encryption keys, 233
 - UltraLite encryption property, 213
 - UltraLite fips creation parameter, 192
 - UltraLite fips property usage, 40
 - UltraLite fips usage, 40
 - UltraLite obfuscate creation parameter, 198
 - UltraLite obfuscate property usage, 40
 - UltraLite performance impact of, 120
 - UltraLite synchronization deployment considerations, 60
 - UltraLite TLS synchronization configuration, 60
- encryption keys
 - UltraLite changing, 192
- encryption property
 - UltraLite description, 213
- END
 - CASE expression for UltraLite, 327
- ENDIF
 - UltraLite IF expressions , 327
- engines
 - (*see also* UltraLite engine)
- entity-relationship diagrams

- UltraLite about, 78
- entity-relationship tab
 - UltraLite using, 78
- environment variables
 - command prompts, xv
 - command shells, xv
 - ERRORLEVEL for UltraLite, 251
 - UltraLite ULSQLCONNECT, 53
 - UltraLite usage, 49
- ER diagram tab
 - UltraLite about, 78
- erasing
 - UltraLite table methods, 75
- error callback
 - UltraLite schema upgrade errors, 68
- error codes
 - UltraLite database creation [ulcreate] utility, 248
 - UltraLite database synchronization [ulsync] utility, 248
 - UltraLite load XML to database [ulload] utility, 248
 - UltraLite unload data to XML [ulunload] utility, 248
 - UltraLite utilities, 248
- ERRORLEVEL environment variable
 - Interactive SQL return code for UltraLite, 251
- errors
 - UltraLite client error -764, 505
 - UltraLite SQLE_DATABASE_ERROR, 500
 - UltraLite SQLE_DEVICE_ERROR, 500
 - UltraLite SQLE_MEMORY_ERROR, 500
- escape sequences
 - UltraLite engine paths, 505
- ESCAPES clause
 - UltraLite LOAD TABLE statement, 482
- ESQL (*see* embedded SQL) (*see* UltraLite SQL)
- event model
 - glossary definition, 517
- event notifications
 - UltraLite working with, 89
- examples
 - (*see also* samples)
 - (*see also* tutorials)
- exclusive OR
 - bitwise operator for UltraLite, 339
- execution plans
 - UltraLite checking for index usage, 110
 - UltraLite how to read, 343
 - UltraLite operations, 343
 - UltraLite overriding, 342
 - UltraLite text of, 342
 - UltraLite working with, 342
- EXISTS conditions
 - UltraLite SQL, 336
- exit codes
 - Interactive SQL [dbisql] utility for UltraLite, 251
 - UltraLite database creation [ulcreate] utility, 248
 - UltraLite database synchronization [ulsync] utility, 248
 - UltraLite load XML to database [ulload] utility, 248
 - UltraLite unload data to XML [ulunload] utility, 248
- EXP function
 - UltraLite syntax, 383
- expansion cards
 - UltraLite writing to, 47
- EXPLANATION function
 - UltraLite syntax, 383
- exponential function
 - UltraLite EXP function, 383
- exponents
 - in UltraLite, 306
- export tools
 - UltraLite ulunload utility, 281
- exporting
 - UltraLite databases with ulunload, 281
- expressions
 - aggregate for UltraLite, 328
 - CASE expressions for UltraLite, 327
 - SQL operator precedence for UltraLite, 340
 - subqueries for UltraLite SQL, 329
 - UltraLite column names, 326
 - UltraLite constants, 326
 - UltraLite IF expressions, 327
 - UltraLite input parameters, 330
 - UltraLite SQL, 325
- expressions in UltraLite
 - about, 325
- extensions
 - UltraLite disallowed, 258
- external logins
 - glossary definition, 517
- extraction
 - glossary definition, 517

F

- failover
 - glossary definition, 517
- failures
 - UltraLite preventing out-of-memory errors, 244
 - UltraLite schema upgrade errors, 68
- features
 - UltraLite comparison list, 2
- Federal Information Processing Standards (*see* FIPS)
- feedback
 - documentation, xvii
 - providing, xvii
 - reporting an error, xvii
 - requesting an update, xvii
- fetches
 - UltraLite, 18
- fetching rows
 - UltraLite concurrency, 18
- FILE
 - glossary definition, 517
- FILE message type
 - glossary definition, 517
- file names
 - UltraLite connection parameters, 47
- file objects
 - UltraLite types, 11
- file property
 - UltraLite description, 213
- file size
 - UltraLite database troubleshooting, 501
- file systems
 - (*see also* VFS)
- file-based downloads
 - glossary definition, 518
- file-definition database
 - glossary definition, 518
- files
 - HotSync conduit, 146
 - transferring with MLFileTransfer, 142
 - UltraLite ActiveSync provider, 64
 - UltraLite CustDB sample location, 95
- filter operation
 - UltraLite execution plans, 343
- filtering
 - UltraLite table methods, 77
- filters
 - UltraLite filenames, 258
- finding out more and requesting technical assistance
 - technical support, xvii
- FIPS
 - UltraLite encrypted database deployment, 59
 - UltraLite fips property usage, 40
 - UltraLite setup and deployment, 192
 - UltraLite TLS-enabled synchronization, 60
- FIPS creation parameter
 - UltraLite description, 192
- fips database property
 - UltraLite deployment steps, 59
 - UltraLite usage, 40
- fips network protocol option
 - UltraLite TLS-enabled synchronization, 60
- FIRST clause
 - UltraLite SELECT statement, 486
- FLOAT data type
 - UltraLite, 312
- FLOOR function
 - UltraLite syntax, 384
- flush count
 - UltraLite formatting, 220
- flush timeout
 - UltraLite formatting, 221
- flushing
 - UltraLite databases, 220
- footprint
 - UltraLite databases, 2
- FOR clause
 - UltraLite SELECT statement, 487
- FOR READ ONLY clause
 - UltraLite direct page scans with, 117
- FORCE ORDER clause
 - UltraLite SELECT statement, 487
- foreign key constraints
 - glossary definition, 518
- foreign key cycles
 - UltraLite about, 138
 - UltraLite issues, 138
- foreign keys
 - glossary definition, 518
 - UltraLite characteristics, 108
 - UltraLite copying method, 77
 - UltraLite foreign keys, 468
 - UltraLite of unnamed, 468
- foreign tables
 - glossary definition, 518
- FORMAT clause

- UltraLite LOAD TABLE statement, 482
- format options (UltraLite)
 - utf8_encoding usage, 37
- FROM clause
 - UltraLite LOAD TABLE statement, 481
 - UltraLite SELECT statement, 486
 - UltraLite syntax, 478
- full backups
 - glossary definition, 518
- function syntax
 - alphabetical list of UltraLite functions A-D, 354
 - alphabetical list of UltraLite functions E-O, 383
 - alphabetical list of UltraLite functions P-Z, 411
- functions
 - alphabetical list of UltraLite functions A-D, 354
 - alphabetical list of UltraLite functions E-O, 383
 - alphabetical list of UltraLite functions P-Z, 411
 - return NULL if you specify NULL argument, 347
 - types of function for UltraLite, 348
 - UltraLite aggregate, 348
 - UltraLite data type conversion SQL, 348
 - UltraLite date and time, 348
 - UltraLite introduction, 347
 - UltraLite miscellaneous, 350
 - UltraLite numeric, 351
 - UltraLite string, 352
 - UltraLite system SQL, 353
- functions, aggregate
 - about, 348
 - UltraLite AVG, 359
 - UltraLite COUNT, 370
 - UltraLite LIST, 395
 - UltraLite MAX, 400
 - UltraLite MIN, 401
 - UltraLite SUM, 432
- functions, data type conversion
 - UltraLite about, 348
 - UltraLite CAST, 361
 - UltraLite CONVERT, 366
 - UltraLite HEXTOINT, 386
 - UltraLite INTTOHEX, 390
 - UltraLite ISDATE, 391
 - UltraLite ISNULL, 392
- functions, date and time
 - UltraLite about, 348
 - UltraLite DATE, 372
 - UltraLite DATEADD, 372
 - UltraLite DATEDIFF, 373
 - UltraLite DATEFORMAT, 375
 - UltraLite DATENAME, 375
 - UltraLite DATEPART, 376
 - UltraLite DATETIME, 377
 - UltraLite DAY, 377
 - UltraLite DAYNAME, 378
 - UltraLite DAYS, 379
 - UltraLite DOW, 382
 - UltraLite GETDATE, 385
 - UltraLite HOUR, 387
 - UltraLite HOURS, 388
 - UltraLite MINUTE, 401
 - UltraLite MINUTES, 402
 - UltraLite MONTH, 405
 - UltraLite MONTHNAME, 406
 - UltraLite MONTHS, 406
 - UltraLite NOW, 408
 - UltraLite QUARTER, 413
 - UltraLite SECOND, 420
 - UltraLite SECONDS, 421
 - UltraLite TODAY, 435
 - UltraLite WEEKS, 442
 - UltraLite YEAR, 444
 - UltraLite YEARS, 444
 - UltraLite YMD, 445
- functions, miscellaneous
 - UltraLite about, 350
 - UltraLite ARGN, 355
 - UltraLite COALESCE, 366
 - UltraLite EXPLANATION, 383
 - UltraLite GREATER, 385
 - UltraLite IFNULL, 389
 - UltraLite LESSER, 395
 - UltraLite NEWID, 407
 - UltraLite NULLIF, 409
 - UltraLite SHORT_PLAN, 422
- functions, numeric
 - DEGREES, 381
 - UltraLite about, 351
 - UltraLite ABS, 354
 - UltraLite ACOS, 354
 - UltraLite ASIN, 357
 - UltraLite ATAN, 357
 - UltraLite ATAN2, 358
 - UltraLite CEILING, 362
 - UltraLite COS, 369
 - UltraLite COT, 369
 - UltraLite EXP, 383

- UltraLite FLOOR, 384
 - UltraLite LOG, 397
 - UltraLite LOG10, 398
 - UltraLite MOD, 404
 - UltraLite PI, 412
 - UltraLite POWER, 412
 - UltraLite RADIANS, 414
 - UltraLite REMAINDER, 414
 - UltraLite ROUND, 419
 - UltraLite SIGN, 423
 - UltraLite SIN, 424
 - UltraLite SQRT, 427
 - UltraLite TAN, 435
 - UltraLite TRUNCNUM, 437
 - functions, string
 - UltraLite about, 352
 - UltraLite ASCII, 356
 - UltraLite BYTE_LENGTH, 360
 - UltraLite BYTE_SUBSTR, 360
 - UltraLite CHAR, 363
 - UltraLite CHAR_LENGTH, 365
 - UltraLite CHARINDEX, 364
 - UltraLite DIFFERENCE, 381
 - UltraLite INSERTSTR, 390
 - UltraLite LCASE, 392
 - UltraLite LEFT, 393
 - UltraLite LENGTH, 394
 - UltraLite LOCATE, 396
 - UltraLite LOWER, 398
 - UltraLite LTRIM, 399
 - UltraLite PATINDEX, 411
 - UltraLite REPEAT, 415
 - UltraLite REPLACE, 416
 - UltraLite REPLICATE, 417
 - UltraLite RIGHT, 418
 - UltraLite RTRIM, 419
 - UltraLite SIMILAR, 424
 - UltraLite SOUNDEX, 425
 - UltraLite SPACE, 426
 - UltraLite STR, 427
 - UltraLite STRING, 428
 - UltraLite STRTOUUID, 429
 - UltraLite STUFF, 430
 - UltraLite SUBSTRING, 430
 - UltraLite TRIM, 436
 - UltraLite UCASE, 438
 - UltraLite UPPER, 439
 - UltraLite UUIDTOSTR, 441
 - functions, system
 - DB_PROPERTY, 380
 - ML_GET_SERVER_NOTIFICATION, 403
 - SUSER_ID, 433
 - SUSER_NAME, 433
 - SYNC_PROFILE_OPTION_VALUE, 434
 - TSEQUAL, 437
 - UltraLite DATALENGTH, 371
 - USER_ID, 440
 - USER_NAME, 441
 - fundamentals
 - UltraLite database management, 13
- ## G
- gateways
 - glossary definition, 518
 - generated join conditions
 - glossary definition, 519
 - generation numbers
 - glossary definition, 519
 - GETDATE function
 - UltraLite syntax, 385
 - GetLastIdentity method
 - UltraLite synchronization, 131
 - getScriptVersion method
 - UltraLite example, 176
 - getStream method
 - UltraLite example, 171
 - getting help
 - technical support, xvii
 - getUploadOK method
 - UltraLite example, 173
 - global autoincrement
 - exhausted range in UltraLite, 129
 - UltraLite clients in MobiLink systems , 129
 - UltraLite global_database_id, 222
 - UltraLite, setting, 129
 - UltraLite, setting defaults, 131
 - global database ID considerations
 - UltraLite about, 222
 - global database identifier
 - UltraLite global_database_id, 222
 - UltraLite, setting, 129
 - global temporary tables
 - glossary definition, 519
 - global_database_id option
 - UltraLite, 222

- UltraLite CREATE TABLE statement, 468
- UltraLite, setting, 129
- global_database_id property
 - UltraLite description, 213
- globally unique identifiers
 - UltraLite clients in MobiLink systems, 129
 - UltraLite SQL syntax for NEWID function, 407
- glossary
 - list of SQL Anywhere terminology, 509
- grant options
 - glossary definition, 519
- graphical plans
 - UltraLite not supported, 342
- GREATER function
 - UltraLite syntax, 385
- GROUP BY clause
 - UltraLite SELECT statement, 486
- group-by operation
 - UltraLite execution plans, 343
- GUIDs
 - UltraLite clients in MobiLink systems, 129
 - UltraLite SQL syntax for NEWID function, 407
 - UltraLite SQL syntax for STRTOUUID function, 429
 - UltraLite SQL syntax for UUIDTOSTR function, 441

H

- hash
 - glossary definition, 519
 - UltraLite configuring size for, 116
 - UltraLite max_hash_size property, 213
 - UltraLite optimal size for, 113
 - UltraLite size considerations, 194
- hashing
 - UltraLite indexes, 194
- help
 - technical support, xvii
- hexadecimal strings
 - UltraLite about, 386
- HEXTOINT function
 - UltraLite syntax, 386
- histograms
 - glossary definition, 519
- host name
 - UltraLite ULSynchronize arguments, 172
- host platforms

- UltraLite Windows supported platforms, 24
- HotSync
 - deploying conduit files, 146
- HotSync conduit
 - installing, 263
 - installing for CustDB, 263
- HotSync configuration overview
 - UltraLite clients, 62
- HotSync synchronization
 - configuring, 148
 - UltraLite architecture , 146
 - UltraLite clients, 147
 - UltraLite Palm OS , 62
- HOUR function
 - UltraLite syntax, 387
- HOURS function
 - UltraLite syntax, 388

I

- iAnywhere developer community
 - newsgroups, xvii
- iAnywhere JDBC driver
 - glossary definition, 520
- icons
 - used in this Help, xvi
- identifiers
 - glossary definition, 520
 - UltraLite SQL, 303
- IDs
 - ml_remote_id option, 223
 - UltraLite global database, 222
 - UltraLite user, 51
- IF expressions
 - UltraLite SQL syntax, 327
- IFNULL function
 - UltraLite syntax, 389
- ignored rows
 - UltraLite synchronization parameter, 160
- ignored_rows synchronization parameter
 - UltraLite reference, 160
- importing data
 - UltraLite troubleshooting, 502
- importing data into databases
 - UltraLite ulload utility, 272
- IN conditions
 - UltraLite SQL, 337
- in-process runtime

- UltraLite about, 23
- incremental backups
 - glossary definition, 520
- index performance considerations
 - UltraLite about, 194
- index scans
 - UltraLite, 108
- index-based UltraLite optimizations
 - UltraLite about, 108
- index-scan operation
 - UltraLite execution plans, 343
- INDEX_COL function
 - syntax, 353
- indexes
 - glossary definition, 520
 - UltraLite bypassing use of, 117
 - UltraLite copying method, 77
 - UltraLite creating, 81
 - UltraLite deleting, 82
 - UltraLite determining which index is scanned, 110
 - UltraLite hash considerations, 194
 - UltraLite hash value, 194
 - UltraLite introduction of, 108
 - UltraLite limitations, 8
 - UltraLite non-unique index characteristics, 80
 - UltraLite page_size usage, 199
 - UltraLite performance enhancements, 111
 - UltraLite primary keys, 32
 - UltraLite sysindex system table, 293
 - UltraLite sysixcol system table, 295
 - UltraLite types, 80
 - UltraLite unique index characteristics, 80
 - UltraLite unique key characteristics, 80
 - UltraLite UNIQUE SQL parameter, 460
 - UltraLite when to create, 81
 - UltraLite when to use, 79
 - UltraLite working with, 79
- indicators
 - UltraLite comments in SQL block, 305
- InfoMaker
 - glossary definition, 520
- initializing
 - UltraLite databases with ulinit, 269
- initializing databases
 - Sybase Central UltraLite plug-in, 30
- INNER JOIN clause
 - UltraLite syntax, 478
- inner joins
 - glossary definition, 520
- inner references
 - subqueries for UltraLite SQL, 329
- input parameters
 - UltraLite about, 330
- INPUT statement
 - UltraLite troubleshooting, 502
- INSERT statement
 - UltraLite Interactive SQL example, 77
 - UltraLite syntax, 480
- inserting
 - UltraLite data using LOAD TABLE statement, 481
 - UltraLite INSERT statement, 480
 - UltraLite rows in bulk, 481
- INSERTSTR function
 - UltraLite syntax, 390
- install-dir
 - documentation usage, xiv
- installation
 - UltraLite troubleshooting, 499
- installing
 - UltraLite on devices, 55
- INTEGER data type
 - UltraLite, 312
- integrated logins
 - glossary definition, 520
- integrity
 - glossary definition, 521
 - UltraLite CREATE TABLE statement, 468
- integrity constraints
 - UltraLite usage, 468
- Interactive SQL
 - alphabetical list of UltraLite functions A-D, 354
 - alphabetical list of UltraLite functions E-O, 383
 - alphabetical list of UltraLite functions P-Z, 411
 - glossary definition, 521
 - UltraLite command line, 249
 - UltraLite displaying plans, 342
 - UltraLite plan interpretation, 343
 - UltraLite plan operations, 343
 - UltraLite text plans, 342
 - UltraLite troubleshooting data imports, 502
- Interactive SQL utility [dbisql]
 - UltraLite exit codes, 251
 - UltraLite syntax, 249
- INTTOHEX function
 - UltraLite syntax, 390
- IS

- logical operators for UltraLite, 333
- ISDATE function
 - UltraLite syntax, 391
- ISNULL function
 - UltraLite syntax, 392
- isolation levels
 - glossary definition, 521
 - UltraLite about, 18
 - UltraLite dirty reads, 18
 - UltraLite isolation levels, 18

J

- JAR files
 - glossary definition, 521
- Java classes
 - glossary definition, 521
- jConnect
 - glossary definition, 521
- JDBC
 - glossary definition, 521
- join conditions
 - glossary definition, 522
- join operation
 - UltraLite execution plans, 343
- join types
 - glossary definition, 522
- joins
 - glossary definition, 522
 - UltraLite FROM clause syntax, 478

K

- keep partial download synchronization parameter
 - UltraLite reference , 161
- key connection parameter
 - UltraLite syntax, 233
- KEY JOIN clause
 - UltraLite syntax, 478
- key joins
 - glossary definition, 519
- keys
 - registry for ulcond11 cache size, 263
 - UltraLite index creation from, 80
 - UltraLite index hash, 194
 - UltraLite primary, 71
- keyset operation
 - UltraLite execution plans, 343
- keywords

- UltraLite SQL, 302

L

- LCASE function
 - UltraLite syntax, 392
- LEFT function
 - UltraLite syntax, 393
- LEFT OUTER JOIN clause
 - UltraLite syntax, 478
- LENGTH function
 - UltraLite syntax, 394
- LESSER function
 - UltraLite syntax, 395
- libraries
 - UltraLite choices, 24
 - UltraLite deploying uleng to Windows Mobile, 57
 - UltraLite FIPS-enabled applications, 192
- like-scan operation
 - UltraLite execution plans, 343
- limitations
 - UltraLite, 8
 - UltraLite data types, 312
- limits
 - UltraLite, 8
- line length
 - UltraLite sqlpp utility output, 252
- linking
 - UltraLite engine libraries, 23
 - UltraLite runtime libraries, 23
- LIST function
 - UltraLite syntax, 395
- Listeners
 - glossary definition, 522
- lists
 - UltraLite LIST function syntax, 395
- literals
 - UltraLite constants, 326
- LOAD TABLE statement
 - UltraLite syntax, 481
- loading
 - UltraLite bulk inserts, 481
 - UltraLite LOAD TABLE statement, 481
- loading data
 - UltraLite LOAD TABLE statement, 481
- loading databases
 - UltraLite databases with ulload, 272
- local temporary tables

- glossary definition, 522
- LOCATE function
 - UltraLite syntax, 396
- locking
 - UltraLite concurrency, 16
- locks
 - glossary definition, 523
- log files
 - glossary definition, 523
 - UltraLite Palm synchronization, 148
- LOG function
 - UltraLite syntax, 397
- LOG10 function
 - UltraLite syntax, 398
- logging
 - UltraLite internal mechanism, 15
- logic
 - UltraLite capturing for synchronization design, 133
- logical indexes
 - glossary definition, 523
- logical operators
 - UltraLite SQL syntax, 333
- lojoin operation
 - UltraLite execution plans, 343
- LONG BINARY data type
 - UltraLite, 312
- LONG VARCHAR data type
 - UltraLite, 312
- LOWER function
 - UltraLite syntax, 398
- lowercase strings
 - UltraLite LCASE function, 392
 - UltraLite LOWER function, 398
- LTM
 - glossary definition, 523
- LTRIM function
 - UltraLite syntax, 399
- M**
- maintaining
 - UltraLite on devices, 55
- maintaining primary key uniqueness
 - UltraLite clients in MobiLink systems, 129
- maintenance releases
 - glossary definition, 523
- management tools
 - UltraLite utilities reference, 247
- managing databases
 - UltraLite data and state, 13
- managing temporary tables
 - UltraLite about, 117
- managing transactions
 - UltraLite increased throughput, 119
- materialized views
 - glossary definition, 523
- mathematical expressions
 - arithmetic operators for UltraLite, 338
- MAX function
 - UltraLite syntax, 400
- max_hash_size creation parameter
 - UltraLite description, 194
- max_hash_size property
 - UltraLite description, 213
- maximum
 - UltraLite date ranges, 312
 - UltraLite physical limits, 8
- mechanism
 - UltraLite application and database deployment, 55
- media failures
 - UltraLite transaction overview , 16
- memory
 - UltraLite limits, 8
- memory failures
 - UltraLite preventing, 244
- memory usage
 - UltraLite database storage, 47
 - UltraLite indexes, 108
 - UltraLite row states, 14
- message log
 - glossary definition, 524
- message stores
 - glossary definition, 524
- message systems
 - glossary definition, 524
- message types
 - glossary definition, 524
- metadata
 - glossary definition, 524
 - UltraLite considering for reserve size, 244
- MIN function
 - UltraLite syntax, 401
- minimum
 - UltraLite date ranges, 312
- MINUTE function
 - UltraLite syntax, 401

MINUTES function
 UltraLite syntax, 402

mirror logs
 glossary definition, 524

MIRROR_FILE connection parameter
 UltraLite syntax , 235

ml_add_connection_script system procedure
 adding, 104

ml_add_table_script system procedure
 adding, 104

ML_GET_SERVER_NOTIFICATION
 syntax, 403

ml_remote_id
 UltraLite property configuration, 265
 UltraLite setting value, 223

ml_remote_id option
 UltraLite, 223

ml_remote_id property
 UltraLite description, 213

mlasdesk.dll
 deploying UltraLite applications, 64

mlasdev.dll
 deploying UltraLite applications, 64

mlasinst utility
 UltraLite deploying with DLL files, 64

mlcecc11.dll
 deploying UltraLite applications, 146

mlcrsa11.dll
 deploying UltraLite applications, 146

mlcrsafips11.dll
 deploying UltraLite applications, 146

mlczlib11.dll
 deploying UltraLite applications, 146

MobiLink
 glossary definition, 525
 UltraLite clients, 125
 UltraLite CREATE PUBLICATION statement, 462
 UltraLite SQL statements, 448
 UltraLite user ID uniqueness, 223

MobiLink clients
 glossary definition, 525
 UltraLite progress counter, 126

MobiLink file transfers
 UltraLite client overview of, 142

MobiLink Monitor
 glossary definition, 525

MobiLink server
 glossary definition, 525
 UltraLite HotSync, 62

MobiLink synchronization
 setting timestamp_increment creation parameter, 210
 UltraLite clients, 125

MobiLink system tables
 glossary definition, 525

MobiLink users
 glossary definition, 525

MOD function
 UltraLite syntax, 404

modeling
 UltraLite databases from MobiLink, 31

modifying
 UltraLite columns, 454
 UltraLite tables, 454

MONEY data type
 UltraLite equivalent, 314

monitoring synchronization
 UltraLite observer synchronization parameter, 163
 UltraLite setObserver method, 163

MONTH function
 UltraLite syntax, 405

MONTHNAME function
 UltraLite syntax, 406

MONTHS function
 UltraLite syntax, 406

multi-process access
 UltraLite engine, 23

multi-table joins
 UltraLite databases, 2

multi-threaded
 UltraLite applications, 13

multi-threaded applications
 UltraLite about , 23

multiple databases
 UltraLite maximum of, 13

multiple devices
 UltraLite connection parameters for, 47

N

name property
 UltraLite description, 213

names
 UltraLite column names, 326

NATURAL JOIN clause

- UltraLite syntax, 478
 - natural joins
 - glossary definition, 519
 - nearest century conversion considerations
 - UltraLite about, 196
 - nearest_century creation parameter
 - UltraLite description, 196
 - nearest_century property
 - UltraLite description, 213
 - network protocols
 - glossary definition, 525
 - UltraLite supported, 22
 - UltraLite Sync Result synchronization parameter, 173
 - UltraLite synchronization using HTTP, 171
 - UltraLite synchronization using HTTPS, 171
 - UltraLite synchronization using TCP/IP, 171
 - network server
 - glossary definition, 525
 - new password
 - UltraLite synchronization parameter, 162
 - new-database-ulcreatefile
 - UltraLite database creation [ulcreate] utility, 255
 - NEWID function
 - UltraLite syntax, 407
 - newmobilinkpwd synchronization parameter
 - UltraLite reference, 162
 - newsgroups
 - technical support, xvii
 - Non-repeatable reads
 - UltraLite about, 18
 - non-unique indexes
 - UltraLite characteristics, 108
 - UltraLite index creation from, 80
 - normalization
 - glossary definition, 525
 - nosync tables
 - UltraLite non-synchronizing tables , 136
 - UltraLite overview, 72
 - NOT
 - bitwise operator for UltraLite, 339
 - logical operators for UltraLite, 333
 - Notifiers
 - glossary definition, 526
 - NOW function
 - UltraLite syntax, 408
 - NT_FILE connection parameter
 - UltraLite syntax , 237
 - NULL
 - UltraLite ISNULL function, 392
 - NULL values
 - UltraLite SQL, 307
 - NULLIF function
 - UltraLite about, 409
 - using with CASE expressions in UltraLite, 328
 - NULLs
 - returned by functions if a NULL argument is specified, 347
 - number of authentication parameters
 - UltraLite synchronization parameter, 162
 - numbers
 - UltraLite SQL, 306
 - NUMERIC data type
 - UltraLite, 312
 - numeric functions
 - UltraLite alphabetical list, 351
 - numeric precision
 - UltraLite precision, 201
- ## O
- obfuscate database creation parameter
 - UltraLite description, 198
 - obfuscate database property
 - UltraLite usage, 40
 - obfuscation
 - UltraLite performance impact of, 120
 - UltraLite usage, 40
 - object trees
 - glossary definition, 526
 - OBJECT_ID function
 - syntax, 353
 - OBJECT_NAME function
 - syntax, 353
 - observer synchronization parameter
 - UltraLite description, 163
 - ODBC
 - glossary definition, 526
 - ODBC Administrator
 - glossary definition, 526
 - ODBC data sources
 - glossary definition, 526
 - on-device creation
 - about, 33
 - online books
 - PDF, xii

-
- operating systems
 - UltraLite Windows supported platforms, 24
 - operator precedence
 - UltraLite SQL syntax, 340
 - operators
 - arithmetic operators for UltraLite, 338
 - bitwise operators for UltraLite, 339
 - comparison operators for UltraLite, 332
 - logical operators for UltraLite, 333
 - precedence of operators for UltraLite, 340
 - string operators for UltraLite, 339
 - UltraLite SQL syntax, 338
 - optimization
 - UltraLite checkpoints, 119
 - UltraLite execution plan access options, 110
 - UltraLite indexes, 113
 - UltraLite queries, 487
 - UltraLite SQL, 342
 - optimizer
 - (*see also* query optimizer)
 - UltraLite execution plan access options, 110
 - UltraLite impact of, 342
 - UltraLite overriding, 342
 - UltraLite plan interpretation, 343
 - UltraLite plan operations, 343
 - UltraLite using, 342
 - optimizing UltraLite
 - query performance, 107
 - options
 - UltraLite browsing, 224
 - options (UltraLite)
 - commit_flush_count, 220
 - commit_flush_timeout, 221
 - DB_PROPERTY function, 380
 - global_database_id, 222
 - ml_remote_id, 223
 - SET OPTION statement, 488
 - OR
 - bitwise operators for UltraLite, 339
 - logical operators for UltraLite, 333
 - Oracle consolidated databases
 - UltraLite issues with, 127
 - ORDER BY clause
 - UltraLite SELECT statement, 487
 - UltraLite troubleshooting with, 504
 - order of operations
 - SQL operator precedence for UltraLite, 340
 - ORDERED_TABLE_SCAN connection parameter (deprecated)
 - UltraLite syntax, 239
 - ordering query results
 - UltraLite about, 117
 - ordering result sets
 - UltraLite troubleshooting, 504
 - outer joins
 - glossary definition, 526
 - outer references
 - subqueries for UltraLite SQL, 329
 - overhead
 - UltraLite considering for reserve size, 244
- P**
- packages
 - glossary definition, 526
 - packed rows
 - UltraLite about, 70
 - packing rows
 - UltraLite effect of, 199
 - page sizes
 - UltraLite page_size property, 213
 - page_size creation database property
 - UltraLite optimization for early Palm OS, 121
 - page_size creation parameter
 - UltraLite description, 199
 - page_size property
 - UltraLite description, 213
 - pages
 - UltraLite size considerations, 199
 - Palm Computing Platform
 - (*see also* Palm OS)
 - Palm data management utility
 - syntax, 258
 - Palm HotSync Conduit installer utility
 - syntax, 263
 - Palm OS
 - configuring UltraLite for early versions of, 121
 - expansion cards, using, 47
 - UltraLite application development for MobiLink, 146
 - UltraLite character sets, 37
 - UltraLite CustDB sample and readme files, 97
 - UltraLite data management utility for, 258
 - UltraLite databases, 47
 - UltraLite HotSync , 62

- UltraLite HotSync deployment, 62
- UltraLite PDB records, 47
- UltraLite publication restrictions, 137
- UltraLite VFS databases, 47
- UltraLite VFS deployment, 256
- PALM_ALLOW_BACKUP connection parameter
 - UltraLite syntax, 240
- PALM_FILE connection parameter
 - UltraLite syntax , 241
- parameters
 - UltraLite connection list, 44
 - UltraLite connection overview, 48
 - UltraLite database creation, 181
 - UltraLite SQL input, 330
- parse trees
 - glossary definition, 526
- partial download retained synchronization parameter
 - UltraLite reference, 163
- partition sizes
 - UltraLite defaults, choosing, 129
 - UltraLite defaults, overriding, 131
 - UltraLite exhausting defaults, 129
- partitioning
 - UltraLite primary keys, 129
 - UltraLite rows publishing, 84
- passthrough mode
 - UltraLite, 92
- password
 - UltraLite synchronization parameter, 164
- Password connection parameter
 - UltraLite syntax, 243
- passwords
 - PWD UltraLite connection parameter, 243
 - UltraLite adding new, 51
 - UltraLite changing, 87
 - UltraLite considerations, 87
 - UltraLite databases, 51
 - UltraLite defaults, 87
 - UltraLite new mobilink password parameter, 162
 - UltraLite password synchronization parameter, 164
 - UltraLite semantics, 51
- paths
 - UltraLite connection parameters, 47
 - UltraLite engine, 498
- PATINDEX function
 - UltraLite syntax, 411
- pattern matching
 - UltraLite PATINDEX function, 411
- PDB
 - (*see also* UltraLite databases and Palm OS)
 - glossary definition, 526
 - UltraLite databases, 47
- PDF
 - documentation, xii
- performance
 - separating commits from checkpoints, 119
 - UltraLite CACHE_SIZE connection parameter, 226
 - UltraLite choosing optimal hash size, 113
 - UltraLite download-only synchronization, 159
 - UltraLite index hashing, 194
 - UltraLite page sizes, 199
 - UltraLite preventing memory failures, 244
 - UltraLite query optimization, 487
 - UltraLite query optimization techniques, 107
 - UltraLite query tuning with index hashing, 111
 - UltraLite specifying FOR READ ONLY clause, 342
 - UltraLite ulcond11 cache, 263
 - UltraLite upload only synchronization, 174
 - UltraLite using index, 79
 - UltraLite using index for applications, 32
 - UltraLite using index for improving query performance, 460
 - UltraLite using index for large tables, 79
 - UltraLite viewing an execution plan, 342
- performance statistics
 - glossary definition, 527
- persistent memory
 - UltraLite database storage, 47
- personal server
 - glossary definition, 527
- phantom rows
 - UltraLite, 18
- physical indexes
 - glossary definition, 527
- physical limitations
 - UltraLite, 8
- PI function
 - UltraLite syntax, 412
- ping
 - UltraLite synchronization parameter, 165
- placeholder
 - UltraLite SQL input parameter, 330
- planning
 - UltraLite synchronization design overview, 133

- plans
 - UltraLite cursors, 383
 - UltraLite operations for, 343
 - UltraLite plan interpretation, 343
 - UltraLite plan operations, 343
 - UltraLite queries, overriding, 342
 - UltraLite queries, reading, 343
 - UltraLite queries, working with, 342
 - UltraLite syntax, 383
 - UltraLite text plans, 342
- platforms
 - UltraLite file storage, 47
 - UltraLite multiple connection parameters for, 47
- plug-in modules
 - glossary definition, 527
- plug-ins
 - UltraLite troubleshooting, 499
- policies
 - glossary definition, 527
- polling
 - glossary definition, 527
- pools
 - UltraLite unused Global IDs, 129
- populating
 - UltraLite databases created with ulinit, 270
- port number
 - UltraLite ULSynchronize arguments, 172
- POWER function
 - UltraLite syntax, 412
- PowerDesigner
 - glossary definition, 527
- PowerJ
 - glossary definition, 527
- precedence
 - SQL operator precedence for UltraLite, 340
- precision creation parameter
 - UltraLite description, 201
- precision property
 - UltraLite description, 213
- predicates
 - ALL for UltraLite SQL, 334
 - ANY for UltraLite SQL, 335
 - BETWEEN for UltraLite SQL, 336
 - comparison operators for UltraLite, 332
 - EXISTS for UltraLite SQL, 336
 - glossary definition, 527
 - IN for UltraLite SQL, 337
 - UltraLite SQL, 331
- prefix
 - UltraLite for Windows Mobile databases , 47
- prepared statements
 - UltraLite input parameters, 330
- primary key constraints
 - glossary definition, 528
- primary key indexes
 - UltraLite bypassing use of, 117
 - UltraLite using, 117
- primary keys
 - glossary definition, 528
 - UltraLite characteristics, 108
 - UltraLite generating unique values, 407
 - UltraLite generating unique values using UUIDs, 407
 - UltraLite indexing, 32
 - UltraLite integrity constraints, 468
 - UltraLite order of columns, 468
 - UltraLite table order, 138
 - UltraLite tables, 71
 - UltraLite troubleshooting data imports, 502
 - UltraLite UUIDs and GUIDs, 407
- primary tables
 - glossary definition, 528
- procedures
 - UltraLite limitations, 2
- programming interfaces
 - UltraLite supported, 24
- progress counter
 - UltraLite offset mismatches with , 126
- properties
 - DB_PROPERTY function, 380
 - UltraLite alphabetical list, 213
 - UltraLite browsing, 218
 - UltraLite database creation parameters, 34
 - UltraLite system table for, 298
- properties (UltraLite)
 - DB_PROPERTY function, 380
- protocol options
 - UltraLite HotSync, 148
- provider files
 - UltraLite deploying ActiveSync, 64
- providers
 - UltraLite ActiveSync files, 64
- proxy tables
 - glossary definition, 528
- public certificate
 - UltraLite application access to, 276

- publication creation wizard
 - UltraLite creating publications, 137
- publication updates
 - glossary definition, 528
- publications
 - glossary definition, 528
 - UltraLite ALTER PUBLICATION statement, 451
 - UltraLite copying method, 77
 - UltraLite CREATE PUBLICATION statement, 462
 - UltraLite design plans with, 133
 - UltraLite dropping , 85
 - UltraLite limit of, 269
 - UltraLite publication synchronization parameter, 166
 - UltraLite publishing overview, 137
 - UltraLite publishing tables, 83
 - UltraLite rows publishing, 84
 - UltraLite schema description for, 296
 - UltraLite SQL CREATE INDEX statement, 460
 - UltraLite SQL DROP INDEX statement, 474
 - UltraLite SQL DROP PUBLICATION statement, 475
 - UltraLite SQL DROP TABLE statement, 477
 - UltraLite synchronization, 166
 - UltraLite synchronization parameter for, 166
 - UltraLite sysarticle system table, 297
 - UltraLite syspublication system table, 296
 - UltraLite table listing in schema , 297
 - UltraLite WHERE clause usage, 84
 - UltraLite working with, 83
- publisher
 - glossary definition, 529
- publishing
 - UltraLite rows, 84
 - UltraLite tables, 83
 - UltraLite whole table, 137
- publishing whole tables
 - UltraLite, 83
- push notifications
 - glossary definition, 529
- push requests
 - glossary definition, 529
- PWD connection parameter
 - UltraLite syntax, 243
- Q**
- QAnywhere
 - glossary definition, 529
- QAnywhere Agent
 - glossary definition, 529
- QUARTER function
 - UltraLite syntax, 413
- queries
 - glossary definition, 529
 - UltraLite optimization, 487
 - UltraLite ordering results with primary key, 117
 - UltraLite troubleshooting unpredictable result sets, 504
- query optimization
 - (*see also* optimizer)
 - UltraLite SQL, 342
- query optimizer
 - (*see also* optimizer)
 - UltraLite, 342
- QUOTES clause
 - UltraLite LOAD TABLE statement, 482
- R**
- RADIANS function
 - UltraLite syntax, 414
- range
 - UltraLite date type, 312
- RDBMS
 - glossary definition, 530
- read-only tables
 - UltraLite databases, 139
 - UltraLite synchronizing, 139
- ReadCommitted
 - UltraLite isolation levels, 18
- reading
 - UltraLite table rows, 18
- reading UltraLite access plans
 - about, 343
- readme files
 - UltraLite CustDB applications, 96
- ReadUncommitted
 - UltraLite isolation levels, 18
- REAL data type
 - UltraLite, 312
- recovery
 - UltraLite, 15
 - UltraLite introduction, 14

- UltraLite transaction overview , 16
- recovery from system failure
 - UltraLite internal mechanism, 15
- Redirector
 - glossary definition, 529
- reference databases
 - creating for UltraLite, 31
 - glossary definition, 529
 - options for UltraLite, 31
- referenced object
 - glossary definition, 530
- referencing object
 - glossary definition, 530
- referential integrity
 - glossary definition, 530
 - UltraLite databases, 2
 - UltraLite indexes, 79
 - UltraLite table order, 138
- registering
 - MobiLink UltraLite applications with ActiveSync, 66
- registry
 - UltraLite Hotsync keys, 62
- registry keys
 - ulcond11 cache size, 263
- regular expressions
 - glossary definition, 530
- reload.sql
 - UltraLite loading, 481
- REMAINDER function
 - UltraLite syntax, 414
- remote databases
 - creating UltraLite clients, 128
 - deleting UltraLite data, 258
 - glossary definition, 530
 - UltraLite synchronization count, 126
- REMOTE DBA authority
 - glossary definition, 530
- remote IDs
 - glossary definition, 531
 - setting in UltraLite databases, 223
- remote servers
 - UltraLite CREATE TABLE statement, 468
- removing
 - UltraLite users, 88
- removing data
 - UltraLite file size impact of, 501
- renaming
 - UltraLite database objects during upgrade, 67
 - UltraLite tables, 454
- REPEAT function
 - UltraLite syntax, 415
- REPLACE function
 - UltraLite syntax, 416
- REPLICATE function
 - UltraLite syntax, 417
- replication
 - glossary definition, 531
- Replication Agent
 - glossary definition, 531
- replication frequency
 - glossary definition, 531
- replication messages
 - glossary definition, 531
- Replication Server
 - glossary definition, 531
- requests
 - UltraLite concurrency, 13
 - UltraLite management of, 13
- RESERVE_SIZE connection parameter
 - UltraLite syntax, 244
- reserved words
 - UltraLite SQL, 302
- restartable downloads
 - UltraLite keep partial download, 161
 - UltraLite partial download retained, 163
 - UltraLite resume partial download, 167
- restoring
 - UltraLite transaction overview , 16
- result sets
 - UltraLite troubleshooting unpredictable changes, 504
- resume partial download synchronization parameter
 - UltraLite reference, 167
- return codes
 - Interactive SQL [dbisql] utility for UltraLite, 251
- RIGHT function
 - UltraLite syntax, 418
- RIGHT OUTER JOIN clause
 - UltraLite syntax, 478
- role names
 - glossary definition, 531
 - UltraLite foreign keys, 468
 - UltraLite role names, 468
- role of user authentication
 - UltraLite about, 51

- roles
 - glossary definition, 531
- rollback logs
 - glossary definition, 532
- ROLLBACK statement
 - UltraLite syntax, 485
- rollbacks
 - UltraLite databases, 14
- rolling back
 - UltraLite transaction overview , 16
 - UltraLite transactions, 485
- ROUND function
 - UltraLite syntax, 419
- rounding
 - UltraLite scale, 203
 - UltraLite scale creation parameter , 203
- row packing
 - UltraLite about, 70
 - UltraLite effect of, 199
 - UltraLite observing, 244
- row-level triggers
 - glossary definition, 532
- rowlimit operation
 - UltraLite execution plans, 343
- rows
 - UltraLite deleting all rows from a table, 493
 - UltraLite fetching, 18
 - UltraLite INSERT statement, 480
 - UltraLite inserting in bulk, 481
 - UltraLite locking of, 16
 - UltraLite publishing, 84
- RSA encryption algorithm
 - UltraLite TLS-enabled synchronization, 60
- RTRIM function
 - UltraLite syntax, 419
- running
 - UltraLite CustDB application, 97
- runtime libraries
 - UltraLite list of, 23
- runtimes
 - (*see also* UltraLite runtime)
- S**
- sample application
 - starting CustDB in UltraLite, 99
- samples
 - (*see also* examples)
 - (*see also* tutorials)
- samples-dir
 - documentation usage, xiv
- scale creation parameter
 - UltraLite description, 203
- scale property
 - UltraLite description, 213
- scan operation
 - UltraLite execution plans, 343
- scanning
 - UltraLite database pages in queries, 117
 - UltraLite indexes in queries, 108
- schema
 - UltraLite catalog of tables for, 11
 - UltraLite changes, precautions, 11
- schemas
 - glossary definition, 532
 - UltraLite SQL ALTER DATABASE SCHEMA FROM FILE syntax, 450
 - UltraLite system tables, 289
- scjview
 - terminating before installing, 499
- script versions
 - glossary definition, 532
 - UltraLite getScriptVersion, 176
 - UltraLite setScriptVersion method , 176
 - UltraLite version synchronization parameter , 176
- script-based uploads
 - glossary definition, 532
- scripted upload
 - UltraLite CREATE PUBLICATION syntax, 462
- scripts
 - glossary definition, 532
- search conditions
 - ALL for UltraLite SQL, 334
 - ANY for UltraLite SQL, 335
 - BETWEEN for UltraLite SQL, 336
 - EXISTS for UltraLite SQL, 336
 - IN for UltraLite SQL, 337
 - UltraLite SQL, 331
- SECOND function
 - UltraLite syntax, 420
- SECONDS function
 - UltraLite syntax, 421
- secured features
 - glossary definition, 532
- security
 - UltraLite, 40

- UltraLite AES FIPS database deployment, 59
- UltraLite FIPS encryption, 192
- UltraLite overview of, 40
- UltraLite TLS-enabled synchronization, 60
- UltraLite user authentication, 51
- SELECT statement
 - UltraLite browsing data example, 76
 - UltraLite copying result sets with, 77
 - UltraLite syntax, 486
 - UltraLite troubleshooting result sets, 504
- selecting
 - UltraLite SELECT statement, 486
- send column names
 - UltraLite synchronization parameter, 168
- send download acknowledgement
 - UltraLite synchronization parameter, 169
- send_column_names synchronization parameter
 - UltraLite reference, 168
- send_download_ack synchronization parameter
 - UltraLite reference, 169
- sensitivity
 - UltraLite case creation parameter, 183
- serialized transaction
 - UltraLite processing of, 16
- server management requests
 - glossary definition, 532
- server message stores
 - glossary definition, 533
- server-initiated synchronization
 - glossary definition, 532
- services
 - glossary definition, 533
- session-based synchronization
 - glossary definition, 533
- SET OPTION statement
 - UltraLite syntax, 488
- setObserver method
 - UltraLite example, 163
- setScriptVersion method
 - UltraLite example, 176
- setStream method
 - UltraLite example, 171
- setStreamParms method
 - UltraLite example, 172
- setting the global database identifier
 - UltraLite clients in MobiLink systems, 129
- setting the hash size
 - about, 116
- setUserData method
 - UltraLite example, 175
- SHORT_PLAN function
 - UltraLite syntax, 422
- SIGN function
 - UltraLite syntax, 423
- SIMILAR function
 - UltraLite syntax, 424
- simple encryption
 - UltraLite performance impact of, 120
- SIN function
 - UltraLite syntax, 424
- SKIP clause
 - UltraLite LOAD TABLE statement, 483
- slash-asterisk
 - UltraLite comment indicator, 305
- SMALLINT data type
 - UltraLite, 312
- SMALLMONEY data type
 - UltraLite equivalent, 314
- snapshot isolation
 - glossary definition, 533
- sort order
 - UltraLite collations, 37
- SOUNDEX function
 - UltraLite syntax, 425
- SPACE function
 - UltraLite syntax, 426
- spaces
 - UltraLite filepath definitions, 498
- special values
 - UltraLite CURRENT DATE, 308
 - UltraLite CURRENT TIME, 308
 - UltraLite CURRENT TIMESTAMP, 309
 - UltraLite SQL, 308
 - UltraLite SQLCODE, 309
- SQL
 - (*see also* UltraLite SQL)
 - comparison operators for UltraLite, 332
 - data types in UltraLite, 312
 - expressions in UltraLite, 325
 - glossary definition, 533
 - operators in UltraLite, 338
 - search conditions in UltraLite, 331
 - UltraLite identifiers , 303
 - UltraLite keywords , 302
 - UltraLite numbers, 306
 - UltraLite reserved words, 302

- UltraLite schema changes with, 11
- UltraLite statement types, 449
- UltraLite strings, 304
- variables in UltraLite, 341
- SQL Anywhere
 - documentation, xii
 - glossary definition, 533
- SQL Anywhere databases
 - database comparison with UltraLite, 2
- SQL code
 - UltraLite schema upgrade errors, 68
- SQL Flagger
 - UltraLite usage of, 252
- SQL functions
 - alphabetical list of UltraLite functions A-D, 354
 - alphabetical list of UltraLite functions E-O, 383
 - alphabetical list of UltraLite functions P-Z, 411
 - return NULL if you specify NULL argument, 347
 - types of function for UltraLite, 348
 - UltraLite aggregate, 348
 - UltraLite data type conversion, 348
 - UltraLite date and time, 348
 - UltraLite introduction, 347
 - UltraLite miscellaneous, 350
 - UltraLite numeric, 351
 - UltraLite string, 352
 - UltraLite system, 353
- SQL passthrough
 - UltraLite support, 92
- SQL preprocessor utility
 - UltraLite syntax , 252
- SQL Remote
 - glossary definition, 533
- SQL statements
 - alphabetical list of UltraLite statements, 448
 - glossary definition, 533
 - UltraLite, 448
- SQL syntax
 - CASE expression for UltraLite, 327
 - UltraLite column names, 326
 - UltraLite comments, 305
 - UltraLite constants, 326
 - UltraLite functions, 348
 - UltraLite IF expressions , 327
 - UltraLite input parameters, 330
 - UltraLite special values, 308
 - UltraLite SQLCODE special value, 309
- SQL-based synchronization
 - glossary definition, 533
- SQLCODE
 - UltraLite concurrency checks, 16
 - UltraLite special value, 309
 - UltraLite SQLE_MAX_ROW_SIZE_EXCEEDED error, 70
- SQLCODE SQLE_LOCKED
 - UltraLite concurrency error, 16
- SQLE_CONVERSION_ERROR
 - UltraLite upgrade warning, 68
- SQLE_DATABASE_ERROR
 - UltraLite data corruption, 500
- SQLE_DEVICE_ERROR
 - UltraLite data corruption, 500
- SQLE_DOWNLOAD_CONFLICT error
 - UltraLite synchronization, 139
- SQLE_MAX_ROW_SIZE_EXCEEDED
 - UltraLite error, 70
- SQLE_MEMORY_ERROR
 - UltraLite data corruption, 500
- SQLE_NOTFOUND
 - UltraLite concurrency error, 16
- SQLE_ROW_DROPPED_DURING_SCHEMA_UPGRADE
 - UltraLite upgrade warning, 67
- SQLE_UNABLE_TO_CONNECT_OR_START
 - troubleshooting, 498
- sqlpp utility
 - UltraLite syntax, 252
- SQRT function
 - UltraLite syntax, 427
- square root function
 - UltraLite SQRT function, 427
- START connection parameter
 - UltraLite for uleng11 on Windows Mobile, 57
 - UltraLite syntax, 245
- START SYNCHRONIZATION DELETE statement
 - UltraLite syntax, 489
- state bytes
 - UltraLite databases, 14
- state management
 - UltraLite overview, 13
- statement-level triggers
 - glossary definition, 534
- statements
 - UltraLite, 448
 - UltraLite prepared, input parameters for, 330
 - UltraLite types, 449

- STOP SYNCHRONIZATION DELETE statement
 - UltraLite syntax, 490
- storage
 - UltraLite databases, 11
 - UltraLite limits, 8
 - UltraLite PALM_ALLOW_BACKUP , 240
 - UltraLite reserve size, 244
- stored procedures
 - glossary definition, 534
 - UltraLite limitations, 2
- STR function
 - UltraLite syntax, 427
- stream error
 - UltraLite synchronization parameter, 169
- stream parameters
 - UltraLite HotSync synchronization, 148
 - UltraLite synchronization parameter, 172
- stream synchronization parameters
 - UltraLite synchronization parameter, 172
- stream type
 - UltraLite synchronization parameter, 171
- stream_error synchronization parameter
 - UltraLite reference, 169
 - UltraLite ul_stream_error structure, 169
- stream_parms synchronization parameter
 - UltraLite reference, 172
 - UltraLite synchronization using HotSync, 146
 - UltraLite usage, 148
- STRING function
 - UltraLite syntax, 428
- string functions
 - UltraLite alphabetical list, 352
- string length
 - UltraLite LENGTH function, 394
- string literal
 - glossary definition, 534
- string literals
 - UltraLite constants, 326
- string operators
 - dynamic SQL syntax for UltraLite, 339
- string position
 - UltraLite LOCATION function, 396
- strings
 - UltraLite case sensitivity, 304
 - UltraLite maximum size, 8
 - UltraLite nearest_century conversions to dates with , 196
 - UltraLite removing trailing blanks , 419
 - UltraLite replacing, 416
 - UltraLite SQL, 304
 - UltraLite SQL functions, 352
- STRIP clause
 - UltraLite LOAD TABLE statement, 483
- strong encryption
 - UltraLite deployment steps, 59
 - UltraLite fips creation parameter, 192
 - UltraLite performance impact of, 120
 - UltraLite usage, 40
- STRTOUUID function
 - UltraLite syntax, 429
- STUFF function
 - UltraLite syntax, 430
- subqueries
 - glossary definition, 534
 - UltraLite SQL, 329
- subquery operation
 - UltraLite execution plans, 343
- SUBSCRIBE BY clause
 - UltraLite synchronization limitations, 137
- subscriptions
 - glossary definition, 534
- SUBSTR function
 - UltraLite syntax, 430
- SUBSTRING function
 - UltraLite syntax, 430
- substrings
 - UltraLite about, 430
 - UltraLite replacing, 416
- SUM function
 - UltraLite syntax, 432
- supplying UltraLite connection parameters
 - about, 45
- support
 - newsgroups, xvii
- SUSER_ID function
 - syntax, 433
- SUSER_NAME function
 - syntax, 433
- Sybase Central
 - browsing CustDB in UltraLite, 104
 - creating UltraLite databases, 30
 - glossary definition, 534
 - troubleshooting UltraLite connections, 499
 - UltraLite column creation methods, 72
 - UltraLite copying database objects method, 77
 - UltraLite creating indexes, 81

- UltraLite creating publications , 137
- UltraLite creating table methods, 71
- UltraLite system table browsing methods, 76
- UltraLite table alteration methods, 74
- UltraLite table browsing methods, 76
- UltraLite table deletion methods, 75
- sync result
 - UltraLite synchronization parameter, 173
- SYNC_PROFILE_OPTION_VALUE function
 - syntax, 434
- synchronization
 - character sets in UltraLite, 36
 - glossary definition, 535
 - setting timestamp_increment creation parameter, 210
 - ulsync utility for UltraLite databases, 275
 - UltraLite application implementation , 139
 - UltraLite client-specific data, 136
 - UltraLite clients, 125
 - UltraLite CustDB tutorial, 102
 - UltraLite design overview, 133
 - UltraLite download_only parameter, 159
 - UltraLite excluding tables with publications, 137
 - UltraLite foreign keys, 138
 - UltraLite HotSync protocol options, 148
 - UltraLite ignored rows, 160
 - UltraLite introduction , 125
 - UltraLite M-Business Anywhere channel, 142
 - UltraLite monitoring, 163
 - UltraLite Palm OS, 62
 - UltraLite progress counting mechanism, 126
 - UltraLite read-only tables , 139
 - UltraLite referential integrity, 138
 - UltraLite schema changes during, 11
 - UltraLite SQLE_DOWNLOAD_CONFLICT error, 139
 - UltraLite stopping, 163
 - UltraLite system table, 291
 - UltraLite task overview , 127
 - UltraLite upload only parameter, 174
- synchronization logic
 - browsing Sybase Central in UltraLite, 104
- synchronization models
 - UltraLite databases, 31
- synchronization parameters
 - UltraLite, 154
 - UltraLite additionalparms, 155
 - UltraLite Authentication Value, 159
 - UltraLite Disable Concurrency overview, 13
 - UltraLite download_only, 159
 - UltraLite getScriptVersion , 176
 - UltraLite getStream method , 171
 - UltraLite getUploadOK method, 173
 - UltraLite keep partial download, 161
 - UltraLite newmobilinkpwd, 162
 - UltraLite observer , 163
 - UltraLite partial download retained, 163
 - UltraLite password, 164
 - UltraLite ping, 165
 - UltraLite publication, 166
 - UltraLite required, 154
 - UltraLite resume partial download, 167
 - UltraLite send_column_names, 168
 - UltraLite send_download_ack , 169
 - UltraLite setObserver method, 163
 - UltraLite setScriptVersion method, 176
 - UltraLite setStream method , 171
 - UltraLite setStreamParms method, 172
 - UltraLite setSynchPublication method, 166
 - UltraLite setUserData method, 175
 - UltraLite stream type , 171
 - UltraLite stream_error, 169
 - UltraLite stream_parms, 172
 - UltraLite Sync Result , 173
 - UltraLite upload_ok , 173
 - UltraLite upload_only, 174
 - UltraLite user_data, 175
 - UltraLite user_name, 175
 - UltraLite version , 176
- synchronization profile options
 - about, 278
 - UltraLite synchronization, 278
- synchronization profiles
 - UltraLite ALTER SYNCHRONIZATION PROFILE statement, 452
 - UltraLite DROP SYNCHRONIZATION PROFILE statement, 464, 476
 - UltraLite SYNCHRONIZE statement, 491
- synchronization scripts
 - browsing the UltraLite sample, 104
- synchronization stream parameters
 - UltraLite stream type , 171
- synchronization streams
 - UltraLite getStream method, 171
 - UltraLite setStream method, 171
 - UltraLite setStreamParms method, 172

- UltraLite setting , 171
- UltraLite stream synchronization parameter , 171
- UltraLite stream_error synchronization parameter, 169
- UltraLite stream_parms synchronization parameter, 172
- UltraLite ULHTTPStream, 171
- UltraLite ULHTTPStream, 171
- SYNCHRONIZE statement
 - UltraLite syntax, 491
- synchronizing
 - UltraLite databases on Windows Mobile, 150
- synchronizing data
 - UltraLite file size impact of, 501
- syntax
 - arithmetic operators for UltraLite, 338
 - bitwise operators for UltraLite, 339
 - CASE expression for UltraLite, 327
 - comparison operators for UltraLite, 332
 - logical operators for UltraLite, 333
 - SQL operator precedence for UltraLite, 340
 - string operators for UltraLite, 339
 - UltraLite column names, 326
 - UltraLite constants, 326
 - UltraLite CURRENT DATE special value, 308
 - UltraLite CURRENT TIMESTAMP special value, 309
 - UltraLite IF expressions , 327
 - UltraLite special values, 308
 - UltraLite SQL comments, 305
 - UltraLite SQL CURRENT TIME special value, 308
 - UltraLite SQL functions, 348
 - UltraLite SQL functions A-D, 354
 - UltraLite SQL functions E-O, 383
 - UltraLite SQL functions P-Z, 411
 - UltraLite SQL input parameters, 330
 - UltraLite SQL operators, 338
 - UltraLite SQLCODE special value, 309
- SYS
 - glossary definition, 535
 - UltraLite system tables, 289
- sysarticle system table [UltraLite]
 - about, 297
- syscolumn system table [UltraLite]
 - about, 292
- sysindex system table [UltraLite]
 - about, 293
- sysixcol system table [UltraLite]
 - about, 295
- syspublication system table [UltraLite]
 - about, 296
- systable system table [UltraLite]
 - about, 291
- system failure
 - UltraLite recover from, 15
- system failures
 - UltraLite transaction overview , 16
- system functions
 - UltraLite about, 353
 - UltraLite limitations, 2
- system objects
 - glossary definition, 535
 - UltraLite displaying methods, 77
- system tables
 - glossary definition, 535
 - UltraLite about, 289
 - UltraLite browsing methods, 76
 - UltraLite hiding and showing, 290
 - UltraLite sysarticle, 297
 - UltraLite syscolumn, 292
 - UltraLite sysindex, 293
 - UltraLite sysixcol, 295
 - UltraLite syspublication, 296
 - UltraLite systable, 291
 - UltraLite sysuldata, 298
- system views
 - glossary definition, 535
- system_error_code values
 - UltraLite synchronization stream errors, 170
- sysuldata system table [UltraLite]
 - about, 298

T

- table constraints
 - UltraLite adding, deleting, or modifying, 454
 - UltraLite CREATE TABLE statement, 468
- table expressions
 - subqueries for UltraLite SQL, 329
- table owners
 - UltraLite, 303
- table scans
 - UltraLite ordering results with primary key, 117
- table size
 - number of rows, 8

- UltraLite limit, 8
- TableOrder
 - UltraLite synchronization parameter, 155
- tableOrder
 - UltraLite ulsync options for, 275
- tables
 - UltraLite allsync, 136
 - UltraLite ALTER TABLE statement, 454
 - UltraLite altering methods, 74
 - UltraLite browsing methods, 76
 - UltraLite bulk loading, 481
 - UltraLite controlling synchronization with publications, 137
 - UltraLite copying method, 77
 - UltraLite copying methods, 77
 - UltraLite CREATE TABLE statement, 468
 - UltraLite creating methods, 71
 - UltraLite deleting methods, 75
 - UltraLite editing methods, 76
 - UltraLite INSERT statement, 480
 - UltraLite intermediate query result storage, 12
 - UltraLite limitations, 8
 - UltraLite nosync, 136
 - UltraLite order of, 138
 - UltraLite publications, 137
 - UltraLite rows publishing, 84
 - UltraLite size limit, 70
 - UltraLite table filtering methods, 77
 - UltraLite temporary table usage, 342
 - UltraLite TRUNCATE TABLE statement, 493
 - UltraLite working with, 70
- TAN function
 - UltraLite syntax, 435
- TCP/IP
 - (*see also* TCP/IP synchronization)
- technical support
 - newsgroups, xvii
- temp operation
 - UltraLite execution plans, 343
- temporary files
 - UltraLite limitation, 12
- temporary tables
 - glossary definition, 535
 - UltraLite about, 12
 - UltraLite limitations, 2, 8
 - UltraLite managing, 117
 - UltraLite queries, 342
 - UltraLite synchronization using, 136
- terminating processes
 - UltraLite troubleshooting upgrades, 499
- TEXT data type
 - UltraLite equivalent, 314
- text plans
 - viewing in UltraLite, 342
- THEN
 - UltraLite IF expressions, 327
- threads
 - UltraLite concurrency, 13
- time considerations
 - UltraLite about, 205
- TIME data type
 - UltraLite , 312
- time functions
 - UltraLite alphabetical list, 348
- time_format creation parameter
 - UltraLite description, 205
- time_format property
 - UltraLite description, 213
- times
 - UltraLite conversion functions, 348
- TIMESTAMP
 - UltraLite column limitations, 2
 - UltraLite TIMESTAMP columns, 468
- TIMESTAMP data type
 - UltraLite, 312
- timestamp_format creation parameter
 - UltraLite description, 207
- timestamp_format property
 - UltraLite description, 213
- timestamp_increment creation parameter
 - UltraLite description, 210
 - using in MobiLink synchronization, 210
- timestamp_increment property
 - UltraLite description, 213
- timestamps
 - UltraLite timestamp_format creation parameter, 207
 - UltraLite timestamp_increment creation parameter, 210
- TINYINT data type
 - UltraLite, 312
- tips
 - UltraLite implementation , 21
- TLS
 - UltraLite client configuration, 60

- UltraLite synchronization deployment considerations, 60
- TLS_TYPE network protocol option
 - UltraLite TLS-enabled synchronization, 60
- TODAY function
 - UltraLite syntax, 435
- TOP clause
 - UltraLite SELECT statement, 486
- topics
 - graphic icons, xvii
- tracking
 - UltraLite synchronization, 126
- transaction log
 - glossary definition, 535
- transaction log mirror
 - glossary definition, 536
- transaction logging
 - UltraLite internal mechanism, 15
- transaction management
 - UltraLite, 13
- transaction processing
 - checkpoints and commits, 119
 - UltraLite COMMIT_FLUSH connection parameter, 229
 - UltraLite commit_flush_count option, 220
 - UltraLite commit_flush_timeout option, 221
- transactional integrity
 - glossary definition, 536
- transactions
 - glossary definition, 535
 - UltraLite checkpoints for, 458
 - UltraLite COMMIT statement, 459
 - UltraLite concurrency, 13
 - UltraLite databases, 14
 - UltraLite flushing, 220
 - UltraLite isolation level, 18
 - UltraLite rolling back, 485
 - UltraLite schema changes impact, 11
- transferring files
 - UltraLite files with MLFileTransfer, 142
- transmission rules
 - glossary definition, 536
- transport-layer security
 - (*see also* TLS)
- triggers
 - glossary definition, 536
 - UltraLite limitations, 2
- TRIM function
 - UltraLite syntax, 436
- troubleshooting
 - avoiding synchronization issues with foreign key cycles, 138
 - newsgroups, xvii
 - UltraLite backing up application, 125
 - UltraLite changing collation, 36
 - UltraLite checksum failures, 184
 - UltraLite connection parameter precedence, 49
 - UltraLite enabling automatic backups, 259
 - UltraLite getUploadOK method , 173
 - UltraLite global ID numbers, 130
 - UltraLite HotSync , 148
 - UltraLite implementing resumable downloads, 140
 - UltraLite maintaining timestamps and timestamp increments, 207
 - UltraLite ping synchronization parameter, 165
 - UltraLite retrieving GLOBAL AUTOINCREMENT value, 131
 - UltraLite running on early Palm OS, 121
 - UltraLite Stream Error synchronization parameter, 169
 - UltraLite Sync Result synchronization parameter , 173
 - UltraLite transmittal of connection parameters, 49
 - UltraLite upload_ok synchronization parameter, 173
- TRUNCATE TABLE statement
 - UltraLite syntax, 493
- truncating
 - UltraLite tables, 493
- TRUNCNUM function
 - UltraLite syntax, 437
- trusted certificates
 - UltraLite application access to encryption information, 276
- TSEQUAL function
 - syntax, 437
- tuning performance
 - UltraLite max_hash_size, 194
 - UltraLite with index hashing, 111
- tutorials
 - UltraLite CustDB database synchronization, 102
 - UltraLite CustDB files, 95
 - UltraLite CustDB introduction, 93

U

UCASE function

- UltraLite syntax, 438

UDB

- (*see also* UltraLite databases)

UID connection parameter

- UltraLite syntax, 246

UL_DEBUG_CONDUIT_LOG environment variable

- UltraLite HotSync troubleshooting, 148

ul_stream_error structure

- UltraLite example, 169

UL_SYNC_ALL macro

- UltraLite publications list, 166

UL_SYNC_ALL_PUBS macro

- UltraLite publications list, 166

ulcond11 utility

- syntax, 263

ULConduitStream function

- UltraLite synchronization stream, 171

ulcreate utility

- syntax, 255

- UltraLite Palm OS VFS considerations, 256

- using, 30

ULDBUtil

- about, 258

uleng11 utility

- deploying to Windows Mobile, 57

- in-process database support, 260

- syntax, 260

ulerase utility

- syntax, 262

ULHTTPSSStream function [UL ESQL]

- UltraLite synchronization stream, 171

ULHTTPStream function [UL ESQL]

- UltraLite synchronization stream, 171

ulinfo utility

- UltraLite syntax, 265

ulinit utility

- syntax, 269

- UltraLite Palm OS VFS considerations, 271

- unsupported collation workaround, 270

- using, 31

ulload utility

- syntax, 272

- UltraLite Palm OS VFS considerations, 274

- using, 32

ULSocketStream function

- UltraLite synchronization stream, 171

ULSQLCONNECT environment variable

- about, 45

- description, 53

ulstop utility

- syntax, 261

ulsync utility

- synchronization profile options, 278

- syntax, 275

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

- collations unsupported, 270

- data conversion, 348

- database creation parameters, 181

- deploying applications and databases, 55

- error codes, 248

- glossary definition, 536

- installing applications and databases, 55

- multi-threaded applications, 23

- SQL functions, aggregate, 348

- SQL functions, data type conversion, 348

- SQL functions, types of, 348

- SQL statement reference, 448

- SQLE_MAX_ROW_SIZE_EXCEEDED error, 70

- system functions, 353

- table owners, 303

- troubleshooting, 497

- utilities reference, 247

UltraLite administration tools

- troubleshooting, 503

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)

- APIs choosing, 24

- building CustDB, 97

- concurrency, 13

- CustDB applications and readme files, 96

- defining engine location at connection time, 245
- deploying ActiveSync provider files, 64
- deploying FIPS-enabled, 192
- deploying HotSync conduit files, 146
- deploying to devices, 57
- development platforms, 24
- error -764, 505
- libraries choosing , 24
- managing multiple requests, 13
- public certificate access, 276
- supported Windows platforms, 24
- synchronization , 125
- TLS-enabled synchronization, 60
- transferring files with MobiLink, 142
- troubleshooting
 - SQLE_UNABLE_TO_CONNECT_OR_START, 498
- UltraLite C/C++
 - building CustDB application, 97
 - CustDB sample and readme files, 97
 - engine support, 23
- UltraLite client synchronization
 - parameters and options about, 154
- UltraLite clients
 - about MobiLink, 125
- UltraLite columns
 - renaming during schema upgrade, 67
- UltraLite connection parameters
 - about, 48
 - list of, 44
- UltraLite connections
 - troubleshooting, 498, 499
- UltraLite creation parameters
 - about, 181
- UltraLite database creation utility
 - syntax, 255
- UltraLite database properties
 - about, 213
- UltraLite database synchronization utility
 - synchronization profile options, 278
- UltraLite databases
 - about, 11
 - backing up on Windows Mobile, 47
 - backups, 15
 - checkpoint usage, 119
 - collation sequences, 36
 - columns adding, 72
 - columns altering, 73
 - composition of, 11
 - concurrency, 13
 - connection optimization strategies, 121
 - connection overview, 43
 - connection parameter list, 44
 - connection parameters overview, 48
 - connections overview, 13
 - counting synchronizations, 126
 - create with the create database wizard, 30
 - creating, 29
 - creating from SQL Anywhere reference database, 31
 - creating from the command prompt, 30
 - creation parameters, 34
 - data and state management, 13
 - data management options, 23
 - database comparison with SQL Anywhere, 2
 - database integrity, 14
 - deadlocks, 16
 - deleting application data from Palm OS device, 258
 - deploying for data encryption, 59
 - deploying for synchronization encryption, 60
 - deploying to devices, 55
 - deployment options, 23
 - describing columns of tables in, 295
 - describing publications in, 297
 - describing tables in, 291
 - desktop creation options, 30
 - displaying system tables, 290
 - embedded engine client files, 23
 - encrypting with the fips creation parameter, 192
 - encryption impact on performance, 120
 - engine and runtime for, 23
 - engine starting, 260
 - engine stopping, 261
 - entity-relationship diagrams, 78
 - environment variables, 49
 - erasing databases, 262
 - fetching rows, 18
 - file internals, 11
 - file path definition, 47
 - file storage on Palm OS, 47
 - filename filter, 258
 - footprint of, 2
 - implementation , 21
 - indexes creating, 460
 - indexes hashing, 194

- indexes hashing overview, 194
- indexes types of, 80
- indexes working with, 79
- indexes, when to create, 81
- indexing primary keys, 32
- initializing from Sybase Central, 30
- initializing from the command prompt, 31
- introduction, 10
- isolation levels, 18
- maintaining row state, 14
- management fundamentals, 13
- management layers of, 10
- managing multiple, 13
- memory usage, 14
- methods of creating, 29
- modeling from MobiLink, 31
- multi-table joins, 2
- obfuscation impact on performance, 120
- objects copy method, 77
- Oracle as reference database for, 127
- page_size creation parameter, 199
- page_size usage, 199
- Palm OS deleting from, 258
- Palm OS support, 47
- populating after running ulinit, 270
- properties browsing, 218
- property optimization strategies, 121
- publications about, 83
- publications dropping, 85
- publishing rows, 84
- publishing tables, 83
- recovery, 14
- recovery from system failure, 15
- renaming objects during schema upgrade, 67
- requests overview, 13
- rollbacks, 14
- row fetching, 18
- row locking, 16
- row packing, 199
- row packing introduction, 70
- row size limit, 70
- rows deleting, 14
- runtime files, 23
- schema, 289
- schema changes, 11
- schema overview, 11
- security overview, 40
- size reduction, 501
- sourcing from XML, 32
- state bytes used, 14
- storing indexes in, 293
- storing properties of, 298
- storing publications in, 296
- supported index types, 108
- supported network protocols, 22
- synchronization introduction, 13
- synchronization profile options, 278
- synchronization utility [ulsync] syntax, 275
- system failure recovery, 15
- table definition tip, 70
- table synchronization suffixes, 72
- tables copying, 77
- tables creating, 71
- tables dropping, 75
- tables filtering, 77
- tables, browsing, 76
- temporary files, 12
- temporary table management, 117
- temporary tables and files, 12
- threads overview, 13
- transactions overview, 13
- troubleshooting, 500, 501, 502
- troubleshooting connections, 499
- ULSQLCONNECT, 49
- UltraLite file property, 213
- UltraLite name property, 213
- UltraLite schema upgrades reference, 450
- UltraLite size limit, 8
- unique keys, 80
- upgrade previous versions, 29
- upgrading schema on devices, 67
- user IDs, 51
- users adding, 87
- users deleting, 88
- UTF8BIN collation for UNICODE characters, 37
- utilities reference , 247
- validating, 15
- viewing option settings, 224
- Windows desktop, 47
- Windows Mobile file paths, 47
- working with, 69
- UltraLite embedded SQL
 - (*see also* UltraLite embedded SQL library functions)
 - character strings, 252
 - CustDB sample and readme files, 97
 - line numbers, 252

-
- preprocessor, 252
 - support, 24
 - UltraLite engine
 - concurrency in, 13
 - contrasting with in-process runtime, 23
 - data management with, 23
 - defining executable location at connection time, 245
 - erase database utility for, 262
 - error codes, 248
 - start utility for, 260
 - stop utility for, 261
 - troubleshooting, 498
 - troubleshooting error -764, 505
 - Windows Mobile deployment, 57
 - UltraLite engine start utility
 - syntax, 260
 - UltraLite engine stop utility
 - syntax, 261
 - UltraLite erase database utility
 - syntax, 262
 - UltraLite for M-Business Anywhere
 - building CustDB application, 98
 - deployment targets, 24
 - engine support, 23
 - UltraLite CustDB sample and readme files, 98
 - UltraLite HotSync conduit
 - deploying, 62
 - UltraLite implementation
 - about, 21
 - UltraLite information utility
 - about, 265
 - UltraLite initialize database utility
 - about, 269
 - UltraLite load XML to database utility
 - syntax, 272
 - UltraLite optimizer
 - execution plan access options, 110
 - UltraLite passwords
 - about, 51
 - UltraLite plug-ins
 - troubleshooting, 499
 - UltraLite publications
 - renaming during schema upgrade, 67
 - UltraLite runtime
 - about, 23
 - concurrency in, 13
 - glossary definition, 536
 - UltraLite schema
 - upgrading on device, 67
 - UltraLite SQL
 - (*see also* UltraLite embedded SQL)
 - comma-separated lists, 395
 - comments, 305
 - CustDB sample and readme files, 97
 - data types, 312
 - dates, 311
 - execution plans for, 342
 - expressions, 325
 - identifiers, 303
 - index-based optimizations, 108
 - keywords, 302
 - NULL values, 307
 - numbers, 306
 - operators, 338
 - page-based optimizations, 117
 - special values, 308
 - SQL functions date and time, 348
 - SQL functions, miscellaneous, 350
 - SQL functions, numeric, 351
 - SQL functions, string, 352
 - strings, 304
 - times, 311
 - troubleshooting queries, 504
 - UltraLite ordering query results with primary key, 117
 - variables, 341
 - UltraLite SQL statements
 - about, 448
 - ALTER DATABASE SCHEMA FROM FILE statement syntax, 450
 - ALTER PUBLICATION statement syntax, 451
 - ALTER SYNCHRONIZATION PROFILE statement syntax, 452
 - ALTER TABLE statement syntax, 454
 - categories of, 449
 - CHECKPOINT statement syntax, 458
 - COMMIT statement syntax, 459
 - CREATE INDEX statement syntax, 460
 - CREATE PUBLICATION statement syntax, 462
 - CREATE SYNCHRONIZATION PROFILE statement syntax, 464
 - CREATE TABLE statement syntax, 468
 - DELETE statement syntax, 473
 - DROP INDEX statement syntax, 474
 - FROM clause, 478

- INSERT statement, 480
- LOAD TABLE statement syntax, 481
- ROLLBACK statement syntax, 485
- SELECT statement syntax, 486
- SET OPTION statement syntax, 488
- START SYNCHRONIZATION DELETE statement syntax, 489
- STOP SYNCHRONIZATION DELETE statement syntax, 490
- SYNCHRONIZE statement syntax, 491
- TRUNCATE TABLE syntax, 493
- UltraLite DROP TABLE statement syntax, 477
- UNION operation syntax, 495
- UNIQUE parameter, 460
- UPDATE statement syntax, 496
- UltraLite synchronization
 - about, 125
 - remote IDs and user IDs, 223
- UltraLite synchronization utility
 - syntax, 275
- UltraLite system table reference
 - about, 289
- UltraLite system tables
 - about, 289
- UltraLite tables
 - excluding from creation process, 269
 - renaming during schema upgrade, 67
- UltraLite temporary files
 - about, 12
- UltraLite temporary tables
 - managing, 117
- UltraLite unload database utility
 - syntax, 281
- UltraLite unload old database utility
 - syntax, 284
- UltraLite user IDs
 - about, 51
 - MobiLink uniqueness, 223
- UltraLite utilities
 - dbisql utility, 249
 - sqlpp utility, 252
 - ulcond11 utility, 263
 - ulcreate utility, 255
 - ULDBUtil utility, 258
 - uleng11 utility, 260
 - ulerase utility, 262
 - ulinfo utility, 265
 - ulinit utility, 269
 - ulload utility, 272
 - ulstop utility, 261
 - ulsync utility, 275
 - ulunload utility, 281
 - ulunloadold utility, 284
 - validate database utility, 286
- UltraLite utilities reference
 - about, 247
- UltraLite Validate Database utility (ulvalid)
 - about, 286
- UltraLite.NET
 - building CustDB application, 98
 - CustDB sample and readme files, 98
 - UltraLite engine support, 23
- ulunload utility
 - syntax, 281
- ulunloadold utility
 - syntax, 284
- ulvalid
 - about, 286
- uncommitted transactions
 - UltraLite isolation level, 18
 - UltraLite overview, 16
- undoing
 - UltraLite transactions, 485
- UNICODE characters
 - UltraLite collation for, 37
- UNION operation
 - UltraLite syntax, 495
- UNION statement
 - UltraLite syntax, 495
- union-all operation
 - UltraLite execution plans, 343
- unions
 - UltraLite multiple select statements, 495
- UNIQUE
 - UltraLite CREATE INDEX parameter, 460
- unique constraints
 - glossary definition, 536
 - UltraLite characteristics, 108
 - UltraLite copying method, 77
 - UltraLite CREATE TABLE statement, 468
- unique indexes
 - UltraLite characteristics, 108
 - UltraLite index creation from, 80
 - UltraLite UNIQUE SQL parameter, 460
- unique keys
 - UltraLite characteristics, 108

- UltraLite index creation from, 80
- UNIQUEIDENTIFIER data type
 - UltraLite, 312
- universally unique identifiers
 - (*see also* UUIDs)
 - UltraLite SQL syntax for NEWID function, 407
- unload
 - glossary definition, 537
- unload old database utility [ulunloadold]
 - syntax, 284
- unloading
 - UltraLite databases, 281
 - UltraLite databases from earlier versions, 284
 - UltraLite databases with ulunload, 281
 - UltraLite databases with ulunloadold, 284
- unnamed foreign keys
 - UltraLite usage, 468
- unpacked rows
 - UltraLite about, 70
- UPDATE statement
 - UltraLite syntax, 496
- updates
 - UltraLite databases, 14
- updating
 - UltraLite updating rows, 496
- UpgradeSchemaFromFile method
 - UltraLite replacement for schema upgrade , 67
- upgrading
 - UltraLite schema error callback, 68
 - UltraLite schema process, 67
 - UltraLite SQL ALTER DATABASE SCHEMA FROM FILE syntax, 450
- upgrading UltraLite
 - troubleshooting connections, 499
- upload ok
 - UltraLite synchronization parameter, 173
- upload only
 - UltraLite synchronization parameter, 174
- upload only synchronization
 - UltraLite databases, 174
 - UltraLite upload_only synchronization parameter, 174
- upload_ok synchronization parameter
 - UltraLite reference, 173
- upload_only synchronization parameter
 - UltraLite reference, 174
- uploads
 - glossary definition, 537
- UPPER function
 - UltraLite syntax, 439
- uppercase characters
 - UltraLite UPPER function, 439
- uppercase strings
 - UltraLite UCASE function, 438
 - UltraLite UPPER function, 439
- user authentication
 - custom MobiLink, 156
 - PWD UltraLite connection parameter, 243
 - UltraLite about, 51
 - UltraLite Authentication Value synchronization parameter, 159
 - UltraLite bypassing, 51
 - UltraLite custom for synchronization, 162
 - UltraLite getUsername method, 175
 - UltraLite password synchronization parameter, 164
 - UltraLite role of, 51
 - UltraLite setup, 51
 - UltraLite synchronization status reports, 157
 - UltraLite user_name synchronization , 175
- user data
 - UltraLite synchronization parameter, 175
- user IDs
 - UltraLite adding new, 51
 - UltraLite changing, 87
 - UltraLite considerations, 87
 - UltraLite databases, 51
 - UltraLite defaults, 87
 - UltraLite semantics, 51
- user name
 - UltraLite synchronization parameter, 175
- user-defined data types
 - glossary definition, 537
 - UltraLite equivalents, 314
 - unsupported in UltraLite, 312
- user_data
 - UltraLite synchronization parameter, 175
- USER_ID function
 - syntax, 440
- user_name
 - UltraLite synchronization parameter, 175
- USER_NAME function
 - syntax, 441
- users
 - UltraLite adding, 87
 - UltraLite deleting, 88
 - UltraLite working with, 87

- utf8_encoding creation parameter
 - UltraLite description, 211
- utf8_encoding database property
 - UltraLite usage, 37
- UTF8BIN collation
 - UltraLite considerations, 37
- utilities
 - ulcreate, 255
 - UltraLite database creation [ulcreate] utility, 255
 - UltraLite engine start [uleng11] utility, 260
 - UltraLite engine stop [ulstop] utility, 261
 - UltraLite erase database [ulerase] utility, 262
 - UltraLite error codes, 248
 - UltraLite HotSync Conduit installer [ulcond11] utility, 263
 - UltraLite information [ulinfo] utility, 265
 - UltraLite initialize database [ulinit] utility, 269
 - UltraLite Interactive SQL [dbisql] syntax, 249
 - UltraLite load XML to database [ulload] utility, 272
 - UltraLite Palm [ULDBUtil] utility, 258
 - UltraLite SQL Preprocessor [sqlpp] utility, 252
 - UltraLite synchronization [ulsync] , 275
 - UltraLite troubleshooting, 503
 - UltraLite unload data to XML [ulunload] utility, 281
 - UltraLite unload old database [ulunloadold] utility, 284
 - Windows Mobile database administration on device, 47
- UUIDs
 - UltraLite SQL syntax for NEWID function, 407
 - UltraLite SQL syntax for STRTOUUID function, 429
 - UltraLite SQL syntax for UUIDTOSTR function, 441
- UUIDTOSTR function
 - UltraLite syntax, 441
- V**
- validate
 - glossary definition, 537
- validate database utility
 - UltraLite about, 286
- validating
 - (*see also* validating databases)
 - UltraLite checksum_level creation parameter, 184
 - UltraLite databases with the validate database wizard, 15
 - UltraLite databases with ulvalid, 286
- validating databases
 - UltraLite validate database utility (ulvalid), 286
- values
 - UltraLite index hash, 194
- VARBINARY data type
 - UltraLite, 312
- VARCHAR data type
 - UltraLite, 312
- variables
 - UltraLite SQL, 341
- version
 - UltraLite synchronization parameter, 176
- version synchronization parameter
 - UltraLite reference, 176
- versions
 - UltraLite troubleshooting utilities, 503
- VFS
 - UltraLite database deployment, 256
 - UltraLite databases, 47
- viewing
 - UltraLite execution plan, 342
 - UltraLite table methods, 76
- viewing UltraLite database settings
 - about, 218
- views
 - glossary definition, 537
- virtual file system (*see* VFS)
- Visual Basic compatibility
 - UltraLite support, 24
- Visual Studio
 - UltraLite building CustDB application, 97
- W**
- WEEKS function
 - UltraLite syntax, 442
- WHEN
 - CASE expression for UltraLite, 327
- WHERE clause
 - UltraLite CREATE PUBLICATION statement, 462
 - UltraLite DELETE statement, 473
 - UltraLite publication usage, 84
 - UltraLite SELECT statement, 486
 - UltraLite synchronization limitations, 137

UltraLite UPDATE statement, 496

whole tables

UltraLite publishing, 137

wildcards

UltraLite PATINDEX function, 411

window (OLAP)

glossary definition, 537

Windows

(*see also* Windows ME)

(*see also* Windows Mobile 5)

(*see also* Windows NT)

(*see also* Windows Vista)

(*see also* Windows XP/200x)

glossary definition, 537

UltraLite character sets, 37

Windows desktop

UltraLite databases, 47

UltraLite engine support, 23

Windows Mobile

glossary definition, 537

troubleshooting error -764, 505

UltraLite ActiveSync deployment, 64

UltraLite building CustDB application using .NET,
98

UltraLite character sets, 37

UltraLite engine deployment , 57

UltraLite engine support, 23

UltraLite file path prefix, 47

UltraLite FIPS enablement, 192

UltraLite MobiLink clients, 150

UltraLite uleng11 deployment, 57

WITH CHECKPOINT clause

UltraLite LOAD TABLE statement, 483

words

UltraLite keywords, 302

UltraLite reserved words, 302

work tables

glossary definition, 537

working with indexes

UltraLite about, 79

X

XML

loading to database, 272

sourcing UltraLite databases from , 32

unloading database to, 275

XML data type

UltraLite equivalent, 314

Y

YEAR function

UltraLite syntax, 444

YEARS function

UltraLite syntax, 444

YMD function

UltraLite syntax, 445

Z

zero-padding

UltraLite date_format creation parameter, 188

UltraLite timestamp_format creation parameter,
208
