



# **MobiLink Server Administration**

**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

<b>About this book .....</b>	<b>xiii</b>
<b>About the SQL Anywhere documentation .....</b>	<b>xiv</b>
<b>Using MobiLink Server Technology .....</b>	<b>1</b>
MobiLink consolidated databases .....	3
Introduction to consolidated databases .....	4
Setting up a consolidated database .....	6
RDBMS-dependent synchronization scripts .....	8
Adaptive Server Enterprise consolidated database .....	10
IBM DB2 LUW consolidated database .....	12
IBM DB2 mainframe consolidated database .....	15
Microsoft SQL Server consolidated database .....	20
MySQL consolidated database .....	22
Oracle consolidated database .....	25
SQL Anywhere consolidated database .....	28
MobiLink server .....	29
Running the MobiLink server .....	30
Stopping the MobiLink server .....	32
Logging MobiLink server actions .....	33
Running the MobiLink server outside the current session .....	35
Running the MobiLink server in a server farm .....	40
Troubleshooting MobiLink server startup .....	41
MobiLink server options .....	43
mlsrv11 syntax .....	45
@data option .....	50
-a option .....	51
-b option .....	52
-bn option .....	53
-c option .....	54
-cm option .....	55
-cn option .....	56
-cr option .....	57

-cs option .....	58
-ct option .....	59
-dl option .....	60
-dr option .....	61
-ds option .....	62
-dsd option .....	63
-dt option .....	64
-e option .....	65
-esu option .....	66
-et option .....	67
-f option .....	68
-fips option .....	69
-fr option .....	70
-ftr option .....	71
-lsc option .....	72
-m option .....	73
-nba option .....	74
-nc option .....	75
-notifier option .....	76
-o option .....	77
-on option .....	78
-oq option .....	79
-os option .....	80
-ot option .....	81
-ppv option .....	82
-q option .....	86
-r option .....	87
-rd option .....	88
-s option .....	89
-sl dnet option .....	90
-sl java option .....	92
-sm option .....	94
-ss option .....	95
-tc option .....	96
-tf option .....	97

---

-tx option .....	98
-ud option .....	99
-ui option .....	100
-ux option .....	101
-v option .....	102
-w option .....	105
-wu option .....	106
-x option .....	107
-xo option .....	113
-zp option .....	118
-zs option .....	119
-zt option .....	120
-zu option .....	121
-zus option .....	122
-zw option .....	123
-zwd option .....	124
-zwe option .....	125
Synchronization techniques .....	127
MobiLink development tips .....	128
Timestamp-based downloads .....	129
Snapshot synchronization .....	133
Partitioning rows among remote databases .....	135
Upload-only and download-only synchronizations .....	138
Maintaining unique primary keys .....	139
Handling conflicts .....	146
Forced conflicts .....	154
Data entry .....	155
Handling deletes .....	156
Handling failed downloads .....	158
Download acknowledgement .....	161
Downloading a result set from a stored procedure call .....	162
Uploading data from self-referencing tables .....	164
MobiLink isolation levels .....	165
MobiLink performance .....	169
Performance tips .....	170
Key factors influencing MobiLink performance .....	174

Monitoring MobiLink performance .....	178
MobiLink Monitor .....	179
Introduction to the MobiLink Monitor .....	180
Starting the MobiLink Monitor .....	181
Using the MobiLink Monitor .....	184
Saving MobiLink Monitor data .....	193
Customizing your statistics .....	194
MobiLink statistical properties .....	195
SQL Anywhere Monitor for MobiLink .....	201
Introducing the SQL Anywhere Monitor .....	202
Monitor quick start .....	205
Tutorial: Using the Monitor .....	206
Start the Monitor .....	211
Exit the Monitor .....	212
Connect to the Monitor .....	213
Disconnect from the Monitor .....	214
Monitoring resources .....	215
Administering resources .....	222
Working with Monitor users .....	228
Alerts .....	232
Installing the SQL Anywhere Monitor on a separate computer .....	236
Troubleshooting the Monitor .....	237
The Relay Server .....	239
Introduction to the Relay Server .....	240
Relay Server configuration file .....	243
Outbound Enabler .....	247
Relay Server State Manager .....	250
Deploying the Relay Server .....	253
Updating a Relay Server farm configuration .....	258
Sybase Relay Server hosting service .....	260
Using MobiLink with the Relay Server .....	262
Redirector (deprecated) .....	265
Introduction to the Redirector (deprecated) .....	266
Setting up the Redirector .....	268
Configuring MobiLink clients and servers for the Redirector .....	269
Configuring Redirector properties .....	271

NSAPI Redirector for Netscape/Sun web servers on Windows (deprecated) ...	277
NSAPI Redirector for Netscape/Sun web servers on Unix (deprecated) .....	280
ISAPI Redirector for Microsoft web servers (deprecated) .....	282
Servlet Redirector (deprecated) .....	284
Apache Redirector (deprecated) .....	287
M-Business Anywhere Redirector (deprecated) .....	289
MobiLink file-based download .....	293
Introduction to file-based download .....	294
Setting up file-based download .....	295
Validation checks .....	298
File-based download examples .....	301

## **MobiLink Events ..... 311**

Writing synchronization scripts .....	313
Introduction to synchronization scripts .....	314
Scripts and the synchronization process .....	317
Script types .....	318
Script parameters .....	320
Script versions .....	324
Required scripts .....	326
Adding and deleting scripts .....	327
Writing scripts to upload rows .....	330
Writing scripts to download rows .....	333
Writing scripts to handle errors .....	338
Synchronization events .....	341
Overview of MobiLink events .....	343
authenticate_file_transfer connection event .....	353
authenticate_parameters connection event .....	355
authenticate_user connection event .....	358
authenticate_user_hashed connection event .....	363
begin_connection connection event .....	367
begin_connection_autocommit connection event .....	368
begin_download connection event .....	369
begin_download table event .....	371
begin_download_deletes table event .....	374

begin_download_rows table event .....	377
begin_publication connection event .....	380
begin_synchronization connection event .....	383
begin_synchronization table event .....	385
begin_upload connection event .....	387
begin_upload table event .....	389
begin_upload_deletes table event .....	391
begin_upload_rows table event .....	394
download_cursor table event .....	396
download_delete_cursor table event .....	400
download_statistics connection event .....	403
download_statistics table event .....	406
end_connection connection event .....	409
end_download connection event .....	411
end_download table event .....	414
end_download_deletes table event .....	417
end_download_rows table event .....	420
end_publication connection event .....	423
end_synchronization connection event .....	426
end_synchronization table event .....	428
end_upload connection event .....	431
end_upload table event .....	433
end_upload_deletes table event .....	436
end_upload_rows table event .....	439
handle_DownloadData connection event .....	442
handle_error connection event .....	446
handle_odbc_error connection event .....	450
handle_UploadData connection event .....	454
modify_error_message connection event .....	460
modify_last_download_timestamp connection event .....	463
modify_next_last_download_timestamp connection event .....	466
modify_user connection event .....	469
nonblocking_download_ack connection event .....	471
prepare_for_download connection event .....	473
publication_nonblocking_download_ack connection event .....	475



---

report_error connection event .....	477
report_odbc_error connection event .....	480
resolve_conflict table event .....	483
synchronization_statistics connection event .....	486
synchronization_statistics table event .....	489
time_statistics connection event .....	492
time_statistics table event .....	495
upload_delete table event .....	498
upload_fetch table event .....	500
upload_fetch_column_conflict table event .....	502
upload_insert table event .....	504
upload_new_row_insert table event .....	506
upload_old_row_insert table event .....	509
upload_statistics connection event .....	512
upload_statistics table event .....	517
upload_update table event .....	522
<b>MobiLink Server APIs .....</b>	<b>525</b>
Writing synchronization scripts in Java .....	527
Introduction to Java synchronization logic .....	528
Setting up Java synchronization logic .....	529
Writing Java synchronization logic .....	531
Java synchronization example .....	538
MobiLink server API for Java reference .....	543
Writing synchronization scripts in .NET .....	589
Introduction to .NET synchronization logic .....	590
Setting up .NET synchronization logic .....	591
Writing .NET synchronization logic .....	593
.NET synchronization techniques .....	600
Loading shared assemblies .....	601
.NET synchronization example .....	604
MobiLink server API for .NET reference .....	606
Direct row handling .....	649
Introduction to direct row handling .....	650
Handling direct uploads .....	654

Handling direct downloads ..... 660

**MobiLink Reference ..... 661**

MobiLink server system procedures ..... 663

    MobiLink system procedures ..... 664

MobiLink utilities ..... 687

    Introduction to MobiLink utilities ..... 688

    MobiLink stop utility (mlstop) ..... 689

    MobiLink user authentication utility (mluser) ..... 690

MobiLink server system tables ..... 693

    Introduction to MobiLink system tables ..... 695

    IBM DB2 mainframe system table name conversions ..... 696

    ml\_active\_remote\_id ..... 697

    ml\_column ..... 698

    ml\_connection\_script ..... 699

    ml\_database ..... 700

    ml\_device ..... 701

    ml\_device\_address ..... 703

    ml\_listening ..... 705

    ml\_passthrough ..... 707

    ml\_passthrough\_repair ..... 708

    ml\_passthrough\_script ..... 709

    ml\_passthrough\_status ..... 711

    ml\_property ..... 712

    ml\_qa\_clients ..... 713

    ml\_qa\_delivery ..... 714

    ml\_qa\_delivery\_archive ..... 716

    ml\_qa\_global\_props ..... 718

    ml\_qa\_notifications ..... 719

    ml\_qa\_repository ..... 720

    ml\_qa\_repository\_archive ..... 721

    ml\_qa\_repository\_props ..... 722

    ml\_qa\_repository\_props\_archive ..... 723

    ml\_qa\_repository\_staging ..... 724

    ml\_qa\_status\_history ..... 725

ml_qa_status_history_archive .....	726
ml_qa_status_staging .....	727
ml_script .....	728
ml_script_version .....	729
ml_scripts_modified .....	730
ml_server .....	731
ml_sis_sync_state .....	732
ml_subscription .....	733
ml_table .....	735
ml_table_script .....	736
ml_user .....	737
MobiLink data mappings between remote and consolidated databases .....	739
Adaptive Server Enterprise data mapping .....	740
IBM DB2 LUW data mapping .....	749
IBM DB2 mainframe data mapping .....	756
Microsoft SQL Server data mapping .....	767
MySQL data mapping .....	774
Oracle data mapping .....	779
Character set considerations .....	789
Character set considerations .....	790
iAnywhere Solutions ODBC drivers for MobiLink .....	793
ODBC drivers supported by MobiLink .....	794
iAnywhere Solutions Oracle driver .....	795
Deploying MobiLink applications .....	799
Introduction to MobiLink deployment .....	800
Deploying the MobiLink server .....	801
Deploying SQL Anywhere MobiLink clients .....	813
Deploying UltraLite MobiLink clients .....	815
Deploying QAnywhere applications .....	816
<b>Glossary .....</b>	<b>821</b>
Glossary .....	823
<b>Index .....</b>	<b>853</b>

---

---

# About this book

## Subject

This book describes how to set up and administer MobiLink servers, consolidated databases and MobiLink applications. It also describes the SQL Anywhere Monitor for MobiLink, a web browser-based administration tool that provides information about the health and availability of MobiLink servers, and the Relay Server, which enables secure communication between mobile devices and MobiLink, Afaria and iAnywhere Mobile Office servers through a web server.

## Audience

This book is for anyone who wants to create distributed information systems. The central data source and remote data stores can be, but are not restricted to, relational database systems.

## Before you begin

For a comparison of MobiLink with other SQL Anywhere synchronization and replication technologies, see [“Comparing synchronization technologies” \[SQL Anywhere 11 - Introduction\]](#).

## 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](mailto: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.

# Using MobiLink Server Technology

This section introduces MobiLink technology and describes how to use it to synchronize data between two or more data sources.

---

MobiLink consolidated databases .....	3
MobiLink server .....	29
MobiLink server options .....	43
Synchronization techniques .....	127
MobiLink performance .....	169
MobiLink Monitor .....	179
SQL Anywhere Monitor for MobiLink .....	201
The Relay Server .....	239
Redirector (deprecated) .....	265
MobiLink file-based download .....	293



---

# MobiLink consolidated databases

## Contents

Introduction to consolidated databases .....	4
Setting up a consolidated database .....	6
RDBMS-dependent synchronization scripts .....	8
Adaptive Server Enterprise consolidated database .....	10
IBM DB2 LUW consolidated database .....	12
IBM DB2 mainframe consolidated database .....	15
Microsoft SQL Server consolidated database .....	20
MySQL consolidated database .....	22
Oracle consolidated database .....	25
SQL Anywhere consolidated database .....	28

---

## Introduction to consolidated databases

Your consolidated database holds system objects that are required by MobiLink. In most cases it also holds your application data, but you can hold all or part of your application data in other forms as well.

MobiLink supports consolidated databases for Windows and Linux on 32-bit and 64-bit environments. Your consolidated database can be one of the following ODBC-compliant RDBMSs:

- Adaptive Server Enterprise (no 64-bit Linux support provided)
- IBM DB2 LUW
- IBM DB2 mainframe
- Microsoft SQL Server
- MySQL
- Oracle
- SQL Anywhere

For version support information, <http://www.sybase.com/detail?id=1002288>.

Your SQL Anywhere installation includes a setup script for each type of RDBMS. You need to run the appropriate setup script to use that RDBMS with MobiLink. The setup script adds tables and stored procedures that are required by MobiLink.

For information about setting up each type of database as a consolidated database, see [“Setting up a consolidated database” on page 6](#).

For information about writing synchronization scripts for particular consolidated databases, see [“RDBMS-dependent synchronization scripts” on page 8](#).

### Synchronizing to other data sources

Your MobiLink environment must have a database that has been set up as a consolidated database. However, you can synchronize data sources other than the consolidated database. The other data sources can be almost anything: a text file, web service, non-relational database, spreadsheet, and so on. You can:

- Create a hybrid application in which you synchronize to both a consolidated database and some other data source.
- Synchronize to only a consolidated database.
- Synchronize to only another data source.

See [“Direct row handling” on page 649](#).

### Restrictions on modifying your consolidated database

Some users find it difficult to change the schema of their consolidated database. For these situations, MobiLink provides solutions, where possible, to keep changes to the consolidated database to a minimum. For example, MobiLink offers a variety of solutions for maintaining unique primary keys, some of which have minimal impact on the consolidated database schema.

In addition, you can avoid almost all impact on your consolidated database by putting your MobiLink system objects in a separate database. See [“MobiLink system database” on page 7](#).



## How remote tables relate to consolidated tables

Your synchronization design specifies mappings between tables and columns in the remote databases with tables and columns in the consolidated database. Typically, tables and columns in remote databases either exactly match the tables and columns in the consolidated database or are subsets of them.

### Arbitrary relationships permitted

Tables in a remote database need not be identical to those in the consolidated database. Synchronized data in one remote application table can be distributed between columns in different tables, and even between tables in different consolidated databases. You specify these relationships using synchronization scripts.

### Direct relationships are simple

The simplest and most common design uses a table structure in the remote database that is a subset of that in the consolidated database. Using this design, every table in the remote database exists in the consolidated database. Corresponding tables have the same structure and foreign key relationships as those in the consolidated database.

The consolidated database frequently contains columns and tables that are not synchronized. Some of these columns or tables may be used for synchronization. For example, a timestamp column can identify new or updated rows in the consolidated database; or a shadow table can be used to track deletes. Non-synchronized columns or tables in the consolidated database can also hold information that is not required at remote sites.

Remote databases also frequently hold tables or columns that aren't synchronized.

### See also

- [“MobiLink data mappings between remote and consolidated databases” on page 739](#)

## Setting up a consolidated database

### Setup scripts

To set up a database so that it can be used as a MobiLink consolidated database, you must run a setup script. Your SQL Anywhere installation includes a script for each of the supported RDBMSs. These scripts are all located in *install-dir\MobiLink\setup*. You can also use the following methods to update the MobiLink system setup:

- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue. See “[Introduction to MobiLink models](#)” [*MobiLink - Getting Started*].

The MobiLink setup script adds MobiLink system tables, stored procedures, triggers, and views to your database. These tables and procedures are required for MobiLink synchronization.

For information about the MobiLink system tables that are installed, see “[MobiLink server system tables](#)” on page 693.

For information about the stored procedures that are installed, see “[MobiLink system procedures](#)” on page 664.

You can view each setup script in a text editor if you want to check what it does.

#### Caution

The database user who runs the setup scripts is given permission to update the MobiLink system tables, which is required to start the MobiLink server and to configure MobiLink. See “[Required permissions](#)” on page 30.

For instructions on how to run the setup scripts, see the section for your RDBMS:

- “[Adaptive Server Enterprise consolidated database](#)” on page 10
- “[IBM DB2 LUW consolidated database](#)” on page 12
- “[IBM DB2 mainframe consolidated database](#)” on page 15
- “[Microsoft SQL Server consolidated database](#)” on page 20
- “[MySQL consolidated database](#)” on page 22
- “[Oracle consolidated database](#)” on page 25
- “[SQL Anywhere consolidated database](#)” on page 28

#### Note

If the consolidated database you are setting up is to be used as a QAnywhere Server Store, the database should be configured to be case insensitive for comparisons and string operations.

## ODBC connection

The MobiLink server needs an ODBC connection to your consolidated database. You must configure the appropriate ODBC driver for your server and create an ODBC data source for the database on the computer where your MobiLink server is running.

For more information about MobiLink ODBC drivers, see [“iAnywhere Solutions ODBC drivers for MobiLink” on page 793](#).

For updated information and complete functional specifications of the ODBC drivers you can use with MobiLink, see [Recommended ODBC Drivers for MobiLink](#).

## MobiLink system database

In some rare cases, you may want to split your consolidated database into two: one database for data and one for the MobiLink system information. When you do this you do not have to add MobiLink system objects to your consolidated database. All MobiLink system objects can be stored in a separate database called the MobiLink system database.

Your MobiLink system database can be any database that is supported as a consolidated database. It does not have to be the same RDBMS as your consolidated database.

It is easy to set up a MobiLink system database. Simply apply MobiLink setup scripts to a database other than your consolidated database. When you start the MobiLink server, connect to both databases:

### Notes

- You can only run the MobiLink server on Windows.
- You cannot use a MobiLink system database with the **MobiLink Create Synchronization Model Wizard** or Model mode.
- There is a performance penalty for storing MobiLink system objects in a separate database.

## RDBMS-dependent synchronization scripts

MobiLink uses synchronization scripts to define the rules you use to synchronize data. Synchronization scripts define:

- How data uploaded from the remote database is to be applied to the consolidated database.
- What data should be downloaded from the consolidated database to the remote database.

See [“Writing synchronization scripts” on page 313](#).

For a complete list of events you can write scripts for, see [“Synchronization events” on page 341](#).

For specific information about each type of consolidated database, see:

- [“SQL Anywhere consolidated database” on page 28](#)
- [“Adaptive Server Enterprise consolidated database” on page 10](#)
- [“IBM DB2 LUW consolidated database” on page 12](#)
- [“IBM DB2 mainframe consolidated database” on page 15](#)
- [“Microsoft SQL Server consolidated database” on page 20](#)
- [“MySQL consolidated database” on page 22](#)
- [“Oracle consolidated database” on page 25](#)

### .NET and Java synchronization scripts

You can write your synchronization logic in the version of the SQL language used by your database. You can also write more portable and powerful scripts using Java or .NET. Both Java and .NET offer flexibility beyond what each RDBMS provides via SQL, while also providing full SQL compatibility. When you use Java or .NET synchronization logic, you can hold session-wide variables, create user-defined procedures, integrate authentication to external servers, and so on.

For information about Java synchronization logic, see [“Writing Java synchronization logic” on page 531](#).

For information about .NET synchronization logic, see [“Writing synchronization scripts in .NET” on page 589](#).

### Invoking procedures from scripts

Some databases, such as Microsoft SQL Server, require that procedure calls with parameters be written using the ODBC syntax.

```
{ CALL procedure_name( {ml param1}, {ml param2}, ... ) }
```

You can return values by defining the parameters as OUT or INOUT in the procedure definition.

### CHAR columns

In many other RDBMSs, CHAR data types are fixed length and blank-padded to the full length of the string. In SQL Anywhere or UltraLite remote MobiLink databases, CHAR is the same as VARCHAR: values are not blank-padded to a fixed width. If you are not using SQL Anywhere as your consolidated database, it is strongly recommended that you use VARCHAR in the consolidated database rather than CHAR. If you must use CHAR, the `mlsrv11 -b` command line option can be used to remove trailing blanks from strings during synchronization. This option is important for string comparisons used to detect conflicts.

See [“-b option” on page 52](#).

### **Data conversion**

For information about the conversion of data that must take place when a MobiLink server communicates with a consolidated database that isn't SQL Anywhere, see [“MobiLink data mappings between remote and consolidated databases” on page 739](#).

## Adaptive Server Enterprise consolidated database

### Setting up Adaptive Server Enterprise as a consolidated database

To set up Adaptive Server Enterprise to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are multiple ways you can do this:

- Run the *syncase.sql* setup script, located in *install-dir\MobiLink\setup*.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue.

#### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, `SELECT * from ml_user`). See [“MobiLink server system tables” on page 693](#).

### ODBC driver

You must set up an ODBC DSN for your Adaptive Server Enterprise consolidated database using the ODBC driver that is provided with your Adaptive Server Enterprise database. See:

- [Recommended ODBC Drivers for MobiLink](#)
- Your Adaptive Server Enterprise documentation

### Adaptive Server Enterprise issues

- **BLOB sizes** To download BLOB data with a data size greater than 32 KB (the default), do the following:
  - On Windows, set **Text Size** on the **Advanced** page of the **Adaptive Server Enterprise ODBC Driver Configuration** window to be greater than the largest expected BLOB.
  - On Linux, set the `TextSize` entry in the *obdc.ini* file to be greater than the largest expected BLOB.
- **CHAR columns** In Adaptive Server Enterprise, CHAR data types are fixed length and blank-padded to the full length of the string. In MobiLink remote databases (SQL Anywhere or UltraLite), CHAR is the same as VARCHAR: values are not blank-padded to a fixed width. It is strongly recommended that you use VARCHAR in the consolidated database rather than CHAR. If you must use CHAR, the `mlsrv11`

-b command line option can be used to remove trailing blanks from strings during synchronization. This option is important for string comparisons used to detect conflicts.

For more information, see [“-b option” on page 52](#).

- **Data type mapping** The data types of columns must map correctly between your consolidated and remote database. See [“Adaptive Server Enterprise data mapping” on page 740](#).
- **Special considerations for version 11.5 and earlier** You cannot use MobiLink system procedures such as `ml_add_connection_script` to add scripts longer than 255 bytes to Adaptive Server Enterprise 11.5 or earlier. To define longer scripts, use Sybase Central or direct insertion.
- **Restrictions on VARBIT** MobiLink does not support synchronizing 0 length (empty) VARBIT or LONG VARBIT values to an Adaptive Server Enterprise consolidated database. Adaptive Server Enterprise does not support a VARBIT type so these types would normally be synchronized to a VARCHAR or TEXT column on the Adaptive Server Enterprise database. On Adaptive Server Enterprise, empty string values are converted into a single space. A space is not allowed in a VARBIT column on SQL Anywhere, so an attempt to download these values causes an error on the remote database.

### Isolation level

See [“MobiLink isolation levels” on page 165](#).

## IBM DB2 LUW consolidated database

MobiLink supports IBM DB2 LUW for Linux, Unix, and Windows. MobiLink does not support IBM DB2 for AS/400.

### Setting up DB2 LUW as a consolidated database

To set up DB2 to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are multiple ways you can do this:

- Run the *syncdb2.sql* setup script, located in *install-dir\MobiLink\setup*. Before running the file, you must copy it to another location and modify it. Instructions follow.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue. Note that you are required to perform step 1 of the following procedure.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue. Note that you are required to perform step 1 of the following procedure.

#### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, `SELECT * from ml_user`). See [“MobiLink server system tables” on page 693](#).

### To run the DB2 setup script

1. To install MobiLink system tables using the setup script, an IBM DB2 LUW tablespace must use a minimum of 8 KB pages. If a tablespace does not use 8 KB pages, complete the following steps:
  - Verify that at least one of your buffer pools has 8 KB pages. If not, create a buffer pool with 8 KB pages.
  - Create a new tablespace and temporary tablespace that use the buffer pool with 8 KB pages.  
For more information, consult your DB2 LUW documentation.
2. Customize *syncdb2.sql* with your connection information:
  - a. Copy *syncdb2.sql* to a new location where it can be modified and stored.
  - b. The *syncdb2.sql* script contains a default connection statement, `connect to DB2Database`. Alter this line to connect to your DB2 database. Use the following syntax:

```
connect to DB2Database user userid using password ~
```



where *DB2Database*, *userid*, and *password* are names you provide. (The *syncdb2.sql* script uses the tilde character (~) as a command delimiter.)

3. Run *syncdb2.sql*:

```
db2 -c -ec -td~ +s -v -f syncdb2.sql
```

## ODBC driver

You must set up an ODBC DSN for your DB2 consolidated database using the ODBC driver that is provided with your DB2 database. See:

- [Recommended ODBC Drivers for MobiLink](#)
- IBM DB2 LUW documentation

## DB2 LUW issues

- **Tablespace capacity** A tablespace and temporary tablespace of any DB2 LUW database that you want to use as a consolidated database must use 8 KB pages.

In addition, there are columns that require a LONG tablespace. If there is no default LONG tablespace, the creation statements for the tables containing these columns must be qualified appropriately, as in the following example:

```
CREATE TABLE ... ( ... )
IN tablespace
LONG IN long-tablespace
```

For an example using the sample application, see [“Exploring the CustDB sample for MobiLink” \[MobiLink - Getting Started\]](#).

- **Session-wide variables** DB2 LUW prior to version 8 does not support session-wide variables. A convenient solution is to use a base table with columns for the MobiLink user name and other session data. The base table has rows representing concurrent synchronizations.
- **User-defined procedures** DB2 LUW prior to version 8.2 requires that you compile SQL procedures into an executable library (such as a DLL). The resulting DLL/shared library must be copied to a special directory on the server. Note that you can write procedures using several different languages, including C/C++ and Java, among others.

For more information about Java and .NET synchronization scripts, see:

- [“Writing synchronization scripts in Java” on page 527](#)
- [“Writing synchronization scripts in .NET” on page 589](#)
- **CHAR columns** In IBM DB2 LUW, CHAR data types are fixed length and blank-padded to the full length of the string. In MobiLink remote databases (SQL Anywhere or UltraLite) CHAR is the same as VARCHAR: values are not blank-padded to a fixed width. It is strongly recommended that you use VARCHAR in the consolidated database rather than CHAR. If you must use CHAR, the `mlsrv11 -b` command line option can be used to remove trailing blanks from strings during synchronization. This option is important for string comparisons used to detect conflicts.

See [“-b option” on page 52](#).

- **Data type mapping** The data types of columns must map correctly between your consolidated and remote database. For details, see [“IBM DB2 LUW data mapping” on page 749](#).
- **Double up the quotation marks in system procedure calls** When you use a MobiLink system procedure to add scripts to your DB2 consolidated database, you need to double up the quotation marks. For example, if the script you are adding with `ml_add_table_script` includes the line `SET "DELETED" = ' 'Y' '` for any other consolidated database, for DB2 you would have to write this as `SET "DELETED" = ' ' 'Y' ' ' ' .`
- **Special considerations for version 5 and earlier** If you are using IBM DB2 LUW prior to version 6, column names and other identifiers are only supported up to 18 characters. This means that you must truncate the names of MobiLink system procedures. For example, to call `ml_add_connection_script`, use the name `ml_add_connection_`.

### Isolation level

See [“MobiLink isolation levels” on page 165](#).

# IBM DB2 mainframe consolidated database

## Setting up DB2 mainframe as a consolidated database

To set up DB2 mainframe to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. You can perform this task using SQL or JCL methods.

### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, SELECT \* from ml\_user). See [“MobiLink server system tables” on page 693](#).

## Common steps to set up the DB2 mainframe environment

1. Create a buffer pool with a page size of 8K or greater and row-level locking for MobiLink schema. Row-level locking is required to handle concurrent synchronizations to the same tables. For this example, name the buffer pool *BP8K*.
2. Create a tablespace named *MLTB8K* with an 8K page size buffer pool for MobiLink schema. For example:

```
create tablespace MLTB8K in IANY bufferpool BP8K locksize row
grant use of tablespace IANY.MLTB8K to public
```

3. If you do not already have one, create a Workload Manager environment for MobiLink schema procedures and name is something like *MLWLM*.
4. Set up an ODBC DSN for your DB2 mainframe consolidated database using the ODBC driver that is provided with your DB2 database. See:
  - [Recommended ODBC Drivers for MobiLink](#)
  - IBM DB2 mainframe documentation

## To create MobiLink system tables using SQL

### Note

The SQL method requires the ability to create stored procedures using DSNTPSMP. If you do not have SQL stored procedures enabled, use the JCL technique.

1. Set up the DB2 mainframe environment using the common steps listed in [“Setting up DB2 mainframe as a consolidated database” on page 15](#).
2. Modify the *syncd2m.sql* setup script, located in *install-dir\MobiLink\setup*.

**Note**

Be sure to make a backup copy of the original *syncd2m.sql* file before you proceed.

In the *syncd2m.sql* file, replace all occurrences of

- {MLTABLESPACE} with MLTB8K, the name of the tablespace.
- {WLMENV} with MLWLM, the Workload Manager.

3. Run the *syncd2m.sql* setup script using the following command line:

```
dbisql -c "uid=user-id;pwd=password;DSN=dsn-name" -nogui syncd2m.sql
```

The message log file, *syncd2m.txt*, gets generated.

4. Open *syncd2m.txt* to verify that the DSNTPSMP calls succeeded.

**To create the MobiLink system tables using JCL**

1. Set up the DB2 mainframe environment using the common steps listed in [“Setting up DB2 mainframe as a consolidated database” on page 15](#).
2. Modify the *syncd2m\_jcl.sql* script, located in *install-dir\MobiLink\setup*.

**Note**

Be sure to make a backup copy of the original *syncd2m\_jcl.sql* file before you proceed.

In the *syncd2m\_jcl.sql* file, replace all occurrences of

- {MLTABLESPACE} with your qualified tablespace, for example MYDB.MYTS.
- {WLMENV} with the name of a Workload Manager associated with your DB2 instance.

3. Start DBISQL and connect to DB2 mainframe.
4. Run the edited copy of the *syncd2m\_jcl.sql* setup script, located in *install-dir\MobiLink\setup*, to create Mobilink tables and define Mobilink procedures in the DB2 mainframe.
5. From the *%SQLANY%\MobiLink\setup* directory, FTP to your mainframe and run the following commands:

```
bin
  hash
  cd xmit
  quote site recfm=fb lrecl=80
  quote site cyl
  put d2mload.xmit
  put d2mdbrm.xmit
  quit
```

6. The two xmit files on the mainframe are as follows:
  - USERID.XMIT.D2MLOAD.XMIT
  - USERID.XMIT.D2MDBRM.XMIT

USERID is the username you gave when connecting via FTP.

7. Open a terminal session and run the following commands from the ISPF Command Shell:

```
RECEIVE INDATASET( 'USERID.XMIT.D2MLOAD.XMIT' )
RECEIVE INDATASET( 'USERID.XMIT.D2MDBRM.XMIT' )
```

The output is as follows:

- USERID.ML.LOADLIB
  - USERID.ML.DBRMLIB
8. Copy the *d2mrelod.jcl* file and modify it as follows:
- Change USERID to your mainframe userid.
  - Change DSNDB0T to your DB2 DSN.
9. Run the edited copy of the *d2mrelod.jcl* script, located in *install-dir\MobiLink\setup*.
10. Copy the *d2mbdpk.jcl* file and modify it as follows:
- Change USERID to your mainframe userid.
  - Change DB0T to your DB2 SSID.
11. Bind all SQL procedures by running the edited copy of *d2mbdpk.jcl*. The following is a reference of SQL procedure mappings to load module names.

Procedure name	Load module name
ml_add_user	mlaub
ml_delete_user	mldub
ml_del_sstate	mldssb
ml_reset_sstate	mlrspb
ml_del_sstate_b4	mldssbb
ml_add_lcs_chk	mlalcsb
ml_add_lcs	mlalcsb
ml_add_cs	mlacsb
ml_add_jcs	mlajcsb
ml_add_dcs	mladcsb
ml_add_lts_chk	mlaltsb
ml_add_lts	mlaltsb
ml_add_ts	mlatsb

Procedure name	Load module name
ml_add_jts	mlajtsb
ml_add_dts	mladtsb
ml_add_property	mlapb
ml_add_column	mlacb
ml_set_device	mlsdb
ml_set_device_nt	mlsdbn
ml_set_dev_addr	mlsdab
ml_set_dev_addr_int	mlsdanb
ml_set_listening	mlslb
ml_set_listen_nt	mlslbn
ml_set_sis_sstate	mlsssb
ml_del_dev_addr	mlddb
ml_del_listen	mlldb
ml_delete_device	mlddb

**DB2 mainframe known issues**

- **DB2 mainframe does not work with Model mode** You can not use DB2 mainframe as your consolidated database when you use the **Create Synchronization Model Wizard**.
- **SELECT statements require the FOR READ ONLY clause** SELECT statements in DB2 mainframe are opened for update by default, meaning that the database acquires write locks with the anticipation of an UPDATE statement after the SELECT statement.

To avoid the write locks and enhance concurrency, append FOR READ ONLY on all SELECT statements that do not precede UPDATE statements. Use FOR READ ONLY in SELECT statements as often as possible, specifically in the download\_cursor and download\_delete\_cursor scripts.

- **Sysplex requires time synchronization** When the DB2 mainframe consolidated database is running in a Sysplex, the clocks of all LPARs in the Sysplex must be synchronized. Failure to synchronize the clocks could result in lost data during database synchronization.
- **Numbers are approximated** Approximate numbers have different possible values. The following is a table of examples.

---

Type	Entered value	DB2 mainframe value	ASA value
Real	123.456	123.4559936523	123.4560012817
Float	123.456	123.45599999999999	123.4560012817
Double	123.456	123.45599999999999	123.456

The recommended approach is to avoid synchronization of double and floating point columns with a DB2 mainframe consolidated database.

**Isolation level**

See [“MobiLink isolation levels”](#) on page 165.

## Microsoft SQL Server consolidated database

### Setting up Microsoft SQL Server as a consolidated database

**Note**

The database user that runs the setup script must be able to create tables, triggers, and stored procedures, so must have the db\_owner role.

To set up Microsoft SQL Server to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are multiple ways you can do this:

- Run the *syncmss.sql* setup script, located in *install-dir\MobiLink\setup*.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin**; connect to your server database; right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue. If you want to use an existing MobiLink system setup, then your default\_schema should be the schema of the MobiLink system setup.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue.

**Note**

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, SELECT \* from ml\_user). See [“MobiLink server system tables” on page 693](#).

### ODBC driver

You must set up an ODBC DSN for your SQL Server consolidated database using the ODBC driver that is provided with your SQL Server database. See:

- [Recommended ODBC Drivers for MobiLink](#)
- Microsoft SQL Server documentation

### SQL Server issues

- **SET NOCOUNT ON** For Microsoft SQL Server, you should specify SET NOCOUNT ON as the first statement in all stored procedures or SQL batches executed via ODBC. Without this option, network buffers can overflow, silently losing data. This is a known SQL Server problem.
- **Procedure calls** Microsoft SQL Server requires that procedure calls with parameters be written using the ODBC syntax:

```
{ CALL procedure_name( {ml param1}, {ml param2}, ... ) }
```



- **CHAR columns** In Microsoft SQL Server, CHAR data types are fixed length and blank-padded to the full length of the string. In MobiLink remote databases (SQL Anywhere or UltraLite) CHAR is the same as VARCHAR: values are not blank-padded to a fixed width. We strongly recommend that you use VARCHAR in the consolidated database rather than CHAR. If you must use CHAR, the mlsrv11 - b command line option can be used to remove trailing blanks from strings during synchronization. This option is important for string comparisons used to detect conflicts.

See [“-b option” on page 52](#).

- **Data type mapping** The data types of columns must map correctly between your consolidated and remote database. For details, see [“Microsoft SQL Server data mapping” on page 767](#).
- **Sample database issues** The SQL Server AdventureWorks sample database contains computed columns. You can't synchronize a computed column. You can set the column to be download-only, or you can exclude the column from synchronization.
- **Implementing conflict detection in an upload\_update script** The behavior of the SQL Server NOCOUNT option means that sometimes the MobiLink server cannot accurately assess how many rows were changed by an upload script. For SQL Server, it is safer to implement a stored procedure in the upload\_update script for conflict detection.

#### Isolation level

See [“MobiLink isolation levels” on page 165](#).

## MySQL consolidated database

The MobiLink server supports MySQL Community and Enterprise servers 5.1.22 or later. QAnywhere and MobiLink models do not support MySQL.

### Setting up MySQL as a consolidated database

To set up MySQL to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are two ways you can do this:

- Using the MySQL command line tool or the MySQL Query Browser, run the *syncmys.sql* setup script, located in *install-dir\MobiLink\setup*. Make sure that your MySQL user ID has privileges to create triggers.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue. Note that if you want to use an existing MobiLink system setup, then your *default\_schema* should be the schema of the MobiLink system setup.

#### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, `SELECT * from ml_user`). See [“MobiLink server system tables” on page 693](#).

### ODBC driver

You must set up an ODBC DSN for your MySQL consolidated database using the ODBC driver that is provided on the MySQL web site. The MobiLink server supports MySQL ODBC driver 5.1.3 or later. See:

- [Recommended ODBC Drivers for MobiLink](#)

To specify your ODBC configuration file in Unix, do one of the following,

- Place the *ODBC.INI* file into the home directory of the current user.
- Create an **ODBCINI** environment variable and set it to the directory location of the *ODBC.INI* file.

If any of your synchronization scripts contain batched SQL commands separated by semicolons, you may need to select the **Allow Multiple Statements** check box on the **Flags 3** page of the **MySQL Connector/ODBC Data Source Configuration** window when you configure a DSN for the MobiLink server to make connections to your MySQL database.

## MySQL issues

- **Storage Engine** The MobiLink server requires the default storage engine to be ACID compliant. If the default storage engine is not ACID compliant, make sure that all MobiLink server tables are created using an ACID compliant storage engine, such as InnoDB and Falcon.

- **Stored Procedures** You cannot use INOUT or OUT parameters in stored procedure calls. Procedures that require these parameters must be implemented as functions that return an OUT value.

Server events that require INOUT parameters, such as `authenticate_user` and `modify_user`, must be implemented as functions and run using a `SELECT` statement instead of a `CALL` statement.

Since user-defined named parameters are not modified after server events run, they are not supported.

- **Cursor Scripts** The events `upload_fetch`, `download_cursor`, and `download_delete_cursor` must be called using a `SELECT` statement, which the MobiLink server runs using a read-committed isolation level. A bug in the MySQL ODBC driver causes the server to read uncommitted operations, such as `INSERT`, `UPDATE`, and `DELETE` statements, which results in inconsistent data between the consolidated database and the remote database.

To work around this problem, affix a `LOCK IN SHARE MODE` clause to your `SELECT` statements. For example,

```
SELECT column1 FROM table1 WHERE id > 0 LOCK IN SHARE MODE
```

This clause protects the `SELECT` statement from uncommitted operations.

- **Bulk upload** The MobiLink server relies on the MySQL ODBC driver, which does not currently support bulk upload.
- **MLSD** The MobiLink server relies on the MySQL ODBC driver, which does not currently support MSDTC.
- **SQLEN Datatypes on the 64-bit MobiLink server for Unix** The MySQL ODBC driver defines `SQLEN` as a 32-bit integer, causing a discrepancy with the 64-bit MobiLink server, which defines `SQLEN` as a 64-bit integer. If you are running MobiLink on a 64-bit Unix environment, you must add the following to your ODBC configuration file,

```
length32=1
```

This entry forces the server to read `SQLEN` as a 32-bit integer. Your configuration should look similar to the following example,

```
[a_mysql_dsn]
Driver=full_path/libmyodbc5.so
server=host_name
uid=user_name
pwd=user_password
database=database_name
length32=1
```

- **MySQL Server Configuration** The MobiLink synchronization scripts are stored in the `mL_script` table as `TEXT` and are retrieved when needed. You may need to set `max_allowed_packet` equal to 16m or greater in the `my.ini` file.

**Isolation level**

See [“MobiLink isolation levels”](#) on page 165.

# Oracle consolidated database

## Setting up Oracle as a consolidated database

To set up Oracle to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are multiple ways you can do this:

- Run the *syncora.sql* setup script, located in *install-dir\MobiLink\setup*.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue. Note that if you have set up aliases for an existing MobiLink system setup, you should connect as the user whose schema has the MobiLink system setup.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue.

### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, `SELECT * from ml_user`). See [“MobiLink server system tables” on page 693](#).

## ODBC driver

You must set up an ODBC DSN for your Oracle consolidated database. See:

- [“iAnywhere Solutions Oracle driver” on page 795](#)
- [Recommended ODBC Drivers for MobiLink](#)

## Oracle issues

- **Stored procedures** If you are using stored procedures in Oracle, you must select the Procedure Returns Results option for the Oracle ODBC driver.  
See [“iAnywhere Solutions Oracle driver” on page 795](#).
- **Session-wide variables** Oracle does not provide session-wide variables. You can store session-wide information in variables within Oracle packages. Oracle packages allow variables to be created, modified and destroyed; these variables may last as long as the Oracle package is current.
- **Autoincrement methods** To maintain primary key uniqueness, you can use an Oracle sequence to generate a list of keys similar to that of an autoincrement field. The CustDB sample database provides coding examples, which can be found in *Samples\MobiLink\CustDB\custora.sql*. Unlike autoincrement, however, you must explicitly reference the sequence. Autoincrement inserts a column value automatically if the column is not referenced in an INSERT statement.

- **Oracle does not support empty strings** In Oracle, an empty string is treated as null. In SQL Anywhere and UltraLite, empty strings have a different meaning from null. Therefore, you should avoid using empty strings in your client databases when you have an Oracle consolidated database.
- **CHAR columns** In Oracle, CHAR data types are fixed length and blank-padded to the full length of the string. In MobiLink remote databases (SQL Anywhere or UltraLite), CHAR is the same as VARCHAR: values are not blank-padded to a fixed width. It is strongly recommended that you use VARCHAR in the consolidated database rather than CHAR. If you must use CHAR, the `mlsrv11 -b` command line option can be used to remove trailing blanks from strings during synchronization. This option is important for string comparisons used to detect conflicts.

See “[-b option](#)” on page 52.

- **Data type mapping** The data types of columns must map correctly between your consolidated and remote database. For details, see “[Oracle data mapping](#)” on page 779.

### Isolation level

See “[MobiLink isolation levels](#)” on page 165.

## Using Oracle varray

The iAnywhere Solutions 11 - ODBC driver for Oracle supports the use of Oracle varray in stored procedures. Using varray in upload scripts (`upload_insert`, `upload_update`, and `upload_delete`) that are written in stored procedures may improve performance of the MobiLink server, compared with upload scripts written in stored procedures that do not use varray. Simple SQL statements such as INSERT, UPDATE and DELETE without stored procedures usually offer the best performance, however using stored procedures, including the varray technique, provides an opportunity to apply business logic that the simple statements do not.

### varray example

The following is a simple example that uses varray:

1. Create a table called `my_table` that contains 3 columns.

```
create table my_table ( pk integer primary key, c1 number(20), c2
varchar2(4000) )
```

2. Create user-defined collection types using varrays.

```
create type my_integer is varray(100) of integer;
create type my_number is varray(100) of number(20);
create type my_varchar is varray(100) of varchar2(8000);
```

`my_varchar` is defined as a varray that contains 100 elements and each element is a data type of `varchar2` and width of 8000. The width is twice as big as that specified for `my_table`.

3. Create stored procedures for insert.

```
create or replace procedure my_insert_proc( pk_v my_integer, c1_v
my_number, c2_v my_varchar )
is
c2_value my_varchar;
begin
```

```
        c2_value := c2_v; -- Work around an Oracle bug
    FORALL i in 1 .. pk_v.COUNT
        insert into my_table ( pk, c1, c2 ) values( pk_v(i), c1_v(i),
c2_value(i) );
    end;
```

## varray restrictions

The following restrictions apply when using varray in stored procedures:

- The DSN must have the **Enable Microsoft distributed transactions** checkbox unchecked.
- BLOB and CLOB varrays are not supported.
- A stored procedure containing a varray must be a standalone procedure, not a packaged procedure.
- If varray is a data type of CHAR, VARCHAR, NCHAR or NVARCHAR, the user-defined varray type must be twice as big as the length specified for the table column.
- The number of rows in the varray that are sent by the MobiLink server to the Oracle consolidated database is set by the -s option, not the size of the varray declared in the varray type. The -s option must not be bigger than the smallest varray type size in use by synchronization scripts. If it is bigger, the MobiLink server issues an error. See [“-s option” on page 89](#).

## SQL Anywhere consolidated database

### Setting up SQL Anywhere as a consolidated database

To set up SQL Anywhere to work as a MobiLink consolidated database, you must run a setup procedure that adds MobiLink system tables, stored procedures, triggers, and views that are required for MobiLink synchronization. There are multiple ways you can do this:

- Run the *syncsa.sql* setup script, located in *install-dir\MobiLink\setup*.
- In the MobiLink plug-in for Sybase Central, choose **Mode » Admin** and connect to your server database. Right-click the database name and choose **Check MobiLink System Setup**. If your database requires setup, you are prompted to continue.
- When you use the **Create Synchronization Model Wizard** or **Deploy Synchronization Model Wizard**, system setup is checked when you connect to your server database. If your database requires setup, you are prompted to continue.

#### Note

The database user who runs the setup script is the only user who has permission to change the MobiLink system tables, which is required for configuring MobiLink applications. See [“Required permissions” on page 30](#).

The RDBMS user that the MobiLink server uses to connect to the consolidated database must be able to use the MobiLink system tables, procedures, and so on, without any qualifiers (for example, `SELECT * from ml_user`). See [“MobiLink server system tables” on page 693](#).

### Setting up the ODBC driver

You must set up an ODBC DSN for your SQL Anywhere consolidated database. The ODBC driver for SQL Anywhere is installed with SQL Anywhere.

For information about the SQL Anywhere ODBC driver, see [“Creating ODBC data sources” \[SQL Anywhere Server - Database Administration\]](#).

### Isolation level

See [“MobiLink isolation levels” on page 165](#).



---

# MobiLink server

## Contents

- Running the MobiLink server ..... 30
- Stopping the MobiLink server ..... 32
- Logging MobiLink server actions ..... 33
- Running the MobiLink server outside the current session ..... 35
- Running the MobiLink server in a server farm ..... 40
- Troubleshooting MobiLink server startup ..... 41

---

## Running the MobiLink server

All MobiLink clients synchronize through the MobiLink server. None connect directly to a database server. You must start the MobiLink server before a MobiLink client synchronizes.

For a list of `mlsrv11` command line options, see [“MobiLink server options” on page 43](#).

The MobiLink server opens connections, via ODBC, with your consolidated database server. It then accepts connections from remote applications and controls the synchronization process.

### To start the MobiLink server

- Run `mlsrv11`. Use the `-c` option to specify the ODBC connection parameters for your consolidated database.

For information about connection parameters, see [“-c option” on page 54](#).

You must specify connection parameters. Other options are available, but are optional. These options allow you to specify how the server works. For example, you can specify a cache size and logging options.

For more information about `mlsrv11` options, see [“mlsrv11 syntax” on page 45](#).

#### Note

The `mlsrv11` options allow you to specify how the MobiLink server works. To control what the server does, you define scripts that are invoked at synchronization events. See [“Synchronization events” on page 341](#).

### Example

The following command starts the MobiLink server using the ODBC data source *SQL Anywhere 11 CustDB* to identify the consolidated database. Enter the entire command on one line.

```
mlsrv11
-c "dsn=SQL Anywhere 11 CustDB;uid=DBA;pwd=sql"
-zs MyServer
-o mlsrv.log
-vcr
-x tcpip(port=3303)
-xo tcpip
```

In this example, the `-c` option provides a connection string that contains an ODBC data source name (DSN) and authentication. The `-zs` option provides a server name. The `-o` option specifies that the log file should be named *mlsrv.log*. The contents of *mlsrv.log* are verbose because of the `-vcr` option. The `-x` option opens a port for version 10 and 11 clients, and the `-xo` option opens a port for version 8 and 9 clients. You must specify a port with either the `-x` option or the `-xo` option, otherwise the command fails because the default port is used for both options.

You can also start the MobiLink server as a Windows service or Unix daemon. See [“Running the MobiLink server outside the current session” on page 35](#).

### Required permissions

You must specify a database user for the MobiLink server to connect to the database server. You specify the database user with the `mlsrv11 -c` option or in the DSN.

This database user must have full select, insert, update, and delete permissions on the MobiLink system tables, and must also have execute permissions on the MobiLink system procedures. By default, the database user who runs the MobiLink setup script has these permissions. If you want to use another database user to run the MobiLink server, you must grant permissions for that user on the ml\_\* tables and the ml\_add\_\*\_script system procedures.

For example, in a SQL Anywhere consolidated database you can grant the required permissions as follows:

```
CREATE USER DBUser IDENTIFIED BY SQL;  
GRANT ALL ON dbo.ml_user to DBUser;  
...  
GRANT EXECUTE ON dbo.ml_add_table_script TO DBUser;  
...
```

You must grant permission for each MobiLink system table and system procedure. For a list of all MobiLink system tables and system procedures, see [“MobiLink server system tables” on page 693](#) and [“MobiLink server system procedures” on page 663](#).

The database user also needs the appropriate permission on all tables referenced in the MobiLink scripts, and execute permissions on any procedures referenced in the MobiLink scripts.

For more information about setting permissions, see [“GRANT statement” \[SQL Anywhere Server - SQL Reference\]](#).

For more information about setup scripts, see [“Setting up a consolidated database” on page 6](#).

## Stopping the MobiLink server

The MobiLink server is stopped from the computer where the server was started. You can stop the MobiLink server in the following ways:

- Use the MobiLink stop utility (mlstop).
- Click **Shut down** on the MobiLink server window.
- On Windows, right-click the icon in the system tray and choose **Shut down**.
- When running on Unix without the MobiLink server window, type Q.
- Use the shutdown method in the MobiLink server API.

### See also

- [“MobiLink stop utility \(mlstop\)” on page 689](#)
- server API for Java: [“shutdown method” on page 573](#)
- server API for .NET: [“ShutDown method” on page 630](#)

## Logging MobiLink server actions

Logging the actions that the server takes is particularly useful during the development process and when troubleshooting. Verbose output is not recommended for normal operation of a production environment because it can slow performance.

### Logging output to a file

Logging output is sent to the MobiLink server messages window. In addition, you can send the output to a message log file using the `-o` option. The following command sends output to a message log file named *mlsrv.log*.

```
mlsrv11 -o mlsrv.log -c ...
```

You can control the size of log files, and specify what you want done when a file reaches its maximum size.

- Use the `-o` option to specify that a log file should be used.
- Use the `-ot` option to specify that a log file should be used when you want the previous contents of the file to be deleted before messages are sent to it.
- In addition to `-o` or `-ot`, use the `-on` option to specify the size at which the log file is renamed with the extension `.old` and a new file is started with the original name.
- In addition to `-o` or `-ot`, use the `-os` option to specify the size at which a new log file is started with a new name based on the date and a sequential number.

See:

- [“-o option” on page 77](#)
- [“-on option” on page 78](#)
- [“-os option” on page 80](#)
- [“-ot option” on page 81](#)

### Controlling the amount of logging output

You can control what information is logged to the message log file and displayed in the MobiLink server window using the `-v` option. See [“-v option” on page 102](#).

### Controlling which warning messages are reported

You can also control which warning messages are reported.

For more information, see:

- [“-zw option” on page 123](#)
- [“-zwd option” on page 124](#)
- [“-zwe option” on page 125](#)

## Viewing MobiLink server logs

You can view MobiLink logs in the following ways:

- In the MobiLink server messages window
- By opening the log file
- Using the MobiLink Server Log File Viewer in Sybase Central

To view log information outside the MobiLink server messages window, you must log the information to a file. See [“Logging output to a file” on page 33](#).

### MobiLink Server Log File Viewer

To view MobiLink server logs, open Sybase Central and choose **Tools » MobiLink 11 » MobiLink Server Log File Viewer**. You are prompted to choose a log file to view. By holding down the shift key, you can open multiple log files at the same time.

The MobiLink Server Log File Viewer reads information that is stored in MobiLink log files. It does not connect to the MobiLink server or change the composition of log files.

The MobiLink Server Log File Viewer allows you to filter the information that you view. In addition, it provides statistics based on the information in the log.

## Running the MobiLink server outside the current session

You can set up the MobiLink server to be available all the time. To make this easier, you can run the MobiLink server for Windows and for Unix so that it remains running when you log off the computer. The way you do this depends on your operating system.

- **Unix daemon** You can run the MobiLink server as a daemon using the `mlsrv11 -ud` option, enabling the MobiLink server to run in the background, and to continue running after you log off.
- **Windows service** You can run the Windows MobiLink server as a service.

To stop a MobiLink server that is running as a service, you can use `mlstop`, `dbsvc`, or the Windows Service Manager.

### See also

- [“-ud option” on page 99](#)
- [“Service utility \(dbsvc\) for Linux” \[SQL Anywhere Server - Database Administration\]](#)
- [“Running the server outside the current session” \[SQL Anywhere Server - Database Administration\]](#)

## Running the Unix MobiLink server as a daemon

To run the MobiLink server in the background on Unix, and to enable it to run independently of the current session, you run it as a daemon.

### To run the Unix MobiLink server as a daemon

- Use the `-ud` option when starting the MobiLink server. For example:

```
mlsrv11 -c "dsn=SQL Anywhere 11 Demo;uid=DBA;pwd=sql" -ud
```

See [“-ud option” on page 99](#).

### See also

- [“Service utility \(dbsvc\) for Linux” \[SQL Anywhere Server - Database Administration\]](#)

## Running the Windows MobiLink server as a service

To run the Windows MobiLink server in the background, and to enable it to run independently of the current session, you run it as a service.

You can run the following service management tasks from the command line, or on the **Services** tab in Sybase Central:

- Add, edit, and remove services.

- Start and stop services.
- Modify the parameters governing a service.

#### See also

- [“Service utility \(dbsvc\) for Windows” \[SQL Anywhere Server - Database Administration\]](#)

## Adding, modifying, and deleting services

The service icons in Sybase Central display the current state of each service using an icon that indicates whether the service is running or stopped.

### To add a new service (Sybase Central)

1. In Sybase Central, in the left pane, click **MobiLink 11**.
2. In the right pane, click the **Services** tab.
3. In the right pane, right-click and choose **New » Service**.
4. Follow the instructions in the **Create Service Wizard**.

You can also use the `dbsvc` utility to create the service. See [“Service utility \(dbsvc\) for Windows” \[SQL Anywhere Server - Database Administration\]](#).

### To delete a service (Sybase Central)

- Choose the service and then click **Edit » Delete**.

### To change the parameters for a service

- Right-click the service and choose **Properties**.

Changes to a service configuration take effect the next time the service is started.

### Setting the startup option

The following options govern startup behavior for MobiLink services. You can set them on the **General** tab of the **Service Properties** window.

- **Automatic** If you choose **Automatic**, the service starts whenever the Windows operating system starts. This setting is appropriate for database servers and other applications running all the time.
- **Manual** If you choose **Manual**, the service starts only when a user with Administrator permissions starts it. For information about Administrator permissions, see your Windows documentation.
- **Disabled** If you choose **Disabled**, the service does not start.

The startup option is applied the next time Windows is started.



## Specifying command line options

The **Configuration** tab of the **Service Properties** window provides a **File name** text box for entering the program executable path and a **Parameters** text box for entering command line options for a service. Do not type the name of the program executable in the **Parameters** box.

For example, to start a MobiLink synchronization service with verbose logging, type the following in the Parameters box:

```
-c "dsn=SQL Anywhere 11 Demo;uid=DBA;pwd=sql"  
-vc
```

The command line options for a service are the same as those for the executable. See [“MobiLink server options” on page 43](#).

## Setting account options

You can choose which account the service runs under. Most services run under the special LocalSystem account, which is the default option for services. You can set the service to log on under another account by opening the **Account** tab on the **Service Properties** window, and entering the account information.

If you choose to run the service under an account other than LocalSystem, that account must have the "log on as a service" privilege. For information about advanced privileges, see your Microsoft Windows documentation.

Whether an icon for the service appears on the taskbar or desktop is dependent on the account you select, and whether **Allow Service To Interact with Desktop** is checked, as follows:

- If a service runs under LocalSystem, and **Allow Service To Interact with Desktop** is checked in the **Service Properties** window, an icon appears on the desktop of every user logged in to Windows XP/200x on the computer running the service. Any user can open the application window and stop the program running as a service.
- If a service runs under LocalSystem, and **Allow Service To Interact with Desktop** is unchecked in the **Service Properties** window, no icon appears on the desktop for any user. Only users with permissions to change the state of services can stop the service.
- If a service runs under another account, no icon appears on the desktop. Only users with permissions to change the state of services can stop the service.

## Changing the executable file

To change the program executable file associated with a service in Sybase Central, click the **Configuration** tab on the **Service Properties** window and type the new path and file name in the **File Name** box.

If you move an executable file to a new directory, you must modify this entry.

## Starting and stopping

### To start or stop a service

1. In Sybase Central, click MobiLink 11 in the left pane, and then open the **Services** tab in the right pane.
2. Right-click the service and choose **Start** or **Stop**.

If you start a service, it keeps running until you stop it. Closing Sybase Central or logging off does not stop the service.

Stopping a service closes all connections to the database and stops the database server. For other applications, the program closes down.

## Running more than one service at a time

Although you can use the Windows Service Manager in the Control Panel for some tasks, you cannot install or configure a MobiLink service from the Windows Service Manager. You can use Sybase Central to perform all the service management for MobiLink.

When you open the Windows Service Manager from the Windows Control Panel, a list of services appears. The names of the SQL Anywhere services are formed from the Service Name you provided when installing the service, prefixed by SQL Anywhere. All the installed services appear together in the list.

This section describes topics specific to running more than one service at a time.

### Service dependencies

In some circumstances you may want to run more than one executable as a service, and these executables may depend on each other. For example, you must run the MobiLink server and the database server to synchronize.

In cases such as these, the services must start in the proper order. If a MobiLink synchronization service starts up before the consolidated database server has started, it fails because it cannot find the consolidated database server. The sequence must be such that the database server is running when you start the MobiLink server. (This does not apply if the consolidated database server is on another computer.)

You can prevent these problems using service groups, which you manage from Sybase Central.

### Service groups

You can assign each service on your system to be a member of a service group. By default, each service belongs to a group. The default group for the MobiLink server is SQLANYMobiLink.

Before you can configure your services to ensure they start in the correct order, you must check that your service is a member of an appropriate group. You can check which group a service belongs to, and change this group, from Sybase Central.

#### To check and change which group a service belongs to

1. In Sybase Central, click MobiLink 11 in the left pane, and then open the **Services** tab in the right pane.
2. Right-click the service and choose **Properties**.
3. Click the **Dependencies** tab. The top text box displays the name of the group the service belongs to.
4. Click **Change** to display a list of available groups on your system.
5. Select one of the groups, or type a name for a new group.
6. Click **OK** to assign the service to that group.

### **Managing service dependencies**

With Sybase Central, you can specify dependencies for a service. For example:

- You can ensure that at least one group has started before the current service.
- You can ensure that any service starts before the current service.

#### **To add a service or group to a list of dependencies**

1. In Sybase Central, click MobiLink 11 in the left pane, and then open the **Services** tab in the right pane.
2. Right-click the service and choose **Properties**.
3. Click the **Dependencies** tab.
4. Click **Add Services** or **Add Service Groups** to add a service or group to the list of dependencies.
5. Select one of the services or groups from the list.
6. Click **OK** to add the service or group to the list of dependencies.

## Running the MobiLink server in a server farm

### Separately licensed component required

Shared state mode is a feature of the MobiLink high availability option, which requires a separate license. See [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

A MobiLink server farm is an environment where there is more than one MobiLink server synchronizing the same set of remote databases with the same consolidated database. This is often required for large scale deployments or for fail-over capability. These MobiLink server farm deployments require MobiLink to run in shared state mode and may require the use of the relay server if an HTTP communication link is used. For TCP based streams, a TCP load balancer should work. When using multiple servers, restartable download does not work.

MobiLink does not run with shared server state by default.

### To enable shared server state

1. Give each MobiLink server a unique name using the `-zs` command line option. These names are used to manage the state of the farm in the consolidated database. See [“-zs option” on page 119](#).
2. Use the `-ss` option to start MobiLink in shared server state mode. If this option is set, MobiLink server prints the following message to the log at startup: `This server is using shared server state for resource locking`. See [“-ss option” on page 95](#).
3. If you are using the notifier with Server Initiated Sync, use the `-lsc` option to specify the local server connect settings. These settings are passed to the other servers in the farm so that they can connect to each other to share the handling of notifications. For example, if running on host **farm\_host22**:  

```
m1srv11 -x tcpip(port=3245) -zs -ss server5 -lsc  
tcpip(host=farm_host22;port=3245) -c ...
```

If you are not using the notifier you don't need the `-lsc` option.

### See also

- [“-lsc option” on page 72](#)
- [“-zs option” on page 119](#)
- [“-ss option” on page 95](#)

## Troubleshooting MobiLink server startup

This section describes some common problems when starting the MobiLink server.

### Ensure that network communication software is running

Appropriate network communication software must be installed and running before you run the MobiLink server. If you are running reliable network software with just one network installed, this should be straightforward. You should confirm that other software requiring network communications is working properly before running the MobiLink server.

If you are running under the TCP/IP protocol, you may want to confirm that ping and telnet are working properly. The ping and telnet applications are provided with many TCP/IP protocol stacks.

### Debug network communications startup problems

If you are having problems establishing a connection across a network, you can use debugging options at both the client and server to diagnose problems. The startup information appears on the server window: you can use the `-o` option to log the results to an output file.

See [“Logging MobiLink server actions” on page 33](#).

---

---

# MobiLink server options

## Contents

mlsrv11 syntax .....	45
@data option .....	50
-a option .....	51
-b option .....	52
-bn option .....	53
-c option .....	54
-cm option .....	55
-cn option .....	56
-cr option .....	57
-cs option .....	58
-ct option .....	59
-dl option .....	60
-dr option .....	61
-ds option .....	62
-dsd option .....	63
-dt option .....	64
-e option .....	65
-esu option .....	66
-et option .....	67
-f option .....	68
-fips option .....	69
-fr option .....	70
-ftr option .....	71
-lsc option .....	72
-m option .....	73
-nba option .....	74
-nc option .....	75
-notifier option .....	76
-o option .....	77
-on option .....	78
-oq option .....	79

-os option .....	80
-ot option .....	81
-ppv option .....	82
-q option .....	86
-r option .....	87
-rd option .....	88
-s option .....	89
-sl dnet option .....	90
-sl java option .....	92
-sm option .....	94
-ss option .....	95
-tc option .....	96
-tf option .....	97
-tx option .....	98
-ud option .....	99
-ui option .....	100
-ux option .....	101
-v option .....	102
-w option .....	105
-wu option .....	106
-x option .....	107
-xo option .....	113
-zp option .....	118
-zs option .....	119
-zt option .....	120
-zu option .....	121
-zus option .....	122
-zw option .....	123
-zwd option .....	124
-zwe option .....	125

---



# mlsrv11 syntax

## Function

Start a MobiLink server.

## Syntax

**mlsrv11 -c "connection-string" [ options ]**

Option	Description
<i>@data</i>	Read in options from the specified environment variable or configuration file. See “ <a href="#">@data option</a> ” on page 50.
<b>-a</b>	Disable automatic reconnection upon synchronization error. See “ <a href="#">-a option</a> ” on page 51.
<b>-b</b>	Trim blank padding of strings. See “ <a href="#">-b option</a> ” on page 52.
<b>-bn</b> <i>size</i>	Specify the maximum number of bytes to consider when comparing BLOBs for conflict detection. See “ <a href="#">-bn option</a> ” on page 53.
<b>-c</b> “ <i>keyword=value; ...</i> ”	Supply ODBC database connection parameters for your consolidated database. See “ <a href="#">-c option</a> ” on page 54.
<b>-cm</b> <i>size</i>	Specify the server memory cache size. See “ <a href="#">-cm option</a> ” on page 55.
<b>-cn</b> <i>connections</i>	Set the maximum number of simultaneous connections with the consolidated database server. See “ <a href="#">-cn option</a> ” on page 56.
<b>-cr</b> <i>count</i>	Set the maximum number of database connection retries. See “ <a href="#">-cr option</a> ” on page 57.
<b>-cs</b> “ <i>keyword=value; ...</i> ”	Supply system database connection parameters for your MobiLink System Database (MLSD).
<b>-ct</b> <i>connection-timeout</i>	Set the length of time a connection may be unused before it is timed out. See “ <a href="#">-ct option</a> ” on page 59.
<b>-dl</b>	Display all log messages in the MobiLink server messages window. See “ <a href="#">-dl option</a> ” on page 60.

Option	Description
<b>-dr</b>	For Adaptive Server Enterprise only. Ensures that tables involved in synchronization do not use the DataRow locking scheme. See <a href="#">“-dr option” on page 61</a> .
<b>-ds</b> <i>size</i>	Specify the maximum amount of data that can be stored for use in all restartable downloads. See <a href="#">“-ds option” on page 62</a> .
<b>-dsd</b>	Disable snapshot isolation, which is the default download isolation level for SQL Anywhere and Microsoft SQL Server consolidated databases. See <a href="#">“-dsd option” on page 63</a> .
<b>-dt</b>	Detect transactions only within the current database. See <a href="#">“-dt option” on page 64</a> .
<b>-e</b> <i>filename</i>	Store remote error logs sent into the named file. See <a href="#">“-e option” on page 65</a> .
<b>-esu</b>	Use snapshot isolation for uploads. See <a href="#">“-esu option” on page 66</a> .
<b>-et</b> <i>filename</i>	Store remote error logs sent into the named file, but delete the contents first if the file exists. See <a href="#">“-et option” on page 67</a> .
<b>-f</b>	Assume synchronization scripts do not change. See <a href="#">“-f option” on page 68</a> .
<b>-fips</b>	Forces all secure MobiLink streams to be FIPS-compliant. See <a href="#">“-fips option” on page 69</a> .
<b>-fr</b>	If table data scripts are missing, synchronization does not terminate but just issues a warning. See <a href="#">“-fr option” on page 70</a> .
<b>-ftr</b> <i>path</i>	Creates a location for files that are to be used by the mlfile-transfer utility. See <a href="#">“-ftr option” on page 71</a> .
<b>-lsc</b> <i>protocol</i> [ <i>protocol-options</i> ]	Specifies the local server connect information. See <a href="#">“-lsc option” on page 72</a> .
<b>-m</b> [ <i>filename</i> ]	Enables QAnywhere messaging. See <a href="#">“-m option” on page 73</a> .
<b>-nba</b> { +   - }	Sets the server download acknowledgement mode of operation. See <a href="#">“-nba option” on page 74</a> .

Option	Description
<b>-nc</b> <i>connections</i>	Sets maximum number of concurrent connections. See “ <a href="#">-nc option</a> ” on page 75.
<b>-notifier</b>	Starts a Notifier for server-initiated synchronization. See “ <a href="#">-notifier option</a> ” on page 76.
<b>-o</b> <i>logfile</i>	Log messages to a file. See “ <a href="#">-o option</a> ” on page 77.
<b>-on</b> <i>size</i>	Set maximum size for log file. See “ <a href="#">-on option</a> ” on page 78.
<b>-oq</b>	Prevent the popup window on startup error. See “ <a href="#">-oq option</a> ” on page 79.
<b>-os</b> <i>size</i>	Maximum size of output file. See “ <a href="#">-os option</a> ” on page 80.
<b>-ot</b> <i>logfile</i>	Log messages to a file, but delete its contents first. See “ <a href="#">-ot option</a> ” on page 81.
<b>-q</b>	Minimize the MobiLink server messages window. See “ <a href="#">-q option</a> ” on page 86.
<b>-r</b> <i>retries</i>	Retry deadlocked uploads at most this many times. See “ <a href="#">-r option</a> ” on page 87.
<b>-rd</b> <i>delay</i>	Set maximum delay, in seconds, before retrying a deadlocked transaction. See “ <a href="#">-rd option</a> ” on page 88.
<b>-s</b> <i>count</i>	Specify the maximum number of rows to be fetched or sent at once. See “ <a href="#">-s option</a> ” on page 89.
<b>-sl dnet</b> <i>script-options</i>	Set the .NET CLR options and force loading of the virtual machine on startup. See “ <a href="#">-sl dnet option</a> ” on page 90.
<b>-sl java</b> <i>script-options</i>	Set the Java virtual machine options and force loading of the virtual machine on startup. See “ <a href="#">-sl java option</a> ” on page 92.
<b>-sm</b> <i>number</i>	Set the maximum number of synchronizations that can be actively worked on. See “ <a href="#">-sm option</a> ” on page 94.
<b>-ss</b>	Puts the server into shared server state mode. See “ <a href="#">-ss option</a> ” on page 95.
<b>-tc</b> <i>minutes</i>	Set the count down timer for SQL script execution. See “ <a href="#">-tc option</a> ” on page 96.

Option	Description
<b>-tf</b>	Fail the SQL script execution when the count down timer expires (not for Oracle). See “-tf option” on page 97.
<b>-tx count</b>	For transactional uploads, batches groups of transactions and commits them together. See “-tx option” on page 98.
<b>-ud</b>	On Unix platforms, run as a daemon. See “-ud option” on page 99.
<b>-ui</b>	For Linux with X window, starts the MobiLink server in shell mode if a usable display isn't available. See “-ui option” on page 100.
<b>-ux</b>	Opens the MobiLink server messages window. See “-ux option” on page 101.
<b>-v [ levels ]</b>	Controls the type of messages written to the log file. See “-v option” on page 102.
<b>-w count</b>	Set the number of database worker threads. See “-w option” on page 105.
<b>-wu count</b>	Set the maximum number of database worker threads permitted to process uploads concurrently. See “-wu option” on page 106.
<b>-x protocol[ (network-parameters) ]</b>	Specify the communications protocol. Optionally, specify network parameters in form <i>parameter=value</i> , with multiple parameters separated by semicolons. See “-x option” on page 107.
<b>-xo protocol[ (network-parameters) ]</b>	For version 8 and 9 clients, specify the communications protocol. Optionally, specify network parameters in form <i>parameter=value</i> , with multiple parameters separated by semicolons. See “-xo option” on page 113.
<b>-zp</b>	In the event of a timestamp conflict between the consolidated and remote database, this option allows timestamp values with a precision higher than the lowest-precision to be used for conflict detection purposes. See “-zp option” on page 118.
<b>-zs name</b>	Specify a server name. See “-zs option” on page 119.
<b>-zt number</b>	Specify the maximum number of processors used to run the MobiLink server. See “-zt option” on page 120.

Option	Description
<b>-zu</b> { +   - }	Controls the automatic addition of users when the <code>authenticate_user</code> script is undefined. See “ <a href="#">-zu option</a> ” on page 121.
<b>-zus</b>	Causes MobiLink to invoke upload scripts for tables for which there is no upload. See “ <a href="#">-zus option</a> ” on page 122.
<b>-zw 1,...5</b>	Controls which levels of warning message to display. See “ <a href="#">-zw option</a> ” on page 123.
<b>-zwd</b> <i>code</i>	Disables specific warning codes. See “ <a href="#">-zwd option</a> ” on page 124.
<b>-zwe</b> <i>code</i>	Enables specific warning codes. See “ <a href="#">-zwe option</a> ” on page 125.

### Description

The MobiLink server opens connections, via ODBC, with your consolidated database server. It then accepts connections from client applications and controls the synchronization process.

You must supply connection parameters for the consolidated database using the `-c` option. The command line options may be specified in any order. The `-c` option is shown here as the first item in a command string as a convention only. It can be anywhere in a list of options, but must precede a connection string.

Unless your ODBC data source is configured to automatically start the consolidated database, the database must be running before you start the MobiLink server.

### See also

- “[MobiLink server](#)” on page 29

## @data option

Reads in options from the specified environment variable or configuration file.

### Syntax

```
mlsrv11 -c "connection-string" @data ...
```

### Remarks

Use this option to read in mlsrv11 command line options from the specified environment variable or configuration file. If both exist with the same name that is specified, the environment variable is used.

For more information about configuration files, see [“Using configuration files” \[SQL Anywhere Server - Database Administration\]](#).

If you want to protect passwords or other information in the configuration file, you can use the File Hiding utility to obfuscate the contents of the configuration file.

See [“File Hiding utility \(dbfhide\)” \[SQL Anywhere Server - Database Administration\]](#).

## -a option

Instructs the MobiLink server to not reconnect on synchronization error.

### Syntax

```
mlsrv11 -c "connection-string" -a ...
```

### Remarks

Should an error occur during synchronization, the MobiLink server automatically disconnects from the consolidated database, and then re-establishes the connection. Reconnecting ensures that the following synchronization starts from a known state. When this behavior is not required, you can use this option to disable it. The maintenance of state information depends on programmer requirements and may vary depending on the ways in which the programmer configures MobiLink scripting to work with the DBMS. This applies even if that database is an Oracle, SQL Anywhere database, or other supported product. Some status information may need to be re-initialized depending on the client application.

## -b option

For columns of type VARCHAR, CHAR, LONG VARCHAR, or LONG CHAR, removes trailing blanks from strings during synchronization.

### Syntax

```
milsrv11 -c "connection-string" -b ...
```

### Remarks

**Note**

It is recommended that you use VARCHAR in the consolidated database rather than CHAR, so that this problem does not occur.

This option helps resolve differences between the SQL Anywhere CHAR data type and the CHAR or VARCHAR data type used by the consolidated database. The SQL Anywhere CHAR data type is equivalent to VARCHAR. However, in most consolidated databases that are not SQL Anywhere, the CHAR(n) data type is blank-padded to n characters.

When -b is specified, the MobiLink server removes trailing blanks from strings for columns of type CHAR, VARCHAR, LONG CHAR, or LONG VARCHAR if the column on the remote is a string. It does this before filtering rows that were uploaded in the current synchronization. The trimmed data is then downloaded to the remote databases.

This option can also be used to detect conflict updates. For each upload update row, the MobiLink server fetches the row from the consolidated database for the given primary key, compares the row with the pre-image of the update, and then determines whether the update is a conflict update. When -b is used, MobiLink trims trailing blanks from columns of type CHAR, VARCHAR, LONG CHAR, or LONG VARCHAR before doing the comparison.

### See also

- [“CHAR columns” on page 8](#)
- [“NVARCHAR data type” \[SQL Anywhere Server - SQL Reference\]](#)

### Example

If the -b option is not used, a primary-key value of 'abc' uploaded from a SQL Anywhere or UltraLite remote to a CHAR(10) column in the consolidated database becomes 'abc' followed by seven blank spaces. If the same row is downloaded, then it appears on the remote as 'abc' followed by seven spaces. If the remote database is not blank-padded, then the remote contains two rows: both 'abc' and 'abc' followed by seven spaces. There is now a duplicate row on the remote.

If the -b option is used, a primary-key value of 'abc' uploaded from a SQL Anywhere or UltraLite remote to a CHAR(10) column in the consolidated database becomes 'abc' followed by seven spaces. Seven spaces still pad the value to ten characters, but if the same row is downloaded, then MobiLink server strips the trailing spaces, and the value appears on the remote as 'abc'. The -b option fixes the duplicate row problem.



## -bn option

Sets the maximum number of BLOB bytes to compare during conflict detection.

### Syntax

```
mlsrv11 -c "connection-string" -bn size ...
```

### Remarks

When two BLOBs contain similar or identical values, the operation of comparing them for filtering or conflict detection can be expensive due to the amount of data involved. This option tells the MobiLink server to consider only the first *size* bytes of two BLOBs when making the comparison. The default is to compare the two BLOBs in their entirety.

Under some situations, limiting the maximum amount of data compared can speed synchronization substantially; however, it can also cause errors. For example, if two large BLOBs differ only in the last few bytes, the MobiLink server may consider them identical when in fact they are not.

## -c option

Specifies connection parameters for the consolidated database.

### Syntax

```
mIsrv11 -c "connection-string" ...
```

### Remarks

The connection string must give the MobiLink server enough information to connect to the consolidated database. The connection string is required.

The connection string must specify connection parameters in the form *keyword=value*, separated by semicolons, with no spaces between parameters.

Connection parameters must be included in the ODBC data source specification if not given in the command line. Check your RDBMS and ODBC data source to determine required connection data.

For a complete list of SQL Anywhere connection parameters, see [“Connection parameters” \[SQL Anywhere Server - Database Administration\]](#).

For information about how to hide the password, see [“File Hiding utility \(dbfhide\)” \[SQL Anywhere Server - Database Administration\]](#).

### Example

```
mIsrv11 -c "dsn=odbcname;uid=DBA;pwd=sql"
```

## -cm option

Sets the maximum size for the server memory cache.

### Syntax

```
m1srv11 -c "connection-string" -cm size[ k | m | g ] ...
```

### Remarks

The maximum amount of memory the server uses for holding table data, network buffers, cached download data, and other structures used for synchronization. When the server has more data than can be held in this memory pool, the data is stored on disk.

The *size* is the amount of memory to reserve in bytes. Use **k**, **m**, or **g** to specify units of kilobytes, megabytes, or gigabytes, respectively. The default is **50M**.

## -cn option

Sets the maximum number of simultaneous consolidated database connections.

### Syntax

```
milsrv11 -c "connection-string" -cn value ...
```

### Remarks

Specifies the maximum number of simultaneous connections that the MobiLink server should make to the consolidated database. The minimum and the default value are one greater than the number of database worker threads. A warning is issued if the supplied value is too small.

A MobiLink database connection is only used for synchronizations using one script version. When the MobiLink server is using all the database connections that it is permitted by the -cn option, if a synchronization is pending but no database connection for its script version currently exists, the MobiLink server disconnects a connection and then creates a new database connection for the pending synchronization's script version.

A value larger than the number of database worker threads may speed performance, particularly if connecting to the consolidated database is slow or if multiple script versions are in use. The optimum maximum number of database connections is the number of script versions times the number of database worker threads, plus one. Connections above this optimum value do not necessarily speed synchronization, and needlessly consumes resources in both the MobiLink server and the consolidated database server.

### See also

- [“-w option” on page 105](#)

## -cr option

Sets the maximum number of database connection retries.

### Syntax

```
mlsrv11 -c "connection-string" -cr value ...
```

### Remarks

Set the maximum number of times that the MobiLink server attempts to connect to the database, before quitting, when a connection goes bad. The default value is three connection retries.

## **-cs option**

Specifies connection parameters for your MobiLink System Database (MLSD).

### **Syntax**

```
mlsrv11 -c "connection-string" -cs "connection-string" ...
```

### **Remarks**

MobiLink server system objects, such as system tables, procedures, triggers, and views can be stored in a database other than the consolidated database. The database that stores the MobiLink system objects is called MLSD.

When this command option is specified on the command line, the MobiLink server makes connections to MLSD to fetch user defined scripts and to maintain synchronization status, such as ML user names, remote IDs, progress offsets, last upload and download timestamps, etc. The MobiLink server uses the original -c command line option connections to the consolidated database and uses these connections to upload data from and download data to the client databases. The consolidated database does not need to have any of the MobiLink server system objects and all the user defined scripts including the error reporting and error handling scripts are fetched from the MLSD and executed in the consolidated database.

When this option is used, the MobiLink server uses the Microsoft Distributed Transaction Coordinator (MSDTC).

The consolidated database and MLSD can be any one of the supported MobiLink consolidated databases. However, the corresponding ODBC drivers must support Microsoft Distributed Transactions.

The consolidated database and MLSD must have a transaction log to use MSDTC.

This option can only be used on Windows operating systems.

## -ct option

Sets the length of time, in minutes, that a connection may be unused before it is timed out and disconnected by the MobiLink server.

### Syntax

```
mlsrv11 -c "connection-string" -ct connection-timeout ...
```

### Remarks

MobiLink database connections that go unused for a specified amount of time are freed by the server. The timeout can be set using the -ct option. A default timeout period of 60 minutes is used.

## **-dl option**

Displays all MobiLink server messages on screen.

### **Syntax**

```
mlsrv11 -c "connection-string" -v -dl ...
```

### **Remarks**

Display all MobiLink server messages in the MobiLink server messages window. By default, only a subset of all messages is shown in the window when a MobiLink server message log file is being output (using -o). In circumstances with many messages, this option can degrade performance.

### **See also**

- [“-o option” on page 77](#)
- [“Logging database server messages to a file” \[SQL Anywhere Server - Database Administration\]](#)



## -dr option

For Adaptive Server Enterprise only. Ensures that tables involved in synchronization do not use the DataRow locking scheme.

### Syntax

```
mIsrv11 -c "connection-string" -dr ...
```

### Remarks

This option should only be used if none of the synchronization tables were created using the DataRow locking scheme.

Use of this option reduces duplicate data sent by the MobiLink server.

### See also

- [“MobiLink isolation levels” on page 165](#)

## -ds option

For use with restartable downloads. Specifies the maximum amount of data that the MobiLink server can use to store all restartable downloads.

### Syntax

```
mlsrv11 -c "connection-string" -ds size[ k | m | g ] ...
```

### Remarks

The MobiLink server holds download data that has not been received by the client for use in a restartable download. This option limits the amount of data that the server holds for all the synchronizations combined.

If *size* is too small the server may release download data, making it impossible to restart a download. The server does not release download data until one of the following occurs:

- The user successfully completes the download.
- The user comes back with a new synchronization request without resume enabled.
- The cache is needed for incoming requests. The oldest unsuccessful download is cleared first.

Use **k**, **m**, or **g** to specify units of kilobytes, megabytes, or gigabytes, respectively. The default is **10m**.

### See also

- [“Resuming failed downloads” on page 158](#)
- [“-dc option” \[MobiLink - Client Administration\]](#)

## -dsd option

Disables snapshot isolation.

### Syntax

```
mlsrv11 -c "connection-string" -dsd ...
```

### Remarks

When the consolidated database is SQL Anywhere (version 10 or later) or Microsoft SQL Server (2005 or later), the default isolation level for downloads is snapshot isolation. If the consolidated database is an earlier version of these databases, the default download isolation level is read committed.

You can also change the default isolation level in a script. However, for SQL Anywhere version 10 and Microsoft SQL Server 2005 and later databases, the isolation level is set at the start of the upload and download transactions. This means that if you set the isolation level in the `begin_connection` script, it may be overridden in the `begin_upload` and `begin_download` scripts.

This option only applies to SQL Anywhere version 10 and Microsoft SQL Server 2005 consolidated databases.

### See also

- [“MobiLink isolation levels” on page 165](#)
- [“-dt option” on page 64](#)
- [“-esu option” on page 66](#)

## -dt option

For Microsoft SQL Server and Adaptive Server Enterprise databases only. Causes MobiLink to detect transactions only within the current database.

### Syntax

```
milsrv11 -c "connection-string" -dt ...
```

### Remarks

This option makes MobiLink ignore all transactions except ones within the current database. It increases throughput and reduces duplication of rows that are downloaded.

Use this option if:

- Your consolidated database is running on Microsoft SQL Server or Adaptive Server Enterprise that is also running other databases.
- You are using snapshot isolation for uploads or downloads with Microsoft SQL Server.
- You are using the DataRow locking scheme for synchronizing tables with Adaptive Server Enterprise.
- Your upload or download scripts do not access any other databases on the server.

This option only applies to Microsoft SQL Server databases using snapshot isolation, and Adaptive Server Enterprise databases using the DataRow locking scheme for tables involved in synchronization.

### See also

- [“MobiLink isolation levels” on page 165](#)
- [“-dsd option” on page 63](#)
- [“-esu option” on page 66](#)

## -e option

Stores error logs sent from SQL Anywhere MobiLink clients.

### Syntax

```
mlsrv11 -c "connection-string" -e filename ...
```

### Remarks

With no `-e` option, error logs from SQL Anywhere MobiLink clients are stored in a file named `mlsrv11.mle`. The `-e` option instructs the MobiLink server to store the error logs in the named file. By default, `dbmlsync` sends, on the occurrence of an error on the remote site, up to 32 kilobytes of remote log messages to a MobiLink server.

This option provides centralized access to remote error logs to help diagnose synchronization issues.

The amount of information delivered from a remote site can be controlled by the `dbmlsync` extended option `ErrorLogSendLimit`.

### See also

- [“-et option” on page 67](#)
- [“ErrorLogSendLimit \(el\) extended option” \[MobiLink - Client Administration\]](#)

## -esu option

Use snapshot isolation for uploads.

### Syntax

```
milsrv11 -c "connection-string" -esu ...
```

### Remarks

By default, MobiLink uses the `SQL_TXN_READ_COMMITTED` isolation level for uploads. In most cases, this is the optimal isolation level.

If you use snapshot isolation for uploads, you may generate conflicts on snapshot transactions during upload updates. If this happens, the MobiLink server rolls back the entire upload and retries it. In this case, you might want to adjust your settings for the MobiLink server options `-r` or `-rd` to specify the delay time between retries and the maximum number of retries.

You can change the default isolation level in a script. To change the upload isolation level, you would typically use the `begin_upload` script.

This option only applies to SQL Anywhere version 10 and Microsoft SQL Server 2005 consolidated databases.

### See also

- [“MobiLink isolation levels” on page 165](#)
- [“-dsd option” on page 63](#)
- [“-dt option” on page 64](#)
- [“-r option” on page 87](#)
- [“-rd option” on page 88](#)

## -et option

Stores error logs sent from SQL Anywhere MobiLink clients in the named file after truncating the existing file.

### Syntax

```
mlsrv11 -c "connection-string" -et filename ...
```

### Remarks

The -et option is the same as the -e option, except that the error log file is truncated before any new errors are added to it.

The amount of information delivered from a remote site can be controlled by the dbmlsync extended option `ErrorLogSendLimit`.

### See also

- [“ErrorLogSendLimit \(el\) extended option” \[MobiLink - Client Administration\]](#)
- [“-e option” on page 65](#)

## **-f option**

Loads synchronization scripts only once, for better performance.

### **Syntax**

```
mIsrv11 -c "connection-string" -f ...
```

### **Remarks**

Without the `-f` option, the MobiLink server checks to see if synchronization scripts have changed during regular operation. This checking is helpful during development, but can have an unnecessary performance impact in a production environment. With the `-f` option, the MobiLink server loads the synchronization scripts only once per MobiLink session.



## -fips option

Forces all secure MobiLink streams to be FIPS-compliant.

### Syntax

```
milsrv11 -c connection-string -fips ...
```

### Remarks

Specifying this option forces all MobiLink encryption to use FIPS-approved algorithms. You can still use unencrypted connections when the -fips option is specified, but you can't use simple encryption.

When you use this option, FIPS-approved algorithms are used for connections regardless of whether you specify them or not. For example, if you start the MobiLink server with the option -fips and the option -x tls(...;fips=no;...), the fips=no setting is ignored and the server starts with fips=yes.

#### **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” \[SQL Anywhere 11 - Introduction\]](#).

For MobiLink transport-layer security, the -fips option causes the server to use the FIPS-approved RSA encryption cipher, even if RSA without FIPS is specified. If ECC is specified, an error occurs because a FIPS-approved elliptic-curve algorithm is not available.

### See also

- [“Encrypting MobiLink client/server communications” \[SQL Anywhere Server - Database Administration\]](#)
- [“FIPS-approved encryption technology” \[SQL Anywhere Server - Database Administration\]](#)

## -fr option

If required table data scripts are missing, synchronization does not abort, but just issues a warning.

### Syntax

```
milsrv11 -c "connection-string" -fr ...
```

### Remarks

By default, the MobiLink server aborts if required synchronization scripts are missing. This option causes MobiLink to issue a warning instead of aborting.

### See also

- [“Required scripts” on page 326](#)
- [“Upload-only and download-only synchronizations” on page 138](#)
- [“Synchronization events” on page 341](#)
- [“Direct row handling” on page 649](#)

## -ftr option

Creates a location for files that are to be used by the mlfiletransfer utility.

### Syntax

```
mlsrv11 -c "connection-string" -ftr path ...
```

### Remarks

This option sets the file transfer root directory. Files that are to be transferred to a user can be placed in the root directory or in a subdirectory with the user name. MobiLink first looks for the requested file in a subdirectory of the file transfer root directory with the user name of the connected client. If the file is not in this subdirectory, MobiLink looks in the file transfer root directory.

This option is required if you want to use the mlfiletransfer utility.

### See also

- [“MobiLink file transfer utility \(mlfiletransfer\)” \[MobiLink - Client Administration\]](#)
- [“authenticate\\_file\\_transfer connection event” on page 353](#)

## -lsc option

Specifies the local server connect information. This information is passed to other servers in the server farm.

### Syntax

```
mlsrv11 -c "connection-string" -lsc protocol[ protocol-options ] ...
```

*protocol* : **tcpip** | **tls** | **http** | **https**

*protocol-options* : ( *option=value*; ... )

### Remarks

This option is only needed when running the notifier in a MobiLink server farm. This information is passed to other servers when they want to connect to the local MobiLink server.

For example if we have a server running on a host named `server_rack10`, the command line could start:

```
mlsrv11 -x tcpip(port=200) -zs -ss server5 -lsc  
tcpip(host=server_rack10;port=200) -c ...
```

In this example another server would use shared state in the consolidated database to get the connect string `tcpip(host=server_rack10;port=200)` and use it to connect to the server just started.

### See also

- [“Running the MobiLink server in a server farm” on page 40](#)
- [“-zs option” on page 119](#)
- [“-ss option” on page 95](#)
- [“Notifiers in a MobiLink server farm” \[\*MobiLink - Server-Initiated Synchronization\*\]](#)

## -m option

Enables QAnywhere messaging.

### Syntax

```
m1srv11 -c "connection-string" -m [ message-properties-file ] ...
```

### Remarks

The optional *message-properties-file* is deprecated. Properties are now specified via Sybase Central and stored in the database, or are specified with server management requests.

In the *message-properties-file*, each property must appear on its own line and consist of a property name, the = character, and then a property value.

For a list of properties you can set, see “[Server properties](#)” [*QAnywhere*].

### See also

- “[Introducing QAnywhere technology](#)” [*QAnywhere*]
- “[Starting QAnywhere with MobiLink enabled](#)” [*QAnywhere*]
- “[Server management requests](#)” [*QAnywhere*]

### Example

To start QAnywhere messaging when you are using the sample server message store (*samples-dir* \QAnywhere\server\qanyserv.db), run the following command:

```
m1srv11 -m -c "dsn=QAnywhere 10 Demo"
```

For information about *samples-dir*, see “[Samples directory](#)” [*SQL Anywhere Server - Database Administration*].

## -nba option

Sets the server download acknowledgement mode of operation.

### Syntax

```
mlsrv11 -c "connection-string" -nba{ + | - } ...
```

### Remarks

When you specify download acknowledgement, you can choose one of two modes: non-blocking (the default, set by -nba+) or blocking (set by -nba-). Blocking download acknowledgement (-nba-) has been deprecated. Use non-blocking whenever possible.

Non-blocking download acknowledgement is recommended because it provides a significant performance advantage over blocking download acknowledgement. However, non-blocking download acknowledgement cannot be used in the following case:

- Clients prior to 10.0.0 do not support non-blocking acknowledgement.

When you enable QAnywhere messaging with the mlsrv11 -m option, you can not use blocking download acknowledgement. You cannot specify both -m and -nba-.

### See also

- dbmlsync: [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Download Acknowledgement synchronization parameter” \[UltraLite - Database Management and Reference\]](#)
- [“nonblocking\\_download\\_ack connection event” on page 471](#)
- [“publication\\_nonblocking\\_download\\_ack connection event” on page 475](#)

## -nc option

Sets the maximum number of concurrent network connections.

### Syntax

```
milsrv11 -c "connection-string" -nc connections ...
```

### Remarks

The MobiLink server rejects new synchronization connections when the limit is reached. On the client, a communication error is issued with a system error code that indicates the connection was refused.

The default is 1024.

To limit the number of concurrent synchronizations for non-persistent HTTP/HTTPS, set -nc significantly higher than -sm. When the -sm limit is reached, the MobiLink server provides an HTTP error 503 (Service Unavailable) to the remote client. If the -nc limit is reached, however, a socket error is issued. The greater the difference between -nc and -sm, the more likely it is that the rejected connections will generate the HTTP 503 error instead of the less descriptive socket error. For example, set -sm to 100 and set -nc to 1000. See [“-sm option” on page 94](#).

## -notifier option

Starts the Notifier for server-initiated synchronization.

### Syntax

```
mlsrv11 -c "connection-string" -notifier [ notifier-properties-file ] ...
```

### Remarks

If you specify a Notifier configuration file name, or if you do not specify a file name but you have a default Notifier properties file called *config.notifier*, the Notifier is configured using that file. This overrides any configuration information that is stored in the *ml\_properties* table in the consolidated database.

Otherwise, MobiLink uses the configuration information that is stored in the *ml\_properties* table in the consolidated database.

When you use the `-notifier` option, you start every Notifier that you have enabled.

For more information about enabling Notifiers, see “Notifier properties” [[MobiLink - Server-Initiated Synchronization](#)].

### See also

- “MobiLink server settings for server-initiated synchronization” [[MobiLink - Server-Initiated Synchronization](#)]
- “Configuring server-side settings using the Notifier configuration file” [[MobiLink - Server-Initiated Synchronization](#)]
- “Notifiers” [[MobiLink - Server-Initiated Synchronization](#)]
- “Notifiers in a MobiLink server farm” [[MobiLink - Server-Initiated Synchronization](#)]



## -o option

Logs output messages to a MobiLink server message log file, and limits the data logged to the MobiLink server messages window.

### Syntax

```
milsrv11 -c "connection-string" -o logfile ...
```

### Remarks

Write all log messages to the specified file. Note that the MobiLink server window, if present, usually shows a subset of all messages logged.

The MobiLink server gives the full error context in its output file if errors occur during synchronization. The error context may include the following information:

- **User Name** This is the actual user name that is provided by MobiLink SQL Anywhere applications during synchronization.
- **Modified User Name** This is the user name as modified by the `modify_user` script.
- **Transaction** This lists the transaction the error occurs in. The transaction could be `authenticate_user`, `begin_synchronization`, `upload`, `prepare_for_download`, `download`, or `end_synchronization`.
- **Table Name** This shows the table name if it is available or null.
- **Row Operation** The operation could be `INSERT`, `UPDATE`, `DELETE` or `FETCH`.
- **Row Data** This shows all the column values of the row that caused the error.
- **Script Version** This is the script version currently used for synchronization.
- **Script** This is the script that caused the error.

Error context information appears in the log regardless of your chosen level of verbosity.

### See also

- [“-os option” on page 80](#)
- [“-dl option” on page 60](#)
- [“-ot option” on page 81](#)
- [“-on option” on page 78](#)
- [“-v option” on page 102](#)

## -on option

Specifies a maximum size for the MobiLink server message log file, after which the file is renamed with the extension `.old` and a new file is started.

### Syntax

```
mlsrv11 -c "connection-string" -on size [ k | m ]...
```

### Remarks

The *size* is the maximum file size for the message log, in bytes. Use the suffix `k` or `m` to specify units of kilobytes or megabytes, respectively. The minimum size limit is 10 KB.

When the log file reaches the specified size, the MobiLink server renames the output file with the extension `.old`, and starts a new one with the original name.

#### Note

If the `.old` file already exists, it is overwritten. To avoid losing old log files, use the `-os` option instead.

This option cannot be used with the `-os` option.

### See also

- [“-o option” on page 77](#)
- [“-ot option” on page 81](#)
- [“-on option” on page 78](#)
- [“-os option” on page 80](#)
- [“-v option” on page 102](#)

## -oq option

On Windows, prevents the appearance of the error window when a startup error occurs.

### Syntax

```
mlsrv11 -c "connection-string" -oq ...
```

### Remarks

By default, the MobiLink server displays a window if a startup error occurs. The -oq option prevents this window from being displayed.

## -os option

Sets the maximum size of the MobiLink server message log file, after which a new log file with a new name is created and used.

### Syntax

```
mlsrv11 -c "connection-string" -os size [ k | m ] ...
```

### Remarks

The *size* is the maximum file size for logging output messages. The default unit is bytes. Use the suffix *k* or *m* to specify units of kilobytes or megabytes, respectively. The minimum size limit is 10 KB.

Before the MobiLink server logs output messages to a file, it checks the current file size. If the log message makes the file size exceed the specified size, the MobiLink server renames the message log file to *yymmddxx.mls*, where *xx* is a number from 00 to 99, and *yymmdd* represents the current year, month, and day.

You can use this option to prune old message log files to free up disk space. The latest output is always appended to the file specified by *-o* or *-ot*.

You cannot use this option with the *-on* option.

#### Note

This option makes an unlimited number of log files. To avoid this situation, use *-o* or *-on*.

### See also

- [“-o option” on page 77](#)
- [“-on option” on page 78](#)
- [“-ot option” on page 81](#)
- [“-v option” on page 102](#)

## -ot option

Logs output messages to the MobiLink server message log file, but deletes the contents first.

### Syntax

```
mlsrv11 -c "connection-string" -ot logfilename ...
```

### Remarks

The default is to send output to the screen.

### See also

- [“-on option” on page 78](#)
- [“-os option” on page 80](#)
- [“-v option” on page 102](#)
- [“-o option” on page 77](#)

## -ppv option

Causes MobiLink to print new periodic monitoring values according to the period specified. Periods are in seconds.

### Syntax

```
mlsrv11 -c "connection-string" -ppv period ...
```

### Remarks

These values can provide insight into the state of the server, and are useful for determining the health and performance of the MobiLink server. For example, one could look at the `DB_CONNECTIONS` and `LONGEST_DB_WAIT` values to look for potential problems with the `-w` option or in the synchronization scripts. The values also provide an easy way to track system wide throughput measures, such as the number of rows uploaded or downloaded per second and the number of successful synchronizations per second.

The suggested period is 60s.

If the period is set too small, the log will grow very quickly.

Each row of output is prefixed with **PERIODIC:** to aid in searching for and filtering out the values.

The printed values can include the following information:

- **CMD\_PROCESSOR\_STAGE\_LEN** The length of the queue for synchronization work.
- **CPU\_USAGE** The amount of CPU time used by the MobiLink server in microseconds.
- **DB\_CONNECTIONS** The number of database connections in use.
- **FREE\_DISK\_SPACE** The disk space available on the temp disk in bytes.
- **HEARTBEAT\_STAGE\_LEN** The length of the queue for periodic, non-sync work.
- **LONGEST\_DB\_WAIT** The longest length of time an active synchronization has been waiting for the database.
- **LONGEST\_SYNC** The age of the oldest synchronization in microseconds.
- **MEMORY\_USED** The bytes of RAM in use (for Windows only).
- **ML\_NUM\_CONNECTED\_CLIENTS** The number of connected synchronization clients.
- **NUM\_COMMITS** The total number of commits.
- **NUM\_CONNECTED\_FILE\_XFERS** The number of mlfiletransfers currently connected.
- **NUM\_CONNECTED\_LISTENERS** The number of listeners currently connected.
- **NUM\_CONNECTED\_MONITORS** The number of monitors currently connected.
- **NUM\_CONNECTED\_PINGS** The number of pinging clients currently connected.
- **NUM\_CONNECTED\_SYNCNS** The number of data synchronizations currently connected.
- **NUM\_ERRORS** The total number of errors.
- **NUM\_FAILED\_SYNCNS** The total number of failed syncs.

- **NUM\_IN\_APPLY\_UPLOAD** The number of synchronizations currently in the apply upload phase.
- **NUM\_IN\_AUTH\_USER** The number of synchronizations currently in the authenticate user phase.
- **NUM\_IN\_BEGIN\_SYNC** The number of synchronizations currently in the begin synchronization phase.
- **NUM\_IN\_CONNECT** The number of synchronizations currently in the connect phase.
- **NUM\_IN\_CONNECT\_FOR\_ACK** The number of synchronizations currently in the connect for download ack phase.
- **NUM\_IN\_END\_SYNC** The number of synchronizations currently in the end synchronization phase.
- **NUM\_IN\_FETCH\_DNLD** The number of synchronizations currently in the fetch download phase.
- **NUM\_IN\_GET\_DB\_WORKER\_FOR\_ACK** The number of synchronizations currently in the get DB worker for ack phase.
- **NUM\_IN\_NON\_BLOCKING\_ACK** The number of synchronizations currently in the non-blocking download ack phase.
- **NUM\_IN\_PREP\_FOR\_DNLD** The number of synchronizations currently in the prepare for download phase.
- **NUM\_IN\_RECVING\_UPLOAD** The number of synchronizations currently in the receive upload phase.
- **NUM\_IN\_SEND\_DNLD** The number of synchronizations currently in the send download phase.
- **NUM\_IN\_SYNC\_REQUEST** The number of synchronizations currently in the synchronization request phase.
- **NUM\_IN\_WAIT\_FOR\_DNLD\_ACK** The number of synchronizations currently in the wait for download ack phase.
- **NUM\_ROLLBACKS** The total number of rollbacks.
- **NUM\_ROWS\_DOWNLOADED** The total number of rows sent to remotes.
- **NUM\_ROWS\_UPLOADED** The total number of rows received from remotes.
- **NUM\_SUCCESS\_SYNCES** The total number of successful syncs.
- **NUM\_UNSUBMITTED\_ERROR\_RPTS** The number of unsubmitted error reports.
- **NUM\_UPLOAD\_CONNS\_IN\_USE** The number of upload connections currently in use.
- **NUM\_WAITING\_CONS** The number of synchronizations currently waiting for the consolidated database.
- **NUM\_WARNINGS** The total number of warnings.
- **PAGES\_IN\_STREAMSTACK** The number of pages held by the network streams.
- **PAGES\_LOCKED** The number of cache pages loaded into memory.
- **PAGES\_LOCKED\_MAX** The number of pages in the memory cache. This is set with -cm. See [“-cm option” on page 55](#).

- **PAGES\_SWAPPED\_IN** The total number of pages ever read from disk.
- **PAGES\_SWAPPED\_OUT** The total number of pages ever swapped to disk.
- **PAGES\_USED** The number of cache pages used. This includes pages swapped to disk so it may be larger than the cache size.
- **RAW\_TCP\_STAGE\_LEN** The length of the network work queue.
- **SERVER\_IS\_PRIMARY** Indicates if the server is primary or secondary. Shows 1 if the server is primary or 0 otherwise.
- **STREAM\_STAGE\_LEN** The length of the high level network processing queue.
- **TCP\_BYTES\_READ** The total number of bytes ever read.
- **TCP\_BYTES\_WRITTEN** The total number of bytes ever written.
- **TCP\_CONNECTIONS** The number of TCP connections currently opened.
- **TCP\_CONNECTIONS\_CLOSED** The total number of connections ever closed.
- **TCP\_CONNECTIONS\_OPENED** The total number of connections ever opened.
- **TCP\_CONNECTIONS\_REJECTED** The total number of connections ever rejected.

### Example

Below is sample output showing the periodic monitoring values.

```

I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_CONNECTIONS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_USED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_LOCKED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_LOCKED_MAX: 12692
I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_CONNECTIONS_OPENED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_CONNECTIONS_CLOSED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_CONNECTIONS_REJECTED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_BYTES_READ: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: TCP_BYTES_WRITTEN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: ML_NUM_CONNECTED_CLIENTS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_SWAPPED_OUT: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_SWAPPED_IN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: PAGES_IN_STREAMSTACK: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: CPU_USAGE: 468750
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_COMMITS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_ROLLBACKS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_SUCCESS_SYNCNS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_FAILED_SYNCNS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_ERRORS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_WARNINGS: 1
I. 2008-10-14 10:34:43. <Main> PERIODIC: DB_CONNECTIONS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: RAW_TCP_STAGE_LEN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: STREAM_STAGE_LEN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: HEARTBEAT_STAGE_LEN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: CMD_PROCESSOR_STAGE_LEN: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_ROWS_DOWNLOADED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_ROWS_UPLOADED: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: FREE_DISK_SPACE: 162552295424
I. 2008-10-14 10:34:43. <Main> PERIODIC: LONGEST_DB_WAIT: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: LONGEST_SYNC: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_UNSUBMITTED_ERROR_RPTS: 247
I. 2008-10-14 10:34:43. <Main> PERIODIC: MEMORY_USED: 94375936
    
```



```
I. 2008-10-14 10:34:43. <Main> PERIODIC: SERVER_IS_PRIMARY: 1
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_CONNECTED_SYNCS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_CONNECTED_PINGS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_CONNECTED_FILE_XFERS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_CONNECTED_MONITORS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_CONNECTED_LISTENERS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_WAITING_CONS: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_SYNC_REQUEST: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_RECVING_UPLOAD: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_CONNECT: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_AUTH_USER: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_BEGIN_SYNC: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_APPLY_UPLOAD: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_PREP_FOR_DNLD: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_FETCH_DNLD: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_END_SYNC: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_SEND_DNLD: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_WAIT_FOR_DNLD_ACK: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_GET_DB_WORKER_FOR_ACK: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_CONNECT_FOR_ACK: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_IN_NON_BLOCKING_ACK: 0
I. 2008-10-14 10:34:43. <Main> PERIODIC: NUM_UPLOAD_CONNS_IN_USE: 0
```

## **-q option**

Instructs MobiLink to run in a minimized window on startup.

### **Syntax**

```
mlsrv11 -c "connection-string" -q ...
```

### **Remarks**

Minimize the MobiLink server window.

## -r option

Sets the maximum number of deadlock retries.

### Syntax

```
mlsrv11 -c "connection-string" -r retries ...
```

### Remarks

By default, MobiLink server retries uploads that are deadlocked, for a maximum of 10 attempts. If the deadlock is not broken, synchronization fails, since there is no guarantee that the deadlock can be overcome. This option allows an arbitrary retry limit to be set. To stop the server from retrying deadlocked transactions, specify **-r 0**. The upper bound on this setting is 2 to the power 32, minus one.

## **-rd option**

Sets the maximum delay time between deadlock retries.

### **Syntax**

```
mlsrv11 -c "connection-string" -rd delay ...
```

### **Remarks**

When upload transactions are deadlocked, the MobiLink server waits a random length of time before retrying the transaction. The random nature of the delay increases the likelihood that future attempts succeed. This option allows you to specify the maximum delay in units of seconds. The value 0 (zero) makes retries instantaneous, but larger values are recommended because they yield more successful retries. The default and maximum delay value is **30**.

## -s option

Sets the maximum number of rows that can be uploaded at the same time.

### Syntax

```
mIsrv11 -c "connection-string" -s count ...
```

### Remarks

Set the maximum number of rows that can be inserted, updated, or deleted at the same time to *count*.

The MobiLink server sends upload rows to the consolidated database through the ODBC driver. This option controls the number of rows sent to the database server in each batch. Increasing this value can speed up processing of the upload stream and reduce network time. However, with a higher setting the MobiLink server may require more resources for applying the upload stream.

The number of rows uploaded at once can be viewed in the log file as **rowset size**.

The default is 10.

## -sl dnet option

Sets the .NET Common Language Runtime (CLR) options and forces the CLR to load on startup.

### Syntax

```
mksrv11 -c "connection-string" -sl dnet options ...
```

### Remarks

Sets options to pass directly to the .NET CLR. The options are:

Option	Description
<b>-Dname=value</b>	Set an environment variable. For example,  <code>-Dsynchtype=far -Dextra_rows=yes</code>  For more information, see the .NET framework class System.Environment.
<b>-MLAutoLoadPath=path</b>	Set the location of base assemblies. Only works with private assemblies. To tell MobiLink where assemblies are located, use this option or -MLDomConfigFile, but not both. When you use -MLAutoLoadPath, you cannot specify a domain in the event script. The default is the current directory.
<b>-MLDomConfigFile=file</b>	Set the location of base assemblies. Use when you have shared assemblies, or you don't want to load all assemblies in the directory, or you can't use MLAutoLoadPath for some other reason. To tell MobiLink where assemblies are located, use -MLDomConfigFile or -MLAutoLoadPath, but not both.
<b>-MLStartClasses=classnames</b>	At server startup, load and instantiate user-defined start classes in the order listed.
<b>-clrConGC</b>	Enable concurrent garbage collection in the CLR.
<b>-clrFlavor=( wks   svr )</b>	Flavor of the .NET CLR to load. The flavor is <b>svr</b> for server and <b>wks</b> for workstation. By default, <b>wks</b> is loaded.
<b>-clrVersion=version</b>	Version of the .NET CLR to load. This must be prefixed with <b>v</b> . For example, <b>v1.0.3705</b> loads the directory <code>\Microsoft.NET\Framework\v1.0.3705</code> .

To display this list of options, run the following command:

```
mksrv11 -sl dnet (?)
```

**See also**

- [“Writing synchronization scripts in .NET” on page 589](#)

## -sl java option

Sets the Java virtual machine options and forces the virtual machine to load on startup.

### Syntax

`milsrv11 -c "connection-string" -sl java ( options ) ...`

### Remarks

Sets -jrepath and other options to pass directly to the Java virtual machine. The options are:

Option	Description
<b>-hotspot</b>   <b>-server</b>   <b>-classic</b>	Override the default choice for the Java VM to use.
<b>-cp</b> <i>location</i> ;...	Specify a set of directories or JAR files in which to search for classes. Instead of -cp, you can also use -classpath.
<b>-D</b> <i>name=value</i>	Set a system property. For example, <code>-Dsynctype=far -Dextra_rows=yes</code>
<b>-DMLStartClasses</b> = <i>classname</i> , ...	At server startup, load and instantiate user-defined start classes in the order listed.
<b>-jrepath</b> <i>path</i>	Override the default JRE path, which is the directory <code>install-dir\Sun\jre160_chip</code> (where chip can be any supported chip, such as x86).
<b>-verbose</b> [ <b>:class</b>   <b>:gc</b>   <b>:jni</b> ]	Enable verbose output.
<b>-X</b> <i>vm-option</i>	Set a VM-specific option as described in the file <code>install-dir\Sun\jre160_chip\bin\client\Xusage.txt</code> (where chip can be any supported chip, such as x86).

To display a list of Java options you can use, type:

```
java
```

### Unix notes

Options must be enclosed in brackets. These can be round brackets, as shown in the syntax above, or curly braces { }.

The -jrepath option is only available on Windows. On Unix, if you want to load a specific JRE, you should set the LD\_LIBRARY\_PATH (LIBPATH on AIX, SHLIB\_PATH on HP-UX) to include the directory containing the JRE. The directory must be listed before any of the SQL Anywhere installation directories.

On Unix, the -cp options must be separated with colons.



### See also

- [“Writing synchronization scripts in Java” on page 527](#)

### Examples

For example, on Windows the following partial `mlsrv11` command line sets the Java virtual machine option that enables system asserts:

```
mlsrv11 -sl java (-cp ;\myclasses; -esa) ...
```

On Windows, the following partial `mlsrv11` command line defines the `LDAP_SERVER` system property:

```
mlsrv11 -sl java ( -cp ;\myclasses; -DLdap_SERVER=huron-ldap ) ...
```

The following partial `mlsrv11` command line works on Unix:

```
mlsrv11 -sl java { -cp .:$CLASSPATH:/opt/myclasses:/opt/my.jar: }
```

## **-sm option**

Sets the maximum number of synchronizations that can be actively worked on, which limits the maximum number of network connections as well.

### **Syntax**

`mlsrv11 -c "connection-string" -sm number ...`

### **Remarks**

The MobiLink server performs the following tasks simultaneously:

1. Read upload data from the network and unpack it.
2. Apply uploads to the consolidated database.
3. Fetch rows to be downloaded from the consolidated database.
4. Pack download data and send it to remote databases.

The number of synchronizations for each task is limited as follows:

- The number of synchronizations doing tasks 2 and 3 is less than or equal to the setting for the `mlsrv11 -w` option.
- The number of synchronizations doing task 2 is less than or equal to the setting for the `mlsrv11 -wu` option.
- The number of synchronizations doing all four tasks is less than or equal to the setting for the `-sm` option.

Higher values for `-sm`, especially when much greater than `-w`, allow the MobiLink server to perform more network tasks (1 and 4) than database tasks (2 and 3). This can help ensure that a database worker doesn't have to wait for tasks when network performance might otherwise be a bottleneck. This can improve throughput. However, if `-sm` is set too high, the MobiLink server can allocate more memory than is directly available, causing the virtual memory paging of the operating system to be activated, which in turn causes memory to be swapped to disk—significantly decreasing throughput.

### **See also**

- [“-w option” on page 105](#)
- [“-wu option” on page 106](#)

## -ss option

**Separately licensed component required**

The -ss option is a feature of the MobiLink high availability option, which requires a separate license. See [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

Enable the MobiLink server to run in a server farm.

### Syntax

```
mlsrv11 -c "connection string" -ss ...
```

### Remarks

By default, a MobiLink server does not run in a server farm. If you would like to enable the current MobiLink server to run in the server farm, you need to specify this option. Note that if the MobiLink server starts with -ss, it adds itself to the ml\_server table in the consolidated database and also stops the other Mobilink servers (if any) that are connected to the same consolidated database, but were not started with -ss. See [“Running the MobiLink server in a server farm” on page 40](#).

When this option is specified the server must have a name specified using the -zs option. For example, the server command line could be:

```
mlsrv11 -c "connection-string" -ss -zs server5
```

### See also

- [“Running the MobiLink server in a server farm” on page 40](#)
- [“-zs option” on page 119](#)
- [“-lsc option” on page 72](#)

## -tc option

Sets a long running statement threshold for SQL script execution.

### Syntax

```
mlsrv11 -c "connection string" -tc minutes ...
```

### Remarks

By default, the MobiLink server watches the execution time of each SQL script and issues a warning message when the execution time of the script reaches 10 minutes (the default). If the -tf option is used and the consolidated database is running on an Oracle server, the MobiLink server fails the script.

The default value can be reset to zero or a positive integer and its units are in minutes. When it is set to zero, the -tc switch is disabled and the MobiLink server does not watch any script execution.

When the warning time is a non-zero value, the MobiLink server shows the warning message in an exponential way. The warning is shown when the execution time first passes the time specified; the warning is shown again when the execution time passes 2 \* the given time, then 4 \* the given time, and so on.

The warning message contains the connection ID used for the current synchronization and timeout warning context that includes the following, if they are available: Remote ID, ML User Name, Modified User Name, Transaction, Table Name, Row Values and Script Version. The timeout warning context is shown regardless of the verbose settings of the MobiLink server.

When the consolidated database is running on an Oracle database server and the timeout warning message occurs, a database user with DBA authority may need to check the consolidated database to determine the cause of the problem. The SID and SERIAL# of the connection used by the synchronization can be found in the warning message. If the synchronization connection is stopped, the MobiLink server terminates the current synchronization.

### See also

- [“-tf option” on page 97](#)

## -tf option

This option is used to let the MobiLink server fail the SQL script if the execution time passes the time specified by -tc. This option is not available when the consolidated database is running on an Oracle server.

### Syntax

```
mlsrv11 -c "connection string" -tf ...
```

### Remarks

If the SQL script fails, the MobiLink server may skip the row (if the script is an upload script and if the `handle_error` script returns 1000) and continue the synchronization or abort the synchronization.

The MobiLink server shows a warning message if this option is specified and it is running against an Oracle server.

This option is ignored if **-tc 0** is specified.

## **-tx option**

When using transactional uploads, this option batches groups of transactions and commits them together.

### **Syntax**

```
mlsrv11 -c "connection-string" -tx count ...
```

### **Remarks**

Use this option to improve performance when doing transactional uploads.

*count* can be any non-negative value. The default is 1, which means commit every transaction separately. Use a value of zero to perform one commit after all transactions have been uploaded.

### **See also**

- “-tu option” [[MobiLink - Client Administration](#)]

## -ud option

Instructs MobiLink to run as a daemon.

### Syntax

```
mlsrv11 -c "connection-string" -ud ...
```

### Remarks

Unix platforms only.

### See also

- [“Running the MobiLink server outside the current session” on page 35](#)

## **-ui option**

For Linux with X window server support, starts the MobiLink server in shell mode if a usable display isn't available.

### **Syntax**

```
mlsrv11 -c "connection-string" -ui ...
```

### **Remarks**

When this option is used, mlsrv11 tries to start with X Windows. If this fails, it starts in shell mode.

When **-ui** is specified, the server attempts to find a usable display. If it cannot find one, for example because the X window server isn't running, then the MobiLink server starts in shell mode.



## -ux option

For Linux, opens the MobiLink server messages window where messages are displayed.

### Syntax

```
mlsrv11 -c "connection-string -ux ...
```

### Remarks

When -ux is specified, the MobiLink server must be able to find a usable display. If it cannot find one, for example because the DISPLAY environment variable is not set or because the X window server is not running, the MobiLink server fails to start.

To run the MobiLink server messages window in quiet mode, use -q.

On Windows, the MobiLink server messages window appears automatically.

### See also

- [“-q option” on page 86](#)

## -v option

Allows you to specify what information is logged to the message log file and displayed in the synchronization window.

### Syntax

```
milsrv11 -c "connection-string" -v[ levels ] ...
```

### Remarks

This option is particularly useful when dbmlsync transaction-level uploads are used.

This option controls the type of messages written to the message log file.

If you specify -v alone, the MobiLink server writes a minimal amount of information about each synchronization.

A high level of verbosity can adversely affect performance and should only be used during development.

The values of levels are as follows. You can use one or more of these options at once; for example, -vnrsu.

- **+** Turn on all logging options that increase verbosity.
- **c** Show the content of each synchronization script when it is invoked. This level implies s.
- **e** Show system event scripts. These system event scripts are used to maintain MobiLink system tables and the SQL scripts that control the upload.
- **f** Show first-read errors. This logs errors caused when load-balancing devices check for server liveness by making connections that don't send any data, and cause failed synchronizations.

For TCP/IP connections, you might be better off using the TCP/IP option **ignore**. For more information, see [“-x option” on page 107](#).

- **h** Show the remote schema as uploaded during synchronization.
- **i** Display the column values of each row uploaded. Use this option instead of -vr, which displays the column values of each row uploaded and downloaded, to reduce the amount of data being logged. Specifying -vi with -vq is the same as specifying -vr.
- **m** Prints the duration of each synchronization and the duration of each synchronization phase to the log whenever a synchronization completes. The synchronization phases are shown below. They are the same as those displayed in the MobiLink Monitor. All times are shown in milliseconds (ms).
  - **Synchronization request** The time taken between creating the network connection between the remote database and the MobiLink server, up to receiving the first bytes of the upload stream. This time is insignificant unless you have set -sm to a smaller value than -nc, in which case this time can include the time that a synchronization is paused, when the number of synchronizations is larger than the maximum number of active synchronizations that were specified with -sm.
  - **Receive upload** The time taken from the first bytes of the upload stream being received by the MobiLink server until the upload stream from the remote database has been completely received. The upload stream includes table definitions and the remote database rows being uploaded, so the

time may be significant even for a download-only synchronization. The time depends on the size of the upload stream and the network bandwidth for the transfer.

- **Get DB worker** The time required to acquire a free database worker thread.
- **Connect** The time required by the database worker thread to make a database connection if a new database connection is needed. For example, after an error or if the script version has changed.
- **Authenticate user** The time required to authenticate the user.
- **Begin synchronization** The time required for the begin\_synchronization event if it is defined, plus the time to fetch the last\_upload\_time for each subscription (using the shared MobiLink administrative connection).
- **Apply upload** The time required for the uploaded data to be applied to the consolidated database.
- **Prepare for download** The time required for the prepare\_for\_download event.
- **Fetch download** The time required to fetch the rows to be downloaded from the consolidated database to create the download stream.
- **End synchronization** The time required for the end\_synchronization event, after which the database worker thread is released. If you are using blocking download acknowledgement, then this phase occurs after the **Wait for download ack** phase. Otherwise, it occurs before the download stream is sent to the remote database.
- **Send download** The time required to send the download stream to the remote database. The time depends on the size of the download stream and the network bandwidth for the transfer. For an upload-only synchronization, the download stream is simply an upload acknowledgement.
- **Wait for download ack** The time spent waiting for the download to be applied to the remote database and for the remote database to send the download acknowledgement. This phase is only shown if the remote database has enabled download acknowledgement.
- **Get DB worker for download ack** The time spent waiting for a free database worker thread after the download acknowledgement has been received. This phase is only shown if the remote database has enabled download acknowledgement and the MobiLink server is using non-blocking download acknowledgement.
- **Connect for download ack** The time required by the database worker thread to make a database connection if a new database connection is needed. This phase is only shown if the remote database has enabled download acknowledgement and the MobiLink server is using non-blocking download acknowledgement.
- **Non-blocking download ack** The time required for the publication\_nonblocking\_download\_ack connection and nonblocking\_download\_ack connection events. This phase is only shown if the remote database has enabled download acknowledgement and the MobiLink server is using non-blocking download acknowledgement.

Each value is prefixed with "PHASE:" to aid in searching for and printing the values.

The following example is sample output showing the durations for the various synchronization phases:

```
I. 2008-06-05 14:48:36. <1> PHASE: start_time: 2008-06-05 14:48:36.048
I. 2008-06-05 14:48:36. <1> PHASE: duration: 175
I. 2008-06-05 14:48:36. <1> PHASE: sync_request: 0
```

```

I. 2008-06-05 14:48:36. <1> PHASE: receive_upload: 19
I. 2008-06-05 14:48:36. <1> PHASE: get_db_worker: 0
I. 2008-06-05 14:48:36. <1> PHASE: connect: 18
I. 2008-06-05 14:48:36. <1> PHASE: authenticate_user: 51
I. 2008-06-05 14:48:36. <1> PHASE: begin_sync: 69
I. 2008-06-05 14:48:36. <1> PHASE: apply_upload: 0
I. 2008-06-05 14:48:36. <1> PHASE: prepare_for_download: 1
I. 2008-06-05 14:48:36. <1> PHASE: fetch_download: 4
I. 2008-06-05 14:48:36. <1> PHASE: wait_for_download_ack: 0
I. 2008-06-05 14:48:36. <1> PHASE: end_sync: 0
I. 2008-06-05 14:48:36. <1> PHASE: send_download: 10
I. 2008-06-05 14:48:36. <1> PHASE: get_db_worker_for_download_ack: 0
I. 2008-06-05 14:48:36. <1> PHASE: connect_for_download_ack: 0
I. 2008-06-05 14:48:36. <1> PHASE: nonblocking_download_ack: 0

```

- **n** Show row-count summaries.
- **o** Show SQL passthrough activity.
- **p** Show progress offsets.
- **q** Display the column values of each row downloaded. Use this option instead of `-vr`, which displays the column values of each row uploaded and downloaded, to reduce the amount of data being logged. Specifying `-vi` with `-vq` is the same as specifying `-vr`.
- **r** Display the column values of each row uploaded or downloaded. To log only the column values of each row uploaded, use `-vi`. To log only the column values of each row downloaded, use `-vq`.
- **s** Show the name of each synchronization script as it is invoked.
- **t** Show the translated SQL that results from scripts that are written in ODBC canonical format. This level implies `c`. The following example shows the automatic translation of a statement for SQL Anywhere.

```

I. 02/11 11:02:14. [102]: begin_upload synch2
{ call SynchLogLine( ?, ?, 'begin_upload' ) }
I. 02/11 11:02:14. [102]: Translated SQL:
call SynchLogLine( ?, ?, 'begin_upload' )

```

The following example shows the translation of the same statement for Microsoft SQL Server.

```

I. 02/11 11:03:21. [102]: begin_upload synch2
{ call SynchLogLine( ?, ?, 'begin_upload' ) }
I. 02/11 11:03:21. [102]: Translated SQL:
EXEC SynchLogLine ?, ?, 'begin_upload'

```

- **u** Show undefined table scripts. This may help new users understand the synchronization process.

**See also**

- [“MobiLink statistical properties” on page 195](#)

---

## -w option

Sets the number of database worker threads.

### Syntax

```
milsrv11 -c "connection-string" -w count ...
```

### Remarks

Each database worker thread accepts synchronization requests one at a time.

Each database worker thread uses one connection to the consolidated database. The MobiLink server opens one additional connection for administrative purposes. So, the minimum number of connections from the MobiLink server to the consolidated database is *count* + 1.

The number of database worker threads has a strong influence on MobiLink synchronization throughput, and you need to run tests to determine the optimum number for your particular synchronization setup. The number of database worker threads determines how many synchronizations can be active in the consolidated database simultaneously; the rest gets queued waiting for database worker threads to become available. Adding database worker threads should increase throughput, but it also increases the possibility of contention between the active synchronizations. At some point adding more database worker threads decrease throughput because the increased contention outweighs the benefit of overlapping synchronizations.

The value set for this option is also the default setting for the -wu option, which can be used to limit the number of threads that can simultaneously upload to the consolidated database. This is useful if the optimum number of database worker threads for downloading is larger than the optimum number for uploading. The best throughput may be achieved with a large number of database worker threads (via -w) with a small number allowed to apply uploads simultaneously (via -wu). In general, the optimum number for -wu depends on the consolidated database, and is relatively independent of the processing or network speeds for the remote databases. Therefore, when you increase the number of threads with -w, you may want to use -wu to restrict the number that can upload simultaneously. For more information, see [“-wu option” on page 106](#).

The default number of database worker threads is **5**.

### See also

- [“-wu option” on page 106](#)
- [“-sm option” on page 94](#)
- [“-cn option” on page 56](#)

## -wu option

Sets the maximum number of database worker threads that can apply uploads to the consolidated database simultaneously.

### Syntax

```
mlsrv11 -c "connection-string" -wu count ...
```

### Remarks

Use the -wu option to limit the number of database worker threads that can simultaneously apply uploads to the consolidated database. When the limit is reached, a database worker thread that is ready to apply its upload to the consolidated database must wait until another finishes its upload.

The most common cause of contention in the consolidated database is having too many database worker threads applying uploads simultaneously. Downloads cause far less contention, so they are limited only by the mlsrv11 -w option. For this reason, the -w setting must be greater than or equal to the -wu setting.

By default, all database worker threads can apply uploads simultaneously. The number of database worker threads that are used is set by the -w option. The default is 5.

### Example

In a pilot setup using a LAN and remote databases on PCs, you find that the optimum number of database worker threads is approximately 10 for both upload-only and download-only synchronizations, and that corresponds to 100% CPU utilization on the consolidated database. With fewer database worker threads you find that throughput is less and the CPU utilization for the consolidated database is lower. With more database worker threads, throughput does not increase because the consolidated database is already processing as fast as it can with 10 workers.

### See also

- [“-w option” on page 105](#)
- [“-sm option” on page 94](#)

## -x option

Sets network protocol and protocol options for MobiLink clients. These are used by the MobiLink server to listen for synchronization requests.

### 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](#)].

### Syntax

```
milsrv11 -c "connection-string" -x protocol[ protocol-options ] ...
```

*protocol* : **tcpip** | **tls** | **http** | **https**

*protocol-options* : ( *option=value*; ... )

### Default

The default is TCPIP with port 2439.

### Parameters

The allowed values of protocol are as follows:

- **tcpip** Accept connections using TCP/IP.
- **tls** Accept connections using TCP/IP using transport-layer security.
- **http** Accept connections using the standard Web protocol.
- **https** Accept connections using a variant of HTTP that handles secure transactions. The HTTPS protocol implements HTTP over SSL/TLS using RSA or ECC encryption.

You can also specify the following network protocol options, in the form *option=value*. You must separate multiple options with semicolons.

- **TCP/IP options** If you specify the tcpip protocol, you can optionally specify the following protocol options (these options are case sensitive):

TCP/IP protocol option	Description
<b>host=hostname</b>	The host name or IP number on which the MobiLink server should listen. The default value is localhost.

TCP/IP protocol option	Description
<b>ignore=</b> <i>hostname</i>	A host name or IP number that gets ignored by the MobiLink server if it makes a connection. This option allows you to ignore requests from load balancers at the lowest possible level, preventing excessive output in the MobiLink server log and MobiLink Monitor output files. You can specify multiple hosts to ignore; for example <code>-x tcpip(ignore=lb1;ignore=123.45.67.89)</code> . If you specify multiple instances of <code>-x</code> on a command line, the host is ignored on all instances; for example, if you specify <code>-x tcpip(ignore=1.1.1.1)-x http</code> , then connections for 1.1.1.1 are ignored on both the TCP/IP and the HTTP streams. However, this does not affect connections via the <code>-xo</code> option.
<b>port=</b> <i>portnumber</i>	The socket port number on which the MobiLink server should listen. The default port is 2439, which is the IANA registered port number for the MobiLink server.

- Options for TCP/IP with transport-layer security** If you specify the `tls` protocol, which is TCP/IP with transport-layer security, you can optionally specify the following protocol options (these options are case sensitive):

TLS protocol options	Description
<b>fips={yes no}</b>	If you specify the TLS protocol with <code>tls_type=rsa</code> , you can specify <code>fips=yes</code> to accept connections using the TCP/IP protocol and FIPS-approved algorithms for encryption. FIPS connections use separate FIPS 140-2 certified software. Servers using RSA encryption without FIPS are compatible with clients using RSA with FIPS, and servers using RSA with FIPS are compatible with clients using RSA without FIPS.
<b>host=</b> <i>hostname</i>	The host name or IP number on which the MobiLink server should listen. The default value is <code>localhost</code> .
<b>ignore=</b> <i>hostname</i>	A host name or IP number that gets ignored by the MobiLink server if it makes a connection. This option allows you to ignore requests from load balancers at the lowest possible level, preventing excessive output in the MobiLink server log and MobiLink Monitor output files. You can specify multiple hosts to ignore; for example <code>-x tcpip(ignore=lb1;ignore=123.45.67.89)</code> .
<b>port=</b> <i>portnumber</i>	The socket port number on which the MobiLink server should listen. The default port is 2439, which is the IANA registered port number for the MobiLink server.



TLS protocol options	Description
<b>tls_type={rsa ecc}</b>	<p>If you specify the TCP/IP protocol as tls, you can specify either elliptic-curve cryptography (ecc) or RSA encryption (rsa). For backward compatibility, ecc can also be specified as certicom. The default tls_type is rsa.</p> <p>When you use TLS, you must specify an identity and an identity password:</p> <ul style="list-style-type: none"> <li>○ <b>identity=identity-file</b> Specify the path and file name of the identity file that is to be used for server authentication.</li> <li>○ <b>identity_password=password</b> Specify the password for the identity</li> </ul> <p>See “<a href="#">Starting the MobiLink server with transport-layer security</a>” [<i>SQL Anywhere Server - Database Administration</i>].</p>
<b>e2ee_type={rsa ecc}</b>	<p>The type of the key used to exchange session keys. Must be either rsa or ecc, and must match the key type in the private key file (see next option). The default e2ee_type is rsa.</p>
<b>e2ee_private_key=file</b>	<p>The PEM-encoded file containing the rsa or ecc private key. This option is required for end-to-end encryption to take effect.</p> <p>PEM-encoded files are created using the createkey utility. See “<a href="#">Key Pair Generator utility (createkey)</a>” [<i>SQL Anywhere Server - Database Administration</i>].</p>
<b>e2ee_private_key_password=password</b>	<p>The password to the private key file. This option is required for end-to-end encryption to take effect.</p>

- **HTTP options** If you specify the http protocol, you can optionally specify the following protocol options (these options are case sensitive):

HTTP options	Description
<b>buffer_size=number</b>	<p>The maximum body size for an HTTP message sent from MobiLink server, in bytes. Changing the option decreases or increases the amount of memory allocated for sending HTTP messages. The default is 65535 bytes.</p>
<b>host=hostname</b>	<p>The host name or IP number on which the MobiLink server should listen. The default value is localhost.</p>
<b>port=portnumber</b>	<p>The socket port number on which the MobiLink server should listen. The port number must match the port the MobiLink server is setup to monitor. The default port is 80.</p>

HTTP options	Description
<b>version</b> = <i>http-version</i>	The MobiLink server automatically detects the HTTP version used by a client. This parameter is a string specifying the default version of HTTP to use in case the server cannot detect the method used by the client. You have a choice of 1.0 or 1.1. The default value is 1.1.

- HTTPS options** The HTTPS protocol uses RSA or ECC digital certificates for transport-layer security. If you specify FIPS encryption, the protocol uses separate FIPS 140-2 certified software that is compatible with https.

For more information, see [“Starting the MobiLink server with transport-layer security” \[SQL Anywhere Server - Database Administration\]](#).

If you specify the **https** protocol, you can optionally specify the following protocol options (these options are case sensitive):

HTTPS options	Description
<b>buffer_size</b> = <i>number</i>	The maximum body size for an HTTPS message sent from MobiLink server, in bytes. Changing the option decreases or increases the amount of memory allocated for sending HTTPS messages. The default is 65535 bytes.
<b>identity</b> = <i>server-identity</i>	The path and file name of the identity file that is to be used for server authentication. For HTTPS, this must be an RSA certificate.
<b>identity_password</b> = <i>password</i>	An optional parameter that specifies a password for the identity file. See <a href="#">“Transport-layer security” [SQL Anywhere Server - Database Administration]</a> .
<b>fips</b> ={ <i>yes no</i> }	You can specify <b>fips=yes</b> to accept connections using the HTTPS protocol and FIPS-approved algorithms for encryption. FIPS connections use separate FIPS 140-2 certified software. Servers using RSA encryption without FIPS are compatible with clients using RSA with FIPS, and servers using RSA with FIPS are compatible with clients using RSA without FIPS.
<b>host</b> = <i>hostname</i>	The host name or IP number on which the MobiLink server should listen. The default value is localhost.
<b>port</b> = <i>portnumber</i>	The socket port number on which the MobiLink server should listen. The port number must match the port the MobiLink server is set up to monitor. The default port is 443.

HTTPS options	Description
<b>tls_type={rsa ecc}</b>	<p>If you specify the TCP/IP protocol as <code>tls</code>, you can specify either elliptic-curve cryptography (<code>ecc</code>) or RSA encryption (<code>rsa</code>). For backward compatibility, <code>ecc</code> can also be specified as <code>certicom</code>. The default <code>tls_type</code> is <code>rsa</code>.</p> <p>When you use transport-layer security, you must specify an identity and an identity password:</p> <ul style="list-style-type: none"> <li>○ <b>identity=identity-file</b> Specify the path and file name of the identity file that is to be used for server authentication.</li> <li>○ <b>identity_password=password</b> Specify the password for the identity file.</li> </ul> <p>See “Starting the MobiLink server with transport-layer security” [<a href="#">SQL Anywhere Server - Database Administration</a>].</p>
<b>version=http-version</b>	<p>The MobiLink server automatically detects the HTTP version used by a client. This parameter is a string specifying the default version of HTTP to use in case the server cannot detect the method used by the client. You have a choice of 1.0 or 1.1. The default value is 1.1.</p>
<b>e2ee_type={rsa ecc}</b>	<p>The type of the key used to exchange session keys. Must be either <code>rsa</code> or <code>ecc</code>, and must match the key type in the private key file (see next option). The default <code>e2ee_type</code> is <code>rsa</code>.</p>
<b>e2ee_private_key=file</b>	<p>The PEM-encoded file containing the <code>rsa</code> or <code>ecc</code> private key. This option is required for end-to-end encryption to take effect.</p> <p>PEM-encoded files are created using the <code>createkey</code> utility. See “Key Pair Generator utility (<code>createkey</code>)” [<a href="#">SQL Anywhere Server - Database Administration</a>].</p>
<b>e2ee_private_key_password=password</b>	<p>The password to the private key file. This option is required for end-to-end encryption to take effect.</p>

### Example

The following command line sets the port to 12345:

```
mksrv11 -c "dsn=SQL Anywhere 11 CustDB;uid=DBA;pwd=sql" -x tcpip(port=12345)
```

The following example specifies the type of security (RSA), the server identity file, and the identity password protecting the server's private key:

```
mksrv11 -c "dsn=my_cons"
-x tls(tls_type=rsa;identity=c:\test\serv_rsa1.crt;identity_password=pwd)
```

The following example is similar to the previous, except that there is a space in the identity file name:

```
mlsrv11 -c "dsn=my_cons"  
-x "tls(tls_type=rsa;identity=c:\Program Files\test  
\serv_rsal.crt;identity_password=pwd)"
```

The following example shows the use of end-to-end encryption over HTTPS:

```
mlsrv11 -c "dsn=my_cons" -x https(tls_type=rsa;identity=my_identity.crt;  
identity_password=my_id_pwd;e2ee_type=rsa;e2ee_private_key=my_pk.pem;  
e2ee_private_key_password=my_pk_pwd)
```

## -xo option

Sets network protocol and protocol options for version 8 and 9 MobiLink clients.

### Syntax

```
mlsrv11 -c "connection-string"
        -xo protocol[ protocol-options ] ...
```

*protocol-options* : ( *keyword=value*; ... )

### Remarks

Specify communications protocol through which to communicate with client applications. The default is TCPIP with port 2439.

The allowed values of protocol are as follows:

- **tcPIP** Accept connections from applications via TCP/IP.
- **http** Accept connections using the standard Web protocol. Client applications can pick their HTTP version and the MobiLink server adjusts on a per-connection basis.
- **https** Accept connections using a variant of HTTP that handles secure transactions. The HTTPS protocol implements HTTP over SSL/TLS using RSA encryption, and is compatible with any other HTTPS server.
- **https\_fips** Accept connections using the HTTPS protocol and FIPS-approved algorithms for encryption. HTTPS\_FIPS uses separate FIPS 140-2 certified software. Servers using `rsa_tls` are compatible with clients using `rsa_tls_fips`, and servers using `rsa_tls_fips` are compatible with clients using `rsa_tls`.

Optionally, you can also specify network protocol options, in the form *option=value*. You must separate multiple options with semicolons. The options you can specify depends on the protocol you choose.

- **TCP/IP options** If you specify the `tcPIP` protocol, you can optionally specify the following protocol options:
  - **backlog=number-of-connections** The maximum number of remote connections before MobiLink server should reject new synchronization requests, causing synchronization to fail on the client side. By specifying a backlog size, you can prevent clients from waiting to synchronize when the server is busy. If you do not specify a backlog size, MobiLink server accepts as many connections as possible, potentially reaching or exceeding the operating system limit on network connections. This may cause slow or erratic behavior.

When a client attempts to synchronize with a MobiLink server that has accepted its maximum number of remote connections, it receives the error code -85 (SQLE\_COMMUNICATIONS\_ERROR). The client application should handle this error and try to connect again in a few minutes.

For more information about SQLE\_COMMUNICATIONS\_ERROR, see [“Communication error” \[Error Messages\]](#).

If you are using MobiLink in an environment where thousands of simultaneous synchronizations are possible, use the backlog option to specify a maximum number of remote connections that is less than the operating system limit. See [“Plan for operating system limitations” on page 172](#).

- **host=hostname** The host name or IP number on which the MobiLink server should listen. The default value is localhost.
- **ignore=hostname** A host name or IP number that is ignored by the MobiLink server if it makes a connection. This option allows you to ignore requests from load balancers at the lowest possible level, preventing excessive output in the MobiLink server log and MobiLink Monitor output files. You can specify multiple hosts to ignore; for example `-x tcpip(ignore=lb1;ignore=123.45.67.89)`.
- **liveness\_timeout=n** The amount of time, in seconds, after the last communication with a client before MobiLink terminates the synchronization. A value of 0 means that there is no timeout. The default is 120 seconds.
- **port=portnumber** The socket port number on which the MobiLink server should listen. The default port is 2439, which is the IANA registered port number for the MobiLink server.

**Note**  
 The `mlsrv11 -x` and `-xo` options use the same default port, and `-x` is started even if you don't specify it. Therefore you must change the port for `-xo` unless you change it in the `-x` option.

- **security=cipher(keyword=value;...)** All communication through this connection is to be encrypted and authenticated using transport-layer security. Cipher can be one of:

Cipher	Description
rsa_tls	RSA encryption.
rsa_tls_fips	RSA encryption that is FIPS-approved. <code>rsa_tls_fips</code> uses separate FIPS 140-2 certified software but is compatible with clients using https (and SQL Anywhere version 9.0.2 or later).
ecc_tls	Elliptic-curve cryptography. For backward compatibility, <code>ecc_tls</code> can also be specified as <code>certicom_tls</code> .

The security parameters are `certificate` (the path and file name of the identity file that is to be used for server authentication), and `certificate_password`. You must use a certificate that matches the cipher suite you choose.

For more information, see [“Starting the MobiLink server with transport-layer security” \[SQL Anywhere Server - Database Administration\]](#).

**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” \[SQL Anywhere 11 - Introduction\]](#).

- **HTTP options** If you specify the http protocol, you can optionally specify the following protocol options:
  - **backlog=number-of-connections** The maximum number of remote connections before MobiLink server should reject new synchronization requests, causing synchronization to fail on the client side. By specifying a backlog size, you can prevent clients from waiting to synchronize when the server is busy. If you do not specify a backlog size, MobiLink server accepts as many connections as possible, potentially reaching or exceeding the operating system limit of network connections. This may cause slow or erratic behavior.

When a client attempts to synchronize with a MobiLink server that has accepted its maximum number of remote connections, it receives the error code -85 (SQLE\_COMMUNICATIONS\_ERROR). The client application should handle this error and try to connect again in a few minutes.

For more information about SQLE\_COMMUNICATIONS\_ERROR, see [“Communication error” \[Error Messages\]](#).

If you are using MobiLink in an environment where thousands of simultaneous synchronizations are possible, use the backlog option to specify a maximum number of remote connections that is less than the operating system limit. For more information, see [“Plan for operating system limitations” on page 172](#).

- **buffer\_size=number** The maximum body size for an HTTP message sent from MobiLink server, in bytes. Changing the option decreases or increase the amount of memory allocated for sending HTTP messages. The default is 65535 bytes.
- **contd\_timeout=seconds** The number of seconds the MobiLink server waits to receive the next part of a partially completed synchronization before the synchronization is abandoned. You can tune this option to free MobiLink server resources when the wait time indicates that the client can not continue the connection. The default value is 30 seconds.
- **host=hostname** The host name or IP number on which the MobiLink server should listen. The default value is localhost.
- **port=portnumber** The socket port number on which the MobiLink server should listen. The port number must match the port that the MobiLink server is monitoring. The default port is 80.

**Note**

The mlsrv11 -x and -xo options use the same default port, and -x is started even if you don't specify it. Therefore, you must change the port for -xo unless you change it in the -x option.

- **session\_key={ cookie | header }** Creates an alternative to the JSESSIONID for tracking connections. This may be necessary if your network already uses the JSESSIONID.
- **unknown\_timeout=seconds** The number of seconds the MobiLink server waits to receive HTTP headers on a new connection before the synchronization is abandoned. You can tune this option to free MobiLink server resources when the wait time indicates that a network failure has occurred. The default value is 30 seconds.
- **version=http-version** The MobiLink server automatically detects the HTTP version used by a client. This parameter is a string specifying the default version of HTTP to use in case the server cannot detect the method used by the client. You have a choice of 1.0 or 1.1. The default value is 1.1.

- **HTTPS or HTTPS\_FIPS options** The https protocol uses RSA digital certificates for transport-layer security. The https\_fips protocol uses separate FIPS 140-2 certified software but is compatible with https.

For more information, see [“Starting the MobiLink server with transport-layer security” \[SQL Anywhere Server - Database Administration\]](#).

**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” \[SQL Anywhere 11 - Introduction\]](#).

If you specify the https protocol, you can optionally specify the following protocol options:

- **backlog=number-of-connections** The maximum number of remote connections before MobiLink should reject new synchronization requests, causing synchronization to fail on the client side. By specifying a backlog size, you can prevent clients from waiting to synchronize when the server is busy. If you do not specify a backlog size, clients attempts to synchronize regardless of the size of the backlog.
- **buffer\_size=number** The maximum body size for an HTTPS message sent from MobiLink server, in bytes. Changing the option decreases or increases the amount of memory allocated for sending HTTPS messages. The default is 65535 bytes.
- **contd\_timeout=seconds** The number of seconds the MobiLink server waits to receive the next part of a partially completed synchronization before the synchronization is abandoned. You can tune this option to free MobiLink server resources when the wait time indicates that the client can not continue the connection. The default value is 30 seconds.
- **host=hostname** The host name or IP number on which the MobiLink server should listen. The default value is localhost.
- **port=portnumber** The socket port number on which the MobiLink server should listen. The port number must match the port that the MobiLink server is set up to monitor. The default port is 443.

**Note**

The mlsrv11 -x and -xo options both use the same default port and -x is started even if you don't specify it, so you must change the port for -xo unless you change it in the -x option.

- **certificate** The path and file name of the certificate file that is to be used for server authentication. This must be an RSA certificate.
- **certificate\_password** An optional parameter that specifies a password for the certificate file. For more information about security, see [“Transport-layer security” \[SQL Anywhere Server - Database Administration\]](#).
- **session\_key={ cookie | header }** Creates an alternative to the JSESSIONID for tracking connections. This may be necessary if your network already uses the JSESSIONID.
- **unknown\_timeout=seconds** The number of seconds the MobiLink server waits to receive HTTP headers on a new connection before the synchronization is abandoned. You can tune this



option to free MobiLink server resources when the wait time indicates that a network failure has occurred. The default value is 30 seconds.

- **version=http-version** The MobiLink server automatically detects the HTTP version used by a client. This parameter is a string specifying the default version of HTTP to use in case the server cannot detect the method used by the client. You have a choice of 1.0 or 1.1. The default value is 1.1.

### Example

The following command line sets the backlog size to 10 connections.

```
m1srv11 -c "dsn=SQL Anywhere 11 CustDB;uid=DBA;pwd=sql" -xo http(backlog=10)
```

## **-zp option**

Adjusts which timestamp values used for conflict detection purposes.

### **Syntax**

```
milsrv11 -c "connection-string" -zp
```

### **Remarks**

In the event of a timestamp conflict between the consolidated and remote database, this option allows timestamp values with a precision higher than the lowest precision to be used for conflict detection purposes. The option is useful when timestamps in the consolidated database are more precise than in the remote, as updated timestamps on the remote can cause conflicts in the next synchronization. The option allows MobiLink to ignore these conflicts. When there is a precision mismatch and `-zp` is not used, a per synchronization and a schema sensitive per table warning are written to the log to advertise the `-zp` option. Another per synchronization warning is also added to tell users to adjust the timestamp precision on the remote database where possible.

## -zs option

Specifies a MobiLink server name for mlstop and shared server state in a server farm.

### Syntax

```
mlsrv11 -c "connection-string" -zs name
```

### Remarks

The name that is specified may include ASCII letters and numbers, but no other characters.

When mlstop is used to shut down a MobiLink server started with the -zs option, you must specify the server name on the mlstop command line. For example, `mlstop myMLserver`. Shutdown may only be initiated from the computer where the MobiLink server is installed.

When MobiLink is running in a server farm, this name must be specified to uniquely identify the server.

### See also

- [“MobiLink stop utility \(mlstop\)” on page 689](#)
- [“Running the MobiLink server in a server farm” on page 40](#)
- [“-lsc option” on page 72](#)
- [“-ss option” on page 95](#)
- [“Notifiers in a MobiLink server farm” \[\*MobiLink - Server-Initiated Synchronization\*\]](#)

## **-zt option**

Specifies the maximum number of processors used to run the MobiLink server.

### **Syntax**

```
mlsrv11 -c "connection-string" -zt number
```

### **Remarks**

This option may be required for some ODBC drivers. It also gives you fine control of processor resources.

This option can only be used on Windows operating systems. The default is the number of processors on the computer.

## -zu option

Controls the automatic addition of users when the `authenticate_user` script is undefined.

### Syntax

```
mlsrv11 -c "connection-string" -zu{ + | - } ...
```

### Remarks

If this is supplied as `-zu+`, then unrecognized MobiLink user names are added to the `ml_user` table on first synchronizing. If the argument is supplied as `-zu-`, or not supplied, unrecognized user names are prevented from synchronizing.

This option is useful during development to register users. It is not recommended for deployed applications.

### See also

- “Synchronizations from new users” [*MobiLink - Client Administration*]
- “MobiLink users” [*MobiLink - Client Administration*]
- “MobiLink user authentication utility (mluser)” on page 690
- “`authenticate_user` connection event” on page 358

## **-zus option**

Causes the MobiLink server to invoke upload scripts for a table even when no rows are uploaded for the table.

### **Syntax**

```
mlsrv11 -c "connection-string" -zus ...
```

### **Remarks**

By default, if no rows are uploaded for a table, the MobiLink server does not invoke upload scripts for that table, even if they are defined. This option overrides the default behavior and causes the MobiLink server to call upload scripts for a table even if no rows are uploaded.

## -zw option

Controls which levels of warning message to display.

### Syntax

`milsrv11 -c "connection-string" -zw levels`

### Remarks

MobiLink has five levels of warning messages:

Level	Description
0	Suppress all warning messages
1	Server and high ODBC level: warning messages when the MobiLink server starts
2	Synchronization and user level: warning messages when a synchronization starts
3	Schema level: warning messages when a MobiLink server is processing a client schema
4	Script and lower ODBC level: warning messages when a MobiLink server fetches, prepares, or executes scripts
5	Table or row level: warning messages when a MobiLink server performs table operations in an upload or download

To specify the level of warning messages you want reported, you can separate levels with a comma, or separate a range with two dots. For example, **-zw 1..3,5** is the same as **-zw 1,2,3,5**.

The reporting of messages has a slight impact on performance. Levels with a higher number tend to produce more messages.

If `-zw` is used more than once in the same command line, MobiLink recognizes only the last instance. If settings of `-zw`, `-zwd`, and `-zwe` conflict, MobiLink gives priority to `-zwe`, then `-zwd`, then `-zw`.

The default is **1,2,3,4,5**, which indicates that all levels of warning message should be displayed.

## **-zwd option**

Disables specific warning codes.

### **Syntax**

```
mIsrv11 -c "connection-string" -zwd code, ...
```

### **Remarks**

You can disable specific warning codes so that they do not get reported, even though other codes of the same level are reported.

For a complete list of warning message codes, see [“MobiLink server warning messages” \[Error Messages\]](#).

If -zwd is used more than once in the same command line, MobiLink accumulates the settings. If settings of -zw, -zwd, and -zwe conflict, MobiLink gives priority to -zwe, then -zwd, then -zw.



## -zwe option

Enables specific warning codes.

### Syntax

```
mlsrv11 -c "connection-string" -zwe code, ...
```

### Remarks

You can enable specific warning codes so that they are reported even though you have disabled other codes of the same level using -zw.

For a complete list of warning message codes, see [“MobiLink server warning messages” \[Error Messages\]](#).

If -zwe is used more than once on the same command line, MobiLink accumulates the settings. If settings of -zw, -zwd, and -zwe conflict, MobiLink gives priority to -zwe, then -zwd, then -zw.

---

---

# Synchronization techniques

## Contents

MobiLink development tips .....	128
Timestamp-based downloads .....	129
Snapshot synchronization .....	133
Partitioning rows among remote databases .....	135
Upload-only and download-only synchronizations .....	138
Maintaining unique primary keys .....	139
Handling conflicts .....	146
Forced conflicts .....	154
Data entry .....	155
Handling deletes .....	156
Handling failed downloads .....	158
Download acknowledgement .....	161
Downloading a result set from a stored procedure call .....	162
Uploading data from self-referencing tables .....	164
MobiLink isolation levels .....	165

---

## MobiLink development tips

Adding synchronization functionality to an application adds an added level of complexity to your application. The following tips may be useful.

When you are adding synchronization to a prototype application, it can be difficult to see which components are causing problems. When developing a prototype, temporarily hard code INSERT statements in your application to provide data for testing and demonstration purposes. Once your prototype is working correctly, enable synchronization and discard the temporary INSERT statements.

Start with straightforward synchronization techniques. Operations such as a simple upload or download require only one or two scripts. Once those are working correctly, you can introduce more advanced techniques, such as timestamps, primary key pools, and conflict resolution.

### MobiLink and primary keys

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, MobiLink applications 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” on page 139](#).

## Timestamp-based downloads

The timestamp method is the most useful general technique for efficient synchronization. This technique involves tracking the last time that each user synchronized and only downloading rows that have changed since then.

MobiLink maintains a timestamp value indicating when each MobiLink user last downloaded data. This value is called the **last download time**.

See [“Using last download times in scripts” on page 130](#).

### To implement timestamp-based synchronization for a table

1. At the consolidated database, add a column that holds the most recent time the row was modified. The column is typically declared as follows:

DBMS	last modified column
Adaptive Server Enterprise	datetime
IBM DB2 LUW	timestamp
IBM DB2 mainframe	timestamp
Microsoft SQL Server	datetime
MySQL	timestamp
Oracle	date
SQL Anywhere	timestamp DEFAULT timestamp

2. In scripts for the `download_cursor` and `download_delete_cursor` events, compare the first parameter to the value in the timestamp column.

### Example

The following table declaration and scripts implement timestamp-based synchronization on the Customer table in the Contact sample:

- Table definition:

```
CREATE TABLE "DBA"."Customer" (
  "cust_id" integer NOT NULL DEFAULT GLOBAL AUTOINCREMENT,
  "name" char(40) NOT NULL,
  "rep_id" integer NOT NULL,
  "last_modified" timestamp NULL DEFAULT timestamp,
  "active" bit NOT NULL,
  PRIMARY KEY ("cust_id") )
```

- `download_delete_cursor` script:

```
SELECT cust_id
FROM Customer JOIN SalesRep
ON Customer.rep_id = SalesRep.rep_id
WHERE Customer.last_modified >= {ml s.last_table_download}
      AND ( SalesRep.ml_username != {ml s.username}
            OR Customer.active = 0 )
```

- `download_cursor` script:

```
SELECT cust_id, Customer.name, Customer.rep_id
FROM Customer KEY JOIN SalesRep
WHERE Customer.last_modified >= {ml s.last_table_download}
      AND SalesRep.ml_username = {ml s.username}
      AND Customer.active = 1
```

See [“Synchronization logic source code”](#) [*MobiLink - Getting Started*] and [“Synchronizing contacts in the Contact sample”](#) [*MobiLink - Getting Started*].

## Using last download times in scripts

The last download timestamp is provided as a parameter to many MobiLink events. The last download timestamp is the value obtained from the consolidated database during the last successful synchronization immediately prior to the download phase. If the current MobiLink user has never synchronized, or has never synchronized successfully, this value is set to 1900-01-01.

See [“How download timestamps are generated and used”](#) on page 131.

If you have multiple publications and have synchronized them at different times, then you can have two different last download timestamps. For this reason, there are two script parameter names for last download timestamps:

- **`last_table_download`** is the last download timestamp for a table.
- **`last_download`** is the last time all tables were synchronized. It is the earliest `last_table_download` value for any table.

When you use question marks instead of named parameters in MobiLink scripts, the correct value is always used.

### Caution

If you are using a SQL Anywhere consolidated database and the column holding last modified information is of type `DEFAULT TIMESTAMP`, then the column should not be synchronized. If your remote databases require such a column, a different column name should be used. Otherwise, the default timestamp value may be overridden by the uploaded value, and does not contain the time that the row was last modified on the consolidated database.

### See also

- [“Script parameters”](#) on page 320

## Example

```
SELECT cust_id, Customer.name, Customer.rep_id
FROM Customer KEY JOIN SalesRep
WHERE Customer.last_modified >= {ml s.last_table_download}
      AND SalesRep.ml_username = {ml s.username}
      AND Customer.active = 1
```

## How download timestamps are generated and used

MobiLink generates and uses a timestamp for timestamp-based downloads as follows:

- After an upload is committed and immediately before invoking the `prepare_for_download` event, the MobiLink server fetches the current time from the consolidated database and saves the value. This timestamp value represents the start time of the current download; the next synchronization should only download data that changes after this time.
- The MobiLink server sends this timestamp value as part of the download, and the client stores it.
- The next time the client synchronizes, it uses the timestamp value for the **last\_download\_timestamp** that it sends with the upload.
- The MobiLink server passes the `last_download_timestamp` that the client just uploaded into your `download_cursor` and `download_delete_cursor`. Your cursor can then select changes with timestamps that are newer or equal to the last `last_download_timestamp` to ensure that only new changes are downloaded.

### Where the last download time is stored

The last download time is stored on the remote database. This is the appropriate place because only the remote knows if the download has been successfully applied.

For SQL Anywhere remotes, the last download time is stored per subscription. See “[SYSSYNC system view](#)” [*SQL Anywhere Server - SQL Reference*].

For UltraLite remotes, the last download time is stored per publication. See “[syspublication system table](#)” [*UltraLite - Database Management and Reference*].

### Changing the last download time

In some rare circumstances you may want to modify the `last_download_timestamp`. For example, if you accidentally delete all the data on a remote database, you can resynchronize it by defining a `modify_last_download_timestamp` connection script to reset the value for the last download timestamp. There is another event, called `modify_next_last_download_timestamp`, which you can use to reset the timestamp not for the current synchronization but for the next synchronization. See:

- “[modify\\_last\\_download\\_timestamp connection event](#)” on page 463
- “[modify\\_next\\_last\\_download\\_timestamp connection event](#)” on page 466

UltraLite also provides functionality to change the last download time from the remote. See:

- C/C++ embedded SQL: “[ULResetLastDownloadTime function](#)” [[UltraLite - C and C++ Programming](#)]
- C++: “[ResetLastDownloadTime function](#)” [[UltraLite - C and C++ Programming](#)]
- .NET 2.0: “[ResetLastDownloadTime method](#)” [[UltraLite - .NET Programming](#)]

**See also**

- “[Using last download times in scripts](#)” on page 130

## Dealing with daylight savings time

Daylight savings time can cause problems in a distributed database system if data is synchronized during the hour that the time changes. In fact, you can lose data. This is only an issue in the autumn when the time goes back and there is a one-hour period that can be ambiguous.

To deal with daylight savings time, you have three possible solutions:

- Ensure that the consolidated database server is using UTC time.
- Turn off daylight savings time on the consolidated database server.
- Shut down for an hour when the time changes.



## Snapshot synchronization

Timestamp-based synchronization is appropriate for most synchronizations. However, occasionally you may want to update a snapshot of your data.

Snapshot synchronization of a table is a complete download of all relevant rows in the table, even if they have been downloaded before. This is the simplest synchronization method, but can involve unnecessarily large data sets being exchanged, which can limit performance.

You can use snapshot synchronization for downloading all the rows of the table, or in conjunction with a partitioning of the rows. See [“Partitioning rows among remote databases” on page 135](#).

### When to use snapshot synchronization

The snapshot method is typically most useful for tables that have both the following characteristics.

- **Relatively few rows** When there are few rows, the overhead for downloading all rows is small.
- **Rows that change frequently** When most rows in a table change frequently, there is little to be gained by explicitly excluding those that have not changed since the last synchronization.

A table that holds a list of exchange rates could be suited to this approach because there are relatively few currencies, but the rates of most change frequently. Depending on the nature of the business, a table that holds prices, a list of interest rates, or current news items could all be candidates.

### To implement snapshot-based synchronization

1. Leave the upload scripts undefined unless remote users update the values.
2. If the table may have rows deleted, write a `download_delete_cursor` script that deletes all the rows from the remote table, or at least all rows no longer required. Do not delete the rows from the consolidated database; rather, mark them for deletion. You must know the row values to delete them from the remote database.  
See [“Writing `download\_delete\_cursor` scripts” on page 335](#).
3. Write a `download_cursor` script that selects all the rows you want to include in the remote table.

### Deleting rows

Rather than deleting rows from the consolidated database, mark them for deletion. You must know the row values to delete them from the remote database. Select only unmarked rows in the `download_cursor` script and only marked rows in the `download_delete_cursor` script.

The `download_delete_cursor` script is executed before the `download_cursor` script. If a row is to be included in the download, you need not include a row with the same primary key in the delete list. When a downloaded row is received at the remote location, it replaces a preexisting row with the same primary key.

See [“Writing scripts to download rows” on page 333](#).

### An alternative deletion technique

Rather than delete rows from the remote database using a `download_cursor` script, you can allow the remote application to delete the rows. For example, immediately following synchronization, you could allow the application to execute SQL statements that delete the unneeded rows.

Rows deleted by the application are ordinarily uploaded to the MobiLink server upon the next synchronization, but you can prevent this upload using the `STOP SYNCHRONIZATION DELETE` statement. For example,

```
STOP SYNCHRONIZATION DELETE;  
DELETE FROM table-name  
  WHERE expiry_date < CURRENT_TIMESTAMP;  
COMMIT;  
START SYNCHRONIZATION DELETE;
```

See [“Writing `download\_delete\_cursor` scripts” on page 335](#).

### Snapshot example

The `ULProduct` table in the sample application is maintained by snapshot synchronization. The table contains relatively few rows, and for this reason, there is little overhead in using snapshot synchronization.

1. There is no upload script. This reflects a business decision that products cannot be added at remote databases.
2. There is no `download_delete_cursor`, reflecting an assumption that products are not removed from the list.
3. The `download_cursor` script selects the product identifier, price, and name of every current product. If the product is pre-existing, the price in the remote table is updated. If the product is new, a row is inserted in the remote table.

```
SELECT prod_id, price, prod_name  
FROM ULProduct
```

For another example of snapshot synchronization in a table with very few rows, see [“Synchronizing sales representatives in the Contact sample”](#) [*MobiLink - Getting Started*].

## Partitioning rows among remote databases

Each MobiLink remote database can contain a different subset of the data in the consolidated database. This means that you can write your synchronization scripts so that data is **partitioned** among remote databases.

The partitioning can be disjoint, or it can contain overlaps. For example, if each employee has their own set of customers, with no shared customers, the partitioning is **disjoint**. If there are shared customers who appear in more than one remote database, the partitioning contains **overlaps**.

Partitioning is implemented in the `download_cursor` and `download_delete_cursor` scripts for the table, which define the rows to be downloaded to the remote database. Each of these scripts takes a MobiLink user name as a parameter. By defining your scripts using this parameter in the WHERE clause, each user gets the appropriate rows.

### Disjoint partitioning

Partitioning is controlled by the `download_cursor` and `download_delete_cursor` scripts for each table involved in synchronization. These scripts take two parameters, a last download timestamp and the MobiLink user name supplied in the call to synchronize.

#### To partition a table among remote databases

1. Include in the table definition a column containing the synchronization user name in the consolidated database. You need not download this column to remote databases.
2. Include a condition in the WHERE clause of the `download_cursor` and `download_delete_cursor` scripts requiring this column to match the script parameter.

The script parameter can be represented by a question mark or a named parameter in the script. For example, the following `download_cursor` script partitions the Contact table by employee ID.

```
SELECT id, contact_name
FROM Contact
WHERE last_modified >= {ml s.last_table_download}
AND emp_id = {ml s.username}
```

See [“download\\_cursor table event” on page 396](#) and [“download\\_delete\\_cursor table event” on page 400](#).

#### Example

The primary key pool tables in the CustDB sample application are used to supply each remote database with its own set of primary key values. This technique is used to avoid duplicate primary keys, and is discussed in [“Using primary key pools” on page 143](#).

A necessary feature of the method is that primary key-pool tables must be partitioned among remote databases in a disjoint fashion.

One key-pool table is ULCustomerIDPool, which holds primary key values for each user to use when they add customers. The table has three columns:

- **pool\_cust\_id** A primary key value for use in the ULCustomer table. This is the only column downloaded to the remote database.
- **pool\_emp\_id** The employee who owns this primary key.
- **last\_modified** This table is maintained using the timestamp technique, based on the last\_modified column.

For information about timestamp synchronization, see [“Timestamp-based downloads” on page 129](#).

The download\_cursor script for this table is as follows.

```
SELECT pool_cust_id
FROM ULCustomerIDPool
WHERE last_modified >= {ml s.last_table_download}
AND pool_emp_id = {ml s.username}
```

When not using a variable or named parameter, you can use a join or sub-selection that includes the ? placeholder.

See [“Synchronizing customers in the Contact sample” \[MobiLink - Getting Started\]](#) and [“Synchronizing contacts in the Contact sample” \[MobiLink - Getting Started\]](#).

## Partitioning with overlaps

Some tables in your consolidated database may have rows that belong to many remote databases. Each remote database has a subset of the rows in the consolidated database and the subset overlaps with other remote databases. This is frequently the case with a customer table. In this case, there is a many-to-many relationship between the table and the remote databases and there is usually a table to represent the relationship. The scripts for the download\_cursor and download\_delete\_cursor events need to join the table being downloaded to the relationship table.

### Example

The CustDB sample application uses this technique for the ULOrder table. The ULEmpCust table holds the many-to-many relationship information between ULCustomer and ULEmployee.

Each remote database receives only those rows from the ULOrder table for which the value of the emp\_id column matches the MobiLink user name.

The SQL Anywhere version of the download\_cursor script for ULOrder in the CustDB application is as follows:

```
SELECT o.order_id, o.cust_id, o.prod_id,
       o.emp_id, o.disc, o.quant, o.notes, o.status
FROM ULOrder o , ULEmpCust ec
WHERE o.cust_id = ec.cust_id
AND ec.emp_id = {ml s.username}
AND ( o.last_modified >= {ml s.last_table_download}
OR ec.last_modified >= {ml s.last_table_download} )
AND ( o.status IS NULL
OR o.status != 'Approved' )
AND ( ec.action IS NULL )
```

This script is fairly complex. It illustrates that the query defining a table in the remote database can include more than one table in the consolidated database. The script downloads all rows in ULOrder for which the following are all true:

- the cust\_id column in ULOrder matches the cust\_id column in ULEmpCust
- the emp\_id column in ULEmpCust matches the synchronization user name
- the last modification of either the order or the employee-customer relationship was later than the most recent synchronization time for this user
- the status is anything other than **Approved**

The action column on ULEmpCust is used to mark columns for delete. Its purpose is not relevant to the current topic.

The download\_delete\_cursor script is as follows.

```
SELECT o.order_id
FROM ULOrder o, ULEmpCust ec
WHERE o.cust_id = ec.cust_id
      AND ec.emp_id = {ml s.username}
      AND ( o.last_modified >= {ml s.last_table_download} OR
            c.last_modified >= {ml s.last_table_download} )
      AND ( o.status IS NULL OR
            o.status != 'Approved' )
      AND ( ec.action IS NULL )
```

This script deletes all approved rows from the remote database.

## Partitioning child tables

The example in the previous section illustrates how to partition tables based on a criterion in some other table. See [“Partitioning with overlaps” on page 136](#).

Some tables in your remote database may have disjoint subsets or overlapping subsets, but do not contain a column that determines the subset. These are child tables that usually have a foreign key (or a series of foreign keys) referencing another table. The referenced table has a column that determines the correct subset.

In this case, the download\_cursor script and the download\_delete\_cursor script need to join the referenced tables and have a WHERE clause that restricts the rows to the correct subset.

For an example, see [“Synchronizing contacts in the Contact sample” \[MobiLink - Getting Started\]](#).

## Upload-only and download-only synchronizations

By default, synchronization is bi-directional: data is both uploaded and downloaded. However, you can choose to do only an upload or only a download.

### Synchronization model note

This topic provides information for how to set up upload-only and download-only synchronization when you create your MobiLink synchronization system in your database. You can also specify upload-only or download-only if you create a synchronization model in Sybase Central.

### SQL Anywhere remote databases

- **Upload** To perform upload-only synchronization, use the dbmlsync option `-uo` or the extended option `UploadOnly`. See:
  - “`-uo` option” [*MobiLink - Client Administration*]
  - “`UploadOnly (uo)` extended option” [*MobiLink - Client Administration*]
- **Download** To perform download-only synchronization, use the dbmlsync option `-ds` or the extended option `DownloadOnly`. See:
  - “`-ds` option” [*MobiLink - Client Administration*]
  - “`DownloadOnly (ds)` extended option” [*MobiLink - Client Administration*]

SQL Anywhere remote databases can also use download-only publications. This approach to downloads is different from download-only synchronizations. See “[Download-only publications](#)” [*MobiLink - Client Administration*].

### UltraLite remote databases

- **Upload** To perform upload-only synchronization, use the `Upload Only` synchronization parameter. See “[Upload Only synchronization parameter](#)” [*UltraLite - Database Management and Reference*].
- **Download** To perform download-only synchronization, use the `Download Only` synchronization parameter. See “[Download Only synchronization parameter](#)” [*UltraLite - Database Management and Reference*].

## Maintaining unique primary keys

Every table that is to be synchronized must have a primary key, and the primary key must be unique across all synchronized databases. The values of primary keys should not be updated.

It is often convenient to use a single column as the primary key for tables. For example, each customer should be assigned a unique identification value. If all the sales representatives work in an environment where they can maintain a direct connection to the database, assigning these numbers is easily accomplished. Whenever a new customer is inserted into the customer table, automatically add a new primary key value that is greater than the last value.

In a disconnected environment, assigning unique values for primary keys when new rows are inserted is not as easy. When a sales representative adds a new customer, she is doing so to a remote copy of the Customer table. You must prevent other sales representatives, working on other copies of the Customer table, from using the same customer identification value.

This section describes the following ways to solve the problem of how to generate unique primary keys:

- [“Using composite keys” on page 139](#)
- [“Using UUIDs” on page 139](#)
- [“Using global autoincrement” on page 140](#)
- [“Using primary key pools” on page 143](#)

## Using composite keys

The MobiLink remote ID uniquely defines a remote database within a synchronization system. Therefore, an easy way to create a unique primary key is to create a composite primary key that includes the MobiLink remote ID as part of its value. If you maintain unique MobiLink user names, you could use the user name instead of the remote ID.

See [“Remote IDs” \[MobiLink - Client Administration\]](#).

## Using UUIDs

You can ensure that primary keys are unique by using the `newid()` function to create universally unique values for your primary key. The resulting UUIDs can be converted to a string using the `uuidtostr()` function, and converted back to binary using the `strtouuid()` function.

UUIDs, also known as GUIDs, are unique across all computers. However, the values are completely random and so cannot be used to determine when a value was added, or the order of values. UUID values are also considerably larger than the values required by other methods (including global autoincrement), and require more table space in both the primary and foreign key tables. Indexes on tables using UUIDs are also less efficient.

## See also

SQL Anywhere databases:

- “The NEWID default” [[SQL Anywhere Server - SQL Usage](#)]
- “NEWID function [Miscellaneous]” [[SQL Anywhere Server - SQL Reference](#)]
- “UNIQUEIDENTIFIER data type” [[SQL Anywhere Server - SQL Reference](#)]

UltraLite databases:

- “Primary key uniqueness in UltraLite” [[UltraLite - Database Management and Reference](#)]
- “NEWID function [Miscellaneous]” [[UltraLite - Database Management and Reference](#)]

## Example

The following SQL Anywhere CREATE TABLE statement creates a primary key that is universally unique:

```
CREATE TABLE customer (  
    cust_key UNIQUEIDENTIFIER NOT NULL  
        DEFAULT NEWID( ),  
    rep_key VARCHAR(5),  
    PRIMARY KEY(cust_key))
```

## Using global autoincrement

In SQL Anywhere and UltraLite databases, you can set the default column value to be GLOBAL AUTOINCREMENT. You can use this default for any column in which you want to maintain unique values, but it is particularly useful for primary keys.

### To use global autoincrement columns

1. Declare the column as a global autoincrement column.

When you specify default global autoincrement, the domain of values for that column is partitioned. Each partition contains the same number of values. For example, if you set the partition size for an integer column in a database to 1000, one partition extends from 1001 to 2000, the next from 2001 to 3000, and so on.

See “[Declaring default global autoincrement](#)” on page 141.

2. Set the global\_database\_id value.

SQL Anywhere supplies default values in a database only from the partition uniquely identified by that database's number. For example, if you assign a database the identity number 10 and the partition size is 1000, the default values in that database would be chosen in the range 10001-11000. Another copy of the database, assigned the identification number 11, would supply default value for the same column in the range 11001-12000.

See “[Setting the global database ID](#)” on page 141.



## Declaring default global autoincrement

You can set default values in your database by selecting the column properties in Sybase Central, or by including the `DEFAULT GLOBAL AUTOINCREMENT` phrase in a `CREATE TABLE` or `ALTER TABLE` statement.

Optionally, the partition size can be specified in parentheses immediately following the `AUTOINCREMENT` keyword. The partition size may be any positive integer, although the partition size is generally chosen so that the supply of numbers within any one partition is rarely, if ever, 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, especially if our column is not of type `INT` or `BIGINT`, it is best to specify the partition size explicitly.

For example, the following SQL statement creates a simple table with two columns: an integer that holds a customer identification number and a character string that holds the customer's name. The partition size is set to 5000.

```
CREATE TABLE customer (
  id   INT           DEFAULT GLOBAL AUTOINCREMENT (5000),
  name VARCHAR(128) NOT NULL,
  PRIMARY KEY (id)
)
```

### See also

- SQL Anywhere: “[CREATE TABLE statement](#)” [[SQL Anywhere Server - SQL Reference](#)]
- UltraLite: “[UltraLite CREATE TABLE statement](#)” [[UltraLite - Database Management and Reference](#)]

## Setting the global database ID

When deploying an application, you must assign a different identification number to each database. You can create and distribute the identification numbers by a variety of means. One method is to place the values in a table and download the correct row to each database based on some other unique property, such as remote ID.

### To set the global database identification number

- In SQL Anywhere, you set the global ID of a database by setting the value of the public option `global_database_id`. The identification number must be a non-negative integer. See “[global\\_database\\_id option \[database\]](#)” [[SQL Anywhere Server - Database Administration](#)].

In UltraLite, you set the global ID of a database by setting the `global_id` option. See “[UltraLite global\\_database\\_id option](#)” [[UltraLite - Database Management and Reference](#)].

### How default values are chosen

The global database ID is set with the public option `global_database_id` in SQL Anywhere, and with the `global_id` option in UltraLite.

The global database id option in each database must be set to a unique, non-negative integer. The range of default values for a particular database is  $pn + 1$  to  $p(n + 1)$ , where  $p$  is the partition size and  $n$  is the value

of the global database ID. For example, if the partition size is 1000 and global database ID is set to 3, then the range is from 3001 to 4000.

SQL Anywhere and UltraLite choose default values by applying the following rules:

- If the column contains no values in the current partition, the first default value is  $pn + 1$ , where  $p$  is the partition size and  $n$  is the value of the global database ID.
- If the column contains values in the current partition, but all are less than  $p(n + 1)$ , the next default value is 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.

If the global database ID is set to the default value of 2147483647, a null value is inserted into the column. Should null values not be permitted, the attempt to insert the row causes an error. This situation arises, for example, if the column is contained in the table's primary key.

Because the global database ID cannot be set to negative values, the values chosen are always positive. The maximum identification number is restricted only by the column data type and the partition size.

Null default values are also generated when the supply of values within the partition has been exhausted. In this case, a new global database ID value should be assigned to the database to allow default values to be chosen from another partition. Attempting to insert the null value causes an error if the column does not permit nulls. To detect that the supply of unused values is low and handle this condition, you can create an event of type GlobalAutoincrement.

Should the values in a particular partition become exhausted, you can assign a new global database ID to that database. You can assign new database ID numbers in any convenient manner. However, one possible technique is to maintain a pool of unused database ID values. This pool is maintained in the same manner as a pool of primary keys.

You can set an event handler to automatically notify the database administrator (or perform some other action) when the partition is nearly exhausted. For SQL Anywhere databases, see [“Defining trigger conditions for events” \[SQL Anywhere Server - Database Administration\]](#).

### See also

- [“Setting the global database ID” on page 141](#)
- SQL Anywhere: [“global\\_database\\_id option \[database\]” \[SQL Anywhere Server - Database Administration\]](#)
- UltraLite: [“UltraLite global\\_database\\_id option” \[UltraLite - Database Management and Reference\]](#)

### Example

In a SQL Anywhere database, the following statement sets the database identification number to 20.

```
SET OPTION PUBLIC.global_database_id = 20
```

If the partition size for a particular column is 5000, default values for this database are selected from the range 100001-105000.

## Using primary key pools

One efficient means of solving this problem of unique primary keys is to assign each user of the database a pool of primary key values that can be used as the need arises. For example, you can assign each sales representative 100 new identification values. Each sales representative can freely assign values to new customers from his or her own pool.

### To implement a primary key pool

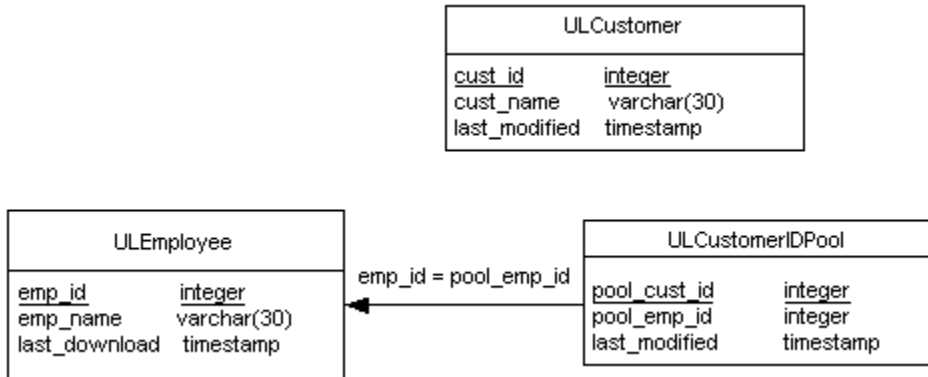
1. Add a new table to the consolidated database and to each remote database to hold the new primary key pool. Apart from a column for the unique value, these tables should contain a column for a user name, to identify who has been given the right to assign the value.
2. Write a stored procedure to ensure that each user is assigned enough new identification values. Assign more new values to remote users who insert many new entries or who synchronize infrequently.
3. Write a `download_cursor` script to select the new values assigned to each user and download them to the remote database.
4. Modify the application that uses the remote database so that when a user inserts a new row, the application uses one of the values from the pool. The application must then delete that value from the pool so it is not used a second time.
5. Write an upload script. The MobiLink server then deletes rows from the consolidated pool of values that a user has deleted from his personal value pool in the remote database.
6. Write an `end_upload` script to call the stored procedure that maintains the pool of values. Doing so has the effect of adding more values to the user's pool to replace those deleted during upload.

### Example

The sample application allows remote users to add customers. It is essential that each new row has a unique primary key value, and yet each remote database is disconnected when data entry is occurring.

The `ULCustomerIDPool` holds a list of primary key values that can be used by each remote database. In addition, the `ULCustomerIDPool_maintain` stored procedure tops up the pool as values are used up. The maintenance procedures are called by a table-level `end_upload` script, and the pools at each remote database are maintained by `upload_insert` and `download_cursor` scripts.

1. The `ULCustomerIDPool` table in the consolidated database holds the pool of new customer identification numbers. It has no direct link to the `ULCustomer` table.



- The ULCustomerIDPool\_maintain procedure updates the ULCustomerIDPool table in the consolidated database. The following sample code is for a SQL Anywhere consolidated database.

```

CREATE PROCEDURE ULCustomerIDPool_maintain ( IN syncuser_id INTEGER )
BEGIN
    DECLARE pool_count INTEGER;

    -- Determine how many ids to add to the pool
    SELECT COUNT(*) INTO pool_count
    FROM ULCustomerIDPool
    WHERE pool_emp_id = syncuser_id;

    -- Top up the pool with new ids
    WHILE pool_count < 20 LOOP
        INSERT INTO ULCustomerIDPool ( pool_emp_id )
        VALUES ( syncuser_id );
        SET pool_count = pool_count + 1;
    END LOOP;
END

```

This procedure counts the numbers that are currently assigned to the current user, and inserts new rows so that this user has enough customer identification numbers.

This procedure is called at the end of the upload, by the end\_upload table script for the ULCustomerIDPool table. The script is as follows:

```
CALL ULCustomerIDPool_maintain( {ml s.username} )
```

- The download\_cursor script for the ULCustomerIDPool table downloads new numbers to the remote database.

```

SELECT pool_cust_id
FROM ULCustomerIDPool
WHERE pool_emp_id = {ml s.username}
AND last_modified >= {ml s.last_table_download}

```

- To insert a new customer, the application using the remote database must select an unused identification number from the pool, delete this number from the pool, and insert the new customer information using this identification number. The following embedded SQL function for an UltraLite application retrieves a new customer number from the pool.

```

bool CDemoDB::GetNextCustomerID( void )
/*****
{

```

```
short ind;

EXEC SQL SELECT min( pool_cust_id )
INTO :m_CustID:ind FROM ULCustomerIDPool;
if( ind < 0 ) {
    return false;
}
EXEC SQL DELETE FROM ULCustomerIDPool
WHERE pool_cust_id = :m_CustID;
return true;
}
```

## Handling conflicts

Conflicts can arise during the upload of rows to the consolidated database. If two users modify the same row on different remote databases, a conflict is detected when the second of the rows arrives at the MobiLink server.

By default,

- If an attempt to insert a row finds that the row has already been inserted, an error is generated.
- If an attempt to delete a row finds that the row has already been deleted, the second attempt to delete is ignored.

If you need different behavior, you can implement it by defining one or more of the upload events that are described in this section.

## About conflicts

### Caution

Never update primary keys in synchronized tables. Updating primary keys defeats the purpose of a primary key because the key is the only way to identify the same row in different databases (remote and consolidated) and the only way to detect conflicts.

Conflicts are not the same as errors. When conflicts can occur, you should define a process to compute the correct values, or at least to log the conflict. Conflict handling is an integral part of a well-designed application.

During the download stage of a synchronization, no conflicts arise in the remote database. If a downloaded row contains a new primary key, the values are inserted into a new row. If the primary key matches that of a pre-existing row, the values in the row are updated.

### Example

User1 starts with an inventory of ten items, and then sells three and updates the Remote1 inventory value to seven items. User2 sells four items and updates the Remote2 inventory to six. When Remote1 synchronizes, the consolidated database is updated to seven. When Remote2 synchronizes, a conflict is detected because the value of the inventory is no longer ten. To resolve this conflict programmatically, you need three row values:

1. The current value in the consolidated database.
2. The new row value that Remote2 uploaded.
3. The old row value that Remote2 obtained during the last synchronization.

In this case, the business logic could use the following to calculate the new inventory value and resolve the conflict:

```
current consolidated - (old remote - new remote)
-> 7 - (10-6) = 3
```

For other examples of how to handle conflicts, see:

- “Synchronizing products in the Contact sample” [[MobiLink - Getting Started](#)]

## Detecting conflicts

When a MobiLink client sends an updated row to the MobiLink server, it includes not only the new updated values (the post-image), but also a copy of the old row values (the pre-image) obtained either in the last download or from the row values existing prior to the first upload of this row. When the pre-image does not match the current values in the consolidated database, a conflict is detected.

There are several scripts provided to detect conflicts. The MobiLink server detects conflicts only if one of the following scripts is applied:

- An `upload_fetch` or `upload_fetch_column_conflict` script.
- An `upload_update` script that includes all non-primary key columns in the WHERE clause.

## Detecting conflicts with upload\_fetch scripts

If you define an `upload_fetch` or `upload_fetch_column_conflict` script for a table, the MobiLink server compares the pre-image of an update to the values of the row returned by the `upload_fetch` script with the same primary key values. If values in the pre-image do not match the current consolidated values, the MobiLink server detects a conflict.

The `upload_fetch` script selects a single row of data from a consolidated database table corresponding to the row being updated. A typical `upload_fetch` script has the following syntax:

```
SELECT pk1, pk2, ...col1, col2, ...
FROM table-name
WHERE pk1 = {ml r.pk1} AND pk2 = {ml r.pk2} ...
      AND col1 = {ml r.col1} AND col2 = {ml r.col2} ...
```

See “[upload\\_fetch table event](#)” on page 500.

The `upload_fetch_column_conflict` event is similar to `upload_fetch`, but it only detects a conflict when two users update the same column. Different users can update the same row, as long as they don't update the same column, without generating a conflict.

See “[upload\\_fetch\\_column\\_conflict table event](#)” on page 502.

You can have only one `upload_fetch` or `upload_fetch_column_conflict` script for each table in the remote database.

## Locking the row on the consolidated database

It is possible that a row might change on the consolidated database after the `upload_fetch` script detects a conflict and before the conflict resolution is completed. To avoid this problem, which could result in incorrect data, you can implement the `upload_fetch` or `upload_fetch_column_conflict` scripts with a row lock.

In SQL Anywhere consolidated databases, you can use either the UPDLOCK or HOLDLOCK keywords, but UPDLOCK is better for concurrency. For example:

```
SELECT column-names from table-name WITH (UPDLOCK)
WHERE where-clause
```

For Oracle, DB2 LUW and DB2 mainframe, use FOR UPDATE. For example:

```
SELECT column-names FROM table-name
WHERE where-clause
FOR UPDATE OF column_name1, column_name3, column_name6
```

### Note

Specifying the column names you want to update, as shown in the example above, preserves computer resources and improves performance.

For Microsoft SQL Server, use HOLDLOCK. For example,

```
SELECT column-names FROM table-name WITH (HOLDLOCK)
WHERE where-clause
```

For Adaptive Server Enterprise, use HOLDLOCK. For example,

```
SELECT column-names FROM table-name
HOLDLOCK
WHERE where-clause
```

## Example

You define an `upload_fetch` script. The MobiLink server uses the script to retrieve the current row in the consolidated database and compares this row to the pre-image of the updated row. If the two rows contain identical values, there is no conflict. If the two rows differ, then a conflict is detected and MobiLink calls the `upload_old_row_insert` and `upload_new_row_insert` scripts, followed by `resolve_conflict`.

See [“Resolving conflicts with resolve\\_conflict scripts”](#) on page 149.

## Detecting conflicts with upload\_update scripts

To use the `upload_update` script to detect conflicts, include all columns in the WHERE clause:

```
UPDATE table-name
SET col1 = {ml r.col1}, col2 = {ml r.col2} ...
WHERE pk1 = {ml o.pk1} AND pk2 = {ml o.pk2} ...
AND col1 = {ml o.col1} AND col2 = {ml o.col2} ...
```

In this statement, `col1`, `col2` and so on are the non-primary key columns, while `pk1`, `pk2` and so on are primary key columns. The values passed to the second set of non-primary key columns (o.) are the pre-image (or old values) of the updated row. The WHERE clause compares old values uploaded from the remote to current values in the consolidated database. If the values do not match, the update is ignored, preserving the values already on the consolidated database.

See [“upload\\_update table event”](#) on page 522.

The `upload_update` script is used for conflict detection only if no conflict is detected by `upload_fetch` or `upload_fetch_column_conflict`.



## Scenario 1

You define scripts for the following events: `upload_update`, `upload_old_row_insert`, `upload_new_row_insert`, and `resolve_conflict`.

You define the following `upload_update` script:

```
UPDATE product
SET name={ml r.name}, description={ml r.description}
WHERE id={ml r.id}
      AND name={ml o.name}
      AND description={ml o.description}
```

MobiLink performs the update and then checks to see how many rows were modified. If no rows were modified, then MobiLink has detected a conflict: no row in the consolidated database matches the pre-image row. MobiLink calls the `upload_old_row_insert` and `upload_new_row_insert` scripts, followed by `resolve_conflict`.

See [“Resolving conflicts with `resolve\_conflict` scripts” on page 149](#).

## Scenario 2

You do not define scripts for `upload_old_row_insert`, `upload_new_row_insert`, and `resolve_conflict`. Instead, you create a stored procedure to handle the conflict detection and resolution and you call it in the `upload_update` script.

See [“Resolving conflicts with `upload\_update` scripts” on page 151](#).

# Resolving conflicts

You have several options for resolving conflicts:

- Resolve conflicts as they occur using temporary or permanent tables and a `resolve_conflict` script.  
See [“Resolving conflicts with `resolve\_conflict` scripts” on page 149](#).
- Resolve conflicts as they occur using an `upload_update` script.  
See [“Resolving conflicts with `upload\_update` scripts” on page 151](#).
- Resolve all conflicts at once using a table's `end_upload` script.  
See [“`end\_upload` table event” on page 433](#).

## Resolving conflicts with `resolve_conflict` scripts

When the MobiLink server detects a conflict using an `upload_fetch` script, the following events take place.

- The MobiLink server inserts old row values uploaded from the remote database as defined by the `upload_old_row_insert` script. Typically, the old values are inserted into a temporary table.  
See [“`upload\_old\_row\_insert` table event” on page 509](#).

- The MobiLink server inserts the new row values uploaded from the remote database as defined by the `upload_new_row_insert` script. Typically, the new values are inserted into a temporary table.  
See “[upload\\_new\\_row\\_insert table event](#)” on page 506.
- The MobiLink server executes the `resolve_conflict` script. In this script you can either call a stored procedure, or define a sequence of steps to resolve the conflict using the new and old row values.

For more information, see “[resolve\\_conflict table event](#)” on page 483.

### Example

In the following example, you create scripts for six events and then you create a stored procedure.

- In the `begin_synchronization` script, you create two temporary tables called `contact_new` and `contact_old`. (You could also do this in the `begin_connection` script.)
- The `upload_fetch` script detects the conflict.
- When there is a conflict, the `upload_old_row_insert` and `upload_new_row_insert` scripts populate the two temporary tables with the new and old data uploaded from the remote database.
- The `resolve_conflict` script calls the stored procedure `MLResolveContactConflict` to resolve the conflict.

Event	Script
<code>begin_synchronization</code>	<pre>CREATE TABLE #contact_new(   id INTEGER,   location CHAR(36),   contact_date DATE); CREATE TABLE #contact_old(   id INTEGER,   location CHAR(36),   contact_date DATE)</pre>
<code>upload_fetch</code>	<pre>SELECT id, location, contact_date FROM contact WHERE id = {ml r.id}</pre>
<code>upload_old_row_insert</code>	<pre>INSERT INTO #contact_new( id, location, contact_date ) VALUES ( {ml r.id}, {ml r.location}, {ml r.contact_date} )</pre>
<code>upload_new_row_insert</code>	<pre>INSERT INTO #contact_old( id, location, contact_date ) VALUES ( {ml r.id}, {ml r.location}, {ml r.contact_date} )</pre>
<code>resolve_conflict</code>	<pre>CALL MLResolveContactConflict( )</pre>
<code>end_synchronization</code>	<pre>DROP TABLE #contact_new; DROP TABLE #contact_old</pre>

The stored procedure `MLResolveContactConflict` is as follows:

```
CREATE PROCEDURE MLResolveContactConflict( )
BEGIN
```

```

--update the consolidated database only if the new contact date
--is later than the existing contact date
UPDATE contact c
  SET c.contact_date = cn.contact_date
  FROM #contact_new cn
  WHERE c.id = cn.id
  AND cn.contact_date > c.contact_date;
--cleanup
DELETE FROM #contact_new;
DELETE FROM #contact_old;
END

```

## Resolving conflicts with upload\_update scripts

Instead of using the `resolve_conflict` script for conflict resolution, you can call a stored procedure in the `upload_update` script. With this technique, you must both detect and resolve conflicts programmatically.

The stored procedure must use the format of the `upload_update` script with a `WHERE` clause that includes all columns but uses the pre-image (old) values.

The `upload_update` script could be as follows:

```

{CALL UpdateProduct(
  {ml o.id}, {ml o.name}, {ml o.desc}, {ml r.name}, {ml r.desc}
)}

```

The `UpdateProduct` stored procedure could be:

```

CREATE PROCEDURE UpdateProduct(
  @id INTEGER,
  @preName VARCHAR(20),
  @preDesc VARCHAR(200),
  @postName VARCHAR(20),
  @postDesc VARCHAR(200) )
BEGIN
  UPDATE product
  SET name = @postName, description = @postDesc
  WHERE id = @id
  AND name = @preName
  AND description = @preDesc
  IF @@rowcount=0 THEN
    // A conflict occurred: handle resolution here.
  END IF
END

```

This approach is often easier to maintain than resolving conflicts with `resolve_conflict` scripts because there is only one script to maintain and all the logic is contained in one stored procedure. However, the code of the stored procedure may be complicated if the tables columns are nullable or if they contain BLOBs or CLOBs. Also, some RDBMSs that are supported MobiLink consolidated databases have limitations on the size of values that can be passed to stored procedures.

See:

- [“Detecting conflicts with upload\\_update scripts” on page 148](#)
- [“upload\\_update table event” on page 522](#)
- [“Resolving conflicts with resolve\\_conflict scripts” on page 149](#)

**Example**

The following stored procedure, `sp_update_my_customer`, contains logic for conflict detection and resolution. It accepts old column values and new column values. This example uses SQL Anywhere features. The script could be implemented as follows.

```
{CALL sp_update_my_customer(
  {ml o.cust_1st_pk},
  {ml o.cust_2nd_pk},
  {ml o.first_name},
  {ml o.last_name},
  {ml o.nullable_col},
  {ml o.last_modified},
  {ml r.first_name},
  {ml r.last_name},
  {ml r.nullable_col},
  {ml r.last_modified}
)}
CREATE PROCEDURE sp_update_my_customer(
  @cust_1st_pk      INTEGER,
  @cust_2nd_pk     INTEGER,
  @old_first_name  VARCHAR(100),
  @old_last_name   VARCHAR(100),
  @old_nullable_col VARCHAR(20),
  @old_last_modified DATETIME,
  @new_first_name  VARCHAR(100),
  @new_last_name   VARCHAR(100),
  @new_nullable_col VARCHAR(20),
  @new_last_modified DATETIME
)
BEGIN
  DECLARE @current_last_modified DATETIME;
  // Detect a conflict by checking the number of rows that are
  // affected by the following update. The WHERE clause compares
  // old values uploaded from the remote to current values in
  // the consolidated database. If the values match, there is
  // no conflict. The COALESCE function returns the first non-
  // NULL expression from a list, and is used in this case to
  // compare values for a nullable column.

  UPDATE my_customer
  SET first_name      = @new_first_name,
      last_name       = @new_last_name,
      nullable_col    = @new_nullable_col,
      last_modified   = @new_last_modified

  WHERE cust_1st_pk   = @cust_1st_pk
     AND cust_2nd_pk = @cust_2nd_pk
     AND first_name   = @old_first_name
     AND last_name    = @old_last_name
     AND COALESCE(nullable_col, '') = COALESCE(@old_nullable_col, '')
     AND last_modified = @old_last_modified;
  ...
  // Use the @@rowcount global variable to determine
  // the number of rows affected by the update. If @@rowcount=0,
  // a conflict has occurred. In this example, the database with
  // the most recent update wins the conflict. If the consolidated
  // database wins the conflict, it retains its current values
  // and no action is taken.

  IF( @@rowcount = 0 ) THEN
  // A conflict has been detected. To resolve it, use business
```

```
// logic to determine which values to use, and update the
// consolidated database with the final values.

SELECT last_modified INTO @current_last_modified
FROM my_customer WITH( HOLDLOCK )
WHERE cust_1st_pk=@cust_1st_pk
      AND cust_2nd_pk=@cust_2nd_pk;

IF( @new_last_modified > @current_last_modified ) THEN
// The remote has won the conflict: use the values it
// uploaded.

UPDATE my_customer
SET first_name   = @new_first_name,
    last_name    = @new_last_name,
    nullable_col = @new_nullable_col,
    last_modified = @new_last_modified
WHERE cust_1st_pk = @cust_1st_pk
      AND cust_2nd_pk = @cust_2nd_pk;

END IF;
END IF;
END;
```

See:

- [“COALESCE function \[Miscellaneous\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [@@rowcount in “Global variables” \[SQL Anywhere Server - SQL Reference\]](#)

## Forced conflicts

Forced conflict resolution is a special technique that forces every uploaded row to be treated as if it were a conflict.

The MobiLink server uses forced conflict resolution when the `upload_insert`, `upload_update`, and `upload_delete` script are all undefined. In this mode of operation, the MobiLink server attempts to insert all uploaded rows from that table using the statements defined by the `upload_old_row_insert` and `upload_new_row_insert` scripts. In essence, all uploaded rows are then treated as conflicts. You can write stored procedures or scripts to process the uploaded values in any way you want.

Without any of the `upload_insert`, `upload_update`, or `upload_delete` scripts, the normal conflict-resolution procedure is bypassed. This technique has two principal uses.

- **Arbitrary conflict detection and resolution** The automatic mechanism only detects errors when updating a row, and only then when the old values do not match the present values in the consolidated database.

You can capture the raw uploaded data using the `upload_old_row_insert` and `upload_new_row_insert` scripts, then process the rows as you see fit.

- **Performance** When the `upload_insert`, `upload_update`, and `upload_delete` are not defined, the MobiLink server is relieved of its normal conflict-detection tasks, which involve querying the consolidated database one row at a time. Instead, it needs only to insert the raw uploaded information using the statements defined by the `upload_old_row_insert` and `upload_new_row_insert` scripts. Performance is improved because the MobiLink server is not fetching rows across the network.

### See also

- [“Forced conflict statistics” on page 198](#)

## Data entry

In some databases, there are tables that are only used for data entry. One way of processing these tables is to upload all inserted rows at each synchronization, and remove them from the remote database on the download. After synchronization, the remote table is empty again, ready for another batch of data.

To achieve this model, you can upload rows into a temporary table and then insert them into a base table using an end\_upload table script. The temporary table can be used in the download\_delete\_cursor to remove rows from the remote database following a successful synchronization.

Alternatively, you can allow the client application to delete the rows, using the STOP SYNCHRONIZATION DELETE statement to stop the deletes being uploaded during the next synchronization.

See “[STOP SYNCHRONIZATION DELETE statement \[MobiLink\]](#)” [*SQL Anywhere Server - SQL Reference*].

## Handling deletes

When rows are deleted from the consolidated database, there needs to be a record of the row so it can be removed from any remote databases that have the row. Two ways to do this are by using logical deletes or shadow tables.

- **Logical deletes** With this method, the row is not deleted. Data that is no longer required is marked as inactive in a status column. The WHERE clause of the `download_cursor` and `download_delete_cursor` can refer to the status of the row.

This technique is used in the `ULEmpCust` table in the `CustDB` sample application, in which the action column holds a D for Delete. The scripts use this value to delete the record from the remote database, and delete the record from the consolidated database at the end of the synchronization. `CustDB` also uses this technique for the `ULOrder` table, and the `Contact` sample uses the technique on the `Customer`, `Contact`, and `Product` tables.

The `MobiLink` synchronization model support for logical deletes assumes that a logical delete column is only on the consolidated database and not on the remote. When copying a consolidated schema to a new remote schema, leave out any columns that match the logical delete column in the model's synchronization settings. For a new model, the default column name is deleted.

To add the logical delete column name to the remote schema:

1. In the **Create Synchronization Model Wizard**, on the **Download Deletes** page, choose **Use logical deletes**.
2. Rename the logical delete column so that it does not match any column names in the consolidated.
3. When the wizard is finished, update the remote schema and keep the default table selection. The logical delete column name appears in the schema change list and be added to remote schema.

### Note

You need to set the column mapping for the remote's logical delete column to the consolidated's logical delete column.

- **Shadow tables** With this method, you create a shadow table that stores the primary key values of deleted rows. When a row is deleted, a trigger can populate the shadow table. The `download_delete_cursor` can use the shadow table to remove rows from remote databases. The shadow table only needs to contain the primary key columns from the real table.

See [“Writing `download\_delete\_cursor` scripts” on page 335](#).

## Temporarily stopping the synchronization of deletes

Ordinarily, SQL Anywhere automatically logs any changes to tables or columns that are part of a publication with a synchronization subscription. These changes are uploaded to the consolidated database during the next synchronization.



You may, however, want to delete rows from synchronized data and not have those changes uploaded. You can do this using `STOP SYNCHRONIZATION DELETE`. This feature can be used to make unusual corrections, but should be used with caution as it effectively disables part of the automatic synchronization functionality. This technique is a practical alternative to deleting the necessary rows using a `download_delete_cursor` script

When a `STOP SYNCHRONIZATION DELETE` statement is executed, none of the delete operations subsequently executed on that connection are synchronized. The effect continues until a `START SYNCHRONIZATION DELETE` statement is executed. The effects do not nest; that is, subsequent executions of `STOP SYNCHRONIZATION DELETE` after the first have no additional effect.

### **To temporarily disable upload of deletes made through a connection**

1. Issue the following statement to stop automatic logging of deletes.

```
STOP SYNCHRONIZATION DELETE
```

2. Delete rows from the synchronized data, as required, using the `DELETE` statement. Commit these changes.
3. Restart logging of deletes using the following statement.

```
START SYNCHRONIZATION DELETE
```

The deleted rows are not sent up to the MobiLink server and are not deleted from the consolidated database.

### **See also**

- SQL Anywhere clients: “[STOP SYNCHRONIZATION DELETE statement \[MobiLink\]](#)” [*SQL Anywhere Server - SQL Reference*]
- UltraLite clients: “[UltraLite STOP SYNCHRONIZATION DELETE statement](#)” [*UltraLite - Database Management and Reference*]

## Handling failed downloads

### Using blocking download acknowledgement

Blocking download acknowledgement has been deprecated. Use non-blocking download acknowledgement whenever possible.

Bookkeeping information about what is downloaded must be maintained in the download transaction. This information is updated atomically with the download being applied to the remote database.

If a failure occurs before the entire download is applied to the remote database, and if you change `SendDownloadAck` to ON, then the MobiLink server does not get confirmation for the download and rolls back the download transaction. Since the bookkeeping information is part of the download transaction, it is also rolled back. Next time the download is built, it uses the original bookkeeping information.

See “[SendDownloadACK \(sa\) extended option](#)” [*MobiLink - Client Administration*] and “[Send Download Acknowledgement synchronization parameter](#)” [*UltraLite - Database Management and Reference*].

When testing your synchronization scripts, you should add logic to your `end_download` script that causes occasional failures. This ensures that your scripts can handle a failed download.

### Using non-blocking download acknowledgement

Bookkeeping information about what is downloaded must be maintained in the nonblocking downloaded acknowledgement transaction. This information should be updated in the `publication_nonblocking_download_ack` or `nonblocking_download_ack` scripts which is called after the remote database successfully applies the download.

If a failure occurs or `SendDownloadAck` is OFF, these non-blocking download acknowledgement scripts are not called and the download timestamp is not updated. When testing your synchronization scripts you should add logic to your `publication_nonblocking_download_ack` or `nonblocking_download_ack` script that causes occasional failures. This ensures that your scripts can handle a failed download.

## Resuming failed downloads

Download failure is caused by a communication error during the download or a user cancelling the download. The MobiLink server holds download data that has not been received by the client for use in a restartable download. You can reduce the probability of download failure by decreasing the maximum amount of data allocated for restartable downloads using the `-ds` option. The server does not release download data until one of the following occurs:

- The user successfully completes the download.
- The user comes back with a new sync request without resume enabled.
- The cache is needed for incoming requests. The oldest unsuccessful download is cleared first.

MobiLink has functionality that can assist with download failure recovery, and prevent retransmission of the entire download. This functionality has separate implementations for SQL Anywhere and UltraLite remote databases. See “[-ds option](#)” on page 62.

## SQL Anywhere remote databases

When synchronization fails during a download, the downloaded data is not applied to the remote database, however, the successfully transmitted download portions are stored in a temporary file on the remote device. Dbmlsync uses this file to avoid lengthy retransmission of data, and to recover from download failure.

There are three ways to implement this functionality. In all cases, dbmlsync aborts and the resumed download fails if there is any new data to be uploaded.

- **-dc** After a download fails, use `-dc` the next time you start dbmlsync to resume the download. If part of the failed download was transmitted, the MobiLink server only transmits the remainder of the download.

For more information, see “[-dc option](#)” [*MobiLink - Client Administration*].

- **ContinueDownload (cd) extended option** When used on the dbmlsync command line, the `cd` extended option works just like the `-dc` option. You can also store this option in the database, or use `sp_hook_dbmlsync_set_extended_options` to set this option in a single synchronization.

See “[ContinueDownload \(cd\) extended option](#)” [*MobiLink - Client Administration*] and “[sp\\_hook\\_dbmlsync\\_set\\_extended\\_options](#)” [*MobiLink - Client Administration*].

- **sp\_hook\_dbmlsync\_end hook** You can use the restart parameter to cause a download to resume. You know a download is resumable if the restartable download parameter is set to true. You can also create logic in the hook to resume a download if a download file exists and is a certain size.

See “[sp\\_hook\\_dbmlsync\\_end](#)” [*MobiLink - Client Administration*].

## UltraLite remote databases

You can control the behavior of UltraLite applications following a failed download as follows:

- If you set the Keep Partial Download synchronization parameter to true when you synchronize, and the download fails before completion, then UltraLite applies that portion of the changes that were downloaded. UltraLite also sets the Partial Download Retained synchronization parameter to true.

The UltraLite database may be in an inconsistent state at this point. Depending on your application, you may want to ensure that synchronization completes successfully or is rolled back before you allow changes to the data. See “[Keep Partial Download synchronization parameter](#)” [*UltraLite - Database Management and Reference*], and “[Partial Download Retained synchronization parameter](#)” [*UltraLite - Database Management and Reference*].

- To resume the download, set the Resume Partial Download synchronization parameter to true and synchronize again. See “[Resume Partial Download synchronization parameter](#)” [*UltraLite - Database Management and Reference*].

The restarted synchronization does not perform an upload, and downloads only those changes that would have been downloaded by the failed download. That is, it completes the failed download but does not synchronize changes made since the previous attempt. To get those changes, you need to synchronize again once the failed download has completed, or call Rollback Partial Download and synchronize with Resume Partial Download set to false.

When you restart the download, many of the synchronization parameters from the failed synchronization are used again automatically. For example, the publications parameter is ignored: the synchronization

downloads those publications requested on the initial download. The only parameters that must be set are the Resume Partial Download parameter (which must be set to true) and the User Name parameter. In addition, settings for the following parameters are obeyed, if set:

- Keep Partial Download (in case of further interruption)
- DisableConcurrency
- Observer
- User Data
- To roll back the changes from the failed download without resuming synchronization, call the function to roll back the changes. This function is `ULRollbackPartialDownload` function for embedded SQL. For UltraLite components, it is a method on the Connection object.
  - **UltraLite.NET** See “[RollbackPartialDownload method](#)” [*UltraLite - .NET Programming*].
  - **Embedded SQL** See “[ULRollbackPartialDownload function](#)” [*UltraLite - C and C++ Programming*].

You may want to roll back the changes from a failed download if synchronization cannot be completed, for example if the server or network is unavailable, and if you want to maintain a consistent set of data while letting the end user continue to work.

For more information about communications errors, see [Error Messages](#).

**Note**

If the `send_download_ack` synchronization parameter is set to true, the setting is ignored for the resumed download.

## Download acknowledgement

Download acknowledgement is not required to ensure that your data is received at the remote.

There are two modes of download acknowledgement: blocking, which has been deprecated, and non-blocking. Non-blocking is the default. There is a performance penalty for using blocking download acknowledgement: the non-blocking mode is recommended.

To use download acknowledgements, there are settings on both the client and server.

On the client, you specify download acknowledgement with the dbmlsync extended option `SendDownloadACK` or the UltraLite synchronization parameter `Send Download Acknowledgement`. If you do not change any settings on the server, the default is non-blocking download acknowledgement.

On the server, you can set the `mlsrv11 -nba` option to specify blocking download acknowledgement. There are two connection events that you can use to record the last successful download time in your consolidated database when using non-blocking download acknowledgement.

### See also

- [“publication\\_nonblocking\\_download\\_ack connection event” on page 475](#)
- [“nonblocking\\_download\\_ack connection event” on page 471](#)
- [“-nba option” on page 74](#)
- dbmlsync: [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Download Acknowledgement synchronization parameter” \[UltraLite - Database Management and Reference\]](#)

## Downloading a result set from a stored procedure call

You can download a result set from a stored procedure call. For example, you might currently have a download\_cursor for the following table:

```
CREATE TABLE MyTable (  
    pk INTEGER PRIMARY KEY NOT NULL,  
    col1 VARCHAR(100) NOT NULL,  
    col2 VARCHAR(20) NOT NULL  
)
```

The download\_cursor table script might look as follows:

```
SELECT pk, col1, col2  
FROM MyTable  
WHERE last_modified >= {ml s.last_table_download}  
AND employee = {ml s.username}
```

If you want your downloads to MyTable to use more sophisticated business logic, you can now create your script as follows, where DownloadMyTable is a stored procedure taking two parameters (last-download timestamp and MobiLink user name) and returning a result set. (This example uses an ODBC calling convention for portability):

```
{call DownloadMyTable( {ml s.last_table_download}, {ml s.username} )}
```

The following are some simple examples for each supported consolidated database. Consult the documentation for your consolidated database for full details.

The following example works with SQL Anywhere, Adaptive Server Enterprise, and Microsoft SQL Server.

```
CREATE PROCEDURE SPDownload  
    @last_dl_ts DATETIME,  
    @u_name VARCHAR( 128 )  
AS  
BEGIN  
    SELECT pk, col1, col2  
    FROM MyTable  
    WHERE last_modified >= @last_dl_ts  
    AND employee = @u_name  
END
```

The following example works with Oracle. Oracle requires that a package be defined. This package must contain a record type for the result set, and a cursor type that returns the record type.

```
Create or replace package SPInfo as  
Type SPRec is record (  
    pk      integer,  
    col1    varchar(100),  
    col2    varchar(20)  
);  
Type SPCursor is ref cursor return SPRec;  
End SPInfo;
```

Next, Oracle requires a stored procedure with the cursor type as the first parameter. Note that the download\_cursor script only passes in two parameters, not three. For stored procedures returning result sets in Oracle, cursor types declared as parameters in the stored procedure definition define the structure of the result set, but do not define a true parameter as such. In this example, the stored procedure also adds the script to the MobiLink system table.

```

Create or replace procedure
  DownloadMyTable( v_spcursor IN OUT SPInfo.SPCursor,
                  v_last_dl_ts IN DATE,
                  v_user_name IN VARCHAR ) As
Begin
  Open v_spcursor For
    select pk, coll, col2
      from MyTable
      where last_modified >= v_last_dl_ts
      and employee = v_user_name;
End;

CALL ml_add_table_script(
  'v1',
  'MyTable',
  'download_cursor',
  '{CALL DownloadMyTable(
    {ml s.last_table_download},{ml s.username} )}'
);

```

The following example works with IBM DB2 LUW.

```

CREATE PROCEDURE DownloadMyTable(
  IN last_dl_ts TIMESTAMP,
  IN u_name VARCHAR( 128 ) )
  LANGUAGE SQL
  MODIFIES SQL DATA
  COMMIT ON RETURN NO
  DYNAMIC RESULT SETS 1

  BEGIN
    DECLARE C1, cursor WITH RETURN FOR
      SELECT pk, coll, col2 FROM MyTable
      WHERE last_modified >= last_dl_ts AND employee = u_name;
    OPEN C1;
  END;

```

The following example works with IBM DB2 mainframe.

```

CREATE PROCEDURE DownloadMyTable(
  IN last_dl_ts TIMESTAMP,
  IN u_name VARCHAR( 128 ) )
  LANGUAGE SQL
  MODIFIES SQL DATA
  EXTERNAL NAME MYDMT
  WLM ENVIRONMENT MYWLM
  COMMIT ON RETURN NO
  DYNAMIC RESULT SETS 1
  BEGIN
    DECLARE C1, cursor WITH RETURN FOR
      SELECT pk, coll, col2 FROM MyTable
      WHERE last_modified >= last_dl_ts AND employee = u_name;
    OPEN C1;
  END;

```

## Uploading data from self-referencing tables

Some tables are self-referencing. For example, an employee table may contain a column that lists employees and a column that lists the manager of each employee, and there may be a hierarchy of managers managing managers. These tables can pose a challenge to synchronization because the MobiLink default behavior is to coalesce all data updates on the remote database, which is efficient but which loses the order of transactions.

There are two techniques for handling this situation:

- If you are using a SQL Anywhere remote database, you can use the `dbmlsync -tu` option to specify that each transaction on the remote should be sent as a separate transaction.  
See “[-tu option](#)” [*MobiLink - Client Administration*].
- Add a mapping table so the order of transactions doesn't matter.



## MobiLink isolation levels

MobiLink connects to a consolidated database at the most optimal isolation level it can, given the isolation levels enabled on the RDBMS. The default isolation levels are chosen to provide the best performance while ensuring data consistency.

In general, MobiLink uses the isolation level `SQL_TXN_READ_COMMITTED` for uploads, and if possible, it uses snapshot isolation for downloads (if that is not possible, it uses `SQL_TXN_READ_COMMITTED`). Snapshot isolation eliminates the problem of downloads being blocked until transactions are closed on the consolidated database.

Snapshot isolation can result in duplicate data being downloaded (if, for example, a long-running transaction causes the same snapshot to be used for a long time), but MobiLink clients automatically handle this, so the only penalty is transmission time and the processing effort at the remote.

Isolation level 0 (`READ UNCOMMITTED`) is generally unsuitable for synchronization and can lead to inconsistent data.

The isolation level is set immediately after a connection to the database occurs. Some other connection setup occurs, and then the transaction is committed. The `COMMIT` is required by most RDBMSs so that the isolation level (and perhaps other settings) can take effect.

### SQL Anywhere version 10

SQL Anywhere version 10 supports snapshot isolation. By default, MobiLink uses the `SQL_TXN_READ_COMMITTED` isolation level for uploads, and snapshot isolation for downloads.

MobiLink can only use snapshot isolation if you enable it in your SQL Anywhere consolidated database. If snapshot isolation is not enabled, MobiLink uses the default `SQL_TXN_READ_COMMITTED`.

Enabling a database to use snapshot isolation can affect performance because copies of all modified rows must be maintained, regardless of the number of transactions that use snapshot isolation. See [“Enabling snapshot isolation”](#) [*SQL Anywhere Server - SQL Usage*].

You can enable snapshot isolation for upload with the `mlsrv11 -esu` option, and disable snapshot isolation with the `mlsrv11 -dsd` option. If you need to change the MobiLink default isolation level in a connection script, you should do so in the `begin_upload` or `begin_download` scripts. If you change the default isolation level in the `begin_connection` script, your setting may be overridden at the start of the upload and download transactions.

See [“-esu option” on page 66](#) and [“-dsd option” on page 63](#).

### SQL Anywhere prior to version 10

If you are using a version of SQL Anywhere prior to version 10, the default MobiLink isolation level is `SQL_TXN_READ_COMMITTED`. You can change the default for the entire MobiLink session in the `begin_connection` script, or change it for the upload and download in the `begin_upload` and `begin_download` scripts, respectively.

### Adaptive Server Enterprise

For Adaptive Server Enterprise, the default MobiLink isolation level is `SQL_TXN_READ_COMMITTED`. You can change the default for the entire MobiLink session in the `begin_connection` script, or change it for the upload and download in the `begin_upload` and `begin_download` scripts, respectively.

### Oracle

Oracle supports snapshot isolation, but calls it `READ COMMITTED`. By default, MobiLink uses the `snapshot/READ COMMITTED` isolation level for upload and download.

You can change the default for the entire MobiLink session in the `begin_connection` script, or change it for the upload and download in the `begin_upload` and `begin_download` scripts, respectively.

For the MobiLink server to be able to make the most effective use of snapshot isolation, the Oracle account used by the MobiLink server must have permission for the `V_$TRANSACTION` Oracle system view. If it does not, a warning is issued and rows may be missed on download. Only `SYS` can grant this access. The Oracle syntax for granting this access is:

```
grant select on SYS.V_$TRANSACTION to user-name
```

### Microsoft SQL Server 2005 and later

Microsoft SQL Server 2005 supports snapshot isolation. By default, MobiLink uses the `SQL_TXN_READ_COMMITTED` isolation level for uploads, and snapshot isolation for download.

MobiLink can only use snapshot isolation if you enable it in your SQL Server consolidated database. If snapshot is not enabled, MobiLink uses the default `SQL_TXN_READ_COMMITTED`. See your SQL Server documentation for details.

You can enable snapshot isolation for upload with the `mlsrv11 -esu` option, and disable snapshot isolation with the `mlsrv11 -dsd` option. If you need to change the MobiLink default isolation level in a connection script, you should do so in the `begin_upload` or `begin_download` scripts. If you change the default isolation level in the `begin_connection` script, your setting may be overridden at the start of the upload and download transactions.

See [“-esu option” on page 66](#) and [“-dsd option” on page 63](#).

To use snapshot isolation on SQL Server, the user ID that you use to connect the MobiLink server to the database must have permission to access the SQL Server system table `SYS.DM_TRAN_ACTIVE_TRANSACTIONS`. If this permission is not granted, MobiLink uses the default level `SQL_TXN_READ_COMMITTED`.

If your consolidated database is running on a Microsoft SQL Server that is also running other databases, if you are using snapshot isolation for uploads or downloads, and if your upload or download scripts do not access any other databases on the server, you should specify the MobiLink server `-dt` option. This option makes MobiLink ignore all transactions except ones within the current database. It increases throughput and reduces duplication of rows that are downloaded.

See [“-dt option” on page 64](#).

### Microsoft SQL Server prior to version 2005

If you are using a version of Microsoft SQL Server prior to version 2005, the default MobiLink isolation level is `SQL_TXN_READ_COMMITTED`. You can change the default for the entire MobiLink session in

the `begin_connection` script, or change it for the upload and download in the `begin_upload` and `begin_download` scripts, respectively.

**See also**

- [“-dsd option” on page 63](#)
- [“-esu option” on page 66](#)
- [“-dt option” on page 64](#)
- [“The synchronization process” \[\*MobiLink - Getting Started\*\]](#)
- [“Isolation levels and consistency” \[\*SQL Anywhere Server - SQL Usage\*\]](#)
- [“Snapshot isolation” \[\*SQL Anywhere Server - SQL Usage\*\]](#)

---

---

# MobiLink performance

## Contents

Performance tips .....	170
Key factors influencing MobiLink performance .....	174
Monitoring MobiLink performance .....	178

---

## Performance tips

The following is a list of suggestions to help you get the best performance out of MobiLink.

### Test

The following all contribute to the throughput of your synchronization system: the type of device running your remote databases, the schema of remote databases, the data volume and synchronization frequency of your remotes, network characteristics (including for HTTP, proxies, web servers, and Redirectors), the hardware where the MobiLink server runs, your synchronization scripts, the concurrent volume of synchronizations, the type of consolidated database you use, the hardware where your consolidated database runs, and the schema of your consolidated database.

Testing is extremely important. Before deploying, you should perform testing using the same hardware and network that you plan to use for production. You should also try to test with the same number of remotes, the same frequency of synchronization, and the same data volume. During this testing you should experiment with the following performance tips.

### Avoid contention

Avoid contention and maximize concurrency in your synchronization scripts.

For example, suppose a `begin_download` script increments a column in a table to count the total number of downloads. If multiple users synchronize at the same time, this script would effectively serialize their downloads. The same counter would be better in the `begin_synchronization`, `end_synchronization`, or `prepare_for_download` scripts because these scripts are called just before a commit so any database locks are held for only a short time.

See [“Contention” on page 174](#).

For information about the transaction structure of synchronization, see [“Transactions in the synchronization process” \[MobiLink - Getting Started\]](#).

### Use an optimal number of database worker threads

Use the MobiLink `-w` option to set the number of MobiLink database worker threads to the smallest number that gives you optimum throughput. You need to experiment to find the best number for your situation.

A larger number of database worker threads can improve throughput by allowing more synchronizations to access the consolidated database at the same time.

Keeping the number of database worker threads small reduces the chance of contention in the consolidated database, the number of connections to the consolidated database, and the memory required for optimal caching.

See:

- [“Number of database worker threads” on page 175](#)
- [“-w option” on page 105](#)
- [“-wu option” on page 106](#)

### **Enable the client-side download buffer for SQL Anywhere clients**

For SQL Anywhere clients, a download buffer allows a MobiLink database worker thread to transmit the download without waiting for the client to apply the download. The download buffer is enabled by default. However, the download buffer cannot be used if download acknowledgement is enabled (see next bullet).

See [“DownloadBufferSize \(dbs\) extended option”](#) [*MobiLink - Client Administration*].

### **Use non-blocking download acknowledgement**

By default, download acknowledgement is not enabled. However, download acknowledgement can be useful for recording remote progress in the consolidated database. Blocking download acknowledgement ties up a database connection while the remote is applying the download. If you use download acknowledgement, you should use non-blocking download acknowledgement. Non-blocking download acknowledgement is the default.

See [“SendDownloadACK \(sa\) extended option”](#) [*MobiLink - Client Administration*] and [“-nba option”](#) on page 74.

### **Avoid synchronizing unnecessary BLOBs**

It is inefficient to include a BLOB in a row that is synchronized frequently. To avoid this, you can create a table that contains BLOBs and a BLOB ID, and reference the ID in the table that needs to be synchronized.

### **Set maximum number of database connections**

Set the maximum number of MobiLink database connections to be your number of synchronization script versions times the number of MobiLink database worker threads, plus one. This reduces the need for MobiLink to close and create database connections. You set the maximum number of connections with the `mlsrv11 -cn` option.

See [“MobiLink database connections”](#) on page 176 and [“-cn option”](#) on page 56.

### **Have enough physical memory**

Ensure that the computer running the MobiLink server has enough physical memory to accommodate the cache in addition to its other memory requirements.

The number of synchronizations being actively processed is not limited by the number of database worker threads. The MobiLink server can unpack uploads and send downloads for a large number of synchronizations simultaneously. For best performance, it is very important that the MobiLink server has a large enough memory cache to process these synchronizations without paging to disk. Use the `-cm` option to set the memory cache for the MobiLink server.

See [“-cm option”](#) on page 55.

### **Use enough processing power**

You should dedicate enough processing power to MobiLink so that the MobiLink server processing is not a bottleneck. Typically the MobiLink server requires significantly less CPU than the consolidated database. However, using Java or .NET row handling adds to the MobiLink server processing requirement. In practice, network limitations or database contention are more likely to be bottlenecks.

### Use minimum logging verbosity

Use the minimum logging verbosity that is compatible with your business needs. By default, verbose logging is off, and MobiLink does not write its log to disk. You can control logging verbosity with the `-v` option, and enable logging to a file with the `-o` or `-ot` options.

As an alternative to verbose log files, you can monitor your synchronizations with the MobiLink Monitor. The MobiLink Monitor does not need to be on the same computer as the MobiLink server, and a Monitor connection has a negligible effect on MobiLink server performance. See [“MobiLink Monitor” on page 179](#).

### Plan for operating system limitations

Operating systems restrict the number of concurrent connections a server can support over TCP/IP. If this limit is reached, which may occur when over 1000 clients attempt to synchronize at the same time, the operating system may exhibit unexpected behavior, such as unexpectedly closing connections and rejecting additional clients that attempt to connect. To prevent this behavior, use the `-sm` option to specify a maximum number of remote connections that is less than the operating system limit.

When a client attempts to synchronize with a MobiLink server that has accepted its maximum number of concurrent synchronizations as specified by the `-sm` option, the client receives the error code -85 (SQLE\_COMMUNICATIONS\_ERROR). The client application should handle this error and try to connect again in a few minutes.

For more information about the `-sm` option, see [“-sm option” on page 94](#).

For more information about SQLE\_COMMUNICATIONS\_ERROR, see [“Communication error” \[Error Messages\]](#).

### Java or .NET vs. SQL synchronization logic

No significant throughput difference has been found between using Java or .NET synchronization logic vs. SQL synchronization logic. However, Java and .NET synchronization logic have some extra overhead per synchronization and require more memory.

In addition, SQL synchronization logic is executed on the computer that runs the consolidated database, while Java or .NET synchronization logic is executed on the computer that runs the MobiLink server. So, Java or .NET synchronization logic may be desirable if your consolidated database is heavily loaded.

Synchronization using direct row handling imposes a heavier processing burden on the MobiLink server, so you may need more RAM, perhaps more disk space, and perhaps more CPU power, depending on how you implement direct row handling.

### Priority synchronization

If you have some tables that you need to synchronize more frequently than others, create a separate publication and subscription for them. When using synchronization models in Sybase Central, you can do this by creating more than one model. You can synchronize this priority publication more frequently than other publications, and synchronize other publications at off-peak times.



**Download only the rows you need**

Take care to download only the rows that are required, for example by using timestamp synchronization instead of snapshot. Downloading unneeded rows is wasteful and adversely affects synchronization performance.

**Optimize script execution**

The performance of your scripts in the consolidated database is an important factor. It may help to create indexes on your tables so that the upload and download cursor scripts can efficiently locate the required rows. However, too many indexes may slow uploads.

When you use the **Create Synchronization Model Wizard** in Sybase Central to create your MobiLink applications, an index is automatically defined for each download cursor when you deploy the model.

**For large uploads, estimate the number of rows**

For SQL Anywhere clients, you can significantly improve the speed of uploading a large number of rows by providing dbmsync with an estimate of the number of rows that are uploaded. You do this with the dbmsync -urc option.

See “-urc option” [*MobiLink - Client Administration*].

## Key factors influencing MobiLink performance

The overall performance of any system, including throughput for MobiLink synchronization, is usually limited by a bottleneck at one point in the system. For MobiLink synchronization, the following might be the bottlenecks limiting synchronization throughput:

- **The performance of the consolidated database** Of particular importance for MobiLink is the speed at which the consolidated database can execute the MobiLink scripts. Multiple database worker threads might execute scripts simultaneously, so for best throughput you need to avoid database contention in your synchronization scripts.
- **The number of MobiLink database worker threads** A smaller number of threads involve fewer database connections, less chance of contention in the consolidated database and less operating system overhead. However, too small a number may leave clients waiting for a free database worker thread, or have fewer connections to the consolidated database than it can overlap efficiently.
- **The bandwidth for client-to-MobiLink communications** For slow connections, such as those over dial-up or wide-area wireless networks, the network may cause clients and MobiLink servers to wait for data to be transferred.
- **The client processing speed** Slow client processing speed is more likely to be a bottleneck in downloads than uploads, since downloads involve more client processing as rows and indexes are written.
- **The bandwidth for MobiLink to consolidated communication** This is unlikely to be a bottleneck if both MobiLink and the consolidated database are running on the same computer, or if they are on separate computers connected by a high-speed network.
- **The speed of the computer running the MobiLink server** If the processing power of the computer running MobiLink is slow, or if it does not have enough memory for the MobiLink database worker threads and buffers, then MobiLink execution speed could be a synchronization bottleneck. The MobiLink server's performance depends little on disk speed as long as the buffers and database worker threads fit in physical memory.

## Tuning MobiLink for performance

The key to achieving optimal MobiLink synchronization throughput is to have multiple synchronizations occurring simultaneously and executing efficiently. To enable multiple simultaneous synchronizations, MobiLink uses pools of database worker threads for different tasks. One pool is dedicated to reading upload data from the network and unpacking it. Another pool of threads, called **database worker threads**, applies the upload to the consolidated database and fetches data to be downloaded from the consolidated database. Another pool of database worker threads is dedicated to packing and sending the download data to the remote databases. Each database worker thread uses a single connection to the consolidated database for applying and fetching changes, using your synchronization scripts.

### Contention

The most important factor is to avoid database contention in your synchronization scripts. Just as with any other multi-client use of a database, you want to minimize database contention when clients are simultaneously accessing a database. Database rows that must be modified by each synchronization can

increase contention. For example, if your scripts increment a counter in a row, then updating that counter can be a bottleneck.

Synchronization requests are accepted (up to the limit specified by the `-sm` option) and the uploaded data is read and unpacked so that it is ready for a database worker thread. If there are more synchronizations than database worker threads, the excess are queued, waiting for a free database worker thread.

You can control the number of database worker threads and connections, but MobiLink always ensures that there is at least one connection per database worker thread. If there are more connections than database worker threads, the excess connections are idle. Excess connections may be useful with multiple script versions, as discussed below.

### Number of database worker threads

Other than contention in your synchronization scripts, the most important factor for synchronization throughput is the number of database worker threads. The number of database worker threads controls how many synchronizations can proceed simultaneously in the consolidated database.

Testing is vital to determine the optimum number of database worker threads.

Increasing the number of database worker threads allows more overlapping synchronizations to access the consolidated database, and increases throughput. However, it also increases resource and database contention between the overlapping synchronizations, and potentially increases the time for individual synchronizations. As the number of database worker threads is increased, the benefit of more simultaneous synchronizations becomes outweighed by the cost of longer individual synchronizations, and adding more database worker threads decreases throughput. Experimentation is required to determine the optimal number of database worker threads for your situation, but the following may help to guide you.

For uploads, performance testing shows that the best throughput happens with a relatively small number of database worker threads: in most cases, three to ten database worker threads. Variation depends on factors like the type of consolidated database, data volume, database schema, the complexity of the synchronization scripts, and the hardware used. The bottleneck is usually due to contention between database worker threads executing the SQL of your upload scripts at the same time in the consolidated database.

For downloads, when blocking download acknowledgement is used the optimum number of database worker threads depends on the client-to-MobiLink bandwidth and the processing speed of clients. For slower clients, more database worker threads are needed to get optimal download performance. This is because downloads involve more client processing and less consolidated database processing than uploads. When blocking download acknowledgement is used, database worker threads block until the remote database finishes applying the download. For non-blocking download acknowledgement, more workers are not needed.

When download acknowledgements are not used (the default), the client-to-MobiLink bandwidth is less influential because a database worker thread is free to process other synchronizations while other threads send the download. So, the number of database worker threads is less critical.

Many downloads can be sent concurrently—far more than the number of database worker threads. For optimal download performance, it is important for the MobiLink server to have enough RAM to buffer these downloads. Otherwise the download is paged to disk and download performance may degrade. To specify the MobiLink server memory cache size, use the `-cm` option.

See [“-cm option” on page 55](#).

If the MobiLink server starts paging to disk (possibly because of too many downloads being processed concurrently), consider using the `-sm` option to either decrease the number of database worker threads or limit the total number of synchronizations being actively processed.

See [“-sm option” on page 94](#).

Leaving download acknowledgement off (the default) can reduce the optimal number of database worker threads for download, because database worker threads do not have to wait for clients to apply downloads.

See [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#).

If you use download acknowledgement, performance is better with non-blocking download acknowledgement (as opposed to blocking). In non-blocking acknowledgement mode the server reuses the database worker thread while the remote database applies the download. This means that the number of database worker threads may not need to be increased, which results in better performance.

To get both the best download throughput and the best upload throughput, MobiLink provides two options. You can specify a total number of database worker threads to optimize downloads. You can also limit the number that can simultaneously apply uploads to optimize upload throughput.

The `-w` option controls the total number of database worker threads. The default is five.

The `-wu` option limits the number of database worker threads that can simultaneously apply uploads to the consolidated database. By default, all database worker threads can apply uploads simultaneously, but that can cause severe contention in the consolidated database. The `-wu` option lets you reduce that contention while still having a larger number of database worker threads to optimize the fetching of download data. The `-wu` option only has an effect if the number is less than the total number of database worker threads.

See [“-w option” on page 105](#) and [“-wu option” on page 106](#).

### MobiLink database connections

MobiLink creates a database connection for each database worker thread. You can use the `-cn` option to specify that MobiLink create a larger pool of database connections, but any excess connections are idle unless MobiLink needs to close a connection or use a different script version.

There are two cases where MobiLink closes a database connection and open a new one. The first case is if an error occurs. The second case is if the client requests a synchronization script version, and none of the available connections have already used that synchronization version.

#### Note

Each database connection is associated with a script version. To change the version, the connection must be closed and reopened.

If you routinely use more than one script version, you can reduce the need for MobiLink to close and open connections by increasing the number of connections. You can eliminate the need completely if the number of connections used for synchronizations is the number of database worker threads times the number of script versions.

An example of tuning MobiLink for two script versions is given in the following command line:

```
mllsrv11 -c "dsn=SQL Anywhere 11 Demo" -w 5 -cn 10
```

Since the maximum number of database connections used for synchronizations is the number of script versions times the number of database worker threads, setting `-cn` to 10 ensures that database connections are not closed and opened excessively.

See [“-cn option” on page 56](#).

## Monitoring MobiLink performance

There are a variety of tools available to help you monitor the performance of your synchronizations.

The MobiLink Monitor is a graphical tool for monitoring synchronizations. It allows you to see the time taken by every aspect of the synchronization.

See [“MobiLink Monitor” on page 179](#).

In addition, there are several MobiLink scripts that are available for monitoring synchronizations. These scripts allow you to use performance statistics in your business logic. You may, for example, want to store the performance information for future analysis, or alert a DBA if a synchronization takes too long. For more information, see:

- [“download\\_statistics connection event” on page 403](#)
- [“download\\_statistics table event” on page 406](#)
- [“synchronization\\_statistics connection event” on page 486](#)
- [“synchronization\\_statistics table event” on page 489](#)
- [“time\\_statistics connection event” on page 492](#)
- [“time\\_statistics table event” on page 495](#)
- [“upload\\_statistics connection event” on page 512](#)
- [“upload\\_statistics table event” on page 517](#)

---

# MobiLink Monitor

## Contents

Introduction to the MobiLink Monitor .....	180
Starting the MobiLink Monitor .....	181
Using the MobiLink Monitor .....	184
Saving MobiLink Monitor data .....	193
Customizing your statistics .....	194
MobiLink statistical properties .....	195

---

## Introduction to the MobiLink Monitor

The MobiLink Monitor is a MobiLink administration tool that provides you with details about the performance of your synchronizations.

When you start the MobiLink Monitor and connect it to a MobiLink server, the MobiLink Monitor begins to collect statistical information about all synchronizations that occur in that monitoring session. The MobiLink Monitor continues to collect data until you disconnect it or shut down the MobiLink server.

You can view the data in tabular or graphical form in the MobiLink Monitor interface. You can also save the data in binary format for viewing with the MobiLink Monitor later, or in *.csv* format to open in another tool, such as Microsoft Excel; or you can export it to an ODBC data source such as a MobiLink-supported relational database.

MobiLink Monitor output allows you to see a wide variety of information about your synchronizations. For example, you can quickly identify synchronizations that result in errors, or that meet other criteria that you specify. You can identify possible contention in synchronization scripts by checking to see if synchronizations of differing durations have phases that end around the same time (because synchronizations are waiting for a previous phase to finish before they can continue).

The MobiLink Monitor can be used routinely in development and production, because monitoring does not degrade performance, particularly when the MobiLink Monitor is run on a different computer from the MobiLink server.

### SQL Anywhere Monitor

The SQL Anywhere Monitor is a browser-based administration tool that provides you with information about the health and availability of SQL Anywhere databases and MobiLink servers. It is useful in assessing overall system health and availability, and for analyzing overall synchronization statistics. The SQL Anywhere Monitor does not provide information about individual synchronizations. For detailed information about individual synchronizations, including timing and other per-synchronization statistics, use the MobiLink Monitor.

For more information about the SQL Anywhere Monitor, see [“SQL Anywhere Monitor” \[SQL Anywhere Server - Database Administration\]](#).



## Starting the MobiLink Monitor

You can have multiple instances of the MobiLink Monitor running for each MobiLink server.

### To start monitoring data

1. Start your consolidated database and MobiLink server, if they are not already running.
2. From the **Start** menu, choose **Programs » SQL Anywhere 11 » MobiLink Monitor**.  
Alternatively, you can type **mlmon** at a command prompt. For details, see below.
3. A MobiLink Monitor connection starts like a synchronization connection to the MobiLink server. For example, if you started the MobiLink server with **-zu+** then it doesn't matter what user ID you use here. For all MobiLink Monitor sessions, the script version is set to **for\_ML\_Monitor\_only**.

The **Connect To MobiLink Server** window should be completed as follows:

- **Host** The network name or IP address of the computer where the MobiLink server is running. By default, it is the computer where the MobiLink Monitor is running. You can use **localhost** if the MobiLink server is running on the same computer as the MobiLink Monitor.
- **Protocol** This should be set to the same network protocol and port that the MobiLink server is using for synchronization requests.
- **Port** This should be set to the same network port that the MobiLink server is using for synchronization requests.
- **Encryption** If you chose HTTPS or TLS for the protocol, this box is enabled. You can choose an encryption type from the dropdown list.
- **Additional Protocol Options** Specify optional parameters. You can set the following parameters, separated by semicolon if you need to specify multiple parameters:
  - **buffer\_size=number**
  - **client\_port=nnnn**
  - **client\_port=nnnn-mmmmm**
  - **persistent={0|1}** (HTTP and HTTPS only)
  - **proxy\_host=proxy\_hostname** (HTTP and HTTPS only)
  - **proxy\_port=proxy\_portnumber** (HTTP and HTTPS only)
  - **url\_suffix=suffix** (HTTP and HTTPS only)
  - **version=HTTP-version-number** (HTTP and HTTPS only)

See “[MobiLink client network protocol options](#)” [*MobiLink - Client Administration*].

4. Start synchronizing.

The data appears in the MobiLink Monitor as it is collected.

## Starting mlmon on the command line

Command line options allow you to have the MobiLink Monitor open a file or connect to a MobiLink server on startup. Use the following syntax:

```
mlmon [ connect-options | inputfile.{ mlm | csv } ]
```

where:

**connect-options** can be one or more of the following:

- **-u** *ml\_username* (required to connect to the MobiLink server)
- **-p** *password*
- **-x** { **tcpip** | **tls** | **http** | **https** } [ ( *keyword=value*;... ) ]

The *keyword=value* pairs can be the host, protocol, and Additional Network Parameters as described above. The **-x** option is required to connect to the MobiLink server.

- **-o** *outputfile*.{ **mlm** | **csv** }

The **-o** option closes the MobiLink Monitor at the end of the connection and saves the session in the specified file.

You can type **mlmon -?** to view the mlmon syntax.

## Starting the MobiLink Monitor on Unix

The following steps can be used if you are using a version of Linux that supports the Linux Desktop icons and if you chose to install them when you installed SQL Anywhere 11.

### To start the MobiLink Monitor (Linux Desktop icons)

1. From the **Applications** menu, choose **SQL Anywhere 11 » MobiLink Monitor**
2. Enter the information to connect to the MobiLink server described in the procedure To start monitoring data.

#### Note

The following steps assume that you have already sourced the SQL Anywhere utilities. See “[Setting environment variables on Unix and Mac OS X](#)” [*SQL Anywhere Server - Database Administration*].

### To start the MobiLink Monitor (Unix command line)

1. In a terminal session, enter the following command:

```
mlmon
```

2. Enter the information to connect to the MobiLink server as described above.

## Stopping the MobiLink Monitor

### To stop the MobiLink Monitor

1. In the MobiLink Monitor, choose **Monitor » Disconnect From MobiLink Server**. This stops the collection of data.

You can also stop collecting data by shutting down the MobiLink server or closing the MobiLink Monitor.

Before closing the MobiLink Monitor, you can save the data for the session. See [“Saving MobiLink Monitor data” on page 193](#).

2. When you are ready to close the MobiLink Monitor, choose **File » Close**.

## Using the MobiLink Monitor

The MobiLink Monitor has the following panes:

- **Details Table** **Details Table** is the top pane. It is a spreadsheet that shows the total time taken by each synchronization, with a breakdown showing the amount of time taken by each part of the synchronization.  
See [“Details Table pane” on page 184](#).
- **Utilization Graph** **Utilization Graph** is the second pane. It provides a graphical representation of queue lengths for different queues on the MobiLink server. The same scale is used for the **Utilization Graph** pane and **Chart** pane. The scale at the bottom of the **Chart** pane represents time. You can select the data that is displayed in the utilization graph by dragging and selecting the data in the **Overview** pane below, or by choosing **View » Go To**.  
See [“Utilization Graph pane” on page 186](#).
- **Chart** **Chart** is the third pane. It provides a graphical representation of synchronizations. The scale at the bottom of this pane represents time. You can select the data that is displayed in the chart by dragging and selecting the data in the **Overview** pane below, or by choosing **View » Go To**.  
See [“Chart pane” on page 188](#).
- **Overview** **Overview** is the bottom pane. It shows an overview of all synchronizations in the session. This pane contains a box outline called the **Marquee Tool** that can select the data appearing in the **Chart** and **Utilization Graph** panes.  
See [“Overview pane” on page 189](#).

In addition, there is an **Options** window that you can use to customize the display, and properties windows for viewing more details. See:

- [“Options window” on page 190](#)
- [“Session properties” on page 190](#)
- [“Sample properties” on page 190](#)
- [“Synchronization properties” on page 191](#)

### Details Table pane

The **Details Table** provides information about the duration of each part of the synchronization. All times are measured by the MobiLink server. Some times may be non-zero even when you do not have the corresponding script defined.

You can choose the columns that appear in the **Details Table** pane by opening **Tools » Options** and then opening the **Table** tab. For a description of the statistics that are available, see [“MobiLink statistical properties” on page 195](#).

The following columns appear by default:

- **sync** Identifies each synchronization. This ID is assigned by the MobiLink server, and not by the MobiLink Monitor, so it does not necessarily start at 1 in any given MobiLink Monitor session and is not received in numerical order. You can see the same IDs in the **Synchronization Properties** window. See [“Synchronization properties” on page 191](#).
- **remote\_id** The ID of the remote database.
- **user** The synchronization user.
- **version** The version of the synchronization script.  
See [“Script versions” on page 324](#).
- **download\_ack** The type of download ack, which can be none, blocking or non-blocking.
- **start\_time** The date and time when the MobiLink server started the synchronization. (This may be later than when the synchronization was requested by the client.)
- **duration** The total duration of the synchronization, in seconds.
- **sync\_request** The time in seconds for MobiLink to receive the uploaded data from the client. This is the portion of the synchronization starting at the connection from the remote and ending just before authentication.
- **receive\_upload** The time taken between creating the network connection between the remote database and the MobiLink server, up to receiving the first bytes of the upload stream. This time is insignificant unless you have set `-sm` to a smaller value than `-nc`, then this time can include the time that a synchronization is paused when the number of synchronizations is larger than the maximum number of active synchronizations that were specified with `-sm`.
- **get\_db\_worker** The time required to acquire a free database worker thread.
- **connect** The time required by the database worker thread to make a database connection if a new database connection is needed. For example, after an error or if the script version has changed.
- **authenticate\_user** The time in seconds for MobiLink to validate the synchronization request, validate the user name, and validate the password (if your synchronization setup requires authentication). This is the length of the authenticate user transaction (from the start of authentication to just before the `begin_synchronization` event).
- **begin\_sync** The time in seconds to run your `begin_synchronization` script, if one was run.
- **apply\_upload** The time in seconds to apply the upload to the consolidated database. This is the time between the `begin_upload` script and the `end_upload` script.
- **prepare\_for\_download** The time in seconds to run your `prepare_for_download` script, if one was run.
- **fetch\_download** The time in seconds to download the data. This is the time between the `begin_download` script and the `end_download` script. If download acknowledgement is enabled, this includes the time to apply the download on the remote database and return acknowledgement.
- **end\_sync** The time in seconds to run the `end_synchronization` script, if one was run.

To sort the table by a specific column, click the column heading. If new data is appearing in the MobiLink Monitor, it gets sorted as it is added.

You can close the **Details Table** pane by clearing **Details Table** option in the **View** menu.

## Utilization Graph pane

The **Utilization Graph** is the second pane from the top. It displays server statistics for several types of work queues.

For more information about the data available in this pane, see [“Using the Utilization Graph” on page 186](#).

The **Utilization Graph** uses the same horizontal scrollbar, horizontal time labels, and zoom level as the **Chart**. This means that an instant in time lines up vertically between the **Graph** pane and the **Chart** pane.

There are multiple ways to select the data that is displayed in the graph:

- From the **View** menu, choose **Go To**.
- In the **Overview** pane, move the **Marquee Tool**. The **Marquee Tool** is the small box that appears in the **Overview** pane. See [“Overview pane” on page 189](#).

You can double-click an area of the **Utilization Graph** to bring up a **Sample Properties** window that shows the details of the sample interval it represents. The sample interval is about a second long. See [“Sample properties” on page 190](#).

### Note

The list of properties is obtained from the MobiLink server or the *.mlm* file that was opened, and it uses the language of the MobiLink server. Some characters may not display properly if the MobiLink Monitor is using a different language.

## Using the Utilization Graph

To see the values for the **Utilization Graph**, and to customize the output, choose **Tools » Options** and open the **Graph** tab. This tab identifies the **Utilization Graph** queues by color, and allows you to customize the graph.

### Properties

- **TCP/IP Work Queue** This queue represents work done by the low-level network layer in the MobiLink server. This layer is responsible for both reading and writing packets from and to the network. The queue is full of read and write requests. It grows when it gets notified of incoming data to read off the network and/or when told by the stream worker to write to the network.

If this queue gets backed up, it is usually due to a backlog of either reads or writes—sometimes both but usually one or the other. Reads can get backed up if the server is using a lot of RAM and memory pages are being swapped in and out a lot. Consider getting more RAM. Writes can get backed up if the network connection between the clients and server is slow. If this queue is the only queue that is backed up, look at the CPU usage. If CPU usage is high, suspect read/memory problems. If CPU usage is low, you may have slow writes. Consider using a faster network.

- **Stream Work Queue** For version 10 clients only. This queue represents work done by the high-level network layer in the MobiLink server. This layer is responsible for higher-level network protocol work such as HTTP, encryption, and compression. This queue grows when lots of reads come in from the TCP/IP layer and/or when lots of write requests come in from the Command processor layer. If this queue is the only queue that is backed up, consider removing some network protocols, such as HTTP or compression. If this is not possible, consider reducing the number of concurrent synchronizations allowed using the `-sm` option.

See “[-sm option](#)” on page 94.

- **Heartbeat Work Queue** This queue represents the layer in MobiLink server that is responsible for sending pulsed events within the server. This layer is responsible, for example, for triggering the one-per-second pulses of samples to the connected MobiLink Monitors.

It is highly unlikely that this queue gets backed up so it is not visible by default.

- **Command Processor Work Queue** This represents work done by the MobiLink server to both interpret internal MobiLink protocol commands and apply these commands to the consolidated database. This queue grows when lots of requests come in. Request types include synchronization requests, Listener requests, `mlfiletransfer` requests, and so on. The queue also grows when the consolidated database is busy working on synchronizations, yet more synchronization requests keep coming in.

If this queue is the only queue that is backed up, look at the CPU usage. If CPU usage is high, the volume of requests may be too high. Consider reducing the number of concurrent synchronizations allowed, using the `-sm` option. If CPU usage is low, look to improving the performance of the consolidated database.

See “[-sm option](#)” on page 94.

- **Busy Database Worker Threads** This value indicates how hard the MobiLink server is pushing the consolidated database. Each unit in the value represents a database worker thread that is doing something in the database. There is no distinction between inserts, updates, deletes, or selects. When this value is zero, the server is not operating on the consolidated database.

When this count is high (when it is close to the maximum set with the `mlsrv11 -w` option), the MobiLink server is pushing the consolidated database as hard as it can. In this case, if your throughput is satisfactory, there is nothing to do. If your throughput is not satisfactory, consider increasing the number of database worker threads via the `-w` option. Note that a higher `-w` value leads to greater contention between connections. This is particularly bad when all connections are performing uploads, so you may need to use the `mlsrv11 -wu` option to set a lower limit for database workers doing uploads. If you cannot seem to find settings for `-w` and `-wu` that provide adequate throughput, examine your synchronization scripts for possible contention issues. Finally, you can consult your RDBMS documentation for ways to improve the overall performance of your consolidated database.

## Scale

This column tells you the current scale for each property.

The vertical scale on the **Utilization Graph** always goes from 0 to 100. This represents zero to one hundred percent of the scale. Each value has its own scale. By default, all scales are 5, meaning that values are expected to be in a range of 0 to 20, scaled (by 5) to the range 0 to 100. If a value becomes greater than 20, the scale is automatically adjusted such that the largest value is at 100.

To determine the maximum value in the display, divide the scale into 100. For example, if the TCP/IP work queue scale is 2.381, then the maximum value is  $(100 / 2.381) = 42$ . The actual maximum isn't usually important. What is important is that values towards the top of the graph are approaching the largest currently-known value for the given property—in other words, the peak load for that property as observed in the current monitoring session.

When the graphs are consistently towards the top of the display and you notice that synchronization throughput is down, you may have a performance problem that needs investigation. Similarly, if one or more values creeps upward over time without diminishing, then there is likely a performance problem. Note that the graphs may often be towards the top of the display with MobiLink server performing well. This just means that MobiLink server is busy and doing its job well.

## Antialiasing

One of your customization choices is antialiasing. Antialiasing makes the graph look better, but can be slower to draw.

## Chart pane

The **Chart** pane presents the same information as the **Details Table**, but in graphical format. The bars in the **Chart** represent the length of time taken by each synchronization, with sub-sections of the bars representing the phases of the synchronization.

### Viewing data

Click a synchronization to select that synchronization in the **Details Table**.

Double-click a synchronization to open the **Synchronization Properties** window. See [“Synchronization properties” on page 191](#).

### Grouping data by remote ID or compactly

You can group the data by user. Choose **View » By Remote ID**.

Alternatively, you can view the data in a compact mode that shows all active synchronizations in as few rows as possible. Choose **View » Compact View**. In **Compact View**, the row numbers are meaningless.

### Zooming in on data

There are several ways to select the data that is visible in the **Chart** pane:

- **Zoom options** There are zoom options in the **View** menu and zoom buttons on the toolbar that allow you to zoom in and out. To have a synchronization fill the available space, use **Zoom To Selection**.
- **Scrollbar** Click the scrollbar at the bottom of the **Chart** pane and slide it.
- **Go To window** To open this window, click **View » Go To**.

**Start Date & Time** lets you specify the start time for the data that appears in the **Chart** pane. If you change this setting, you must specify at least the year, month, and date of the date-time.



**Chart Range** lets you specify the duration of time that is displayed. The chart range can be specified in milliseconds, seconds, minutes, hours, or days. The chart range determines the granularity of the data: a smaller length of time means that more detail is visible.

- **Marquee Tool** In the **Overview** pane, move the **Marquee Tool**. The **Marquee Tool** is the small box that appears in the **Overview** pane. See [“Overview pane” on page 189](#).

### Time axis

At the bottom of the **Chart** pane there is a scale showing time periods. The format of the time is readjusted automatically depending on the span of time that is displayed. You can always see the complete date-time by hovering your cursor over the scale.

### Default color scheme

You can view or set the colors in the **Chart** pane by opening the **Options** window (available from the **Tools** menu). The default color scheme for the **Chart** pane uses lime green for uploads, coral red for downloads, and blue for begin and end phases, with a darker shade for earlier parts of a phase.

For information about setting colors, see [“Options window” on page 190](#).

## Overview pane

The **Overview** pane shows an overview of the entire MobiLink Monitor session. You can navigate through the session using the **Marquee Tool**, which is the box inside the **Overview** pane.

Active synchronizations, completed synchronizations, and failed synchronizations are represented with colors. To set the colors, open the MobiLink Monitor, choose **Tools » Options**, and then click the **Overview** tab.

See [“Options window” on page 190](#).

You can close the **Overview** pane by deselecting it in the **View** menu.

You can also separate the **Overview** pane from the rest of the MobiLink Monitor window. In the **Options** window, open the **Overview** tab and clear the **Keep overview window attached to main window** checkbox.

### Marquee tool

The **Marquee Tool** is the small box that appears in the **Overview** pane. You can use the **Marquee Tool** to see different data, or to see data at different granularity. The area represented within the box is displayed in the chart and graph panes. You can use the **Marquee Tool** as follows:

- Click in the **Overview** pane to move the **Marquee Tool** and the start time of the data shown in the chart or utilization graph.
- Drag in the **Overview** pane to redraw the **Marquee Tool** to change the **Marquee Tool**'s location and size and change the start time and the range of data. If you make the marquee box smaller, you shorten the interval of the visible data in the chart, which makes more detail visible.

### To change the color of the Marquee Tool

1. Choose **Tools » Options**.
2. Click the **Overview** tab.
3. Select a new color in the **Marquee** field.
4. Click **OK**.

## Options window

Options allow you to specify many settings, including colors and patterns for the graphical display in the **Chart** pane and the **Overview** pane.

To open the **Options** window, open the MobiLink Monitor and choose **Tools » Options**.

### Restoring defaults

To restore default settings, delete the file *mlMonitorSettings11*. This file is stored in your user profiles directory.

## Session properties

The **Session Properties** window provides statistics about the session. It provides property values for the entire time that the MobiLink Monitor has been running. To open the **Session Properties** window, open the MobiLink Monitor and choose **File » Properties**.

The **General** tab provides basic information about the session.

The **Statistics** tab shows the same statistics as **Sample Properties**. See [“Sample properties” on page 190](#).

## Sample properties

The **Sample Properties** window provides detailed statistics for time intervals. Each time interval is about one second long. Samples are numbered by the MobiLink Monitor to reflect the order in which they were received.

You can customize the appearance of the graph to hide properties but all properties appear in the **Sample Properties** window. If you have hidden a property, it is identified as **Hidden** in the **Sample Properties** window; otherwise the color is shown.

To open the **Sample Properties** window, click in the **Graph** pane for the time period that you want to examine.

The **Sample** tab provides information for the one-second time interval it represents. The properties that are displayed are for the end of the time interval.

The **Range** tab shows averages for the entire range of samples that were visible when the properties window was opened (the horizontal range that is visible in the **Overview**). The range statistics are not calculated until you click **Calculate** in the **Range** tab.

**Sample Properties** contains the following information:

- **Sample** Samples are numbered by the MobiLink Monitor to reflect the order in which they were received. On the **Sample** tab, this shows the sample number. On the **Range** tab it shows the range of samples.
- **Start time and end time** On the **Sample** tab, this reflects a sample time period of approximately one second.
- **Statistics - Color column** The color used in the graph for this property.
- **Statistics - Property column** Shows the queue length for the sample or average queue for the range. This column displays the following types of property:
  - **TCP/IP work queue** This lists the number of buffers waiting to be filled by reading from the network plus the number of buffers waiting to be written to the network. The actual number is not very meaningful, but large numbers may indicate network-related bottlenecks.
  - **Stream work queue** For version 10 clients only. This queue represents work done by the high-level network layer in the MobiLink server. This layer is responsible for higher-level network protocol work such as HTTP, encryption, and compression.
  - **Heartbeat work queue** This is the queue length for periodic internal MobiLink server tasks other than synchronizations.
  - **Command processor work queue** This is the queue length for performing database tasks and interpreting or preparing communications with MobiLink clients. The actual number is not very meaningful, but large numbers may indicate database-related bottlenecks.
  - **Busy database worker threads** This value indicates how hard the MobiLink server is pushing the consolidated database. Each unit in the value represents a database worker thread that is doing something in the database. There is no distinction between inserts, updates, deletes, or selects. When this value is zero, the server is not operating on the consolidated database.

Note: The names of the properties are obtained from the MobiLink server or .mlm file and are in the language of the MobiLink server. Some characters may not display properly if the MobiLink Monitor is using a different language.

- **Statistics - Value column** The property value.
- **Statistics - Limit column** The maximum allowed value for the property. This is useful so that the graph can use a scale of 0-100% for all properties. For properties such as page faults that are essentially unbounded, the limit is ignored.

## Synchronization properties

Double-click a synchronization in either the **Details Table** pane or the **Chart** pane to see properties for that synchronization.

You can choose to see statistics for all tables (which is the sum for all tables in the synchronization), or for individual tables. The dropdown list provides a list of the tables that were involved in the synchronization.

For descriptions of the quantities displayed on any page of the **Synchronization Properties** window, click **Help**.

For an explanation of the statistics in **Synchronization Properties** window, see [“MobiLink statistical properties” on page 195](#).

---

## Saving MobiLink Monitor data

You can save the data from a MobiLink Monitor session as a binary file (.mlm), as a text file with comma-separated values (.csv), as tables in a relational database, or as a Microsoft Excel file.

### Saving to file

To save the data as a file, choose **File » Save As**.

- Save the data as a binary (.mlm) file if you want to view the saved data in the MobiLink Monitor. To reopen, choose **File » Open**. The binary file format is the only format that preserves all monitored information.
- Save the data as a comma separated file (.csv) if you want to view it in another tool, such as Microsoft Excel. This only saves the information in the synchronization properties windows, except per table information and the session end time. You can also open a .csv file in the MobiLink Monitor.

In the .csv file format, time durations are stored in milliseconds.

You can specify that you want data to be saved automatically to a file. To do this, choose **Tools » Options**, and enter an output file name on the **General** tab. The output file is overwritten by new data.

### Exporting to a relational database or Excel

You can also export MobiLink Monitor data using an ODBC connection. You can export to any relational database that is supported by MobiLink, and to a Microsoft Excel spreadsheet.

When you export data, all the columns in your MobiLink Monitor session are exported, and a column named `export_time` that identifies the time the export was performed. Data from the graph is not exported.

The data source must have quoted identifiers enabled, because some of the columns are reserved words. The MobiLink Monitor enables quoted identifiers automatically for SQL Anywhere, Adaptive Server Enterprise, and Microsoft SQL Server databases. If the quoted identifiers option is not enabled, the export fails.

### To export the data to a database or Excel

1. After collecting MobiLink Monitor information, disconnect from the MobiLink server.
2. In the MobiLink Monitor, choose **File » Export To Database**.
3. Select options for the output.
  - You can name the two tables that are created to hold the data, or use the defaults. If the tables do not exist, they are created by the MobiLink Monitor. For Excel output, the two table names identify the two worksheets that are created.
  - Choose whether you want to overwrite data in existing tables. If you do not choose to overwrite the data, new data is appended to existing data.
4. Click **OK** to open the **Connect** window and connect to the database or Excel spreadsheet using ODBC.

## Customizing your statistics

The **Watch Manager** allows you to visibly distinguish synchronizations that meet criteria that you specify. For example, you might want to highlight big synchronizations, long synchronizations, small synchronizations that take a long time, or synchronizations that receive warnings.

To open the **Watch Manager**, open the MobiLink Monitor and then click **Tools » Watch Manager**.

The left pane of the **Watch Manager** contains a list of all available watches. The right pane contains a list of active watches. To add or remove a watch from the active list, select a watch in the left pane and click the appropriate button.

There are three predefined watches (**Active**, **Completed**, and **Failed**). You can edit predefined watches to change the way they are displayed, and you can deactivate them by removing them from the right pane.

No synchronizations are displayed in the chart unless they meet the conditions of a watch. If you disable all watches (by removing them from the **Current Watches** list), then no synchronizations are shown in the **Chart** pane or the **Overview** pane.

The order of watches in the right pane is important. Watches that are closer to the top of the list are processed first. Use the **Move Up** and **Move Down** buttons to organize the order of watches in the right pane.

You can use the predefined watches, and create other watches. To edit a watch condition, remove it and then add the new watch condition.

When a new MobiLink Monitor connects to the same MobiLink server, it shows up as a short synchronization in any MobiLink Monitors that are already connected. The MobiLink Monitor synchronization has the version name for\_ML\_Monitor\_only. You can hide this MobiLink Monitor synchronization with a watch.

### To create a new watch

1. In the **Watch Manager**, click **New**.
2. Give the watch a name in the **Name** box.
3. Select a **Property**, comparison **Operator**, and **Value**.  
For a complete list of properties, see [“MobiLink statistical properties” on page 195](#).
4. Click **Add**. (You must click **Add** to save the settings.)
5. If desired, select another **Property**, **Operator**, and **Value**, and click **Add**.
6. Select a **Chart Pattern** for the watch in the **Chart** pane.
7. Select an **Overview Color** for the watch in the **Overview** pane.

## MobiLink statistical properties

The following is a list of the properties that are available in the MobiLink Monitor. These statistics can be viewed in the **New Watch** window, the **Details Table** pane, or the **Synchronization Properties** window. In **Synchronization Properties**, the property names do not contain underscores.

For more information about the **New Watch** window, see [“Customizing your statistics” on page 194](#).

For more information about the **Details Table**, see [“Details Table pane” on page 184](#).

For more information about the **Synchronization Properties** window, see [“Synchronization properties” on page 191](#).

### Synchronization statistics

MobiLink statistical properties return the following information for synchronizations when not using forced conflict mode.

For information about forced conflict mode, see [“Forced conflict statistics” on page 198](#).

Property	Description
active	True if the synchronization is in progress.
apply_upload	The time required for the uploaded data to be applied to the consolidated database.
authenticate_user	Total time to perform user authentication, including executing the authenticate_* events.
begin_sync	Total time for the begin_synchronization event.
completed	True if the synchronization completed successfully.
conflicted_deletes	Always zero.
conflicted_inserts	Always zero.
conflicted_updates	Number of update rows that caused conflict. A row is included only when a resolve conflict script was successfully called for it.
connect	The time required by the database worker thread to make a database connection if a new database connection is needed. For example, after an error or if the script version has changed.
connect_for_download_ack	The time required by the database worker thread to make a database connection if a new database connection is needed.
connection_retries	Number of times the MobiLink server retried the connection to the consolidated database.

Property	Description
download_bytes	Amount of memory used within the MobiLink server to store the download. This provides a good indication of the impact on server memory of a synchronization.
download_deleted_rows	Number of row deletions fetched from the consolidated database by the MobiLink server (using download_delete_cursor scripts).
download_errors	Number of errors that occurred during the download.
download_fetched_rows	Number of rows fetched from the consolidated database by the MobiLink server (using download_cursor scripts).
download_filtered_rows	Number of fetched rows that were not downloaded to the MobiLink client because they matched rows that the client uploaded.
download_warnings	Number of warnings that occurred during the download.
duration	Total time for the synchronization, as measured by the MobiLink server.
end_sync	Total time for the end_synchronization event.
fetch_download	The time required to fetch the rows to be downloaded from the consolidated database to create the download stream.
get_db_worker	The time required to acquire a free database worker thread.
get_db_worker_for_download_ack	The time spent waiting for a free database worker thread after the download acknowledgement has been received.
ignored_deletes	Number of upload delete rows that caused errors while the upload_delete script was invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_delete script defined for the given table.
ignored_inserts	The total number of upload insert rows that were ignored. They were ignored because 1) there is no upload_insert script in normal mode or no upload_new_row_insert script in forced conflict mode; or 2) errors occurred when the MobiLink server was invoking the corresponding script and the handle_error or handle_odbc_error event returned 1000.
ignored_updates	Number of upload update rows that caused conflict but a resolve conflict script was not successfully called or no upload_update script was defined.



Property	Description
nonblocking_download_ack	The time required for the publication_nonblocking_download_ack connection and nonblocking_download_ack connection events.
prepare_for_download	Total time for the prepare_for_download event.
remote_id	The remote ID that uniquely identifies the remote database.
send_download	The time required to send the download stream to the remote database. The time depends on the size of the download stream and the network bandwidth for the transfer. For an upload-only synchronization, the download stream is simply an upload acknowledgement.
start_time	Date-time (in ISO-8601 extended format) for the start of the synchronization.
sync	A number uniquely identifying the synchronization within the MobiLink Monitor session.
sync_deadlocks	Number of deadlocks in the consolidated database that were detected for the synchronization.
sync_errors	Total number of errors that occurred for the synchronization.
sync_request	The time taken between creating the network connection between the remote database and the MobiLink server, up to receiving the first bytes of the upload stream.
sync_tables	Number of client tables that were involved in the synchronization.
sync_warnings	Number of warnings that occurred for the synchronization.
upload_bytes	Amount of memory used within the MobiLink server to store the upload. This provides a good indication of the impact on server memory of a synchronization.
upload_deadlocks	Number of deadlocks in the consolidated database that were detected during the upload.
upload_deleted_rows	Number of rows that were successfully deleted from the consolidated database.
upload_errors	Number of errors that occurred during the upload.

Property	Description
upload_inserted_rows	Number of rows that were successfully inserted in the consolidated database.
upload_updated_rows	Number of rows that were successfully updated in the consolidated database.
upload_warnings	Number of warnings that occurred during the upload.
user	MobiLink user name.
version	Name of the synchronization version.
wait_for_download_ack	The time spent waiting for the download to be applied to the remote database and for the remote database to send the download acknowledgement.

### Forced conflict statistics

When you are in forced conflict mode, MobiLink statistical properties return the following information.

Statistical property	Description
conflicted_deletes	Number of upload delete rows that were successfully inserted into the consolidated database using the upload_old_row_insert script.
conflicted_inserts	Number of upload insert rows that were inserted into the consolidated database using the upload_new_row_insert script.
conflicted_updates	Number of upload update rows that were successfully applied using the upload_new_row_insert or upload_old_row_insert scripts.
ignored_deletes	Number of upload delete rows that caused errors while the upload_old_row_insert script was invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_old_row_insert script defined for the given table.
ignored_inserts	Number of upload insert rows that caused errors while the upload_new_row_insert script was invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_new_row_insert script defined for the given table.
ignored_updates	Number of upload update rows that caused errors while the upload_new_row_insert or upload_old_row_insert scripts were invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_new_row_insert and upload_old_row_insert script defined for the given table.

---

<b>Statistical property</b>	<b>Description</b>
upload_deleted_rows	Always zero.
upload_inserted_rows	Always zero.
upload_updated_rows	Always zero.

---

---

# SQL Anywhere Monitor for MobiLink

## Contents

Introducing the SQL Anywhere Monitor .....	202
Monitor quick start .....	205
Tutorial: Using the Monitor .....	206
Start the Monitor .....	211
Exit the Monitor .....	212
Connect to the Monitor .....	213
Disconnect from the Monitor .....	214
Monitoring resources .....	215
Administering resources .....	222
Working with Monitor users .....	228
Alerts .....	232
Installing the SQL Anywhere Monitor on a separate computer .....	236
Troubleshooting the Monitor .....	237

---

## Introducing the SQL Anywhere Monitor

The SQL Anywhere Monitor, also referred to as the Monitor, is a web browser-based administration tool that provides you with information about the health and availability of SQL Anywhere databases and MobiLink servers.

This chapter describes how to use the Monitor to collect metrics about MobiLink servers. For information about using the Monitor with SQL Anywhere databases, see “SQL Anywhere Monitor” [[SQL Anywhere Server - Database Administration](#)].

The Monitor provides the following functionality:

- **Constant data collection** Unlike many of the other administration tools available with SQL Anywhere 11, the Monitor collects metrics all the time, even when you are not logged in to the web browser. The Monitor collects metrics until you shut it down.
- **Email alert notification** As the metrics are collected, the Monitor examines the metrics and can send email alerts when it detects conditions that indicate something is wrong with a MobiLink server.
- **Browser-based interface** At any time, you can connect to the Monitor using a web browser to review alerts and metrics that have been collected.
- **Monitor multiple databases and MobiLink servers** From one tool, you can simultaneously monitor SQL Anywhere databases and MobiLink servers running on the same or different computers.

For information about monitoring SQL Anywhere databases see “SQL Anywhere Monitor” [[SQL Anywhere Server - Database Administration](#)].

- **Minimal performance impact** The Monitor can be used routinely in development and production environments because monitoring does not degrade performance.

### Requirements

- It is recommended that install the latest version of Adobe Flash Player that is available for your operating system. The Monitor is backwards compatible with version 9 of Adobe Flash Player. To determine the correct version, visit <http://www.adobe.com/products/flashplayer/systemreqs/>.
- You must enable JavaScript in your web browser.
- You must have SQL Anywhere 11.0.1 installed.

### Running the Monitor in a production environment

You can install and run the Monitor on a separate computer. This prevents the Monitor resources and configuration from being overwritten during subsequent SQL Anywhere upgrades or updates. Installing on a separate computer is recommended if you want to use the Monitor in a production environment. See “Installing the SQL Anywhere Monitor on a separate computer” [[SQL Anywhere Server - Database Administration](#)].

## Limitations

- You can use the Monitor to collect metrics about the following types of SQL Anywhere databases and MobiLink servers:
  - SQL Anywhere 9.0.2, 10.0.0, 10.0.1, 11.0.0, and 11.0.1
  - MobiLink 11.0.0 with at least the first EBF applied and 11.0.1
- You can only run one Monitor on a computer.
- The Monitor does not provide information about individual synchronizations. For detailed information about individual synchronizations, including timing and other per-synchronization statistics, use the MobiLink Monitor. See [“Introduction to the MobiLink Monitor” on page 180](#).

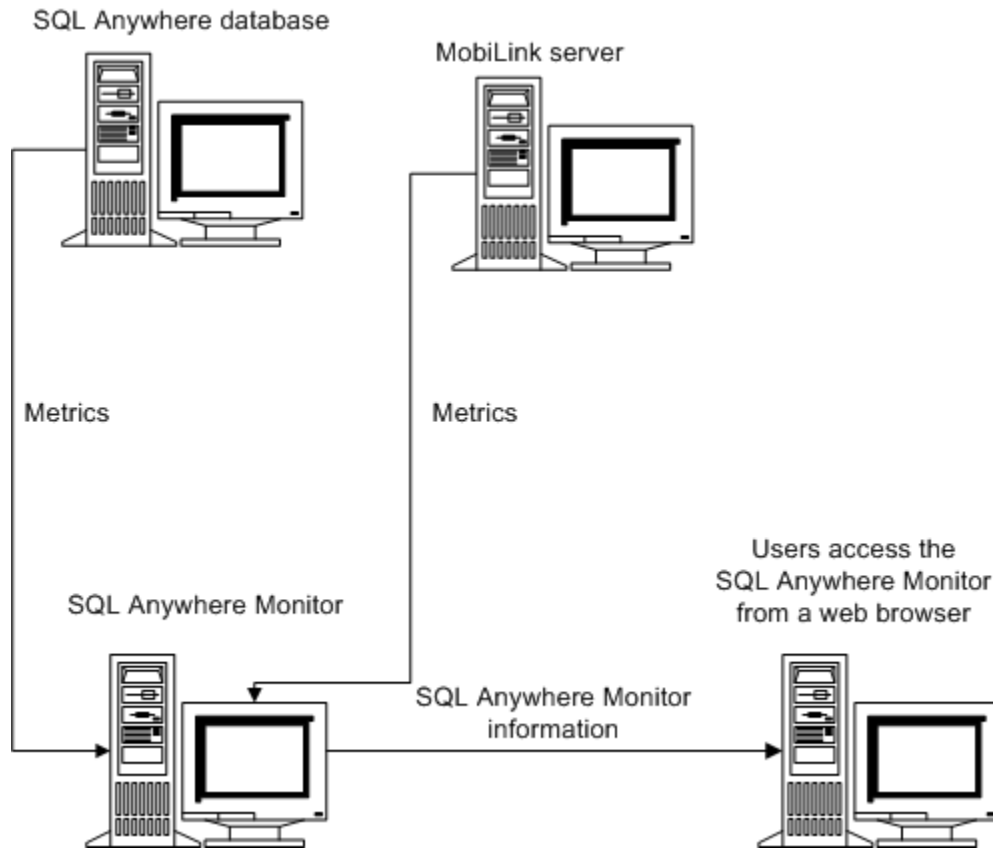
## See also

For information about other administration and performance tools that are available for MobiLink servers, see:

- [“MobiLink Monitor” on page 179](#)

## Monitor architecture

The Monitor collects metrics and performance data from SQL Anywhere databases and MobiLink servers running on other computers, while a separate computer accesses the Monitor via a web browser.



The Monitor is designed to help any type of user, whether they are a DBA or not, who is responsible for such tasks as:

- Ensuring that a MobiLink server is connected to the network.
- Ensuring that there is enough disk space or memory available for a MobiLink server.
- Reviewing the number of synchronizations a MobiLink server performs over a specified time.

**See also**

- [“Monitor quick start” on page 205](#)



---

## Monitor quick start

The following steps are required to set up MobiLink server monitoring:

1. Install SQL Anywhere 11.0.1 on a computer that is always connected to your network. The Monitor uses SQL Anywhere to monitor MobiLink servers.

The Monitor can run on the same computer as the resources it is monitoring, but it is recommended, particularly in production environments, that you run the Monitor on a different computer to minimize the impact on the MobiLink server, or other applications.

2. Ensure that your web browser has the appropriate version of Adobe Flash Player installed and that JavaScript is enabled. See [“Requirements” on page 202](#).
3. Start your MobiLink server (if it is not already running).
4. Start the Monitor and open it in your web browser. See [“Start the Monitor” on page 211](#).

The computer where you are using a web browser to access the Monitor must be connected to the network where the Monitor is running.

5. Log in as an administrator. The default user name is **admin** and the default password is also **admin**.
6. Click the **Administration** tab and add a MobiLink server as a resource to be monitored. See [“Add resources” on page 222](#).
7. Add new users and change the password for the admin user. See [“Create Monitor users” \[SQL Anywhere Server - Database Administration\]](#).
8. Configure alerts for the MobiLink server you want to monitor. See [“Alerts” \[SQL Anywhere Server - Database Administration\]](#).
9. Click the **Monitoring** tab to see the collected metrics for your MobiLink server. See [“Monitoring resources” on page 215](#).

# Tutorial: Using the Monitor

Use this tutorial to set up monitoring of the MobiLink Synchronization Server Sample.


## Lesson 1: Start the Monitor

### To start and open the Monitor

1. Start the Monitor. Choose **Start » Programs » SQL Anywhere 11 » SQL Anywhere Monitor » SQL Anywhere Monitor**.


You do not have to perform this step if you installed the Monitor on a separate computer. When the Monitor is installed on a separate computer than the one SQL Anywhere is running on, it runs as a service and is automatically started when the computer starts.

2. Browse data. This step is different depending on whether the Monitor is installed on a separate computer. In the system tray, right-click the SQL Anywhere Monitor icon and choose **Browse Data**.



If the Monitor is installed on a separate computer, choose **Start » Programs » SQL Anywhere Monitor 11 » Browse Data**. No icon appears in the system tray.

Alternatively, you can open a web browser and browse to *http://localhost:4950*.



The top pane of the **Monitoring** tab lists the resources that are being monitored. When you first open the Monitor, it is only monitoring itself.

**See also**

- [“Lesson 2: Set up the Monitor to monitor a MobiLink server” on page 207](#)

## Lesson 2: Set up the Monitor to monitor a MobiLink server

The Monitor collects metrics from databases and MobiLink servers. In this section, you start the MobiLink Synchronization Server Sample, and then add the server as a resource to be monitored. To collect metrics from a SQL Anywhere database, see [“Lesson 2: Set up the Monitor to monitor a database” \[SQL Anywhere Server - Database Administration\]](#)

**To add a resource to monitor**

1. Start the MobiLink Synchronization Server Sample.

From the **Start** menu, choose **Programs » SQL Anywhere 11 » MobiLink » Synchronization Server Sample**.

2. Log in as the default administrator to the Monitor:
  - a. Click **Login**.
  - b. In the **User Name** field, type **admin**, and in the **Password** field, type **admin**.
  - c. Click **Login**.
3. Click the **Administration** tab.
4. Click the **Resources** tab.
5. Click **Add**.
6. Select **MobiLink Server**, and then click **Next**.
7. Name the resource **MobiLinkServerSample**, and then click **Next**.
8. In the **Host** field, type **localhost**, and then click **Next**.
9. When you are prompted for the required authorization, in the **User ID** field, type a user name such as **user1**, and in the **Password** field, type a password, such as **sql**.

These credentials are used to connect to the MobiLink server. The user ID and password are kept by the Monitor.
10. Click **Create**.
11. The new resource, **MobiLinkServerSample**, is created and monitoring starts.
12. Click **OK**.
13. Click the **Monitoring** tab.

The **MobiLinkServerSample** resource appears on the **Monitoring** tab, and the collected metrics appear on the tabs in the bottom pane.

## Lesson 3: Test an alert

In this lesson, you intentionally trigger an alert so you can practice handling alerts.

### To view and resolve an alert

1. Trigger an alert by shutting down the MobiLink Synchronization Server Sample.
  - a. On Windows, double-click the MobiLink server icon in the system tray for the MobiLink server.
  - b. Click **Shut Down** in the MobiLink Server window.
  - c. Click **Yes**.

2. In the Monitor, click the **Monitoring** tab.

The **State** for the MobiLinkServerSample resource changes to **Server Down** and the **Status** for the MobiLinkServerSample resource changes to **Needs Attention!**.

It can take a few seconds for these changes in state and status to occur. By default, the Monitor collects information from the resource every 30 seconds.

3. In the bottom pane, click **Alerts**.
4. Select the **Availability Alert** and click **Details** to read the description.
5. Click **OK**.
6. Restart the Synchronization Server Sample.

From the **Start** menu, choose **Programs » SQL Anywhere 11 » MobiLink » Synchronization Server Sample**.

The **State** for the MobiLinkServerSample resource changes to **Alive**, but the **Status** remains unchanged. It can take a few minutes for the change to appear.

7. Delete the alert by selecting the alert and clicking **Delete**.

The **Status** changes to **Healthy**.

## Lesson 4: Set up the Monitor to send emails when alerts occur

When an alert occurs, it is always listed in the **Alerts** tab in the lower pane of the **Monitoring** tab. In the following procedure, you set up the Monitor to send you an email whenever an alert occurs.

### To set up email notification

1. Create a user who can receive emails.
  - a. Click the **Administration** tab.
  - b. Click the **Users** tab.
  - c. Click **New**.

- d. In the **User Name** field, type **JoeSmith**.
  - e. In the **Password** and the **Confirm Password** fields, type **sql**.
  - f. In the **Email** field, enter a valid email address.
  - g. Choose **English** in the **Preferred Language** field.
  - h. Select **Operator** for the **User Type**.

An operator can receive alerts via email and can resolve and delete alerts. This user can access the **Monitoring** tab, but cannot access the **Administration** tab.

For information about the different types of users, see [“Working with Monitor users” on page 228](#).
  - i. Click **Save**.

The new user is created.
2. Associate the user with the MobiLinkServerSample resource.
    - a. Click the **Resources** tab.
    - b. Select the **MobiLinkServerSample** resource and click **Configure**.
    - c. In the **Configure Resource** window, click **Operators**.
    - d. In the **Available Operators** list, select **JoeSmith** and click **Add**.
    - e. Click **Save**.
    - f. Click **OK**.
  3. Configure email alert notification.
    - a. Click the **Administration** tab.
    - b. Click the **Configuration** tab.
    - c. Click **Edit**.
    - d. Select **Send Alert Notifications By Email**.
    - e. Configure the other settings as required.
    - f. Test that you have properly configured email notification.

Click **Send Test Email**.
    - g. When prompted, enter an email address to send the test email to and click **OK**.

A test email is sent to the email address specified.
    - h. Click **Save**.

When an alert occurs, an email is sent to the specified user with information about the alert. For information about setting up alerts, see [“Lesson 3: Test an alert” on page 208](#).

## Lesson 5: Cleanup

The following procedure removes the MobiLinkServerSample resource, which deletes the collected metrics and stops data collection. In a production environment when you want to continue monitoring your MobiLink server, you leave both the MobiLink server and the Monitor running.

### To stop monitoring

1. Remove the MobiLinkServerSample resource.
  - a. Click the **Administration** tab.
  - b. Click the **Resources** tab.
  - c. Select the MobiLinkServerSample resource, and click **Stop**.
  - d. Click **Remove**.
  - e. Click **Yes** to confirm that you want to remove the resource.
2. Log out of the Monitor.

Click **Logout**.
3. Close the web browser window where you are viewing the Monitor.
4. Exit the Monitor.

In the system tray, right-click the SQL Anywhere Monitor icon and choose **Exit SQL Anywhere Monitor**.
5. Shut down the MobiLink server.
  - a. Double-click the MobiLink Server icon in the system tray for the MobiLink Synchronization Server Sample.
  - b. Click **Shut Down** in the MobiLink Server messages window.
  - c. Click **Yes**.

---

## Start the Monitor

Starting the Monitor causes the Monitor to start collecting metrics for *all* resources in the Monitor.

The procedure for starting the Monitor is different depending on whether the Monitor is running on a separate computer.

### To start the Monitor

1. Choose **Start » Programs » SQL Anywhere 11 » SQL Anywhere Monitor » Start SQL Anywhere Monitor**.

The SQL Anywhere Monitor icon appears in the system tray.

2. Connect to the Monitor. See [“Connect to the Monitor” on page 213](#).

### To start the Monitor on a separate computer

1. The Monitor runs automatically as a service when installed on a separate computer. However, if you stop monitoring, you can restart it. To do so, browse to *install-dir\bin32*.
2. On Windows, run the following:

```
samonitor.bat start service
```

On Linux, run the following:

```
samonitor.sh start service
```

When the Monitor runs as a service, no SQL Anywhere Monitor icon appears the system tray.

3. Connect to the Monitor. See [“Connect to the Monitor” on page 213](#).

### See also

- [“Exit the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Connect to the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Disconnect from the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Monitoring resources” \[SQL Anywhere Server - Database Administration\]](#)

## Exit the Monitor

Exiting the Monitor stops the collection of metrics for all resources. It is recommended that you leave the Monitor running, but close the web browser. To stop monitoring a specific MobiLink server, see [“Stop monitoring resources” on page 225](#).

The procedure for exiting the Monitor is different depending on whether the Monitor is running on a separate computer.

### To exit the Monitor

- In the system tray, right-click the SQL Anywhere Monitor icon and choose **Exit SQL Anywhere Monitor**.

### To exit the Monitor on a separate computer

1. Browse to *install-dir\bin32*.
2. On Windows, run the following:

```
samonitor.bat stop service
```

On Linux, run the following:

```
samonitor.sh stop service
```

### See also

- [“Start the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Connect to the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Disconnect from the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Monitoring resources” \[SQL Anywhere Server - Database Administration\]](#)



## Connect to the Monitor

The computer that you are using to connect to the Monitor must be connected to the network where the Monitor is running.

### To connect to the Monitor

1. Start the Monitor, if it isn't already running. See [“Start the Monitor” on page 211](#).
2. Browse data. This step is different depending on whether the Monitor is installed on a separate computer.

From the **Start** menu, choose **Programs » SQL Anywhere 11 » SQL Anywhere Monitor » Browse Data**.

If the Monitor is installed on a separate computer, choose **Start » Programs » SQL Anywhere Monitor 11 » Browse Data**.

A web browser opens the default URL for connecting to the Monitor: **http://computer-name:4950**, where *computer-name* is the name of the computer the Monitor is running. For example, *http://localhost:4950*.

3. If prompted, enter your user name and password for the Monitor. The user name and password for the Monitor are case sensitive. See [“Working with Monitor users” on page 228](#).

### See also

- [“Start the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Exit the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Disconnect from the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Monitoring resources” \[SQL Anywhere Server - Database Administration\]](#)

## Disconnect from the Monitor

You can disconnect from the Monitor by logging out or closing the web browser.

Disconnecting from the Monitor has no effect on the collection of metrics. If you want to stop collecting metrics, then stop monitoring the resource or exit the monitor. See [“Stop monitoring resources” on page 225](#), or [“Exit the Monitor” on page 212](#).

### To disconnect from the Monitor

- Click **Logout**.

### See also

- [“Start the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Exit the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Connect to the Monitor” \[SQL Anywhere Server - Database Administration\]](#)
- [“Monitoring resources” \[SQL Anywhere Server - Database Administration\]](#)

## Monitoring resources

In the Monitor, the **Monitoring** tab provides an overview of the health and availability of the MobiLink servers being monitored.

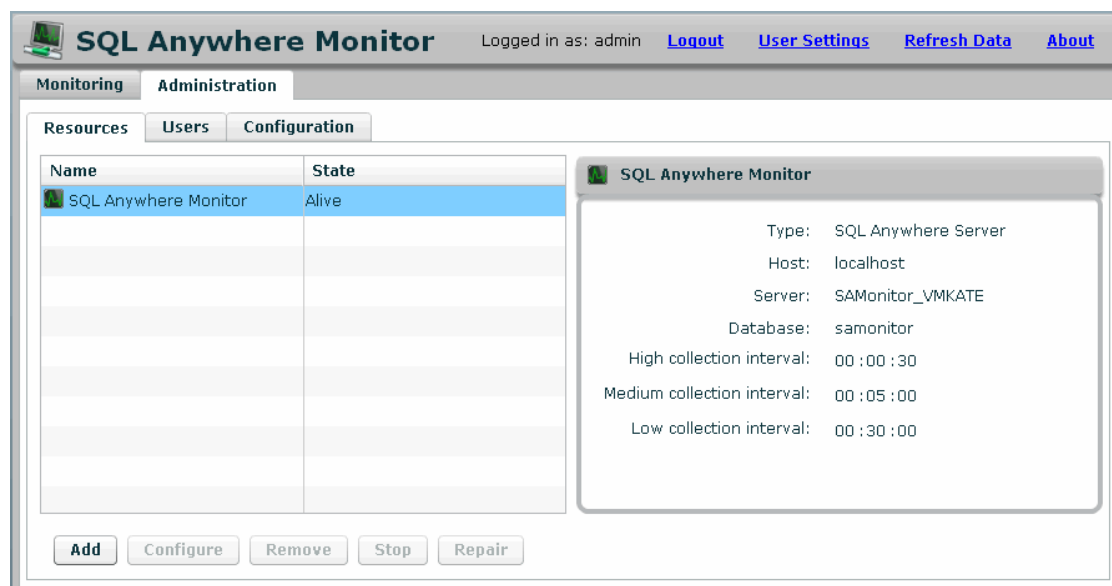
### Monitoring tab

The top pane contains a table that lists the resources that are being monitored. A **resource** is a MobiLink server. This table also indicates whether the resources are currently running and whether they require a user to perform any actions on them. See [“Interpreting resource states and status” on page 215](#).

The bottom pane of the **Monitoring** tab contains the alerts and a variety of current metrics for the selected MobiLink server. Most of these tabs contain links to graphs. You can change the range of the graphs with the dropdown list and arrows at the top right of each graph.

### Administration tab

The **Administration** tab is reserved for administrators. On it, you can select the MobiLink servers that you want to monitor, add and edit users, and configure the Monitor.



### See also

- [“Working with Monitor users” on page 228](#)
- [“Monitor metrics” on page 216](#)

## Interpreting resource states and status

The top pane of the **Monitoring** tab contains a table that lists the MobiLink servers that are being monitored. In this table, the **State** column provides information about the connections between the Monitor and its

resources. The **Status** column indicates whether the resources require an operator or an administrator user to perform actions on them. See [“Working with Monitor users” on page 228](#).

### Resource state

A resource is always in one of the following states:

- **Alive** The resource is connected and the Monitor is collecting metrics.
- **Blackout** The Monitor is waiting for the blackout period to end before it resumes monitoring of the resource.
- **Server Down** The MobiLink server being monitored is stopped.
- **Host Down** The Monitor cannot locate the computer that is hosting the resource.
- **Unknown** The Monitor is not monitoring the resource.

### Resource status

A resource has one of the following statuses:

- **Healthy** There are no unresolved alerts for the resources.
- **Needs Attention** There are one or more alerts for the resource.
- **Monitoring Stopped** The resource is not being monitored.
- **Unknown** The resource is not alive and there are no alerts for it.

## Monitor metrics

The Monitor collects and stores metrics from MobiLink servers, including, but not limited to:

- Whether the resource is running.
- Whether the computer that the resource is running on is running properly and is connected to the network.
- Whether the resource is listening and processing requests.
- The number of synchronizations that the MobiLink server performs over a period of time.

The rate at which metrics are collected is determined by the collection interval settings that are set by administrators. See [“Collection intervals” on page 223](#).

Which metrics are collected and what thresholds should be used to issue alerts are determined by the metric settings that are set by the administrators. See [“Specify metrics to collect” on page 223](#).

### Displaying metrics

The Monitor display is automatically refreshed every minute. You can change the refresh the interval by clicking **User Settings**. This setting is independent of the collection interval rate for a resource, which specifies how often the Monitor collects metrics from the resource being monitored.

### To set the refresh rate

1. Click **User Settings** in the top, right corner.
2. Set a time for the **Refresh Interval**. The default is one minute.
3. Click **OK**.

When you click **Refresh Data** on the **Monitoring** tab, the Monitor retrieves and displays the latest metrics.

### To refresh metrics

- Click **Refresh Data**.

When you press F5, the Monitor reloads the web browser and retrieves and displays the metrics that the Monitor has collected to date.

### To reload the Monitor

- Press F5.

## Metric tab descriptions

The following tabs are used by both SQL Anywhere and MobiLink server resources.

- “Monitoring tab: Alerts tab” on page 218
- “Monitoring tab: Server tab” on page 218

The following tabs are used only by SQL Anywhere resources.

- “Monitoring tab: CPU tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Unscheduled Requests tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Memory tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Disk tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: HTTP tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Connections tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Failed Connections tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Queries tab” [*SQL Anywhere Server - Database Administration*]
- “Monitoring tab: Mirror tab” [*SQL Anywhere Server - Database Administration*]

The following tabs are used only by MobiLink server resources.

- “Monitoring tab: Synchronization tab” on page 219
- “Monitoring tab: Consolidated Database tab” on page 220

- [“Monitoring tab: Machine Resources tab” on page 220](#)

## Monitoring tab: Alerts tab

Lists the fifty most recent alerts. Once the list exceeds 50 alerts, old alerts are removed as new alerts arrive. See [“Alerts” on page 232](#).

## Monitoring tab: Server tab

### MobiLink Server

- **Server Name** The name of the MobiLink server as specified by the `-zs` option for the connected server. The default value is `<default>`. See [“-zs option” on page 119](#).
- **Version** Shows the version of the software being run.
- **Start Time** Shows the time when the MobiLink server started.
- **Unsubmitted error reports** Shows the number of unsubmitted error reports for the server. An error report is submitted when SQL Anywhere software crashes. See [“Suppress alerts for unsubmitted error reports from resources” on page 235](#).

### License

- **Name Of Licensed Company** Shows the name of the licensed company.
- **Name of Licensed User** Shows the name of the licensed user.

### Host

- **Name** Shows the name of the computer running the MobiLink server. Typically, this is the computer's host name.
- **Operating System Platform** Shows the operating system on which the software is running.
- **Processor Architecture** Shows a string that identifies the processor type.
- **Number of CPUs** Shows the number of CPUs that the machine running the software has.

### Additional Information

- **Consolidated Database Type** Shows the type of consolidated database. For example, SQL Anywhere.
- **Maximum Concurrent Uploads To Database** Shows the maximum number of concurrent uploads to the database. See [“-wu option” on page 106](#).
- **Database Worker Threads** Shows the number of database worker threads. See [“-w option” on page 105](#).
- **Maximum Cache Size** The maximum size for the MobiLink server memory cache, as set by the `-cm` option for `mlsrv11`. See [“-cm option” on page 55](#).

- **Maximum Number Of Pages In Cache** The number of pages in the MobiLink memory cache. This is implicitly set with the `-cm` option for `mlsrv11`. See [“-cm option” on page 55](#).
- **Maximum Number Of Database Connections** The maximum number of database connections, as set by the `-cn` option or the `-w` option for `mlsrv11`. See [“-cn option” on page 56](#) and [“-w option” on page 105](#).
- **Maximum Number Of TCP Connections** The maximum number of TCP connections, as set by the `-nc` option for `mlsrv11`. See [“-nc option” on page 75](#).
- **Maximum Number Of Clients** The maximum number of clients. See [“-sm option” on page 94](#).
- **Consolidated Version** Shows the version of the consolidated database.
- **Driver Version** Shows the version of the driver for the consolidated database.
- **Driver Name** Shows the name of the driver for the consolidated database.
- **Primary Server In Farm** Indicates if the server is primary or secondary.

#### See also

- [“Monitoring tab: Server tab” on page 218](#)

## Monitoring tab: Synchronization tab

This tab is used when monitoring MobiLink servers.

- **Completed Synchronization Rate** Shows the rate of completed synchronizations for the server, expressed in synchronizations per second.
- **Failed Synchronization Rate** Shows the rate of failed synchronizations for the server, expressed in synchronizations per second.
- **Synchronization Error Rate** Shows the rate of synchronization errors, expressed in errors per second.
- **Synchronization Warning Rate** Shows the rate of synchronization warnings for the server, expressed in warnings per second.
- **Longest Active Synchronization Time** Shows the elapsed time of the oldest active synchronization for the server, in seconds.
- **Active Requests** Shows the number of active requests in the server.
- **Applying Upload** Shows the number of requests in the server that are currently in the apply upload or begin synchronization phase. For a description of the synchronization phases, see [“-v option” on page 102](#).
- **Generating Download** Shows the number of requests in the server that are currently in the prepare for download, fetch download, wait for download acknowledgement or end synchronization phase. For a description of the synchronization phases, see [“-v option” on page 102](#).
- **Active Authentications** Shows the number of requests in the server that are currently in the authenticate user phase. For a description of the synchronization phases, see [“-v option” on page 102](#).

## Monitoring tab: Consolidated Database tab

This tab is used when monitoring MobiLink servers.

- **Connections In Use** Shows the number of database connections currently in use by this MobiLink server.
- **Longest Active Wait For Database Worker Thread** Shows the longest length of time an active request has been waiting for a database worker thread in this server.
- **Waiting For Database Worker Thread** Shows the number of requests in the server that are currently waiting for a database worker thread.
- **Number Of Upload Connections In Use** Shows the number of upload connections currently in use in the server.

## Monitoring tab: Machine Resources tab

This tab is used when monitoring MobiLink servers.

- **CPU Usage** Shows the percentage of CPU time used by the MobiLink server.
- **Total CPU Time** Shows the total amount of CPU time used by the MobiLink server in seconds.
- **MobiLink Cache Pages Used** Shows the percentage of MobiLink cache pages used by the server. This is set implicitly with the `-cm` option for `mlsrv11`. See [“-cm option” on page 55](#).
- **MobiLink Cache Pages Locked** Shows the percentage of MobiLink cache pages loaded into the server memory. This is set implicitly with the `-cm` option for `mlsrv11`. See [“-cm option” on page 55](#).
- **MobiLink Cache Pages In** Shows the number of MobiLink cache pages read from disk per second by the server.
- **MobiLink Cache Pages Out** Shows the number of MobiLink cache pages swapped to disk per second.
- **Memory Usage** Shows the bytes of RAM in use by the server (for Windows servers only).
- **Free Disk Space for MobiLink Cache** Shows the disk space available on the temp disk for MobiLink cache in bytes.
- **Open Network Connections** Shows the number of TCP connections currently open by the server.
- **Rejected Network Connections** Shows the total number of network connections rejected per second by the server.

## Delete old Monitor metrics

You can customize how long the Monitor keeps historical metrics. You can choose to use any or all of the settings. By default, the Monitor performs maintenance on itself once a day at midnight. Maintenance affects metrics, not alerts.



**To configure the deletion of historical metrics**

1. Click **Administration**.
2. Click the **Configuration** tab.
3. Click **Edit**.
4. Click **Maintenance**.
5. Specify a time when the Monitor should perform maintenance. By default, it performs maintenance at midnight. The time is local to the computer where the Monitor is running.
6. Customize the **Data Reduction** settings:
  - **Take A Daily Average Of Values Older Than** When you select this option, an average is taken for all numeric metrics that are older than the specified number of days, and then the numeric metrics are deleted. Non-numeric metrics are not deleted.
  - **Delete Values Older Than** When you select this option, all metrics that are older than the specified length of time are deleted.
  - **Delete Old Values When The Total Disk Space Used By The SQL Anywhere Monitor Becomes Greater Than X(MB)** When you select this option, you specify the maximum amount of space that can be used to store the metrics. When the amount of disk space used reaches or exceeds the amount specified, the Monitor deletes metrics, starting with the oldest metrics, preventing the Monitor from using more disk space for its metrics. Metrics are deleted until a sufficient amount of free space exists to store new metrics.
7. Click **Save**.

## Administering resources

A **resource** is a MobiLink server. You add resources to the Monitor, and then you start monitoring them.

The default resource, named **SQL Anywhere Monitor**, reports on the health of the Monitor itself. You cannot modify this resource, nor can you stop monitoring it.

### Start monitoring resources

When you start monitoring a resource, the Monitor starts collecting metrics.

Monitoring of a resource starts:

- Automatically when you add a resource. See [“Add resources” on page 222](#).
- Automatically when you start the Monitor. By default, all existing resources are started automatically when you start the Monitor.
- Automatically at the end of a blackout period. The Monitor automatically attempts to connect to the resource and resume monitoring.
- When an administrator opens the **Administration** tab, clicks **Resources**, selects a resource from the list, and clicks **Start**.

## Add resources

To monitor a MobiLink server, you must first add the resource to the Monitor.

When you add a MobiLink server as a resource to be monitored, the server is not modified in any way. When you add the resource, you supply a user ID and password to connect to the MobiLink server. The user ID and password are kept by the Monitor.

Only administrators can add resources. By default, resource monitoring starts when the resource is added.

### To add a resource to monitor

1. Log in to the Monitor.
2. Click the **Administration** tab.
3. On the **Resources** tab, click **Add**.
4. Follow the instructions in the **Add Resource** window to add a resource to monitor a MobiLink server.

When you add a MobiLink server, you must supply a user ID and password for the resource. These credentials are used to connect to the MobiLink server. The user ID and password are kept by the Monitor.

5. Click **Create**.

The resource is added and monitoring of the resource starts.

6. Click **OK**.

---

## Collection intervals

There are three types of collection intervals:

- **High collection interval** This rate is used for information that changes frequently, such as the completed synchronization rate.
- **Medium collection interval** This rate is used for information that changes less frequently, such as the amount of available disk space.
- **Low collection interval** This rate is used for information that changes infrequently, such as unsubmitted error reports.

Administrators can configure how often the Monitor collects a resource's metrics. Collection intervals are set per resource. You cannot configure the default resource, the SQL Anywhere Monitor.

### To edit the collection intervals

1. Click the **Administration** tab.
2. Click the **Resources** tab, and select a resource from the list.
3. Click **Configure**.
4. Click **Collection Intervals**.
5. Configure the other settings as required, and then click **Save**.
6. Click **OK**.

### See also

- [“Monitor metrics” on page 216](#)
- [“Specify metrics to collect” on page 223](#)

## Specify metrics to collect

Administrators can configure what metrics the Monitor collects and when alerts should be issued. You cannot configure the default resource, the SQL Anywhere Monitor.

### To configure what metrics are collected

1. Click the **Administration** tab.
2. Click the **Resources** tab, and select a resource from the list.
3. Click **Configure**.
4. Click **Metrics**. Select the metrics and alerts. For definitions of the metrics and alerts, see [“Types of metrics and alerts” on page 224](#).
5. Configure the other settings as required.

6. Click **Save**.
7. Click **OK**.

**See also**

- [“Monitor metrics” on page 216](#)
- [“Collection intervals” on page 223](#)

## Types of metrics and alerts

The following list describes the metrics that are available for your resource in the **Configure Resource** window: **Metrics** tab. Many of the default settings are arbitrary because each synchronization system has different behaviors and constraints, so the defaults may be inappropriate for your environment. You should carefully consider each metric and set each of them according to your needs.

- **CPU Usage (High Collection Interval)**
  - **Alert When CPU Usage Exceeds The Given Threshold For The Given Number Of Seconds** The **Threshold** default is 100 percent. The **Seconds** default is 300.
- **Memory Usage (Medium Collection Interval)**
  - **Alert When The Percentage Of Cache Pages Used Is Greater Than X Percent** The default is 100.
  - **Alert When The Percentage Of Locked Cache Pages Is Greater Than X** The default is 80.
  - **Alert When The Number Of Pages Being Swapped In And Out Per Second Exceeds The Given Threshold For The Given Number Of Seconds** The threshold default is 256. The seconds default is 120.
- **Network Usage (High Collection Interval)** Select this option to collect metrics about network usage in the server. You can view these metrics on the **Machine Resources** tab. See [“Monitoring tab: Machine Resources tab” on page 220](#).
- **Synchronizations (High Collection Interval)** Select this option to collect metrics about synchronizations in the server. You can view these metrics on the **Synchronization** tab. See [“Monitoring tab: Synchronization tab” on page 219](#).
  - **Alert When Longest Active Synchronization Time Is Greater Than X(Seconds)** The default is 600.
  - **Alert When The Number Of Failed Synchronizations Exceeds The Given Threshold For The Given Number Of Minutes** The threshold of failed synchronizations default is 20. The minutes default is 60.
- **Synchronization Throughput (High Collection Interval)** Select this option to collect metrics about synchronization throughput in the server. You can view these metrics on the **Synchronization** tab. See [“Monitoring tab: Synchronization tab” on page 219](#)
- **Error Rate (High Collection Interval)** Select this option to collect metrics about error rates in the server.

- **Alert When The Number Of Errors Exceeds The Given Threshold For The Given Number Of Minutes** The Threshold (Errors) default is 50. The Minutes default is 60.
- **Warning Rate (High Collection Interval)** Select this option to collect metrics about warning rates.
- **Database Connections In Use (High Collection Interval)** Select this option to collect metrics about the number of database connections in use in the server. You can view these metrics on the **Consolidated Database** tab. See [“Monitoring tab: Consolidated Database tab” on page 220](#).
- **Free Disk Space For MobiLink Cache (Medium Collection Interval)** Select this option to collect metrics on the disk space available for MobiLink cache on the server. You can view these metrics on the **Machine Resources** tab. See [“Monitoring tab: Machine Resources tab” on page 220](#).
  - **Alert When Free Disk Space For MobiLink Cache Is Less Than X (MB)** The default is 100.
- **Longest Active Wait For Database Worker Thread (High Collection Interval)** Select this option to collect metrics on the longest active wait time for database worker threads in the server.
  - **Alert When The Longest Active Wait For Database Worker Thread Is Greater Than X (Seconds)** The default is 300.
- **Longest Active Synchronization Time** Select this option to collect metrics on the longest active synchronization time in the server. You can view these metrics on the **Synchronization** tab. This metric should be configured for a high collection interval. See [“Monitoring tab: Synchronization tab” on page 219](#).
- **Suppress Alerts For The Same Condition That Occur Within Minutes** Select this option to prevent receiving duplicate alerts within a specified time. The default is 30 minutes.

## Stop monitoring resources

You stop monitoring resources when you do not want the Monitor to collect metrics from a MobiLink server. For example, you want to stop monitoring when you know that the resource will be unavailable; otherwise, you receive alerts until the resource is available. Except for the default Monitor resource, you can stop monitoring any resource at any time.

When you stop monitoring a resource, the Monitor:

- Stops collecting metrics for the resource.
- Stops issuing alerts for the resource.

There are two ways to stop monitoring a resource:

- **Schedule a regular, repeating, blackout period** This method is a good choice when the following conditions apply:
  - You must repeatedly stop monitoring the MobiLink server. For example, you perform regular maintenance at the end of each month.
  - You know in advance how long the MobiLink server is unavailable. For example, you know that your regular maintenance takes four hours.

- You need monitoring to automatically restart. When a blackout completes, the Monitor attempts to reconnect to the resource and to continue collecting data.

To use this method, you create blackouts to make the Monitor stop monitoring at specified times. See [“Automatically stop monitoring resources using blackouts” on page 226](#).

- **Manually stop the monitoring** This method is a good choice when the following conditions are met:
  - You need to stop monitoring for infrequent or one-time tasks. For example, you need to stop monitoring because the computer that the resource is running on needs to be taken off-line for special maintenance.
  - You are available to restart the monitoring afterwards. When a resource has been stopped manually, the Monitor waits for you to restart the monitoring.

To use this method, see [“Manually stop monitoring resources” on page 226](#).

If you want to permanently stop monitoring a resource, you can remove it from the Monitor. See [“Remove resources” on page 227](#).

## Manually stop monitoring resources

The following procedure describes how to manually stop a resource. For information about what happens when you stop a resource, see [“Stop monitoring resources” on page 225](#).

### To manually stop a resource

1. Click the **Administration** tab.
2. Select the resource to stop.
3. On the **Resources** tab, click **Stop**.

### See also

- [“Start monitoring resources” on page 222](#)
- [“Automatically stop monitoring resources using blackouts” on page 226](#)

## Automatically stop monitoring resources using blackouts

The following procedure describes how to stop a resource using blackouts. For information about what happens when you stop a resource and about when you should use blackouts, see [“Stop monitoring resources” on page 225](#).

Blackouts are times when you do not want the Monitor to collect metrics. When a blackout completes, the Monitor attempts to reconnect to the resources and to continue collecting data.

Blackouts occur in the local time of the resource.

### To configure the blackout time

1. Log in to the Monitor as an administrator.
2. Click the **Administration** tab.
3. On the **Resources** tab, select the resource you want to specify the blackout time for.
4. Click **Configure**.
5. Click the **Blackouts** tab.
6. Click **New**.
7. In the **New Blackout Period** window, specify the date and time for the blackout.  
The time is local to the computer where the resource MobiLink server resides.
8. Click **Save**.
9. Click **Save**.
10. Click **OK**.

### See also

- [“Start monitoring resources” on page 222](#)
- [“Manually stop monitoring resources” on page 226](#)

## Remove resources

You should only remove resources when you are certain that you don't need to monitor them; for example, if the server is no longer being used.

Removing a resource causes the Monitor to:

- Permanently stop monitoring the resource.
- Discard the metrics collected for the resource.

Only administrators can remove resources. You cannot delete the **SQL Anywhere Monitor** resource.

### To remove a resource

1. Click the **Administration** tab.
2. On the **Resources** tab, select a resource, and then click **Remove**.
3. Click **Yes**.

### See also

- [“Stop monitoring resources” on page 225](#)

## Working with Monitor users

The Monitor supports three types of users:

- **Read-only user** Has read-only access to monitor resources. Read-only users can view the metrics on the **Monitoring** tab, but cannot access the **Administration** tab. A user name and password are required.
- **Operator** Has read-only access to monitor resources and can receive alerts. These users can view the metrics on the **Monitoring** tab, can receive email alerts, and can resolve and delete alerts. However, operators cannot access the **Administration** tab. A user name and password are required.
- **Administrator** Has the same access as an operator, and can also configure resources and add users. Administrators can also access the **Administration** tab. The default user, **admin**, is an administrator. A user name and password are required.

The user name and password for logging in to the Monitor are case sensitive.

### Default user

By default, when you first start the Monitor, it has one administrator user, named **admin**, with password **admin**. By default, this user has full permissions. It is recommended that you change the default administrator password to restrict access to the Monitor. See [“Edit Monitor users” on page 229](#).

### Read-only access without a user name

By default, the Monitor does not require anyone to log in to have read-only access. However, for security and other reasons, the administrator can require that users log in. See [“Require Monitor users to login” on page 230](#).

## Create Monitor users

You must be an administrator to add Monitor users.

### To add a new Monitor user

1. Click the **Administration** tab.
2. Click the **Users** tab.
3. Click **New**.
4. Fill in the information for the new user. An email address is only required for users who should receive email alerts from the Monitor.  
  
Click **Save**.
5. If you create an operator or an administrator, you can associate the user with a resource. See [“Associate Monitor users with resources” on page 229](#).



**See also**

- [“Edit Monitor users” \[SQL Anywhere Server - Database Administration\]](#)

## Associate Monitor users with resources

You must associate a user with a resource if you want the user to receive email alerts about the associated resource. You can only associate an operator or an administrator with a resource.

**To associate an operator or administrator with a resource**

1. Click the **Administration** tab.
2. Click the **Resources** tab.
3. Select the resource and click **Configure**.
4. Click **Operators**.
5. From the **Available Operators** list, select the user and click **Add**.
6. Click **Save**.
7. Click **OK**.
8. Verify that the Monitor is set up to send alert notifications by email. See [“Send alert emails” on page 233](#).

**See also**

- [“Working with Monitor users” on page 228](#)

## Edit Monitor users

As an administrator, you can edit Monitor users to change their:

- Passwords
- Email addresses
- Language settings
- User types

**To edit an existing Monitor user**

1. Click the **Administration** tab.
2. Click the **Users** tab.
3. Select the user to edit.
4. Click **Edit**.

5. Change the settings for the user as required.
6. Click **Save**.
7. If you are editing an operator or an administrator, you can associate the user with a resource. See [“Associate Monitor users with resources” on page 229](#).

**See also**

- [“Working with Monitor users” on page 228](#)
- [“Create Monitor users” on page 228](#)
- [“Delete Monitor users” on page 230](#)

## Delete Monitor users

Deleting a user removes the user from the Monitor and disassociates the user from any resource.

You must be an administrator to delete Monitor users.

**To delete an existing Monitor user**

1. Click the **Administration** tab.
2. Click the **Users** tab.
3. Select the user to delete.
4. Click **Delete**.
5. Click **Yes** to delete the selected user. Click **Delete All** to delete all users.

The user is deleted from the Monitor.

**See also**

- [“Create Monitor users” on page 228](#)
- [“Edit Monitor users” on page 229](#)
- [“Associate Monitor users with resources” on page 229](#)

## Require Monitor users to login

By default, anyone can have read-only access to the Monitor. You can change this behavior so that whenever a user opens the Monitor in a web browser, they must provide a user name and password before they can see any monitoring data.

**To restrict access to the Monitor**

1. Click the **Administration** tab.
2. On the **Configuration** tab, click **Edit**.
3. Click **Authentication**.

4. Clear the **Allow Anyone Read-only Access To The SQL Anywhere Monitor** option.
5. Click **Save**.

**See also**

- [“Create Monitor users” on page 228](#)
- [“Edit Monitor users” on page 229](#)

## Alerts

An **alert** is a condition or state of interest that should be brought to an administrator's or operator's attention. Alerts include information about the cause of the problem, and provide advice for resolving the problem.

There are several predefined alerts for conditions such as low disk space, critical software updates, failed login attempts, and high memory usage. When an alert condition is met, the alert is listed in the bottom pane on the **Monitoring** tab. In the top pane, the MobiLink server **Status** changes to indicate that an alert exists. You can configure the Monitor to send an email to operators and administrators when an alert occurs. See [“Send alert emails” on page 233](#).

Alerts are detected by the Monitor based on metrics that are collected. They are not detected at the MobiLink server being monitored. You can change the default threshold values and choose which alerts are enabled by editing the resource. See [“Monitor metrics” on page 216](#).

## View alerts

Any user can view alerts; however, only operators and administrators can resolve and delete alerts.

### To view an alert

1. Click the **Monitoring** tab.
2. Select a resource from the list.
3. In the bottom pane, click the **Alerts** tab.
4. Select a row in the alerts list.
5. Click **Details**.
6. Click **OK**.

### See also

- [“Resolve alerts” on page 232](#)
- [“Delete alerts” on page 233](#)
- [“Send alert emails” on page 233](#)

## Resolve alerts

Once the issue that triggered an alert has been addressed, you can mark an alert as resolved. Resolving an alert causes the Monitor to change the alert's status column, but leave the alert in the alert list. If you want to remove the alert, you must delete it. See [“Delete alerts” on page 233](#).

Only operators and administrators can resolve alerts.

### To resolve an alert

1. Click the **Monitoring** tab.

2. Select a resource from the list.
3. In the bottom pane, click the **Alerts** tab.
4. Select the row in the alerts list.
5. Click **Mark Resolved** to resolve the selected alert. Click **Mark All Resolved** to resolve all alerts in the list.

The value in the **Status** column on the **Alerts** tab changes to **Resolved**.

If this was the resource's only unresolved alert, the resource's status changes to **Healthy**.

#### See also

- [“Delete alerts” on page 233](#)
- [“Resolve alerts” on page 232](#)
- [“Send alert emails” on page 233](#)
- [“View alerts” on page 232](#)
- [“Alerts” on page 232](#)

## Delete alerts

The Monitor keeps only the most recent 50 alerts in the alert list. If you do not want an alert to appear in the alerts list any more, you can delete the alert. You can delete alerts, regardless of their status.

Only operators and administrators can delete alerts.

#### To delete alerts

1. Click the **Monitoring** tab.
2. Select a resource from the list.
3. In the bottom pane, click the **Alerts** tab.
4. Select a row in the alerts list.
5. Click **Delete**.

The alert is removed from the alerts list.

#### See also

- [“Resolve alerts” on page 232](#)
- [“Send alert emails” on page 233](#)
- [“View alerts” on page 232](#)
- [“Alerts” on page 232](#)

## Send alert emails

You can configure the Monitor to send an email to operators and administrators when an alert occurs.

To have the Monitor send alert notifications by email, you must:

1. Create an administrator or operator with an email address. See [“Create Monitor users” on page 228](#).
2. Associate the administrator or operator with a resource. See [“Associate Monitor users with resources” on page 229](#).
3. Enable the Monitor to send emails. See [“Enable the Monitor to send alert emails” on page 234](#).

## Enable the Monitor to send alert emails

As an administrator, you can configure the Monitor to send emails when an alert occurs. The Monitor supports the SMTP and MAPI protocols for sending emails.

### To enable the Monitor to send alert notifications by email

1. Click the **Administration** tab.
2. Click the **Configuration** tab.
3. Click **Edit**.
4. Click **Alert Notification**.
5. Select **Send Alert Notifications By Email**.
6. Choose either SMTP or MAPI for the **Which Protocol Do You Want To Use To Send Alerts By Email?** field.
7. Configure the other settings as required.
  - **MAPI**
    - **User Name** Type the user name for the MAPI server.
    - **Password** Type the password for the MAPI server.
  - **SMTP**
    - **Server** Specify which SMTP server to use. Type the server name or the IP address for the SMTP server. For example, *SMTP.yourcompany.com*.
    - **Port** Specify the port number to connect to on the SMTP server. The default is 25.
    - **Sender Name** Specify an alias for the sender's email address. For example, *JoeSmith*.
    - **Sender Address** Specify the email address of the sender. For example, *jsmith@emailaddress.com*.
    - **This SMTP Server Requires Authentication** Select this option if your SMTP server requires authentication.
      - **User Name** Specify the user name to provide to SMTP servers requiring authentication.
      - **Password** Specify the password to provide to SMTP servers requiring authentication.
8. Test that you have properly configured email notification.

---

Click **Send Test Email**.

9. When prompted, enter an email address to send the test email to and click **OK**.

A test email is sent to the email address specified.

10. Click **Save**.

#### See also

- [“Resolve alerts” on page 232](#)
- [“Delete alerts” on page 233](#)
- [“View alerts” on page 232](#)

## Suppress alerts for unsubmitted error reports from resources

As an administrator, you can configure whether the Monitor sends out alerts when resources have unsubmitted error reports. By default, the Monitor does not send these alerts. For information about error reports and about how to submit them, see [“Error reporting in SQL Anywhere” \[SQL Anywhere Server - Database Administration\]](#).

#### To suppress alerts for unsubmitted error reports

1. Click the **Administration** tab
2. Click the **Configuration** tab.
3. Click **Edit**.
4. Click **Options**.
5. Click **Save**.

## Installing the SQL Anywhere Monitor on a separate computer

These instructions explain how to install SQL Anywhere Monitor on a separate computer than the one that SQL Anywhere is running on.

Some advantages to running the SQL Anywhere Monitor on a separate computer include:

- The Monitor runs in the background as a service.
- The Monitor starts automatically when the computer starts.
- Upgrades and updates of SQL Anywhere do not overwrite the Monitor when it is installed on a separate computer. This is important if the separate computer is in a production environment.

### To install the Monitor on a separate computer

- Run the *setup.exe* file from the *Monitor* directory on your installation media, and follow the instructions provided.



## Troubleshooting the Monitor

Problem	Recommendation
When you press F5 to refresh the web browser window, you are required to log in to the Monitor.	Enable JavaScript in your web browser.
You receive a network communication error when you try to connect to the Monitor.	Start the Monitor. See <a href="#">“Start the Monitor” on page 211</a> .
After upgrading to the latest version of Adobe Flash Player you continue to receive instructions to upgrade Adobe Flash Player.	Verify that the installed version Adobe Flash Player is supported by your operating system. The Monitor is backwards compatible with version 9 of Adobe Flash Player. To determine the correct version, visit: <a href="http://www.adobe.com/products/flashplayer/systemreqs/">http://www.adobe.com/products/flashplayer/systemreqs/</a> .
The Monitor is unable to start monitoring a SQL Anywhere database resource.	Verify that the resource's password verification functions and login procedures allow the user sa_monitor_user to connect to the resource.
You are not receiving any alert emails.	<p>Verify that the Monitor is properly configured to send emails and send a test email. See <a href="#">“Enable the Monitor to send alert emails” on page 234</a>.</p> <p>Verify that the alert emails from the Monitor are not being blocked by a virus scanner. See <a href="#">“xp_startsmtp system procedure” [SQL Anywhere Server - SQL Reference]</a>.</p>
The number of unscheduled requests reported by the Monitor appears to be less than the actual number of unscheduled requests.	<p>When collecting metrics about the number of unscheduled requests, the Monitor executes query on the resource. This query could be an unscheduled request.</p> <p>Unscheduled queries are processed sequentially as they arrive. Therefore, if there are unscheduled requests when the Monitor attempts to execute its query, then this query must wait for the existing unscheduled requests to complete before it can execute.</p> <p>As a result, when the Monitor collects the number of unscheduled requests, this number does not include the unscheduled requests that existed between the time when the Monitor issued its query and the query executed.</p>

Problem	Recommendation
<p>You are not receiving alerts when the database disk space surpasses the specified threshold.</p>	<p>Between Monitor collection intervals, it is possible for a database to exceed the specified disk space alert threshold and the amount of space available. In such a case, the database would stop responding before the Monitor could collect the disk usage metrics and issue an alert</p> <p>If your database grows quickly, set the disk space alert threshold to a higher number so that you can receive an alert before the database runs out of space. See <a href="#">“Types of metrics and alerts” on page 224</a>.</p>
<p>When you open the Monitor in a Firefox web browser from a non-English computer, the Monitor appears in English.</p>	<p>Firefox does not correctly use your computer's preferred locale. You can use Internet Explorer or try the following Firefox workaround:</p> <ol style="list-style-type: none"> <li>1. In Firefox, open a new tab.</li> <li>2. In the address bar, type the following: <ul style="list-style-type: none"> <li><code>about:config</code></li> </ul> <p>Press <b>Enter</b>.</p> <p>If prompted, click <b>I'll Be careful, I Promise!</b></p> </li> <li>3. In the <b>Filter</b> field, type the following: <ul style="list-style-type: none"> <li><code>general.useragent.locale</code></li> </ul> </li> <li>4. In the preference list, double-click <b>general.user-agent.locale</b>.</li> <li>5. In the <b>Enter String Value</b> window, enter your locale. For example, type fr-FR for French, de-DE for German, zh-CN for Chinese, and ja-JP for Japanese.</li> <li>6. Click <b>OK</b>.</li> </ol>

---

# The Relay Server

## Contents

Introduction to the Relay Server .....	240
Relay Server configuration file .....	243
Outbound Enabler .....	247
Relay Server State Manager .....	250
Deploying the Relay Server .....	253
Updating a Relay Server farm configuration .....	258
Sybase Relay Server hosting service .....	260
Using MobiLink with the Relay Server .....	262

---

## Introduction to the Relay Server

The Relay Server enables secure, load-balanced communication between mobile devices and MobiLink, Afaria and iAnywhere Mobile Office servers through a web server. The Relay Server provides the following:

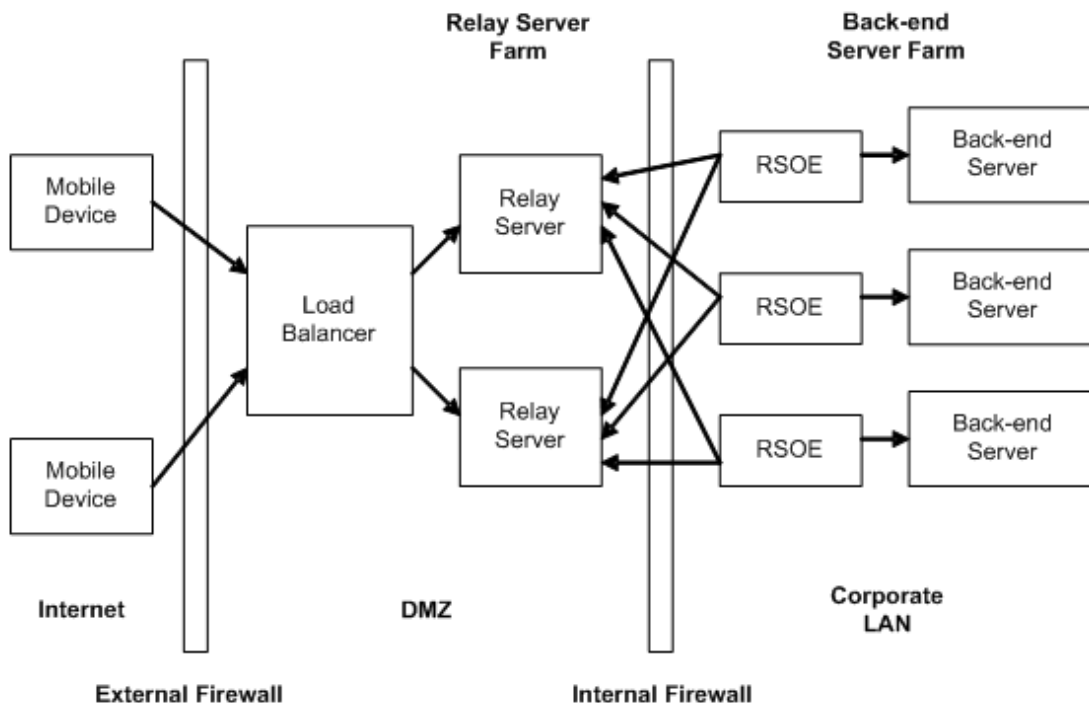
- A common communication architecture for mobile devices communicating with MobiLink, Afaria and iAnywhere Mobile Office servers.
- A mechanism to enable a load-balanced and fault-tolerant environment for MobiLink, Afaria and iAnywhere Mobile Office servers.
- A way to help communication between mobile devices and MobiLink, Afaria and iAnywhere Mobile Office servers in a way that integrates easily with existing corporate firewall configurations and policies.

## Relay Server architecture

A Relay Server deployment consists of the following:

- Mobile devices running client applications and services that need to communicate with back-end servers running in a corporate LAN.
- Optional load balancer to direct requests from the mobile devices to a group of Relay Servers.
- One or more Relay Servers running in the corporate DMZ.
- Back-end servers running in a corporate LAN that are responsible for servicing client requests.
- One Relay Server Outbound Enabler (RSOE) per back-end server. The Outbound Enabler manages all communication between a back-end server and the Relay Server farm.

The following diagram shows the Relay Server architecture.



The Relay Server consists of a set of web extensions, a background process for maintaining state information, and a web server.

Because the Relay Server is a web extension running in a web server, all communication is performed using HTTP or HTTPS. Using HTTP easily integrates with existing corporate firewall configurations and policies. The Relay Server requires that the connection from the corporate LAN to the Relay Server be initiated from inside the corporate LAN. This provides a more secure deployment environment because it does not require inbound connections from the DMZ into the corporate LAN.

The Relay Server contains two web extensions: a client extension and a server extension. The client extension handles client requests made from applications running on mobile devices. The server extension handles requests made by the Outbound Enabler on behalf of a back-end server.

## The Relay Server farm

A Relay Server farm is any number of Relay Servers with a front-end load balancer. It is possible to set up a Relay Server farm with a single Relay Server, in which case a load balancer is not required. In this case, mobile devices can connect directly to the Relay Server.

## Back-end server farm

A back-end server farm is a group of homogeneous back-end servers. A client making a request through the Relay Server farm must specify the back-end server farm it is targeting.

## Load balancer

The load balancer directs requests from the mobile devices to a Relay Server running in the Relay Server farm. The load balancer is not required if there is only one Relay Server.

## Relay Server Outbound Enabler

The Relay Server Outbound Enabler runs on the same computer as the back-end server. Its primary function is to initiate an outbound connection to all Relay Servers in the Relay Server farm on behalf of the back-end server. There is one Outbound Enabler per back-end server. See [“Outbound Enabler” on page 247](#).

---

## Relay Server configuration file

A Relay Server configuration file is used to define both a Relay Server farm and the back-end server farms connecting to the Relay Server farm. The Relay Server configuration file is divided into sections:

- [“Relay Server section” on page 243](#)
- [“Backend farm section” on page 244](#)
- [“Backend server section” on page 245](#)
- [“Options section” on page 245](#)

Each section starts with a section tag. A section tag is formed by enclosing a keyword that identifies the section name in square brackets. For example, `[relay_server]` denotes the start of the Relay Server section.

The section tag is followed by several lines defining various properties related to the section being defined. A property is defined by specifying the property name on the left-hand side of an equal sign and its value on the right-hand side of the equal sign. For example, `property name = value`. All section and property names are case insensitive. Comments are marked with pound sign (#) character at the beginning of a line.

The configuration file should contain only 7-bit ASCII characters. The sections can be specified in any order.

## Relay Server section

The Relay Server section is used to define a single Relay Server, so there must be a Relay Server section for each Relay Server in the farm. This section is identified by the `relay_server` keyword.

### Relay Server section properties

The following properties can be specified in a Relay Server section:

- **enable** Specifies whether this Relay Server is to be included in the Relay Server farm. Possible values are:
  - **Yes** Indicates that this Relay Server is to be included in the Relay Server farm.
  - **No** Indicates that this Relay Server should not be included in the Relay Server farm.

The default is Yes. This property is optional.

- **host** The hostname or IP address that should be used by the Outbound Enabler to make a direct connection to the Relay Server.
- **http\_port** The HTTP port that should be used by the Outbound Enabler to make a direct connection to the Relay Server. A value of **0** or **off** disables HTTP connections. By default, this property is enabled and set to 80.
  - **0 or off** Disable HTTP access from Outbound Enabler.
  - **1 to 65535** Enable HTTP at the specified port.

- **https\_port** The HTTPS port that should be used by the Outbound Enabler to make a direct connection to the Relay Server. A value of **0** or **off** disables HTTPS connections. By default, this property is enabled and set to 443.
  - **0 or off** Disable HTTPS access from Outbound Enabler.
  - **1 to 65535** Enable HTTPS at the specified port.
- **description** Enter a custom description to a maximum of 2048 characters. This property is optional.

## Backend farm section

The backend farm section specifies the properties of a back-end server farm. A back-end server farm is a group of homogenous back-end servers. A client making a request through the Relay Server farm must specify the back-end server farm it is targeting. There is one backend farm section for each back-end server farm.

This section is identified by the `backend_farm` keyword.

### Backend farm section properties

The following properties can be specified in a backend farm section:

- **enable** Specifies whether to allow connections from this back-end server farm. Possible values are:
  - **Yes** Allow connections from this back-end server farm.
  - **No** Disallow connections from this back-end server farm.

The default is Yes. This property is optional.

- **id** The name assigned to the back-end server farm, to a maximum of 2048 characters.
- **client\_security** Specifies the level of security the back-end server farm requires of its clients. The possible values are:
  - **on** Indicates that clients must connect using HTTPS.
  - **off** Indicates that clients must connect using HTTP.

This property is optional. If no value is specified, clients can connect using either HTTP or HTTPS.

- **backend\_security** Specifies the level of security required of an Outbound Enabler in the back-end server farm to connect to the Relay Server farm. The possible values are:
  - **on** Indicates that all connections from the back-end farm must be made using HTTPS.
  - **off** Indicates that all connections from the back-end farm must be made using HTTP.

This property is optional. If no value is specified, either HTTP or HTTPS can be used to connect.

- **description** Enter a custom description to a maximum of 2048 characters. This property is optional.



---

## Backend server section

The backend server section defines a back-end server connection. It specifies the information that is used by the Outbound Enabler when it connects to the Relay Server farm on behalf of a back-end server. There is a backend server section for each Outbound Enabler connecting to the Relay Server farm. The backend server section also assigns a back-end server to a back-end server farm.

This section is identified by the `backend_server` keyword.

### Backend server section properties

The following properties can be specified in a backend server section:

- **enable** Specifies whether to allow connections from this back-end server. Possible values are:
  - **Yes** Allows connections from this back-end server.
  - **No** Disallows connections from this back-end server.

The default is Yes. This property is optional.

- **id** The name assigned to the back-end server connection, to a maximum of 2048 characters.
- **farm** The name of the back-end server farm that this back-end server belongs to.
- **MAC** The MAC address of the network adapter used by the Outbound Enabler to communicate with the Relay Server. The address is specified using the IEEE 802 MAC-48 format. To get the MAC address in the correct format, look in the Relay Server Outbound Enabler console or log. This property is optional. If it is not specified, MAC address checking does not occur.
- **token** A security token that is used by the Relay Server to authenticate the back-end server connection, to a maximum of 2048 characters. This property is optional.
- **description** Enter a custom description to a maximum of 2048 characters. This property is optional.

## Options section

The options section is used to specify properties that apply to each Relay Server in the farm. Only one options section is allowed.

This section is identified by the `options` keyword.

### Options section properties

The following properties can be specified in an options section:

- **start** The method used to start the State Manager. The possible values are:
  - **auto** The State Manager is started automatically using the State Manager command line defaults.
  - **no** The State Manager is started externally as a Windows service.
  - **full path** Specify the full path to the State Manager executable (*rshost*).

The default is auto. This property is optional.

- **shared\_mem** Specifies the maximum amount of shared memory that the Relay Server uses for state tracking. The default is 10 megabytes. This property is optional.
- **verbosity** You can set verbosity to the following levels:
  - **0** Log errors only. Use this logging level for deployment. This is the default.
  - **1** Request level logging. All HTTP requests are written to the log file.

Errors are displayed regardless of the log level specified, and warnings are displayed only if the log level is greater than 0.

## Relay Server configuration file format

This is the basic format of a Relay Server configuration file:

```
#
# Options
#
[options]
# List of Relay Server properties that apply to all Relay Servers
option = value

#
# Define a Relay Server section, one for each
# Relay Server in the Relay Server farm
#
[relay_server]
# List of properties for the Relay Server
property = value

#
# Define a backend server farm section, one for each backend
# server farm
#
[backend_farm]
# List of properties for a backend server farm
property = value

#
# Define a backend server section, one for each
# Outbound Enabler connecting to the Relay Server farm
#
[backend_server]
# List of properties for the backend server connection
property = value
```

## Outbound Enabler

The Outbound Enabler runs on the same computer as the back-end server. Its purpose is to:

- Open an outbound connection from the computer running in the corporate LAN to the Relay Server farm running in the DMZ.
- Forward client requests received from the Relay Server to the back-end server and forward back-end server responses back to the client via the Relay Server.

When the Outbound Enabler starts, it makes an HTTP request to retrieve the list of Relay Servers running in the farm. This is done using the server URL that maps to the web server extension component of the Relay Server. The server URL can map directly to a Relay Server or it can map to a load balancer. If the server URL maps to a load balancer, the load balancer forwards the request to one of the Relay Servers running in the farm. The Relay Server that receives the request from the Outbound Enabler returns the connection information for all Relay Servers in the farm. The Outbound Enabler then creates two outbound connections, called channels, to each Relay Server returned. One channel, called the up channel, is created using an HTTP request with an essentially infinite response. The response is a continuous stream of client requests from the Relay Server to the Outbound Enabler. The second channel, called the down channel, is created using an HTTP request with an essentially infinite content length. The request is formed by a continuous stream of server responses to client requests.

When the Outbound Enabler receives a client request on the up channel from one of the Relay Servers it has connected to, it forwards it to the back-end server that the Outbound Enabler is servicing. Once a response is received from the back-end server, it gets forwarded to the Relay Server from which it received the corresponding request using the down channel.

### Outbound Enabler syntax

```
rsoe [option]+
```

```
rsoe @{ filename | environment-variable } ...
```

### Parameters

**Options** The following options can be used with the Outbound Enabler. They are all optional.

rsoe options	Description
@data	Reads options from the specified environment variable or configuration file. 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 <a href="#">“File Hiding utility (dbfhide)” on page 248</a> .
-f farm	The name of the farm that the back-end server belongs to.
-id id	The name assigned to the back-end server.
-cs "connection-string"	The host and port used to connect to the back-end server. The default is "host=localhost;port=80".

rsoe options	Description
<b>-cr</b> <i>"connection-string"</i>	<p>The Relay Server connection string. The format of the Relay Server connection string is a semicolon separated list of name-value pairs. The name-value pairs consist of the following:</p> <ul style="list-style-type: none"> <li>● <b>host</b> IP address or hostname of the Relay Server. The default is localhost.</li> <li>● <b>port</b> The port the Relay Server is listening on. This is required.</li> <li>● <b>url_suffix</b> URL path to the server extension of the Relay Server. The default for Windows is <i>/ias_relay_server/server/rs_server.dll</i>. The default for Linux is <i>/srv/iarelayserver</i>.</li> <li>● <b>https</b> 0 - HTTP (default) 1 - HTTPS</li> </ul> <p>For <b>https=1</b>, the following options can also be specified:</p> <ul style="list-style-type: none"> <li>● <b>tls_type</b> RSA</li> <li>● <b>certificate_name</b> Common name field of the certificate.</li> <li>● <b>certificate_company</b> Organization name field of the certificate.</li> <li>● <b>certificate_unit</b> Organization unit field of the certificate.</li> <li>● <b>trusted_certificates</b> File containing a list of trusted root certificates.</li> </ul>
<b>-t</b> <i>token</i>	The security token to be passed to the Relay Server.
<b>-v</b> <i>level</i>	<p>Set the verbosity level to use for logging. The <i>level</i> can be <b>0</b>, <b>1</b>, or <b>2</b>:</p> <ul style="list-style-type: none"> <li>● <b>0</b> Log errors only. Use this logging level for deployment.</li> <li>● <b>1</b> Session level logging. This is a higher level view of a synchronization session.</li> <li>● <b>2</b> Request level logging. Provides a more detailed view of HTTP requests within a synchronization session.</li> </ul>
<b>-d</b> <i>seconds</i>	The Relay Server connection retry interval. The default is 5 seconds.
<b>-s</b>	Stop the Outbound Enabler.

### File Hiding utility (dbfhide)

The File Hiding utility (dbfhide) uses simple encryption to obfuscate the contents of configuration files and initialization files.

## Syntax

**dbfhide** *original-configuration-file encrypted-configuration-file*

Option	Description
<i>original-configuration-file</i>	Specifies the name of the original file.
<i>encrypted-configuration-file</i>	Specifies a name for the new obfuscated file.

The Relay Server and Outbound Enabler detect that a configuration file has been obfuscated using dbfhide and process it accordingly.

This utility does not accept the @data parameter to read in options from a configuration file.

## Deployment considerations

The following considerations should be noted when using the Outbound Enabler:

- **Outbound Enabler as a Windows service** The Outbound Enabler may also be set up and maintained as a Windows service using the Service Utility. See [“Relay Server State Manager” on page 250](#).
- **Authentication** You cannot use simple or digest authentication. The *rsoe.exe* does not support simple or digest authentication with web servers, regardless of the web server type or operating system.

## Relay Server State Manager

The Relay Server State Manager is a process that is responsible for maintaining Relay Server state information across client requests and Outbound Enabler sessions. The State Manager is also responsible for managing the log file used by the Relay Server. The State Manager can either be started automatically by the Relay Server or started as a Windows service (on Windows only).

The default log file name is *ias\_relay\_server\_host.log*. On Windows, this file is located in the directory specified by the TEMP environment variable. On Linux, the file is located in the directory specified by the TMP, TEMP, or TMPDIR environment variables. If none of those variables are set, a log file is created on the root.

**Note**

In all cases, the Apache user process must have write permissions to the *tmp* directory location you choose.

On a graceful shutdown, the State Manager renames the log file to a file of the form *<yymmdd><nn>.log* where *<yymmdd>* represents the date on which the log file was renamed and *<nn>* is the sequential version number of the log file for that day.

Starting the State Manager as a Windows service is the recommended method. Note that starting the State Manager manually on a command line is not supported.

It is possible to specify the options that are used by the Relay Server to start the State Manager. To change the options, set the **start** property in the options section of the Relay Server configuration file. For example:

```
[options]
start = "rshost -o c:\temp\myrshost.log"
```

Note that you must specify the name of the Relay Server State Manager executable (*rshost*) before the options.

## Starting the Relay Server State Manager as a Windows service

For Windows only, the State Manager can be started as a Windows service by using the Service utility *dbsvc.exe*. The start property in the options section of the Relay Server configuration file should be set to **no**. See [“Options section” on page 245](#).

The Service utility (*dbsvc*) is used to create, modify and delete services. For a full listing of usage information, run *dbsvc.exe* without any options.

### To set up an auto started State Manager service named rs:

```
dbsvc -as -s auto -w rs "C:\inetpub\wwwroot\ias_relay_server\server
\rshost.exe" -q -qc -f c:\inetpub\wwwroot\ias_relay_server\server
\rstest.config -o c:\temp\rs.log
```

### To start the service:

```
dbsvc.exe -u rs
```

**To stop the service:**

```
dbsvc.exe -x rs
```

**To uninstall the service:**

```
dbsvc.exe -d rs
```

## Starting the Relay Server State Manager automatically

The State Manager process is started automatically when the first Outbound Enabler connects to the Relay Server. This is the default behavior when the start property in the options section of the Relay Server configuration file is not specified or is explicitly specified as auto. The default log file location is *%temp%\ias\_relay\_server\_host.log*. See [“Options section” on page 245](#).

## Starting the Relay Server State Manager automatically with customized options

When auto start is desired but you want to override some default behavior such as verbosity level or log file location, you can use the start property in the options section of the Relay Server configuration file to explicitly specify your State Manager command line. The -f option cannot be used in this case and the configuration file must be named *rs.config* and be placed in the same directory as the server extension. See [“Relay Server State Manager command line syntax” on page 251](#).

**Note**

Do not specify a log file location under the wwwroot directory. IIS does not allow a worker process to create a file under the published tree.

## Relay Server State Manager command line syntax

```
rshost [option]+
```

**Parameters**

**Options** The following options can be used to configure the State Manager. They are all optional.

rshost options	Description
-f <i>filename</i>	File name of the Relay Server configuration file.
-o <i>filename</i>	File name to use for logging.
-oq	Prevent popup window on startup error.
-q	Run in minimized window.

<b>rshost options</b>	<b>Description</b>
<b>-qc</b>	Close window on completion.
<b>-u</b>	Update configuration of a running Relay Server.
<b>-ua</b>	Archive the log file to <yymmdd><nn>.log and truncate.



## Deploying the Relay Server

The following is an overview of how to deploy the Relay Server to IIS on Windows:

1. Deploy the Relay Server components. See [“Deploying the Relay Server components to IIS on Windows” on page 253](#).
2. Deploy the web server extensions and State Manager. See [“Deploying the web server extensions and State Manager” on page 254](#).
  - a. Create an application pool. See [“Creating an application pool” on page 254](#).
  - b. Enable the Relay Server web extensions and deploy the Relay Server configuration file. See [“Deploying the web server extensions and State Manager” on page 254](#).
3. Deploy Relay Server configuration updates, as necessary. See [“Updating a Relay Server configuration for IIS on Windows” on page 258](#).

The following is an overview of how to deploy the Relay Server to Apache on Linux:

1. Deploy the Relay Server components. See [“Deploying the Relay Server components to Apache on Linux” on page 256](#).
2. Deploy the web extension files and State Manager. See [“Deploying the web extension files and State Manager” on page 256](#).
3. Deploy Relay Server configuration updates, as necessary. See [“Updating a Relay Server configuration for Apache on Linux” on page 259](#).

## Deploying the Relay Server components to IIS on Windows

The Relay Server for Windows consists of the following executables:

- *rs\_client.dll*
- *rs\_server.dll*
- *rshost.exe*
- *dbngen11.dll*
- *dbsvc.exe*
- *dbfhide.exe*
- *dbicu11.dll*
- *dbicudt11.dll*
- *dbsupport.exe*
- *dbghelp.dll*

### See also

- [“Relay Server State Manager” on page 250](#)
- [“File Hiding utility \(dbfhide\)” on page 248](#)

## Deploying the web server extensions and State Manager

### To deploy the Relay Server files

1. Create the following directories under the web site home directory that you use for the Relay Server:
  - *ias\_relay\_server*
  - *ias\_relay\_server\client*
  - *ias\_relay\_server\server*
2. Copy *rs\_client.dll* to the *ias\_relay\_server\client* directory.
3. Create the Relay Server configuration file *rs.config*. See [“Relay Server configuration file” on page 243](#).
4. Copy *rs\_server.dll*, *rshost.exe* and *rs.config* to the *ias\_relay\_server\server* directory.
5. Ensure that the connection timeout property of the web site home directory is set to 60 seconds or more.

## Creating an application pool

A dedicated application pool must be created for the *rs\_server.dll* and *rs\_client.dll* web server extensions. All worker recycling options need to be turned off as the Relay Server utilizes long running worker processes.

### To create the application pool

1. Start IIS Manager Console.
2. Right-click **Application Pools** and create a new application pool, for example RS\_POOL.
3. Edit the properties for the application pool you created:
  - a. Select the **Recycling** tab and turn off all the recycling options.
  - b. Select the **Performance** tab and do the following:
    - i. Turn off **Shutdown Worker Processes After Being Idle**.
    - ii. Set the number of worker processes to the total number of processing cores. You can further adjust this number depending on your usage and performance preferences. See the IIS performance notes about Web garden size for more information.

## Enabling the Relay Server web extensions

The steps below describe the process to enable the Relay Server web extensions.

### To edit the properties of *ias\_relay\_server* and enable the Relay Server web extensions

1. Select the **Directory** tab and do the following:
  - a. Set execute permissions to **Scripts And Executables**.

- b. Click **Create** under **Application Settings**. Select the application pool you created in [“Creating an application pool” on page 254](#) as the associated application pool.
2. Select the **Directory Security** tab and do the following:
  - a. Click **Edit** in **Authentication and Access Control**.
  - b. Enable anonymous access and fill in the user name and password for an account belonging to the Administrators group.

Alternatively, you may leave the setting as the built-in user **IUSR\_%computername%** and execute the following command to grant permission to access the IIS metabase.

```
C:\Windows\Microsoft.Net\Framework\<Version>\aspnet_regiis.exe -ga IUSR_%computername%
```
3. Under **Web Server Extensions** in the IIS manager, allow both *rs\_server.dll* and *rs\_client.dll* to be run as ISAPI.
4. Deploy the Relay Server configuration file by creating a Relay Server configuration file and copying it to the *ias\_relay\_server\server* directory.

#### See also

- [“Relay Server configuration file” on page 243](#)

## Performance tips

Keep the following in mind when deploying the Relay Server to IIS on Windows:

- The Relay Server web extension does not rely on ASP.NET. Removing the ASP.NET ISAPI filter yields better performance. The filter gets turned on by default in a standard IIS install. To turn off the filter, do the following:
  1. Start IIS Manager Console.
  2. Edit the properties of **Default Web Site**.
  3. Under the **ISAPI Filters** tab, remove the ASP.NET filter.
- For better performance, you can turn off the IIS access log. To turn off the access log, do the following:
  1. Start IIS Manager Console.
  2. Edit the properties of the *ias\_relay\_server* directory under **Default Web Site**.
  3. Under the **Directory** tab, clear the **Log Visits** selection.
- In a production environment, Relay Server verbosity can be set to 0 via the Relay Server configuration file. This yields better performance under high loads.
- The Relay Server does not impose restrictions on the Web garden size. One worker process may serve requests from all Outbound Enablers as well as from all the clients. However, the number of threads that can be created in the process is limited by the process heap space left available for thread creation. The thread created by IIS has a 256k stack size. If your machine has adequate resources, experiment with a

higher number of processes if you suspect you are hitting a concurrency limit when the server is loaded with thousands of concurrent requests.

## Deploying the Relay Server components to Apache on Linux

The Relay Server for Linux consists of the following executables:

- *mod\_rs\_ap\_client.so*
- *mod\_rs\_ap\_server.so*
- *rshost*
- *dblgen11.res*
- *libdbtasks11\_r.so*
- *libdbicudt11.so*
- *libdbicu11\_r.so*
- *libdblib11\_r.so*
- *dbsupport*
- *dbfhide*

### See also

- [“Relay Server State Manager” on page 250](#)
- [“File Hiding utility \(dbfhide\)” on page 248](#)

## Deploying the web extension files and State Manager

### To deploy the Relay Server files

1. Copy the above files into your Apache install *modules* directory.
2. Create the Relay Server configuration file *rs.config*. See [“Relay Server configuration file” on page 243](#).
3. Copy *rs.config* into the *modules* directory. The server module expects the *rshost* executable to be in the same directory where you copied the *rs.config* file.
4. Set the `PATH` and `LD_LIBRARY_PATH` environment variables to include the Apache *modules* directory.
5. Edit the Apache *conf/httpd.conf* file.
  - a. Add the following lines to load the Relay Server client and server modules:

```
LoadModule iarelayserver_client_module modules/mod_rs_ap_client.so
LoadModule iarelayserver_server_module modules/mod_rs_ap_server.so
```

The client and server modules are invoked using different URLs. The client module explicitly looks for the string *iarelayserver* in the URL path. That part of the URL need not change.

- b. Add the following line to create a `<location>` section for the client module:

```
<LocationMatch /cli/iarelayserver/* >
  SetHandler iarelayserver-client-handler
</LocationMatch>
```

- c. Add the following line to create a `<location>` section for the server module:

```
<Location /srv/iarelayserver/* >
  SetHandler iarelayserver-server-handler
  RSConfigFile "/<apache-install>/modules/rs.config"
</Location>
```

You must specify an `RSConfigFile` directive which specifies the location of the Relay Server configuration file, `rs.config`. The `rs.config` file must reside in the same directory where the `rshost` executable is deployed.

- d. If the `Timeout` directive is set, ensure it is set to at least 60 seconds.
6. On Linux, if any of the following environment variables are set globally when Apache spawns a process, then there is nothing further needed for the configuration of Apache: `$TMP`, `$TMPDIR` or `$TEMP`.

If any of the above environment variables are not set globally, or if you want the default Relay Server log file to go in a specific temporary directory (for example, when the State Manager is started automatically but without customizations), then edit the file `/<apache-dir>/bin/envvars` to set and then export `TMP`.

For example, to edit `$TMP` in the `envvars` file, do the following:

```
set TMP="/tmp"
export TMP
```

This sets the environment variable in the shell that Apache creates before it spawns its processes.

**Note**

In all cases, the Apache user process must have write permissions to the `tmp` directory location you choose.

## Updating a Relay Server farm configuration

A Relay Server farm configuration is defined by the contents of the Relay Server configuration file. Each Relay Server in a Relay Server farm shares the same Relay Server configuration file, so when you update a Relay Server farm configuration you must update the Relay Server configuration file at each Relay Server in the farm. Updates include any of the following:

- Adding a new Relay Server to the Relay Server farm.
- Creating a new backend server farm and allowing it access to the Relay Server farm.
- Adding a new backend server to an existing backend server farm.
- Changing the properties of a Relay Server, backend server farm, or a backend server.
- Changing options.

One way to update a Relay Server configuration is to shutdown all Relay Servers, replace the Relay Server configuration file with the updated version, and restart all the Relay Servers. However, shutting down and restarting the Relay Servers means that users of the Relay Server may incur a service interruption.

The preferred method of updating a Relay Server configuration is to use the Relay Server State Manager to update the configuration while a Relay Server farm is running without interrupting service.

Updating a Relay Server configuration is done by launching a new instance of the Relay Server State Manager using the following command line format:

```
rshost -u -f <filename>
```

The `-u` option instructs the Relay Server State Manager to perform an update operation. The `-f` option specifies the name of the configuration file containing the updated configuration. See [“Relay Server State Manager” on page 250](#).

Below is an overview of the steps required to update a Relay Server farm configuration:

1. Make your changes to the master copy of the Relay Server configuration file.
2. On each computer running an instance of a Relay Server that belongs to the Relay Server farm being updated, do the following:
  - a. Replace the old configuration file with the updated configuration file.
  - b. Run the Relay Server State Manager with the updated configuration file.

## Updating a Relay Server configuration for IIS on Windows

### To update a Relay Server configuration for IIS on Windows

1. For each computer that belongs to the Relay Server farm you are updating, copy the updated configuration file to the `ias_relay_server\server` directory under the Relay Server web site home directory. The configuration file must be called `rs.config` if auto start is used.

2. From the `ias_relay_server\server` directory, run the following command line to apply the configuration update:

```
rshost -u -f rs.config
```

3. Repeat the previous steps for each computer in the Relay Server farm that is being updated.

## Updating a Relay Server configuration for Apache on Linux

### To update a Relay Server configuration for Apache on Linux

1. Copy the updated configuration file to the `/modules` directory under the Apache install directory. The configuration file must be called `rs.config` if auto start is used.
2. From the `/<Apache-install>/modules` directory, run the following command line to apply the configuration update:

```
rshost -u -f rs.config
```

3. Repeat the previous steps for each computer in the Relay Server farm that is being updated.

## Sybase Relay Server hosting service

The Sybase Relay Server hosting service is a farm of Relay Servers hosted by Sybase. It is intended to ease the development of mobile applications that use MobiLink data synchronization and to simplify the evaluation process for developers, especially where data is sent using public wireless networks. You do not need to ask your IT department to install anything or open any holes in your corporate firewall. All communication between MobiLink and the hosting service uses HTTP(S) via an outbound connection initiated by MobiLink.

The Sybase Relay Server hosting service is not intended for production deployments. Before deploying your production application, you must first install the Relay Server in your own corporate infrastructure.

## Using the Relay Server hosting service

The following sections describe how to perform some basic tasks.

### Subscribing to the Relay Server hosting service

To use the Sybase Relay Server hosting service you must first subscribe to it.

#### To subscribe to the Sybase Relay Server hosting service

1. From your web browser, go to <http://relayserver.sybase.com/account>. This takes you to the Sybase Relay Server hosting service home page.
2. Create an account by clicking **Register**.
3. You are asked to specify a **Subscription ID** (choose one that is unique to your organization) and **Password**, provide contact information for your self and your organization, and agree to the **Hosted Relay Service Terms of Service**. Click **Submit**.

Once you have successfully registered, an email is sent to you confirming your registration.

### Logging in to the Relay Server hosting service

#### To log in to the Relay Server hosting service

1. Log in to your newly created account by clicking **Log In**.
2. Enter the **Subscription ID** and **Password** you entered during the registration process. Once logged in, you are taken to the **Account Information** page. The account information page allows you to modify subscriber information and specify the back-end server farm(s) that will be accessing this service.

### Adding a server farm

#### To add a server farm

1. Click **Add New Farm**.



2. Enter a unique **Farm Name** to describe the server farm.
3. Choose a **Type** from the dropdown list.
4. Provide a unique name for each server in the farm. You can specify a maximum of two servers.
5. Click **Create Farm**. A confirmation is displayed if the farm was successfully added.
6. Click **Configuration Instructions** to learn more about using the service. The instructions are based on the information you provided.
7. Click **Log Out** when you are done.

## Using MobiLink with the Relay Server

The following sections provides information about using the Relay Server with MobiLink.

For information about which operating systems and browsers are supported for the Relay Server, see <http://www.sybase.com/detail?id=1002288>.

For information about deploying the Outbound Enabler, see “[Deploying the MobiLink server](#)” on page 801.

## Connecting a client to the Relay Server farm

Once a Relay Server farm has been properly configured, a client connects to the Relay Server farm using the following URL:

`http://<Relay Server client extension URI>/<farm>`

### Options

Option	Description
<code>&lt;Relay Server client extension URI&gt;</code>	For IIS on Windows, <code>&lt;domain name&gt;/ias_relay_server/client/rs_client.dll</code> For Apache on Linux, <code>&lt;domain name&gt;/cli/iarelayserver</code>
<code>&lt;farm&gt;</code>	Identifies the back-end farm (a group of back-end servers) that Relay Server forwards the client request to.

### SQL Anywhere MobiLink client connection example

A SQL Anywhere MobiLink client should specify the following options to connect to server farm **F1**:

```
-e "ctp=http;
    adr='host=relayserver.sybase.com;
    url_suffix=/ias_relay_server/client/rs_client.dll/F1'"
```

For HTTPS, change http to https.

### UltraLite/UltraLiteJ MobiLink client connection example

An UltraLite/UltraLiteJ MobiLink client should set the following properties in the ULSyncParms class to connect to server farm **F1**:

- Set the stream type to HTTPS.
- Set the stream parameters to the following:

```
"host=testrelay.iAnywhere.com; url_suffix=/ias_relay_server/client/
rs_client.dll/F1"
```

## QAnywhere client connection example

A QAnywhere client should specify the following options to connect to server farm **F1**:

```
-x "http(host=relayserver.sybase.com;url_suffix=/ias_relay_server/client/  
rs_client.dll/F1"
```

## Sample scenario

Suppose company ABC has developed a mobile application and now wants to set up the deployment runtime to service the mobile application. Initially, the mobile deployment consists of 10000 devices and grows in the future. The customer therefore wants a fault tolerant and load-balanced environment that is able to handle the load today and be easily extended to handle more mobile deployments in the future. Based on the data synchronization characteristics of the mobile application, the customer has determined that the following configuration is needed:

- 2 MobiLink servers
- 2 Relay Servers
- 1 load balancer

Since the company uses IIS as its web server, the IIS version of the Relay Server is used.

### Notes

- Each Relay Server is deployed on its own computer. Two computers, with host names **rs1.abc.com** and **rs2.abc.com** are used.
- Each MobiLink server is deployed on its own computer. The two MobiLink servers are assigned names **ml1** and **ml2** and belong to the back-end server farm called **abc.mobilink**.
- The load balancer is addressable using the host name **www.abc.com**.
- For maximum security, HTTPS is used by all clients and Outbound Enablers connecting to the Relay Servers. We assume all Web servers are equipped with a certificate from a well known Certificate Authority (CA), and the back-end server computers all have the corresponding trusted root certificates in their standard certificate store.

### To set up the Relay Server farm

1. The first step is to create the Relay Server configuration file.

The filename containing the configuration must be called **rs.config**. For this particular scenario, the following configuration file is used:

```
#  
# Options  
#  
[options]  
verbosity = 1  
  
#  
# Define the Relay Server farm  
#
```

```
[relay_server]
host = rs1.abc.com

[relay_server]
host = rs2.abc.com

#
# Define the MobiLink backend server farm
#
[backend_farm]
id = abc.mobilink
client_security = on
backend_security = on

#
# List MobiLink servers that are connecting to the Relay Server farm
#
[backend_server]
farm = abc.mobilink
id = ml1
token = mltoken1

[backend_server]
farm = abc.mobilink
id = ml2
token=mltoken2
```

2. Deploy the configuration file *rs.config* along with the Relay Server components to the two computers that are running the Relay Server.
3. Start MobiLink server on the two computers that are running the MobiLink servers, and then start the corresponding Outbound Enabler using the following commands.

On the computer running MobiLink server with id ml1:

```
mlsrv11 -x http -z ml1 -ss <other ML options>
rsoe -f abc.mobilink -id ml1 -t mltoken1 -cr
"host=www.abc.com;port=443;https=1"
```

On the computer running MobiLink server with id ml2:

```
mlsrv11 -x http -z ml2 -ss <other ML options>
rsoe -f abc.mobilink -id ml2 -t mltoken2 -cr
"host=www.abc.com;port=443;https=1"
```

4. Once all servers and Outbound Enablers are running, MobiLink clients are able to connect to the farm using the following connection information:
  - **HTTPS** protocol
  - **host** www.abc.com
  - **url\_suffix** /ias\_relay\_server/client/rs\_client.dll/abc.mobilink

---

# Redirector (deprecated)

## Contents

Introduction to the Redirector (deprecated) .....	266
Setting up the Redirector .....	268
Configuring MobiLink clients and servers for the Redirector .....	269
Configuring Redirector properties .....	271
NSAPI Redirector for Netscape/Sun web servers on Windows (deprecated) .....	277
NSAPI Redirector for Netscape/Sun web servers on Unix (deprecated) .....	280
ISAPI Redirector for Microsoft web servers (deprecated) .....	282
Servlet Redirector (deprecated) .....	284
Apache Redirector (deprecated) .....	287
M-Business Anywhere Redirector (deprecated) .....	289

---

**Note**

The Redirector has been deprecated. In its place, use the Relay Server. See [“The Relay Server” on page 239](#).

## Introduction to the Redirector (deprecated)

### Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

MobiLink includes a web server extension called the **Redirector** that routes requests and responses between a client and the MobiLink server. A plug-in such as this is also commonly called a **reverse proxy**.

The main reason for routing requests through a web server is to use existing web server and firewall configurations for HTTP or HTTPS synchronization. However, a web server can operate as a proxy without the Redirector. The Redirector is most useful when you have more than one MobiLink server.

See “[Options when using a web server](#)” on page 267.

Using the Redirector, you can configure your web server to route specific URL requests to one or more computers running the MobiLink server.

Web servers can be configured to pass requests with specific URLs or ranges of URLs to extension programs commonly written in the form of Perl CGI scripts, DLLs, or other extension mechanisms. These extension programs may access external data sources and provide responses for the web server to deliver to its clients.

### Load balancing and failover

The Redirector implements load balancing and failover using a simple round robin algorithm (servers are chosen in a fixed cyclic order). The Redirector pings each MobiLink server and stops sending requests to a server that is not responding. The Redirector detects when a MobiLink server is running again and resumes sending requests at that time.

### HTTPS synchronization

When you specify the HTTPS protocol on a MobiLink client, HTTPS is used for the connection between the remote database and web server: HTTP headers are encrypted over TLS using RSA encryption before being sent to or from the web server, and the web server decrypts the HTTPS and sends HTTP to MobiLink via the Redirector. All Redirectors support this version of HTTPS, in which HTTPS is only used for the connection between the MobiLink client and the web server.

The HTTPS protocol is slower than other secure protocols.

### Full HTTPS

For some Redirectors (such as the Apache Redirector, the ISAPI Redirector, and the NSAPI Redirector on Windows), the Redirector offers an option to re-encrypt the stream as HTTPS and send it to the MobiLink server.

For a list of Redirectors that support HTTPS from the Redirector to the MobiLink server, see <http://www.sybase.com/detail?id=1061837>.

### Supported web servers

Plug-ins are provided for the following web servers:

Redirector plug-in	...supports
ISAPI Redirector	Microsoft web servers
NSAPI Redirector	Sun One (Netscape) and iPlanet web servers on Windows and Unix
Servlet Redirector	Web servers that support the Java Servlet API 2.3, including Apache Tomcat and Sun One web servers on Unix
Native Apache Redirector	Apache web server
M-Business Anywhere Redirector	M-Business Anywhere web server

For more information about Redirector support, see:

- <http://www.sybase.com/detail?id=1061837>

## Options when using a web server

The Redirector is one way to route MobiLink synchronization through a web server. The Redirector is particularly useful for synchronizing across a firewall or with multiple MobiLink servers.

The main reason for routing requests through a web server is to use existing web server and firewall configurations for HTTP or HTTPS synchronization. The Redirector is most useful when you have more than one MobiLink server.

You can also route synchronizations through a web server without using the Redirector. In this case, you might configure your web server as a proxy to route synchronizations to a MobiLink server. For more information about how to do this with your web server, see your web server documentation.

The following table contains recommendations to help you decide how best to route your MobiLink synchronizations.

	Direct connection possible	Direct connection not possible
<b>One MobiLink server</b>	Use TCP/IP instead of HTTP	Use an HTTP or HTTPS proxy to pass messages through the web server to the MobiLink server
<b>Multiple MobiLink servers</b>	Use the Redirector with HTTP or HTTPS	Use the Redirector with HTTP or HTTPS

See “Redirector (deprecated)” on page 265.

## Setting up the Redirector

**Note**

The Redirector has been deprecated. In its place, use the Relay Server. See [“The Relay Server” on page 239](#).

The following sections describe how to configure your web server to manage synchronization requests.

### Overview of the configuration process

1. Configure the MobiLink server.  
See [“Configuring MobiLink clients and servers for the Redirector” on page 269](#).
2. Modify the Redirector configuration file. There are two ways to do this, depending on whether you are using a Redirector that supports server groups or one that does not support server groups. See:
  - [“Configuring Redirector properties \(for Redirectors that support server groups\)” on page 273](#)
  - [“Configuring Redirector properties \(for Redirectors that don't support server groups\)” on page 275](#)
3. Perform web server-specific configuration.  
See one of the following:
  - [“NSAPI Redirector for Netscape/Sun web servers on Windows \(deprecated\)” on page 277](#)
  - [“ISAPI Redirector for Microsoft web servers \(deprecated\)” on page 282](#)
  - [“Servlet Redirector \(deprecated\)” on page 284](#)
  - [“Apache Redirector \(deprecated\)” on page 287](#)
  - [“M-Business Anywhere Redirector \(deprecated\)” on page 289](#)
4. Configure MobiLink clients.  
See [“Configuring MobiLink clients and servers for the Redirector” on page 269](#).



# Configuring MobiLink clients and servers for the Redirector

## Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

This section describes how to configure MobiLink clients and the MobiLink server for synchronization through a web server. The following procedures set the parameters required for requests directed through web servers.

## MobiLink clients

### To configure MobiLink clients (SQL Anywhere and UltraLite)

1. Specify the communication type for MobiLink clients to HTTP or HTTPS.

For more information about setting the communication type for SQL Anywhere clients, see “[CommunicationType \(ctp\) extended option](#)” [*MobiLink - Client Administration*].

For more information about setting the communication type for UltraLite clients, see “[Network protocol options for UltraLite synchronization streams](#)” [*UltraLite - Database Management and Reference*].

2. Specify the following HTTP/HTTPS synchronization protocol options for MobiLink clients:

- **host** the name or IP address of the web server.
- **port** the web server port accepting HTTP or HTTPS requests.
- **url\_suffix** This setting depends on the type of Redirector you are using:

- For the ISAPI Redirector:

```
exe_dir/iaredirect.dll/ml/[server-group/]
```

where *exe\_dir* is the location of *iaredirect.dll*, and *server-group* is optionally the name of the group.

- For NSAPI Redirectors:

```
mlredirect/ml/[server-group/]
```

where *mlredirect* is a name mapped in your *obj.conf* file.

- For the servlet Redirector:

```
iaredirect/ml/
```

- For the native Redirector for Apache, set this to whatever you chose in the Redirector's <location> tag in the *httpd.conf* file. For example, if the location is <Location /iaredirect/ml>, then the url\_suffix is:

```
iaredirect/ml/
```

- For the M-Business Anywhere Redirector, set this to whatever you chose in the Redirector's <location> tag in the *sync.conf* file. For example, if the location is <Location / iaredirect/ml>, then the url\_suffix is:

```
iaredirect/ml/
```

See “url\_suffix” [*MobiLink - Client Administration*].

For more information about setting protocol options for UltraLite clients, see “Network protocol options for UltraLite synchronization streams” [*UltraLite - Database Management and Reference*].

For more information about setting protocol options for SQL Anywhere clients, see “MobiLink client network protocol option summary” [*MobiLink - Client Administration*].

### MobiLink server

#### To configure MobiLink servers

1. For Redirectors that support HTTPS, you can start the MobiLink server with the HTTPS protocol. For a list of Redirectors that support HTTPS, see <http://www.sybase.com/detail?id=1061837>.

For Redirectors that do not support HTTPS, the MobiLink server must be started with the HTTP protocol to use HTTP or HTTPS for communication between the client and the proxy. These Redirectors cannot use HTTPS directly.

For example, the HTTP protocol may be specified on the mlsrvl1 command line as follows:

```
mlsrvl1 -x http
```

See “-x option” on page 107.

2. In addition, you may want to set the following parameter for the MobiLink server:
  - **port** for the HTTP protocol, MobiLink defaults to port 80. For the HTTPS protocol, MobiLink defaults to port 443. If the MobiLink server is running on the same computer as the web server, port 80 is normally in use by the web server. If this is the case you must specify a different port. For example, you could use port 2439, which is the Internet Assigned Numbers Authority (IANA)-registered port number for the MobiLink server.

For more information about port, see “-x option” on page 107.

## Configuring Redirector properties

### Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

There are different properties for Redirectors that support server groups and Redirectors that don't support server groups.

## MobiLink server groups

You can partition your MobiLink servers into server groups. This allows you to have groups of clients that access distinct groups of MobiLink servers.

For a list of Redirectors that support server groups, see <http://www.sybase.com/detail?id=1061837>.

You create a server group by adding a section to your Redirector configuration file (*redirector\_server\_group.config*) with the group name in square brackets followed by settings for that group. At a minimum, a group must specify one ML directive. You can also set the ML\_CLIENT\_TIMEOUT option for a group. You reference the group in the url\_suffix option on your client.

You can create a default server group by specifying a group with no name before you specify any named groups in the file. This default group is useful for backward compatibility. It is used when the client does not specify a server group name in their url\_suffix option.

See “url\_suffix” [[MobiLink - Client Administration](#)].

You can also specify default settings for all server groups for the SLEEP and LOG\_LEVEL properties. These can be specified anywhere in the configuration file.

## Supporting old and new clients

If your MobiLink server needs to support version 8 or 9 remote databases and version 10 and later remote databases, then you need to open a minimum of two ports: you use the mlsrv11 -x option to open a port for new clients, and you use the mlsrv11 -xo option to open a port for old clients. If you are also using the Redirector, you need to set up server groups so that the Redirector directs clients to the appropriate port.

In a typical Redirector setup, you would start multiple MobiLink servers. In the simplest case, you have one MobiLink server running with two ports, opened with -x and -xo, and you create two server groups, one for each. The following is a partial mlsrv11 command line that opens two ports for the MobiLink server:

```
mlsrv11 -c "dsn=YourDSN" -x http(port=111) -xo http(port=222)
```

You add sections to your Redirector configuration file for the two server groups:

```
[v10service]
ML="host=mySrv.myCorp.com;port=111"
[v9service]
ML="host=mySrv.myCorp.com;port=222"
```

When you start your clients, you specify the `url_suffix` option with the name of the server group. For example, for SQL Anywhere clients and an ISAPI web server, part of the `dbmlsync` command line for the version 10 clients would be:

```
dbmlsync -e "adr='host=somehost;port=5001;url_suffix=scripts/iaredirect.dll/ml/v10service' "...
```

Part of the `dbmlsync` command line for your version 9 clients would be:

```
dbmlsync -e "adr='host=somehost;port=5001;url_suffix=scripts/iaredirect.dll/ml/v9service' "...
```

### See also

- [“-x option” on page 107](#)
- [“-xo option” on page 113](#)
- [“Configuring Redirector properties \(for Redirectors that support server groups\)” on page 273](#)
- [“url\\_suffix” \[MobiLink - Client Administration\]](#)

### Example

The following is a sample `redirector_server_group.config` file, showing some typical settings and the creation of server groups.

```
#
# Set up the default server group:
#
ML="host=mySrv1.myCorp.com;port=222"
ML="host=mySrv2.myCorp.com;port=222"
#
# Set up a server group named myOldGroup:
#
[myOldGroup]
ML="host=myOldSrv1.myCorp.com;port=111"
ML="host=myOldSrv2.myCorp.com;port=111"
ML_CLIENT_TIMEOUT=30
#
# Set up a server group named myNewGroup:
#
[myNewGroup]
ML="host=myNewSrv1.myCorp.com;port=333"
ML="host=myNewSrv2.myCorp.com;port=555"
ML_CLIENT_TIMEOUT=240
#
# Set up a server group named mlSecureGroup:
#
[theirSecureGroup]
ML="https=true;Srv1.Corp.com;trusted_certificates=c:\Corp\publicRoot.crt"
ML="https=true;Srv2.Corp.com;trusted_certificates=c:\Corp\publicRoot.crt"
#
# Set global properties:
#
LOG_LEVEL=5
SLEEP=15
```

## Configuring Redirector properties (for Redirectors that support server groups)

This section describes generic web server configuration steps to configure Redirector properties. It applies to Redirectors that support server groups.

For a list of Redirectors that support server groups, see <http://www.sybase.com/detail?id=1061837>.

For information about server groups, see “MobiLink server groups” on page 271.

### To configure Redirector properties

1. Complete the steps in “Configuring MobiLink clients and servers for the Redirector” on page 269.
2. Configure a **Redirector configuration file**. A template file named *redirector\_server\_group.config* is located in *install-dir\MobiLink\redirector*. The easiest way to configure a Redirector configuration file is to modify *redirector\_server\_group.config*.

The following rules apply to the Redirector configuration file:

- The maximum line length is 2000 characters.
- Comments start with the hash character (#).
- For the ISAPI Redirector, the configuration file must be named *redirector.config* and must be in the same directory as *iaredirect.dll*.

You can set the following directives in this file:

- **Server groups** To create server groups, you create sections in *redirector\_server\_group.config* that start with a server group name in square brackets, and then define the server group.

See “MobiLink server groups” on page 271.

- **LOG\_LEVEL** Used to control the amount of output written to the log file. Values are 0 to 7, with higher numbers generating more output. By default, the log file is called *redirector.log* and is located in the same place as the Redirector configuration file. For NSAPI Redirectors, you can change the name and location in *magnus.conf* using the *logFile* directive.

- **ML** There are two ways that you can use the ML directive:
  - For Redirectors that do not support HTTPS from the Redirector to the MobiLink server or when you are not using HTTPS, you can use the ML directive to specify the list the computers running the MobiLink server, in the form `ML=host:port`. To specify multiple computers, you repeat this syntax on separate lines. For example:

```
ML=209.123.123.1:8080
ML=myCompany.com:8081
```

- If your Redirector supports HTTPS from the Redirector to the MobiLink server and you are using HTTPS, you should specify MobiLink client network protocol options in a semicolon-separated list, as follows:

```
ML="https=true;network-client-options;..."
```

For example,

```
ML="https=true;host=My-pc;port=82;trusted_certificates=rsaroot.crt"
```

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

For a description of HTTPS support in the Redirector, see “[Full HTTPS](#)” on page 266.

For a list of Redirectors that support HTTPS from the Redirector to the MobiLink server, see <http://www.sybase.com/detail?id=1061837>.

Your MobiLink server must be started with the same protocol and port number as it is given in your ML directive. If there is a difference, you must stop the MobiLink server and restart it with the correct information.

- **ML\_CLIENT\_TIMEOUT** Used to ensure that the MobiLink server can detect duplicate synchronizations from the same remote database. The timeout should be set to the maximum timeout of any client using the same server group. If you set this property to 0, resynchronization to a different server is allowed immediately. The default value is 240 seconds.
  - **SLEEP** Used to set the interval in seconds at which the Redirector checks that the servers are functioning. The Redirector checks one server, waits for the amount of time set in this option, checks the next server, and so on in a cycle. For example, `SLEEP=10`. `SLEEP` is case sensitive. The default is 20 seconds.
3. Copy the Redirector configuration file to the web server.

If the MobiLink server is not installed on the same computer as the web server, copy the Redirector configuration file to the computer that holds the web server (or to a drive that computer has access to).

For ISAPI web servers, copy the Redirector configuration file to the directory `Inetpub\scripts` and rename it `redirector.config`.

For other web servers, you can copy the Redirector configuration file to any directory.
  4. Complete web server-specific configuration in one of the following sections:
    - “[ISAPI Redirector for Microsoft web servers \(deprecated\)](#)” on page 282
    - “[NSAPI Redirector for Netscape/Sun web servers on Windows \(deprecated\)](#)” on page 277

### Example

The following is a sample Redirector configuration file. This file specifies the following:

- The Redirector should sleep for 10 seconds after checking that a server is functioning.
- The three computers running MobiLink servers that are able to process requests.

```
SLEEP=10
ML=myServ-pc:80
ML=209.123.123.1:8080
ML=myCompany.com:8081
```

## Configuring Redirector properties (for Redirectors that don't support server groups)

This section describes generic web server configuration steps to configure Redirector properties. It applies to Redirectors that don't support server groups.

For a list of Redirectors that support server groups, see <http://www.sybase.com/detail?id=1061837>.

### To configure Redirector properties

1. Complete the steps in “Configuring MobiLink clients and servers for the Redirector” on page 269.
2. Copy *redirector.config* to the web server.

The file *redirector.config* is located in *install-dir\MobiLink\redirector*.

If the MobiLink server is not installed on the same computer as the web server, copy *redirector.config* to the computer that holds the web server.

3. Configure the Redirector configuration file.

To configure communications between the web server and MobiLink server, you must edit the file *redirector.config* on the computer that holds the web server.

The following rules apply to *redirector.config*:

- The maximum line length is 300 characters.
- Comments start with the hash character (#).
- You cannot include spaces or tabs in the directive definitions.

You can set the following directives in this file:

- **LOG\_LEVEL** Used to control the amount of output written to the log file. Values are 0, 1, and 2, with 1 being the default and 2 generating the most output. For the Apache Redirector, this setting has no effect; set the log level in the LogLevel section of the Apache configuration file, *httpd.conf*.
- **ML** ML is case sensitive. There are two ways that you can use the ML directive.

For Redirectors that do not support HTTPS or when you are not using HTTPS, you can use the ML directive to specify the list the computers running the MobiLink server, in the form `ML=host:port`. To specify multiple computers, you repeat this syntax on separate lines. For example:

```
ML=209.123.123.1:8080
ML=myCompany.com:8081
```

If your Redirector supports HTTPS from the Redirector to the MobiLink server, you can specify MobiLink client network protocol options in a semicolon-separated list, as follows:

```
ML="https=true;network-client-options;..."
```

For example,

```
ML="https=true;host=My-pc;port=82;trusted_certificates=rsaroot.crt"
```

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

For a list of Redirectors that support HTTPS from the Redirector to the MobiLink server, see <http://www.sybase.com/detail?id=1061837>.

Your MobiLink server must be started with the same protocol and port number as it is given in your ML directive. If there is a difference, you must stop the MobiLink server and restart it with the correct information.

- **ML\_CLIENT\_TIMEOUT** Used to ensure that each step of a single synchronization is directed to the same MobiLink server. The Redirector maintains an association between client and server for the duration of ML\_CLIENT\_TIMEOUT. This value is also used to ensure that the MobiLink server can detect duplicate synchronizations from the same remote database. The value of this parameter should be greater than the longest step in any user's synchronization.

The default value is 600 seconds (ten minutes).

- **SLEEP** Used to set the interval in seconds at which the Redirector checks that the servers are functioning. The default is 1800 (30 minutes). For example, SLEEP=3600. SLEEP is case sensitive.

4. Complete web server-specific configuration in one of the following sections:

- “[NSAPI Redirector for Netscape/Sun web servers on Unix \(deprecated\)](#)” on page 280
- “[Servlet Redirector \(deprecated\)](#)” on page 284
- “[Apache Redirector \(deprecated\)](#)” on page 287
- “[M-Business Anywhere Redirector \(deprecated\)](#)” on page 289

### Example

The following is a sample *redirector.config* file. This file specifies the following:

- The Redirector should check every 1800 seconds that the servers are functioning.
- The three computers running MobiLink servers that are able to process requests. When you specify multiple servers, load balancing is automatically enabled.

```
SLEEP=1800
ML=myServ-pc:80
ML=209.123.123.1:8080
ML=myCompany.com:8081
```



# NSAPI Redirector for Netscape/Sun web servers on Windows (deprecated)

## Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

The NSAPI Redirector is provided for the Sun Java System web server, which was previously known as Sun One and the Netscape iPlanet Enterprise Edition web server.

For information about version support, see <http://www.sybase.com/detail?id=1061837>.

To use this Redirector on Unix, see “[NSAPI Redirector for Netscape/Sun web servers on Unix \(deprecated\)](#)” on page 280.

To use the Redirector with Netscape/Sun web servers on other platforms, you can use the servlet Redirector. See “[Servlet Redirector \(deprecated\)](#)” on page 284.

## To configure the NSAPI Redirector

1. Complete the steps in “[Configuring Redirector properties \(for Redirectors that support server groups\)](#)” on page 273.
2. If necessary, copy the file *iaredirect.dll* to the computer that holds the web server. This file is located in *install-dir\MobiLink\redirector\web-server*, where *web-server* is the name of your NSAPI web server.
3. If your web server is on a separate computer from the Redirector, you must copy the following files to that computer and ensure that they are in your path. The files you need depends on what, if any, encryption you are using.

The following file locations are relative to *install-dir*:

Setup	Files required
All	<ul style="list-style-type: none"> <li>• <i>bin32\dblgen11.dll</i><sup>1</sup></li> <li>• <i>bin32\dbicu11.dll</i></li> <li>• <i>bin32\dbicudt11.dll</i></li> </ul>
ECC encryption	<ul style="list-style-type: none"> <li>• <i>bin32\mlcecc11.dll</i></li> </ul>
RSA encryption	<ul style="list-style-type: none"> <li>• <i>bin32\mlcrsa11.dll</i></li> </ul>
RSA encryption with FIPS	<ul style="list-style-type: none"> <li>• <i>bin32\mlcrsafips11.dll</i></li> <li>• <i>bin32\sbgse2.dll</i></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

For information about how to change the language, see “[Understanding the locale language](#)” [*SQL Anywhere Server - Database Administration*].

4. Update the NSAPI *magnus.conf* and *obj.conf* web server configuration files as follows.

### Sample file provided

Sample copies of *magnus.conf* and *obj.conf*, preconfigured for the MobiLink server, are located in *install-dir\MobiLink\redirector\web-server*, where *web-server* is the name of your NSAPI web server.

Update the following sections of the *magnus.conf* and *obj.conf*.

- In *magnus.conf*, specify where *iaredirect.dll* and the Redirector configuration file are located.

At the end of the Init section, add the following text, where *location* is the actual location of the files (*iaredirect.dll* and the Redirector configuration file can be in different locations, although both must be on the same computer as the web server or a drive that is accessible to the web server):

```
Windows:
Init fn="load-modules" shlib="location/iaredirect.dll"
funcs="redirector,initialize_redirector"
Init fn="initialize_redirector" configFile="location/redirector.config"
```

- In *obj.conf*, specify the name of the Redirector to be used in URLs.

At the beginning of the default object section, add the following text. This section should appear exactly as provided below, except that you can change *mlredirect* to whatever you want. All requests of the form *http://host:port/mlredirect/ml/\** are sent to one of the MobiLink servers running with the Redirector.

```
<Object name=default>
NameTrans fn="assign-name" from="/mlredirect/ml/*" name="redirectToML"
```

- In *obj.conf*, specify the objects that are called by the Redirector. After the default object section, add the following section:

```
<Object name="redirectToML">
Service fn="redirector" serverType="ml"
</Object>
```

## Example

The following is an example of the section of *magnus.conf* that you need to customize.

```
Init fn="load-modules" shlib="D:/iaredirect.dll"
funcs="redirector,initialize_redirector"
Init fn="initialize_redirector" configFile="D:/redirector.config"
```

The following is an example of the sections of *obj.conf* that are important to the Redirector:

```
<Object name=default>
NameTrans fn="assign-name" from="/mlredirect/ml/*" name="redirectToML"
...
<Object name="redirectToML">
Service fn="redirector" serverType="ml"
</Object>
```

## To test your configuration

1. Call the Redirector using the following syntax:

```
http://host:port/mlredirect/ml/
```

2. Check the log file to see if the Redirector logged a request.

**Note:** This test does not make a connection to the MobiLink server.

## NSAPI Redirector for Netscape/Sun web servers on Unix (deprecated)

### Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

The NSAPI Redirector is provided for the Sun Java System web server, which was previously known as Sun One and the Netscape iPlanet Enterprise Edition web server.

For information about version support, see <http://www.sybase.com/detail?id=1061837>.

To use this Redirector on Windows, see “[NSAPI Redirector for Netscape/Sun web servers on Windows \(deprecated\)](#)” on page 277.

To use the Redirector with Netscape/Sun web servers on other platforms, you can use the servlet Redirector. See “[Servlet Redirector \(deprecated\)](#)” on page 284.

### To configure the NSAPI Redirector

1. Complete the steps in “[Configuring Redirector properties \(for Redirectors that don't support server groups\)](#)” on page 275.
2. If necessary, copy the file *iaredirect.so* to the computer that holds the web server. This file is located in *install-dir\MobiLink\redirector\web-server*, where *web-server* is the name of your NSAPI web server.
3. Update the NSAPI web server configuration files *magnus.conf* and *obj.conf* as follows.

### Sample file provided

Sample copies of *magnus.conf* and *obj.conf* are located in *install-dir\MobiLink\redirector\web-server*, where *web-server* is the name of your NSAPI web server. You can use these sample files to confirm where the following sections fit in to the file.

Update the following sections of the files *magnus.conf* and *obj.conf*.

- In *magnus.conf*, specify where *iaredirect.so* and *redirector.config* are located.

At the end of the Init section, add the following text, where *location* is the actual location of the files. (*iaredirect.so* and *redirector.config* can be in different locations, although both must be on the same computer as the web server.)

```
Solaris:
Init fn="load-modules" shlib="location/iaredirect.so"
funcs="redirector,initialize_redirector"
Init fn="initialize_redirector" configFile="location/redirector.config"
```

- In *obj.conf*, specify the name of the Redirector to be used in URLs.

At the beginning of the "default object" section, add the following text. This section should appear exactly as provided below, except that you can change *mlredirect* to whatever you want. All requests of the form *http://host:port/mlredirect/ml/\** are sent to one of the MobiLink servers running with the Redirector.

```
<Object name=default>
NameTrans fn="assign-name" from="/mlredirect/ml/*" name="redirectToML"
```

- In *obj.conf*, specify the objects that are called by the Redirector. After the "default object" section, add the following section:

```
<Object name="redirectToML">
Service fn="redirector" serverType="ml"
</Object>
```

## Example

The following is an example of the section of *magnus.conf* that you need to customize.

```
Init fn="load-modules" shlib="location/iaredirect.so"
funcs="redirector,initialize_redirector"
Init fn=" initialize_redirector " configFile="location/redirector.config"
```

The following is an example of the sections of *obj.conf* that are important to the Redirector.

```
<Object name=default>
NameTrans fn="assign-name" from="/mlredirect/ml/*" name="redirectToML"
...
<Object name="redirectToML">
Service fn="redirector" serverType="ml"
</Object>
```

## To test your configuration

1. Call the Redirector using the following syntax:

```
http://host:port/mlredirect/ml/
```

2. Check the log file to see if the Redirector logged a request.

**Note:** This test does not make a connection to the MobiLink server.

## ISAPI Redirector for Microsoft web servers (deprecated)

**Note**

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

If you are using a Microsoft web server, you can use the ISAPI version of the Redirector.

For information about version support, see <http://www.sybase.com/detail?id=1061837>.

### To configure ISAPI Redirector for Microsoft web servers

1. Complete the steps in “[Configuring Redirector properties \(for Redirectors that support server groups\)](#)” on page 273.

2. Copy the file *iaredirect.dll* to *Inetpub\scripts* on the computer that holds the web server.

The file *iaredirect.dll* is located in *install-dir\MobiLink\redirector\IIS5*.

The directory *Inetpub\scripts* is in the Microsoft web server installation directory.

3. If your web server is on a separate computer from the Redirector, you must copy the following files to that computer and ensure that they are in your path. What files you need depends on what, if any, encryption you are using.

The following file locations are relative to *install-dir*:

Setup	Files required
All	<ul style="list-style-type: none"> <li>• <i>bin32\dblgen11.dll</i><sup>1</sup></li> <li>• <i>bin32\dbicu11.dll</i></li> <li>• <i>bin32\dbicudt11.dll</i></li> </ul>
ECC encryption	<ul style="list-style-type: none"> <li>• <i>bin32\mlcecc11.dll</i></li> </ul>
RSA encryption	<ul style="list-style-type: none"> <li>• <i>bin32\mlcrsa11.dll</i></li> </ul>
RSA encryption with FIPS	<ul style="list-style-type: none"> <li>• <i>bin32\mlcrsafips11.dll</i></li> <li>• <i>bin32\sbgse2.dll</i></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

For information about how to change the language, see “[Understanding the locale language](#)” [*SQL Anywhere Server - Database Administration*].

**Note****To test your configuration**

1. Call the ISAPI Redirector using the following syntax:

```
protocol://host[:port]/exec_dir/iaredirect.dll/ml/
```

where:

- **protocol** is http or https.
- **host** is the host name of the web server.
- **port** is the port on which the web server is listening, if it is not the default port.
- **exec\_dir** is the directory where you installed the Redirector DLL, *iaredirect.dll*. The default directory is *scripts*.

For example,

```
http://server:8080/scripts/iaredirect.dll/ml/
```

2. Check the log file to see if the Redirector logged a request.

**Note:** This test does not make a connection to the MobiLink server.

3. If your configuration is not successful, check the following:

- The directory *Inetpub\scripts* should be created during the web server installation with execute permissions.
- You can put your Redirector configuration file and *iaredirect.dll* in a different directory only if you use Internet Information Services to give execute permissions to the directory.
- You must have a virtual directory that points to the *Inetpub\scripts* directory. If you do not, you must open Internet Information Services and manually create a virtual directory. This virtual directory should point to *Inetpub\scripts* and have Execute Permissions set to Scripts and Executables. See the IIS online help for instructions.

## Servlet Redirector (deprecated)

**Note**

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

The servlet Redirector is provided for web servers that support the Java servlet specification version 2.3 and later. The following procedure is an example of how to set up the servlet Redirector for Tomcat version 5.5.9 and Apache 2.0.55.

For information about version support, see <http://www.sybase.com/detail?id=1061837>.

There is also a native Redirector for Apache web servers. For more information, see “[Apache Redirector \(deprecated\)](#)” on page 287.

### Overview

This section describes how to install the servlet version of the Redirector to work on an Apache web server in conjunction with the Tomcat servlet container.

Installation requires the following steps:

#### To configure the servlet Redirector for Apache Tomcat

1. Complete the steps in “[Configuring Redirector properties \(for Redirectors that don't support server groups\)](#)” on page 275.
2. Install the servlet version of the Redirector in Tomcat.
3. Configure the Apache web server to run as a proxy.

### Install the servlet Redirector in Tomcat

In the following procedure, `%CATALINA_HOME%` is the root directory of your Tomcat installation.

#### To install the servlet Redirector in Tomcat

1. Install Tomcat as a standalone server.
2. Optionally, set the required Tomcat HTTP port.

Tomcat binds to port 8080 by default. If there is a conflict, perhaps because another web server is using this port,

  - Open the file: `%CATALINA_HOME%/conf/server.xml`.
  - Search for 8080 (which is in a `<Connector>` tag).
  - Change it to a port that is not in use.
3. Install the servlet Redirector as a web application.
  - Copy `iaredirect.war` file to `%CATALINA_HOME%/webapps`.



- Shut down and restart Tomcat.  
Tomcat expands the war file and creates the directory *iaredirect* for the Redirector web application.
- Edit the file `%CATALINA_HOME%/webapps/iaredirect/WEB-INF/web.xml`. Search for **redirector.config** (in an `<init-param>` tag), and correct the path for the *redirector.config* file.  
Change the entry **redirector.config** to read `drive:/path/redirector.config`. Even on Windows operating systems, use a forward slash as a path separator, as in `d:/redirector.config`.
- Shut down and restart Tomcat for the changes to take effect.  
Once the changes have taken effect, you no longer need the war file in the deployed location.
- The Redirector can now be invoked through the following URL:  
`http://tc-host:tc-port/iaredirect/ml/`  
where *tc-host* is the computer and *tc-port* the port on which Tomcat is listening.

### Configure the Apache web server as a proxy

In the following procedure, `%APACHE_HOME%` is the root directory of your Apache installation.

#### To configure the Apache web server as a proxy

1. Install the Apache web server.
2. Optionally, change the Apache web server port:
  - Edit the file `%APACHE_HOME%/conf/httpd.conf` and change the **Port** setting.
3. Configure Apache to run as a proxy:

In `%APACHE_HOME%/conf/httpd.conf`, add the following directives:

```
LoadModule proxy_module module-path/mod_proxy.so
LoadModule proxy_connect_module module-path/mod_proxy_connect.so
LoadModule proxy_http_module module-path/mod_proxy_http.so
```

where *module-path* is the location of the module. For example, the path may be `modules/mod_proxy.so` (the default).

4. Configure Apache to forward Redirector URLs to Tomcat.  
In `%APACHE_HOME%/conf/httpd.conf`, add the following directive so that Apache forwards URLs of the form `http://localhost/iaredirect/*` to the Tomcat 5 Connector listening on port 8080:

```
ProxyPass /iaredirect http://localhost:8080/iaredirect
```

The port number must match the port number used for Tomcat. If Tomcat and Apache are not running on the same computer, provide the computer name where Tomcat is running instead of **localhost**.

### Verifying your setup

#### To check your configuration

1. Call the Redirector using the following syntax:

```
http://host:port/iaredirect/ml/
```

2. Check the log file to see if the Redirector logged a request.

**Note:** This test does not make a connection to the MobiLink server.

# Apache Redirector (deprecated)

## Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

The following setup instructions are written for the Apache web server.

For information about version support, see <http://www.sybase.com/detail?id=1061837>.

If you are using Tomcat, you can also use the servlet Redirector. For more information, see “[Servlet Redirector \(deprecated\)](#)” on page 284.

## To configure the Apache Redirector

- Complete the steps in “[Configuring Redirector properties \(for Redirectors that don't support server groups\)](#)” on page 275.
- Copy the file *mod\_iaredirect.dll* or *mod\_iaredirect.so* to the appropriate directory in your web server, as follows:
  - For Apache on Windows, the file *mod\_iaredirect.dll* is located in *install-dir\MobiLink\redirector\apache\v20\*. Copy this file to the *%apache-home%\modules* directory on the computer that holds the web server.
  - For Apache for Solaris or Linux, the file *mod\_iaredirect.so* is located in *install-dir\MobiLink\redirector\apache\v20\*. Copy it to the *\$APACHE\_HOME/modules* directory on the computer that holds the web server.
- If your web server is on a separate computer from the Redirector, you must copy the following files to that computer and ensure that they are in your path (Windows) or shared path (Unix). What files you need depends on what, if any, encryption you are using.

The following file locations are relative to *install-dir*:

Setup	Files required
ECC encryption	<ul style="list-style-type: none"> <li>Windows: <i>bin32\mlcecc11.dll</i></li> <li>Unix: <i>lib32/libmlcecc11_r.so</i></li> </ul>
RSA encryption	<ul style="list-style-type: none"> <li>Windows: <i>bin32\mlcrsa11.dll</i></li> <li>Unix: <i>lib32/libmlcrsa11_r.so</i></li> </ul>
RSA encryption with FIPS	<ul style="list-style-type: none"> <li>Windows: <i>bin32\mlcrsafips11.dll</i> and <i>bin32\sbgse2.dll</i></li> <li>Unix: <i>lib32/libmlcrsafips11_r.so</i> and <i>libsbgse2_r.so</i></li> </ul>

- Update the Apache web server configuration file *httpd.conf* as follows.
  - In the LoadModule section, add the following line for Windows:

```
LoadModule iaredirect_module modules/mod_iaredirect.dll
```

or the following line for Solaris and Linux:

```
LoadModule iaredirect_module modules/mod_iaredirect.so
```

- Add the following section to the file:

```
<Location /iaredirect/ml>
    SetHandler iaredirect-handler
    iaredirectorConfigFile location/redirector.config
</Location>
```

where */iaredirect/ml* is the relative URL path that you use to invoke the Redirector, and *location* is the directory where *redirector.config* is located.

- If you are using Apache on Solaris or Linux, you may also want to add the following optional directives to the <Location> section you just created:
  - **MaxSyncUsers *number*** The maximum number of MobiLink users synchronizing through the Redirector. This number is used to allocate necessary resources to the Redirector. This number cannot be less than 60. The default is 1000. Only change this setting if the default number of users is less than the actual number.
  - **ShmemDiagnosis on|off** If set to on, debugging of the memory resource is allowed. The default is off.
- 5. To help with debugging, you may want to increase the amount of logging information that the Redirector outputs. To do this, modify the `LogLevel` directive in *httpd.conf* and set it to **LogLevel info**. The log level can be (from most to least verbose): debug, info, notice, warn, error, crit, alert, and emerg.

### Example

The following are examples of the sections of *httpd.conf* that configure the Apache web server to route requests to the MobiLink server. This example works for Windows. For Unix and Linux, change *mod\_iaredirect.dll* to *mod\_iaredirect.so*.

```
LoadModule iaredirect_module modules/mod_iaredirect.dll
...

<Location /iaredirect/ml>
    SetHandler iaredirect-handler
    iaredirectorConfigFile c:/redirector.config
</Location>
```

### To test your configuration

1. Call the Redirector using the following syntax:

```
http://host:port/iaredirect/ml/
```

where *iaredirect/ml* is the relative URL path you specified in the <Location> tag of *httpd.conf*.

2. Check the log file to see if the Redirector logged a request. The default location of the log file is `$APACHE_HOME/logs/error.log`.

**Note:** This test does not make a connection to the MobiLink server.

## M-Business Anywhere Redirector (deprecated)

### Note

The Redirector has been deprecated. In its place, use the Relay Server. See “[The Relay Server](#)” on page 239.

The following setup instructions are written for M-Business Anywhere on Windows, Solaris, or Linux. For information about version support, see <http://www.sybase.com/detail?id=1061837>.

### To configure the M-Business Anywhere Redirector

1. Complete the steps in “[Configuring Redirector properties \(for Redirectors that don't support server groups\)](#)” on page 275.
2. Copy the file *mod\_iaredirect.dll* or *mod\_iaredirect.so* to the M-Business Anywhere *\bin* directory on the computer that holds the web server. This file is located in *install-dir\MobiLink\redirector\MBusinessAnywhere*.
3. If your web server is on a separate computer from the Redirector, you must copy the following files to that computer and ensure that they are in your path (Windows) or shared path (Unix). The files you need depends on what, if any, encryption you are using.

The following file locations are relative to *install-dir*:

Setup	Files required
ECC encryption	<ul style="list-style-type: none"> <li>• Windows: <i>bin32\mlcecc11.dll</i></li> <li>• Unix: <i>lib32/libmlcecc11_r.so</i></li> </ul>
RSA encryption	<ul style="list-style-type: none"> <li>• Windows: <i>bin32\mlcrsa11.dll</i></li> <li>• Unix: <i>lib32/libmlcrsa11_r.so</i></li> </ul>
RSA encryption with FIPS	<ul style="list-style-type: none"> <li>• Windows: <i>bin32\mlcrsafips11.dll</i> and <i>bin32\sbgse2.dll</i></li> <li>• Unix: <i>lib32/libmlcrsafips11_r.so</i> and <i>libsbgse2_r.so</i></li> </ul>

4. For Windows, update the M-Business Anywhere *sync.conf.default* web server configuration file as follows:

- In the LoadModule section, add the following line:

```
LoadModule iaredirect_module @@ServerRoot@@/bin/mod_iaredirect.dll
```

- In the SyncLoadFile section, add the following line:

```
SyncLoadFile @@ServerRoot@@/bin/mod_iaredirect.dll
```

- Add the following section to the file:

```
<Location /iaredirect/ml>
  SetHandler iaredirect-handler
  iaredirectorConfigFile @@ServerRoot@@/conf/redirector.config
</Location>
```

- Run the `setup_defaults.bat` command file to build these changes into the `sync.conf` file.
5. For Solaris and Linux, update the M-Business Anywhere `sync.conf` web server configuration file as follows:

- In the LoadModule section, add the following line:

```
LoadModule iaredirect_module path/bin/mod_iaredirect.so
```

where `path` is the location of the M-Business Anywhere `bin` directory.

- Add the following section to the file:

```
<Location /iaredirect/ml>
  SetHandler iaredirect-handler
  iaredirectorConfigFile location/redirector.config
</Location>
```

where `/iaredirect/ml` is the relative URL path that you use to invoke the Redirector, and `location` is the directory where `redirector.config` is located.

- You may also want to add the following optional directives to the `<Location>` section you just created:
    - **MaxSyncUsers *n*** The maximum number of MobiLink users synchronizing through the Redirector. This number is used to allocate necessary resources to the Redirector. This number cannot be less than 60. The default is 1000. Only change this setting if the default number of users is less than the actual number.
    - **ShmemDiagnosis on|off** If set to on, allows debugging of the memory resource. The default is off.
6. To help with debugging, you may want to increase the amount of logging information that the Redirector outputs. To do this, modify the `LogLevel` directive in `sync.conf` and set it to **LogLevel info**. The log level can be (from most to least verbose): debug, info, notice, warn, error, crit, alert, and emerg.
7. Restart your M-Business Sync Server for the changes to take effect.

### Example

The following are examples of `sync.conf` sections that configure the M-Business Anywhere web server to route requests to the MobiLink server.

This example works on Windows:

```
LoadModule iaredirect_module "c:\program files\M-Business Anywhere/bin/
mod_iaredirect.dll"
...
SyncLoadFile "c:\program files\M-Business Anywhere/bin/mod_iaredirect.dll"
...
<Location \iaredirect\ml>
  SetHandler iaredirect-handler
  iaredirectorConfigFile "c:\AvantGoServer\conf/redirector.config"
</Location>
```

The following example works on Unix and Linux:

```
LoadModule iaredirect_module modules/mod_iaredirect.so
...
<Location /iaredirect/ml>
  SetHandler iaredirect-handler
```

```
    iaredirectorConfigFile "/redirector.config"  
  </Location>
```

### To test your configuration

1. Call the Redirector using the following syntax:

```
http://host:port/iaredirect/ml/
```

where *iaredirect* is the path you specified in the <Location> tag of *sync.conf*.

2. Check the log files *sync\_access.log* and *sync\_error.log* to verify that the Redirector logged a request.

**Note:** This test does not make a connection to the MobiLink server.

---



---

# MobiLink file-based download

## Contents

Introduction to file-based download .....	294
Setting up file-based download .....	295
Validation checks .....	298
File-based download examples .....	301

---

## Introduction to file-based download

File-based download is an alternative way to download data to SQL Anywhere remote databases: downloads can be distributed as files, enabling offline distribution of synchronization changes. This allows you to create a file once and distribute it to many remote databases.

With file-based download, you can put download synchronization changes in a file and transfer it to SQL Anywhere remote databases in any way a file can be transferred. For example, you can:

- broadcast the data by satellite multicast
- apply the update using Sybase Afaria
- email or FTP the file to users

You choose the users you want to receive the file. Full synchronization integrity is preserved in file-based download, including conflict detection and resolution. You can ensure that the file is secure by applying third-party encryption on the file.

### When to use

File-based downloads are useful when a large amount of data changes on the consolidated database, but the remote database does not update the data frequently or does not do any updates at all. For example, price lists, product lists, and code tables.

File-based downloads are not useful when the downloaded data is updated frequently on the remote database or when you are running frequent upload-only synchronizations. In these situations, the remote sites may be unable to apply download files because of integrity checks that are performed when download files are applied.

File-based downloads currently can be used only with SQL Anywhere remote databases.

### Download-only publications

In most cases, you should use a download-only publication for your file-based download. Use a regular publication only when you need to perform uploads with the same publication as you perform file-based downloads.

See [“Download-only publications”](#) [*MobiLink - Client Administration*].

If you use a regular publication, file-based downloads cannot be used as the sole means of updating remote databases. In that case you still need to regularly perform full synchronizations or upload-only synchronizations. Full or upload-only synchronizations are required to advance log offsets and to maintain the log file, which otherwise grows large and slows down synchronization. A full synchronization may also be required to recover from errors.

---

## Setting up file-based download

The following steps provide an overview of the tasks required to set up file-based download, assuming that you already have MobiLink synchronization set up.

In most cases, you should use a download-only publication with your file-based download. Use a regular publication only when you need to do uploads with the same publication as you do file-based downloads.

### Overview of setting up file-based download

1. Create a file-definition database.  
See [“Creating the file-definition database” on page 295](#).
2. At the consolidated database, create scripts with a new script version.  
See [“Changes at the consolidated database” on page 295](#).
3. Create a download file.  
See [“Creating the download file” on page 296](#).
4. Apply the download file.  
See [“Synchronizing new remotes” on page 296](#).

### Other resources for getting started

- [“File-based download examples” on page 301](#)

## Creating the file-definition database

To set up file-based download, you create a **file-definition database**. This is a SQL Anywhere database that has the same synchronization tables and publications as your remote databases. It can be located anywhere. This database contains no data or state information. It does not have to be backed up or maintained; in fact, you can delete it and recreate it as needed.

The file-definition database must include the following:

- the same publications as the remote databases, the tables and columns used in the publication, the foreign key relationships and constraints of those tables and columns, and the tables required by those foreign key relationships.
- a MobiLink user name that identifies the group of remote databases that are to apply the download file. You use this group MobiLink user name in your synchronization scripts to identify the group of remote databases.

## Changes at the consolidated database

On the consolidated database, create a new script version for your file-based download, and implement any scripts required by your existing synchronization system into it. Upload scripts are not required. This script

version is used only for file-based download. For this script version, all scripts that take MobiLink user names as parameters, instead, take a MobiLink user name that refers to a group of remote databases. This is the user name that is defined in the file-definition database.

For each script version that you have defined, implement a `begin_publication` script.

For timestamp-based downloads, implement a `modify_last_download_timestamp` script for each script version. How you implement this script depends on how much data you intend to send in each download file. For example, one approach is to use the earliest time that any user from the group last downloaded successfully. Remember that the `ml_username` parameter passed to this script is actually the group name.

### See also

- [“Script versions” on page 324](#)
- [“begin\\_publication connection event” on page 380](#)
- [“modify\\_last\\_download\\_timestamp connection event” on page 463](#)

## Creating the download file

The download file contains the data to be synchronized. To create the download file, set up your file-definition database and consolidated database as described above. Run `dbmlsync` with the `-bc` option and supply a file name with the extension `.df`. For example,

```
dbmlsync -c "uid=DBA;pwd=sql;eng=fbd1_eng;dbf=fdef.db" -v+  
-e "sv=filebased" -bc file1.df
```

You can also choose to specify options when you create the download file:

- **-be option** Use `-be` to add a string to the download file that can be accessed at the remote database using the `sp_hook_dbmlsync_validate_download_file` stored procedure.  
See [“-be option” \[MobiLink - Client Administration\]](#) and [“sp\\_hook\\_dbmlsync\\_validate\\_download\\_file” \[MobiLink - Client Administration\]](#).
- **-bg option** Use the `-bg` option to create a download file that can be used by remotes that have never synchronized.

## Synchronizing new remotes

If you want to apply a download file to a remote database that has never synchronized using MobiLink, then before you apply the download file you need to either perform a normal synchronization on the remote database or use the `dbmlsync -bg` option when creating the download file.

For timestamp-based synchronization, doing either of these two things causes the download of an initial snapshot of the data. For both timestamp and snapshot based synchronization, this step sets the generation number to the value that is generated by the `begin_publication` script on the consolidated database.

## Perform a normal synchronization

You can prepare a remote database to receive download files by performing a synchronization that does not use a download file.

## Use the `-bg` option

Alternatively, you can create a download file with the `-bg` option to use with remotes that haven't yet synchronized. You apply this initial download file to prepare the remote database for file-based synchronization.

- **Snapshot downloads** If you are performing snapshot downloads, then the initial download file just needs to set the generation number. You may choose to include an initial snapshot of the data in this file, but since each snapshot download contains all the data and does not depend on previous downloads, this is not required.

For snapshot downloads, using the `-bg` option is straightforward. Just specify `-bg` in the `dbmlsync` command line when you create the download file. You can use the same script version to create the initial download file as you use for subsequent download files.

- **Timestamp-based downloads** If you are performing timestamp-based downloads, then the initial download must set the generation number on the remote database and include a snapshot of the data. With timestamp-based downloads, each download builds on previous ones. Each download file contains a last download timestamp. All rows changed on the consolidated after the file's last download timestamp are included in the file. To apply a file, a remote database must already have received all the changes that occurred before the file's last download timestamp. This is confirmed by checking that the file's last download timestamp is greater than or equal to the remote database's last download timestamp (the time up to which the remote database has received all changes from the consolidated database).

Before a remote can apply its first normal download file, it must receive all data changed before that file's last download timestamp and after January 1, 1900. The initial download file created with the `-bg` option must contain this data. The easiest way to select this data is to create a separate script version that uses the same `download_cursor`'s as your normal file-based synchronization script version but does not have a `modify_last_download_timestamp` script. If no `modify_last_download_timestamp` script is defined, then the last download timestamp for a file-based download defaults to January 1, 1900.

If you apply download files built with the `-bg` option to remote databases that have already synchronized, the `-bg` option causes the generation numbers on the remote database to be updated with the value on the consolidated database at the time the download file was created. This defeats the purpose of generation numbers, which is to prevent you from applying further file-based downloads until an upload has been performed in situations such as when recovering a consolidated database that is lost or corrupted.

See [“MobiLink generation numbers” on page 299](#).

## Validation checks

Before applying a download file to a remote database, dbmsync does several things to ensure that the synchronization is valid.

- dbmsync checks the download file to ensure that the file-definition database that was used to create it has:
  - the same publication as the remote database
  - the same tables and columns used in the publication
  - the same foreign key relationships and constraints as those tables and columns
- dbmsync checks to see if there is any data in the publication that has not been uploaded from the remote. If there is, the download file is not applied, because applying the download file could cause pending upload data to be lost.
- dbmsync checks the last download timestamp, next last download timestamp, and creation time of the download file to ensure that:
  - newer data on the remote database is not overwritten by older data contained in the download file.
  - a download file is not applied if applying it means that the remote database would miss some changes that have occurred on the consolidated database. This situation might occur if the remote did not apply previous file-based downloads.  
See [“Automatic validation” on page 298](#).
- Optionally, dbmsync checks the generation number in the remote database to ensure it matches the generation number in the download file.  
See [“MobiLink generation numbers” on page 299](#).
- Optionally, you can create custom validation logic with the `sp_hook_dbmsync_validate_download_file` stored procedure.  
For more information, see [“Custom validation” on page 300](#).

## Automatic validation

Before applying a download file, dbmsync performs special checks on the last download timestamp, next last download timestamp, download file creation time, and transaction log.

### Last download timestamp and next last download timestamp

Each download file contains all changes to be downloaded that occurred on the consolidated database between the file's last download timestamp, and its next last download timestamp. Both times are expressed in terms of the time at the consolidated database. By default the file's last download time is Jan 1, 1900 12:00 AM and the file's next last download timestamp is the time the download file was created. These defaults can be overridden by implementing the `modify_last_download_timestamp` and `modify_next_last_download_timestamp` scripts on the consolidated database.

A remote site can apply a download file only if the file's last download timestamp is less than or equal to the remote's last download timestamp. This ensures that a remote never misses operations that occur on the consolidated database. Usually when a file-based download fails based on this check, the remote has missed one or more download files. The situation can be corrected by applying the missing download files or by performing a full or download-only synchronization.

In addition, a remote site can apply a download file only if the file's next last download timestamp is greater than the remote's last download timestamp. The remote's last download timestamp is the time (at the consolidated database) up to which the remote has received all changes that are to be downloaded. The remote database's last download time is updated each time the remote successfully applies a download (normal or file-based). This check ensures that a download file is not applied if more recent data has already been downloaded. A common case where this could happen occurs when download files are applied out of order. For example, suppose a download file *F1.df* is created, and another file *F2.df* is created later. This check ensures that *F1.df* cannot be applied after *F2.df*, because that could allow newer data in *F2.df* to be overwritten with older data in *F1.df*.

When a file-based download fails based on the next last download timestamp, no additional action is required other than to delete the file. Synchronization succeeds once a new file is received.

### Creation time

The download file's creation time indicates the time at the consolidated database when creation of the file began. A download file can only be applied if the file's creation time is greater than the remote database's last upload time. The remote's last upload time is the time at the consolidated database when the remote's last successful upload was committed. This check ensures that data that has been uploaded after the creation of the download (and is newer than the download) is not overwritten by older data in the download file.

When a download file is rejected based on this check, no action is required. The remote site should be able to apply the next download file.

When an upload fails because dbmsync did not receive an acknowledgement after sending an upload to the MobiLink server, the remote database's last upload time may be incorrect. In this case, the creation time check cannot be performed and the remote is unable to apply download files until it completes a normal synchronization.

### Transaction log

Before applying a download file, dbmsync scans the remote database's transaction log and builds up a list of all changes that must be uploaded. Dbmsync only applies a download file if it does not contain any operations that affect rows with changes that must be uploaded.

## MobiLink generation numbers

Generation numbers provide a mechanism for forcing remote databases to upload data before applying any more download files. This is especially useful when a problem on the consolidated database has resulted in data loss and you must recover lost data from the remote databases.

On the remote database, a separate generation number is automatically maintained for each subscription. On the consolidated database, the generation number for each subscription is determined by the

begin\_publication script. Each time a remote performs a successful upload, it updates the remote generation number with the value set by the begin\_publication script in the consolidated database.

Each time a download file is created, the generation number set by the begin\_publication script is stored in the download file. A remote site only applies a download file if the generation number in the file is equal to the generation number stored in the remote database.

**Note**

Whenever the generation number generated by the begin\_publication script for a file-based download changes, the remote databases must perform a successful upload before they can apply any new download files.

The sp\_hook\_dbmsync\_validate\_download\_file stored procedure can be used to override the default checking of the generation number.

For more information about managing MobiLink generation numbers, see:

- [“begin\\_publication connection event” on page 380](#)
- [“end\\_publication connection event” on page 423](#)
- [“sp\\_hook\\_dbmsync\\_validate\\_download\\_file” \[MobiLink - Client Administration\]](#)

## Custom validation

You can create custom validation logic to determine if a download file should be applied to a remote database. You do this with the sp\_hook\_dbmsync\_validate\_download\_file stored procedure. With this stored procedure, you can both reject a download file and override the default checking of the generation number.

You can use the dbmsync -be option to embed a string in the file. You use the -be option against the file-definition database when you create the download file. This string is passed to the sp\_hook\_dbmsync\_validate\_download\_file through the #hook\_dict table, and can be used in your validation logic.

For more information, see [“sp\\_hook\\_dbmsync\\_validate\\_download\\_file” \[MobiLink - Client Administration\]](#).



## File-based download examples

This section contains two examples. Each sets up a file-based download synchronization using a consolidated database with only one table. The first is a simple snapshot example and the second is a slightly more involved timestamp-based example.

### Snapshot example

This example implements file-based download for snapshot synchronization. It sets up the three databases that are required by the file-based download, and then demonstrates how to download data. This example is presented in such a way that you can either just read through it, or you can cut and paste the text to run the sample.

#### Create databases for the sample

The following commands create the three databases used in the example: a consolidated database, a remote database, and a file-definition database.

```
dbinit scon.sdb
dbinit sremote.db
dbinit sfdef.db
```

The following commands start the three databases and create a data source name for MobiLink to use to connect to the consolidated database.

```
dbeng11 -n sfdef_eng sfdef.db
dbeng11 -n scon_eng scon.sdb
dbeng11 -n sremote_eng sremote.db
dbdsn -y -w fbd_demo -c "eng=scon_eng;dbf=scon.sdb;uid=DBA;
pwd=sql;astart=off;astop=off"
```

Open Interactive SQL, connect to *scon.sdb* and run the MobiLink setup script. For example:

```
read "c:\Program Files\SQL Anywhere 11\MobiLink\setup\syncsa.sql"
```

Start the MobiLink server:

```
start mlsv11 -v+ -c "dsn=fbd_demo" -zu+ -ot scon.txt
```

#### Set up the snapshot example consolidated database

In this example, the consolidated database has one table, called T1. After connecting to the consolidated database, you can run the following SQL to create table T1:

```
CREATE TABLE T1 (
  pk INTEGER PRIMARY KEY,
  c1 INTEGER
);
```

The following code creates a script version called *filebased* and creates a download script for that script version.

```
CALL ml_add_table_script( 'filebased',
  'T1', 'download_cursor',
  'SELECT pk, c1 FROM T1' );
```

The following code creates a script version called normal and creates upload and download scripts for that script version.

```
CALL ml_add_table_script ( 'normal', 'T1',
  'upload_insert',
  'INSERT INTO T1 VALUES ({ml r.pk}, {ml r.c1})');

CALL ml_add_table_script( 'normal', 'T1',
  'upload_update',
  'UPDATE T1 SET c1 = {ml r.c1} WHERE pk = {ml r.pk} ' );

CALL ml_add_table_script( 'normal', 'T1',
  'upload_delete',
  'DELETE FROM T1 WHERE pk = {ml r.pk}' );

CALL ml_add_table_script( 'normal', 'T1',
  'download_cursor',
  'SELECT pk, c1 FROM T1' );

COMMIT;
```

The following command creates the stored procedure begin\_pub and specifies that begin\_pub is the begin\_publication script for both the "normal" and "filebased" script versions:

```
CREATE PROCEDURE begin_pub (
  INOUT generation_num integer,
  IN   username          varchar(128),
  IN   pubname           varchar(128) )
BEGIN
  SET generation_num=1;
END;

CALL ml_add_connection_script(
  'filebased',
  'begin_publication',
  '{ call begin_pub(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name} ) }' );

CALL ml_add_connection_script( 'normal',
  'begin_publication',
  '{ call begin_pub(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name} ) }' );
```

### Create the snapshot example remote database

In this example, the remote database also contains one table, called T1. Connect to the remote database and run the following SQL to create the table T1, a publication called P1, and a user called U1. The SQL also creates a subscription for U1 to P1.

```
CREATE TABLE T1 (
  pk INTEGER PRIMARY KEY,
  c1 INTEGER
);
```

```

CREATE PUBLICATION P1 (
    TABLE T1
);

CREATE SYNCHRONIZATION USER U1;

CREATE SYNCHRONIZATION SUBSCRIPTION
TO P1
FOR U1;

```

The following code creates an `sp_hook_dbmsync_validate_download_file` hook to implement user-defined validation logic in the remote database:

```

CREATE PROCEDURE sp_hook_dbmsync_validate_download_file()
BEGIN
    DECLARE udata varchar(256);
    SELECT value
        INTO udata
        FROM #hook_dict
        WHERE name = 'user data';
    IF udata <> 'ok' THEN
        UPDATE #hook_dict
        SET value = 'FALSE'
        WHERE name = 'apply file';
    END IF;
END

```

### Create the snapshot example file-definition database

A file-definition database is required in MobiLink systems that use file-based download. This database has the same schema as the remote databases being updated by file-based download, and it contains no data or state information. The file-definition database is used solely to define the structure of the data that is to be included in the download file. One file-definition database can be used for many groups of remote databases, each defined by its own MobiLink group user name.

The following code defines the file-definition database for this sample. It creates a schema that is identical to the remote database, and also creates:

- a publication called P1 that publishes all rows of the T1 table. The same publication name must be used in the file-definition database and the remote databases.
- a MobiLink user called G1. This user represents all the remotes that are to be updated in the file-based download.
- a subscription to the publication.

You must connect to `sfdef.db` before running this code.

```

CREATE TABLE T1 (
    pk INTEGER PRIMARY KEY,
    c1 INTEGER
);

CREATE PUBLICATION P1 (
    TABLE T1
);

CREATE SYNCHRONIZATION USER G1;

CREATE SYNCHRONIZATION SUBSCRIPTION

```

```
TO P1
FOR G1;
```

### Prepare for initial synchronization

To prepare your new remote database so that you can apply a download file, you need to either perform a normal synchronization or create the download file with the `dbmsync -bg` option. This example shows you how to initialize your new remote database by performing a normal synchronization.

You can perform an initial synchronization of the remote database with the script version called `normal` that was created earlier:

```
dbmsync -c "uid=DBA;pwd=sql;eng=sremote_eng;
dbf=sremote.db" -v+ -e "sv=normal"
```

### Demonstrate the snapshot example file-based download

Connect to the consolidated database and insert some data that is synchronized by file-based download, such as the following:

```
INSERT INTO T1 VALUES( 1, 1 );
INSERT INTO T1 VALUES( 2, 4 );
INSERT INTO T1 VALUES( 3, 9 );
COMMIT;
```

The following command must be run on the computer that holds the file-definition database. It does the following:

- The `dbmsync -bc` option creates the download file, and names it *file1.df*.
- The `-be` option includes the string "OK" in the download file that is accessible to the `sp_dbmsync_validate_download_file` hook.

```
dbmsync -c
"uid=DBA;pwd=sql;eng=sfdef_eng;dbf=sfdef.db"
-v+ -e "sv=filebased" -bc file1.df -be ok -ot fdef.txt
```

To apply the download file, run `dbmsync` with the `-ba` option on the remote database, supplying the name of the download file you want to apply:

```
dbmsync -c "uid=DBA;pwd=sql;eng=sremote_eng;
dbf=sremote.db" -v+ -ba file1.df -ot remote.txt
```

The changes are now applied to the remote database. Open Interactive SQL, connect to the remote database, and run the following SQL statement to verify that the remote has the data:

```
SELECT * FROM T1
```

### Clean up the snapshot example

The following commands stop all three database servers and erase the files.

```
del file1.df
mlstop -h -w
dbstop -y -c "eng=sfdef_eng; uid=DBA; pwd=sql"
dbstop -y -c "eng=scons_eng; uid=DBA; pwd=sql"
dbstop -y -c "eng=sremote_eng; uid=DBA; pwd=sql"
dberase -y sfdef.db
```

```
dberase -y scons.db
dberase -y sremote.db
```

## Timestamp-based example

This example implements file-based download for timestamp-based synchronization. It sets up the three databases and then demonstrates how to download data by file. This example is presented in such a way that you can either just read through it, or you can cut and paste the text to run the sample.

### Create databases for the sample

The following commands create the three databases used in the example: a consolidated database, a remote database, and a file-definition database.

```
dbinit tcons.db
dbinit tremote.db
dbinit tfdef.db
```

The following commands start the three databases and create a data source name for MobiLink to use to connect to the consolidated database.

```
dbeng11 -n tfdef_eng tfdef.db
dbeng11 -n tcons_eng tcons.db
dbeng11 -n tremote_eng tremote.db
dbdsn -y -w tfbd_demo -c "eng=tcons_eng;dbf=tcons.db;uid=DBA;
pwd=sql;astart=off;astop=off"
```

Open Interactive SQL, connect to *tcons.db* and run the MobiLink setup script. For example:

```
read "c:\Program Files\SQL Anywhere 11\MobiLink\setup\syncsa.sql"
```

Start the MobiLink server:

```
start mlsrv11 -v+ -c "dsn=tfbd_demo" -zu+ -ot tcons.txt
```

### Set up the timestamp example consolidated database

In this example, the consolidated database has one table, called T1. After connecting to the consolidated database, you can run the following code to create T1:

```
CREATE TABLE T1 (
  pk INTEGER PRIMARY KEY,
  c1 INTEGER,
  last_modified TIMESTAMP DEFAULT TIMESTAMP
);
```

The following code defines a script version called normal with a minimal number of scripts. This script version is used for synchronizations that do **not** use file-based download.

```
CALL ml_add_table_script( 'normal', 'T1',
  'upload_insert',
  'INSERT INTO T1( pk, c1) VALUES( {ml r.pk}, {ml r.c1} )' );

CALL ml_add_table_script( 'normal', 'T1',
  'upload_update',
  'UPDATE T1 SET c1 = {ml r.c1} WHERE pk = {ml r.pk} ' );

CALL ml_add_table_script( 'normal', 'T1',
```

```
'upload_delete',
'DELETE FROM T1 WHERE pk = {ml r.pk}' );

CALL ml_add_table_script( 'normal', 'T1',
'download_cursor',
'SELECT pk, c1 FROM T1
WHERE last_modified >= {ml s.last_table_download}' );
```

The following code sets the generation number for all subscriptions to 1. It is good practice to use generation numbers in case your consolidated database ever becomes lost or corrupted and you need to force an upload.

```
CREATE PROCEDURE begin_pub (
    INOUT generation_num integer,
    IN   username          varchar(128),
    IN   pubname           varchar(128) )
BEGIN
    SET generation_num = 1;
END;

CALL ml_add_connection_script( 'normal',
'begin_publication',
'{ call begin_pub(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name},
    {ml s.last_publication_upload},
    {ml s.last_publication_download} ) }' );

COMMIT;
```

The following code defines the script version called filebased. This script version is used to create file-based download.

```
CALL ml_add_connection_script( 'filebased',
'begin_publication',
'{ call begin_pub(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name} ) }' );

CALL ml_add_table_script( 'filebased', 'T1',
'download_cursor',
'SELECT pk, c1 FROM T1
WHERE last_modified >= {ml s.last_table_download}' );
```

The following code sets the last download time so that all changes that occurred within the last five days are included in download files. Any remote that has missed all the download files created in the last five days have to perform a normal synchronization before being able to apply any more file-based downloads.

```
CREATE PROCEDURE ModifyLastDownloadTimestamp(
    INOUT last_download_timestamp TIMESTAMP,
    IN   ml_username              VARCHAR(128) )
BEGIN
    SELECT dateadd( day, -5, CURRENT TIMESTAMP )
    INTO last_download_timestamp;
END;

CALL ml_add_connection_script( 'filebased',
'modify_last_download_timestamp',
'CALL ModifyLastDownloadTimestamp(
    {ml s.last_download}, {ml s.username} )' );
```

```
COMMIT;
```

### Create the timestamp example remote database

In this example, the remote database also contains one table, called T1. After connecting to the remote database, run the following code to create table T1, a publication called P1, and a user called U1. The code also creates a subscription for U1 to P1.

```
CREATE TABLE T1 (
  pk INTEGER PRIMARY KEY,
  c1 INTEGER
);

CREATE PUBLICATION P1 (
  TABLE T1
);

CREATE SYNCHRONIZATION USER U1;

CREATE SYNCHRONIZATION SUBSCRIPTION
TO P1
FOR U1;
```

The following code defines an `sp_hook_dbmsync_validate_download_file` stored procedure. This stored procedure prevents the application of download files that do not have the string "ok" embedded in them.

```
CREATE PROCEDURE sp_hook_dbmsync_validate_download_file()
BEGIN
  DECLARE udata varchar(256);

  SELECT value
  INTO udata
  FROM #hook_dict
  WHERE name = 'user data';

  IF udata <> 'ok' THEN
    UPDATE #hook_dict
    SET value = 'FALSE'
    WHERE name = 'apply file';
  END IF;
END
```

### Create the timestamp example file-definition database

The following code defines the file-definition database for the timestamp example. It creates a table, a publication, a user, and a subscription for the user to the publication.

```
CREATE TABLE T1 (
  pk INTEGER PRIMARY KEY,
  c1 INTEGER
);

CREATE PUBLICATION P1 (
  TABLE T1
);

CREATE SYNCHRONIZATION USER G1;

CREATE SYNCHRONIZATION SUBSCRIPTION
TO P1
FOR G1;
```

## Prepare for initial synchronization

To prepare your new remote database so that you can apply a download file, you need to either perform a normal synchronization or create the download file with the `dbmlsync -bg` option. This example shows you how to use `-bg`.

The following code defines a script version called `filebased_init` for the consolidated database. This script version has a single `begin_publication` script.

```
CALL ml_add_table_script(
  'filebased_init', 'T1', 'download_cursor',
  'SELECT pk, c1 FROM T1' );

CALL ml_add_connection_script(
  'filebased_init',
  'begin_publication',
  '{ call begin_pub(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name} ) }' );

COMMIT;
```

The following two command lines create and apply an initial download file using the script version called `filebased_init` and the `-bg` option.

```
dbmlsync -c "uid=DBA;pwd=sql;eng=tfdef_eng;dbf=tfdef.db"
-v+ -e "sv=filebased_init" -bc tfile1.df -be ok -bg
-ot tfdef1.txt

dbmlsync -c "uid=DBA;pwd=sql;eng=tremote_eng;dbf=tremote.db"
-v+ -ba tfile1.df -ot tremote.txt
```

## Demonstrate the timestamp example file-based download

Connect to the consolidated database and insert some data that is synchronized by file-based download, such as the following:

```
INSERT INTO T1(pk, c1) VALUES( 1, 1 );
INSERT INTO T1(pk, c1) VALUES( 2, 4 );
INSERT INTO T1(pk, c1) VALUES( 3, 9 );
commit;
```

The following command line creates a download file containing the new data.

```
dbmlsync -c
"uid=DBA;pwd=sql;eng=tfdef_eng;dbf=tfdef.db"
-v+ -e "sv=filebased" -bc tfile2.df -be ok -ot tfdef2.txt
```

The following command line applies the download file to the remote database.

```
dbmlsync -c "uid=DBA;pwd=sql;eng=tremote_eng;dbf=tremote.db"
-v+ -ba tfile2.df -ot tfdef3.txt
```

The changes are now applied to the remote database. Open Interactive SQL, connect to the remote database, and run the following SQL statement to verify that the remote has the data:

```
SELECT * FROM T1
```



**Clean up the timestamp example**

The following commands stop all three database servers and then erase the files.

```
del tfile1.df
mlstop -h -w
dbstop -y -c "eng=tfdef_eng; uid=DBA; pwd=sql"
dbstop -y -c "eng=tcons_eng; uid=DBA; pwd=sql"
dbstop -y -c "eng=tremote_eng; uid=DBA; pwd=sql"
dberase -y tfdef.db
dberase -y tcons.db
dberase -y tremote.db
```

---

# MobiLink Events

This section describes how to write scripts for MobiLink events.

---

Writing synchronization scripts .....	313
Synchronization events .....	341



---

# Writing synchronization scripts

## Contents

Introduction to synchronization scripts .....	314
Scripts and the synchronization process .....	317
Script types .....	318
Script parameters .....	320
Script versions .....	324
Required scripts .....	326
Adding and deleting scripts .....	327
Writing scripts to upload rows .....	330
Writing scripts to download rows .....	333
Writing scripts to handle errors .....	338

---

## Introduction to synchronization scripts

You control the synchronization process by writing synchronization scripts and storing or referencing them in MobiLink system tables in the consolidated database. You can write scripts in SQL, Java, or .NET.

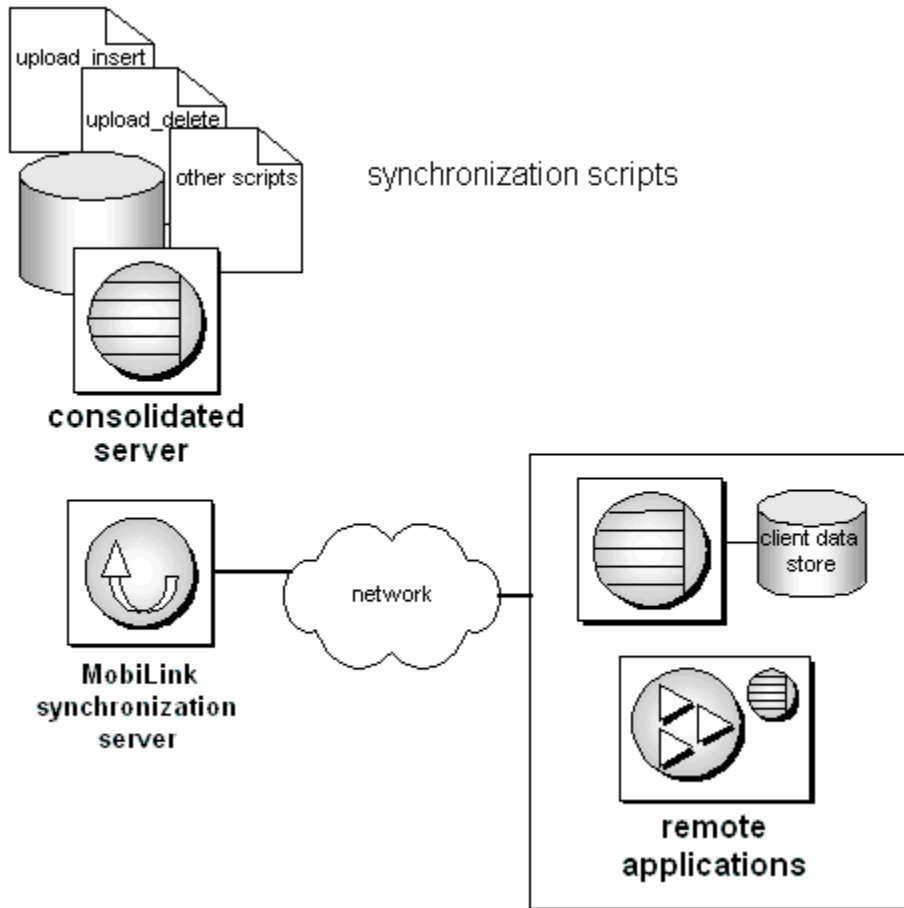
**MobiLink synchronization logic** is specified with synchronization scripts. Scripts define:

- how data that is uploaded from the remote database should be applied to the consolidated database
- what data should be downloaded from the consolidated database

Scripts can be individual statements or stored procedure calls. They are stored or referenced in your consolidated database. To add scripts to the consolidated database, you can use Sybase Central or you can use system procedures.

There should be no implicit or explicit commit or rollback in your SQL synchronization scripts or the procedures or triggers that are called from your SQL synchronization scripts. COMMIT or ROLLBACK statements within SQL scripts alter the transactional nature of the synchronization steps. If you use them, MobiLink cannot guarantee the integrity of your data in the event of a failure.

During synchronization, the MobiLink server reads the scripts and executes them against the consolidated database.



The synchronization process has multiple steps. A unique event identifies each step. You control the synchronization process by writing scripts associated with some of these events. You write a script only when some particular action must occur at a particular event. The MobiLink server executes each script when its associated event occurs. If you do not define a script for a particular event, the MobiLink server simply proceeds to the next step.

For example, one event is `begin_upload_rows`. You can write a script and associate it with this event. The MobiLink server reads this script when it is first needed, and executes it during the upload phase of synchronization. If you write no script, the MobiLink server proceeds immediately to the next step, which is processing the uploaded rows.

Some scripts, called table scripts, are associated not only with an event, but also with a particular table in the remote database. The MobiLink server performs some tasks on a table-by-table basis; for example, downloading rows. You can have many scripts associated with the same event, but each with different application tables. Alternatively, you can define many scripts for some application tables, but none for others.

For an overview of events, see [“The synchronization process”](#) [*MobiLink - Getting Started*].

For a description of every script you can write, see [“Synchronization events”](#) on page 341.

You can write scripts in SQL, Java, or .NET. This chapter applies to all kinds of scripts, but focuses on how to write synchronization scripts in SQL.

For a description and comparison of SQL, Java, and .NET, see [“Options for writing server-side synchronization logic” \[MobiLink - Getting Started\]](#).

For information about writing scripts in .NET, see [“Writing synchronization scripts in .NET” on page 589](#).

For information about writing scripts in Java, see [“Writing synchronization scripts in Java” on page 527](#).

For information about how to implement synchronization scripts, see [“Synchronization techniques” on page 127](#).

## Simple synchronization script

MobiLink provides many events that you can exploit, but it is not mandatory to provide scripts for each event. In a simple synchronization model, you may need only a few scripts.

Downloading all the rows from the table to each remote database synchronizes the ULProduct table in the CustDB sample application. In this case, no additions are permitted at the remote databases. You can implement this simple form of synchronization with a single script; in this case only one event has a script associated with it.

The MobiLink event that controls the rows to be downloaded during each synchronization is named the `download_cursor` event. Cursor scripts must contain `SELECT` statements. The MobiLink server uses these queries to define a cursor. In the case of a `download_cursor` script, the cursor selects the rows to be downloaded to one particular table in the remote database.

In the CustDB sample application, there is a single `download_cursor` script for the ULProduct table in the sample application, which consists of the following query:

```
SELECT prod_id, price, prod_name
FROM ULProduct
```

This query generates a result set. The rows that make up this result set are downloaded to the client. In this case, all the rows of the table are downloaded.

The MobiLink server knows to send the rows to the ULProduct application table because this script is associated with both the `download_cursor` event and ULProduct table by the way it is stored in the consolidated database. Sybase Central allows you to make these associations.

In this example, the query selects data from a consolidated table also named ULProduct. The names need not match. You could, instead, download data to the ULProduct application table from any table, or any combination of tables, in the consolidated database by rewriting the query.

You can write more complicated synchronization scripts. For example, you could write a script that downloads only recently modified rows, or one that provides different information to each remote database.



## Scripts and the synchronization process

Each script corresponds to a particular event in the synchronization process. You write a script only when some action must occur. All unnecessary events can be left undefined.

The two principal parts of the process are the processing of uploaded information and the preparation of rows for downloading.

The MobiLink server reads and prepares each script once, when it is first needed. The script is then executed whenever the event is invoked.

### The sequence of events

For information about the full sequence of MobiLink events, see [“Overview of MobiLink events” on page 343](#).

For the details of upload processing, see [“Writing scripts to upload rows” on page 330](#).

For the details of download processing, see [“Writing scripts to download rows” on page 333](#).

### Notes

- MobiLink technology allows multiple clients to synchronize concurrently. In this case, each client uses a separate connection to the consolidated database.
- The `begin_connection` and `end_connection` events are independent of any one synchronization as one connection can handle many synchronization requests. These scripts have no parameters. These are examples of connection-level scripts.
- Some events are invoked only once for each synchronization and have a single parameter. This parameter is the user name, which uniquely identifies the MobiLink client that is synchronizing. These are also examples of connection-level scripts.
- Some events are invoked once for each table being synchronized. Scripts associated with these events are called table-level scripts. They provide two parameters. The first is the user name supplied in the call to the synchronization function, and the second is the name of the table in the remote database being synchronized.

While each table can have its own table scripts, you can also write table-level scripts that are shared by several tables.

- Some events, such as `begin_synchronization`, occur at both the connection level and the table level. You can supply both connection and table scripts for these events.
- The COMMIT statements illustrate how the synchronization process is broken up into distinct transactions.
- Errors are a separate event that can occur at any point within the synchronization process. Errors are handled using the following script.

```
handle_error( error_code, error_message, user_name, table_name )
```

For reference material, including details about each script and its parameters, see [“Synchronization events” on page 341](#).

## Script types

Synchronization scripts can apply to the entire connection or to specific tables.

- **connection-level scripts** These scripts perform actions that are connection-specific or synchronization-specific and that are independent of any one remote table. These scripts are used in conjunction with other scripts when implementing more complex synchronization schemes.

See [“Connection scripts” on page 318](#).

- **table-level scripts** These scripts perform actions specific to one synchronization and one particular remote table. These scripts are used in conjunction with other scripts when implementing more complex synchronization schemes such as conflict resolution.

See [“Table scripts” on page 318](#).

## Connection scripts

Connection-level 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.

Connection scripts control actions centered on connecting and disconnecting, and synchronization-level event actions such as beginning and ending the upload or download process. Some connection scripts have related table scripts. These connection scripts are always invoked regardless of the tables being synchronized.

You only need to write a connection-level script when some action must occur at a particular event. You may need to create scripts for only a few events. The default action at any event is for the MobiLink server to perform no actions. Some simple synchronization schemes need no connection scripts.

### **ml\_global script version**

To save you from defining the same scripts multiple times, you can define connection-level scripts once and then re-use them. You do this by defining a script version called `ml_global`.

See [“ml\\_global script version” on page 325](#).

## Table scripts

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.

The synchronization scripts for a given table can refer to any table (or a combination of tables) in the consolidated database. You can use this feature to fill a particular remote table with data stored in one or more consolidated tables, or to store data uploaded from a single remote table into multiple tables in the consolidated database.

**Table names need not match**

The names of tables in the remote databases need not match the names of the tables in the consolidated database. The MobiLink server determines which scripts are associated with a table by looking up the remote table name in the ml\_table system table.

## Script parameters

Most synchronization scripts can receive parameters from the MobiLink server. For details about the parameters you can use in each script, see [“Synchronization events” on page 341](#).

You can specify parameters in your SQL scripts in one of two ways:

- question marks
- named script parameters

### Script parameters represented by question marks

Representing parameters with question marks is an ODBC convention. To use question marks in your MobiLink SQL scripts, place a single question mark in your script for each parameter. The MobiLink server replaces each question mark with the value of a parameter. It substitutes values in the order the parameters appear in the script definition.

The parameters must be in the order specified in [“Synchronization events” on page 341](#). Some parameters are optional. A parameter is optional only if no subsequent parameters are specified. For example, you must use parameter 1 if you want to use parameter 2.

### Named script parameters

MobiLink provides named parameters that you can use instead of question marks in your scripts. Named parameters have the following advantages:

- Named parameters allow you to specify any subset of the available parameters in any order.
- With the exception of in/out parameters, you can specify the same named parameter more than once within a script.
- When you use named parameters, you can specify the remote ID in your scripts. This is the only way to specify the remote ID in scripts.
- You can create your own named parameters. See [“User-defined named parameters” on page 322](#).

You cannot mix named parameters and question marks in a single script.

There are four types of MobiLink named parameters. To specify a named parameter, you must prefix it with its type, as follows:

Type of named parameter	Prefix	Examples
System parameters.	s.	{ml s.remote_id}
Row parameters. (The column name. If the column contains spaces, enclose it in double quotes or square brackets.)	r.	{ml r.cust_id} {ml r."Column name"}

Type of named parameter	Prefix	Examples
Old row parameters. (Only used in upload_update scripts to specify the pre-image column values. If the column name contains spaces, enclose it in double quotes or square brackets.)	o.	{ml o.cust_name} {ml o."Column name"}
Authentication parameters. See <a href="#">“Authentication parameters” on page 323</a> .	a.	{ml a.1}
User-defined parameters. See <a href="#">“User-defined named parameters” on page 322</a> .	ui.	{ml ui.varname}

To reference a script parameter by name, enclose the parameter in curly braces and prefix it with ml, as in **{ml parameter}**. For example, **{ml s.action\_code}**. The curly brace notation is an ODBC convention.

For convenience, you can enclose a larger section of code in the curly braces, as long as the section of code does not contain any schema names with the same name as a MobiLink script parameter. For example, both of the following upload\_insert scripts are valid and equivalent:

```
INSERT INTO t ( id, c0 ) VALUES( {ml r.id}, {ml r.c0} )
```

and

```
{ml INSERT INTO t ( id, c0 ) VALUES( r.id, r.c0 ) }
```

#### Note

To use named row parameters when the columns in your remote database were not generated by the MobiLink **Create Synchronization Model Wizard**, you need to use the ml\_add\_column system procedure to store column information in the consolidated database. See [“ml\\_add\\_column system procedure” on page 666](#).

## Commenting script parameters

The following forms of comments are recognized:

- Double hyphen prefix (--)
- Double forward slash prefix (//)
- Block commenting (/\* \*/)

The first two forms cause the script text to be ignored until the end of a line. The third form causes all script text between the /\* prefix and the \*/suffix to be ignored. Block commenting cannot be nested.

Any other type of vendor-specific comment is not recognized and should not be used to comment references to a named parameter.

## User-defined named parameters

You can also define your own parameters. These are especially useful for RDBMSs that don't allow user-defined variables.

User-defined parameters are defined (and set to null) when first referenced. They must start with `ui` and a period (`ui.`). A user-defined parameter lasts for one synchronization—it is set to null at the start of every synchronization. User-defined parameters are in/out.

A typical use of user-defined parameters is to access state information without having to store it in a table (requiring potentially complex joins).

### Example

For example, assume you create a stored procedure called `MyCustomProc` that sets a variable called `var1` to `custom_value`:

```
CREATE PROCEDURE MyCustomProc(  
  IN username (VARCHAR 128), INOUT var1 (VARCHAR 128)  
)  
begin  
  SET var1 = 'custom_value';  
end
```

The following `begin_connection` script defines the user-defined parameter `var1` and sets the value to `custom_value`:

```
CALL ml_add_connection_script (  
  'version1',  
  'begin_synchronization',  
  '{call MyCustomProc( {ml s.username}, {ml ui.var1} )}' );
```

The following `begin_upload` script references `var1`, whose value is `custom_value`:

```
CALL ml_add_connection_script (  
  'version1',  
  'begin_upload',  
  'update SomeTable set some_column = 123 where some_other_column = {ml  
  ui.var1}' );
```

Assume you have another stored procedure called `MyPFDPProc` that defines its first parameter to in/out. The following `prepare_for_download` script changes the value of `var1` to `pdf_value`:

```
CALL ml_add_connection_script (  
  'version1',  
  'prepare_for_download',  
  '{call MyPFDPProc( {ml ui.var1} )}' );
```

The following `begin_download` script references `var1`, whose value is now `pdf_value`:

```
CALL ml_add_connection_script (  
  'version1',  
  'begin_download',  
  'insert into SomeTable values( {ml s.username}, {ml ui.var1} )' );
```

## Authentication parameters

In MobiLink scripts, authentication parameters are named parameters that are prefaced with the letter a, such as {ml a.1}. The parameters must be numbers starting at 1, with a limit of 255. The values are sent up from MobiLink clients.

When used in the `authenticate_*` scripts, authentication parameters pass authentication information.

Authentication parameters can be used in all other events (except `begin_connection` and `end_connection`) to pass information from MobiLink clients. This technique is a convenient way to do something that you could otherwise do by creating and populating a table.

On SQL Anywhere remotes, you pass the information with the `dbmlsync -ap` option. On UltraLite remotes, you pass the information with `auth_parms` and `num_auth_parms`.

### See also

- “Script parameters” on page 320
- `dbmlsync`: “-ap option” [*MobiLink - Client Administration*]
- UltraLite: “Authentication Parameters synchronization parameter” [*UltraLite - Database Management and Reference*] and “Number of Authentication Parameters parameter” [*UltraLite - Database Management and Reference*]

### Example

For UltraLite remote databases, pass the parameters using the `num_auth_parms` and `auth_parms` fields in the `ul_synch_info` struct. `num_auth_parms` is a count of the number of parameters, from 0 to 255. `auth_parms` is a pointer to an array of strings. To prevent the strings from being viewed as plain text, the strings are sent in the same way as passwords. If `num_auth_parms` is 0, set `auth_parms` to null. The following is an example of passing parameters in UltraLite:

```
ul_char * Params[ 3 ] = { UL_TEXT( "param1" ),
    UL_TEXT( "param2" ), UL_TEXT( "param3" ) };

...
info.num_auth_parms = 3;
info.auth_parms = Params;
```

For SQL Anywhere remote databases, you pass authentication parameters using the `dbmlsync -ap` option, in a comma-separated list. For example, the following command line passes three parameters:

```
dbmlsync -ap "param1,param2,param3"
```

On the server, you reference the scripts in the order in which they were sent up. In this example, the `authenticate_parameters` script could be:

```
CALL my_auth_parm (
    {ml s.authentication_status},
    {ml s.remote_id},
    {ml s.username},
    {ml a.1},
    {ml a.2},
    {ml a.3}
)
```

## Script versions

Scripts are organized into groups called **script versions**. By specifying a particular version, MobiLink clients can select which set of synchronization scripts are used to process the upload and prepare the download.

For information about how to add a script version to the consolidated database, see [“Adding a script version” on page 325](#).

### Application of script versions

Script versions allow you to organize your scripts into sets, which are run under different circumstances. This ability provides flexibility and is especially useful in the following circumstances:

- **Customizing applications** Using a different set of scripts to process information from different types of remote users. For example, you could write a different set of scripts for use when managers synchronize their databases than would be used for other people in the organization. Although you could achieve the same functionality with one set of scripts, these scripts would be more complicated.
- **Upgrading applications** When you want to upgrade a database application, new scripts may be needed because the new version of your application may handle data differently. New scripts are almost always necessary when the remote database changes. It is usually impossible to upgrade all users simultaneously. MobiLink clients can request that a new set of scripts be used during synchronization. Since both old and new scripts can coexist on the server, all users can synchronize no matter which version of your application they are using.
- **Maintaining multiple applications** A single MobiLink server may need to synchronize two entirely different applications. For example, some employees may use a sales application, whereas others require an application designed for inventory control. When two applications require different sets of data, you can create two versions of the synchronization scripts, one version for each application.
- **Setting properties for the script version** You can set properties for your script version that can be referenced from classes in .NET or Java synchronization logic. See [“ml\\_add\\_property system procedure” on page 677](#).

### Assigning version names

A script version name is a string. You specify this name when you add a script to the consolidated database. For example, if you add your scripts with the `ml_add_connection_script` and the `ml_add_table_script` stored procedures, the script version name is the first parameter. Alternatively, if you add your scripts using Sybase Central, you are prompted for the version name.

You cannot use the following names for script versions: `ml_sis_1` or `ml_qa_1`. These names are used internally by MobiLink.

#### Caution

It is strongly recommended that your script version names do not start with `ml_`. Script versions starting with `ml_` are reserved for internal use.

### Specifying a version for a synchronization

If no script version is specified at the remote site when synchronization is initiated, the synchronization fails.



## ml\_global script version

You can create a script version called **ml\_global** that is used differently from other script versions. If you create a script version called `ml_global`, you define it once and then the connection scripts associated with are automatically used in all synchronizations. You never explicitly specify `ml_global` as a script version.

If you define a script in the `ml_global` script version and then you define a script for the same event in the script version that you specify for the synchronization, the specified script version is used. Scripts in the `ml_global` script version are only used if they are not defined in the primary script version that is being synchronized.

The `ml_global` script version can only contain connection-level scripts. It is not required, and may not be useful if you are using only one script version.

## Adding a script version

All scripts are associated with a script version. When working in Sybase Central Admin mode, you must add a version name to your consolidated database before you can add any connection scripts. When adding scripts with system procedures, a new version name is automatically added with the script. In Sybase Central Model mode, only one script version is allowed and it is by default given the same name as the model.

See [“Script versions” on page 324](#).

### To add a script version to a database (Sybase Central Admin mode)

1. In Sybase Central, choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Right-click the **Versions** folder and choose **File » New » Version**.
3. Follow the instructions in the **Create Script Version Wizard**.

### To remove a script version from a database (Sybase Central Admin mode)

1. In Sybase Central, choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Click the **Versions** folder.
3. In the right pane, right-click the version name and select **Delete**.
4. Click **Yes**.

### To add a script version to a database (system procedures)

- You can add a script version in the same operation as adding a connection script or table script.  
For more information, see [“System procedures to add or delete scripts” on page 664](#).

## Required scripts

When you run the MobiLink server, certain scripts are required. Which scripts are required is determined by whether you are doing a bi-directional, upload-only, or download-only synchronization.

For bi-directional or upload-only synchronization, MobiLink requires at least one of the following table scripts:

- `upload_delete`
- `upload_insert`
- `upload_new_row_insert`
- `upload_old_row_insert`
- `upload_update`
- Or, if you are processing the upload by direct row handling, MobiLink requires a script for the `handle_UploadData` connection event.

For bi-directional or download-only synchronization, MobiLink expects every table in the synchronization to have a download table script (`download_cursor` or `download_delete_cursor`). Or, if you are processing the download by direct row handling, MobiLink requires that you specify a `handle_DownloadData` connection script. Note that this script can be empty and you can process the download in any other event.

By default, if a required script is missing the synchronization aborts. You can override this behavior using the MobiLink server `-fr` option.

See [“-fr option” on page 70](#).

## Adding and deleting scripts

When you use the **Create Synchronization Model Wizard**, scripts are automatically added to the consolidated database when you deploy the model.

When you create synchronization scripts outside Sybase Central Model mode, you must add them to MobiLink system tables in the consolidated database. In the case of SQL scripts, the entire script is saved in the MobiLink system table. In the case of Java or .NET scripts, the method name is registered in the system table. The method for storing scripts and method names is similar.

See [“MobiLink server system tables” on page 693](#).

If you are using Sybase Central, you must add a synchronization version to the database before you can add individual scripts. See [“Adding a script version” on page 325](#).

### To add a connection script (Sybase Central Admin mode)

1. Choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Right-click **Connection Scripts** and choose **New » Connection Script**.
3. Follow the instructions in the **Create Connection Script Wizard**.

### To delete a connection script (Sybase Central Admin mode)

1. Choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Expand **Connection Scripts**.
3. Right-click a connection script and choose **Delete**.
4. Click **Yes**.

### To add a table script (Sybase Central Admin mode)

1. Choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Expand **Synchronized Tables**.
3. Right-click the table and choose **New » Table Script**.
4. Follow the instructions in the **Create Table Script Wizard**.

### To delete a table script (Sybase Central Admin mode)

1. Choose **Connections » Connect With MobiLink 11** and connect to the consolidated database.
2. Expand **Synchronized Tables**.
3. Expand the table.
4. Right-click the table script and choose **Delete**.
5. Click **Yes**.

### To add or delete all types of scripts (system procedures)

- You can add scripts to a consolidated database or delete scripts from a consolidated database using stored procedures that are installed when you set up your consolidated database.

For a description of the stored procedures that you can use to add or delete scripts, see:

- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)

## Direct inserts of scripts

In most cases, it is recommended that you use stored procedures or Sybase Central to insert scripts into the system tables. However, in some rare cases you may need to use an INSERT statement to directly insert the scripts. For example, older versions of some RDBMSs may have length limitations that make it difficult to use stored procedures.

For a complete description of the MobiLink system tables, see [“MobiLink server system tables” on page 693](#).

The format of the INSERT statements that are required to directly insert scripts can be found in the source code for the ml\_add\_connection\_script and ml\_add\_table\_script stored procedures. The source code for these stored procedures is located in the MobiLink setup scripts. There is a different setup script for each supported RDBMS. The setup scripts are all located in *install-dir\MobiLink\setup* and are called:

Consolidated database	Setup file
Adaptive Server Enterprise	<i>syncase.sql</i>
IBM DB2 mainframe	<i>syncd2m.sql</i> or <i>syncd2m_jcl.sql</i>
IBM DB2 LUW	<i>syncdb2.sql</i>
Microsoft SQL Server	<i>syncmss.sql</i>
MySQL	<i>syncmys.sql</i>
Oracle	<i>syncora.sql</i>
SQL Anywhere	<i>syncsa.sql</i>

## Ignoring scripts

If an upload stream contains insert, update, or delete data for a table that has no `upload_insert`, `upload_update`, and `upload_delete` script in the consolidated database, or if there is no download script (`download_cursor` and `download_delete_cursor` scripts) for the table, then the MobiLink server would either:

- complain about the missed scripts and abort the synchronization, if the MobiLink server was not started with `-fr`
- show messages to warn the users about data inconsistency, if the MobiLink server was started with `-fr`

The warning messages can be suppressed with the `-zwd` MobiLink server command option, however, this option suppresses the warning messages for all the synchronization tables.

Now, the MobiLink server treats any connection and table scripts that contain the prefix `--{ml_ignore}` differently. The MobiLink server recognizes these scripts as intentionally ignored scripts. More precisely, if an upload stream contains insert, update, or delete data for a synchronization table that has `upload_insert`, `upload_update`, or `upload_delete` script with the prefix `--{ml_ignore}`, the MobiLink server does not execute these scripts against the consolidated database and continues the synchronization without showing any error or warning messages, regardless of whether the server was started with or without the `-fr` option.

This logic will apply to downloads too. However, the MobiLink server aborts the synchronization or shows warning messages if the upload stream contains delete (insert or update) data for a synchronization table that has no intentionally ignored or actual `upload_delete` (`upload_insert` or `upload_update`) scripts, even if there are intentionally ignored or actual `upload_insert` and `upload_update` scripts for this table.

## Writing scripts to upload rows

To inform the MobiLink server on how to process the upload stream data received from the remote databases, you define upload scripts. You write separate scripts to handle rows that are updated, inserted, or deleted at the remote database. A simple implementation would perform corresponding actions (update, insert, delete) at the consolidated database.

The MobiLink server uploads data in a single transaction. For a description of the upload process, see [“Events during upload” on page 349](#).

For techniques for uploading rows in .NET synchronization logic, see [“Uploading or downloading rows” on page 600](#).

### Notes

- The begin\_upload and end\_upload scripts for each remote table hold logic that is independent of the individual rows being updated.
- The upload consists of single row inserts, updates, and deletes. These actions are typically performed using upload\_insert, upload\_update, and upload\_delete scripts.
- To prepare the upload for SQL Anywhere clients, the dbmlsync utility requires access to all transaction logs written since the last successful synchronization. See [“Transaction log files” \[MobiLink - Client Administration\]](#).
- When synchronizing a remote database using MobiLink client version 9.0 or earlier, or when using question marks instead of named parameters as placeholders in upload\_insert, upload\_new\_row\_insert, or upload\_old\_row\_insert events, the MobiLink server uses the column order of the table as it is defined in the remote database. The column order in the event statement must match the column order as it is defined in the remote database but table and column names in the consolidated database can be different from those in the remote database.

The following is an INSERT statement used only when emp\_name is defined before emp\_id in the remote database.

```
INSERT INTO emp (emp_name, emp_id)
VALUES (?, ?);
```

## Writing upload\_insert scripts

The MobiLink server uses this event during processing of the upload to handle rows inserted into the remote database.

The following is an INSERT statement used in an upload\_insert script.

```
INSERT INTO emp ( emp_id, emp_name )
VALUES ( { ml r.emp_id }, { ml r.emp_name } );
```

### Notes

- The upload\_new\_row\_insert and upload\_old\_row\_insert events accept remote\_id and user\_name as extra parameters. These parameters must appear prior to the full column list of the table.

**See also**

- [“Writing scripts to upload rows” on page 330](#)
- [“upload\\_insert table event” on page 504](#)

## Writing upload\_update scripts

The MobiLink server uses this event during processing of the upload to handle rows updated at the remote database. The following UPDATE statement illustrates the use of the upload\_update statement.

```
UPDATE emp
SET emp_name = {ml r.emp_name}
WHERE emp_id = {ml o.emp_id};
```

**Notes**

- When synchronizing a remote database using MobiLink client version 9.0 or earlier, or when using question marks instead of named parameters as placeholders, the number of parameters can be equal to one of the following:
  - The number of non-primary key columns + primary key columns.
  - $2 * (\text{the number of non-primary key columns} + \text{primary key columns})$ .

The column order must consist of non-primary key columns first, followed by one of the following:

- The primary key columns.
- All the columns.

**See also**

- [“Writing scripts to upload rows” on page 330](#)
- [“upload\\_update table event” on page 522](#)

## Writing upload\_delete scripts

The MobiLink server uses this event during processing of the upload to handle rows deleted from the remote database. The following statement shows how to use the upload\_delete statement.

```
DELETE FROM emp
WHERE emp_id = {ml r.emp_id};
```

**Notes**

- When synchronizing a remote database using MobiLink client version 9.0 or earlier, or when using question marks instead of named parameters as placeholders, the number of parameters must be equal to one of the following:
  - The number of primary key columns.
  - The number of all columns.

**See also**

- [“Writing scripts to upload rows” on page 330](#)
- [“upload\\_delete table event” on page 498](#)

## Writing upload\_fetch scripts

The upload\_fetch script is a SELECT statement that defines a cursor in the consolidated database table. This cursor is used to compare the old values of updated rows, as received from the remote database, against the value in the consolidated database. In this way, the upload\_fetch script identifies conflicts when updates are being processed.

Given a synchronized table defined as:

```
CREATE TABLE uf_example (  
  pk1 integer NOT NULL,  
  pk2 integer NOT NULL,  
  val varchar(200),  
  PRIMARY KEY( pk1, pk2 ));
```

Then one possible upload\_fetch script for this table is:

```
SELECT pk1, pk2, val  
FROM uf_example  
WHERE pk1 = {ml r.pk1} and pk2 = {ml r.pk2}
```

See [“upload\\_fetch table event” on page 500](#).

The MobiLink server requires the WHERE clause of the query in the upload\_fetch script to identify exactly one row in the consolidated database to be checked for conflicts.



## Writing scripts to download rows

There are two scripts that can be used for processing each table during the download transaction. These are the `download_cursor` script, which performs inserts and updates, and the `download_delete_cursor` script, which performs deletes.

These scripts are either `SELECT` statements or calls to procedures that return result sets. The MobiLink server downloads the result set of the script to the remote database. The MobiLink client automatically inserts or updates rows based on the `download_cursor` script result set, and deletes rows based on the `download_delete_cursor` event.

For more information about using stored procedures, see [“Downloading a result set from a stored procedure call” on page 162](#).

The MobiLink server downloads data in a single transaction. For a description of the download process, see [“Events during download” on page 351](#).

### Notes

- Like the upload, the download starts and ends with connection events. Other events are table-level events.
- If you change the `SendDownloadAck` setting to `ON`, the server behavior depends on the download acknowledgement mode you are using. For blocking download acknowledgement, if no confirmation of the download is received from the client, the entire download transaction is rolled back in the consolidated database. For non-blocking download acknowledgement, the download transaction is committed but the download timestamp update and acknowledgement scripts are not executed until the acknowledgement is received.

By default, `SendDownloadAck` is set to `OFF`.

See [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#), [“Send Download Acknowledgement synchronization parameter” \[UltraLite - Database Management and Reference\]](#), [“-nba option” on page 74](#), [“nonblocking\\_download\\_ack connection event” on page 471](#) and [“publication\\_nonblocking\\_download\\_ack connection event” on page 475](#).

- The `begin_download` and `end_download` scripts for each remote table hold logic that is independent of the individual rows being updated.
- For timestamp-based downloads, you specify the `last_download_timestamp` parameter to ensure that only changes since the last synchronization are downloaded. For example, the `download_cursor` or `download_delete_cursor` SQL script could include the line:

```
WHERE Customer.last_modified >= {ml s.last_table_download}
```

See [“Using last download times in scripts” on page 130](#).

- The download does not distinguish between inserts and updates. The script associated with the `download_cursor` event is a `SELECT` statement that defines the rows to be downloaded. The client detects whether the row exists or not and then it performs the appropriate insert or update operation.
- At the end of the download processing, the client automatically deletes rows if necessary to avoid referential integrity violations.

**Caution**

Do not synchronize shadow tables that were created by previous deployments (for example, tables ending with `_mod` or `_del` should not be synchronized). These tables are only needed by the consolidated database to track modified or deleted rows.

See “Referential integrity and synchronization” [*MobiLink - Getting Started*].

## Writing download\_cursor scripts

You write `download_cursor` scripts to download information from the consolidated database to your remote database. You must write one of these scripts for each table in the remote database for which you want to download changes. You can use other scripts to customize the download process, but no others are necessary.

- Each `download_cursor` script must contain a `SELECT` statement or a call to a procedure that contains a `SELECT` statement. The MobiLink server uses this statement to define a cursor in the consolidated database.
- The script must select all columns that correspond to the columns in the corresponding table in the remote database. The columns in the consolidated database can have different names than the corresponding columns in the remote database, but they must be of compatible types.

### Example

The following script could serve as a `download_cursor` script for a remote table that holds employee information. The MobiLink server would use this SQL statement to define the download cursor. This script downloads information about all the employees.

```
SELECT emp_id, emp_fname, emp_lname
FROM employee;
```

The MobiLink server passes specific parameters to some scripts. To use these parameters, you can use named parameters, or you can include a question mark in your SQL statement. In the latter case, the MobiLink server substitutes the value of the parameter before executing the statement against the consolidated database. The following script shows how you can use named parameters:

```
CALL ml_add_table_script(
    'Lab',
    'ULOrder',
    'download_cursor',
    'SELECT o.order_id, o.cust_id, o.prod_id, o.emp_id, o.disc, o.quant,
o.notes, o.status
FROM ULOrder o
WHERE o.last_modified >= {ml s.last_table_download}
AND o.emp_name = {ml s.username}' )
```

### Notes

- Row values can be selected from a single table or from a join of multiple tables.
- The script itself need not include the name of the remote table. The remote table need not have the same name as the table in the consolidated database. The name of the remote table is identified by an entry in

the `ml_table` MobiLink system table. In Sybase Central, the remote tables are listed together with their scripts.

- The rows in the remote table must contain the values of `emp_id`, `emp_fname`, and `emp_lname`. The remote columns must be in that order, although they can have different names. The columns in the remote database are in the same order as those in the reference database.
- All cursor scripts must select the columns in the same order as the columns are defined in the remote database. Where column names or table structure is different in the consolidated database, columns should be selected in the correct order for the remote database, or equivalently, the reference database. Columns are assigned to columns in the remote database based on their order in the `SELECT` statement.
- When you build an UltraLite application, the UltraLite generator creates a sample download script for each table in your UltraLite application. It inserts these sample scripts into your reference database. The example scripts assume that the consolidated database contains the same tables as your application. You must modify the sample scripts if your consolidated database differs in design, but these scripts provide a starting point.
- The `download_cursor` script must contain all columns in the same order as they are defined in the remote database.

#### See also

- [“download\\_cursor table event” on page 396](#)
- [“Partitioning rows among remote databases” on page 135](#)
- [“Writing download\\_delete\\_cursor scripts” on page 335](#)

## Writing download\_delete\_cursor scripts

You write `download_delete_cursor` scripts to delete rows from your remote database. You must write one of these scripts for each table in the remote database from which you want to delete rows during synchronization.

You cannot just delete rows from the consolidated database and have them disappear from remote databases. You need to keep track of the primary keys for deleted rows, so that you can select those primary keys with your `download_delete_cursor`. There are two common techniques for achieving this:

- **Logical deletes** Do not physically delete the row in the consolidated database. Instead, have a status column that keeps track of whether rows are valid. This simplifies the `download_delete_cursor`. However, the `download_cursor` and other applications may need to be modified to recognize and use the status column. If you have a last modified column that holds the time of deletion, and if you also keep track of the last download time for each remote, then you can physically delete the row once all remote download times are newer than the time of deletion.
- **Shadow table** For each table for which you want to track deletes, create a shadow table with two columns, one holding the primary key for the table, and the other holding a timestamp. Create a trigger that inserts the primary key and timestamp into the shadow table whenever a row is deleted. Your `download_delete_cursor` can then select from this shadow table. As with logical deletes, you can delete the row from the shadow table once all remote databases have downloaded the corresponding data.

The MobiLink server deletes rows in the remote database by selecting primary key values from the consolidated database and passing those values to the remote database. If the values match those of a primary key in the remote database, then that row is deleted.

- Each `download_delete_cursor` script must contain a `SELECT` statement or a call to a stored procedure that returns a result set. The MobiLink server uses this statement to define a cursor in the consolidated database.
- This statement must select all the columns that correspond to the primary key columns in the table in the remote database. The columns in the consolidated database can have different names than the corresponding columns in the remote database, but they must be of compatible types.
- The values must be selected in the same order as the corresponding columns are defined in the remote database. That order is the order of the columns in the `CREATE TABLE` statement used to make the table, not the order they appear in the statement that defines the primary key.
- If you delete a parent record, the child records are automatically deleted as well.

For more information about deleting child records, see [“Referential integrity and synchronization” \[MobiLink - Getting Started\]](#).

While each `download_delete_cursor` script must select all the column values present in the primary key of the corresponding remote table, it may also select all the other columns. This feature is present only for compatibility with older clients. Selecting the additional columns is less efficient, as the database server must retrieve more data. Unless the client is of an old design, the MobiLink server discards the extra values immediately. The extra values are downloaded only to older clients.

### Deleting all the rows in a table

When MobiLink detects a `download_delete_cursor` with a row that contains all nulls, it deletes all the data in the remote table. The number of nulls in the `download_delete_cursor` can be the number of primary key columns or the total number of columns in the table.

For example, the following `download_delete_cursor` SQL script deletes every row in a table in which there are two primary key columns. This example works for SQL Anywhere, Adaptive Server Enterprise, and Microsoft SQL Server databases.

```
SELECT NULL, NULL
```

In IBM DB2 LUW and Oracle consolidated databases, you must specify a dummy table to select null. For IBM DB2 LUW 7.1, you can use the following syntax:

```
SELECT NULL, NULL FROM SYSIBM.SYSDUMMY1
```

For Oracle consolidated databases, you can use the following syntax:

```
SELECT NULL, NULL FROM DUAL
```

### Examples

The following example is a `download_delete_cursor` script for a remote table that holds employee information. The MobiLink server uses this SQL statement to define the delete cursor. This script deletes information about all employees who are both in the consolidated and remote databases at the time the script is executed.

```
SELECT emp_id
FROM employee
```

The `download_delete_cursor` accepts the parameters `last_download` and `ml_username`. The following script shows how you can use each parameter to narrow your selection.

```
SELECT order_id
FROM ULOrder
WHERE last_modified >= {ml s.last_table_download}
AND status = 'Approved'
AND user_name = {ml s.username}
```

**Note**

For some consolidated databases, you may need to cast to the appropriate data type. See “[CAST function \[Data type conversion\]](#)” [*SQL Anywhere Server - SQL Reference*].

These examples could be inefficient in an organization with many employees. You can make the delete process more efficient by selecting only rows that could be present in the remote database. For example, you could limit the number of rows by selecting only those people who have recently been assigned a new manager. Another strategy is to allow the client application to delete the rows itself. This method is possible only when a rule identifies the unneeded rows. For example, rows might contain a timestamp that indicates an expiry date. Before you delete the rows, use the `STOP SYNCHRONIZATION DELETE` statement to stop these deletes being uploaded during the next synchronization. Be sure to execute `START SYNCHRONIZATION DELETE` immediately afterwards if you want other deletes to be synchronized in the normal fashion.

**Notes**

- The `download_delete_cursor` script must contain primary key columns in the same order as they are defined in the remote database.

**See also**

You can use the referential integrity checking built into all MobiLink clients to delete rows in an efficient manner. See “[Referential integrity and synchronization](#)” [*MobiLink - Getting Started*].

For more information about using `download_delete_cursor` scripts, see:

- “[download\\_cursor table event](#)” on page 396
- “[download\\_delete\\_cursor table event](#)” on page 400
- “[Handling deletes](#)” on page 156
- “[Temporarily stopping the synchronization of deletes](#)” on page 156
- “[STOP SYNCHRONIZATION DELETE statement \[MobiLink\]](#)” [*SQL Anywhere Server - SQL Reference*]
- “[Partitioning rows among remote databases](#)” on page 135
- “[Snapshot synchronization](#)” on page 133

## Writing scripts to handle errors

An error in a synchronization script occurs when an operation in the script fails while the MobiLink server is executing it. The DBMS returns a `SQLCODE` to the MobiLink server indicating the nature of the error. Each consolidated database DBMS has its own set of `SQLCODE` values.

When an error occurs, the MobiLink server invokes the `handle_error` event. You should provide a connection script associated with this event to handle errors. The MobiLink server passes several parameters to this script to provide information about the nature and context of the error, and requires an output value to tell it how to respond to the error.

### Error handling actions

Some actions you may want to take in an error-handling script are:

- Log the error in a separate table.
- Instruct the MobiLink server whether to ignore the error and continue, or rollback the synchronization, or rollback the synchronization and shut down the MobiLink server.
- Send an email message.

For more information, see [“handle\\_error connection event” on page 446](#).

## Reporting errors

Since errors can disrupt the natural progression of the synchronization process, it can be difficult to create a log of errors and their resolutions. The `report_error` script provides a means of accomplishing this task. The MobiLink server executes this script whenever an error occurs. If the `handle_error` script is defined, it is executed immediately prior to the reporting script.

The parameters to the `report_error` script are identical to those of the `handle_error` script, except that the `report_error` script can not modify the action code. Since the value of the action code is that returned by the `handle_error` script, this script can be used to debug error-handling problems.

This script typically consists of an insert statement, which records the values, perhaps with other data, such as the time or date, in a table for later reference. To ensure that this data is not lost, the MobiLink server always runs this script in a separate transaction and automatically commits the changes when this script completes.

See [“report\\_error connection event” on page 477](#).

### Example

The following `report_error` script, which consists of a single insert statement, adds the script parameters into a table, along with the current date and time. The script does not commit this change because the MobiLink server always does so automatically.

```
INSERT INTO errors
VALUES(
  CURRENT_DATE,
  {ml s.action_code},
```

```
{ml s.error_code},  
{ml s.error_message},  
{ml s.username},  
{ml s.table} );
```

## Handling multiple errors in a single SQL statement

ODBC allows multiple errors per SQL statement, and some RDBMSs make use of this feature. Microsoft SQL Server, for example, can have two errors for a single statement. The first is the actual error, and the second is usually an informational message telling you why the current statement has been terminated.

When a single SQL statement causes multiple errors, the `handle_error` script is invoked once per error. The MobiLink server uses the most severe action code (that is, the numerically greatest) to determine the action to take. The same applies to the `handle_error` script.

If the `handle_error` script itself causes a SQL error, then the default action code (3000) is assumed.

---



---

# Synchronization events

## Contents


Overview of MobiLink events .....	343
authenticate_file_transfer connection event .....	353
authenticate_parameters connection event .....	355
authenticate_user connection event .....	358
authenticate_user_hashed connection event .....	363
begin_connection connection event .....	367
begin_connection_autocommit connection event .....	368
begin_download connection event .....	369
begin_download table event .....	371
begin_download_deletes table event .....	374
begin_download_rows table event .....	377
begin_publication connection event .....	380
begin_synchronization connection event .....	383
begin_synchronization table event .....	385
begin_upload connection event .....	387
begin_upload table event .....	389
begin_upload_deletes table event .....	391
begin_upload_rows table event .....	394
download_cursor table event .....	396
download_delete_cursor table event .....	400
download_statistics connection event .....	403
download_statistics table event .....	406
end_connection connection event .....	409
end_download connection event .....	411
end_download table event .....	414
end_download_deletes table event .....	417
end_download_rows table event .....	420
end_publication connection event .....	423
end_synchronization connection event .....	426
end_synchronization table event .....	428
end_upload connection event .....	431

end_upload table event .....	433
end_upload_deletes table event .....	436
end_upload_rows table event .....	439
handle_DownloadData connection event .....	442
handle_error connection event .....	446
handle_odbc_error connection event .....	450
handle_UploadData connection event .....	454
modify_error_message connection event .....	460
modify_last_download_timestamp connection event .....	463
modify_next_last_download_timestamp connection event .....	466
modify_user connection event .....	469
nonblocking_download_ack connection event .....	471
prepare_for_download connection event .....	473
publication_nonblocking_download_ack connection event .....	475
report_error connection event .....	477
report_odbc_error connection event .....	480
resolve_conflict table event .....	483
synchronization_statistics connection event .....	486
synchronization_statistics table event .....	489
time_statistics connection event .....	492
time_statistics table event .....	495
upload_delete table event .....	498
upload_fetch table event .....	500
upload_fetch_column_conflict table event .....	502
upload_insert table event .....	504
upload_new_row_insert table event .....	506
upload_old_row_insert table event .....	509
upload_statistics connection event .....	512
upload_statistics table event .....	517
upload_update table event .....	522

---

## Overview of MobiLink events

When a synchronization request occurs and the MobiLink server decides that a new connection must be created, the `begin_connection` event is fired and synchronization starts.



Following the synchronization, the connection is placed in a connection pool, and MobiLink again waits for a synchronization request. Before a connection is eventually dropped from the connection pool, the `end_connection` event is fired. But if another synchronization request for the same version is received, then MobiLink handles the next synchronization request on the same connection. There are several events that affect the current synchronization.

### Transactions

Within each synchronization, the following transactions may occur. Each transaction is optional.

- authentication
- begin synchronization
- upload
  - You can specify multiple upload transactions with the `dbmlsync -tu` option.
- prepare for download
- download
- end synchronization
- non-blocking download acknowledgement

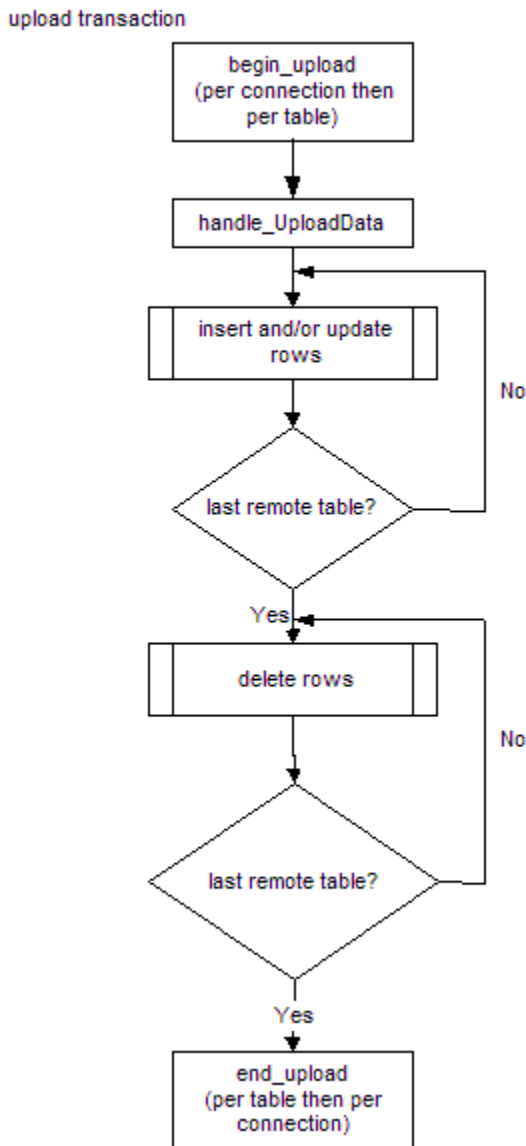
In addition, you can have two connection transactions. A begin connection transaction occurs right after a connection is made, and an end connection transaction occurs when the connection is closed.

The primary phases of a synchronization are the upload and download transactions. The events contained in the upload and download transactions are outlined below.

### The upload transaction

The upload transaction applies changes uploaded from a remote database.

The begin\_upload event marks the beginning of the upload transaction. The upload transaction is a two-part process. First, inserts and updates are uploaded for all remote tables, and second, deletes are uploaded for all remote tables.



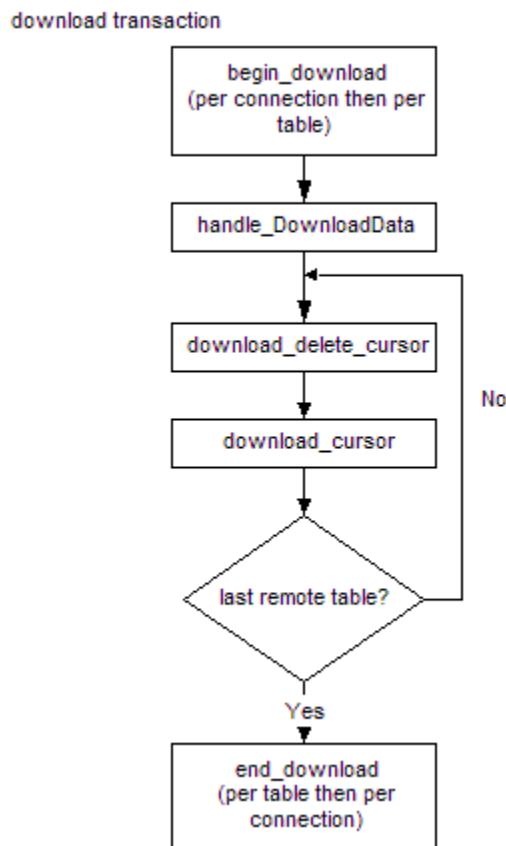
The `end_upload` event marks the end of the upload transaction.

See [“Writing scripts to upload rows”](#) on page 330.

### The download transaction

The download transaction fetches rows from the consolidated database. It begins with the `begin_download` event.

The download transaction is a two-part process. For each table, first deletes are downloaded, and then update/insert rows (upserts) are downloaded. The `end_download` event ends the download transaction.



See [“Writing scripts to download rows”](#) on page 333.

### The non-blocking download acknowledgement transaction

The non-blocking download acknowledgement transaction is only performed when MobiLink is in non-blocking download acknowledgement mode and a download acknowledgement is received. This transaction has two purposes. The scripts `publication_nonblocking_download_ack` and `nonblocking_download_ack` are run in this transaction; they help download status tracking. Secondly, download timestamps in the MobiLink system tables are updated during this transaction.

Note that this transaction is not performed on the same database connection as the other events for the target synchronization. This means that no connection level variables may be referenced in this transaction.

### Event overview in pseudocode

The following pseudocode provides an overview of the sequence in which events, and the scripts of the same names, are invoked. This representation of the MobiLink event model assumes a full synchronization (not upload-only or download-only) with no errors.

### Notes

- In most cases, if you have not defined a script for a given event, the default action is to do nothing.
- The `begin_connection` and `end_connection` events are **connection-level events**. They are independent of any single synchronization and have no parameters.
- Some events are invoked once per synchronization for each table being synchronized. Scripts associated with these events are called **table-level scripts**.

While each table can have its own table scripts, you can also write table-level scripts that are shared by several tables.

- Some events, such as `begin_synchronization`, occur at both the connection level and the table level. You can supply both connection and table scripts for these events.
- The COMMIT statements illustrate how the synchronization process is broken up into distinct transactions.
- A database error can occur at any point within the synchronization process. Database errors are handled using the `handle_error` or `handle_odbc_error` scripts.

There should be no implicit or explicit commit or rollback in your SQL synchronization scripts or the procedures or triggers that are called from your SQL synchronization scripts. COMMIT or ROLLBACK statements within SQL scripts alter the transactional nature of the synchronization steps. If you use them, MobiLink cannot guarantee the integrity of your data in the event of a failure.

## MobiLink complete event model

```
-----  
MobiLink complete event model.
```

```
Legend:
```

```
- // This is a comment.
```

```
- <name>
```

```
    The pseudo code for <name> is listed separately  
    in a later section, under a banner:
```

```
        -----  
        name  
        -----
```

```
- VariableName <- value
```

```
    Assign the given value to the given variable name.  
    Variable names are in mixed case.
```

```
- event_name
```

```
    If you have defined a script for the given event name,
```

```

    it is invoked.
-----

CONNECT to consolidated database
begin_connection_autocommit
begin_connection
COMMIT
for each synchronization request with
    the same script version {
    <synchronize>
}
end_connection
COMMIT
DISCONNECT from consolidated database

-----

synchronize
-----

<authenticate>
<begin_synchronization>
<upload>
<prepare_for_download>
<download>
<end_synchronization>

-----

authenticate
-----

Status <- 1000
UseDefaultAuthentication <- TRUE
if( authenticate_user script is defined ) {
    UseDefaultAuthentication <- FALSE
    TempStatus <- authenticate_user
    if( TempStatus > Status ) {
        Status <- TempStatus
    }
}
if( authenticate_user_hashed script is defined ) {
    UseDefaultAuthentication <- FALSE
    TempStatus <- authenticate_user_hashed
    if( TempStatus > Status ) {
        Status <- TempStatus
    }
}
if( authenticate_parameters script is defined )
{
    TempStatus <- authenticate_parameters
    if( TempStatus > Status ) {
        Status <- TempStatus
    }
}
if( UseDefaultAuthentication ) {
    if( the user exists in the ml_user table ) {
        if( ml_user.hash_password column is not NULL ) {
            if( password matches ml_user.hash_password ) {
                Status <- 1000
            } else {
                Status <- 4000
            }
        }
    } else {
        Status <- 1000
    }
} else if( -zu+ was on the command line ) {

```

```
        Status <- 1000
      } else {
        Status <- 4000
      }
    }
  if( Status >= 3000 ) {
    // Abort the synchronization.
  } else {
    // UserName defaults to MobiLink user name
    // sent from the remote.
    if( modify_user script is defined ) {
      UserName <- modify_user
      // The new value of UserName is later passed to
      // all scripts that expect the MobiLink user name.
    }
  }
}
COMMIT

-----
begin_synchronization
-----

begin_synchronization // Connection event.
for each table being synchronized {
  begin_synchronization // Call the table level script.
}
for each publication being synchronized {
  begin_publication
}
COMMIT

-----
end_synchronization
-----

for each publication being synchronized {
  if( begin_publication script was called ) {
    end_publication
  }
}
for each table being synchronized {
  if( begin_synchronization table script was called ) {
    end_synchronization // Table event.
  }
}
if( begin_synchronization connection script was called ) {
  end_synchronization // Connection event.
}
for each table being synchronized {
  synchronization_statistics // Table event.
}
synchronization_statistics // Connection event.
for each table being synchronized {
  time_statistics // Table event.
}
time_statistics // Connection event.

COMMIT
```

For the details of upload processing, see [“Events during upload” on page 349](#).

For the details of download processing, see [“Events during download” on page 351](#).



## Events during upload

The following pseudocode illustrates how upload events and upload scripts are invoked.

These events take place at the upload location in the complete event model. See [“Overview of MobiLink events” on page 343](#).

### Overview of the upload

```

-----
upload
-----

begin_upload // Connection event
for each table being synchronized {
  begin_upload // Table event
}
  handle_UploadData
  for each table being synchronized {
    begin_upload_rows
    for each uploaded INSERT or UPDATE for this table {
      if( INSERT ) {
        <upload_inserted_row>
      }
      if( UPDATE ) {
        <upload_updated_row>
      }
    }
    end_upload_rows
  }
  for each table being synchronized IN REVERSE ORDER {
    begin_upload_deletes
    for each uploaded DELETE for this table {
      <upload_deleted_row>
    }
    end_upload_deletes
  }
}
For each table being synchronized {
  if( begin_upload table script is called ) {
    end_upload // Table event
  }
}
if( begin_upload connection script was called ) {
  end_upload // Connection event

  for each table being synchronized {
    upload_statistics // Table event.
  }
  upload_statistics // Connection event.

  COMMIT

```

### Upload inserts

```

-----
<upload_inserted_row>
-----
// NOTES:
// - Only table scripts for the current table are involved.

ConflictsAreExpected <- (

```

```
        upload_new_row_insert script is defined
    or upload_old_row_insert script is defined
    or resolve_conflict script is defined )
if( upload_insert script is defined ) {
    upload_insert
} else if( ConflictsAreExpected
    and upload_update script is not defined
    and upload_insert script is not defined
    and upload_delete script is not defined ) {
    // Forced conflict.
    upload_new_row_insert
    resolve_conflict
} else {
    // Ignore the insert.
}
```

### Upload updates

```
-----
upload_updated_row
-----
// NOTES:
// - Only table scripts for the current table are involved.
// - Both the old (original) and new rows are uploaded for
//   each update.

ConflictsAreExpected <- (
    upload_new_row_insert script is defined
    or upload_old_row_insert script is defined
    or resolve_conflict script is defined )
Conflicted <- FALSE
if( upload_update script is defined ) {
    if( ConflictsAreExpected
        and upload_fetch script is defined ) {
        FETCH using upload_fetch INTO current_row
        if( current_row <> old_row ) {
            Conflicted <- TRUE
        }
    }
    if( not Conflicted ) {
        upload_update
    }
} else if( upload_update script is not defined
    and upload_insert script is not defined
    and upload_delete script is not defined ) {
    // Forced conflict.
    Conflicted <- TRUE
}
if( ConflictsAreExpected and Conflicted ) {
    upload_old_row_insert
    upload_new_row_insert
    resolve_conflict
}
```

### Upload deletes

```
-----
upload_deleted_row
-----
// NOTES:
// - Only table scripts for the current table are involved.

ConflictsAreExpected <- (
    upload_new_row_insert script is defined
```

```

    or upload_old_row_insert script is defined
    or resolve_conflict script is defined )
if( upload_delete is defined ) {
    upload_delete
} else if( ConflictsAreExpected
and upload_update script is not defined
and upload_insert script is not defined
and upload_delete script is not defined ) {
    // Forced conflict.
    upload_old_row_insert
    resolve_conflict
} else {
    // Ignore this delete.
}

```

## Events during download

The following pseudocode provides an overview of the sequence in which download events, and the script of the same name, are invoked.

These events take place at the download location in the complete event model provided in [“Overview of MobiLink events” on page 343](#).

```

-----
prepare_for_download
-----

modify_last_download_timestamp
fetch the next download timestamp from consolidated
prepare_for_download
if( modify_last_download_timestamp script is defined
    or prepare_for_download script is defined ) {
    COMMIT
}

-----
download
-----

begin_download // Connection event.
for each table being synchronized {
    begin_download // Table event.
}
handle_DownloadData
for each table being synchronized {
    begin_download_deletes
    for each row in download_delete_cursor {
        if( all primary key columns are NULL ) {
            send TRUNCATE to remote
        } else {
            send DELETE to remote
        }
    }
    end_download_deletes
    begin_download_rows
    for each row in download_cursor {
        send INSERT ON EXISTING UPDATE to remote
    }
    end_download_rows
}

```

```
    modify_next_last_download_timestamp
  for each table being synchronized {
    if( begin_download table script is called ) {
      end_download // Table event
    }
  }
if( begin_download connect script is called ) {
  end_download // Connection event
}
  for each table being synchronized {
    download_statistics // Table event.
  }
  download_statistics // Connection event.

COMMIT
```

### Notes

- If an acknowledgement is expected, and if no confirmation of the downloads is received from the client, the entire download transaction is rolled back in the consolidated database.

For SQL Anywhere remotes, see [“SendDownloadACK \(sa\) extended option”](#) [*MobiLink - Client Administration*]. For UltraLite remotes, see [“Send Download Acknowledgement synchronization parameter”](#) [*UltraLite - Database Management and Reference*].

- The download stream does not distinguish between inserts and updates. The script associated with the download\_cursor event is a SELECT statement that defines the rows to be downloaded. The client detects whether the row exists and then it performs the appropriate insert or update operation.
- At the end of the download processing, the client automatically deletes rows that violate referential integrity.

See [“Referential integrity and synchronization”](#) [*MobiLink - Getting Started*].

## authenticate\_file\_transfer connection event

Implements custom authentication for file transfers using the mlfiletransfer utility or the MLFileTransfer method.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.file_authentication_code	INTEGER. Required. This is an INOUT parameter. It indicates the overall success of the authentication.  If this value is 1000-1999, file transfer is allowed. If this value is 2000-2999, file transfer is not allowed.	1
s.filename	VARCHAR(128). This optional parameter is the name of the file that is being transferred that is to be authenticated. Do not include a path. The file must be located in the root transfer directory that you specified with the mlsrv11 -ftr option, or in one of the subdirectories that are automatically created.	2
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	3

### Remarks

The MobiLink server executes this event before allowing any file transfer using the mlfiletransfer utility or MLFileTransfer method. It is executed after the user has authenticated using regular authentication. If this script is not defined, the file transfer is allowed.

The MLFileTransfer method can only be used by UltraLite clients.

**See also**

- [“Adding and deleting scripts” on page 327](#)
- [“-ftr option” on page 71](#)
- [“MobiLink file transfer utility \(mlfiletransfer\)” \[\*MobiLink - Client Administration\*\]](#)
- [UltraLite: “Using MobiLink file transfers” \[\*UltraLite - Database Management and Reference\*\]](#)
- [“Using remote IDs and MobiLink user names in scripts” \[\*MobiLink - Client Administration\*\]](#)

## authenticate\_parameters connection event

Receives values from the remote that can be used to authenticate beyond a user ID and password. The values can also be used to arbitrarily customize each synchronization.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.authentication_status	INTEGER. This is an INOUT parameter.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
a.N (one or more)	VARCHAR(128). For example, named parameters could be a.1 a.2.	3...

### Parameter Description

- authentication\_status** The authentication\_status parameter is required. It indicates the overall success of the authentication, and can be set to one of the following values:

Returned Value	authentication_status	Description
V <= 1999	1000	Authentication succeeded.
1999 < V <= 2999	2000	Authentication succeeded, but password expiring soon.
2999 < V <= 3999	3000	Authentication failed as password has expired.
3999 < V <= 4999	4000	Authentication failed.

Returned Value	authentication_status	Description
4999 < V <= 5999	5000	Authentication failed as user is already synchronizing.
5999 < V	4000	If the returned value is greater than 5999, MobiLink interprets it as a returned value of 4000 (authentication failed).

- **username** This parameter is the MobiLink user name. VARCHAR(128).
- **remote\_ID** The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.  
See “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*].
- **remote\_parameters** The number of remote parameters must match the number expected or an error results. An error also occurs if parameters are sent from the client and there is no script for this event.

### Remarks

You can send strings (or parameters in the form of strings) from both SQL Anywhere and UltraLite clients. This allows you to have authentication beyond a user ID and password. It also means that you can customize your synchronization based on the value of parameters, and do this in a pre-synchronization phase, during authentication.

The MobiLink server executes this event upon starting each synchronization. It is executed in the same transaction as the `authenticate_user` event.

You can use this event to replace the built-in MobiLink authentication mechanism with a custom mechanism. You may want to call into the authentication mechanism of your DBMS, or you may want to implement features not present in the MobiLink built-in mechanism.

If the `authenticate_user` or `authenticate_user_hashed` scripts are invoked and return an error, this event is not called.

SQL scripts for the `authenticate_parameters` event must be implemented as stored procedures.



**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “Authentication parameters” on page 323
- “MobiLink users” [*MobiLink - Client Administration*]
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Custom user authentication” [*MobiLink - Client Administration*]
- “authenticate\_user connection event” on page 358
- “authenticate\_user\_hashed connection event” on page 363
- “begin\_synchronization connection event” on page 383
- dbmlsync: “-ap option” [*MobiLink - Client Administration*]
- UltraLite: “Authentication Parameters synchronization parameter” [*UltraLite - Database Management and Reference*] and “Number of Authentication Parameters parameter” [*UltraLite - Database Management and Reference*]

**Examples**

For UltraLite remote databases, pass the parameters using the num\_auth\_parms and auth\_parms fields in the ul\_synch\_info struct. num\_auth\_parms is a count of the number of parameters, from 0 to 255. auth\_parms is a pointer to an array of strings. To prevent the strings from being viewed as plain text, the strings are sent in the same way as passwords. If num\_auth\_parms is 0, set auth\_parms to null. The following is an example of passing parameters in UltraLite:

```
ul_char * Params[ 3 ] = { UL_TEXT( "param1" ),
    UL_TEXT( "param2" ), UL_TEXT( "param3" ) };

...
info.num_auth_parms = 3;
info.auth_parms = Params;
```

For SQL Anywhere remote databases, you pass parameters using the dbmlsync -ap option, in a comma-separated list. For example, the following command line passes three parameters:

```
dbmlsync -ap "param1,param2,param3"
```

In this example, the authenticate\_parameters script could be:

```
CALL my_auth_parm (
    {ml s.authentication_status},
    {ml s.remote_id},
    {ml s.username},
    {ml a.1},
    {ml a.2},
    {ml a.3}
)
```

## authenticate\_user connection event

Implements custom user authentication.

### Parameters

In the following table, the description indicates the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.authentication_status	INTEGER. This is an INOUT parameter.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.password	VARCHAR(128). The password for authentication purposes. If the user does not supply a password, this value is null.	3
s.new_password	VARCHAR(128). The new password, if this is being used to reset the user password. If the user does not change their password, this value is null.	4

### Default action

Use MobiLink built-in user authentication mechanism.

### Remarks

The MobiLink server executes this event upon starting each synchronization. It is executed in a transaction before the `begin_synchronization` transaction.

You can use this event to replace the built-in MobiLink authentication mechanism with a custom mechanism. You may want to call into the authentication mechanism of your DBMS, or you may want to implement features not present in the MobiLink built-in mechanism, such as password expiry or a minimum password length.

The parameters used in an `authenticate_user` event are as follows:

- **authentication\_status** The authentication\_status parameter is required. It indicates the overall success of the authentication, and can be set to one of the following values:

Returned Value	authentication_status	Description
V <= 1999	1000	Authentication succeeded.
1999 < V <= 2999	2000	Authentication succeeded: password expiring soon.
2999 < V <= 3999	3000	Authentication failed: password expired.
3999 < V <= 4999	4000	Authentication failed.
4999 < V <= 5999	5000	Authentication failed as user is already synchronizing.
5999 < V	4000	If the returned value is greater than 5999, MobiLink interprets it as a returned value of 4000.

- **username** This optional parameter is the MobiLink user name.  
See [“Using remote IDs and MobiLink user names in scripts”](#) [*MobiLink - Client Administration*].
- **remote\_id** The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.
- **password** This optional parameter indicates the password for authentication purposes. If the user does not supply a password, this is null.
- **new\_password** This optional parameter indicates a new password. If the user does not change their password, this is null.

SQL scripts for the authenticate\_user event must be implemented as stored procedures.

When the two authentication scripts are both defined, and both scripts return different authentication\_status codes, the higher value is used.

The authenticate\_user script is executed in a transaction along with all authentication scripts. This transaction always commits.

There are predefined scripts that you can use for the authenticate\_user event to simplify authentication using LDAP, IMAP and POP3 servers.

See [“Authenticating to external servers”](#) [*MobiLink - Client Administration*].

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “MobiLink users” [*MobiLink - Client Administration*]
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Custom user authentication” [*MobiLink - Client Administration*]
- “Authenticating to external servers” [*MobiLink - Client Administration*]
- “authenticate\_user\_hashed connection event” on page 363
- “authenticate\_parameters connection event” on page 355
- “modify\_user connection event” on page 469
- “begin\_synchronization connection event” on page 383

### SQL example

A typical `authenticate_user` script is a call to a stored procedure. The order of the parameters in the call must match the order above. The following example uses `ml_add_connection_script` to assign the event to a stored procedure called `my_auth`.

```
CALL ml_add_connection_script(
    'ver1', 'authenticate_user', 'call my_auth ( {ml s.username} )'
)
```

The following SQL Anywhere stored procedure uses only the user name to authenticate—it has no password check. The procedure ensures only that the supplied user name is one of the employee IDs listed in the `ULEmployee` table.

```
CREATE PROCEDURE my_auth( in @user_name varchar(128) )
BEGIN
    IF EXISTS
        ( SELECT * FROM ulemployee
          WHERE emp_id = @user_name )
    THEN
        MESSAGE 'OK' type info to client;
        RETURN 1000;
    ELSE
        MESSAGE 'Not OK' type info to client;
        RETURN 4000;
    END IF
END
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `authenticateUser` as the script for the `authenticate_user` event when synchronizing the script version `ver1`. This syntax is for SQL Anywhere consolidated databases.

```
CALL ml_add_java_connection_script(
    'ver1', 'authenticate_user',
    'ExamplePackage.ExampleClass.authenticateUser'
)
```

The following is the sample Java method `authenticateUser`. It calls Java methods that check and, if needed, change the user's password.

```
public String authenticateUser(
    ianywhere.ml.script.InOutInteger authStatus,
```

```

String user,
String pwd,
String newPwd )
throws java.sql.SQLException {
// A real authenticate_user handler would
// handle more authentication code states.
_curUser = user;
if( checkPwd( user, pwd ) ) {
// Authentication successful.
if( newPwd != null ) {
// Password is being changed.
if( changePwd( user, pwd, newPwd ) ) {
// Authentication OK and password change OK.
// Use custom code.
authStatus.setValue( 1001 );
} else {
// Authentication OK but password
// change failed. Use custom code.
java.lang.System.err.println( "user: "
+ user + " pwd change failed!" );
authStatus.setValue( 1002 );
}
} else {
authStatus.setValue( 1000 );
}
} else {
// Authentication failed.
authStatus.setValue( 4000 );
}
return ( null );
}
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called AuthUser as the script for the authenticate\_user connection event when synchronizing the script version ver1. This syntax is for SQL Anywhere consolidated databases.

```

CALL ml_add_dnet_connection_script(
    'ver1', 'authenticate_user',
    'TestScripts.Test.AuthUser'
)

```

The following is the sample .NET method AuthUser. It calls .NET methods that check and, if needed, change the user's password.

```

public string AuthUser(
    ref int authStatus,
    string user,
    string pwd,
    string newPwd ) {
// A real authenticate_user handler would
// handle more authentication code states.
_curUser = user;
if( CheckPwd( user, pwd ) ) {
// Authentication successful.
if( newPwd != null ) {
// Password is being changed.
if( ChangePwd( user, pwd, newPwd ) ) {
// Authentication OK and password change OK.
// Use custom code.
authStatus = 1001;
}
}
}
}

```

```
    } else {
        // Authentication OK but password
        // change failed. Use custom code.
        System.Console.WriteLine( "user: "
            + user + " pwd change failed!" );
        authStatus = 1002;
    }
} else {
    authStatus = 1000 ;
}
} else {
    // Authentication failed.
    authStatus = 4000;
}
return ( null );
}
```

For a more detailed example of an `authenticate_user` script written in C# in .NET, see [“.NET synchronization example” on page 604](#).

## authenticate\_user\_hashed connection event

Implements a custom user authentication mechanism.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.authentication_status	INTEGER. This is an INOUT parameter.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.hash_password	BINARY(20). If the user does not supply a password, this value is null.	3
s.hash_new_password	BINARY(20). If this event is not being used to change the user's password, this value is null.	4

### Default action

Use MobiLink built-in user authentication mechanism.

### Remarks

This event is identical to `authenticate_user` except for the passwords, which are in the same hashed form as those stored in the `ml_user.hash_password` column. Passing the passwords in hashed form provides increased security.

A one-way hash is used. A one-way hash takes a password and converts it to a byte sequence that is (essentially) unique to each possible password. The one-way hash lets password authentication take place without having to store the actual password in the consolidated database.

This script can be called multiple times during an authentication sequence for a user.

When `authenticate_user` and `authenticate_user_hashed` are both defined, and both scripts return different `authentication_status` codes, the higher value is used.

**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “MobiLink users” [*MobiLink - Client Administration*]
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Custom user authentication” [*MobiLink - Client Administration*]
- “authenticate\_user connection event” on page 358
- “authenticate\_parameters connection event” on page 355

**SQL example**

A typical authenticate\_user\_hashed script is a call to a stored procedure. The order of the parameters in the call must match the order above. The following example calls ml\_add\_connection\_script to assign the event to a stored procedure called my\_auth.

```
CALL ml_add_connection_script(  
  'ver1', 'authenticate_user_hashed',  
  'call my_auth (  
    {ml s.authentication_status},  
    {ml s.username},  
    {ml s.hashed_password})'  
)
```

The following SQL Anywhere stored procedure uses both the user name and password to authenticate. The procedure ensures only that the supplied user name is one of the employee IDs listed in the ULEmployee table. The procedure assumes that the Employee table has a binary(20) column called hashed\_pwd.

```
CREATE PROCEDURE my_auth(  
  inout @authentication_status integer,  
  in @user_name varchar(128),  
  in @hpwd binary(20) )  
BEGIN  
  IF EXISTS  
  ( SELECT * FROM ulemLOYEE  
    WHERE emp_id = @user_name  
      and hashed_pwd = @hpwd )  
  THEN  
    message 'OK' type info to client;  
    RETURN 1000;  
  ELSE  
    message 'Not OK' type info to client;  
    RETURN 4000;  
  END IF  
END
```

**Java example**

The following call to a MobiLink system procedure registers a Java method called authUserHashed as the script for the authenticate\_user\_hashed event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(  
  'ver1', 'authenticate_user_hashed',  
  'ExamplePackage.ExampleClass.authUserHashed')
```

The following is the sample Java method authUserHashed. It calls Java methods that check and, if needed, change the user's password.



```

public String authUserHashed(
    anywhere.ml.script.InOutInteger authStatus,
    String user,
    byte pwd[],
    byte newPwd[] )
throws java.sql.SQLException {
    // A real authenticate_user_hashed handler
    // would handle more auth code states.
    _curUser = user;
    if( checkPwdHashed( user, pwd ) ) {
        // Authorization successful.
        if( newPwd != null ) {
            // Password is being changed.
            if( changePwdHashed( user, pwd, newPwd ) ) {
                // Authorization OK and password change OK.
                // Use custom code.
                authStatus.setValue( 1001 );
            } else {
                // Auth OK but password change failed.
                // Use custom code
                java.lang.System.err.println( "user: " + user
                    + " pwd change failed!" );
                authStatus.setValue( 1002 );
            }
        } else {
            authStatus.setValue( 1000 );
        }
    } else {
        // Authorization failed.
        authStatus.setValue( 4000 );
    }
    return ( null );
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called AuthUserHashed as the script for the authenticate\_user\_hashed connection event when synchronizing the script version ver1. This syntax is for SQL Anywhere consolidated databases.

```

CALL ml_add_dnet_connection_script(
    'ver1',
    'authenticate_user_hashed',
    'TestScripts.Test.AuthUserHashed'
)

```

The following is the sample .NET method AuthUserHashed.

```

public string AuthUserHashed(
    ref int authStatus,
    string user,
    byte[] pwd,
    byte[] newPwd ) {
    // A real authenticate_user_hashed handler
    // would handle more auth code states.
    _curUser = user;
    if( CheckPwdHashed( user, pwd ) ) {
        // Authorization successful.
        if( newPwd != null ) {
            // Password is being changed.
            if( ChangePwdHashed( user, pwd, newPwd ) ) {
                // Authorization OK and password change OK.
                // Use custom code.
            }
        }
    }
}

```

```
    authStatus = 1001;
  } else {
    // Auth OK but password change failed.
    // Use custom code
    System.Console.WriteLine( "user: " + user
      + " pwd change failed!" );
    authStatus = 1002;
  }
  } else {
    authStatus = 1000;
  }
  } else {
    // Authorization failed.
    authStatus = 4000;
  }
  return ( null );
}
```

## begin\_connection connection event

Invoked at the time the MobiLink server connects to the consolidated database server.

### Parameters

None.

### Default action

None.

### Remarks

The MobiLink synchronization opens connections on demand as synchronization requests come in. When an application forms or reforms a connection with the MobiLink server, the MobiLink server temporarily allocates one connection with the database server for the duration of that synchronization. This event may not be called if the MobiLink server is using a connection from the pool.

#### Note

This script is not generally used in Java or .NET, because instead of database variables you would use member variables in this class instance, and prepare the members in the constructor.

### See also

- [“Adding and deleting scripts” on page 327](#)
- [“end\\_connection connection event” on page 409](#)
- [“-cn option” on page 56](#)
- [“-w option” on page 105](#)

### SQL example

The following SQL script works with a SQL Anywhere consolidated database. Two variables are created, one for the last\_download timestamp, and one for employee ID.

```
CALL ml_add_connection_script(  
    'custdb',  
    'begin_connection',  
    'create variable @LastDownload timestamp;  
    create variable @EmployeeID integer;')
```

## **begin\_connection\_autocommit connection event**

Turns on autocommit.

### **Parameters**

None.

### **Default action**

Autocommit is off.

### **Remarks**

When the MobiLink server connects to the consolidated database, it turns off autocommit so that it can roll back the upload and download if an error occurs.

However, if you are using an Adaptive Server Enterprise consolidated database, you cannot perform DDL functions such as creating temporary tables unless autocommit is on. If you are using an Adaptive Server Enterprise consolidated database, run your DDL commands in the `begin_connection_autocommit` event. When the event is finished, autocommit is turned off.

`begin_connection_autocommit` scripts must be written so that they are repeatable. This is because if an error or deadlock occurs, the MobiLink server needs to be able to retry the script (since it can't roll it back).

This event only executes if a script has been defined for the event.

### **See also**

- [“Adding and deleting scripts” on page 327](#)

## begin\_download connection event

Processes any statements just before the MobiLink server commences preparing the download.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_download	TIMESTAMP. The last download time of any synchronized table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

The MobiLink server executes this event as the first step in the processing of downloaded information. Download information is processed in a single transaction. The execution of this event is the first action in this transaction.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_download connection event” on page 411](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

The following example calls `ml_add_connection_script` to assign the event to a stored procedure called `SetDownloadParameters`.

```
CALL ml_add_connection_script (
    'Lab',
```

```
'begin_download',  
'CALL SetDownloadParameters( {ml s.last_table_download}, {ml  
s.username} )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginDownloadConnection` as the script for the `begin_download` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(  
  'example_ver',  
  'begin_download',  
  'ExamplePackage.ExampleClass.beginDownloadConnection' )
```

The following is the sample Java method `beginDownloadConnection`. It calls a Java method (`prepDeleteTables`) that prepares the delete tables using a JDBC synchronization that was set earlier.

```
public String beginDownloadConnection(  
  Timestamp ts,  
  String user )  
  throws java.sql.SQLException {  
  prepDeleteTables ( _syncConn, ts, user );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginDownload` as the script for the `begin_download` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'begin_download',  
  'TestScripts.Test.BeginDownload'  
)
```

The following is the sample .NET method `BeginDownload`. It calls a .NET method (`prepDeleteTables`) that prepares the delete tables using a JDBC synchronization that was set earlier.

```
public string BeginDownload(  
  DateTime timestamp,  
  string user ) {  
  prepDeleteTables ( _syncConn, ts, user );  
  return ( null );  
}
```

## begin\_download table event

Processes statements related to a specific table just before preparing the download inserts, updates, and deletions.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.table	VARCHAR (128). The table name.	3

### Default action

None.

### Remarks

The MobiLink server executes this event as the first step in preparing download information for a specific table. The download information is prepared in its own transaction. The execution of this event is the first table-specific action in the transaction.

You can have one begin\_download script for each table in the remote database. The script is only invoked when that table is synchronized.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_download table event” on page 414](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

The following call to the MobiLink system procedure `ml_add_table_script` calls the `BeginTableDownload` procedure. This syntax is for a SQL Anywhere 11 consolidated database.

```
CALL ml_add_table_script(  
    'version1',  
    'Leads',  
    'begin_download',  
    'CALL BeginTableDownload(  
        {ml s.last_table_download},  
        {ml s.username},  
        {ml s.table} )' );
```

The following SQL statements create the `BeginTableDownload` procedure.

```
CREATE PROCEDURE BeginTableDownload(  
    LastDownload timestamp,  
    MLUser varchar(128),  
    TableName varchar(128) )  
BEGIN  
    EXECUTE IMMEDIATE 'update ' || TableName ||  
    ' set last_download_check = CURRENT_TIMESTAMP  
    WHERE Owner = ' || MLUser;  
END
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginDownloadTable` as the script for the `begin_download` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1',  
    'begin_download',  
    'ExamplePackage.ExampleClass.beginDownloadTable' )
```

The following is the sample Java method `beginDownloadTable`. It saves the name of the current table for use in a later method call.

```
public String beginDownloadTable(  
    Timestamp ts,  
    String user,  
    String table ) {  
    _curTable = table;  
    return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginTableDownload` as the script for the `begin_download` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script(  
    'ver1', 'table1', 'begin_download',  
    'TestScripts.Test.BeginTableDownload'  
)
```

The following is the sample .NET method `BeginTableDownload`. It saves the name of the current table for use in a later method call.



```
public string BeginTableDownload(
    DateTime timestamp,
    string user,
    string table ) {
    _curTable = table;
    return ( null );
}
```

## begin\_download\_deletes table event

Processes statements related to a specific table just before fetching a list of rows to be deleted from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR (128). The MobiLink user name.	2
s.table	VARCHAR (128). The table name.	3

### Default action

None.

### Remarks

This event is executed immediately before fetching a list of rows to be deleted from the named table in the remote database.

You can have one begin\_download\_deletes script for each table in the remote database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_download\\_rows table event” on page 377](#)
- [“end\\_download\\_rows table event” on page 420](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

## SQL example

To minimize the amount of data on remotes, you can use this event to flag data that is deleted when the `download_delete_cursor` is executed. The following example flags for deletion sales leads from the remote device that are over 10 weeks old. The example can be used on a SQL Anywhere 11 database.

The following call to a MobiLink system procedure assigns the `BeginDownloadDeletes` stored procedure to the `begin_download_deletes` event when synchronizing the script version `ver1`.

```
CALL ml_add_table_script (
  'ver1',
  'Leads',
  'begin_download_deletes',
  'CALL BeginDownloadDeletes (
    {ml s.last_table_download},
    {ml s.username},
    {ml s.table})' );
```

The following SQL statement creates the `BeginDownloadDeletes` stored procedure.

```
CREATE PROCEDURE BeginDownloadDeletes(
  LastDownload timestamp,
  MLUser varchar(128),
  TableName varchar(128) )
BEGIN
  execute immediate 'update ' || TableName ||
  ' set delete_flag = 1 where
  days(creation_time, CURRENT DATE) > 70 and Owner = '
  || MLUser;
END;
```

## Java example

The following call to a MobiLink system procedure registers a Java method called `beginDownloadDeletes` as the script for the `begin_download_deletes` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'begin_download_deletes',
  'ExamplePackage.ExampleClass.beginDownloadDeletes' )
```

The sample Java method `beginDownloadDeletes` saves the name of the current table for use in a later method call.

```
public String beginDownloadDeletes (
  Timestamp ts,
  String user,
  String table ) {
  _curTable = table;
  return ( null );
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginDownloadDeletes` as the script for the `begin_download_deletes` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script (
  'ver1', 'table1', 'begin_download_deletes',
```

```
    'TestScripts.Test.BeginDownloadDeletes'  
  )
```

The sample .NET method `BeginDownloadDeletes` saves the name of the current table for use in a later method call.

```
public string BeginDownloadDeletes(  
    DateTime timestamp,  
    string user,  
    string table ) {  
    _curTable = table;  
    return ( null );  
}
```

## begin\_download\_rows table event

Processes statements related to a specific table just before fetching a list of rows to be inserted or updated in the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR (128). The MobiLink user name.	2
s.table	VARCHAR (128). The table name.	3

### Default action

None.

### Remarks

This event is executed immediately before fetching the stream of rows to be inserted or updated in the named table in the remote database.

You can have one begin\_download\_rows script for each table in the remote database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_download\\_deletes table event” on page 374](#)
- [“end\\_download\\_deletes table event” on page 417](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

## SQL example

You can use the `begin_download_rows` table event to flag rows that you no longer want to download for this table. The following example archives sales leads that are over seven days old.

The following call to a MobiLink system procedure registers the `BeginDownloadRows` stored procedure for the `begin_download_rows` event.

```
CALL ml_add_table_script(  
  'version1',  
  'Leads',  
  'begin_download_rows',  
  'CALL BeginDownloadRows (  
    {ml s.last_table_download},  
    {ml s.username},  
    {ml s.table})' ); )
```

The following SQL statement creates the `BeginDownloadRows` stored procedure.

```
CREATE PROCEDURE BeginDownloadRows (  
  LastDownload timestamp, MLUser varchar(128),  
  TableName varchar(128) )  
BEGIN  
  execute immediate 'update ' || TableName ||  
  ' set download_flag = 0 where  
  days(creation_time, CURRENT DATE) > 7 and Owner = '  
  || MLUser;  
END;
```

## Java example

The following call to a MobiLink system procedure registers a Java method called `beginDownloadRows` as the script for the `begin_download_rows` table event when synchronizing the script version `ver1`. This syntax is for SQL Anywhere consolidated databases.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'begin_download_rows',  
  'ExamplePackage.ExampleClass.beginDownloadRows' )
```

The following is the sample Java method `beginDownloadRows`. It generates an `UPDATE` statement using the table and user for MobiLink to execute.

```
public String beginDownloadRows(  
  Timestamp ts,  
  String user,  
  String table ) {  
  return( "update " + table + " set download_flag = 0 "  
  + " where days(creation_time, CURRENT DATE) > 7 " +  
  " and Owner = '" + user + "' " );  
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginDownloadRows` as the script for the `begin_download_rows` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script(  
  'ver1', 'table1', 'begin_download_rows',
```

```
'TestScripts.Test.BeginDownloadRows'  
)
```

The following is the sample .NET method BeginDownloadRows. It generates an UPDATE statement using the table and user for MobiLink to execute.

```
public string BeginDownloadRows(  
    DateTime timestamp,  
    string user,  
    string table ) {  
    return( "update " + table + " set download_flag = 0 "  
        + " where days(creation_time, CURRENT DATE) > 7 " +  
        " and Owner = '" + user + "'" );  
}
```

## begin\_publication connection event

Provides useful information about the publication(s) being synchronized. This script may also be used to manage generation numbers for file-based downloads.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.generation_number	INTEGER. This is an INOUT parameter. If your deployment does not use file-based downloads, this parameter can be ignored. The default is 1.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.publication_name	VARCHAR(128). The name of the publication.	3
s.last_publication_upload	TIMESTAMP. The time of the last successful upload of this publication.	4
s.last_publication_download	TIMESTAMP. The last download time for the publication.	5
s.subscription_id	VARCHAR(128). The subscription ID.	6

### Default action

The default generation number is 1. If no script is defined for this event, the generation number sent to the remote is always 1.

### Remarks

This event lets you design synchronization logic based on the publications currently being synchronized. This event is invoked in the same transaction as the begin\_synchronization event, and is invoked after the begin\_synchronization event. It is invoked once per publication being synchronized.

One potential use for this event is to affect what is downloaded based on the publication used. For example, consider a table that is part of both a priority publication (PriorityPub) and a publication for all tables (AllTablesPub). A script for the begin\_publication event could store the publication names in a Java class or a SQL variable or package. Download scripts could then behave differently based on whether the publication being synchronized is PriorityPub or AllTablesPub.



If an UltraLite remote is synchronizing with UL\_SYNC\_ALL, this event is invoked once with the name 'unknown'.

### Generation number

The `generation_number` parameter is specifically for file-based downloads. The output value of the generation number is passed from the `begin_synchronization` script to the `end_synchronization` script. The meaning of the `generation_number` depends on whether the current synchronization is being used to create a download file, or whether the current synchronization has an upload.

In file-based downloads, generation numbers are used to force an upload before the download. The number is stored in the download file. During a synchronization that has an upload, one generation number is output for every subscription to a publication. They are sent to the remote database in the upload acknowledgement, and stored in `SYSSYNC.generation_number`.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_publication connection event” on page 423](#)
- [“MobiLink file-based download” on page 293](#)
- [“MobiLink generation numbers” on page 299](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

You may want to record the information for each publication being synchronized. The following example calls `ml_add_connection_script` to assign the event to a stored procedure called `RecordPubSync`.

```
CALL ml_add_connection_script(
  'version1',
  'begin_publication',
  '{CALL RecordPubSync(
    {ml s.generation_number},
    {ml s.username},
    {ml s.publication_name},
    {ml s.last_publication_upload},
    {ml s.last_publication_download},
    {ml s.subscription_id} )}' );
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginPublication` as the script for the `begin_publication` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(
  'ver1',
  'begin_publication',
  'ExamplePackage.ExampleClass.beginPublication' )
```

The following is the sample Java method `beginPublication`. It saves the name of each publication for later use.

```
public String beginPublication(
  anywhere.ml.script.InOutInteger generation_number,
  String user,
```

```
String pub_name,  
Timestamp last_publication_upload,  
Timestamp last_download ) {  
  _publicationNames[ _numPublications++ ] = pub_name;  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called BeginPub as the script for the begin\_publication connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'begin_publication',  
  'TestScripts.Test.BeginPub'  
)
```

The following is the sample .NET method BeginPub. It saves the name of each publication for later use.

```
public string BeginPub(  
  ref int generation_number,  
  string user,  
  string pub_name,  
  DateTime last_publication_upload,  
  DateTime last_download ) {  
  _publicationNames[ _numPublications++ ] = pub_name;  
  return ( null );  
}
```

## begin\_synchronization connection event

Processes any statements at the time an application connects to the MobiLink server in preparation for the synchronization process.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1

### Default action

None.

### Remarks

The MobiLink server executes this event immediately after an application preparing to synchronize has formed a connection with the MobiLink server. It is executed within a separate transaction before the upload transaction.

The begin\_synchronization script is useful for maintaining statistics. This is because the end\_synchronization script is invoked even if there is an error or conflict, so while the upload transaction is rolled back, things like statistics are maintained.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_synchronization connection event” on page 426](#)
- [“begin\\_synchronization table event” on page 385](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

You may want to store the username value in a temporary table or variable if you want to reference that value many times in subsequent scripts.

```
CALL ml_add_connection_script (
  'version1',
  'begin_synchronization',
  'set @EmployeeID = {ml s.username}' );
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginSynchronizationConnection` as the script for the `begin_synchronization` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'begin_synchronization',  
  'ExamplePackage.ExampleClass.beginSynchronizationConnection'  
)
```

The following is the sample Java method `beginSynchronizationConnection`. It saves the name of the synchronizing user for later use.

```
public String beginSynchronizationConnection(  
  String user ) {  
  _curUser = user;  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginSync` as the script for the `begin_synchronization` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script( 'ver1',  
  'begin_synchronization',  
  'TestScripts.Test.BeginSync'  
)
```

The following is the sample .NET method `BeginSync`. It saves the name of the synchronizing user for later use.

```
public string BeginSync(  
  string user ) {  
  _curUser = user;  
  return ( null );  
}
```

## begin\_synchronization table event

Processes statements related to a specific table at the time an application connects to the MobiLink server in preparation for the synchronization process.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR (128). The MobiLink user name.	1
s.table	VARCHAR (128). The table name.	2

### Default action

None.

### Remarks

The MobiLink server executes this event after an application that is preparing to synchronize has formed a connection with the MobiLink server, and after the begin\_synchronization connection-level event.

You can have one begin\_synchronization script for each table in the remote database. The event is only invoked when the table is synchronized.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_synchronization table event” on page 428](#)
- [“begin\\_synchronization connection event” on page 383](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The begin\_synchronization table event is used to set up the synchronization of a particular table. The following SQL script registers a script that creates a temporary table for storing rows during synchronization. This syntax is for a SQL Anywhere consolidated database.

```
CALL ml_add_table_script(  
  'ver1',  
  'sales_order',  
  'begin_synchronization',  
  'CREATE TABLE #sales_order (  
    id          integer NOT NULL default autoincrement,  
    cust_id     integer NOT NULL,  
    order_date  date NOT NULL,  
    fin_code_id char(2) NULL,  
    region      char(7) NULL,  
    sales_rep   integer NOT NULL,  
    PRIMARY KEY (id),  
  )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginSynchronizationTable` as the script for the `begin_synchronization` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'begin_synchronization',  
  'ExamplePackage.ExampleClass.beginSynchronizationTable' )
```

The following is the sample Java method `beginSynchronizationTable`. It adds the current table name to a list of table names contained in this instance.

```
public String beginSynchronizationTable(  
  String user,  
  String table ) {  
  _tableList.add( table );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginTableSync` as the script for the `begin_synchronization` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script (   
  'ver1',  
  'table1',  
  'begin_synchronization',  
  'TestScripts.Test.BeginTableSync' )
```

The following is the sample .NET method `BeginTableSync`. It adds the current table name to a list of table names contained in this instance.

```
public string BeginTableSync(  
  string user,  
  string table ) {  
  _tableList.Add( table );  
  return ( null );  
}
```

## begin\_upload connection event

Processes any statements just before the MobiLink server commences processing the stream of uploaded inserts, updates, and deletes.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR (128). The MobiLink user name.	1

### Default action

None.

### Remarks

The MobiLink server executes this event as the first step in the processing of uploaded information. Upload information is processed in a single transaction. The execution of this event is the first action in this transaction.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_upload connection event” on page 431](#)
- [“begin\\_upload table event” on page 389](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The begin\_upload connection event is used to perform whatever steps you need performed prior to uploading rows. The following SQL script creates a temporary table for storing old and new row values for conflict processing of the sales\_order table. This example works with a SQL Anywhere consolidated database.

```
CALL ml_add_connection_script(
  'version1',
  'begin_upload',
  'CREATE TABLE #sales_order_conflicts (
    id          integer NOT NULL default autoincrement,
    cust_id     integer NOT NULL,
    order_date  date NOT NULL,
    fin_code_id char(2) NULL,
```

```
region      char(7) NULL,  
sales_rep   integer NOT NULL,  
new_value   char(1) NOT NULL,  
PRIMARY KEY (id) )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `beginUploadConnection` as the script for the `begin_upload` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'begin_upload',  
  'ExamplePackage.ExampleClass.beginUploadConnection ' )
```

The following is the sample Java method `beginUploadConnection`. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String beginUploadConnection( String user ) {  
  java.lang.System.out.println(  
    "Starting upload for user: " + user );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `BeginUpload` as the script for the `begin_upload` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'begin_upload',  
  'TestScripts.Test.BeginUpload'  
)
```

The following C# example saves the current user name for use in a later event.

```
public string BeginUpload( string curUser ) {  
  user = curUser;  
  return ( null );  
}
```



## begin\_upload table event

Processes statements related to a specific table just before the MobiLink server commences processing the stream of uploaded inserts, updates, and deletes.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

The MobiLink server executes this event as the first step in the processing of uploaded information. Upload information is processed in a separate transaction. The execution of this event is the first table-specific action in this transaction.

You can have one begin\_upload script for each table in the remote database. The script is only invoked when the table is actually synchronized.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_upload table event” on page 433](#)
- [“begin\\_upload connection event” on page 387](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

When uploading rows from a remote you may want to place the changes in an intermediate table and manually process changes yourself. You can populate a global temporary table in this event.

```
CALL ml_add_table_script(  
  'version1',  
  'Leads',  
  'begin_upload',  
  'INSERT INTO T_Leads  
  SELECT * FROM Leads  
  WHERE Owner = @EmployeeID' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called beginUploadTable as the script for the begin\_upload table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'begin_upload',  
  'ExamplePackage.ExampleClass.beginUploadTable'  
)
```

The following is the sample Java method beginUploadTable. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String beginUploadTable(  
  String user,  
  String table ) {  
  java.lang.System.out.println("Beginning to process upload for: " + table);  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called BeginTableUpload as the script for the begin\_upload table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
  'ver1',  
  'table1',  
  'begin_upload',  
  'TestScripts.Test.BeginTableUpload'  
)
```

The following is the sample .NET method BeginTableUpload. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string BeginTableUpload(  
  string user,  
  string table ) {  
  System.Console.WriteLine("Beginning to process upload for: " + table);  
  return ( null );  
}
```

## begin\_upload\_deletes table event

Processes statements related to a specific table just before uploading deleted rows from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

This event occurs immediately before applying the changes that result from rows deleted in the client table named in the second parameter.

You can have one begin\_upload\_deletes script for each table in the remote database. The script is only invoked when the table is actually synchronized.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_upload\\_deletes table event” on page 436](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The begin\_upload\_deletes connection event is used to perform whatever steps you need performed after uploading inserts and updates for a particular table, but before deletes are uploaded for that table. The

following SQL script creates a temporary table for storing deletes temporarily during upload. This syntax is for a SQL Anywhere consolidated database.

```
CALL ml_add_table_script(  
  'ver1',  
  'sales_order',  
  'begin_upload_deletes',  
  'CREATE TABLE #sales_order_deletes (  
    id          integer NOT NULL default autoincrement,  
    cust_id     integer NOT NULL,  
    order_date  date NOT NULL,  
    fin_code_id char(2) NULL,  
    region      char(7) NULL,  
    sales_rep   integer NOT NULL,  
    PRIMARY KEY (id) )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called beginUploadDeletes as the script for the begin\_upload\_deletes table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'begin_upload_deletes',  
  'ExamplePackage.ExampleClass.beginUploadDeletes' )
```

The following is the sample Java method beginUploadDeletes. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String beginUploadDeletes(  
  String user,  
  String table )  
  throws java.sql.SQLException {  
  java.lang.System.out.println( "Starting upload  
    deletes for table: " + table );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called BeginUploadDeletes as the script for the begin\_upload\_deletes table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
  'ver1',  
  'table1',  
  'begin_upload_deletes',  
  'TestScripts.Test.BeginUploadDeletes'  
)
```

The following is the sample .NET method BeginUploadDeletes. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string BeginUploadDeletes(  
  string user,  
  string table ) {
```

```
System.Console.WriteLine(  
    "Starting upload deletes for table: " + table );  
return ( null );  
}
```

## begin\_upload\_rows table event

Processes statements related to a specific table just before uploading inserts and updates from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

This event occurs immediately prior to applying the changes that result from inserts and deletes to the client table named in the second parameter.

You can have one begin\_upload\_rows script for each table in the remote database. The script is only invoked when the table is actually synchronized.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_upload\\_rows table event” on page 439](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The begin\_upload\_rows connection event is used to perform whatever steps you need performed before uploading inserts and updates for a particular table. The following script calls a stored procedure that prepares the consolidated database for inserts and updates into the Inventory table:

```
CALL ml_add_table_script(
  'MyCorp 1.0',
  'Inventory',
  'begin_upload_rows',
  'CALL PrepareForUpserts()' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called beginUploadRows as the script for the begin\_upload\_rows table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'begin_upload_rows',
  'ExamplePackage.ExampleClass.beginUploadRows' )
```

The following is the sample Java method beginUploadRows. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String beginUploadRows(
  String user,
  String table )
  throws java.sql.SQLException {
  java.lang.System.out.println(
    "Starting upload rows for table: " +
    table + " and user: " + user );
  return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called BeginUploadRows as the script for the begin\_upload\_rows table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(
  'ver1',
  'table1',
  'begin_upload_rows',
  'TestScripts.Test.BeginUploadRows'
)
```

The following is the sample .NET method BeginUploadRows. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string BeginUploadRows(
  string user,
  string table ) {
  System.Console.WriteLine(
    "Starting upload rows for table: " +
    table + " and user: " + user );
  return ( null );
}
```

## download\_cursor table event

Defines a cursor to select rows to download and insert or update in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

The MobiLink server opens a read-only cursor with which to fetch a list of rows to download to the remote database. This script should contain a suitable SELECT statement.

You can have one download\_cursor script for each table in the remote database.

To optimize performance of the download stage of synchronization to UltraLite clients, when the range of primary key values is outside the current rows on the device, you should order the rows in the download cursor by primary key. Downloads of large reference tables, for example, can benefit from this optimization.

Each download\_cursor script must contain a SELECT statement or a call to a procedure that contains a SELECT statement. The MobiLink server uses this statement to define a cursor in the consolidated database.

The script must select all columns that correspond to the columns in the corresponding table in the remote database. The columns in the consolidated database can have different names than the corresponding columns in the remote database, but they must be of compatible types.

The columns must be selected in the order that the corresponding columns are defined in the remote database.



Note that `download_cursor` allows for cascading deletes. So, you can delete records from a database.

To avoid downloading unnecessary rows, you should include the following line in the `WHERE` clause of your `download_cursor` script:

```
AND last_table_download > '1900/1/1'
```

For Java and .NET applications, this script must return valid SQL.

If you are considering using `READPAST` table hints in `download_cursor` scripts because you are doing lots of updates that affect download performance, consider using snapshot isolation for downloads instead. The `READPAST` table hint can cause problems if used in `download_cursor` scripts. When using timestamp-based downloads, the `READPAST` hint can cause rows to be missed, and can cause a row to never be downloaded to a remote database. For example:

- A row is added to the consolidated database and committed. The row has a `last_modified` column with a time of yesterday.
- The same row is updated but not committed.
- A remote database with a `last_download` time of last week synchronizes.
- A `download_cursor` script attempts to select the row using `READPAST`, and skips the row.
- The transaction that updated the row is rolled back. The next last download time for the remote is advanced to today.

From this point on, the row is never downloaded unless it is updated. A possible workaround is to implement a `modify_next_last_download_timestamp` script and set the last download time to be the start time of the oldest open transaction.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Writing scripts to download rows” on page 333](#)
- [“Writing download\\_cursor scripts” on page 334](#)
- [“Partitioning rows among remote databases” on page 135](#)
- [“download\\_delete\\_cursor table event” on page 400](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)
- ["Using READPAST with MobiLink synchronization" in “FROM clause” \[SQL Anywhere Server - SQL Reference\]](#)

### SQL example

The following example comes from an Oracle installation, although the statement is valid against all supported databases. This example downloads all rows that have been changed since the last time the user downloaded data, and that match the user name in the `emp_name` column.

```
CALL ml_add_table_script(  
  'Lab',  
  'ULOrder',  
  'download_cursor',  
  'SELECT order_id,
```

```
    cust_id,  
    prod_id,  
    emp_id,  
    disc,  
    quant,  
    notes,  
    status  
FROM ULOrder  
WHERE last_modified >= {ml s.last_table_download}  
    AND emp_name = {ml s.username}' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `downloadCursor` as the script for the `download_cursor` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'ULCustomer',  
    'download_cursor',  
    'ExamplePackage.ExampleClass.downloadCursor ' )
```

The following is the sample Java method `downloadCursor`. It returns a SQL statement to download rows where the `last_modified` column is greater than or equal to the last download time.

```
public String downloadCursor(  
    java.sql.Timestamp ts,  
    String user ) {  
    return( "SELECT cust_id, cust_name FROM ULCustomer  
            WHERE last_modified >= ' "  
        + ts + " ' " );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `DownloadCursor` as the script for the `download_cursor` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'download_cursor',  
    'TestScripts.Test.DownloadCursor'  
)
```

The following is the sample .NET method `DownloadCursor`. It populates a temporary table with the contents of a file called `rows.txt`. It then returns a cursor that causes MobiLink to send the rows in the temporary table to the remote database. This syntax is valid for SQL Anywhere consolidated databases.

```
public string DownloadCursor(  
    DateTime ts,  
    string user ) {  
    DBCommand stmt = curConn.CreateCommand();  
    StreamReader input = new StreamReader( "rows.txt" );  
    string sql = input.ReadLine();  
    stmt.CommandText = "DELETE FROM dnet_dl_temp";  
    stmt.ExecuteNonQuery();  
    while( sql != null ){  
        stmt.CommandText = "INSERT INTO dnet_dl_temp VALUES " + sql;  
        stmt.ExecuteNonQuery();  
        sql = input.ReadLine();  
    }
```

```
    }  
    return( "SELECT * FROM dnet_dl_temp" );  
}
```

## download\_delete\_cursor table event

Defines a cursor to select rows that are to be deleted in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

The MobiLink server opens a read-only cursor with which to fetch a list of rows to download, and then insert or update in the remote database. This script must contain a SELECT statement that returns the primary key values of the rows to be deleted from the table in the remote database.

You can have one download\_delete\_cursor script for each table in the remote database.

If the download\_delete\_cursor has nulls for the primary key columns for one or more rows in a table, then MobiLink tells the remote to delete all the data in the table. See [“Deleting all the rows in a table” on page 336](#).

Note that rows deleted from the consolidated database do not appear in a result set defined by a download\_delete\_cursor event, and so are not automatically deleted from the remote database. One technique for identifying rows to be deleted from remote databases is to add a column to the consolidated database table identifying a row as inactive.

To avoid downloading unnecessary rows, you should include the following line in the WHERE clause of your download\_delete\_cursor script:

```
AND last_modified > '1900/1/1'
```

For Java and .NET applications, this script must return valid SQL.

It can be problematic using READPAST table hints in a download\_delete\_cursor. For details, see the download\_cursor event.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “download\_cursor table event” on page 396
- “Writing scripts to download rows” on page 333
- “Partitioning rows among remote databases” on page 135
- “Writing download\_delete\_cursor scripts” on page 335
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Using last download times in scripts” on page 130
- “Using READPAST with MobiLink synchronization” in “FROM clause” [*SQL Anywhere Server - SQL Reference*]

### SQL example

This example is taken from the Contact sample and can be found in *Samples\MobiLink>Contact\build\_consol.sql*. It deletes from the remote database any customers who have been changed since the last time this user downloaded data (`Customer.last_modified >= {ml s.last_table_download}`), and either

- do not belong to the synchronizing user (`SalesRep.username != {ml s.username}`), or
- are marked as inactive in the consolidated database (`Customer.active = 0`).

```
CALL ml_add_table_script(
  'ver1',
  'table1',
  'download_delete_cursor',
  'SELECT cust_id FROM Customer key join SalesRep
   WHERE Customer.last_modified >= {ml s.last_table_download} AND
   ( SalesRep.username != {ml s.username} OR Customer.active = 0 )')
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `downloadDeleteCursor` as the script for the `download_delete_cursor` event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'download_delete_cursor',
  'ExamplePackage.ExampleClass.downloadDeleteCursor' )
```

The following is the sample Java method `downloadDeleteCursor`. It calls a Java method that generates the SQL for the download delete cursor.

```
public String downloadDeleteCursor(
  Timestamp ts,
  String user ) {
```

```
    return( getDownloadCursor( _curUser, _curTable ) );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called DownloadDeleteCursor as the script for the download\_delete\_cursor table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'download_delete_cursor',  
    'TestScripts.Test.DownloadDeleteCursor'  
)
```

The following is the sample .NET method DownloadDeleteCursor. It calls a .NET method that generates the SQL for the download delete cursor.

```
public string DownloadDeleteCursor(  
    DateTime timestamp,  
    string user ) {  
    return( GetDownloadCursor( _curUser, _curTable ) );  
}
```

## download\_statistics connection event

Tracks synchronization statistics for download operations.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name as specified in your SYNCHRONIZATION USER definition.	1
s.warnings	INTEGER. The number of warnings issued.	2
s.errors	INTEGER. The number of errors, including handled errors, that occurred.	3
s.fetched_rows	INTEGER. The number of rows fetched by the download_cursor script.	4
s.deleted_rows	INTEGER. The number of rows fetched by the download_delete_cursor script.	5
s.filtered_rows	INTEGER. The number of rows from the deleted_rows parameter actually sent to the remote. This reflects download filtering of uploaded values.	6
s.bytes	INTEGER. The number of bytes sent to the remote as the download.	7

### Default action

None.

### Remarks

The `download_statistics` event allows you to gather, for any user, statistics on downloads. The `download_statistics` connection script is called just prior to the commit at the end of the download transaction.

#### Note

Depending on the command line, not all warnings or errors are logged, so the warnings and errors counts may be more than the number of warnings or errors logged.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“download\\_statistics table event” on page 406](#)
- [“upload\\_statistics connection event” on page 512](#)
- [“upload\\_statistics table event” on page 517](#)
- [“synchronization\\_statistics connection event” on page 486](#)
- [“synchronization\\_statistics table event” on page 489](#)
- [“time\\_statistics connection event” on page 492](#)
- [“time\\_statistics table event” on page 495](#)
- [“MobiLink Monitor” on page 179](#)
- [“Using remote IDs and MobiLink user names in scripts” \[\*MobiLink - Client Administration\*\]](#)

### SQL example

The following example inserts synchronization statistics into a table called `download_audit`.

```
CALL ml_add_connection_script(  
  'ver1',  
  'download_statistics',  
  'INSERT INTO download_audit(  
    user_name,  
    warnings,  
    errors,  
    deleted_rows,  
    fetched_rows,  
    download_rows,  
    bytes )  
VALUES (  
  {ml s.username},  
  {ml s.warnings},  
  {ml s.errors},  
  {ml s.fetched_rows},  
  {ml s.deleted_rows},  
  {ml s.filtered_rows},  
  {ml s.bytes})')
```

Once vital statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

### Java example

The following call to a MobiLink system procedure registers a Java method called `downloadStatisticsConnection` as the script for the `download_statistics` event when synchronizing the script version `ver1`.



```
CALL ml_add_java_connection_script(  
  'ver1',  
  'download_statistics',  
  'ExamplePackage.ExampleClass.downloadStatisticsConnection' )
```

The following is the sample Java method `downloadStatisticsConnection`. It prints the number of fetched rows to the MobiLink message log. (Note that printing the number of fetched rows to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String downloadStatisticsConnection(  
  String user,  
  int warnings,  
  int errors,  
  int fetchedRows,  
  int deletedRows,  
  int bytes ) {  
  java.lang.System.out.println(  
    "download connection stats fetchedRows: "  
    + fetchedRows );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `DownloadStats` as the script for the `download_statistics` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'download_statistics',  
  'TestScripts.Test.DownloadStats'  
)
```

The following is the sample .NET method `DownloadStats`. It prints the number of fetched rows to the MobiLink message log. (Note that printing the number of fetched rows to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string DownloadStats(  
  string user,  
  int warnings,  
  int errors,  
  int deletedRows,  
  int fetchedRows,  
  int downloadRows,  
  int bytes ) {  
  System.Console.WriteLine(  
    "download connection stats fetchedRows: "  
    + fetchedRows );  
  return ( null );  
}
```

## download\_statistics table event

Tracks synchronization statistics for download operations by table.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name as specified in your SYNCHRONIZATION USER definition.	1
s.table	VARCHAR(128). The table name.	2
s.warnings	INTEGER. The number of warnings issued.	3
s.errors	INTEGER. The number of errors, including handled errors, that occurred.	4
s.fetched_rows	INTEGER. The number of rows fetched by the download_cursor script.	5
s.deleted_rows	INTEGER. The number of rows fetched by the download_delete_cursor script.	6
s.filtered_rows	INTEGER. The number of rows from (6) actually sent to the remote. This reflects download filtering of uploaded values.	7
s.bytes	INTEGER. The number of bytes sent to the remote as the download.	8

### Default action

None.

## Remarks

The download\_statistics event allows you to gather, for any user and table, statistics on downloads as they apply to that table. The download\_statistics table script is called just prior to the commit at the end of the download transaction.

## See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “download\_statistics connection event” on page 403
- “upload\_statistics connection event” on page 512
- “upload\_statistics table event” on page 517
- “synchronization\_statistics connection event” on page 486
- “synchronization\_statistics table event” on page 489
- “time\_statistics connection event” on page 492
- “time\_statistics table event” on page 495
- “MobiLink Monitor” on page 179
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

## SQL example

The following example inserts synchronization statistics into a table called download\_audit. Once vital statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

```
CALL ml_add_table_script(
  'ver1',
  'table1',
  'download_statistics',
  'INSERT INTO download_audit (
    user_name,
    table, warnings,
    errors,
    deleted_rows,
    fetched_rows,
    download_rows,
    bytes)
  VALUES (
    {ml s.username},
    {ml s.table},
    {ml s.warnings},
    {ml s.errors},
    {ml s.fetched_rows},
    {ml s.deleted_rows},
    {ml s.filtered_rows},
    {ml s.bytes})')
```

## Java example

The following call to a MobiLink system procedure registers a Java method called downloadStatisticsTable as the script for the download\_statistics table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'download_statistics',
  'ExamplePackage.ExampleClass.downloadStatisticsTable' )
```

The following is the sample Java method `downloadStatisticsTable`. It prints some statistics for this table to the MobiLink message log. (Note that printing statistics for a table to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String downloadStatisticsTable(
    String user,
    String table,
    int warnings,
    int errors,
    int fetchedRows,
    int deletedRows,
    int bytes ) {
    java.lang.System.out.println( "download table stats "
        + "table: " + table + "bytes: " + bytes );
    return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `DownloadTableStats` as the script for the `download_statistics` table event when synchronizing the script version `ver1` and the table `table1`.

```
CALL ml_add_dnet_table_script(
    'ver1',
    'table1',
    'download_statistics',
    'TestScripts.Test.DownloadTableStats'
)
```

The following is the sample .NET method `DownloadTableStats`. It prints some statistics for this table to the MobiLink message log. (Note that printing statistics for a table to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string DownloadTableStats(
    string user,
    string table,
    int warnings,
    int errors,
    int deletedRows,
    int fetchedRows,
    int downloadRows,
    int bytes ) {
    System.Console.WriteLine( "download table stats "
        + "table: " + table + "bytes: " + bytes );
    return ( null );
}
```

---

## end\_connection connection event

Processes any statements just before the MobiLink server closes a connection with the consolidated database server, either in preparation to shut down or when a connection is removed from the connection pool.

### Parameters

None.

### Default action

None.

### Remarks

You can use the end\_connection script to perform an action of your choice just prior to closing of a connection between the MobiLink server and the consolidated database server.

This script is normally used to complete any actions started by the begin\_connection script and free any resources acquired by it.

### See also

- [“begin\\_connection connection event” on page 367](#)
- [“Adding and deleting scripts” on page 327](#)

### SQL example

The following SQL script drops a temporary table that was created by the begin\_connection script. This syntax is for a SQL Anywhere consolidated database. Strictly speaking, this table doesn't need to be dropped explicitly, since SQL Anywhere does this automatically when the connection is destroyed. Whether a temporary table needs to be dropped explicitly depends on your consolidated database type.

```
CALL ml_add_connection_script(  
    'version 1.0',  
    'end_connection',  
    'DROP TABLE #sync_info' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endConnection as the script for the end\_connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(  
    'ver1',  
    'end_connection',  
    'ExamplePackage.ExampleClass.endConnection' )
```

The following is the sample Java method endConnection. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String endConnection() {  
    java.lang.System.out.println( "Ending connection." );  
    return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndConnection as the script for the end\_connection connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'end_connection',  
    'TestScripts.Test.EndConnection'  
)
```

The following is the sample .NET method EndConnection. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string EndConnection() {  
    System.Console.WriteLine( "Ending connection." );  
    return ( null );  
}
```

## end\_download connection event

Processes any statements just after the MobiLink server concludes preparation of the download data.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_download	TIMESTAMP. The last download times of any synchronized table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

The MobiLink server executes this script after all rows have been downloaded. If you are using blocking download acknowledgement, the script is executed after the confirmation of receipt has been received. Download information is processed in a single transaction. The execution of this script is the last non statistical action in this transaction.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_download connection event” on page 369](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

The following example shows one possible use of an end\_download connection script. The ULEmpCust table has an action column. The following script uses the value in this column to delete records from the remote database.

```
CALL ml_add_connection_script(
  'ver1',
  'end_download',
  'DELETE FROM ULEmpCust ec
   WHERE ec.emp_id = {ml s.username} AND action = 'D''')
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endDownloadConnection as the script for the end\_download connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'end_download',
  'ExamplePackage.ExampleClass.endDownloadConnection' )
```

The following is the sample Java method endDownloadConnection. The ULEmpCust table has an action column. The following script uses the value in this column to delete records from the remote database. It also uses the current MobiLink connection (saved earlier) to perform an update before the download ends. The SQL syntax is for SQL Anywhere consolidated databases.

```
public String endDownloadConnection(
  Timestamp ts,
  String user )
  throws java.sql.SQLException {
  String del_sql = "DELETE FROM ULEmpCust ec " +
    "WHERE ec.emp_id = '" + user + "' " +
    "AND action = 'D' ";
  execUpdate( _syncConn, del_sql );
  return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndDownload as the script for the end\_download connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(
  'ver1',
  'end_download',
  'TestScripts.Test.EndDownload' )
```

The following is the sample .NET method EndDownload. The ULEmpCust table has an action column. The following script uses the value in this column to delete records from the remote database. It also uses the current MobiLink connection (saved earlier) to perform an update before the download ends. The SQL syntax is for SQL Anywhere consolidated databases.

```
public string EndDownload(
  DateTime timestamp,
  string user ) {
  string del_sql = "DELETE FROM ULEmpCust ec " +
    "WHERE ec.emp_id = '" + user + "' " +
    "AND action = 'D' ";
```



```
    execUpdate( _syncConn, del_sql );  
    return ( null );  
}
```

## end\_download table event

Processes statements related to a specific table just after the MobiLink server concludes preparing the stream of downloaded inserts, updates, and deletes.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.table	VARCHAR(128). The table name.	3

### Default action

None.

### Remarks

The MobiLink server executes this script after all rows have been downloaded and confirmation of receipt has been received. The download information is prepared in a separate transaction. The execution of this script is the last table-specific, non-statistical action in this transaction.

You can have one end\_download script for each table in the remote database.

**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “begin\_download table event” on page 371
- “end\_download connection event” on page 411
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Using last download times in scripts” on page 130

**SQL example**

The end\_download table event is used to perform whatever steps you need performed after downloading a particular table. The following SQL Anywhere SQL script drops a temporary table created by a prepare\_for\_download script to hold download rows for the sales\_summary table.

```
CALL ml_add_table_script(
  'MyCorp 1.0',
  'sales_summary',
  'end_download',
  'DROP TABLE #sales_summary_download' )
```

**Java example**

The following call to a MobiLink system procedure registers a Java method called endDownloadTable as the script for the end\_download table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script (
  'ver1',
  'table1',
  'end_download',
  'ExamplePackage.ExampleClass.endDownloadTable' )
```

The following is the sample Java method endDownloadTable. It resets the current table member variable.

```
public String endDownloadTable(
  Timestamp ts,
  String user,
  String table ) {
  _curTable = null;
  return ( null );
}
```

**.NET example**

The following call to a MobiLink system procedure registers a .NET method called EndTableDownload as the script for the end\_download table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(
  'ver1',
  'table1',
  'end_download',
  'TestScripts.Test.EndTableDownload'
)
```

The following is the sample .NET method EndTableDownload. It resets the current table member variable.

```
public string EndTableDownload
  DateTime timestamp,
  string user,
```

```
string table ) {  
  _curTable = null;  
  return ( null );  
}
```

## end\_download\_deletes table event

Processes statements related to a specific table just after preparing a list of rows to be deleted from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.table	VARCHAR(128). The table name.	3

### Default action

None.

### Remarks

This script is executed immediately after preparing a list of rows to be deleted from the named table in the remote database.

You can have one end\_download\_deletes script for each table in the remote database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_download\\_deletes table event” on page 374](#)
- [“end\\_download\\_connection event” on page 411](#)
- [“begin\\_download\\_rows table event” on page 377](#)
- [“end\\_download\\_rows table event” on page 420](#)
- [“download\\_delete\\_cursor table event” on page 400](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

You may want to mark a row as deleted on the remote database. The following script updates a column in the consolidated database called OnRemote.

**Note**

The WHERE clause on the UPDATE matches the WHERE clause used for your download\_delete\_cursor event script.

```
CALL ml_add_table_script(  
  'version1',  
  'Leads',  
  'end_download_deletes',  
  'UPDATE Leads SET OnRemote = 0  
  WHERE LastModified >= {ml s.last_table_download}  
  AND Owner = {ml s.username} AND DeleteFlag=1');
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endDownloadDeletes as the script for the end\_download\_deletes table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'end_download_deletes',  
  'ExamplePackage.ExampleClass.endDownloadDeletes' )
```

You may want to mark a row as deleted on the remote database. The following is the sample Java method endDownloadDeletes. It updates a column in the consolidated database called OnRemote to indicate the record no longer resides on the remote database.

**Note**

The WHERE clause on the UPDATE matches the WHERE clause used for your download\_delete\_cursor event script.

```
public String endDownloadDeletes(  
  Timestamp ts,  
  String user,  
  String table ) {  
  return( "UPDATE Leads SET OnRemote = 0  
  WHERE LastModified >= {ml s.last_table_download}" );  
}
```

```
        AND Owner = {ml s.username} AND DeleteFlag=1" );  
    }
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndDownloadDeletes as the script for the end\_download\_deletes table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'end_download_deletes',  
    'TestScripts.Test.EndDownloadDeletes'  
)
```

You may want to mark a row as deleted on the remote database. The following is the sample .NET method EndDownloadDeletes. It updates a consolidated database column called OnRemote to indicate that the record no longer resides on the remote database. The WHERE clause on the UPDATE matches the WHERE clause used for your download\_delete\_cursor event script.

```
public string EndDownloadDeletes(  
    DateTime timestamp,  
    string user,  
    string table) {  
    return( "UPDATE Leads SET OnRemote = 0  
        WHERE LastModified >= {ml s.last_table_download}  
        AND Owner = {ml s.username} AND DeleteFlag=1" );  
}
```

## end\_download\_rows table event

Processes statements related to a specific table just after preparing a list of rows to be inserted or updated in the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_table_download	TIMESTAMP. The last download time for the table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.table	VARCHAR(128). The table name.	3

### Default action

None.

### Remarks

This script is executed immediately after preparing the stream of rows to be inserted or updated in the named table in the remote database.

You can have one end\_download\_rows script for each table in the remote database.



**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “begin\_download\_rows table event” on page 377
- “end\_download connection event” on page 411
- “end\_download\_deletes table event” on page 417
- “begin\_download\_deletes table event” on page 374
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Using last download times in scripts” on page 130

**SQL example**

You may want to mark a row as successfully downloaded to the remote database. The following script updates a column in the consolidated database called OnRemote.

**Note**

The WHERE clause on the UPDATE matches the WHERE clause used for your download\_delete\_cursor event script.

```
CALL ml_add_table_script(
  'version1',
  'Leads',
  'end_download_rows',
  'UPDATE Leads SET OnRemote = 1
   WHERE LastModified >= {ml s.last_table_download}
   AND Owner = {ml s.username}
   AND DownloadFlag=1' );
```

**Java example**

The following call to a MobiLink system procedure registers a Java method called endDownloadRows as the script for the end\_download\_rows table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'end_download_rows',
  'ExamplePackage.ExampleClass.endDownloadRows' )
```

The following is the sample Java method endDownloadRows. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String endDownloadRows(
  Timestamp ts,
  String user,
  String table ) {
  java.lang.System.out.println(
    "Done downloading inserts and updates for table "
    + table );
  return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndDownloadRows as the script for the end\_download\_rows table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'end_download_rows',  
    'TestScripts.Test.EndDownloadRows'  
)
```

The following is the sample .NET method EndDownloadRows. It prints a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string EndDownloadRows(  
    DateTime timestamp,  
    string user,  
    string table ) {  
    System.Console.WriteLine(  
        "Done downloading inserts and updates for table "  
        + table );  
    return ( null );  
}
```

## end\_publication connection event

Provides useful information about the publication(s) being synchronized.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.generation_number	INTEGER. If your deployment does not use file-based downloads, this parameter can be ignored. The default value is 1.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.publication_name	VARCHAR(128)	3
s.last_publication_upload	TIMESTAMP. Last successful upload time of this publication.	4
s.last_publication_download	TIMESTAMP. The last download time of this publication.	5
s.subscription_id	VARCHAR(128). The subscription ID.	6

### Default action

None.

### Remarks

This event lets you design synchronization logic based on the publications currently being synchronized. This event is invoked in the same transaction as the end\_synchronization event, and is invoked before the end\_synchronization event. It is invoked once per publication being synchronized.

If the current synchronization successfully applied an upload, the last\_upload parameter contains the time this latest upload was applied. If you are using blocking download acknowledgement and the current synchronization has a successful download acknowledgement, the last\_download time contains the time this

latest download was generated. This is the same value that was passed to the download scripts as the last download time.

If an UltraLite remote is synchronizing with UL\_SYNC\_ALL, this event is invoked once with the name 'unknown'.

### Generation number

The generation\_number parameter is specifically for file-based downloads.

The output value of the generation number is passed from the begin\_publication script to the end\_publication script. The meaning of the generation\_number depends on whether the current synchronization is being used to create a download file, or whether the current synchronization has an upload.

In file-based downloads, generation numbers are used to force an upload before the download. The number is stored in the download file.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_publication connection event” on page 380](#)
- [“MobiLink file-based download” on page 293](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

You may want to record the information for each publication being synchronized. The following example calls ml\_add\_connection\_script to assign the event to a stored procedure called RecordPubEndSync.

```
CALL ml_add_connection_script(  
  'version1',  
  'end_publication',  
  'CALL RecordPubEndSync(  
    {ml s.generation_number},  
    {ml s.username},  
    {ml s.publication_name},  
    {ml s.last_publication_upload},  
    {ml s.last_publication_download} )' );
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endPublication as the script for the end\_publication connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'end_publication',  
  'ExamplePackage.ExampleClass.endPublication' );
```

The following is the sample Java method endPublication. It outputs a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String endPublication(  
  anywhere.ml.script.InOutInteger generation_number,
```

```
String user,  
String pub_name,  
Timestamp last_publication_upload,  
Timestamp last_publication_download ) {  
java.lang.System.out.println(  
    "Finished synchronizing publication " + pub_name );  
return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndPub as the script for the end\_publication connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'end_publication',  
    'TestScripts.Test.EndPub'  
)
```

The following is the sample .NET method endPub. It outputs a message to the MobiLink message log. (Note that printing a message to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string EndPub(  
    ref int generation_number,  
    string user,  
    string pub_name,  
    DateTime last_publication_upload,  
    DateTime last_publication_download ) {  
    System.Console.Write(  
        "Finished synchronizing publication " + pub_name );  
    return ( null );  
}
```

## end\_synchronization connection event

Processes any statements at the time an application disconnects from the MobiLink server upon completion of the synchronization process.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.synchronization_ok	INTEGER. This value is 1 for a successful synchronization and 0 for an unsuccessful synchronization.	2

### Default action

None.

### Remarks

The MobiLink server executes this script after synchronization is complete. If you are using blocking download acknowledgement, the script is executed after the MobiLink client has returned confirmation of receipt of the download.

This script is executed within a separate transaction after the download transaction.

The end\_synchronization script is useful for maintaining statistics. This is because if the begin\_synchronization script is called, the end\_synchronization script is invoked even if there is an error or conflict, so while the upload transaction is rolled back, statistics are maintained.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_synchronization connection event” on page 383](#)
- [“begin\\_synchronization table event” on page 385](#)
- [“end\\_synchronization table event” on page 428](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

**SQL example**

The following SQL script calls a system procedure that records the end time of the synchronization attempt along with its success or failure status. This syntax is for SQL Anywhere consolidated databases.

```
CALL ml_add_connection_script(
  'ver1',
  'end_synchronization',
  'CALL RecordEndOfSyncAttempt(
    {ml s.username},
    {ml s.synchronization_ok} )' )
```

**Java example**

The following call to a MobiLink system procedure registers a Java method called endSynchronizationConnection as the script for the end\_synchronization event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'end_synchronization',
  'ExamplePackage.ExampleClass.endSynchronizationConnection'
)
```

The following is the sample Java method endSynchronizationConnection. It uses a JDBC connection to execute an update. This syntax is for SQL Anywhere consolidated databases.

```
public String endSynchronizationConnection(
  String user )
  throws java.sql.SQLException {
  execUpdate( _syncConn,
    "UPDATE sync_count set count = count + 1 where user_id = '"
    + user + "' ");
  return ( null );
}
```

**.NET example**

The following call to a MobiLink system procedure registers a .NET method called EndSync as the script for the end\_synchronization connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(
  'ver1',
  'end_synchronization',
  'TestScripts.Test.EndSync'
)
```

The following is the sample .NET method EndSync. It updates the table sync\_count. This syntax is for SQL Anywhere consolidated databases.

```
public string EndSync(
  string user ) {
  return(
    "UPDATE sync_count set count = count + 1 where user_id = '"
    + user + "' ");
  return ( null );
}
```

## end\_synchronization table event

Processes statements related to a specific table at the time an application disconnects from the MobiLink server upon completion of the synchronization process.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2
s.synchronization_ok	INTEGER. This value is 1 for a successful synchronization and 0 for an unsuccessful synchronization.	3

### Default action

None.

### Remarks

The MobiLink server executes this script after an application has synchronized and is about to disconnect from the MobiLink server, and before the connection level script of the same name.

You can have one end\_synchronization script for each table in the remote database.



**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “begin\_synchronization table event” on page 385
- “end\_synchronization connection event” on page 426
- “end\_synchronization table event” on page 428
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

**SQL example**

The following SQL Anywhere SQL script drops a temporary table created by the begin\_synchronization script.

```
CALL ml_add_table_script(
  'ver1',
  'sales_order',
  'end_synchronization',
  'DROP TABLE #sales_order' )
```

**Java example**

The following call to a MobiLink system procedure registers a Java method called endSynchronizationTable as the script for the end\_synchronization table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'end_synchronization',
  'ExamplePackage.ExampleClass.endSynchronizationTable' )
```

The following is the sample Java method endSynchronizationTable. It returns a SQL statement that drops a temporary table created by the begin\_synchronization script.

```
public String endSynchronizationTable(
  String user,
  String table ) {
  return( "DROP TABLE #sales_order" );
}
```

**.NET example**

The following call to a MobiLink system procedure registers a .NET method called EndTableSync as the script for the end\_synchronization table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(
  'ver1',
  'table1',
  'end_synchronization',
  'TestScripts.Test.EndTableSync'
)
```

The following is the sample .NET method EndTableSync. It returns a SQL to statement that drops a temporary table created by the begin\_synchronization script.

```
public string EndTableSync(
  string user,
  string table ) {
```

```
    return( "DROP TABLE #sales_order" );  
}
```

## end\_upload connection event

Processes any statements just after the MobiLink server concludes processing uploaded inserts, updates, and deletes.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1

### Default action

None.

### Remarks

The MobiLink server executes this script as the last step in the processing of uploaded information. Upload information is processed in a single transaction. The execution of this script is the last action in this transaction before statistical scripts.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_upload connection event” on page 387](#)
- [“end\\_upload table event” on page 433](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The following SQL Anywhere SQL script calls the EndUpload stored procedure.

```
CALL ml_add_connection_script(
  'ver1',
  'sales_order',
  'end_upload',
  'CALL EndUpload({ml s.username});' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endUploadConnection as the script for the end\_upload connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(  
    'ver1',  
    'end_upload',  
    'ExamplePackage.ExampleClass.endUploadConnection' )
```

The following is the sample Java method `endUploadConnection`. It calls a method to perform operations on the database.

```
public String endUploadConnection( String user ) {  
    // Clean up new and old tables.  
    Iterator two_iter = _tables_with_ops.iterator();  
    while( two_iter.hasNext() ) {  
        TableInfo cur_table = (TableInfo)two_iter.next();  
        dumpTableOps( _sync_conn, cur_table );  
    }  
    _tables_with_ops.clear();  
    return ( null );  
}
```

### **.NET example**

The following call to a MobiLink system procedure registers a .NET method called `EndUpload` as the script for the `end_upload` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'end_upload',  
    'TestScripts.Test.EndUpload'  
)
```

The following is the sample .NET method `EndUpload`. It returns a SQL statement that calls the `EndUpload` stored procedure.

```
public string EndUpload( string user ) {  
    return ( "CALL EndUpload({ml s.username});" );  
}
```

## end\_upload table event

Processes statements related to a specific table just after the MobiLink server concludes processing the stream of uploaded inserts, updates, and deletions.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

The MobiLink server executes this script as the last step in the processing of uploaded information. Upload information is processed in a separate transaction. The execution of this script is the last table-specific action in this transaction.

You can have one end\_upload script for each table in the remote database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_upload table event” on page 389](#)
- [“end\\_upload connection event” on page 431](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

## SQL example

The following call to a MobiLink system procedure assigns the end\_upload event to a stored procedure called ULCustomerIDPool\_maintain.

```
CALL ml_add_table_script(  
    'custdb',  
    'ULCustomerIDPool',  
    'end_upload',  
    'CALL ULCustomerIDPool_maintain( username );' );
```

The following SQL statements create the ULCustomerIDPool\_maintain stored procedure.

```
CREATE OR REPLACE PROCEDURE ULCustomerIDPool_maintain(  
    SyncUserID IN integer )  
AS  
    pool_count INTEGER;  
    pool_max    INTEGER;  
BEGIN  
    -- Determine how many ids to add to the pool  
    SELECT COUNT(*)  
        INTO pool_count  
        FROM ULCustomerIDPool  
        WHERE pool_emp_id = SyncUserID;  
    -- Determine the current Customer id max  
    SELECT MAX(pool_cust_id)  
        INTO pool_max  
        FROM ULCustomerIDPool;  
    -- Top up the pool with new ids  
    WHILE pool_count < 20 LOOP  
        pool_max := pool_max + 1;  
        INSERT INTO ULCustomerIDPool(  
            pool_cust_id, pool_emp_id )  
            VALUES ( pool_max, SyncUserID );  
        pool_count := pool_count + 1;  
    END LOOP;  
END;
```

## Java example

The following call to a MobiLink system procedure registers a Java method called endUploadTable as the script for the end\_upload table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1'  
    'end_upload',  
    'ExamplePackage.ExampleClass.endUploadTable' )
```

The following is the sample Java method endUploadTable. It generates a delete for a table with a name related to the passing-in table name. This syntax is for SQL Anywhere consolidated databases.

```
public String endUploadTable(  
    String user,  
    String table ) {  
    return( "DELETE from '" + table + "_temp" );  
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called EndUpload as the script for the end\_upload table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'end_upload',  
    'TestScripts.Test.EndUpload'  
)
```

The following .NET example moves rows inserted into a temporary table into the table passed into the script.

```
public string EndUpload( string user, string table ) {  
    DBCommand stmt = curConn.CreateCommand();  
    // Move the uploaded rows to the destination table.  
    stmt.CommandText = "INSERT INTO "  
        + table  
        + " SELECT * FROM dnet_ul_temp";  
    stmt.ExecuteNonQuery();  
    stmt.Close();  
    return ( null );  
}
```

## end\_upload\_deletes table event

Processes statements related to a specific table just after applying deletes uploaded from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

This script is run immediately after applying the changes that result from rows deleted in the remote table named in the second parameter.

You can have one end\_upload\_deletes script for each table in the remote database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_upload\\_deletes table event” on page 391](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

You can use this event to process rows deleted during the upload on an intermediate table. You can compare the rows in the base table with rows in the intermediate table and decide what to do with the deleted row.



The following call to a MobiLink system procedure assigns the EndUploadDeletesLeads stored procedure to the end\_upload\_deletes event.

```
CALL ml_add_table_script(
  'version1',
  'Leads',
  'end_upload_deletes',
  'call EndUploadDeletesLeads()');
```

The following SQL statement creates the EndUploadDeletes stored procedure.

```
CREATE PROCEDURE EndUploadDeletesLeads ( )
Begin
  FOR names AS curs CURSOR FOR
  SELECT LeadID
  FROM Leads
  WHERE LeadID NOT IN (SELECT LeadID FROM T_Leads);
  DO
  CALL decide_what_to_do( LeadID );
  END FOR;
end
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endUploadDeletes as the script for the end\_upload\_deletes table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'end_upload_deletes',
  'ExamplePackage.ExampleClass.endUploadDeletes' )
```

The following is the sample Java method a endUploadDeletes. It calls a Java method that manipulates the database.

```
public String endUploadDeletes(
  String user,
  String table )
throws java.sql.SQLException {
  processUploadedDeletes( _syncConn, table );
  return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndUploadDeletes as the script for the end\_upload\_deletes table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(
  'ver1',
  'table1',
  'end_upload_deletes',
  'TestScripts.Test.EndUploadDeletes'
)
```

The following is the sample .NET method a EndUploadDeletes. It calls a .NET method that manipulates the database.

```
public string EndUploadDeletes(  
    string user,  
    string table ) {  
    processUploadedDeletes( _syncConn, table );  
    return ( null );  
}
```

## end\_upload\_rows table event

Processes statements related to a specific table just after applying uploaded inserts and updates from the specified table in the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

Uploaded information can require inserting or updating rows in the consolidated database. This script is run immediately after applying the changes that result from modifications to the remote table named in the second parameter.

You can have one end\_upload\_rows script for each table in the remote database.

### See also

- [“Adding and deleting scripts” on page 327](#)
- [“begin\\_upload\\_rows table event” on page 394](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL Example

The following call to a MobiLink system procedure registers a SQL method called endUploadRows as the script for the EndUploadRows table event when synchronizing the script version ver1.

```
CALL ml_add_table_script(  
  'version1',  
  'table1',  
  'end_upload_rows',  
  'CALL EndUploadRows(  
    { ml s.username },  
    { ml s.table } )' )
```

The following is the sample SQL method EndUploadRows. It calls a SQL method that manipulates the database.

```
CREATE PROCEDURE EndUploadRows (  
  IN user VARCHAR(128)  
  IN table VARCHAR{128} )  
BEGIN  
  CALL decide_what_to_do(table);  
END;
```

### Java example

The following call to a MobiLink system procedure registers a Java method called endUploadRows as the script for the end\_upload\_rows table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'end_upload_rows',  
  'ExamplePackage.ExampleClass.endUploadRows' )
```

The following is the sample Java method endUploadRows. It calls a Java method that manipulates the database.

```
public String endUploadRows(  
  String user,  
  String table )  
  throws java.sql.SQLException {  
  processUploadedRows( _syncConn, table );  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called EndUploadRows as the script for the end\_upload\_rows table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
  'ver1',  
  'table1',  
  'end_upload_rows',  
  'TestScripts.Test.EndUploadRows'  
)
```

The following is the sample .NET method endUploadRows. It calls a .NET method that manipulates the database.

```
public string EndUploadRows(  
  string user,  
  string table ) {  
  processUploadedRows( _syncConn, table );  
}
```

```
    return ( null );  
}
```

## handle\_DownloadData connection event

Used by direct row handling to create a set of rows to download.

### Parameters

None.

### Default action

None.

### Remarks

The handle\_DownloadData event allows you to determine what operations to download to MobiLink clients using direct row handling.

Direct row handling is used to synchronize to data sources other than MobiLink supported consolidated databases. See [“Direct row handling” on page 649](#).

To create the direct download, you can use the DownloadData and DownloadTableData classes in the MobiLink server API for Java or .NET.

For Java, the DBConnectionContext getDownloadData method returns a DownloadData instance for the current synchronization. DownloadData encapsulates all download operations to send to a remote client. You can use the DownloadData getDownloadTables and getDownloadTableByName methods to obtain a DownloadTableData instance. DownloadTableData encapsulates download operations for a particular table. You can use the getUpsertPreparedStatement method to obtain prepared statements for insert and update operations. You can use the DownloadTableData getDeletePreparedStatement method to obtain prepared statements for delete operations.

For .NET, the DBConnectionContext GetDownloadData method returns a DownloadData instance for the current synchronization. DownloadData encapsulates all download operations to send to a remote client. You can use the DownloadData GetDownloadTables and GetDownloadTableByName methods to obtain a DownloadTableData instance. DownloadTableData encapsulates download operations for a particular table. You can use the GetUpsertCommand method to obtain commands for insert and update operations. You can use the DownloadTableData getDeleteCommand method to obtain commands for delete operations.

For Java, see [“DBConnectionContext interface” on page 543](#). For .NET, see [“DBConnectionContext interface” on page 609](#).

You can create the download in handle\_DownloadData or another synchronization event. MobiLink provides this flexibility so that you can set the download when data is uploaded or when particular events occur. If you want to create the direct download in an event other than handle\_DownloadData, you must create a handle\_DownloadData script whose method does nothing. MobiLink requires this script to be defined to enable direct row handling. Except in upload-only synchronization, the MobiLink server requires that at a minimum, a handle\_DownloadData script be defined.

If you create the direct download in an event other than handle\_DownloadData, the event must not be before the begin\_synchronization event and cannot be after the end\_download event.

**Note**

This event cannot be implemented as SQL.

**See also**

- [“Direct row handling” on page 649](#)
- [“Handling direct downloads” on page 660](#)
- [Java: “DownloadData interface” on page 548](#)
- [Java: “DownloadTableData interface” on page 550](#)
- [.NET: “DownloadData interface” on page 622](#)
- [.NET: “DownloadTableData interface” on page 623](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Required scripts” on page 326](#)
- [“Adding and deleting scripts” on page 327](#)

**Java example**

The following call to a MobiLink system procedure registers a Java method called handleDownload for the handle\_DownloadData connection event when synchronizing the script version ver1. You run this system procedure against your MobiLink consolidated database.

```
CALL ml_add_java_connection_script(
  'ver1',
  'handle_DownloadData',
  'MyPackage.MyClass.handleDownload' )
```

See [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#).

The following example shows you how to use the handleDownload method to create a download.

The following code sets up a class level DBConnectionContext instance in the constructor for a class called MobiLinkOrders.

```
import ianywhere.ml.script.*;
import java.io.*;
import java.sql.*;
import java.lang.System;

public class MobiLinkOrders{

    DBConnectionContext _cc;

    public MobiLinkOrders( DBConnectionContext cc ) {
        _cc = cc;
    }
}
```

In your HandleDownload method, you use the DBConnectionContext getDownloadData method to return a DownloadData instance for the current synchronization. The DownloadData getDownloadTableByName method returns a DownloadTableData instance for the remoteOrders table. The DownloadTableData getUpsertPreparedStatement method returns a java.sql.PreparedStatement. To add an operation to the download, you set all column values and call the executeUpdate method.

The following is the handleDownload method of the MobiLinkOrders class. It adds two rows to the download for a table called remoteOrders.

```
// Method used for the handle_DownloadData event.
public void handleDownload() throws SQLException {
    // Get DownloadData instance for current synchronization.
    DownloadData downloadData = _cc.getDownloadData();

    // Get a DownloadTableData instance for the remoteOrders table.
    DownloadTableData td = downloadData.getDownloadTableByName("remoteOrders");

    // Get a java.sql.PreparedStatement for upsert (update/insert) operations.
    PreparedStatement upsertPS = td.getUpsertPreparedStatement();

    // Set values for one row.
    upsertPS.setInt( 1, 2300 );
    upsertPS.setInt( 2, 100 );

    // Add the values to the download.
    int updateResult = upsertPS.executeUpdate();

    // Set values for another row.
    upsertPS.setInt( 1, 2301 );
    upsertPS.setInt( 2, 50 );
    updateResult = upsertPS.executeUpdate();

    upsertPS.close();

    // ...
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called HandleDownload as the script for the handle\_DownloadData connection event when synchronizing the script version ver1. This syntax is for SQL Anywhere consolidated databases.

```
CALL ml_add_dnet_connection_script(
    'ver1', 'handle_DownloadData',
    'TestScripts.Test.HandleDownload'
)
```

The following is the sample .NET method HandleDownload:

```
using System;
using System.Data;
using System.IO;
using iAnywhere.MobiLink.Script;
using iAnywhere.MobiLink;

namespace MyScripts
{
    /// <summary>
    /// Tests that scripts are called correctly for most sync events.
    /// </summary>
    public class MobiLinkOrders
    {
        private DBConnectionContext _cc;

        public MobiLinkOrders( DBConnectionContext cc )
        {
            _cc = cc;
        }

        ~MobiLinkOrders()
        {
        }
    }
}
```



```
}  
  
public void handleDownload()  
{  
    // Get DownloadData instance for current synchronization.  
    DownloadData my_dd = _cc.GetDownloadData();  
  
    // Get a DownloadTableData instance for the remoteOrders table.  
    DownloadTableData td = my_dd.GetDownloadTableByName("remoteOrders");  
  
    // Get an IDbCommand for upsert (update/insert) operations.  
    IDbCommand upsert_stmt = td.GetUpsertCommand();  
  
    IDataParameterCollection parameters = upsert_stmt.Parameters;  
  
    // Set values for one row.  
    parameters[ 0 ] = 2300;  
    parameters[ 1 ] = 100;  
  
    // Add the values to the download.  
    int update_result = upsert_stmt.ExecuteNonQuery();  
  
    // Set values for another row.  
    parameters[ 0 ] = 2301;  
    parameters[ 1 ] = 50;  
    update_result = upsert_stmt.ExecuteNonQuery();  
  
    // ...  
}  
}  
}
```

## handle\_error connection event

Executed whenever the MobiLink server encounters a SQL error.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.action_code	INTEGER. This is an INOUT parameter.	1
s.error_code	INTEGER	2
s.error_message	TEXT	3
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	4
s.table	VARCHAR(128). If the script is not a table script, the table name is null.	5

### Default action

When no handle\_error script is defined or this script causes an error, the default action code is 3000: roll back the current transaction and cancel the current synchronization.

### Remarks

The MobiLink server sends in the current action code. Initially, this is set to 3000 for each set of errors caused by a single SQL operation. Usually, there is only one error per SQL operation, but there may be more. This handle\_error script is called once per error in the set. The action code passed into the first error is 3000. Subsequent calls are passed in the action code returned by the previous call. MobiLink uses the highest numerical value returned from multiple calls.

You can modify the action code in the script, and return a value instructing MobiLink how to proceed. The action code tells the MobiLink server what to do next. Before it calls this script, the MobiLink server sets the action code to a default value, which depends upon the severity of the error. Your script may modify this value. Your script must return or set an action code.

The action code parameter takes one of the following values:

- **1000** Skip the current row and continue processing.
- **3000** Roll back the current transaction and cancel the current synchronization. This is the default action code, and is used when no handle\_error script is defined or this script causes an error.
- **4000** Roll back the current transaction, cancel the synchronization, and shut down the MobiLink server.

The error codes and message allow you to identify the nature of the error. If the error happened as part of synchronization, the user name is supplied. Otherwise, this value is null.

The MobiLink server executes this script if an ODBC error occurs while MobiLink is processing an insert, update, or delete script during the upload transaction or is fetching download rows. If an ODBC error occurs at another time, the MobiLink server calls the report\_error or report\_ODBC\_error script and aborts the synchronization.

If the error happened while manipulating a particular table, the table name is supplied. Otherwise, this value is null. The table name is the name of a table in the client application. This name may or may not have a direct counterpart in the consolidated database, depending upon the design of the synchronization system.

SQL scripts for the handle\_error event must be implemented as stored procedures.

You can return a value from the handle\_error script one of the following ways:

- Pass the action parameter to an OUTPUT parameter of a procedure:

```
CALL my_handle_error( {ml s.action_code}, {ml s.error_code}, {ml
s.error_message}, {ml s.username}, {ml s.table} )
```

- Set the action code via a procedure or function return value:

```
{ml s.action_code} = CALL my_handle_error( {ml s.error_code}, {ml
s.error_message}, {ml s.username}, {ml s.table} )
```

Most RDBMSs use the RETURN statement to set the return value from a procedure or function.

The CustDB sample application contains error handlers for various database-management systems.

## See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“report\\_error connection event” on page 477](#)
- [“report\\_odbc\\_error connection event” on page 480](#)
- [“handle\\_odbc\\_error connection event” on page 450](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

## SQL example

The following example works with a SQL Anywhere consolidated database. It allows your application to ignore redundant inserts.

The following call to a MobiLink system procedure assigns the ULHandleError stored procedure to the handle\_error event.

```
CALL ml_add_connection_script(
    'ver1',
    'handle_error',
    'CALL ULHandleError(
        {ml s.action_code},
        {ml s.error_code},
        {ml s.error_message},
        {ml s.username} ) ,
    {ml s.table} )'
```

The following SQL statement creates the ULHandleError stored procedure.

```
CREATE PROCEDURE ULHandleError(
    INOUT action integer,
    IN error_code integer,
    IN error_message varchar(1000),
    IN user_name varchar(128),
    IN table_name varchar(128) )
BEGIN
    -- -196 is SQLE_INDEX_NOT_UNIQUE
    -- -194 is SQLE_INVALID_FOREIGN_KEY
    IF error_code = -196 or error_code = -194 then
        -- ignore the error and keep going
        SET action = 1000;
    ELSE
        -- abort the synchronization
        SET action = 3000;
    END IF;
END
```

## Java example

The following call to a MobiLink system procedure registers a Java method called handleError as the script for the handle\_error connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
    'ver1',
    'handle_error',
    'ExamplePackage.ExampleClass.handleError' )
```

The following is the sample Java method handleError. It processes an error based on the data that is passed in. It also determines the resulting error code.

```
public String handleError(
    anywhere.ml.script.InOutInteger actionCode,
    int errorCode,
    String errorMessage,
    String user,
    String table ) {
    int newAC;
    if( user == null ) {
        newAC = handleNonSyncError( errorCode,
            errorMessage );
    }
    else if( table == null ) {
```

```

        newAC = handleConnectionError( errorCode,
        errorMessage, user ); }
    else {
        newAC = handleTableError( errorCode,
        errorMessage, user, table );
    }
    // Keep the most serious action code.
    if( actionCode.getValue() < newAC ) {
        actionCode.setValue( newAC );
    }
}
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called HandleError as the script for the handle\_error connection event when synchronizing the script version ver1.

```

CALL mll_add_dnet_connection_script(
'ver1',
'handle_error',
'TestScripts.Test.HandleError' )

```

The following is the sample .NET method HandleError.

```

public string HandleError() (
    ref int actionCode,
    int errorCode,
    string errorMessage,
    string user,
    string table ) {
    int new_ac;
    if( user == null ) {
        new_ac = HandleNonSyncError( errorCode,
        errorMessage ); }
    else if( table == null ) {
        new_ac = HandleConnectionError( errorCode,
        errorMessage, user ); }
    else {
        new_ac = HandleTableError( errorCode,
        errorMessage, user, table );
    }
    // Keep the most serious action code.
    if( actionCode < new_ac ) {
        actionCode = new_ac;
    }
}
}

```

## handle\_odbc\_error\_connection\_event

Executed whenever the MobiLink server encounters an error triggered by the ODBC Driver Manager.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.action_code	INTEGER. This is an INOUT parameter.	1
s.ODBC_state	VARCHAR(5)	2
s.error_message	TEXT	3
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	4
s.table	VARCHAR(128)	5

### Default action

The MobiLink server selects a default action code. You can modify the action code in the script, and return a value instructing MobiLink how to proceed. The action code parameter takes one of the following values:

- **1000** Skip the current row and continue processing.
- **3000** Rollback the current transaction and cancel the current synchronization. This is the default action code, and is used when no handle\_error script is defined or this script causes an error.
- **4000** Rollback the current transaction, cancel the synchronization, and shut down the MobiLink server.

### Remarks

The MobiLink server executes this script whenever it encounters an error flagged by the ODBC Driver Manager if the error occurs while MobiLink is processing an insert, update, or delete script during the upload transaction or is fetching download rows. If an ODBC error occurs at another time, the MobiLink server calls the report\_error or report\_ODBC\_error script and aborts the synchronization.

The error codes allow you to identify the nature of the error.

The action code tells the MobiLink server what to do next. Before it calls this script, the MobiLink server sets the action code to a default value, which depends upon the severity of the error. Your script may modify this value. Your script must return or set an action code.

The handle\_odbc\_error script is called after the handle\_error and report\_error scripts, and before the report\_odbc\_error script.

When only one, but not both, error-handling script is defined, the return value from that script decides error behavior. When both error-handling scripts are defined, the MobiLink server uses the numerically highest action code. If both handle\_error and handle\_ODBC\_error are defined, MobiLink uses the action code with the highest numerical value returned from all calls.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“handle\\_error connection event” on page 446](#)
- [“report\\_error connection event” on page 477](#)
- [“report\\_odbc\\_error connection event” on page 480](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The following example works with a SQL Anywhere consolidated database. It allows your application to ignore ODBC integrity constraint violations.

The following call to a MobiLink system procedure assigns the HandleODBCError stored procedure to the handle\_odbc\_error event.

```
CALL ml_add_connection_script(
  'ver1',
  'handle_odbc_error',
  'CALL HandleODBCError(
    {ml s.action_code},
    {ml s.ODBC_state},
    {ml s.error_message},
    {ml s.username},
    {ml s.table} )' )
```

The following SQL statement creates the HandleODBCError stored procedure.

```
CREATE PROCEDURE HandleODBCError(
  INOUT action integer,
  IN odbc_state varchar(5),
  IN error_message varchar(1000),
  IN user_name varchar(128),
  IN table_name varchar(128) )
BEGIN
  IF odbc_state = '23000' then
    -- Ignore the error and keep going.
    SET action = 1000;
  ELSE
    -- Abort the synchronization.
    SET action = 3000;
  END IF;
END
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `handleODBCError` as the script for the `handle_odbc_error` event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'handle_odbc_error',  
  'ExamplePackage.ExampleClass.handleODBCError'  
)
```

The following is the sample Java method `handleODBCError`. It processes an error based on the data that is passed in. It also determines the resulting error code.

```
public String handleODBCError(  
  anywhere.ml.script.InOutInteger actionCode,  
  String ODBCState,  
  String errorMessage,  
  String user,  
  String table ) {  
  int newAC;  
  if( user == null ) {  
    newAC = handleNonSyncError( ODBCState,  
      errorMessage );  
  }  
  else if( table == null ) {  
    newAC = handleConnectionError( ODBCState,  
      errorMessage, user );  
  } else {  
    newAC = handleTableError( ODBCState,  
      errorMessage, user, table );  
  }  
  // Keep the most serious action code.  
  if( actionCode.getValue() < newAC ) {  
    actionCode.setValue( newAC );  
  }  
  return( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `HandleODBCError` as the script for the `handle_odbc_event` when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'handle_odbc_error',  
  'TestScripts.Test.HandleODBCError' )
```

The following is the sample .NET method `HandleODBCError`.

```
public string HandleODBCError (  
  ref int actionCode,  
  string ODBCState,  
  string errorMessage,  
  string user,  
  string table ) {  
  int new_ac;  
  if( user == null ) {  
    new_ac = HandleNonSyncError( ODBCState,  
      errorMessage );  
  }  
}
```



```
else if( table == null ) {
    new_ac = HandleConnectionError( ODBCState,
        errorMessage, user );
} else {
    new_ac = HandleTableError( ODBCState,
        errorMessage, user, table );
}
// Keep the most serious action code.
if( actionCode < new_ac ) {
    actionCode = new_ac;
}
return( null );
}
```

## handle\_UploadData connection event

Used by direct row handling to process uploaded rows.

### Parameters

Parameter name for SQL scripts	Description	Order
UploadData	A .NET or Java class encapsulating table operations uploaded by a MobiLink client. This class is defined in the MobiLink server API for Java and .NET.	1

### Default action

None.

### Remarks

The handle\_UploadData event allows you to process the upload for MobiLink direct row handling. This event fires once for each upload transaction in a synchronization, unless you are using transaction-level uploads, in which case it fires for each transaction.

See [“Direct row handling” on page 649](#).

This event takes a single UploadData parameter. Your Java or .NET method can use the UploadData getUploadedTables or getUploadedTableByName methods to obtain UploadedTableData instances. UploadedTableData allows you to access insert, update, and delete operations uploaded by a MobiLink client in the current synchronization.

For more information about the UploadData and UploadedTableData classes, see [“Handling direct uploads” on page 654](#).

If you want to read column name metadata, you must specify the SendColumnNames MobiLink client extended option or property. Otherwise you can refer to columns by index, as defined at the remote database.

See [“SendColumnNames \(scn\) extended option” \[MobiLink - Client Administration\]](#) and [“Send Column Names synchronization parameter” \[UltraLite - Database Management and Reference\]](#).

**Note**

This event cannot be implemented as SQL.

**See also**

- [“Direct row handling” on page 649](#)
- [“Handling direct uploads” on page 654](#)
- [Java: “UploadData interface” on page 579](#)
- [Java: “UploadedTableData interface” on page 581](#)
- [.NET: “UploadData interface” on page 641](#)
- [.NET: “UploadedTableData interface” on page 643](#)
- [dbmsync: “SendColumnNames \(scn\) extended option” \[\*MobiLink - Client Administration\*\]](#)
- [UltraLite: “Send Column Names synchronization parameter” \[\*UltraLite - Database Management and Reference\*\]](#)
- [“handle\\_DownloadData connection event” on page 442](#)
- [“Required scripts” on page 326](#)
- [“Adding and deleting scripts” on page 327](#)

**Java examples**

The following call to a MobiLink system procedure registers a Java method called handleUpload for the handle\_UploadData connection event when synchronizing the script version ver1. You run this system procedure against your MobiLink consolidated database.

```
CALL ml_add_java_connection_script(
    'ver1',
    'handle_UploadData',
    'MyPackage.MyClass.handleUpload' )
```

For more information about ml\_add\_java\_connection\_script, see [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#).

The following Java method processes the upload for the remoteOrders table. The UploadData.getUploadedTableByName method returns an UploadedTableData instance for the remoteOrders table. The UploadedTableData.getInserts method returns a java.sql.ResultSet instance representing new rows.

```
import ianywhere.ml.script.*;
import java.sql.*;
import java.io.*;
// ...

public void handleUpload( UploadData ut )
    throws SQLException, IOException {
    // Get an UploadedTableData instance representing the
    // remoteOrders table.
    UploadedTableData remoteOrdersTable =
    ut.getUploadedTableByName("remoteOrders");
    // Get inserts uploaded by the MobiLink client.
    java.sql.ResultSet results = remoteOrdersTable.getInserts();
    while( results.next() ) {
        // You can reference column names here because SendColumnNames is on
        // Get the primary key.
        int pk = results.getInt("pk");

        // Get the uploaded num_ordered value.
        int numOrdered = results.getInt("num_ordered");

        // The current insert row is now ready to be uploaded to wherever
        // you want it to go (a file, a web service, and so on).
```

```
    }  
    results.close();  
}
```

The following example outputs insert, update and delete operations uploaded by a MobiLink remote database. The UploadData getUploadedTables method returns UploadedTableData instances representing all tables uploaded by a remote. The order of the tables in this array is the order in which they were uploaded by the remote. The UploadedTableData getInserts, getUpdates, and getDeletes methods return standard JDBC result sets. You can use the println method or output data to a text file or another location.

```
import ianywhere.ml.script.*;  
import java.sql.*;  
import java.io.*;  
// ...  
  
public void handleUpload( UploadData ud )  
    throws SQLException, IOException {  
    UploadedTableData tables[] = ud.getUploadedTables();  
    for( int i = 0; i < tables.length; i++ ) {  
        UploadedTableData currentTable = tables[i];  
        println( "table " + java.lang.Integer.toString( i ) +  
            " name: " + currentTable.getName() );  
        // Print out delete result set.  
        println( "Deletes" );  
        printRSInfo( currentTable.getDeletes() );  
        // Print out insert result set.  
        println( "Inserts" );  
        printRSInfo( currentTable.getInserts() );  
        // print out update result set  
        println( "Updates" );  
        printUpdateRSInfo( currentTable.getUpdates() );  
    }  
}
```

The printRSInfo method prints out an insert, update, or delete result set and accepts a single java.sql.ResultSet object. Detailed column information, including column labels, is provided by the ResultSetMetaData object returned by the ResultSet getMetaData method. Column labels are available only if the client has the SendColumnNames option turned on. The printRow method prints out each row in a result set.

```
public void printRSInfo( ResultSet results )  
    throws SQLException, IOException {  
  
    // Obtain the result set metadata.  
    ResultSetMetaData metaData = results.getMetaData();  
    String columnHeading = "";  
  
    // Print out column headings.  
    for( int c = 1; c <= metaData.getColumnCount(); c++ ) {  
        columnHeading += metaData.getColumnLabel(c);  
        if( c < metaData.getColumnCount() ) {  
            columnHeading += ", ";  
        }  
    }  
  
    println( columnHeading );  
    while( results.next() ) {  
  
        // Print out each row.  
        printRow( results, metaData.getColumnCount() );  
    }  
}
```

```

    }

    // Close the java.sql.ResultSet.
    results.close();
}

```

The printRow method, shown below, uses the ResultSet getString method to obtain each column value.

```

public void printRow( ResultSet results, int colCount )
    throws SQLException, IOException {
    String row = "( ";

    for( int c = 1; c <= colCount; c++ ) {
        // Get a column value.
        String currentColumn = results.getString( c );

        // Check for null values.
        if( currentColumn == null ) {
            currentColumn = "<NULL>";
        }

        // Add the column value to the row string.
        row += cur_col;
        if( c < colCount ) {
            row += ", ";
        }
    }

    row += " )";

    // Print out the row.
    println( row );
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called HandleUpload for the handle\_UploadData connection event when synchronizing the script version ver1. You run this system procedure against your MobiLink consolidated database.

```

CALL ml_add_dnet_connection_script(
    'ver1',
    'handle_UploadData',
    'TestScripts.Test.HandleUpload' )

```

The following .NET method processes the upload for the remoteOrders table.

```

using System;
using System.Data;
using System.IO;
using iAnywhere.MobiLink.Script;
using iAnywhere.MobiLink;

namespace MyScripts
{
    public class MyUpload
    {
        public MyUpload( DBConnectionContext cc )
        {
        }

        ~MyUpload()
    }
}

```

```
{
}

public void handleUpload( UploadData ut )
{
    int i;
    UploadedTableData[] tables = ut.GetUploadedTables();

    for( i=0; i<tables.Length; i+=1 ) {
        UploadedTableData cur_table = tables[i];
        Console.WriteLine( "table " + i + " name: " + cur_table.GetName() );
        // Print out delete result set.
        Console.WriteLine( "Deletes" );
        printRSInfo( cur_table.GetDeletes() );

        // Print out insert result set.
        Console.WriteLine( "Inserts" );
        printRSInfo( cur_table.GetInserts() );

        // print out update result set
        Console.WriteLine( "Updates" );
        printUpdateRSInfo( cur_table.GetUpdates() );
    }
}

public void printRSInfo( IDataReader dr )
{
    // Obtain the result set metadata.
    DataTable dt = dr.GetSchemaTable();
    DataColumnCollection cc = dt.Columns;
    DataColumn dc;
    String columnHeading = "";

    // Print out column headings.
    for( int c=0; c < cc.Count; c = c + 1 ) {
        dc = cc[ c ];
        columnHeading += dc.ColumnName;
        if( c < cc.Count - 1 ) {
            columnHeading += ", ";
        }
    }
    Console.WriteLine( columnHeading );

    while( dr.Read() ) {
        // Print out each row.
        printRow( dr, cc.Count );
    }

    // Close the java.sql.ResultSet.
    dr.Close();
}

public void printUpdateRSInfo( UpdateDataReader utr )
{
    // Obtain the result set metadata.
    DataTable dt = utr.GetSchemaTable();
    DataColumnCollection cc = dt.Columns;
    DataColumn dc;
    String columnHeading = "TYPE, ";

    // Print out column headings.
    for( int c = 0; c < cc.Count; c = c + 1 ) {
        dc = cc[ c ];
        columnHeading += dc.ColumnName;
    }
}
```

```
        if( c < cc.Count - 1 ) {
            columnHeading += ", ";
        }
    }
    Console.Write( columnHeading );

    while( utr.Read() ) {
        // Print out the new values for the row.
        utr.SetNewRowValues();
        Console.Write( "NEW:" );
        printRow( utr, cc.Count );

        // Print out the old values for the row.
        utr.SetOldRowValues();
        Console.Write( "OLD:" );
        printRow( utr, cc.Count );
    }

    // Close the java.sql.ResultSet.
    utr.Close();
}

public void printRow( IDataReader dr, int col_count )
{
    String row = "( ";
    int c;

    for( c = 0; c < col_count; c = c + 1 ) {
        // Get a column value.
        String cur_col = dr.GetString( c );

        // Check for null values.
        if( cur_col == null ) {
            cur_col = "<NULL>";
        }

        // Add the column value to the row string.
        row += cur_col;
        if( c < col_count ) {
            row += ", ";
        }
    }

    row += " )";

    // Print out the row.
    Console.Write( row );
}
}
```

## modify\_error\_message connection event

The script can be used to customize the message text (error, warning, and information) that is sent to remote databases.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.error_message	VARBINARY(1024). This is an IN-OUT parameter.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.error_code	INT.	3

### Default action

None.

### Remarks

SQL scripts for the modify\_error\_message event must be implemented as stored procedures.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The following example downloads everything from one day ago, regardless of whether the databases were synchronized since then.

The following SQL statement creates the ModifyLastErrorMessage stored procedure:



```

CREATE PROCEDURE ModifyLastErrorMessage(
    inout error_message VARBINARY(1024),
    in username VARCHAR(128),
    in error_code INT )
BEGIN
    SELECT dateadd(day, -1, last_download_time )
    INTO last_download_time
END
    
```

The following call to a MobiLink system procedure assigns ModifyLastErrorMessage to the modify\_error\_message connection event for the script version modify\_ts\_test:

```

CALL ml_add_connection_script(
    'modify_ts_test',
    'modify_error_message',
    'CALL ModifyLastErrorMessage (
        {ml s.error_message},
        {ml s.username},
        {ml s.error_code} )' );
    
```

### Java example

The following call to a MobiLink system procedure registers a Java method called modifyLastErrorMessage as the script for the modify\_error\_message connection event when synchronizing the script version ver1.

```

CALL ml_add_java_connection_script(
    'ver1',
    'modify_error_message',
    'ExamplePackage.ExampleClass.modifyLastErrorMessage' )
    
```

The following is the sample Java method modifyLastErrorMessage. It prints the current error message and error code.

```

public String modifyLastErrorMessage(
    String lastErrorMessage,
    String userName,
    int errorCode ) {
    java.lang.System.out.println( "error message: " +
        lastErrorMessage );
    java.lang.System.out.println( "error code: " +
        String.valueOf(errorCode) );
    return( null );
}
    
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called ModifyLastErrorMessage as the script for the modify\_error\_message connection event when synchronizing the script version ver1.

```

CALL ml_add_dnet_connection_script(
    'ver1',
    'modify_error_message',
    'TestScripts.Test.ModifyLastErrorMessage' )
    
```

The following is a sample .NET method ModifyLastErrorMessage. It prints the current error code and error message.

```

public string ModifyLastErrorMessage (
    string errorMessage,
    
```

```
string userName,  
string errorCode ) {  
System.Console.WriteLine( "error message: " + errorMessage );  
System.Console.WriteLine( "error code: " + errorCode );  
return ( null );  
}
```

## modify\_last\_download\_timestamp connection event

The script can be used to modify the last download time for the current synchronization.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_download	TIMESTAMP. The last download time for any synchronized table. This is an INOUT parameter.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

Use this script when you want to modify the last\_download timestamp for the current synchronization. If this script is defined, the MobiLink server uses the modified last\_download timestamp as the last\_download timestamp passed to the download scripts. A typical use of this script is to recover from losing data on the remote; you can reset the last\_download timestamp to an early time such as 1900-01-01 00:00 so that the next synchronization downloads all the data.

SQL scripts for the modify\_last\_download\_timestamp event must be implemented as stored procedures. The MobiLink server passes in the last\_download timestamp as the first parameter to the stored procedure, and replaces the timestamp by the first value passed out by the stored procedure.

This script is executed just before the prepare\_for\_download script, in the same transaction.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]
- “Using last download times in scripts” on page 130
- “How download timestamps are generated and used” on page 131
- “modify\_next\_last\_download\_timestamp connection event” on page 466

### SQL example

The following SQL statement creates a stored procedure. The following syntax is for Oracle consolidated databases:

```
CREATE PROCEDURE ModifyDownloadTimestamp(  
    download_timestamp OUT DATE,  
    user_name IN VARCHAR )  
AS  
BEGIN  
    -- N is the maximum replication latency in the consolidated cluster  
    download_timestamp := download_timestamp - N;  
END;
```

The following syntax is for SQL Anywhere, Adaptive Server Enterprise, and SQL Server consolidated databases:

```
CREATE PROCEDURE ModifyDownloadTimestamp  
    @download_timestamp DATETIME OUTPUT,  
    @t_name VARCHAR( 128 )  
AS  
BEGIN  
    -- N is the maximum replication latency in consolidated cluster  
    SELECT @download_timestamp = @download_timestamp - N  
END
```

The following syntax is for DB2 mainframe consolidated databases:

```
CREATE PROCEDURE modify_ldts(  
    OUT ts          TIMESTAMP,  
    IN t_name      VARCHAR(128) )  
LANGUAGE SQL  
BEGIN  
    set ts = TIMESTAMP_FORMAT('2000-01-02 03:04:05', 'YYYY-MM-DD  
HH24:MI:SS');  
END
```

The following call to a MobiLink system procedure assigns the ModifyDownloadTimestamp stored procedure to the modify\_last\_download\_timestamp event. The following syntax is for a SQL Anywhere consolidated database:

```
CALL ml_add_connection_script(  
    'my_version',  
    'modify_last_download_timestamp',  
    '{CALL ModifyDownloadTimestamp(  
    {ml s.last_download},  
    {ml s.username} ) }' )
```

## Java example

The following call to a MobiLink system procedure registers a Java method called `modifyLastDownloadTimestamp` as the script for the `modify_last_download_timestamp` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'modify_last_download_timestamp',  
  'ExamplePackage.ExampleClass.modifyLastDownloadTimestamp' )
```

The following is the sample Java method `modifyLastDownloadTimestamp`. It prints the current and new timestamp and modifies the timestamp that is passed in.

```
public String modifyLastDownloadTimestamp(  
    Timestamp lastDownloadTime,  
    String userName ) {  
    java.lang.System.out.println( "old date: " +  
        lastDownloadTime.toString() );  
    lastDownloadTime.setDate(  
        lastDownloadTime.getDate() -1 );  
    java.lang.System.out.println( "new date: " +  
        lastDownloadTime.toString() );  
    return( null );  
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called `ModifyLastDownloadTimestamp` as the script for the `modify_last_download_timestamp` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'modify_last_download_timestamp',  
  'TestScripts.Test.ModifyLastDownloadTimestamp' )
```

The following is the sample .NET method `ModifyLastDownloadTimestamp`.

```
public string ModifyLastDownloadTimestamp(  
    DateTime lastDownloadTime,  
    string userName ) {  
    System.Console.WriteLine( "old date: " +  
        last_download_time.ToString() );  
    last_download_time = DateTime::Now;  
    System.Console.WriteLine( "new date: " +  
        last_download_time.ToString() );  
    return( null );  
}
```

## modify\_next\_last\_download\_timestamp connection event

The script can be used to modify the last download time for the next synchronization.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.next_last_download	TIMESTAMP. This is an INOUT parameter. The MobiLink server generates this value immediately after the upload is committed.	1
s.last_download	TIMESTAMP. This is an IN parameter. This is the last download time for the current synchronization. It can be useful in avoiding duplication, by making sure you don't modify the next_last_download timestamp to be earlier than the last_download timestamp.	2
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	3

### Default action

None.

## Remarks

This script allows you to change the next\_last\_download timestamp, which effectively changes the last\_download timestamp for the next synchronization. This allows you to reset the next synchronization without affecting the current synchronization.

SQL scripts for the modify\_next\_last\_download\_timestamp event must be implemented as stored procedures. The MobiLink server passes in the next\_last\_download timestamp as the first parameter to the stored procedure, and replaces the timestamp by the first value passed out by the stored procedure.

This script is executed in the download transaction, after downloading user tables.

## See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)
- [“How download timestamps are generated and used” on page 131](#)
- [“modify\\_last\\_download\\_timestamp connection event” on page 463](#)

## SQL example

The following example shows one application of this script. Create a stored procedure. The following syntax is for a SQL Anywhere consolidated database:

```
CREATE PROCEDURE ModifyNextDownloadTimestamp(
    inout download_timestamp TIMESTAMP ,
    in last_download TIMESTAMP ,
    in user_name VARCHAR(128) )
BEGIN
    SELECT dateadd(hour, -1, download_timestamp )
    INTO download_timestamp
END
```

Install the script into your SQL Anywhere consolidated database:

```
CALL ml_add_connection_script(
    'modify_ts_test',
    'modify_next_last_download_timestamp',
    'CALL ModifyNextDownloadTimestamp (
        {ml s.next_last_download},
        {ml s.last_download},
        {ml s.username} )' )
```

## Java example

The following call to a MobiLink system procedure registers a Java method called modifyNextDownloadTimestamp as the script for the modify\_next\_last\_download\_timestamp connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
    'ver1',
    'modify_next_last_download_timestamp',
    'ExamplePackage.ExampleClass.modifyNextDownloadTimestamp' )
```

The following is the sample Java method modifyNextDownloadTimestamp. It sets the download timestamp back an hour.

```
public String modifyNextDownloadTimestamp(
    Timestamp downloadTimestamp,
    Timestamp lastDownload,
    String userName ) {
    downloadTimestamp.setHours(
        downloadTimestamp.getHours() -1 );
    return( null );
}
```

### **.NET example**

The following call to a MobiLink system procedure registers a .NET method called `ModifyNextDownloadTimestamp` as the script for the `modify_next_last_download_timestamp` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(
    'ver1',
    'modify_next_last_download_timestamp',
    'TestScripts.Test.ModifyNextDownloadTimestamp' )
```

The following is the sample .NET method `ModifyNextDownloadTimestamp`. It sets the download timestamp back an hour.

```
public string ModifyNextDownloadTimestamp (
    DateTime download_timestamp,
    DateTime last_download,
    string user_name ) {
    download_timestamp = download_timestamp.AddHours( -1 );
    return ( null );
}
```



## modify\_user connection event

Provide the MobiLink user name.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name. This is an INOUT parameter.	1

### Default action

None.

### Remarks

The MobiLink server provides the user name as a parameter when it calls scripts; the user name is sent by the MobiLink client. In some cases, you may want to have an alternate user name. This script allows you to modify the user name used in calling MobiLink scripts.

The username parameter must be long enough to hold the user name.

SQL scripts for the modify\_user event must be implemented as stored procedures.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“authenticate\\_user connection event” on page 358](#)
- [“authenticate\\_user\\_hashed connection event” on page 363](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The following example maps a remote database user name to the ID of the user using the device, by using a mapping table called user\_device. This technique can be used when the same person has multiple remotes (such as a PDA and a laptop) requiring the same synchronization logic (based on the user's name or id).

The following call to a MobiLink system procedure assigns the ModifyUser stored procedure to the modify\_user event. This syntax is for a SQL Anywhere consolidated database.

```
CALL ml_add_connection_script(
  'ver1',
```

```
'modify_user',  
'call ModifyUser( {ml s.username} )' )
```

The following SQL statement creates the ModifyUser stored procedure.

```
CREATE PROCEDURE ModifyUser( INOUT u_name varchar(128) )  
BEGIN  
  SELECT user_name  
  INTO u_name  
  FROM user_device  
  WHERE device_name = u_name;  
END
```

### Java example

The following call to a MobiLink system procedure registers a Java method called modifyUser as the script for the modify\_user connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(  
  'ver1',  
  'modify_user',  
  'ExamplePackage.ExampleClass.modifyUser' )
```

The following is the sample Java method modifyUser. It gets the user ID from the database and then uses it to set the user name.

```
public String modifyUser(  
  InOutString ioUserName )  
  throws SQLException {  
  Statement uidSelect = curConn.createStatement();  
  ResultSet uidResult = uidSelect.executeQuery(  
    "SELECT rep_id FROM SalesRep WHERE name = '" +  
    ioUserName.getValue() + "' " );  
  uidResult.next();  
  ioUserName.setValue(  
    java.lang.Integer.toString(uidResult.getInt( 1 ));  
  uidResult.close();  
  uidSelect.close();  
  return ( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called ModUser as the script for the modify\_user connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
  'ver1',  
  'modify_user',  
  'TestScripts.Test.ModUser'  
)
```

The following is the sample .NET method ModUser.

```
public string ModUser(  
  string user ) {  
  return ( "SELECT rep_id FROM SalesRep WHERE name = '" + user + "' " );  
}
```

## nonblocking\_download\_ack connection event

When you use non-blocking download acknowledgement, this script provides a place to record the information that a download has been applied successfully.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.last_download	TIMESTAMP. This is an IN parameter. This is the last download time for the current synchronization. It can be useful in avoiding duplication, by making sure you don't modify the next_last_download timestamp to be earlier than the last_download timestamp.	3

### Remarks

This event lets you record the time when the download was successfully applied at the remote database.

This event is only called when using non-blocking download acknowledgement. When in non-blocking mode, the download transaction is committed and the synchronization ends when the download is sent. This event is called when the synchronization client acknowledges a successful download. This event is called on a new connection, after the end\_synchronization script of the original synchronization. The actions of this event are committed along with an update to the download time in the MobiLink system tables.

Due to the special nature of this script, any connection-level variables set during the synchronization are not available when this event is executed.

### See also

- [“publication\\_nonblocking\\_download\\_ack connection event” on page 475](#)
- [“-nba option” on page 74](#)
- dbmsync: [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Download Acknowledgement synchronization parameter” \[UltraLite - Database Management and Reference\]](#)

### SQL example

The following script adds a record to the table `download_pubs_acked`. The record contains the remote ID, the first authentication parameter, and the download timestamp.

```
INSERT INTO download_pubs_acked( rem_id, auth_parm, last_download )
VALUES( {ml s.remote_id}, {ml a.1}, {ml s.last_publication_download} )
```

## prepare\_for\_download connection event

Processes any required operations between the upload and download transactions.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.last_download	TIMESTAMP. The last download time of any synchronized table.	1
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2

### Default action

None.

### Remarks

The MobiLink server executes this script as a separate transaction, between the upload transaction and the start of the download transaction.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“end\\_upload connection event” on page 431](#)
- [“begin\\_download connection event” on page 369](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)
- [“Using last download times in scripts” on page 130](#)

### SQL example

The following call to a MobiLink system procedure registers a SQL method called prepareForDownload as the script for the prepare\_for\_download event when synchronizing the script version ver1.

```
CALL ml_add_connection_script(
  'ver1',
  'prepare_for_download',
  'CALL prepareForDownload(
    { ml s.current_time },
    { ml s.username } )' )
```

The following is the sample SQL method prepareForDownload. It calls a SQL method to modify some rows in the database.

```
CREATE PROCEDURE prepareForDownload (
  IN ts TIMESTAMP,
  IN user VARCHAR(128))
BEGIN
  CALL adjustUploadedRows(user)
END;
```

### Java example

The following call to a MobiLink system procedure registers a Java method called prepareForDownload as the script for the prepare\_for\_download event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'prepare_for_download',
  'ExamplePackage.ExampleClass.prepareForDownload' )
```

The following is the sample Java method prepareForDownload. It calls a Java method to modify some rows in the database.

```
public String prepareForDownload(
  Timestamp ts,
  String user ) {
  adjustUploadedRows( _syncConn, user );
  return( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called PrepareForDownload as the script for the prepare\_for\_download connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(
  'ver1',
  'prepare_for_download',
  'TestScripts.Test.PrepareForDownload'
)
```

The following is the sample .NET method PrepareForDownload. It calls a .NET method to modify some rows in the database.

```
public string PrepareForDownload(
  DateTime timestamp,
  string user ) {
  AdjustUploadedRows ( _syncConn, user );
  return ( null );
}
```

## publication\_nonblocking\_download\_ack connection event

When you use non-blocking download acknowledgement, this script provides a place to record the information that a publication has been successfully downloaded.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	2
s.last_publication_download	TIMESTAMP. The last download time of any synchronized table.	3
s.publication name	VARCHAR(128). The name of the publication.	4
s.subscription_id	VARCHAR(128). The publication ID.	5

### Remarks

This event lets you record the time when the download of this publication was successfully applied at the remote database.

This event is only called when using non-blocking download acknowledgement. When in non-blocking mode, the download transaction is committed and the synchronization ends when the download is sent. When the synchronization client acknowledges a successful download, this event is called once per publication in the download. This event is called on a new connection and after the end\_synchronization script of the original synchronization. The actions of this event are committed along with an update to the download time in the MobiLink system tables.

Due to the special nature of this script, any connection-level variables set during the synchronization are not available when this event is executed.

### See also

- [“nonblocking\\_download\\_ack connection event” on page 471](#)
- [“-nba option” on page 74](#)
- dbmsync: [“SendDownloadACK \(sa\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Download Acknowledgement synchronization parameter” \[UltraLite - Database Management and Reference\]](#)

### SQL example

The following script adds a record to a table called `download_pubs_acked`. The record contains the publication name, the first authentication parameter, and a download timestamp.

```
INSERT INTO download_pubs_acked( pub_name, auth_parm, last_download )
VALUES( {ml s.publication_name}, {ml a.1}, {ml
s.last_publication_download} )
```



## report\_error connection event

Allows you to log errors and to record the actions selected by the handle\_error script.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.action_code	INTEGER. This is an INOUT parameter. This parameter is mandatory.	1
s.error_code	INTEGER.	2
s.error_message	TEXT.	3
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	4
s.table	VARCHAR(128).	5

### Default action

None.

### Remarks

This script allows you to log errors and to record the actions selected by the handle\_error script. This script is executed after the handle\_error event, whether a handle\_error script is defined. It is always executed in its own transaction, on a different database connection than the synchronization connection (the administrative/information connection).

The error code and error message allow you to identify the nature of the error. The action code value is returned by the last call to an error handling script for the SQL operation that caused the current error.

If the error happened as part of synchronization, the user name is supplied. Otherwise, this value is null.

If the error happened while manipulating a particular table, the table name is supplied. Otherwise, this value is null. The table name is the name of a table in the remote database. This name may or may not have a direct counterpart in the consolidated database, depending on the design of the synchronization system.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“handle\\_error connection event” on page 446](#)
- [“handle\\_odbc\\_error connection event” on page 450](#)
- [“report\\_odbc\\_error connection event” on page 480](#)
- [“Using remote IDs and MobiLink user names in scripts” \[MobiLink - Client Administration\]](#)

### SQL example

The following example works with a SQL Anywhere consolidated database. It inserts a row into a table used to record synchronization errors.

```
CALL ml_add_connection_script(
  'ver1',
  'report_error',
  'INSERT INTO sync_error(
    action_code,
    error_code,
    error_message,
    user_name,
    table_name )
VALUES (
  {ml s.action_code},
  {ml s.error_code},
  {ml s.error_message},
  {ml s.username},
  {ml s.table} )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called reportError as the script for the report\_error connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'report_error',
  'ExamplePackage.ExampleClass.reportError' )
```

The following is the sample Java method reportError. It logs the error to a table using the JDBC connection provided by MobiLink. It also sets the action code.

```
public String reportError(
  anywhere.ml.script.InOutInteger actionCode,
  int errorCode,
  String errorMessage,
  String user,
  String table )
throws java.sql.SQLException {
  // Insert error information in a table,
  JDBCLogError( _syncConn, errorCode, errorMessage,
    user, table );
  actionCode.setValue( getActionCode( errorCode ) );
}
```

```
    return( null );  
}
```

### **.NET example**

The following call to a MobiLink system procedure registers a .NET method called ReportError as the script for the report\_error connection event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'report_error',  
    'TestScripts.Test.ReportError' )
```

The following is the sample .NET method ReportError. It logs the error to a table using a .NET method.

```
public string ReportError(  
    ref int actionCode,  
    int errorCode,  
    string errorMessage,  
    string user,  
    string table ) {  
    LogError(_syncConn, errorCode, errorMessage, user, table);  
}
```

## report\_odbc\_error\_connection\_event

Allows you to log errors and to record the actions selected by the handle\_odbc\_error script.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.action_code	INTEGER. This is an INOUT parameter. This parameter is mandatory.	1
s.ODBC_state	VARCHAR(5).	2
s.error_message	TEXT.	3
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	4
s.table	VARCHAR(128).	5

### Default action

None.

### Remarks

This script allows you to log errors and to record the actions selected by the handle\_odbc\_error script. This script is executed after the handle\_odbc\_error event, whether a handle\_odbc\_error script is defined. It is always executed in its own transaction, on a different database connection than the synchronization connection (the administrative/information connection).

The error code and error message allow you to identify the nature of the error. The action code value is returned by the last call to an error handling script for the SQL operation that caused the current error.

If the error happened as part of synchronization, the user name is supplied. Otherwise, this value is null.

If the error happened while manipulating a particular table, the table name is supplied. Otherwise, this value is null. The table name is the name of a table in the remote database. This name may or may not have a direct counterpart in the consolidated database, depending on the design of the synchronization system.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “handle\_error connection event” on page 446
- “handle\_odbc\_error connection event” on page 450
- “report\_error connection event” on page 477
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

### SQL example

The following example works with a SQL Anywhere consolidated database. It inserts a row into a table used to record synchronization errors.

```
CALL ml_add_connection_script(
  'ver1',
  'report_odbc_error',
  'INSERT INTO sync_error(
    action_code,
    odbc_state,
    error_message,
    user_name,
    table_name )
VALUES(
  {ml s.action_code},
  {ml s.ODBC_state},
  {ml s.error_message},
  {ml s.username},
  {ml s.table} )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called reportODBCError as the script for the report\_odbc\_error event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'report_odbc_error',
  'ExamplePackage.ExampleClass.reportODBCError' )
```

The following is the sample Java method reportODBCError. It logs the error to a table using the JDBC connection provided by MobiLink. It also sets the action code.

```
public String reportODBCError(
  ianywhere.ml.script.InOutInteger actionCode,
  String ODBCState,
  String errorMessage,
  String user,
  String table )
throws java.sql.SQLException {
  JDBCLogError( _syncConn, ODBCState, errorMessage,
    user, table );
  actionCode.setValue( getActionCode( ODBCState ) );
  return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called ReportODBCError as the script for the report\_odbc\_error event when synchronizing the script version ver1.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'report_odbc_error',  
    'TestScripts.Test.ReportODBCError' )
```

The following is the sample .NET method ReportODBCError. It logs the error to a table using a .NET method.

```
public string ReportODBCError (  
    ref int actionCode,  
    string ODBCState,  
    string errorMessage,  
    string user,  
    string table ) {  
    LogError(_syncConn, ODBCState, errorMessage, user, table);  
    return ( null );  
}
```

## resolve\_conflict table event

Defines a process for resolving a conflict in a specific table.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2

### Default action

None.

### Remarks

When a row is updated on a remote database, the MobiLink client saves a copy of the original values. The client sends both old and new values to the MobiLink server.

When the MobiLink server receives an updated row, it compares the original values with the present values in the consolidated database. The comparison is done using the `upload_fetch` script.

If the old uploaded values do not match the current values in the consolidated database, the row conflicts. Instead of updating the row, the MobiLink server inserts both old and new values into the consolidated database. The old and new rows are handled using the `upload_old_row_insert` and `upload_new_row_insert` scripts, respectively.

Once the values have been inserted, the MobiLink server executes the `resolve_conflict` script. It provides the opportunity to resolve the conflict. You can implement any scheme of your choosing.

This script is executed once per conflict.

Alternatively, instead of defining the `resolve_conflict` script, you can resolve conflicts in a set-oriented fashion by putting conflict-resolution logic either in your `end_upload_rows` script or in your `end_upload` table script.

You can have one `resolve_conflict` script for each table in the remote database.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “upload\_old\_row\_insert table event” on page 509
- “upload\_new\_row\_insert table event” on page 506
- “upload\_update table event” on page 522
- “end\_upload\_rows table event” on page 439
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

### SQL example

The following statement defines a `resolve_conflict` script suited to the CustDB sample application for an Oracle installation. It calls a stored procedure `ULResolveOrderConflict`.

```
exec ml_add_table_script(
  'custdb', 'ULOrder', 'resolve_conflict',
  'begin ULResolveOrderConflict();
end; ')
CREATE OR REPLACE PROCEDURE ULResolveOrderConflict()
AS
  new_order_id integer;
  new_status   varchar(20);
  new_notes    varchar(50);
BEGIN
  -- approval overrides denial
  SELECT order_id, status, notes
    INTO new_order_id, new_status, new_notes
    FROM ULNewOrder
   WHERE syncuser_id = SyncUserID;
  IF new_status = 'Approved' THEN
    UPDATE ULOrder o
      SET o.status = new_status, o.notes =
        new_notes
    WHERE o.order_id = new_order_id;
  END IF;
  DELETE FROM ULOldOrder;
  DELETE FROM ULNewOrder;
END;
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `resolveConflict` as the script for the `resolve_conflict` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'resolve_conflict',
  'ExamplePackage.ExampleClass.resolveConflict' )
```

The following is the sample Java method `resolveConflict`. It calls a Java method that uses the JDBC connection provided by MobiLink to resolve the conflict.



```
public String resolveConflict(  
    String user,  
    String table) {  
    resolveRows(_syncConn, user );  
}
```

### **.NET example**

The following call to a MobiLink system procedure registers a .NET method called ResolveConflict as the script for the resolve\_conflict table event when synchronizing the script version ver1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'resolve_conflict',  
    'TestScripts.Test.ResolveConflict' )
```

The following is the sample .NET method ResolveConflict. It calls a .NET method that resolves the conflict.

```
public string ResolveConflict(  
    String user,  
    String table) {  
    ResolveRows(_syncConn, user );  
}
```

## synchronization\_statistics connection event

Tracks synchronization statistics.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.warnings	INTEGER. The number of warnings issued during the synchronization.	2
s.errors	INTEGER. The number of errors that occurred during the synchronization.	3
s.deadlocks	INTEGER. The number of deadlocks in the consolidated database that were detected for the synchronization.	4
s.synchronized_tables	INTEGER. The number of client tables that were involved in the synchronization.	5
s.connection_retries	INTEGER. The number of times the MobiLink server retried the connection to the consolidated database.	6

**Default action**

None.

**Remarks**

The synchronization\_statistics event allows you to gather, for any user and connection, various statistics about the current synchronization. The synchronization\_statistics connection script is called just prior to the commit at the end of the end synchronization transaction.

**See also**

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“download\\_statistics connection event” on page 403](#)
- [“download\\_statistics table event” on page 406](#)
- [“upload\\_statistics connection event” on page 512](#)
- [“upload\\_statistics table event” on page 517](#)
- [“synchronization\\_statistics table event” on page 489](#)
- [“time\\_statistics connection event” on page 492](#)
- [“time\\_statistics table event” on page 495](#)
- [“MobiLink Monitor” on page 179](#)
- [“Using remote IDs and MobiLink user names in scripts” \[\*MobiLink - Client Administration\*\]](#)

**SQL example**

The following example inserts synchronization statistics into the sync\_con\_audit table.

```
CALL ml_add_connection_script(
  'ver1',
  'synchronization_statistics',
  'INSERT INTO sync_con_audit(
    ml_user,
    warnings,
    errors,
    deadlocks,
    synchronized_tables,
    connection_retries)
VALUES (
  {ml s.username},
  {ml s.warnings},
  {ml s.errors},
  {ml s.deadlocks},
  {ml s.synchronized_tables},
  {ml s.connection_retries})' )
```

Once statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

**Java example**

The following call to a MobiLink system procedure registers a Java method called synchronizationStatisticsConnection as the script for the synchronization\_statistics connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
```

```
'synchronization_statistics',  
'ExamplePackage.ExampleClass.synchronizationStatisticsConnection'  
)
```

The following is the sample Java method `synchronizationStatisticsConnection`. It logs some of the statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String synchronizationStatisticsConnection(  
    String user,  
    int warnings,  
    int errors,  
    int deadlocks,  
    int synchronizedTables,  
    int connectionRetries ) {  
    java.lang.System.out.println(  
        "synch statistics number of deadlocks: "  
        + deadlocks ;  
    return( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `SyncStats` as the script for the `synchronization_statistics` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(  
    'ver1',  
    'synchronization_statistics',  
    'TestScripts.Test.SyncStats'  
)
```

The following is the sample .NET method `SyncStats`. It logs some of the statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string SyncStats(  
    string user,  
    int warnings,  
    int errors,  
    int deadLocks,  
    int syncedTables,  
    int connRetries ) {  
    System.Console.WriteLine( "synch statistics  
    number of deadlocks: " + deadlocks ;  
    return( null );  
}
```

## synchronization\_statistics table event

Tracks synchronization statistics.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2
s.warnings	INTEGER. The number of warnings that occurred for the table during the synchronization.	3
s.errors	INTEGER. The number of errors that were related to the table during the synchronization.	4

### Default action

None.

### Remarks

The synchronization\_statistics event allows you to gather, for any user and table, the number of warnings and errors that occurred during synchronization. The synchronization\_statistics table script is called just prior to the commit at the end of the end synchronization transaction.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “download\_statistics connection event” on page 403
- “download\_statistics table event” on page 406
- “upload\_statistics connection event” on page 512
- “upload\_statistics table event” on page 517
- “synchronization\_statistics connection event” on page 486
- “time\_statistics connection event” on page 492
- “time\_statistics table event” on page 495
- “MobiLink Monitor” on page 179
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

### SQL example

The following example inserts synchronization statistics into the sync\_tab\_audit table.

```
CALL ml_add_table_script(  
  'ver1',  
  'table1',  
  'INSERT INTO sync_tab_audit (  
    ml_user,  
    table,  
    warnings,  
    errors)  
VALUES (  
  {ml s.username},  
  {ml s.table},  
  {ml s.warnings},  
  {ml s.errors} ) ' )
```

Once synchronization statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

### Java example

The following call to a MobiLink system procedure registers a Java method called synchronizationStatisticsTable as the script for the synchronization\_statistics table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
  'ver1',  
  'table1',  
  'synchronization_statistics',  
  'ExamplePackage.ExampleClass.synchronizationStatisticsTable'  
)
```

The following is the sample Java method synchronizationStatisticsTable. It logs some of the statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String synchronizationStatisticsTable(  
  String user,  
  String table,  
  int warnings,  
  int errors ) {  
  java.lang.System.out.println( "synch statistics for
```

```
        table: " + table + " errors: " + errors );  
    return( null );  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called SyncTableStats as the script for the synchronization\_statistics table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'synchronization_statistics',  
    'TestScripts.Test.SyncTableStats'  
)
```

The following is the sample .NET method SyncTableStats. It logs some of the statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string SyncTableStats(  
    string user,  
    string table,  
    int warnings,  
    int errors ) {  
    System.Console.WriteLine( "synch statistics for  
        table: " + table + " errors: " + errors );  
    return( null );  
}
```

## time\_statistics connection event

Tracks time statistics by user and event.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.event_name	VARCHAR(128)	2
s.num_of_calls	INTEGER. The number of times the script was called.	3
s.minimum_time	INTEGER. Milliseconds. The shortest time it took to execute a script during this synchronization.	4
s.maximum_time	INTEGER. Milliseconds. The longest time it took to execute a script during this synchronization.	5
s.total_time	INTEGER. Milliseconds. The total time it took to execute all scripts in the synchronization. (This is not the same as the length of the synchronization.)	6

### Default action

None.



## Remarks

The time\_statistics event allows you to gather time statistics for any user during synchronization. The statistics are gathered only for those events for which there is a corresponding script. The script gathers aggregate data for occasions where a single event occurs multiple times. The script can be especially useful for time comparisons across users, events and tables.

## See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“time\\_statistics table event” on page 495](#)
- [“download\\_statistics connection event” on page 403](#)
- [“download\\_statistics table event” on page 406](#)
- [“upload\\_statistics connection event” on page 512](#)
- [“upload\\_statistics table event” on page 517](#)
- [“synchronization\\_statistics connection event” on page 486](#)
- [“synchronization\\_statistics table event” on page 489](#)
- [“MobiLink Monitor” on page 179](#)
- [“Using remote IDs and MobiLink user names in scripts” \[\*MobiLink - Client Administration\*\]](#)

## SQL example

The following example inserts statistical information into the time\_statistics table.

```
CALL ml_add_connection_script(
  'ver1',
  'time_statistics',
  'INSERT INTO time_statistics (
    id,
    ml_user,
    table,
    event_name,
    number_of_calls,
    minimum_time,
    maximum_time,
    total_time)
VALUES (
  ts_id.nextval,
  {ml s.username},
  {ml s.event_name},
  {ml s.number_of_calls},
  {ml s.minimum_time},
  {ml s.maximum_time},
  {ml s.total_time} ) ' )
```

## Java example

The following call to a MobiLink system procedure registers a Java method called timeStatisticsConnection as the script for the time\_statistics connection event when synchronizing the script version ver1.

```
CALL ml_add_java_connection_script(
  'ver1',
  'time_statistics',
  'ExamplePackage.ExampleClass.timeStatisticsConnection' )
```

The following is the sample Java method `timeStatisticsConnection`. It prints statistics for the `prepare_for_download` event. (Note that printing statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String timeStatisticsConnection(
    String username,
    String tableName,
    String eventName,
    int numberOfCalls,
    int minimumTime,
    int maximumTime,
    int totalTime ) {
    if( eventName.equals( "prepare_for_download" ) ) {
        java.lang.System.out.println(
            "prepare_for_download num_calls: " + numCalls +
            "total_time: " + totalTime );
    }
    return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `TimeStats` as the script for the `time_statistics` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(
    'ver1',
    'time_statistics',
    'TestScripts.Test.TimeStats'
)
```

The following is the sample .NET method `TimeStats`. It prints statistics for the `prepare_for_download` event. (Note that printing statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string TimeStats(
    string user,
    string eventName,
    int numberOfCalls,
    int minimumTime,
    int maximumTime,
    int totTime ) {
    if( event_name=="prepare_for_download" ) {
        System.Console.WriteLine(
            "prepare_for_download num_calls: " + num_calls +
            "total_time: " + total_time );
    }
    return ( null );
}
```

## time\_statistics table event

Tracks time statistics.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2
s.event_name	VARCHAR(128)	3
s.number_of_calls	INTEGER. The number of times the script was called.	4
s.minimum_time	INTEGER. Milliseconds. The shortest time it took to execute a script during the synchronization of this table.	5
s.maximum_time	INTEGER. Milliseconds. The longest time it took to execute a script during the synchronization of this table.	6
s.total_time	INTEGER. Milliseconds. The total time it took to execute all scripts in the synchronization of the table. (This is not the same as the length of the synchronization.)	7

### Default action

None.

### Remarks

The `time_statistics` table event allows you to gather time statistics for any user and table during synchronization. The statistics are gathered only for those events for which there is a corresponding script. The script gathers aggregate data for occasions where a single event occurs multiple times. The script can be especially useful for time comparisons across users, events and tables.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“time\\_statistics connection event” on page 492](#)
- [“download\\_statistics connection event” on page 403](#)
- [“download\\_statistics table event” on page 406](#)
- [“upload\\_statistics connection event” on page 512](#)
- [“upload\\_statistics table event” on page 517](#)
- [“synchronization\\_statistics connection event” on page 486](#)
- [“synchronization\\_statistics table event” on page 489](#)
- [“MobiLink Monitor” on page 179](#)
- [“Using remote IDs and MobiLink user names in scripts” \[\*MobiLink - Client Administration\*\]](#)

### SQL example

The following example inserts statistical information into the `time_statistics` table.

```
CALL ml_add_table_script (
  'ver1',
  'table1',
  'time_statistics',
  'INSERT INTO time_statistics(
    ml_user,
    table,
    event_name,
    number_of_calls,
    minimum_time,
    maximum_time,
    total_time)
VALUES (
  {ml s.username},
  {ml s.table},
  {ml s.event_name},
  {ml s.number_of_calls},
  {ml s.minimum_time},
  {ml s.maximum_time},
  {ml s.total_time} )' );
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `timeStatisticsTable` as the script for the `time_statistics` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
```

```
'time_statistics',
'ExamplePackage.ExampleClass.timeStatisticsTable' )
```

The following is the sample Java method timeStatisticsTable. It prints statistics for the upload\_old\_row\_insert event.

```
public String timeStatisticsTable(
    String username,
    String tableName,
    String eventName,
    int numberOfCalls,
    int minimumTime,
    int maximumTime,
    int totalTime ) {
    if( eventName.equals( "upload_old_row_insert" ) ) {
        java.lang.System.out.println(
            "upload_old_row_insert num_calls: " + numCalls +
            "total_time: " + totalTime );
    }
    return ( null );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called TimeTableStats as the script for the time\_statistics table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(
    'ver1',
    'table1',
    'time_statistics',
    'TestScripts.Test.TimeTableStats'
)
```

The following is the sample .NET method TimeTableStats. It prints statistics for the upload\_old\_row\_insert event.

```
public string TimeTableStats(
    string user,
    string table,
    string eventName,
    int numberOfCalls,
    int minimumTime,
    int maximumTime,
    int totTime ) {
    if( event_name == "upload_old_row_insert" ) {
        System.Console.WriteLine(
            "upload_old_row_insert num_calls: " + num_calls +
            "total_time: " + total_time );
    }
    return ( null );
}
```

## upload\_delete table event

Provides an event that the MobiLink server uses during processing of the upload to handle rows deleted from the remote database.

### Parameters

Parameter name for SQL scripts	Order
<i>r.pk-column-1</i>	1
...	...
<i>r.pk-column-N</i>	<i>N</i>
<i>r.column-1</i>	<i>N + 1</i>
...	...
<i>r.column-M</i>	<i>N + M</i>

### Default action

None.

### Remarks

The statement-based `upload_delete` script handles rows that are deleted on the remote database. The action taken at the consolidated database can be a `DELETE` statement, but need not be.

You can have one `upload_delete` script for each table in the remote database.

For Java and .NET applications, this script must return valid SQL.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“upload\\_insert table event” on page 504](#)
- [“upload\\_update table event” on page 522](#)

### SQL example

This example is taken from the Contact sample and can be found in `Samples\MobiLink>Contact\build_consol.sql`. It marks customers that are deleted from the remote database as inactive.

```
CALL ml_add_table_script(  
  'ver1',  
  'table1',  
  'upload_delete',  
  'UPDATE Customer  
   SET active = 0  
   WHERE cust_id={ml r.cust_id}' )
```

## Java example

The following call to a MobiLink system procedure registers a Java method called uploadDeleteTable as the script for the upload\_delete table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1',  
    'upload_delete',  
    'ExamplePackage.ExampleClass.uploadDeleteTable' )
```

The following is the sample Java method uploadDeleteTable. It calls genUD which dynamically generates an UPLOAD statement.

```
public String uploadDeleteTable() {  
    return( genUD(_curTable) );  
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called UploadDelete as the script for the upload\_delete table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'upload_delete',  
    'TestScripts.Test.UploadDelete'  
)
```

The following is the sample .NET method UploadDelete. It calls genUD which dynamically generates an UPLOAD statement.

```
public string UploadDelete( object pk1 ) {  
    return( genUD(_curTable) );  
}
```

## upload\_fetch table event

Fetches rows from a synchronized table in the consolidated database for the purpose of row-level conflict detection.

### Parameters

Parameter name for SQL scripts	Order
<i>r.primary-key-1</i>	1
<i>r.primary-key-2</i>	2
...	...
<i>r.primary-key-N</i>	N

### Default action

None.

### Remarks

The statement-based `upload_fetch` script fetches rows from a synchronized table for the purposes of conflict detection. It is a companion to the `upload_update` event.

The columns of the result set must match the number of columns being uploaded from the remote database for this table. If the values returned do not match the pre-image in the uploaded row, a conflict is identified.

Do not use `READPAST` table hints in `upload_fetch` scripts. If the script skips a locked row using `READPAST`, the synchronization logic thinks that the row was deleted. Depending on what scripts you have defined, this either causes the uploaded update to be ignored or it triggers conflict resolution. Ignoring the update is likely to be unacceptable behavior, and may be harmful. Triggering conflict resolution may not be a problem, depending on the resolution logic you have implemented.

You can have only one `upload_fetch` or `upload_fetch_column_conflict` script for each table in the remote database.

This script may be ignored if none of the following scripts are defined: `upload_new_row_insert`, `upload_old_row_insert`, and `resolve_conflict`.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Detecting conflicts” on page 147](#)
- [“resolve\\_conflict table event” on page 483](#)
- [“upload\\_delete table event” on page 498](#)
- [“upload\\_insert table event” on page 504](#)
- [“upload\\_update table event” on page 522](#)
- ["Using READPAST with MobiLink synchronization" in “FROM clause” \[SQL Anywhere Server - SQL Reference\]](#)



## SQL example

The following SQL script is taken from the Contact sample and can be found in *samples-dir\MobiLink\Contact\build\_consol.sql*. It is used to identify conflicts that occur when rows updated in the remote database Product table are uploaded. This script selects rows from a table also named Product, but depending on your consolidated and remote database schema, the two table names may not match.

```
CALL ml_add_table_script(
  'ver1',
  'Product',
  'upload_fetch',
  'SELECT id, name, size, quantity, unit_price
   FROM Product
   WHERE id={ml r.id}' )
```

## Java example

This script must return valid SQL.

The following call to a MobiLink system procedure registers a Java method called uploadFetchTable as the script for the upload\_fetch table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'Product',
  'upload_fetch',
  'ExamplePackage.ExampleClass.uploadFetchTable' )
```

The following is the sample Java method uploadFetchTable. It calls genUF to dynamically generate an UPLOAD statement.

```
public String uploadFetchTable() {
    return( genUF(_curTable) );
}
```

## .NET example

This script must return valid SQL.

The following call to a MobiLink system procedure registers a .NET method called UploadFetchTable as the script for the upload\_fetch table event when synchronizing the script version ver1.

```
CALL ml_add_dnet_table_script(
  'ver1',
  'Product',
  'upload_fetch',
  'TestScripts.Test.UploadFetchTable' )
```

The following is the sample .NET method UploadFetchTable. It calls GenUF to dynamically generate an UPLOAD statement.

```
public string UploadFetchTable() {
    return( GenUF(_curTable) );
}
```

## upload\_fetch\_column\_conflict table event

Fetches rows from a synchronized table in the consolidated database for the purpose of column-level conflict detection.

### Parameters

Parameter name for SQL scripts	Order
<i>r.pk-column-1</i>	1
...	...
<i>r.pk-column-N</i>	<i>N</i>
<i>r.column-1</i>	<i>N + 1</i>
...	...
<i>r.column-M</i>	<i>N + M</i>

### Default action

None.

### Remarks

The statement-based `upload_fetch_column_conflict` script fetches columns from a synchronized table for the purposes of conflict detection. It is a companion to the `upload_update` event.

This script only detects a conflict when two users update the same column. Different users can update the same row, as long as they don't update the same column, without generating a conflict.

For example, using the `upload_fetch_column_conflict` script, you could avoid detecting a conflict when one of your remote users updated the `quant` column of the `ULOrder` table, and another remote user updated the `notes` column for the same row. You would only detect a conflict if they both updated the `quant` column.

You can have only one `upload_fetch` or `upload_fetch_column_conflict` script for each table in the remote database.

This script may be ignored if none of the following scripts are defined: `upload_new_row_insert`, `upload_old_row_insert`, and `resolve_conflict`.

**See also**

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Detecting conflicts” on page 147](#)
- [“upload\\_fetch table event” on page 500](#)
- [“resolve\\_conflict table event” on page 483](#)
- [“upload\\_delete table event” on page 498](#)
- [“upload\\_insert table event” on page 504](#)
- [“upload\\_update table event” on page 522](#)

## upload\_insert table event

Provides an event that the MobiLink server uses during processing of the upload to handle rows inserted into the remote database.

### Parameters

Parameter name for SQL scripts	Order
<i>r.pk-column-1</i>	1
...	...
<i>r.pk-column-N</i>	<i>N</i>
<i>r.column-1</i>	<i>N+1</i>
...	...
<i>r.column-M</i>	<i>N+M</i>

### Default action

None.

### Remarks

The statement based `upload_insert` script performs direct inserts of column values.

You can have one `upload_insert` script for each table in the remote database.

For Java and .NET applications, this script must return valid SQL.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“upload\\_delete table event” on page 498](#)
- [“upload\\_update table event” on page 522](#)
- [“upload\\_fetch table event” on page 500](#)

### SQL example

This example handles inserts that were made on the Customer table in the remote database. The script inserts the values into a table named Customer in the consolidated database. The final column of the table identifies the Customer as active. The final column does not appear in the remote database.

```
CALL ml_add_table_script(  
  'ver1',  
  'Customer',  
  'upload_insert',  
  'INSERT INTO Customer(  
    cust_id,  
    name,
```

```

    rep_id,
    active )
VALUES (
  {ml r.cust_id},
  {ml r.name},
  {ml r.rep_id},
  1 )' );

```

### Java example

The following call to a MobiLink system procedure registers a Java method called uploadInsertTable as the script for the upload\_insert table event when synchronizing the script version ver1.

```

CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'upload_insert',
  'ExamplePackage.ExampleClass.uploadInsertTable' )

```

The following is the sample Java method uploadInsertTable. It dynamically generates an INSERT statement. This syntax is for SQL Anywhere consolidated databases.

```

public String uploadInsertTable() {
    return("INSERT INTO " + _curTable + getCols(_curTable)
        + " VALUES " + getQM(_curTable));
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called UploadInsert as the script for the upload\_insert table event when synchronizing the script version ver1 and the table table1. This syntax is for SQL Anywhere consolidated databases.

```

CALL ml_add_dnet_table_script(
  'ver1',
  'table1',
  'upload_insert',
  'TestScripts.Test.UploadInsert'
)

```

The following is the sample .NET method UploadInsert. It returns an INSERT statement for the ULCustomer table.

```

public static string UploadInsert() {
    return("INSERT INTO ULCustomer( cust_id, cust_name ) VALUES ( {ml
r.cust_id}, {ml r.cust_name} )");
}

```

## upload\_new\_row\_insert table event

Conflict resolution scripts for statement-based uploads commonly require access to the old and new values of rows uploaded from the remote database. This event allows you to handle the new, updated values of rows uploaded from the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See “[SQL-Java data types](#)” on page 532 and “[SQL-.NET data types](#)” on page 594.

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name. This parameter is optional.	Optional
<b>r.pk-column-1</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	1 (2 if username is referenced)
...		...
<b>r.pk-column-N</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N$ ( $N+1$ if username is referenced)
<b>r.column-1</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N + 1$ ( $N+2$ if username is referenced)
...	...	...
<b>r.column-M</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N + M$ ( $N+M+1$ if username is referenced)

### Default action

None.

## Remarks

When a MobiLink client sends an updated row to the MobiLink server, it includes not only the new values (the post-image), but also a copy of the old row values (the pre-image). When the pre-image does not match the current values in the consolidated database, a conflict is detected.

This event allows you to save post-image values to a table. You can use this event to assist in developing conflict resolution procedures for statement-based updates. The parameters for this event hold new row values from the remote database before the update is performed on the corresponding consolidated database table. This event is also used to insert rows in statement-based, forced-conflict mode.

The script for this event is usually an insert statement that inserts the new row into a temporary table for use by a `resolve_conflict` script.

You can have one `upload_new_row_insert` script for each table in the remote database.

For Java and .NET applications, this script must return valid SQL.

## See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “Handling conflicts” on page 146
- “`resolve_conflict` table event” on page 483
- “`upload_old_row_insert` table event” on page 509
- “`upload_update` table event” on page 522
- “Forced conflicts” on page 154
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

## SQL example

This example handles updates made on the `product` table in the remote database. The script inserts the new value of the row into a global temporary table named `product_conflict`. The final column of the table identifies the row as a new row.

```
CALL ml_add_table_script(
  'ver1',
  'table1',
  'upload_new_row_insert',
  'INSERT INTO DBA.product_conflict(
    id,
    name,
    size,
    quantity,
    unit_price,
    row_type )
VALUES(
  {ml r.id},
  {ml r.name},
  {ml r.size},
  {ml r.quantity},
  {ml r.unit_price},
  'New' )' )
```

### Java example

The following call to a MobiLink system procedure registers a Java method called `uploadNewRowInsertTable` as the script for the `upload_new_row_insert` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1',  
    'upload_new_row_insert',  
    'ExamplePackage.ExampleClass.uploadNewRowInsertTable'  
)
```

The following is the sample Java method `uploadNewRowInsertTable`. It dynamically generates an INSERT statement. This syntax is for SQL Anywhere consolidated databases.

```
public String uploadNewRowInsertTable() {  
    return("insert into" + _curTable + "_new" +  
        getCols(_curTable) + "values" + getNamedParams(_curTable));  
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called `UploadNewRowInsertTable` as the script for the `upload_new_row_insert` table event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'upload_new_row_insert',  
    'TestScripts.Test.UploadNewRowInsertTable'  
)
```

The following is the sample .NET method `UploadNewRowInsertTable`. It dynamically generates an INSERT statement. This syntax is for SQL Anywhere consolidated databases.

```
public string UploadNewRowInsertTable() {  
    return("insert into" + _curTable + "_new" +  
        GetCols(_curTable) + "values" + GetNamedParams(_curTable));  
}
```



## upload\_old\_row\_insert table event

Conflict resolution scripts for statement-based uploads commonly require access to the old and new values of rows uploaded from the remote database. This event allows you to handle the old values of rows uploaded from the remote database.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name. This parameter is optional.	Optional
<b>r.pk-column-1</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	1 (2 if username is referenced)
...		...
<b>r.pk-column-N</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N$ ( $N+1$ if username is referenced)
<b>r.column-1</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N + 1$ ( $N+2$ if username is referenced)
...	...	...
<b>r.column-M</b>	A column value from the old (pre-image) row, where the named parameter is specified as a column name prefaced by <b>r</b> .	$N + M$ ( $N+M+1$ if username is referenced)

### Default action

None.

### Remarks

When a MobiLink client sends an updated row to the MobiLink server, it includes not only the new values (the post-image), but also a copy of the old row values (the pre-image). When the pre-image does not match the current values in the consolidated database, a conflict is detected.

This event allows you to save pre-image values to a table. You can use this event to assist in developing conflict resolution procedures for statement-based updates. The parameters for this event hold old row values from the remote database before the update is performed on the corresponding consolidated database table. This event is also used to insert rows in statement-based, forced-conflict mode.

The script for this event is usually an insert statement that inserts the old row into a temporary table for use by a `resolve_conflict` script.

You can have one `upload_old_row_insert` script for each table in the remote database.

For Java and .NET applications, this script must return valid SQL.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “Handling conflicts” on page 146
- “resolve\_conflict table event” on page 483
- “upload\_new\_row\_insert table event” on page 506
- “upload\_update table event” on page 522
- “Forced conflicts” on page 154
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

### SQL example

This example handles updates made on the product table in the remote database. The script inserts the old value of the row into a global temporary table named `product_conflict`. The final column of the table identifies the row as an old row.

```
CALL ml_add_table_script(  
  'ver1',  
  'table1',  
  'upload_old_row_insert',  
  'INSERT INTO DBA.product_conflict (  
    id,  
    name,  
    size,  
    quantity,  
    unit_price,  
    row_type )  
VALUES (  
  {ml r.id},  
  {ml r.name},  
  {ml r.size},  
  {ml r.quantity},  
  {ml r.unit_price},  
  'Old' )' )
```

## Java example

The following call to a MobiLink system procedure registers a Java method called `uploadOldRowInsertTable` as the script for the `upload_old_row_insert` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1',  
    'upload_old_row_insert',  
    'ExamplePackage.ExampleClass.uploadOldRowInsertTable'  
)
```

The following is the sample Java method `uploadOldRowInsertTable`. It dynamically generates an INSERT statement.

```
public String uploadOldRowInsertTable() {  
    return( "old" + getCols(_curTable) +  
           "values" + getNamedParams(_curTable));  
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called `UploadOldRowInsertTable` as the script for the `upload_old_row_insert` table event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'upload_old_row_insert',  
    'TestScripts.Test.UploadOldRowInsertTable'  
)
```

The following is the sample .NET method `UploadOldRowInsertTable`. It dynamically generates an UPLOAD statement.

```
public string UploadOldRowInsertTable() {  
    return( "old" + GetCols(_curTable) +  
           "values" + GetNamedParams(_curTable));  
}
```

## upload\_statistics connection event

Tracks synchronization statistics for upload operations.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.warnings	INTEGER. The number of warnings that occurred.	2
s.errors	INTEGER. The number of errors that occurred.	3
s.inserted_rows	INTEGER. The number of rows that were successfully inserted in the consolidated database.	4
s.deleted_rows	INTEGER. The number of rows that were successfully deleted from the consolidated database.	5
s.updated_rows	INTEGER. The number of rows that were successfully updated in the consolidated database.	6
s.conflicted_inserts	INTEGER. Always zero.	7
s.conflicted_deletes	INTEGER. Always zero.	8

Parameter name for SQL scripts	Description	Order
s.conflicted_updates	INTEGER. The number of update rows that caused conflict. A row is included only when a resolve conflict script was successfully called for it.	9
s.ignored_inserts	INTEGER. The total number of upload insert rows that were ignored. They were ignored because 1) there is no upload_insert script in normal mode or no upload_new_row_insert script in forced conflict mode; or 2) errors occurred when the MobiLink server was invoking the corresponding script and the handle_error or handle_odbc_error event returned 1000.	10
s.ignored_deletes	INTEGER. The number of upload delete rows that caused errors while the upload_delete script was invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_delete script defined for the given table.	11
s.ignored_updates	INTEGER. The number of upload update rows that caused conflict but a resolve conflict script was not successfully called or no upload_update script was defined.	12
s.bytes	INTEGER. The amount of memory used within the MobiLink server to store the upload.	13
s.deadlocks	INTEGER. The number of deadlocks in the consolidated database that were detected for the synchronization.	14

**Default action**

None.

**Remarks**

The upload\_statistics event allows you to gather, for any user, statistics on uploads. The upload\_statistics connection script is called just prior to the commit at the end of the upload transaction.

### See also

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “download\_statistics connection event” on page 403
- “download\_statistics table event” on page 406
- “upload\_statistics table event” on page 517
- “synchronization\_statistics connection event” on page 486
- “synchronization\_statistics table event” on page 489
- “time\_statistics connection event” on page 492
- “time\_statistics table event” on page 495
- “MobiLink Monitor” on page 179
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

### SQL example

The following example inserts synchronization statistics for upload operations into the table `upload_summary_audit`.

```
CALL ml_add_connection_script (
  'ver1',
  'upload_statistics',
  'INSERT INTO upload_summary_audit (
    ml_user,
    warnings,
    errors,
    inserted_rows,
    deleted_rows,
    updated_rows,
    conflicted_inserts,
    conflicted_deletes,
    conflicted_updates,
    bytes,
    ignored_inserts,
    ignored_deletes,
    ignored_updates,
    bytes, deadlocks )
VALUES (
  {ml s.username},
  {ml s.warnings},
  {ml s.errors},
  {ml s.inserted_rows},
  {ml s.deleted_rows},
  {ml s.updated_rows},
  {ml s.conflicted_inserts},
  {ml s.conflicted_deletes},
  {ml s.conflicted_updates},
  {ml s.ignored_inserts},
  {ml s.ignored_deletes},
  {ml s.ignored_updates},
  {ml s.bytes},
  {ml s.deadlocks} ) ' )
```

Once statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

## Java example

The following call to a MobiLink system procedure registers a Java method called `uploadStatisticsConnection` as the script for the `upload_statistics` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(
  'ver1',
  'upload_statistics',
  'ExamplePackage.ExampleClass.uploadStatisticsConnection' )
```

The following is the sample Java method `uploadStatisticsConnection`. It logs some statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String uploadStatisticsConnection(
  String user,
  int warnings,
  int errors,
  int insertedRows,
  int deletedRows,
  int updatedRows,
  int conflictedInserts,
  int conflictedDeletes,
  int conflictedUpdates,
  int ignoredInserts,
  int ignoredDeletes,
  int ignoredUpdates,
  int bytes,
  int deadlocks ) {
  java.lang.System.out.println( "updated rows: " +
    updatedRows );
  return ( null );
}
```

## .NET example

The following call to a MobiLink system procedure registers a .NET method called `UploadStats` as the script for the `upload_statistics` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_dnet_connection_script(
  'ver1',
  'upload_statistics',
  'TestScripts.Test.UploadStats'
)
```

The following is the sample .NET method `UploadStats`. It logs some statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public string UploadStats (
  string user,
  int warnings,
  int errors,
  int insertedRows,
  int deletedRows,
  int updatedRows,
  int conflictInserts,
  int conflictDeletes,
  int conflictUpdates,
  int ignoredInserts,
```

```
int ignoredDeletes,  
int ignoredUpdates,  
int bytes,  
int deadlocks ) {  
System.Console.WriteLine( "updated rows: " +  
    updatedRows );  
return ( null );  
}
```



## upload\_statistics table event

Tracks synchronization statistics for upload operations for a specific table.

### Parameters

In the following table, the description provides the SQL data type. If you are writing your script in Java or .NET, you should use the appropriate corresponding data type. See [“SQL-Java data types” on page 532](#) and [“SQL-.NET data types” on page 594](#).

In SQL scripts, you can specify event parameters by name or with a question mark, but you cannot mix names and question marks within a script. If you use question marks, the parameters must be in the order shown below and are optional only if no subsequent parameters are specified (for example, you must use parameter 1 if you want to use parameter 2). If you use named parameters, you can specify any subset of the parameters in any order.

Parameter name for SQL scripts	Description	Order
s.remote_id	VARCHAR(128). The MobiLink remote ID. You can only reference the remote ID if you are using named parameters.	Not applicable
s.username	VARCHAR(128). The MobiLink user name.	1
s.table	VARCHAR(128). The table name.	2
s.warnings	INTEGER. The number of warnings issued in the upload of the table.	3
s.errors	INTEGER. The number of errors, including handled errors, that occurred in the upload of the table.	4
s.inserted_rows	INTEGER. The number of rows that were successfully inserted in the consolidated database.	5
s.deleted_rows	INTEGER. The number of rows that were successfully deleted from the consolidated database.	6
s.updated_rows	INTEGER.	7
s.conflicted_inserts	INTEGER. Always zero.	8
s.conflicted_deletes	INTEGER. Always zero.	9

Parameter name for SQL scripts	Description	Order
s.conflicted_updates	INTEGER. The number of update rows that caused conflict. A row is included only when a resolve conflict script was successfully called for it.	10
s.ignored_inserts	INTEGER. The total number of upload insert rows that were ignored. They were ignored because 1) there is no upload_insert script in normal mode or no upload_new_row_insert script in forced conflict mode; or 2) errors occurred when the MobiLink server was invoking the corresponding script and the handle_error or handle_odbc_error event returned 1000.	11
s.ignored_deletes	INTEGER. The number of upload delete rows that caused errors while the upload_delete script was invoked, when the handle_error or handle_odbc_error are defined and returned 1000, or when there is no upload_delete script defined for the given table.	12
s.ignored_updates	INTEGER. The number of upload update rows that caused conflict but a resolve conflict script was not successfully called or no upload_update script was defined.	13
s.bytes	INTEGER. The amount of memory used within the MobiLink server to store the upload.	14
s.deadlocks	INTEGER. The number of deadlocks in the consolidated database that were detected for the synchronization.	15

**Default action**

None.

**Remarks**

The upload\_statistics event allows you to gather, for any user, vital statistics on synchronization happenings as they apply to any table. The upload\_statistics table script is called just prior to the commit at the end of the upload transaction.

**See also**

- “Script parameters” on page 320
- “Adding and deleting scripts” on page 327
- “download\_statistics connection event” on page 403
- “upload\_statistics connection event” on page 512
- “upload\_statistics table event” on page 517
- “synchronization\_statistics connection event” on page 486
- “synchronization\_statistics table event” on page 489
- “time\_statistics connection event” on page 492
- “time\_statistics table event” on page 495
- “MobiLink Monitor” on page 179
- “Using remote IDs and MobiLink user names in scripts” [*MobiLink - Client Administration*]

**SQL Example**

The following example inserts a row into a table used to track upload statistics.

```
CALL ml_add_connection_script(
  'ver1',
  'upload_statistics',
  'INSERT INTO my_upload_statistics (
    user_name,
    table_name,
    num_warnings,
    num_errors,
    inserted_rows,
    deleted_rows,
    updated_rows,
    conflicted_inserts,
    conflicted_deletes,
    conflicted_updates,
    ignored_inserts,
    ignored_deletes,
    ignored_updates, bytes,
    deadlocks )
VALUES(
  {ml s.username},
  {ml s.table},
  {ml s.warnings},
  {ml s.errors},
  {ml s.inserted_rows},
  {ml s.deleted_rows},
  {ml s.updated_rows},
  {ml s.conflicted_inserts},
  {ml s.conflicted_deletes},
  {ml s.conflicted_updates},
  {ml s.ignored_inserts},
  {ml s.ignored_deletes},
  {ml s.ignored_updates},
  {ml s.bytes},
  {ml s.deadlocks} )' )
```

The following example works with an Oracle consolidated database.

```
CALL ml_add_connection_script(
  'ver1',
  'upload_statistics',
  'INSERT INTO upload_tables_audit (
    id,
```

```
user_name,  
table,  
warnings,  
errors,  
inserted_rows,  
deleted_rows,  
updated_rows,  
conflicted_inserts,  
conflicted_deletes,  
conflicted_updates,  
ignored_inserts,  
ignored_deletes,  
ignored_updates,  
bytes,  
deadlocks )  
VALUES (  
    ut_audit.nextval,  
    {ml s.username},  
    {ml s.table},  
    {ml s.warnings},  
    {ml s.errors},  
    {ml s.inserted_rows},  
    {ml s.deleted_rows},  
    {ml s.updated_rows},  
    {ml s.conflicted_inserts},  
    {ml s.conflicted_deletes},  
    {ml s.conflicted_updates},  
    {ml s.ignored_inserts},  
    {ml s.ignored_deletes},  
    {ml s.ignored_updates},  
    {ml s.bytes},  
    {ml s.deadlocks} )' )
```

Once statistics are inserted into the audit table, you may use these statistics to monitor your synchronizations and make optimizations where applicable.

### Java example

The following call to a MobiLink system procedure registers a Java method called `uploadStatisticsTable` as the script for the `upload_statistics` table event when synchronizing the script version `ver1`.

```
CALL ml_add_java_table_script(  
    'ver1',  
    'table1',  
    'upload_statistics',  
    'ExamplePackage.ExampleClass.uploadStatisticsTable' )
```

The following is the sample Java method `uploadStatisticsTable`. It logs some statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```
public String uploadStatisticsTable(  
    String user,  
    int warnings,  
    int errors,  
    int insertedRows,  
    int deletedRows,  
    int updatedRows,  
    int conflictedInserts,  
    int conflictedDeletes,  
    int conflictedUpdates,  
    int ignoredInserts,
```

```

    int ignoredDeletes,
    int ignoredUpdates,
    int bytes,
    int deadlocks ) {
    java.lang.System.out.println( "updated rows: " +
        updatedRows );
    return ( null );
}

```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called UploadTableStats as the script for the upload\_statistics table event when synchronizing the script version ver1 and the table table1.

```

CALL ml_add_dnet_table_script(
    'ver1',
    'table1',
    'upload_statistics',
    'TestScripts.Test.UploadTableStats'
)

```

The following is the sample .NET method uploadStatisticsTable. It logs some statistics to the MobiLink message log. (Note that logging statistics to the MobiLink message log might be useful at development time but would slow down a production server.)

```

public string UploadTableStats(
    string user,
    string table,
    int warnings,
    int errors,
    int insertedRows,
    int deletedRows,
    int updatedRows,
    int conflictInserts,
    int conflictDeletes,
    int conflictUpdates,
    int ignoredInserts,
    int ignoredDeletes,
    int ignoredUpdates,
    int bytes,
    int deadlocks ) {
    System.Console.WriteLine( "updated rows: " +
        updatedRows );
    return ( null );
}

```

## upload\_update table event

Provides an event that the MobiLink server uses during processing of the upload to handle rows updated at the remote database.

### Parameters

Parameter	Order
<i>r.column-1</i>	1
...	...
<i>r.column-M</i>	M
<i>r.pk-column-1</i>	M + 1
...	...
<i>r.pk-column-N</i>	M + N
<i>o.column-N</i>	M + N + 1
...	...
<i>o.column-M</i>	M + N + M

### Default action

None.

### Remarks

The statement-based `upload_update` script may perform direct updates of column values as specified in the `UPLOAD` statement.

The `WHERE` clause must include all the primary key columns being synchronized. The `SET` clause must contain all the non-primary key columns being synchronized.

You use as many non-primary key columns in your `SET` clause as exist in the table, and MobiLink sends the correct number of column values. Similarly, in the `WHERE` clause, you can have any number of primary keys, but all must be specified here, and MobiLink sends the correct values. MobiLink sends these column values and primary key values in the order the columns or primary keys appear in a MobiLink report of your schema. You can use the `-vh` option to determine the column ordering for this table schema.

For example, in the following `upload_update` script, the question marks are in good order:

```
UPDATE MyTable
SET column_1 = ?, ..., column_M = ?
WHERE pk_column_1 = ? AND ... AND pk_column_N = ?
```

You can have one `upload_update` script for each table in the remote database.

For Java and .NET applications, this script must return valid SQL.

To use the upload\_update script to detect conflicts, include all non-primary key columns in the WHERE clause:

```
UPDATE table-name
SET col1 = {ml r.col1}, col2 = {ml r.col2} ...
WHERE pk1 = {ml r.pk1} AND pk2 = {ml r.pk2} ...
  AND col1 = {ml o.col1} AND col2 = {ml o.col2} ...
```

In this statement, col1 and col2 are the non-primary key columns, while pk1 and pk2 are primary key columns. The values passed to the second set of non-primary key columns are the pre-image of the updated row. The WHERE clause compares old values uploaded from the remote to current values in the consolidated database. If the values do not match, the update is ignored, preserving the values already on the consolidated database.

### See also

- [“Script parameters” on page 320](#)
- [“Adding and deleting scripts” on page 327](#)
- [“Detecting conflicts with upload\\_update scripts” on page 148](#)
- [“Resolving conflicts with upload\\_update scripts” on page 151](#)
- [“upload\\_delete table event” on page 498](#)
- [“upload\\_fetch table event” on page 500](#)
- [“upload\\_insert table event” on page 504](#)

### SQL example

This example handles updates made to the Customer table in the remote database. The script updates the values in a table named Customer in the consolidated database.

```
CALL ml_add_table_script(
  'ver1',
  'table1',
  'upload_update',
  'UPDATE Customer
   SET name = {ml r.name}, rep_id = {ml r.rep_id}
   WHERE cust_id = {ml o.cust_id}')
```

### Java example

The following call to a MobiLink system procedure registers a Java method called uploadUpdateTable as the script for the upload\_update table event when synchronizing the script version ver1.

```
CALL ml_add_java_table_script(
  'ver1',
  'table1',
  'upload_update',
  'ExamplePackage.ExampleClass.uploadUpdateTable' )
```

The following is the sample Java method uploadUpdateTable. It calls a method called genUU to dynamically generate an UPLOAD statement.

```
public String uploadUpdateTable() {
  return( genUU(_curTable) );
}
```

### .NET example

The following call to a MobiLink system procedure registers a .NET method called UploadUpdate as the script for the upload\_update table event when synchronizing the script version ver1 and the table table1.

```
CALL ml_add_dnet_table_script(  
    'ver1',  
    'table1',  
    'upload_update',  
    'TestScripts.Test.UploadUpdate'  
)
```

The following is the sample .NET method UploadUpdate. It calls a method called GenUU to dynamically generate an UPLOAD statement.

```
public string UploadUpdate() {  
    return ( genUU(_curTable) );  
}
```



# MobiLink Server APIs

This section describes the MobiLink Server APIs for Java and .NET.

---

Writing synchronization scripts in Java .....	527
Writing synchronization scripts in .NET .....	589
Direct row handling .....	649



---

# Writing synchronization scripts in Java

## Contents

Introduction to Java synchronization logic .....	528
Setting up Java synchronization logic .....	529
Writing Java synchronization logic .....	531
Java synchronization example .....	538
MobiLink server API for Java reference .....	543

---

## Introduction to Java synchronization logic

You control the actions of the MobiLink server by writing synchronization scripts. You can implement these scripts in SQL, .NET or Java. Java synchronization logic can function just as SQL logic functions: the MobiLink server can make calls to Java methods on the occurrence of MobiLink events just as it accesses SQL scripts on the occurrence of MobiLink events. A Java method can return a SQL string to MobiLink.

This section tells you how to set up, develop, and run Java synchronization logic. It includes a sample application and the MobiLink server API for Java Reference.

### See also

- [“Tutorial: Using Java synchronization logic” \[MobiLink - Getting Started\]](#)
- [“Options for writing server-side synchronization logic” \[MobiLink - Getting Started\]](#)
- [“Writing synchronization scripts” on page 313](#)

## Setting up Java synchronization logic

When you install SQL Anywhere, the installer automatically sets the location of the MobiLink server API for Java classes. When you start the MobiLink server, it automatically includes these classes in your classpath. The MobiLink server API for Java classes are located in *install-dir\java\mlscript.jar*.

### To implement synchronization scripts in Java

1. Create your own class or classes. Write a method for each required synchronization script. These methods must be public. The class must be public in the package.

See [“Methods” on page 533](#).

Each class with non-static methods should have a public constructor. The MobiLink server automatically instantiates each class the first time a method in that class is called.

See [“Constructors” on page 532](#).

2. When compiling the class, you must include the JAR file *java\mlscript.jar*.

For example,

```
javac MyClass.java -classpath "c:\Program Files\SQL Anywhere 11\java
\mlscript.jar"
```

3. In the MobiLink system tables on your consolidated database, specify the name of the package, class, and method to call for each synchronization script. One class is permitted per script version.

For example, you can add this information to the MobiLink system tables using the `ml_add_java_connection_script` stored procedure or the `ml_add_java_table_script` stored procedure.

For example, the following SQL statement, when run in a SQL Anywhere database, specifies that for the script version `ver1`, `myPackage.myClass.myMethod` should be run whenever the `authenticate_user` connection-level event occurs. The method that is specified must be the fully qualified name of a public Java method, and the name is case sensitive.

```
call ml_add_java_connection_script('ver1',
'authenticate_user', 'myPackage.myClass.myMethod')
```

For more information about adding scripts, see:

- [“System procedures to add or delete scripts” on page 664](#)
  - [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
  - [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)
4. Instruct the MobiLink server to load classes. A vital part of setting up Java synchronization logic is to tell the virtual machine where to look for Java classes. There are two ways to do this:
    - Use the `mlsrv11 -sl java -cp` option to specify a set of directories or jar files in which to search for classes. For example, run the following command:

```
mlsrv11 -c "dsn=consolidated1" -sl java (-cp %classpath%;c:\local\Java
\myclasses.jar)
```

The MobiLink server automatically appends the location of the MobiLink server API for Java classes (java\mlscript.jar) to the set of directories or jar files. The -sl java option also forces the Java VM to load on server startup.

For more information about the available Java options, see [“-sl java option” on page 92](#).

- Explicitly set the classpath. To set the classpath for user-defined classes, use a statement such as the following:

```
SET classpath=%classpath%;c:\local\Java\myclasses.jar
```

If your system classpath includes your Java synchronization logic classes, you do not need to make changes to your MobiLink server command line.

You can use the -sl java option to force the Java virtual machine to load at server startup. Otherwise, the Java virtual machine is started when the first Java method is executed.

For more information about the available Java options, see [“-sl java option” on page 92](#).

5. On Unix, if you want to load a specific JRE, you should set the LD\_LIBRARY\_PATH (LIBPATH on AIX, SHLIB\_PATH on HP-UX) to include the directory containing the JRE. The directory must be listed before any of the SQL Anywhere installation directories.

### See also

- [“Writing Java synchronization logic” on page 531](#)
- [“Java synchronization example” on page 538](#)
- [“Tutorial: Using Java synchronization logic” \[MobiLink - Getting Started\]](#)
- [“MobiLink server API for Java reference” on page 543](#)
- [“Options for writing server-side synchronization logic” \[MobiLink - Getting Started\]](#)
- [“Writing synchronization scripts” on page 313](#)

---

## Writing Java synchronization logic

Writing Java synchronization logic requires knowledge of MobiLink events, some knowledge of Java, and knowledge of the MobiLink server API for Java.

For a complete description of the API, see [“MobiLink server API for Java reference” on page 543](#).

Java synchronization logic can be used to maintain state information, and implement logic around the upload and download events. For example, a `begin_synchronization` script written in Java could store the MobiLink user name in a variable. Scripts called later in the synchronization process can access this variable. Also, you can use Java to access rows in the consolidated database, before or after they are committed.

Using Java reduces dependence on the consolidated database. Behavior is affected less by upgrading the consolidated database to a new version or switching to a different database management system.

### Direct row handling

You can use MobiLink direct row handling to communicate remote data to any central data source, application, or web service. Direct row handling uses special classes in the MobiLink server APIs for Java or .NET for direct access to synchronized data.

See [“Direct row handling” on page 649](#).

## Class instances

The MobiLink server instantiates your classes at the connection level. When an event is reached for which you have written a non-static Java method, the MobiLink server automatically creates an instance of the class, if it has not already done so on the present connection.

See [“Constructors” on page 594](#).

All methods directly associated with a connection-level or table-level event for one script version **must belong to the same class**.

For each database connection, once a class has been instantiated, the class persists until that connection is closed. So, the same instance may be used for multiple consecutive synchronization sessions. Unless it is explicitly cleared, information present in public or private variables persists across synchronizations that occur on the same connection.

You can also use static classes or variables. In this case, the values are available across all connections.

The MobiLink server automatically deletes your class instances only when the connection to the consolidated database is closed.

## Transactions

The normal rules regarding transactions apply to Java methods. The start and duration of database transactions is critical to the synchronization process. Transactions must be started and ended only by the

MobiLink server. Explicitly committing or rolling back transactions on the synchronization connection within a Java method violates the integrity of the synchronization process and can cause errors.

These rules apply only to the database connections created by the MobiLink server and, in particular, to SQL statements returned by methods. If your classes create other database connections, use existing management rules to manage classes created by other database connections.

## SQL-Java data types

The following table shows SQL data types and the corresponding Java data types.

SQL data type	Corresponding Java data type
VARCHAR	java.lang.String
CHAR	java.lang.String
INTEGER	int or Integer
BINARY	byte[ ]
TIMESTAMP	java.sql.Timestamp
INOUT INTEGER	ianywhere.ml.script.InOutInteger
INOUT VARCHAR	ianywhere.ml.script.InOutString
INOUT CHAR	ianywhere.ml.script.InOutString
INOUT BYTEARRAY	ianywhere.ml.script.InOutByteArray

The MobiLink server automatically adds the `ianywhere.ml.script` package to your classpath if it is not already present. However, when you compile your class you need to add the path of `install-dir\java\mlscript.jar`.

## Constructors

The constructor of your class may have one of two possible signatures.

```
public MyScriptClass(ianywhere.ml.script.DBConnectionContext sc)
```

or

```
public MyScriptClass()
```

The synchronization context passed to you is for the connection through which the MobiLink server is synchronizing the current user.



The `DBConnectionContext.getConnection` method returns the same database connection that MobiLink is using to synchronize the present user. You can execute statements on this connection, but you must not commit or roll back the transaction. The MobiLink server manages the transactions.

The MobiLink server prefers to use constructors with the first signature. It only uses the non-argument constructor if a constructor with the first signature is not present.

See [“DBConnectionContext interface” on page 543](#).

## Methods

In general, you implement one method for each synchronization event. These methods must be public. If they are private, the MobiLink server cannot use them and fails to recognize that they exist.

The names of the methods are not important, as long as the names match the names specified in the `ml_script` table in the consolidated database. In the examples included in the documentation, however, the method names are the same as those of the MobiLink events because this naming convention makes the Java code easier to read.

The signature of your method should match the signature of the script for that event, except that you can truncate the parameter list if you do not need the values of parameters at the end of the list. You should accept only the parameters you need, because overhead is associated with passing the parameters.

You cannot, however, overload the methods. Only one method prototype per class may appear in the `ml_script` system table.

### Registering methods

After creating a method, you must register it. Registering the method creates a reference to the method in the MobiLink system tables on the consolidated database, so that the method is called when the event occurs. You register methods in the same way that you add synchronization scripts, except instead of adding the entire SQL script to the MobiLink system table, you add only the method name.

See [“Adding and deleting scripts” on page 327](#).

### Return values

Methods called for a SQL-based upload or download must return a valid SQL language statement. The return type of these methods must be `java.lang.String`. No other return types are allowed.

The return type of all other scripts must either be `java.lang.String` or `void`. No other types are allowed. If the return type is a string and not null, the MobiLink server assumes that the string contains a valid SQL statement and executes this statement in the consolidated database as it would an ordinary SQL-language synchronization script. If a method ordinarily returns a string but does not want to execute a SQL statement against the database upon its return, it can return null.

## Debugging Java classes

MobiLink provides various information and facilities that you may find helpful when debugging your Java code. This section describes where you can find this information and how you can exploit these capabilities.

### Information in the MobiLink server's log file

The MobiLink server writes messages to a message log file. The server log file contains the following information:

- The Java Runtime Environment. You can use the `-jrepath` option to request a particular JRE when you start the MobiLink server. The default path is the path of the JRE installed with SQL Anywhere 11.
- The path of the standard Java classes loaded. If you did not specify these explicitly, the MobiLink server automatically adds them to your classpath before invoking the Java virtual machine.
- The fully specified names of the specific methods invoked. You can use this information to verify that the MobiLink server is invoking the correct methods.
- Any output written in a Java method to `java.lang.System.out` or `java.lang.System.err` is redirected to the MobiLink server log file.
- The `mlsrv11` command line option `-verbose` can be used.

See [“-v option” on page 102](#).

### Using a Java debugger

You can debug your Java classes using a standard Java debugger. Specify the necessary parameters using the `-sl java` option on the `mlsrv11` command line.

See [“-sl java option” on page 92](#).

Specifying a debugger causes the Java virtual machine to pause and wait for a connection from a Java debugger.

### Printing information from Java

Alternatively, you may choose to add statements to your Java methods that print information to the MobiLink message log, using `java.lang.System.err` or `java.lang.System.out`. Doing so can help you track the progress and behavior of your classes.

#### Performance tip

Printing information in this manner is a useful monitoring tool, but is not recommended in a production scenario.

The same technique can be exploited to log arbitrary synchronization information or collect statistical information on how your scripts are used.

### Writing your own test driver

You may want to write your own driver to exercise your Java classes. This approach can be helpful because it isolates the actions of your Java methods from the rest of the MobiLink system.

## Handling MobiLink server errors in Java

When scanning the log is not enough, you can monitor your applications programmatically. For example, you can send messages of a certain type in an email.

You can write methods that are passed a class representing every error or warning message that is printed to the log. This may help you monitor and audit a MobiLink server.

The following code installs a LogListener for all warning messages, and writes the information to a file.

```
class TestLogListener implements LogListener {
    FileOutputStream _out_file;

    public TestLogListener(FileOutputStream out_file) {
        _out_file = out_file;
    }

    public void messageLogged(ServerContext sc, LogMessage msg) {
        String type;
        String user;

        try {
            if (msg.getType() == LogMessage.ERROR) {
                type = "ERROR";
            } else if (msg.getType() == LogMessage.WARNING) {
                type = "WARNING";
            }
            else {
                type = "UNKNOWN!!!";
            }

            user = msg.getUser();
            if (user == null) {
                user = "NULL";
            }
            _out_file.write(("Caught msg type="
                + type
                + " user=" + user
                + " text=" + msg.getText()
                + "\n").getBytes()
            );
            _out_file.flush();
        }
        catch(Exception e) {
            // Print some error output to the MobiLink log.
            e.printStackTrace();
        }
    }
}
```

The following code registers TestLogListener to receive warning messages. Call this code from anywhere that has access to the ServerContext such as a class constructor or synchronization script.

```
// ServerContext serv_context;
serv_context.addWarningListener(
    new MyLogListener(ll_out_file)
);
```

### See also

- [“addErrorListener method” on page 567](#)
- [“removeErrorListener method” on page 572](#)
- [“addWarningListener method” on page 568](#)
- [“removeWarningListener method” on page 573](#)
- [“LogListener interface” on page 560](#)
- [“LogMessage class” on page 561](#)

## User-defined start classes

You can define start classes that are loaded automatically when the server is started. The purpose of this feature is to allow you to write Java code that executes at the time the MobiLink server starts the JVM—before the first synchronization. This means you can create connections or cache data before a user synchronization request.

You do this with the `DMLStartClasses` option of the `mlsrv11 -sl java` option. For example, the following is part of a `mlsrv11` command line. It causes `mycl1` and `mycl2` to be loaded as start classes.

```
-sl java(-DMLStartClasses=com.test.mycl1,com.test.mycl2)
```

Classes are loaded in the order in which they are listed. If the same class is listed more than once, more than one instance is created.

All start classes must be public and must have a public constructor that either accepts no arguments or accepts one argument of type `ianywhere.ml.script.ServerContext`.

The names of loaded start classes are output to the MobiLink log with the message "Loaded JAVA start class: *classname*".

For more information about Java virtual machine options, see [“-sl java option” on page 92](#).

To see the start classes that are constructed at server start time, see [“getStartClassInstances method” on page 571](#).

### Example

The following is a start class template. It starts a daemon thread that processes events and creates a database connection. (Not all start classes need to create a thread but if a thread is spawned it should be a daemon thread.)

```
import ianywhere.ml.script.*;
import java.sql.*;

public class StartTemplate extends
    Thread implements ShutdownListener {
    ServerContext    _sc;
    Connection      _conn;
    boolean         _exit_loop;

    public StartTemplate(ServerContext sc) throws SQLException {

        // Perform setup first so that an exception
        // causes MobiLink startup to fail.
        _sc = sc;
    }
}
```

```
        // Create a connection for use later.
        _conn = _sc.makeConnection();

        _exit_loop = false;
        setDaemon(true);
        start();
    }

    public void run() {
        _sc.addShutdownListener(this);

        // run() cannot throw exceptions.
        try {
            handlerLoop();
            _conn.close();
            _conn = null;
        }
        catch(Exception e) {

            // Print some error output to the MobiLink log.
            e.printStackTrace();

            // This thread shuts down and so does not
            // need to be notified of shutdown.
            _sc.removeShutdownListener(this);

            // Ask server to shutdown so that this fatal
            // error is fixed.
            _sc.shutdown();
        }
        // Shortly after return this thread no longer exists.
        return;
    }

    // stop our event handler loop
    public void shutdownPerformed(ServerContext sc) {
        try {
            // Wait max 10 seconds for thread to die.
            join(10*1000);
        }
        catch(Exception e) {
            // Print some error output to the MobiLink log.
            e.printStackTrace();
        }
    }

    private void handlerLoop() throws InterruptedException {
        while (!_exit_loop) {
            // Handle events in this loop. Sleep not
            // needed, block on event queue.
            sleep(1 * 1000);
        }
    }
}
```

## Java synchronization example

Java synchronization logic works with MobiLink and common Java classes to provide you with flexibility in deploying applications using MobiLink server. The following section introduces you to this extended range of functionality using a simple example.

This section describes a working example of Java synchronization logic. Before you try to use this class or write your own class, use the following checklist to ensure you have all the pieces in place before compiling the class.

- Plan your functionality using, for example, pseudocode.
- Create a map of database tables and columns.
- Configure the consolidated database for Java synchronization by ensuring you have specified in the MobiLink system tables the language type and location of the Java synchronization methods.

See [“Setting up Java synchronization logic” on page 529](#).

- Create a list of associated Java classes that are called during the running of your Java class.
- Store your Java classes in a location that is in the classpath for MobiLink server.

### Plan

The Java synchronization logic for this example points to the associated Java files and classes that contain functionality needed for the example to work. It shows you how to create a class `CustEmpScripts`. It shows you how to set up a synchronization context for the connection. Finally, the example provides Java methods to

- Authenticate a MobiLink user
- Perform download and upload operations using cursors for each database table.

### Schema

The tables to be synchronized are `emp` and `cust`. The `emp` table has three columns called `emp_id`, `emp_name` and `manager`. The `cust` table has three columns called `cust_id`, `cust_name` and `emp_id`. All columns in each table are synchronized. The mapping from consolidated to remote database is such that the table names and column names are identical in both databases. One additional table, an audit table, is added to the consolidated database.

### Java class files

The files used in the example are included in the *Samples\MobiLink\JavaAuthentication* directory.

### Setup

The following code sets up the Java synchronization logic. The import statements tell the Java virtual machine the location of needed files. The public class statement declares the class.

```
// Use a package when you create your own script.  
import ianywhere.ml.script.InOutInteger;  
import ianywhere.ml.script.DBConnectionContext;  
import ianywhere.ml.script.ServerContext;
```

```

import java.sql.*;

public class CustEmpScripts {
    // Context for this synchronization connection.
    DBConnectionContext _conn_context;

    // Same connection MobiLink uses for sync.
    // Do not commit or close this.
    Connection _sync_connection;
    Connection _audit_connection;

    //Get a user id given the user name. On audit connection.
    PreparedStatement _get_user_id_pstmt;

    // Add record of user logins added. On audit connection.
    PreparedStatement _insert_login_pstmt;

    // Prepared statement to add a record to the audit table.
    // On audit connection.
    PreparedStatement _insert_audit_pstmt;

    // ...
}

```

The CustEmpScripts constructor sets up all the prepared statements for the authenticateUser method. It sets up member data.

```

public CustEmpScripts(DBConnectionContext cc) throws SQLException {
    try {
        _conn_context = cc;
        _sync_connection = _conn_context.getConnection();

        ServerContext serv_context =
            _conn_context.getServerContext();
        _audit_connection = serv_context.makeConnection();

        // Get the prepared statements ready.
        _get_user_id_pstmt =
            _audit_connection.prepareStatement(
                "select user_id from ml_user where name = ?"
            );

        _insert_login_pstmt =
            _audit_connection.prepareStatement(
                "INSERT INTO login_added(ml_user, add_time) "
                + "VALUES (?, { fn CONVERT({ fn NOW() }, SQL_VARCHAR) })"
            );

        _insert_audit_pstmt =
            _audit_connection.prepareStatement(
                "INSERT INTO login_audit(ml_user_id, audit_time,
audit_action) "
                + "VALUES (?, { fn CONVERT({ fn NOW() },
SQL_VARCHAR) }, ?)"
            );
    }
    catch(SQLException e) {
        freeJDBCResources();
        throw e;
    }
    catch(Error e) {
        freeJDBCResources();
        throw e;
    }
}

```

```
    }  
}
```

The finalize method cleans up JDBC resources if end\_connection is not called. It calls the freeJDBCResources method, which frees allocated memory and closes the audit connection.

```
protected void finalize() throws SQLException, Throwable {  
    super.finalize();  
    freeJDBCResources();  
}  
  
private void freeJDBCResources() throws SQLException {  
    if (_get_user_id_pstmt != null) {  
        _get_user_id_pstmt.close();  
    }  
    if (_insert_login_pstmt != null) {  
        _insert_login_pstmt.close();  
    }  
    if (_insert_audit_pstmt != null) {  
        _insert_audit_pstmt.close();  
    }  
    if (_audit_connection != null) {  
        _audit_connection.close();  
    }  
    _conn_context      = null;  
    _sync_connection   = null;  
    _audit_connection  = null;  
    _get_user_id_pstmt = null;  
    _insert_login_pstmt = null;  
    _insert_audit_pstmt = null;  
}
```

The endConnection method cleans up resources once the resources are not needed.

```
public void endConnection() throws SQLException {  
    freeJDBCResources();  
}
```

The authenticateUser method below approves all user logins and logs user information to database tables. If the user is not in the ml\_user table they are logged to login\_added. If the user id is found in ml\_user then they are logged to login\_audit. In a real system you would not ignore the user\_password, but this sample approves all users for simplicity. The endConnection method throws SQLException if any of the database operations fail with an exception.

```
public void authenticateUser(  
    InOutInteger authentication_status,  
    String user_name) throws SQLException  
{  
  
    boolean new_user;  
    int user_id;  
  
    // Get ml_user id.  
    _get_user_id_pstmt.setString(1, user_name);  
  
    ResultSet user_id_rs =  
    _get_user_id_pstmt.executeQuery();  
    new_user = !user_id_rs.next();  
    if (!new_user) {  
        user_id = user_id_rs.getInt(1);  
    }  
}
```



```

else {
    user_id = 0;
}

user_id_rs.close();
user_id_rs = null;

// In this tutorial always allow the login.
authentication_status.setValue(1000);

if (new_user) {
    _insert_login_pstmt.setString(1, user_name);
    _insert_login_pstmt.executeUpdate();
    java.lang.System.out.println("user: " + user_name + " added. ");
}
else {
    _insert_audit_pstmt.setInt(1, user_id);
    _insert_audit_pstmt.setString(2, "LOGIN ALLOWED");
    _insert_audit_pstmt.executeUpdate();
}
_audit_connection.commit();
return;
}

```

The following methods use SQL code to act as cursors on the database tables. Since these are cursor scripts, they must return a SQL string.

```

public static String empUploadInsertStmt() {
    return("INSERT INTO emp(emp_id, emp_name) VALUES(?, ?)");
}

public static String empUploadDeleteStmt() {
    return("DELETE FROM emp WHERE emp_id = ?");
}

public static String empUploadUpdateStmt() {
    return("UPDATE emp SET emp_name = ? WHERE emp_id = ?");
}

public static String empDownloadCursor() {
    return("SELECT emp_id, emp_name FROM emp");
}

public static String custUploadInsertStmt() {
    return("INSERT INTO cust(cust_id, emp_id, cust_name) VALUES (?, ?, ?)");
}

public static String custUploadDeleteStmt() {
    return("DELETE FROM cust WHERE cust_id = ?");
}

public static String custUploadUpdateStmt() {
    return("UPDATE cust SET emp_id = ?, cust_name = ? WHERE cust_id = ?");
}

public static String custDownloadCursor() {
    return("SELECT cust_id, emp_id, cust_name FROM cust");
}

```

Use the following command to compile the code:

```
javac -cp %sqlany11%\java\mlscript.jar CustEmpScripts.java
```

Run the MobiLink server with the location of CustEmpScripts.class in the classpath. The following is a partial command line:

```
mlsrv11 ... -sl java (-cp <class_location>)
```

# MobiLink server API for Java reference

This section explains the MobiLink Java interfaces and classes, and their associated methods and constructors. To use these classes, reference the *mlscript.jar* assembly, located in *install-dir\java\*.

## DBConnectionContext interface

### Syntax

```
public ianywhere.ml.script.DBConnectionContext
```

### Remarks

Interface for obtaining and accessing information about the current database connection. A `DBConnectionContext` instance is passed to the constructor of classes containing scripts. If context is required for a background thread or beyond the lifetime of a connection, use the `ServerContext` class.

### See also

- [“Constructors” on page 532](#)
- [“ServerContext interface” on page 565](#)

### Members

All members of `ianywhere.ml.script.DBConnectionContext`, including all inherited members.

- [“getConnection method” on page 544](#)
- [“getDownloadData method” on page 544](#)
- [“getProperties method” on page 545](#)
- [“getRemoteID method” on page 546](#)
- [“getServerContext method” on page 546](#)
- [“getVersion method” on page 547](#)

### Example

The following example shows you how to create a class level `DBConnectionContext` instance to use in your synchronization scripts. The `DBConnectionContext` `getConnection` method obtains a `java.sql.Connection` instance representing the current connection with the MobiLink consolidated database.

```
import ianywhere.ml.script.*;  
import java.io.*;  
import java.sql.*;  
  
public class OrderProcessor {  
    DBConnectionContext _cc;  
  
    public OrderProcessor(DBConnectionContext cc) {  
        _cc = cc;  
    }  
  
    // The method used for the handle_DownloadData event.  
    public void HandleEvent() throws SQLException {  
        java.sql.Connection my_connection = _cc.getConnection();
```

```
    } // ...  
} // ...
```

**Caution**

A DBConnectionContext instance should not be used outside the thread that calls into your Java code.

## getConnection method

### Syntax

```
public java.sql.Connection getConnection()  
throws java.sql.SQLException
```

### Remarks

Returns the existing connection with the MobiLink consolidated database as a JDBC connection. The java.sql.Connection object returned by this method represents same connection that the MobiLink server uses to execute SQL scripts.

This connection must not be committed, closed or altered in any way that would affect the MobiLink server use of this connection. The connection returned is only valid for the lifetime of the underlying MobiLink connection. Do not use the connection after the end\_connection event has been called for that connection.

If an error occurs binding the existing connection as a JDBC connection then it throws java.sql.SQLException

If a server connection with full access is required, use ServerContext.makeConnection().

### Returns

The existing connection with the MobiLink consolidated database as a JDBC connection.

### Example

See [“DBConnectionContext interface” on page 543](#).

## getDownloadData method

### Syntax

```
public DownloadData getDownloadData()
```

### Remarks

Returns a DownloadData instance for the current synchronization. Use the DownloadData class to create the download for direct row handling.

## Returns

A DownloadData instance for the current synchronization.

## See also

- [“DownloadData interface” on page 548](#)
- [“Direct row handling” on page 649](#)

## Example

The following example shows you how to obtain a DownloadData instance for the current synchronization using the DBConnectionContext getDownloadData method.

### Note

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used for the handle_DownloadData event.
public void HandleDownload() throws SQLException {
    // get the DownloadData for the current synchronization
    DownloadData my_dd = _cc.getDownloadData();

    // ...
}

// ...
```

## getProperties method

### Syntax

```
public java.util.Properties getProperties()
```

### Remarks

Returns the properties for this connection, based on this connection's script version. Properties are stored in the ml\_property table.

Consult your Java SDK documentation for more information about java.util.Properties.

### Returns

The properties for this connection.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

### Example

The following example shows you how to output the properties for a DBConnectionContext.

**Note**

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used to output the connection properties.
public void outputProperties() {
    // output the Properties for the current synchronization
    java.util.Properties properties = _cc.getProperties();
    System.out.println(properties.toString());
}
```

## getRemoteID method

### Syntax

```
public java.lang.String getRemoteID( )
```

### Remarks

Returns the remote ID of the database currently synchronizing on this connection. If your remote database is prior to version 10, it returns the MobiLink user name.

### Returns

The remote ID.

### See also

- “Remote IDs” [[MobiLink - Client Administration](#)]

### Example

The following example shows you how to output the remote ID for a DBConnectionContext.

**Note**

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used to output the remote ID.
public void outputRemoteID() {
    // output the Remote ID for the current synchronization
    String remoteID = _cc.getRemoteID();
    System.out.println(remoteID);
}
```

## getServerContext method

### Syntax

```
public ServerContext getServerContext( )
```

### Remarks

Returns the ServerContext for this MobiLink server.

## Returns

The ServerContext for this MobiLink server.

## See also

- [“ServerContext interface” on page 565](#)

## Example

The following example shows you how to get the ServerContext instance for a DBConnectionContext and shut down the server.

### Note

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// A method that uses an instance of the ServerContext to shut down the
server
public void shutDownServer() {
    ServerContext context = _cc.getServerContext();
    context.shutdown();
}
```

## getVersion method

### Syntax

```
public java.lang.String getVersion( )
```

### Remarks

Returns the script version string.

### Returns

The script version string.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

## Example

The following example shows you how to get the script version and use it to make decisions.

### Note

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// A method that uses the script version
public void handleEvent() {
    // ...

    String version = _cc.getVersion();
    if (version.equals("My Version 1")) {
```

```
        // ...
    } else if (version.equals("My Version 2")) {
        // ...
    }
}
// ...
```

## DownloadData interface

### Syntax

```
public ianywhere.ml.script.DownloadData
```

### Remarks

Encapsulates download data operations for direct row handling. To obtain a DownloadData instance, use the DBConnectionContext `getDownloadData` method.

Use the DownloadData.`getDownloadTables` and `getDownloadTableByName` methods to return DownloadTableData instances.

This download data is available through DBConnectionContext. It is not valid to access the download data before the `begin_synchronization` event or after the `end_download` event. It is not valid to access DownloadData in an upload-only synchronization.

### See also

- [“DownloadTableData interface” on page 550](#)
- [“handle\\_DownloadData connection event” on page 442](#)
- DBConnectionContext [“getDownloadData method” on page 544](#)
- [“Direct row handling” on page 649](#)

### Members

All members of **ianywhere.ml.script.DownloadData**, including all inherited members.

- [“getDownloadTableByName method” on page 549](#)
- [“getDownloadTables method” on page 550](#)

### Example

The following example shows you how to obtain a DownloadData instance for the current synchronization using the DBConnectionContext `getDownloadData` method.

```
DBConnectionContext _cc;

// Your class constructor.
public OrderProcessor(DBConnectionContext cc) {
    _cc = cc;
}

// The method used for the handle_DownloadData event.
public void handleDownload() throws SQLException {
    // Get the DownloadData for the current synchronization.
```



```
        DownloadData my_dd = _cc.getDownloadData();  
    } // ...
```

## getDownloadTableByName method

### Syntax

```
public DownloadTableData getDownloadTableByName(  
    string table-name);
```

### Remarks

Gets the named download table for this synchronization. Returns null if there is no table with the given name in this synchronization.

### Parameters

- **table\_name** The name of the table for which you want the download data.

### Returns

A DownloadTableData instance representing the specified table, or null if a table of the given name does not exist for the current synchronization.

### See also

- [“DownloadData interface” on page 548](#)
- [“DownloadTableData interface” on page 550](#)
- [“DBConnectionContext interface” on page 543](#)
- [“Direct row handling” on page 649](#)

### Example

The following example uses the getDownloadTableByName method to return a DownloadTableData instance for the remoteOrders table.

**Note**

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used for the handle_DownloadData event.  
public void handleDownload() throws SQLException {  
    // Get the DownloadData for the current synchronization.  
    DownloadData my_dd = _cc.getDownloadData();  
  
    // Get the DownloadTableData for the remoteOrders table.  
    DownloadTableData my_download_table =  
    my_dd.getDownloadTableByName("remoteOrders");  
  
    // ...  
}
```

## getDownloadTables method

### Syntax

```
public DownloadTableData[] getDownloadTables( )
```

### Remarks

Gets an array of all the tables for download data in this synchronization. The operations performed on this table are sent to the remote database.

### Returns

An array of DownloadTableData objects for the current synchronization. The order of tables in the array is the same as the upload order of the remote.

### See also

- [“DownloadData interface” on page 548](#)
- [“DownloadTableData interface” on page 550](#)
- [“DBConnectionContext interface” on page 543](#)
- [“Direct row handling” on page 649](#)

### Example

The following example uses the DownloadData.getDownloadTables method to obtain an array of DownloadTableData objects for the current synchronization. The example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used for the handle_DownloadData event.
public void handleDownload() throws SQLException {
    // Get the DownloadData for the current synchronization.
    DownloadData my_dd = _cc.getDownloadData();

    // Get an array of tables to set download operations.
    DownloadTableData[] download_tables = my_dd.getDownloadTables();

    // Get the first table in the DownloadTableData array.
    DownloadTableData my_download_table = download_tables[0];

    // ...
}
```

## DownloadTableData interface

### Syntax

```
public anywhere.ml.script.DownloadTableData
```

### Remarks

Encapsulates table operations for MobiLink direct downloads. Use this interface to set the data operations that are downloaded to the client. To obtain DownloadTableData instances for the current synchronization, use the DownloadData interface. You can use the DownloadTableData.getUpsertPreparedStatement and

`getDeletePreparedStatement` methods to obtain Java prepared statements for insert and update, and delete operations, respectively. The `java.sql.PreparedStatement.executeUpdate` method registers an operation for download.

**Note**

You must set all column values for insert and update prepared statements. For delete operations you set primary key values.

You cannot have both the delete and upsert prepared statements open at the same time.

Consult your Java SDK documentation for more information about `java.sql.PreparedStatement`.

**See also**

- [“DownloadData interface” on page 548](#)
- [“handle\\_DownloadData connection event” on page 442](#)
- [“Direct row handling” on page 649](#)

**Members**

All members of `ianywhere.ml.script.DownloadTableData`, including all inherited members.

- [“getDeletePreparedStatement method” on page 552](#)
- [“getUpsertPreparedStatement method” on page 553](#)
- [“getName method” on page 554](#)
- [“getMetaData method” on page 555](#)
- [“getLastDownloadTime method” on page 556](#)

**Example**

Assume you use a table called `remoteOrders` in MobiLink client databases.

```
CREATE TABLE remoteOrders (
    pk INT NOT NULL,
    coll VARCHAR(200),
    PRIMARY KEY (pk)
);
```

The following example uses the `DownloadData.getDownloadTableByName` method to return a `DownloadTableData` instance representing the `remoteOrders` table.

**Note**

This example assumes you have a `DBConnectionContext` instance called `_cc`.

```
// The method used for the handle_DownloadData event
public void handleDownload() throws SQLException {

    // Get the DownloadData for the current synchronization.
    DownloadData my_dd = _cc.getDownloadData();

    // Get the DownloadTableData for the remoteOrders table.
    DownloadTableData td = my_dd.getDownloadTableByName("remoteOrders");

    // User defined-methods to set download operations.
    setDownloadInserts(td);
```

```
        setDownloadDeletes(td);  
    }  
    // ...  
}
```

In this example, the `setDownloadInserts` method uses the `DownloadTableData.getUpsertPreparedStatement` to obtain a prepared statement for rows you want to insert or update. The `PreparedStatement.setInt` and `PreparedStatement.setString` methods set the column values you want to insert into the remote database.

```
void setDownloadInserts(DownloadTableData td) {  
    java.sql.PreparedStatement insert_ps = td.getUpsertPreparedStatement();  
  
    // The following method calls are the same as the following SQL  
    statement:  
    // INSERT INTO remoteOrders(pk, coll) values(2300, "truck");  
    insert_ps.setInt(1, 2300);  
    insert_ps.setString(2, "truck");  
  
    int update_result = insert_ps.executeUpdate();  
    if (update_result == 0) {  
        // Insert was filtered because it was uploaded  
        // in the same synchronization.  
    }  
    else {  
        // Insert was not filtered.  
    }  
}
```

The `setDownloadDeletes` method uses the `DownloadTableData.getDeletePreparedStatement` to obtain a prepared statement for rows you want to delete. The `java.sql.PreparedStatement.setInt` method sets the primary key values for rows you want to delete in the remote database and the `java.sql.PreparedStatement.executeUpdate` method registers the row values for download.

```
void setDownloadDeletes(DownloadTableData td) {  
    java.sql.PreparedStatement delete_ps = td.getDeletePreparedStatement();  
  
    // The following method calls are the same as the following SQL  
    statement:  
    // DELETE FROM remoteOrders where pk=2300;  
    delete_ps.setInt(1, 2300);  
    delete_ps.executeUpdate();  
}
```

## getDeletePreparedStatement method

### Syntax

```
public java.sql.PreparedStatement getDeletePreparedStatement() throws SQLException
```

### Remarks

Returns a `java.sql.PreparedStatement` instance that allows the user to add delete operations to the download. The prepared statement applies to the download table and contains a parameter for each primary key column in the table.

The prepared statement applies to the `DownloadTableData` instance and contains a parameter for each primary key column in the table.

To include a delete operation in the download, set all columns in your `java.sql.PreparedStatement` and then call the `java.sql.PreparedStatement.executeUpdate` method.

**Note**

You must set all primary key values for download delete operations.

**Returns**

A `java.sql.PreparedStatement` instance for adding delete operations to the download.

**Exceptions**

- **SQLException** Thrown if there is a problem retrieving the delete `java.sql.PreparedStatement` instance.

**See also**

- [“DownloadTableData interface” on page 550](#)
- [“Direct row handling” on page 649](#)

**Example**

In the following example, the `setDownloadDeletes` method uses the `DownloadTableData.getDeletePreparedStatement` to obtain a prepared statement for rows you want to delete. The `java.sql.PreparedStatement.setInt` method sets the primary key values for rows you want to delete in the remote database and the `java.sql.PreparedStatement.executeUpdate` method sets the row values in the download.

```
void setDownloadDeletes(DownloadTableData td) {
    java.sql.PreparedStatement delete_ps = td.getDeletePreparedStatement();
    // This is the same as executing the following SQL statement:
    // DELETE FROM remoteOrders where pk=2300;
    delete_ps.setInt(1, 2300);
    delete_ps.executeUpdate();
    delete_ps.close();
}
```

## getUpsertPreparedStatement method

**Syntax**

```
public java.sql.PreparedStatement getUpsertPreparedStatement( ) throws SQLException
```

**Remarks**

Returns a `java.sql.PreparedStatement` instance which allows the user to add upsert (insert or update) operations to the download of a synchronization. The prepared statement applies to the `DownloadTableData` instance and contains a parameter for each column in the table.

To include an insert or update operation in the download, set all column values in your `java.sql.PreparedStatement` and then call the `java.sql.PreparedStatement.executeUpdate` method. Calling `java.sql.PreparedStatement.executeUpdate` on the prepared statement returns 0 if the insert or update

operation was filtered and returns 1 if the operation was not filtered. An operation is filtered if it was uploaded in the same synchronization.

**Note**

You must set all column values for download insert and update operations.

**Returns**

A `java.sql.PreparedStatement` instance for adding upsert operations to the download.

**Exceptions**

- **SQLException** Thrown if there is a problem retrieving the upsert `java.sql.PreparedStatement` instance.

**See also**

- [“DownloadTableData interface” on page 550](#)
- [“Direct row handling” on page 649](#)

**Example**

In the following example, the `setDownloadInserts` method uses the `DownloadTableData.getUpsertPreparedStatement` to obtain a prepared statement for rows you want to insert or update. The `java.sql.PreparedStatement.setInt` and `PreparedStatement.setString` methods set the column values, and the `PreparedStatement.executeUpdate` method sets the row values in the download.

```
void setDownloadInserts(DownloadTableData td) {
    java.sql.PreparedStatement insert_ps = td.getUpsertPreparedStatement();

    // This is the same as executing the following SQL statement:
    // INSERT INTO remoteOrders(pk, coll) VALUES (2300, "truck");
    insert_ps.setInt(1, 2300);
    insert_ps.setString(2, "truck");

    int update_result = insert_ps.executeUpdate();
    if (update_result == 0) {
        // Insert was filtered because it was uploaded
        // in the same synchronization.
    }
    else {
        // Insert was not filtered.
    }
    insert_ps.close();
}
```

## getName method

**Syntax**

```
public java.lang.String getName()
```

## Remarks

Returns the table name for the DownloadTableData instance. You can also access the table name using the java.sql.ResultSetMetaData instance returned by the DownloadTableData.getMetaData method.

## Returns

The table name for the DownloadTableData instance.

## See also

- [“DownloadTableData interface” on page 550](#)
- [DownloadTableData “getMetaData method” on page 555](#)
- [“Direct row handling” on page 649](#)

## Example

The following example shows you how to output the table name for the DownloadTableData instance.

**Note**

This example assumes you have a DBConnectionContext instance called `_cc`.

```
// The method used for the handle_DownloadData event
public void handleDownload() throws SQLException {

    // Get the DownloadData for the current synchronization.
    DownloadData my_dd = _cc.getDownloadData();

    // Get the DownloadTableData for the remoteOrders table.
    DownloadTableData td = my_dd.getDownloadTableByName("remoteOrders");

    // Print the table name to standard output (remoteOrders)
    System.out.println(td.getName());

    // User defined-methods to set download operations.
    setDownloadInserts(td);
    setDownloadDeletes(td);

    // ...
}
```

## getMetaData method

### Syntax

```
public java.sql.ResultSetMetaData getMetaData()
```

### Remarks

Gets the metadata for the DownloadTableData instance. The metadata is a standard java.sql.ResultSetMetaData object.

If you want the metadata to contain column name information, specify in your client that column names should be sent with the upload.

Consult your Java SDK documentation for more information about java.sql.ResultSetMetaData.

## Returns

The metadata for the DownloadTableData instance.

## See also

- [“DownloadTableData interface” on page 550](#)
- [“Direct row handling” on page 649](#)
- SQL Anywhere clients: [“SendColumnNames \(scn\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Column Names synchronization parameter” \[UltraLite - Database Management and Reference\]](#)

## Example

The following example shows you how get the number of columns used in the query for the DownloadTableData instance.

```
import java.sql.ResultSetMetaData;

// The method used to return the number of columns in a DownloadTableData
instance query
public int getNumColumns(DownloadTableData td) {
    ResultSetMetaData rsmd = td.getMetaData();
    return rsmd.getColumnCount();
}
```

## getLastDownloadTime method

### Syntax

```
public java.sql.Timestamp getLastDownloadTime()
```

### Remarks

Returns last download time for this table. This is the same last download time passed to several of the per table download events.

The last download time is useful for generating the table download data for a particular synchronization.

### Returns

The last download time for this download table.

### See also

- [“DownloadTableData interface” on page 550](#)
- [“Direct row handling” on page 649](#)

### Example

The following example is a snippet of code that populates a table in the download with inserts using the last download time. Note that this example assumes you have a DBConnectionContext instance called `_cc`.



```
// The method used for the handle_DownloadData event
public void handleDownload() throws SQLException {

    // Get the DownloadData for the current synchronization.
    DownloadData my_dd = _cc.getDownloadData();

    // Get the DownloadTableData for the remoteOrders table.
    DownloadTableData td = my_dd.getDownloadTableByName("remoteOrders");

    // get the inserts given a last download time
    ResultSet inserts_rs =
makeInsertsFromTimestamp(td.getLastDownloadTime());

    // fill the DownloadTableData using the inserts resultset.
    setDownloadInsertsFromRS(td, inserts_rs);
    inserts_rs.close();

    // ...
}
```

## InOutInteger interface

### Syntax

```
public ianywhere.ml.script.InOutInteger
```

### Remarks

Passed into methods to enable the functionality of an in/out parameter passed to a SQL script.

### Members

All members of **ianywhere.ml.script.InOutInteger**, including all inherited members.

- [“getValue method” on page 558](#)
- [“setValue method” on page 558](#)

### Example

The following call to a MobiLink system procedure registers a Java method called `handleError` as the script for the `handle_error` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(
    'ver1',
    'handle_error',
    'ExamplePackage.ExampleClass.handleError'
)
```

The following is the sample Java method `handleError`. It processes an error based on the data that is passed in. It also determines the resulting error code.

```
public String handleError(
    ianywhere.ml.script.InOutInteger actionCode,
    int errorCode,
    String errorMessage,
    String user,
    String table)
{
```

```
int new_ac;
if (user == null) {
    new_ac = handleNonSyncError(errorCode, errorMessage);
} else if (table == null) {
    new_ac = handleConnectionError(errorCode, errorMessage, user);
}
else {
    new_ac = handleTableError(errorCode, errorMessage, user, table);
}

// Keep the most serious action code.
if (actionCode.getValue() < new_ac) {
    actionCode.setValue(new_ac);
}
}
```

## getValue method

### Syntax

```
public int getValue( )
```

### Remarks

Returns the value of this integer parameter.

### Returns

The value of this integer parameter.

### Example

See: [“InOutInteger interface” on page 557](#).

## setValue method

### Syntax

```
public void setValue( int new_value )
```

### Remarks

Sets the value of this integer parameter.

### Parameters

- **new\_value** The value for this integer to take.

### Example

See: [“InOutInteger interface” on page 557](#).

## InOutString interface

### Syntax

```
public ianywhere.ml.script.InOutString
```

### Remarks

Passed into methods to enable the functionality of an in/out parameter passed to a SQL script.

### Members

All members of **ianywhere.ml.script.InOutString**, including all inherited members.

- [“getValue method” on page 559](#)
- [“setValue method” on page 560](#)

### Example

The following call to a MobiLink system procedure registers a Java method called `modifyUser` as the script for the `modify_user` connection event when synchronizing the script version `ver1`.

```
CALL ml_add_java_connection_script(
    'ver1',
    'modify_user',
    'ExamplePackage.ExampleClass.modifyUser'
)
```

The following is the sample Java method `modifyUser`. It gets the user ID from the database and then uses it to set the user name.

```
public String modifyUser(InOutString io_user_name) throws SQLException {
    Statement uid_select = curConn.createStatement();
    ResultSet uid_result = uid_select.executeQuery(
        "SELECT rep_id FROM SalesRep WHERE name = '"
        + io_user_name.getValue() + "' "
    );
    uid_result.next();
    io_user_name.setValue( java.lang.Integer.toString(uid_result.getInt(1)) );
    uid_result.close();
    uid_select.close();
    return (null);
}
```

## getValue method

### Syntax

```
public java.lang.String getValue()
```

### Remarks

Returns the value of this string parameter.

### Returns

The value of this string parameter.

### Example

See: [“InOutString interface” on page 559](#)

## setValue method

### Syntax

```
public void setValue( java.lang.String new_value )
```

### Remarks

Sets the value of this String parameter.

### Parameters

- **new\_value** The value for this String to take.

### Example

See: [“InOutString interface” on page 559](#)

## LogListener interface

### Syntax

```
public ianywhere.ml.script.LogListener
```

### Remarks

The listener interface for catching messages that are printed to the log.

### See also

- [“Handling MobiLink server errors in Java” on page 534](#)

### Members

All members of **ianywhere.ml.script.LogListener**, including all inherited members.

- [“messageLogged method” on page 561](#)

### Example

See: [“Handling MobiLink server errors in Java” on page 534](#)

## messageLogged method

### Syntax

```
public void messageLogged(  
    ServerContext sc,  
    LogMessage message )
```

### Remarks

Invoked when a message is printed to the log.

### Parameters

- **sc** The context for the server that is printing the message.
- **message** The LogMessage that has been sent to the MobiLink log.

### Example

See: [“Handling MobiLink server errors in Java” on page 534](#)

## LogMessage class

### Syntax

```
public ianywhere.ml.script.LogMessage
```

### Remarks

Holds the data associated with a log message.

Extends java.lang.Object.

### See also

- [“Handling MobiLink server errors in Java” on page 534](#)

### Members

All members of **ianywhere.ml.script.LogMessage**, including all inherited members.

- [“ERROR variable” on page 563](#)
- [“INFO variable” on page 564](#)
- [“WARNING variable” on page 564](#)
- [“getType method” on page 564](#)
- [“getUser method” on page 565](#)
- [“getText method” on page 565](#)

### Example

The following example installs a LogListener for all warning, error, and info messages, then writes the information to a file. The following code installs a LogListener for all warning messages.

```
class WarningLogListener implements LogListener {
    FileOutputStream _outFile;

    public WarningLogListener(FileOutputStream outFile) {
        _outFile = outFile;
    }

    public void messageLogged(ServerContext sc, LogMessage msg) {
        String user;

        try {
            if (msg.getType() != LogMessage.WARNING) {
                //this class deals exclusively with warnings
                return;
            }
            user = msg.getUser();

            if (user == null) {
                user = "NULL";
            }

            _outFile.write(("Caught warning"
                + " user=" + user
                + " text=" + msg.getText()
                + "\n").getBytes());
            _outFile.flush();
        }
        catch(Exception e) {
            // Print some error output to the MobiLink log.
            e.printStackTrace();
        }
    }
}
```

The following code installs a LogListener for all error messages.

```
class ErrorLogListener implements LogListener {
    FileOutputStream _outFile;

    public ErrorLogListener(FileOutputStream outFile) {
        _outFile = outFile;
    }

    public void messageLogged(ServerContext sc, LogMessage msg) {
        String user;

        try {
            if (msg.getType() != LogMessage.ERROR) {
                //this class deals exclusively with errors
                return;
            }

            user = msg.getUser();
            if (user == null) {
                user = "NULL";
            }

            _outFile.write(("Caught error"
                + " user=" + user
                + " text=" + msg.getText()
                + "\n").getBytes());
            _outFile.flush();
        }
    }
}
```

```

    }
    catch(Exception e) {
        // Print some error output to the MobiLink log.
        e.printStackTrace();
    }
}

```

The following code installs a `LogListener` for all info messages.

```

class InfoLogListener implements LogListener {
    FileOutputStream _outFile;

    public InfoLogListener(FileOutputStream outFile) {
        _outFile = outFile;
    }

    public void messageLogged(ServerContext sc, LogMessage msg) {
        String user;

        try {
            if (msg.getType() != LogMessage.ERROR) {
                // this class deals exclusively with info
                return;
            }

            user = msg.getUser();
            if (user == null) {
                user = "NULL";
            }

            _outFile.write(("Caught info"
                + " user=" + user
                + " text=" + msg.getText()
                + "\n").getBytes());
            _outFile.flush();
        }
        catch(Exception e) {
            // Print some error output to the MobiLink log.
            e.printStackTrace();
        }
    }
}

```

The following code registers `WarningLogListener`, `ErrorLogListener`, and `InfoLogListener` to receive warning, error, and info messages respectively. Call this code from anywhere that has access to the `ServerContext` such as a class constructor or synchronization script.

```

// ServerContext serv_context;
// FileOutputStream outFile
serv_context.addWarningListener(new WarningLogListener(outFile));
serv_context.addErrorListener(new ErrorLogListener(outFile));
serv_context.addInfoListener(new InfoLogListener(outFile));

```

## ERROR variable

### Syntax

```
int ERROR
```

### Remarks

The log message is an error.

### Example

See: [“LogMessage class” on page 561.](#)

## INFO variable

### Syntax

int **INFO**

### Remarks

The message log displays information.

### Example

See: [“addInfoListener method” on page 566.](#)

## WARNING variable

### Syntax

int **WARNING**

### Remarks

The log message is a warning.

### Example

See: [“LogMessage class” on page 561.](#)

## getType method

### Syntax

public int **getType()**

### Remarks

Accessor for this message type.

### Returns

The type of this message, which can be either `LogMessage.ERROR`, `LogMessage.INFO`, or `LogMessage.WARNING`.



**Example**

See: [“LogMessage class” on page 561](#).

## getUser method

**Syntax**

```
public java.lang.String getUser()
```

**Remarks**

Accessor for this message user. If the message has no user, then the user is null.

**Returns**

The user associated with this message.

**Example**

See: [“LogMessage class” on page 561](#).

## getText method

**Syntax**

```
public java.lang.String getText()
```

**Remarks**

Accessor for the message text.

**Returns**

The main text of this message.

**Example**

See: [“LogMessage class” on page 561](#).

## ServerContext interface

**Syntax**

```
public ianywhere.ml.script.ServerContext
```

**Remarks**

An instantiation of all the context that is present for the duration of the MobiLink server. This context can be held as static data and used in a background thread. It is valid for the duration of the Java virtual machine invoked by MobiLink.

To access a `ServerContext` instance, use the `DBConnectionContext.getServerContext` method.

### Members

All members of `ianywhere.ml.script.ServerContext`, including all inherited members.

- [“addInfoListener method” on page 566](#)
- [“addErrorListener method” on page 567](#)
- [“addShutdownListener method” on page 568](#)
- [“addWarningListener method” on page 568](#)
- [“getProperties method” on page 569](#)
- [“getPropertiesByVersion method” on page 569](#)
- [“getPropertySetNames method” on page 570](#)
- [“getStartClassInstances method” on page 571](#)
- [“makeConnection method” on page 571](#)
- [“removeErrorListener method” on page 572](#)
- [“removeInfoListener method” on page 572](#)
- [“removeShutdownListener method” on page 572](#)
- [“removeWarningListener method” on page 573](#)
- [“shutdown method” on page 573](#)

## addInfoListener method

### Syntax

```
public void addInfoListener( LogListener // )
```

### Remarks

Adds the specified `LogListener` from the list of listeners to receive a notification when info is printed. The method `LogListener.messageLogged` (`ianywhere.ml.script.ServerContext`) is called.

### Parameters

- **//** The `LogListener` to be notify.

### Example

The following code registers a listener of type `MyLogListener` to receive notifications of info messages.

```
// ServerContext serv_context;  
serv_context.addInfoListener(new MyLogListener(ll_out_file));  
  
// The following code shows an example of processing those messages:  
class MyLogListener implements LogListener {  
    FileOutputStream _out_file;  
    public TestLogListener(FileOutputStream out_file) {  
        _out_file = out_file;  
    }  
  
    public void messageLogged(ServerContext sc, LogMessage msg) {  
        String type;  
        String user;
```

```

try {
    if (msg.getType() == LogMessage.ERROR) {
        type = "ERROR";
    } else if (msg.getType() == LogMessage.WARNING) {
        type = "WARNING";
    } else if (msg.getType() == LogMessage.INFO) {
        type = "INFO";
    } else {
        type = "UNKNOWN!!!";
    }

    user = msg.getUser();
    if (user == null) {
        user = "NULL";
    }
    _out_file.write(("Caught msg type="
        + type
        + " user=" + user
        + " text=" +msg.getText()
        + "\n").getBytes()
    );
    _out_file.flush();
}
catch(Exception e) {

    // if we print the exception from processing an info message,
    // we may recurse indefinitely
    if (msg.getType() != LogMessage.ERROR) {
        // Print some error output to the MobiLink log.
        e.printStackTrace();
    }
}
}
}
}
}

```

## addErrorListener method

### Syntax

```
public void addErrorListener( LogListener ll )
```

### Remarks

Adds the specified LogListener to receive a notification when an error is printed.

When an error is printed, the following method is called

```
LogListener.messageLogged(ianywhere.ml.script.ServerContext, ianywhere.ml.script.LogMessage).
```

### Parameters

- ll The LogListener to notify.

### See also

- [“messageLogged method” on page 561](#)

## Example

See: [“LogMessage class” on page 561](#).

## addShutdownListener method

### Syntax

```
public void addShutdownListener( ShutdownListener sl )
```

### Remarks

Adds the specified ShutdownListener that is to receive notification before the server context is destroyed. On shutdown, the method ShutdownListener.shutdownPerformed (iAnywhere.ml.script.ServerContext) is called.

### Parameters

- **sl** The ShutdownListener to notify on shutdown.

### Example

See: [“ShutdownListener interface” on page 575](#).

## addWarningListener method

### Syntax

```
public void addWarningListener( LogListener ll )
```

### Remarks

Adds the specified LogListener to receive a notification when a warning is printed.

The following method is called: LogListener.messageLogged(iAnywhere.ml.script.ServerContext, iAnywhere.ml.script.LogMessage).

### Parameters

- **ll** The LogListener to notify.

### Example

See: [“LogMessage class” on page 561](#).

## getProperties method

### Syntax

```
public java.util.Properties getProperties(  
    java.lang.String component,  
    java.lang.String set)
```

### Remarks

Returns the set of properties associated with a given component and property set. These are stored in the MobiLink system table `ml_property`.

### Parameters

- **component** The component.
- **set** The property set.

### Returns

The set of properties, which may be empty.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

### Example

The following code lists all the ServerContext's Properties.

```
import java.util.*;  
// ServerContext serverContext;  
// PrintStream out  
Properties prop = serverContext.getProperties();  
prop.list(out);
```

## getPropertiesByVersion method

### Syntax

```
public java.util.Properties getPropertiesByVersion( java.lang.String script_version )
```

### Remarks

Returns the set of properties associated with the script version. These are stored in the MobiLink system table `ml_property`. The script version is stored in the `property_set_name` column when the `component_name` is `ScriptVersion`.

### Parameters

- **script\_version** The script version for which to return associated properties.

## Returns

The set of properties associated with the given script version.

## See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

## Example

The following code lists all the ServerContext's Properties associated with a given script version.

```
import java.util.*;
// ServerContext serverContext;
// PrintStream out
Properties prop = serverContext.getPropertiesByVersion("MyScriptVersion");
prop.list(out);
```

## getPropertySetNames method

### Syntax

```
public Iterator getPropertySetNames(
    java.lang.String component_name )
```

### Remarks

Returns the list of property set names for a given component. These are stored in the MobiLink system table ml\_property.

### Parameters

- **component\_name** The name of the component for which to list property names.

### Returns

The list of property set names for the given component.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

### Example

The following code lists all the ServerContext's Properties associated with a given component.

```
import java.util.*;
// ServerContext serverContext;
// PrintStream out
Properties prop = serverContext.getPropertySetNames("Component Name");
prop.list(out);
```

## getStartClassInstances method

### Syntax

```
public java.lang.Object[] getStartClassInstances()
```

Gets an array of the start classes that were constructed at server start time. The array length is zero if there are no start classes.

### Returns

An array of start classes that were constructed at server start time, or an array of length zero if there are no start classes.

### Example

The following is an example that uses `getStartClassInstances()`:

```
Object objs[] = sc.getStartClassInstances();
int i;
for (i=0; i < objs.length; i += 1) {
    if (objs[i] instanceof MyClass) {
        // Use class.
    }
}
```

### See also

- [“User-defined start classes” on page 536](#)

## makeConnection method

### Syntax

```
public java.sql.Connection makeConnection()
throws java.sql.SQLException
```

### Remarks

Creates a new server connection. If an error occurs when opening a new connection, the method throws `java.sql.SQLException`.

### Returns

The new server connection.

### Exceptions

- **SQLException** Thrown if an error occurred opening the new connection.

## removeErrorListener method

### Syntax

```
public void removeErrorListener( LogListener //)
```

### Remarks

Removes the specified LogListener from the list of listeners that are to receive a notification when an error is printed.

### Parameters

- **ll** The LogListener to no longer notify.

### Example

The following code removes a LogListener from the list of error listeners.

```
// ServerContext serverContext;  
// LogListener myErrorListener  
serverContext.removeErrorListener(myErrorListener);
```

## removeInfoListener method

### Syntax

```
public void removeInfoListener( LogListener s/)
```

### Remarks

Remove the specified LogListener from the list of listeners to receive a notification when info is printed.

### Parameters

- **ll** The listener to no longer notify.

### Example

The following code removes a LogListener from the list of info listeners.

```
// ServerContext serverContext;  
// LogListener myInfoListener  
serverContext.removeInfoListener(myInfoListener);
```

## removeShutdownListener method

### Syntax

```
public void removeShutdownListener( ShutdownListener s/)
```



**Remarks**

Removes the specified ShutdownListener from the list of listeners that are to receive notification before the server context is destroyed.

**Parameters**

- **sl** The ShutdownListener to no longer notify on shutdown.

**Example**

The following code removes a ShutdownListener from the list of listeners that are to receive notification before the server context is destroyed.

```
// ServerContext serverContext;  
// ShutdownListener myShutdownListener  
serverContext.removeShutdownListener(myShutdownListener);
```

## removeWarningListener method

**Syntax**

```
public void removeWarningListener( LogListener // )
```

**Remarks**

Removes the specified LogListener from the list of listeners that are to receive a notification when a warning is printed.

**Parameters**

- **ll** The LogListener to no longer notify.

**Example**

The following code removes a LogListener from the list of warning listeners.

```
// ServerContext serverContext;  
// LogListener myWarningListener  
serverContext.removeWarningListener(myWarningListener);
```

## shutdown method

**Syntax**

```
public void shutdown( )
```

**Remarks**

Forces the server to shut down. Any registered ShutdownListener instances have their shutdownPerformed method called.

## Example

The following code forces the server to shut down.

```
// ServerContext serverContext;  
serverContext.shutdown();
```

## ServerException class

### Syntax

```
public ianywhere.ml.script.ServerException
```

### Remarks

Thrown to indicate that there is an error condition that makes any further synchronization on the server impossible. Throwing this exception causes the MobiLink server to shut down.

### Members

All members of **ianywhere.ml.script.ServerException**, including all inherited members.

- [“ServerException constructors” on page 575](#)

### Example

The following code is a function that can throw a `ServerException` if a fatal problem occurs, which causes the MobiLink server to shut down.

```
public void handleUpload(UploadData ud)  
    throws SQLException, IOException, ServerException  
{  
  
    UploadedTableData tables[] = ud.getUploadedTables();  
    if (tables == null) {  
        throw new ServerException("Failed to read uploaded tables");  
    }  
  
    for (int i = 0; i < tables.length; i++) {  
        UploadedTableData currentTable = tables[i];  
        println("table " + java.lang.Integer.toString(i)  
            + " name: " + currentTable.getName());  
  
        // Print out delete result set.  
        println("Deletes");  
        printRSInfo(currentTable.getDeletes());  
  
        // Print out insert result set.  
        println("Inserts");  
        printRSInfo(currentTable.getInserts());  
  
        // print out update result set  
        println("Updates");  
        printUpdateRSInfo(currentTable.getUpdates());  
    }  
}
```

## ServerException constructors

### Syntax

```
public ServerException( )
```

### Remarks

Constructs a `ServerException` with no detail message.

### Syntax

```
public ServerException( java.lang.String s )
```

### Remarks

Constructs a `ServerException` with a specified detail message.

### Parameters

- **s** The detailed message.

### Example

See: [“ServerException class” on page 574](#).

## ShutdownListener interface

### Syntax

```
public ianywhere.ml.script.ShutdownListener
```

### Remarks

The listener interface for catching server shutdowns. Use this interface to ensure that all threads, connections, and other resources are cleaned up before the server exits

### Members

All members of `ianywhere.ml.script.ShutdownListener`, including all inherited members.

- [“shutdownPerformed method” on page 576](#)

### Example

The following code installs a `ShutdownListener` for the `ServerContext`.

```
class MyShutdownListener implements ShutdownListener {
    FileOutputStream _outFile;
    public MyShutdownListener(FileOutputStream outFile) {
        _outFile = outFile;
    }

    public void shutdownPerformed(ServerContext sc) {
        // Add shutdown code
        try {
```

```
        _outFile.write(("Shutting Down" + "\n").getBytes());
        _outFile.flush();
    }
    catch(Exception e) {
        // Print some error output to the MobiLink log.
        e.printStackTrace();
    }
    // ...
}
}
```

The following code registers `MyShutdownListener`. Call this code from anywhere that has access to the `ServerContext` such as a class constructor or synchronization script.

```
// ServerContext serv_context;
// FileOutputStream outFile
serv_context.addShutdownListener(new MyShutdownListener(outFile));
```

## shutdownPerformed method

### Syntax

```
public void shutdownPerformed( ServerContext sc)
```

### Remarks

Invoked before the `ServerContext` is destroyed due to server shutdown.

### Parameters

- **sc** The context for the server that is being shut down.

### Example

See: [“ShutdownListener interface” on page 575](#).

## SynchronizationException class

### Syntax

```
public ianywhere.ml.script.SynchronizationException
```

### Remarks

Thrown to indicate that there is an error condition that makes the completion of the current synchronization impossible. Throwing this exception forces the MobiLink server to rollback.

### Members

All members of `ianywhere.ml.script.SynchronizationException`, including all inherited members.

- [“SynchronizationException constructors” on page 577](#)

**Example**

The following code is a function that can throw a `SynchronizationException` if a problem occurs, which causes the MobiLink server to rollback.

```
public void handleUpload(UploadData ud)
    throws SQLException, IOException, SynchronizationException
{
    UploadedTableData tables[] = ud.getUploadedTables();

    for (int i = 0; i < tables.length; i++) {
        UploadedTableData currentTable = tables[i];
        println("table " + java.lang.Integer.toString(i)
            + " name: " + currentTable.getName());

        // Print out delete result set.
        println("Deletes");
        printRSInfo(currentTable.getDeletes());

        // Print out insert result set.
        println("Inserts");
        printRSInfo(currentTable.getInserts());

        // print out update result set
        println("Updates");
        printUpdateRSInfo(currentTable.getUpdates());

        if (/* Reason for Sync failure */) {
            throw new SynchronizationException("Sync Failed");
        }
    }
}
```

## SynchronizationException constructors

**Syntax**

```
public SynchronizationException( )
```

**Remarks**

Constructs a `SynchronizationException` with no detail message.

**Syntax**

```
public SynchronizationException( java.lang.String s )
```

**Remarks**

Constructs a `SynchronizationException` with the specified detail message.

**Parameters**

- **s** A detail message.

## UpdateResultSet

### Syntax

```
public ianywhere.ml.script.UpdateResultSet
```

### Remarks

A result set object including special methods for accessing the pre-image (old) and post-image (new) values of a specified row. To obtain an UpdateResultSet instance, use the DownloadTableData.getUpdates method.

UpdateResultSet extends java.sql.ResultSet and adds the setNewRowValues and setOldRowValues methods. Otherwise it can be used as a regular resultset. Consult your Java SDK documentation for more information about java.sql.ResultSet

### See also

- DownloadTableData [“getUpdates method” on page 584](#)
- [“Handling conflicts for direct uploads” on page 655](#)
- [“Direct row handling” on page 649](#)

### Members

All members of **ianywhere.ml.script.UpdateResultSet**, including all inherited members.

- [“setNewRowValues method” on page 578](#)
- [“setOldRowValues method” on page 579](#)

### Example

The following code shows how to obtain an UpdateResultSet instance from a DownloadTableData instance.

```
// DownloadTableData tableData  
UpdateResultSet results = tableData.getUpdates();
```

## setNewRowValues method

### Syntax

```
public void setNewRowValues( )
```

### Remarks

Sets the mode of this result set to return new column values ( the post update row). The result set represents the latest updated values in the remote client database. This is the default mode.

### See also

- [“UpdateResultSet” on page 578](#)
- [“Handling conflicts for direct uploads” on page 655](#)
- [“Direct row handling” on page 649](#)

## Example

The following code shows how to set the mode of the UpdateResultSet to return new column values.

```
// UpdateResultSet results
results.setNewRowValues();
```

## setOldRowValues method

### Syntax

```
public void setOldRowValues()
```

### Remarks

Sets the mode of this result set to return old column values (the pre update row). In this mode, the UpdateResultSet represents old column values obtained by the client in the last synchronization.

### See also

- [“UpdateResultSet” on page 578](#)
- [“Handling conflicts for direct uploads” on page 655](#)
- [“Direct row handling” on page 649](#)

## Example

The following code shows how to set the mode of the UpdateResultSet to return old column values.

```
// UpdateResultSet results
results.setOldRowValues();
```

## UploadData interface

### Syntax

```
public ianywhere.ml.script.UploadData
```

### Remarks

Encapsulates upload operations for direct row handling. An UploadData instance representing a single upload transaction is passed to the handle\_UploadData synchronization event.

#### Caution

You must handle direct row handling upload operations in the method registered for the handle\_UploadData event. The UploadData is destroyed after each call to the registered method. Do not create a new instance of UploadData to use in subsequent events.

Use the UploadData.getUploadedTables or UploadData.getUploadedTableByName methods to obtain UploadedTableData instances.

A synchronization has one UploadData unless the remote database is using transactional upload.

### See also

- [“UploadedTableData interface” on page 581](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Direct row handling” on page 649](#)
- [“Handling direct uploads” on page 654](#)

### Members

All members of **ianywhere.ml.script.UploadData**, including all inherited members.

- [“getUploadedTableByName method” on page 580](#)
- [“getUploadedTables method” on page 581](#)

### Example

See [“handle\\_UploadData connection event” on page 454](#).

## getUploadedTableByName method

### Syntax

```
public UploadedTableData getUploadedTableByName(  
    java.lang.String table_name  
)
```

### Remarks

Returns a `UploadedTableData` instance representing the specified table.

### Parameters

- **table\_name** The name of the uploaded table for which you want the uploaded data.

### Returns

An `UploadedTableData` instance representing the specified table, or null if a table of the given name does not exist for the current synchronization.

### See also

- [“UploadData interface” on page 579](#)
- [“UploadedTableData interface” on page 581](#)
- [“Direct row handling” on page 649](#)

### Example

Assume you use a method called `HandleUpload` for the `handle_UploadData` synchronization event. The following example uses the `getUploadedTableByName` method to return an `UploadedTableData` instance for the `remoteOrders` table.

```
// The method used for the handle_UploadData event.  
  
public void handleUpload(UploadData ut)  
    throws SQLException, IOException
```



```

{
    UploadedTableData uploaded_t1 =
    ut.getUploadedTableByName("remoteOrders");
    // ...
}

```

## getUploadedTables method

### Syntax

```
public UploadedTableData[] getUploadedTables()
```

### Remarks

Returns an array of UploadedTableData objects for the current synchronization. The order to the tables in the array is the same order that MobiLink uses for SQL row handling, and is the optimal order for preventing referential integrity violations. If your data source is a relational database, use this table order.

### Returns

An array of UploadedTableData objects for the current synchronization. The order of tables in the array is the same as the upload order of the client.

### See also

- [“UploadData interface” on page 579](#)
- [“UploadedTableData interface” on page 581](#)
- [“Direct row handling” on page 649](#)

### Example

Assume you use a method called HandleUpload for the handle\_UploadData synchronization event. The following example uses the getUploadedTables method to return UploadedTableData instances for the current upload transaction.

```

// The method used for the handle_UploadData event.

public void handleUpload(UploadData ud)
    throws SQLException, IOException
{
    UploadedTableData tables[] = ud.getUploadedTables();
    //...
}

```

## UploadedTableData interface

### Syntax

```
public ianywhere.ml.script.UploadedTableData
```

## Remarks

Encapsulates table operations for direct row handling uploads. You can use an `UploadedTableData` instance to obtain a table's insert, update, and delete operations for a single upload transaction. Use the `UploadedTableData.getInserts`, `UploadedTableData.getUpdates`, and `UploadedTableData.getDeletes` methods to return standard JDBC `java.sql.ResultSet` objects.

Consult your Java SDK documentation for more information about `java.sql.ResultSet` and `java.sql.ResultSetMetaData`.

Table metadata can be accessed using the `UploadedTableData.getMetaData` method or the result sets returned by `getInserts`, `getUpdates`, and `getDeletes`. The delete result set only includes primary key columns for a table.

## See also

- [“UploadData interface” on page 579](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Direct row handling” on page 649](#)

## Members

All members of `ianywhere.ml.script.UploadedTableData`, including all inherited members.

- [“getDeletes method” on page 582](#)
- [“getInserts method” on page 583](#)
- [“getUpdates method” on page 584](#)
- [“getName method” on page 585](#)
- [“getMetaData method” on page 586](#)

## Example

See: [“UploadData interface” on page 579](#).

## getDeletes method

### Syntax

```
public java.sql.ResultSet getDeletes()
```

### Remarks

Returns a `java.sql.ResultSet` object representing delete operations uploaded by a MobiLink client. The result set contains primary key values for deleted rows.

### Returns

A `java.sql.ResultSet` object that represents delete operations uploaded by a MobiLink client.

**See also**

- [“UploadedTableData interface” on page 581](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Direct row handling” on page 649](#)

**Example**

Assume your remote client contains a table called remoteOrders. The following example uses the DownloadTableData.getDeletes method to obtain a result set of deleted rows. In this case, the delete result set includes a single primary key column.

```
import ianywhere.ml.script.*;
import java.sql.*;

// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut)
    throws SQLException, IOException
{
    // Get an UploadedTableData for the remoteOrders table.
    UploadedTableData remoteOrdersTable =
    ut.getUploadedTableByName("remoteOrders");

    // Get deletes uploaded by the MobiLink client.
    java.sql.ResultSet delete_rs = remoteOrdersTable.getDeletes();

    while (delete_rs.next()) {
        // Get primary key values for deleted rows.
        int deleted_id = delete_rs.getInt(1);

        // ...
    }
    delete_rs.close();
}
```

## getInserts method

**Syntax**

```
public java.sql.ResultSet getInserts( )
```

**Remarks**

Returns a java.sql.ResultSet object representing insert operations uploaded by a MobiLink client. Each insert is represented by one row in the result set.

**Returns**

A java.sql.ResultSet object representing insert operations uploaded by a MobiLink client.

**See also**

- [“UploadedTableData interface” on page 581](#)
- [“Direct row handling” on page 649](#)

## Example

Assume your remote client contains a table called remoteOrders. The following example uses the DownloadTableData.getInserts method to obtain a result set of inserted rows. The code obtains the order amount for each row in the current upload transaction.

```
import ianywhere.ml.script.*;
import java.sql.*;

// The method used for the handle_UploadData event
public void HandleUpload(UploadData ut)
    throws SQLException, IOException
{
    // Get an UploadedTableData instance representing the remoteOrders table.
    UploadedTableData remoteOrdersTable =
    ut.getUploadedTableByName("remoteOrders");

    // Get inserts uploaded by the MobiLink client.
    java.sql.ResultSet rs = remoteOrdersTable.getInserts();
    while (rs.next()) {
        // get the uploaded order_amount
        double order_amount = rs.getDouble("order_amount");

        // ...
    }
    rs.close();
}
```

## getUpdates method

### Syntax

```
public UpdateResultSet getUpdates()
```

### Remarks

Returns a UpdateResultSet object representing update operations uploaded by a MobiLink client. Each update is represented by one row including all column values. UpdateResultSet extends java.sql.ResultSet to include special methods for MobiLink conflict detection.

### Returns

An UpdateResultSet object representing update operations uploaded by a MobiLink client.

### See also

- [“UploadedTableData interface” on page 581](#)
- [“UpdateResultSet” on page 578](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Handling conflicts for direct uploads” on page 655](#)
- [“Direct row handling” on page 649](#)

## Example

Assume your remote client contains a table called remoteOrders. The following example uses the `UploadedTableData.getUpdates` method to obtain a result set of updated rows. The code obtains the order amount for each row.

```
import ianywhere.ml.script.*;
import java.sql.*;

// the method used for the handle_UploadData event
public void HandleUpload(UploadData ut)
    throws SQLException, IOException
{
    // Get an UploadedTableData instance representing the remoteOrders table.
    UploadedTableData remoteOrdersTable =
    ut.getUploadedTableByName("remoteOrders");

    // Get inserts uploaded by the MobiLink client.
    java.sql.ResultSet rs = remoteOrdersTable.getUpdates();
    while (rs.next()) {
        // Get the uploaded order_amount.
        double order_amount = rs.getDouble("order_amount");

        // ...
    }
    rs.close();
}
```

## getName method

### Syntax

```
public java.lang.String getName( )
```

### Remarks

Returns the table name for the `UploadedTableData` instance. You can also access the table name using the `java.sql.ResultSetMetaData` instance returned by the `getMetaData` method.

### Returns

The table name for the `UploadedTableData` instance.

### See also

- [“UploadedTableData interface” on page 581](#)
- [UploadedTableData “getMetaData method” on page 586](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Direct row handling” on page 649](#)

## Example

The following example obtains the name of each uploaded table in a single upload transaction.

```
import ianywhere.ml.script.*;
import java.sql.*;

// The method used for the handle_UploadData event.
```

```
public void HandleUpload(UploadData ud) {
    throws SQLException, IOException
    {
        int i;

        // Get UploadedTableData instances.
        UploadedTableData tables[] = ud.getUploadedTables();

        for (i=0; i<tables.length; i+=1) {
            // Get the table name.
            String table_name = tables[i].getName();

            // ...
        }
    }
}
```

## getMetaData method

### Syntax

```
public java.sql.ResultSetMetaData getMetaData()
```

### Remarks

Gets the metadata for the UploadedTableData instance. The metadata is a standard java.sql.ResultSetMetaData instance.

If you want the ResultSetMetaData to contain column name information, you must specify the client option to send column names.

Consult your Java SDK documentation for more information about java.sql.ResultSetMetaData.

### Returns

The metadata for the UploadedTableData instance.

### See also

- dbmlsync: “SendColumnNames (scn) extended option” [[MobiLink - Client Administration](#)]
- UltraLite: “Send Column Names synchronization parameter” [[UltraLite - Database Management and Reference](#)]

### Example

The following example obtains a java.sql.ResultSetMetaData instance for an uploaded table called remoteOrders. The code uses the ResultSetMetaData.getColumnCount and getColumnLabel methods to compile a list of column names.

```
import ianywhere.ml.script.*;
import java.sql.*;

// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut) {
    throws SQLException, IOException
    {
        // Get an UploadedTableData instance representing the remoteOrders table.
        UploadedTableData remoteOrdersTable =
```

```
ut.getUploadedTableByName("remoteOrders");

// get inserts uploaded by the MobiLink client
java.sql.ResultSet rs = remoteOrdersTable.getInserts();

// Obtain the result set metadata.
java.sql.ResultSetMetaData md = rs.getMetaData();
String columnHeading = "";

// Compile a list of column names.
for (int c=1; c <= md.getColumnCount(); c += 1) {
    columnHeading += md.getColumnLabel();

    if (c < md.getColumnCount()) {
        columnHeading += ", ";
    }
}
//...
```

In this case, a method called `HandleUpload` handles the `handle_UploadData` synchronization event.

**See also**

- [“UploadedTableData interface” on page 581](#)
- [“Direct row handling” on page 649](#)
- [“SendColumnNames \(scn\) extended option” \[MobiLink - Client Administration\]](#)
- [“handle\\_UploadData connection event” on page 454](#)

---



---

# Writing synchronization scripts in .NET

## Contents

Introduction to .NET synchronization logic .....	590
Setting up .NET synchronization logic .....	591
Writing .NET synchronization logic .....	593
.NET synchronization techniques .....	600
Loading shared assemblies .....	601
.NET synchronization example .....	604
MobiLink server API for .NET reference .....	606

---

## Introduction to .NET synchronization logic

MobiLink supports the Visual Studio programming languages for writing synchronization scripts. To write MobiLink scripts in .NET, you can use any language that lets you create valid .NET assemblies. In particular, the following languages are tested and documented:

- C#
- Visual Basic .NET
- C++

.NET synchronization logic can function just as SQL logic functions: the MobiLink server can make calls to .NET methods on the occurrence of MobiLink events just as it accesses SQL scripts on the occurrence of MobiLink events. A .NET method can return a SQL string to MobiLink.

This section tells you how to set up, develop, and run .NET synchronization logic for C#, Visual Basic .NET, and C++. It includes a sample application and the MobiLink server API for .NET Reference.

### See also

- [“Tutorial: Using .NET synchronization logic”](#) [*MobiLink - Getting Started*]
- [“Options for writing server-side synchronization logic”](#) [*MobiLink - Getting Started*]
- [“Writing synchronization scripts”](#) on page 313

## Setting up .NET synchronization logic

When you implement synchronization scripts in .NET, you must tell MobiLink where to find the packages, classes, and methods that are contained in your assemblies.

### To implement synchronization scripts in .NET

1. Create your own class or classes. Write a method for each required synchronization event. These methods must be public.

For more information about methods, see [“Methods” on page 595](#).

Each class with non-static methods should have a public constructor. The MobiLink server automatically instantiates each class the first time a method in that class is called for a connection.

See [“Constructors” on page 594](#).

2. Create one or more assemblies. While compiling, reference *iAnywhere.MobiLink.Script.dll*, which contains a repository of MobiLink server API classes to use in your own .NET methods. *iAnywhere.MobiLink.Script.dll* is located in *install-dir\Assembly\v2*.

You can compile your class on the command line, or using Visual Studio or another .NET development environment.

See [“MobiLink server API for .NET reference” on page 606](#).

3. Compile your project.

For example, compile from Visual Studio as follows:

- a. On the **VS.NET Project** menu, choose **Add Existing Item**.
- b. Locate *iAnywhere.MobiLink.Script.dll*.  
In the **Open** list, choose **Link File**.

#### Note

For Visual Studio, always use the Link File method. Do not use the Add Reference option to reference *iAnywhere.MobiLink.Script.dll*. The Add Reference option duplicates *iAnywhere.MobiLink.Script.dll* in the same physical directory as your class assembly, creating problems for the MobiLink server.

- c. Use the **Build** menu to build your assembly.

You can also compile from the command line, as follows:

Replace *dll-path* with the path to *iAnywhere.MobiLink.Script.dll*. for example, in C#:

```
csc /out:dll-pathout.dll /target:library /reference:dll-  
pathiAnywhere.MobiLink.Script.dll sync_v1.cs
```

4. In the MobiLink system tables in your consolidated database, specify the name of the package, class, and method to call for each synchronization script. No more than one class is permitted per script version.

For example, you can add this information to the MobiLink system tables using the `ml_add_dnet_connection_script` stored procedure or the `ml_add_dnet_table_script` stored procedure.

The following SQL statement, when run in a SQL Anywhere database, specifies that `myNamespace.myClass.myMethod` should be run whenever the `authenticate_user` connection-level event occurs.

```
CALL ml_add_dnet_connection_script(  
    'version1',  
    'authenticate_user',  
    'myNamespace.myClass.myMethod'  
)
```

**Note**

The fully qualified method name is case sensitive.

As a result of this procedure call, the `script_language` column of the `ml_script` system table contains the word **dnet**. The `script` column contains the qualified name of a public .NET method.

See [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#) and [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#).

You can also add this information using Sybase Central.

See [“Adding and deleting scripts” on page 327](#).

5. Instruct the MobiLink server to load assemblies and start the CLR. You tell MobiLink where to locate these assemblies using options in the `mlsrv11` command line. There are two options to choose from:
  - **Use `-sl dnet (-MLAutoLoadPath)`** This sets the given path to the application base directory and loads all the private assemblies within it. You should use this option in most cases. For example, to load all assemblies located in *dll-path*, enter:

```
mlsrv11 -c "dsn=consolidated1" -sl dnet(-MLAutoLoadPath=dll-path)
```

When you use the `-MLAutoLoadPath` option you cannot specify a domain when entering the fully qualified method name for the event script.

See [“Loading assemblies” on page 601](#) and [“-sl dnet option” on page 90](#).

- **Use `-sl dnet (-MLDomConfigFile)`** This option requires a configuration file that contains domain and assembly settings. You should use this option when you have shared assemblies, when you don't want to load all the assemblies in a directory, or when for some other reason you need to use a configuration file.

For more information about loading shared assemblies, see [“Loading assemblies” on page 601](#).

For more information about the `mlsrv11` option `-sl dnet`, see [“-sl dnet option” on page 90](#).

**Note**

You can use the `-MLAutoLoadPath` option or the `-MLDomConfigFile` option, but not both.

## Writing .NET synchronization logic

To write .NET synchronization logic, you require knowledge of MobiLink events, some knowledge of .NET, and familiarity with the MobiLink server API for .NET.

For a complete description of the API, see [“MobiLink server API for .NET reference” on page 606](#).

.NET synchronization logic can be used to maintain state information, and implement logic around the upload and download events. For example, a `begin_synchronization` script written in .NET could store the MobiLink user name in a variable. Scripts called later in the synchronization process can access this variable. Also, you can use .NET to access rows in the consolidated database, before or after they are committed.

Using .NET also reduces dependence on the consolidated database. Behavior is affected less by upgrading the consolidated database to a new version or switching to a different database management system.

### Direct row handling

You can use MobiLink direct row handling to communicate remote data to any central data source, application, or web service. Direct row handling uses special classes in the MobiLink server APIs for Java or .NET for direct access to synchronized data.

See [“Direct row handling” on page 649](#).

## Class instances

The MobiLink server instantiates your classes at the database connection level. When an event is reached for which you have written a non-static .NET method, the MobiLink server automatically instantiates the class, if it has not already done so on the present database connection.

See [“Constructors” on page 594](#).

### Note

All methods directly associated with a connection-level or table-level event for one script version must belong to the same class.

For each database connection, once a class has been instantiated, the class persists until that connection is closed. So, the same instance may be used for multiple consecutive synchronization sessions. Unless it is explicitly cleared, information present in public or private variables persists across synchronizations that occur on the same connection.

You can also use static classes or variables. In this case, the values are available across all connections in the same domain.

The MobiLink server automatically deletes your class instances only when the connection to the consolidated database is closed.

## Transactions

The normal rules regarding transactions apply to .NET methods. The start and duration of database transactions is critical to the synchronization process. Transactions must be started and ended only by the MobiLink server. Explicitly committing or rolling back transactions on the synchronization connection within a .NET method violates the integrity of the synchronization process and can cause errors.

These rules apply only to the database connections created by the MobiLink server and, in particular, to SQL statements returned by methods.

## SQL-.NET data types

The following table shows SQL data types and the corresponding .NET data types for MobiLink script parameters.

SQL data type	Corresponding .NET data type
VARCHAR	string
CHAR	string
INTEGER	int
BINARY	byte [ ]
TIMESTAMP	DateTime
INOUT INTEGER	ref int
INOUT VARCHAR	ref string
INOUT CHAR	ref string
INOUT BYTEARRAY	ref byte [ ]

## Constructors

The constructor of your class takes no parameters or takes one `iAnywhere.MobiLink.Script.DBConnectionContext` parameter. For example:

```
public ExampleClass(iAnywhere.MobiLink.Script.DBConnectionContext cc)
```

or

```
public ExampleClass()
```

The synchronization context passed to you is for the connection through which the MobiLink server is synchronizing the current user.

The `DBConnectionContext.GetConnection` method returns the same database connection that MobiLink is using to synchronize the present user. You can execute statements on this connection, but you must not commit or roll back the transaction. The MobiLink server manages the transactions.

The MobiLink server uses the constructor that takes a `iAnywhere.MobiLink.Script.DBConnectionContext` parameter if it exists. If it does not, it uses the void constructor.

See [“DBConnectionContext interface” on page 609](#).

## Methods

In general, you implement one method for each synchronization event. These methods must be public. If they are private, the MobiLink server cannot use them and fails to recognize that they exist.

The names of the methods are not important, as long as the names match the names specified in the `ml_script` table in the consolidated database. In the examples included in the documentation, however, the method names are the same as those of the MobiLink events. This naming convention makes the .NET code easier to read.

The signature of your method should match the signature of the script for that event, except that you can truncate the parameter list if you do not need the values of parameters at the end of the list. You should accept only the parameters you need, because overhead is associated with passing the parameters.

You cannot, however, overload the methods. Only one method prototype per class may appear in the `ml_script` system table.

### Registering methods

After creating a method, you must register it. Registering the method creates a reference to the method in the MobiLink system tables on the consolidated database, so that the method is called when the event occurs. You register methods in the same way that you add synchronization scripts, except instead of adding the entire SQL script to the MobiLink system table, you add only the qualified method name.

See [“Adding and deleting scripts” on page 327](#).

### Return values

Methods called for a SQL-based upload or download must return a valid SQL language statement. The return type of these methods must be `String`. No other return types are allowed.

The return type of all other scripts must either be `string` or `void`. No other types are allowed. If the return type is a string and not null, the MobiLink server assumes that the string contains a valid SQL statement and executes this statement in the consolidated database as it would an ordinary SQL-language synchronization script. If a method ordinarily returns a string but does not want to execute a SQL statement against the database upon its return, it can return null.

## User-defined start classes

You can define start classes that are loaded automatically when the server is started. The purpose of this feature is to allow you to write .NET code that executes at the time the MobiLink server starts the CLR—

before the first synchronization. This means you can create connections or cache data before the first user synchronization request in the server instance.

You do this with the `MLStartClasses` option of the `mlsrv11 -sl dnet` option. For example, the following is part of an `mlsrv11` command line. It causes `mycl1` and `mycl2` to be loaded as start classes.

```
-sl dnet(-MLStartClasses=MyNameSpace.MyClass.mycl1,MyNameSpace.MyClass.mycl2)
```

Classes are loaded in the order in which they are listed. If the same class is listed more than once, more than one instance is created.

All start classes must be public and must have a public constructor that either accepts no arguments or accepts one argument of type `MobiLink.Script.ServerContext`.

The names of loaded start classes are output to the MobiLink log with the message "Loaded .NET start class: *classname*".

For more information about .NET CLR, see [“-sl dnet option” on page 90](#).

To see the start classes that are constructed at server start time, see [“GetStartClassInstances method” on page 629](#).

### Example

The following is a start class template. It starts a daemon thread that processes events and creates a database connection. (Not all start classes need to create a thread but if a thread is spawned it should be a daemon thread.)

```
using System;
using System.IO;
using System.Threading;
using iAnywhere.MobiLink.Script;

namespace TestScripts {
    public class MyStartClass {
        ServerContext    _sc;
        bool              _exit_loop;
        Thread            _thread;
        OdbcConnection   _conn;

        public MyStartClass(ServerContext sc) {

            // Perform setup first so that an exception
            // causes MobiLink startup to fail.
            _sc = sc;

            // Create connection for use later.
            _conn = _sc.makeConnection();
            _exit_loop = false;
            _thread = new Thread(new ThreadStart(run)) ;
            _thread.IsBackground = true;
            _thread.Start();
        }

        public void run() {
            ShutdownCallback callback = new
            ShutdownCallback(shutdownPerformed);
            _sc.ShutdownListener += callback;

            // run() can't throw exceptions.
        }
    }
}
```



```
try {
    handlerLoop();
    _conn.close();
    _conn = null;
}
catch(Exception e) {
    // Print some error output to the MobiLink log.
    Console.Error.Write(e.ToString());

    // There is no need to be notified of shutdown.
    _sc.ShutdownListener -= callback;

    // Ask server to shut down so this fatal error can be fixed.
    _sc.Shutdown();
}

// Shortly after return, this thread no longer exists.
return;
}

public void shutdownPerformed(ServerContext sc) {
    // Stop the event handler loop.
    try {
        _exit_loop = true;

        // Wait a maximum of 10 seconds for thread to die.
        _thread.Join(10*1000);
    }
    catch(Exception e) {
        // Print some error output to the MobiLink log.
        Console.Error.Write(e.ToString());
    }
}

private void handlerLoop() {
    while (!_exit_loop) {
        // Handle events in this loop.
        Thread.Sleep(1*1000);
    }
}
}
```

## Printing information from .NET

You may choose to add statements to your .NET methods that print information to the MobiLink log using `System.Console`. Doing so can help you track the progress and behavior of your classes.

### Performance tip

Printing information in this manner to the MobiLink log is a useful monitoring tool, but is not recommended in a production scenario.

The same technique can be exploited to log arbitrary synchronization information or collect statistical information on how your scripts are used.

## Handling MobiLink server errors with .NET

When scanning the log is not enough, you can monitor your applications programmatically. For example, you can send messages of a certain type in an email.

You can write methods that are passed a class representing every error or warning message that is printed to the log. This may help you monitor and audit a MobiLink server.

The following code installs a listener for all error messages and prints the information to a StreamWriter.

```
class TestLogListener {
    public TestLogListener(StreamWriter output_file) {
        _output_file = output_file;
    }

    public void errCallback(ServerContext sc, LogMessage lm) {
        string type;
        string user;

        if (lm.Type == LogMessage.MessageType.ERROR) {
            type = "ERROR";
        } else if (lm.Type==LogMessage.MessageType.WARNING) {
            type = "WARNING";
        }
        else {
            type = "INVALID TYPE!!";
        }
        if (lm.User == null) {
            user = "null";
        }
        else {
            user = lm.User;
        }

        _output_file.WriteLine("Caught msg type=" + type
            + " user=" + user
            + " text=" + lm.Text);
        _output_file.Flush();
    }
    StreamWriter _output_file;
}
```

The following code registers the TestLogListener. Call this code from somewhere that has access to the ServerContext such as a class constructor or synchronization script.

```
// ServerContext serv_context;
TestLogListener etll = new TestLogListener(log_listener_file);
serv_context.ErrorListener += new LogCallback(etll.errCallback);
```

### See also

- [“LogCallback delegate” on page 627](#)
- [ErrorListener and WarningListener in “ServerContext interface” on page 629](#)
- [“LogMessage class” on page 627](#)
- [“MessageType enumeration” on page 627](#)

## Debugging .NET synchronization logic

The following procedure describes one way you can debug your .NET scripts using Visual Studio.

### To debug .NET scripts

1. Compile your code with debugging information turned on using one of the following methods:
  - On the `csc` command line, set the `/debug+` option.
  - Use Microsoft Visual Studio settings to set debug output.
    - Choose **File** » **Build** » **Configuration Manager**.  
In the **Active Solution Configuration** list, choose **Debug**.
    - Build your assembly.
2. Close running instances of Visual Studio that contain your source files.
3. In this step, you start a new Visual Studio instance to debug the MobiLink server and your .NET synchronization scripts. Start Visual Studio using a command line option to debug the MobiLink server.
  - At a command prompt, navigate to the `Common7\IDE` subdirectory of your Visual Studio installation.
  - Start `devenv` (the Visual Studio IDE) using the `/debugexe` option.

For example, run the following command to debug the MobiLink server. Remember to specify `mlsrv11` options, including the connection string and the option to load .NET assemblies.

For 32-bit Windows environments:

```
devenv /debugexe %sqlany11%\bin32\mlsrv11.exe -c ...
```

For 64-bit Windows environments:

```
devenv /debugexe %sqlany11%\bin64\mlsrv11.exe -c ...
```

Visual Studio starts and `mlsrv11.exe` appears in the Solution Explorer window.

4. Set up Microsoft Visual Studio for debugging .NET code:
  - In the Visual Studio Solution Explorer window, right-click `mlsrv11.exe` and choose Properties.
  - Change Debugger Type from Auto to Mixed or Managed Only to ensure that Visual Studio only debugs your .NET synchronization scripts.
5. Open the associated .NET source files and set break points.

Note: Open the source files individually in the `mlsrv11` solution. Do not open the original solution or project file.
6. Start MobiLink from the Debug menu or by pressing F5.

If prompted, save `mlsrv11.sln`.

If the No Symbolic Information window appears, click **OK** to debug anyway. You are debugging the managed .NET synchronization scripts that MobiLink calls, not the MobiLink server itself.
7. Perform a synchronization that causes the code with a breakpoint to be executed by MobiLink.

## .NET synchronization techniques

This section describes techniques you can use to tackle common .NET synchronization tasks.

### Uploading or downloading rows

For information about how to upload or download rows via .NET, see [“Direct row handling” on page 649](#).

---

## Loading shared assemblies

This section details options to load .NET assemblies and details the process to load shared assemblies.

### Loading assemblies

A .NET assembly is a package of types, metadata, and executable code. In .NET applications, all code must be in an assembly. Assembly files have the extension *.dll* or *.exe*.

There are the following types of assemblies:

- **Private assemblies** A private assembly is a file in the file system.
- **Shared assemblies** A shared assembly is an assembly that is installed in the global assembly cache.

Before MobiLink can load a class and call a method of that class, it must locate the assembly that contains the class. MobiLink only needs to locate the assembly that it calls directly. The assembly can then call any other assemblies it needs.

For example, MobiLink calls MyAssembly, and MyAssembly calls UtilityAssembly and NetworkingUtilsAssembly. In this situation, MobiLink only needs to be configured to find MyAssembly.

MobiLink provides the following ways to load assemblies:

- **Use -sl dnet ( -MLAutoLoadPath )** This option only works with private assemblies. It sets the path to the application base directory and loads all the assemblies within it.

When you use the `-MLAutoLoadPath` option you cannot specify a domain when entering the fully qualified method name for the event script.

When you specify a path and directory with `-MLAutoLoadPath`, MobiLink does the following:

- sets this path as the application base path
- loads all classes in all files ending with *.dll* or *.exe* in the directory that you specified
- creates one application domain and loads into that domain all user classes that do not have a domain specified

Assemblies in the global assembly cache cannot be called directly with this option. To call these shared assemblies, use `-MLDomConfigFile`.

- **Use -sl dnet ( -MLDomConfigFile )** This option works with both private and shared assemblies. It requires a configuration file that contains domain and assembly settings. You should use this option when you have shared assemblies, when you don't want to load all the assemblies in the application base path, or when for some other reason you need to use a configuration file.

With this option, MobiLink reads the settings in the specified domain configuration file. A domain configuration file contains configuration settings for one or more .NET domains. If there is more than one domain represented in the file, the first one that is specified is used as the default domain. (The default domain is used when scripts do not have a domain specified.)

When loading assemblies, MobiLink tries to load the assembly first as private, and then attempts to load the assembly from the global assembly cache. Private assemblies must be located in the application base directory. Shared assemblies are loaded from the global assembly cache.

With the `-MLDomConfigFile` option, only assemblies that are specified in the domain configuration file can be called directly from event scripts.

### Sample domain configuration file

A sample domain configuration file called `mlDomConfig.xml` is installed with MobiLink. You can write your own file from scratch, or edit the sample to suit your needs. The sample file is located in the SQL Anywhere path, in

`MobiLink\setup\dnet\mlDomConfig.xml`

The following is the content of the sample domain configuration file `mlDomConfig.xml`:

```
<?xml version="1.0" encoding="utf-8"?>
<config xmlns="iAnywhere.MobiLink.mlDomConfig"
xsi:schemaLocation='iAnywhere.MobiLink.mlDomConfig mlDomConfig.xsd'
xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance' >
  <domain>
    <name>SampleDomain1</name>
    <appBase>C:\scriptsDir</appBase>
    <configFile></configFile>
    <assembly name="Assembly1" />
    <assembly name="Assembly2" />
  </domain>

  <domain>
    <name>SampleDomain2</name>
    <appBase>\Dom2assembly</appBase>
    <configFile>\Dom2assembly\AssemblyRedirects.config</configFile>
    <assembly name="Assembly3" />
    <assembly name="Assembly4" />
  </domain>
</config>
```

The following is an explanation of the contents of `mlDomConfig.xml`:

- **name** is the domain name, used when specifying the domain in an event script. An event script with the format "DomainName:Namespace.Class.Method" would require that the DomainName domain be in the domain configuration file.  
You must specify at least one domain name.
- **appBase** is the directory that the domain should use as its application base directory. All private assemblies are loaded by the .NET CLR based on this directory. You must specify appBase.
- **configFile** is the .NET application configuration file that should be used for the domain. This can be left blank. It is usually used to modify the default assembly binding and loading behavior. Refer to your .NET documentation for more information about application configuration files.
- **assembly** is the name of an assembly that MobiLink should load and search when resolving type references in event scripts. You must specify at least one assembly. If an assembly is used in more than one domain, it must be specified as an assembly in each domain. If the assembly is private, it must be in the application base directory for the domain.

For more information about the mlsrv11 option -sl dnet, see [“-sl dnet option” on page 90](#).

## .NET synchronization example

This example modifies an existing application to describe how to use .NET synchronization logic to handle the `authenticate_user` event. It creates a C# script for `authenticate_user` called `AuthUser.cs`. This script looks up the user's password in a table called `user_pwd_table` and authenticates the user based on that password.

### To create your .NET synchronization script

1. Add the table `user_pwd_table` to the database. Execute the following SQL statements in Interactive SQL:

```
CREATE TABLE user_pwd_table (  
    user_name  varchar(128) PRIMARY KEY NOT NULL,  
    pwd        varchar(128)  
)
```

2. Add a user and password to the table:

```
INSERT INTO user_pwd_table VALUES('user1', 'myPwd')
```

3. Create a directory for your .NET assembly. For example, `c:\mlexample`.
4. Create a file called `AuthUser.cs` with the following contents:

See [“authenticate\\_user connection event” on page 358](#).

```
using System;  
using iAnywhere.MobiLink.Script;  
namespace MLEExample {  
  
    public class AuthClass {  
        private DBConnection _conn;  
  
        /// AuthClass constructor.  
        public AuthClass(DBConnectionContext cc) {  
            _conn = cc.GetConnection();  
        }  
  
        /// The DoAuthenticate method handles the 'authenticate_user'  
        /// event.  
        /// Note: This method does not handle password changes for  
        /// advanced authorization status codes.  
        public void DoAuthenticate(  
            ref int authStatus,  
            string user,  
            string pwd,  
            string newPwd)  
        {  
            DBCommand pwd_command = _conn.CreateCommand();  
            pwd_command.CommandText = "select pwd from user_pwd_table"  
                + " where user_name = ? ";  
            pwd_command.Prepare();  
  
            // Add a parameter for the user name.  
            DBParameter user_param = new DBParameter();  
            user_param.DbType = SQLType.SQL_CHAR;  
  
            // Set the size for SQL_VARCHAR.  
            user_param.Size = (uint) user.Length;  
            user_param.Value = user;  
            pwd_command.Parameters.Add(user_param);  
        }  
    }  
}
```



```

        // Fetch the password for this user.
        DBRowReader rr      = pwd_command.ExecuteReader();
        object[] pwd_row    = rr.NextRow();

        if (pwd_row == null) {
            // User is unknown.
            authStatus = 4000;
        }
        else {
            if (((string) pwd_row[0]) == pwd) {
                // Password matched.
                authStatus = 1000;
            }
            else {
                // Password did not match.
                authStatus = 4000;
            }
        }
        pwd_command.Close();
        rr.Close();
        return;
    }
}

```

The `MLEExample.AuthClass.DoAuthenticate` method handles the `authenticate_user` event. It accepts the user name and password and returns an authorization status code indicating the success or failure of the validation.

5. Compile the file `AuthUser.cs`. You can do this on the command line or in Visual Studio.

For example, the following command line compiles `AuthUser.cs` and generate an Assembly named `example.dll` in `c:\mlexample`.

```

csc /out:c:\mlexample\example.dll /target:library /reference:"%SQLANY11%
\Assembly\v2\iAnywhere.MobiLink.Script.dll" AuthUser.cs

```

6. Register .NET code for the `authenticate_user` event. The method you need to execute (`DoAuthenticate`) is in the namespace `MLEExample` and class `AuthClass`. Execute the following SQL:

```

CALL ml_add_dnet_connection_script('ex_version', 'authenticate_user',
'MLEExample.AuthClass.DoAuthenticate')
COMMIT

```

7. Run the MobiLink server with the following option. This option causes MobiLink to load all assemblies in `c:\myexample`:

```

-sl dnet (-MLAutoLoadPath=c:\mlexample)

```

Now, when a user synchronizes with the version `ex_version`, they are authenticated with the password from the table `user_pwd_table`.

## MobiLink server API for .NET reference

This section explains the MobiLink .NET interfaces and classes, and their associated methods, properties, and constructors. To use these classes, reference the *iAnywhere.MobiLink.Script.dll* assembly, located in *install-dir\Assembly\v2*.

This section focuses on C#, but there are equivalents in Visual Basic.NET and C++.

### DBCommand interface

#### Syntax

```
interface DBCommand
    Member of iAnywhere.MobiLink.Script
```

#### Remarks

Represents a SQL statement or database command. DBCommand can represent an update or query.

#### Example

For example, the following C# code uses the DBCommand interface to execute two queries:

```
DBCommand stmt = conn.CreateCommand();
stmt.CommandText = "SELECT t1a1, t1a2 FROM table1 ";

DBRowReader rs = stmt.ExecuteReader();
printResultSet(rs);
rs.Close();

stmt.CommandText = "SELECT t2a1 FROM table2 ";
rs = stmt.ExecuteReader();
printResultSet(rs);
rs.Close();
stmt.Close();
```

The following C# example uses DBCommand to execute an update with parameters:

```
public void prepare_for_download(
    DateTime last_download,
    String ml_username)
{
    DBCommand cstmt = conn.CreateCommand();
    cstmt.CommandText = "CALL myProc(?, ?, ?, ?)";
    cstmt.Prepare();

    DBParameter param = new DBParameter();
    param.DbType = SqlDbType.SQL_CHAR;
    param.Value = "10000";
    cstmt.Parameters.Add(param);

    param = new DBParameter();
    param.DbType = SqlDbType.SQL_INTEGER;
    param.Value = 20000;
    cstmt.Parameters.Add(param);
}
```

```

        param                = new DBParameter();
        param.DbType         = SQLType.SQL_DECIMAL;
        param.Precision      = 5;
        param.Value          = new Decimal(30000);
        pstmt.Parameters.Add(param);

        param                = new DBParameter();
        param.DbType         = SQLType.SQL_TIMESTAMP;
        param.Precision      = 19;
        param.Value          = last_download;
        pstmt.Parameters.Add(param);

        // Execute update
        DBRowReader rset     = pstmt.ExecuteNonQuery();
        pstmt.Close();
    }

```

## Prepare method

### Syntax

```
void Prepare()
```

### Remarks

Prepares the SQL statement stored in CommandText for execution.

## ExecuteNonQuery method

### Syntax

```
int ExecuteNonQuery()
```

### Remarks

Executes a non-query statement. Returns the number of rows in the database affected by the SQL statement.

## ExecuteReader method

### Syntax

```
DBRowReader ExecuteReader()
```

### Remarks

Executes a query statement returning the result set. Returns a DBRowReader for retrieving results returned by the SQL statement.

## Close method

### Syntax

```
void Close()
```

### Remarks

Closes the current SQL statement or command.

## CommandText property

### Syntax

```
string CommandText
```

### Remarks

The value is the SQL statement to be executed.

## Parameters property

### Syntax

```
DBParameterCollection Parameters
```

### Remarks

Gets the `iAnywhere.MobiLink.Script.DBParameterCollection` for this `DBCommand`.

## DBConnection interface

### Syntax

```
interface DBConnection  
Member of iAnywhere.MobiLink.Script
```

### Remarks

Represents a MobiLink ODBC connection.

This interface allows user-written synchronization logic to access an ODBC connection created by MobiLink.

## Commit method

### Syntax

```
void Commit( )
```

### Remarks

Commits the current transaction.

## Rollback method

### Syntax

```
void Rollback( )
```

### Remarks

Rolls back the current transaction.

## Close method

### Syntax

```
void Close( )
```

### Remarks

Closes the current connection.

## CreateCommand method

### Syntax

```
DBCommand CreateCommand( )
```

### Remarks

Creates a SQL statement or command on this connection. Returns the newly generated DBCommand.

## DBConnectionContext interface

### Syntax

```
interface DBConnectionContext  
Member of iAnywhere.MobiLink.Script
```

## Remarks

Interface for obtaining and accessing information about the current database connection. This is passed to the constructor of classes containing scripts. If context is required for a background thread or beyond the lifetime of a connection, use a `ServerContext`.

For more information about constructors, see [“Constructors” on page 594](#).

### Caution

A `DBConnectionContext` instance should not be used outside the thread that calls into your .NET code.

## GetConnection method

### Syntax

```
iAnywhere.MobiLink.Script.DBConnection GetConnection( )  
Member of iAnywhere.MobiLink.Script.DBConnectionContext
```

### Remarks

Returns the existing connection to the MobiLink consolidated database. The connection is the same connection that MobiLink uses to execute SQL scripts.

This connection must not be committed, closed or altered in any way that would affect the MobiLink server use of the connection. The connection returned is only valid for the lifetime of the underlying MobiLink connection. Do not use the connection after the `end_connection` event has been called for the connection.

If a server connection with full access is required, use `ServerContext.makeConnection()`.

## GetDownloadData method

### Syntax

```
DownloadData GetDownloadData();
```

### Remarks

Returns a `DownloadData` instance for the current synchronization. Use the `DownloadData` instance to create the download for direct row handling.

### Returns

A `DownloadData` instance for the current synchronization.

### Example

The following example assumes you have a `DBConnectionContext` instance called `_cc`.

```
// The method used for the handle_DownloadData event.  
public void HandleDownload() {  
    // Get the DownloadData for the current synchronization.  
    DownloadData my_dd = _cc.GetDownloadData();  
}
```

```
// Get an array of tables to set download operations.
DownloadTableData[] download_tables = my_dd.GetDownloadTables();

// Get the first table in the DownloadTableData array.
DownloadTableData my_download_table = download_tables[0];

// ...
}
```

## GetServerContext method

### Syntax

```
public iAnywhere.MobiLink.Script.ServerContext.GetServerContext()
Member of iAnywhere.MobiLink.Script.DBConnectionContext
```

### Remarks

Returns the ServerContext for this MobiLink server.

## GetProperties method

### Syntax

```
NameValueCollection getProperties()
```

### Remarks

Returns the properties for this connection, based on this connection's script version. Properties are stored in the ml\_property table.

For more information, see [“ml\\_property” on page 712](#) and [“ml\\_add\\_property system procedure” on page 677](#).

## GetRemoteID method

### Syntax

```
string GetRemoteID()
```

### Remarks

Returns the remote ID of the database currently synchronizing on this connection. If your remote database is prior to version 10, it returns the MobiLink user name.

### Returns

The remote ID.

### See also

- “Remote IDs” [[MobiLink - Client Administration](#)]

## GetVersion method

### Syntax

```
string getVersion( )
```

### Remarks

Returns the name of the script version.

### See also

- “ml\_property” on page 712
- “ml\_add\_property system procedure” on page 677

## DBParameter class

### Syntax

```
class DBParameter  
Member of iAnywhere.MobiLink.Script
```

### Remarks

Represents a bound ODBC parameter.

DBParameter is required to execute commands with parameters. All parameters must be in place before the command is executed.

### Example

For example, the following C# code uses DBCommand to execute an update with parameters:

```
public void handleUpload(UploadData ud) {  
    UploadedTableData UTDAdmin = ud.GetUploadedTableByName( "Admin" );  
    IDataReader      AdminIns  = UTDAdmin.GetInserts( );  
    DBCommand        stmt1     = _conn.CreateCommand( );  
    DBParameter      p_id      = new DBParameter( );  
    DBParameter      p_data    = new DBParameter( );  
  
    stmt1.CommandText      = "INSERT INTO Admin(admin_id,data) VALUES  
(?,?)";  
    p_id.DbType            = SQLType.SQL_BIGINT;  
    stmt1.Parameters.Add(p_id);  
    p_data.DbType          = SQLType.SQL_VARCHAR;  
    p_data.Size            = 30;  
    stmt1.Parameters.Add(p_data);  
  
    stmt1.Prepare( );  
  
    while (AdminIns.Read()) {
```



```
        p_id.Value = AdminIns.GetInt64(0);  
        p_data.Value = AdminIns.GetString(1);  
        stmt1.ExecuteNonQuery();  
    }  
    stmt1.Close();  
}
```

## DbType property

### Syntax

**SQLTYPE DbType**

### Remarks

The value is the SQLType of this parameter.

The default value is SQLType.SQL\_TYPE\_NULL.

## Direction property

### Syntax

**System.Data.ParameterDirection Direction**

### Remarks

The value is the Input/Output direction of this parameter.

The default value is ParameterDirection.Input.

## IsNullable property

### Syntax

bool **IsNullable**

### Remarks

The value Indicates whether this parameter can be null.

The default value is false.

## ParameterName property

### Syntax

string **ParameterName**

### Remarks

The value is the name of this parameter.

The default value is null.

## Precision property

### Syntax

uint **Precision**

### Remarks

The value is the decimal precision of this parameter. Only used for `SQLType.SQL_NUMERIC` and `SQLType.SQL_DECIMAL` parameters.

The default value is 0.

## Scale property

### Syntax

short **Scale**

### Remarks

The value is the resolvable digits of this parameter. Only used for `SQLType.SQL_NUMERIC` and `SQLType.SQL_DECIMAL` parameters.

The default value is 0.

## Size property

### Syntax

uint **Size**

### Remarks

The value is the size in bytes of this parameter.

The default value is inferred from `DbType`.

## Value property

### Syntax

object **Value**

### Remarks

The value is the value of this parameter.

The default value is null.

## DBParameterCollection class

### Syntax

class **DBParameterCollection**  
inherits from **IDataParameterCollection**, **IList**, **ICollection**, **IEnumerable**  
Member of **iAnywhere.MobiLink.Script**

### Remarks

Collection of DBParameters. When DBCommand creates a DBParameterCollection it is empty and must be filled with appropriate parameters before the DBCommand executes.

## DBParameterCollection method

### Syntax

**DBParameterCollection**( )

### Remarks

Creates an empty list of DBParameters.

## Contains( string parameterName ) method

### Syntax

bool **Contains**( string *parameterName* )

### Remarks

Returns true if the collection contains a parameter with the specified name.

### Parameters

- **parameterName** The name of the parameter to check for.

## IndexOf( string parameterName ) method

### Syntax

int IndexOf( string *parameterName* )

### Remarks

Returns index of the parameter, or -1 if there is no parameter with the given name.

### Parameters

- **parameterName** The name of the parameter to find.

## RemoveAt( string parameterName ) method

### Syntax

void RemoveAt( string *parameterName* )

### Remarks

Removes the parameter with the given name from the collection.

### Parameters

- **parameterName** The name of the parameter to remove.

## Add( object value ) method

### Syntax

int Add( object *value* )

### Remarks

Adds the given parameter to the collection.

### Parameters

- **value** The iAnywhere.MobiLink.Script.DBParameter instance to add to the collection.

### Returns

The index of the added parameter in the collection.

## Clear method

### Syntax

void Clear( )

**Remarks**

Removes all parameters from the collection.

## Contains( object value ) method

**Syntax**

bool **Contains**( object *value* )

**Remarks**

Returns true if this collection contains the given iAnywhere.MobiLink.Script.DBParameter.

**Parameters**

- **value** The iAnywhere.MobiLink.Script.DBParameter to check for.

## IndexOf( object value ) method

**Syntax**

int **IndexOf**( object *value* )

**Remarks**

Returns the index of the given iAnywhere.MobiLink.Script.DBParameter in the collection.

**Parameters**

- **value** The iAnywhere.MobiLink.Script.DBParameter to find.

## Insert( int index, object value ) method

**Syntax**

void **Insert**( int *index*, object *value* )

**Remarks**

Inserts the given iAnywhere.MobiLink.Script.DBParameter into the collection at the specified index.

**Parameters**

- **value** The iAnywhere.MobiLink.Script.DBParameter to insert.
- **index** The index at which to insert the DBParameter.

## Remove( object value ) method

### Syntax

```
void Remove( object value )
```

### Remarks

Removes the given `iAnywhere.MobiLink.Script.DBParameter` from the collection.

### Parameters

- **value** The `iAnywhere.MobiLink.Script.DBParameter` to remove.

## RemoveAt( int index) method

### Syntax

```
void RemoveAt( int index )
```

### Remarks

Removes the `iAnywhere.MobiLink.Script.DBParameter` at the given index in the collection.

### Parameters

- **index** The index of the `iAnywhere.MobiLink.Script.DBParameter` to remove.

## CopyTo(Array array, int index) method

### Syntax

```
void CopyTo( Array array, int index )
```

### Remarks

Copies the contents of the collection into the given array starting at the specified index.

### Parameters

- **array** The array to which to copy the contents of the collection.
- **index** The index in the array at which to begin copying the contents of the collection.

## GetEnumerator method

### Syntax

```
IEnumerator GetEnumerator( )
```

**Remarks**

Returns an enumerator for the collection.

## IsFixedSize property

**Syntax**

bool **IsFixedSize**

**Remarks**

Returns false.

## IsReadOnly property

**Syntax**

bool **IsReadOnly**

**Remarks**

Returns false.

## Count property

**Syntax**

int **Count**

**Remarks**

The number of parameters in the collection.

## IsSynchronized property

**Syntax**

bool **IsSynchronized**

**Remarks**

Returns false.

## SyncRoot property

### Syntax

object **SyncRoot**

### Remarks

Object that can be used to synchronize the collection.

## this[ string parameterName ] property

### Syntax

object **this[ string parameterName ]**

### Remarks

Gets or sets the `iAnywhere.MobiLink.Script.DBParameter` with the given name in the collection.

### Parameters

- **parameterName** The name of the `iAnywhere.MobiLink.Script.DBParameter` to get or set.

## this[ int index ] property

### Syntax

object **this[ int index ]**

### Remarks

Gets or sets the `iAnywhere.MobiLink.Script.DBParameter` at the given index in the collection.

### Parameters

- **index** The index of the `iAnywhere.MobiLink.Script.DBParameter` to get or set.

## DBRowReader interface

### Syntax

interface **DBRowReader**  
Member of **iAnywhere.MobiLink.Script**

### Remarks

Represents a set of rows being read from a database. Executing the method `DBCommand.executeReader()` creates a `DBRowReader`.



The following example is a C# code fragment. It calls a function with the rows in the result set represented by the given `DBRowReader`.

```
DBCommand stmt = conn.CreateCommand();
stmt.CommandText = "select intCol, strCol from table1 ";
DBRowReader rset = stmt.ExecuteReader();
object[] values = rset.NextRow();

while (values != null) {
    handleRow((int) values[0], (String) values[1]);
    values = rset.NextRow();
}

rset.Close();
stmt.Close();
```

## NextRow method

### Syntax

```
object[ ] NextRow()
```

### Remarks

Retrieves and returns the next row of values in the result set. If there are no more rows in the result set, it returns null.

See [“SQLType enumeration” on page 633](#).

## Close method

### Syntax

```
void Close()
```

### Remarks

Cleans up resources used by this `MLDBRowReader`. After `Close()` is called, this `MLDBRowReader` cannot be used again.

## ColumnNames property

### Syntax

```
string[ ] ColumnNames
```

### Remarks

Gets the names of all columns in the result set. The value is an array of strings corresponding to the column names in the result set.

## ColumnTypes property

### Syntax

SQLType[ ] **ColumnTypes**

### Remarks

Gets the types of all columns in the result set. The value is an array of SQLTypes corresponding to the column types in the result set.

## DownloadData interface

### Syntax

interface **DownloadData**

### Remarks

Encapsulates download data operations for direct row handling. To obtain a DownloadData instance, use the DBConnectionContext GetDownloadData method. Use the DownloadData.GetDownloadTables and GetDownloadTableByName methods to return DownloadTableData instances.

This download data is available through DBConnectionContext. It is not valid to access the download data before the begin\_synchronization event or after the end\_download event. It is not valid to access DownloadData in an upload-only synchronization.

### See also

- [“DownloadTableData interface” on page 623](#)
- [“handle\\_DownloadData connection event” on page 442](#)
- [DBConnectionContext “GetDownloadData method” on page 610](#)
- [“Direct row handling” on page 649](#)

## GetDownloadTables method

### Syntax

DownloadTableData[ ] **GetDownloadTables**( );

### Remarks

Gets an array of all the tables for download data in this synchronization. The operations performed on this table are sent to the remote database.

### Returns

An array of download table data. The order of tables in the array is the same as the upload order for the remote.

## Example

The following example uses the `DownloadData.GetDownloadTables` method to obtain an array of `DownloadTableData` objects for the current synchronization. The example assumes you have a `DBConnectionContext` instance called `_cc`.

```
// The method used for the handle_DownloadData event.
public void HandleDownload() {
    // Get the DownloadData for the current synchronization.
    DownloadData my_dd = _cc.GetDownloadData();

    // Get an array of tables to set download operations.
    DownloadTableData[] download_tables = my_dd.GetDownloadTables();

    // Get the first table in the DownloadTableData array.
    DownloadTableData my_download_table = download_tables[0];

    // ...
}
```

## GetDownloadTableByName method

### Syntax

```
DownloadTableData GetDownloadTableByName(
    string table-name);
```

### Remarks

Gets the named download table for this synchronization. Returns null if there is no table with the given name in this synchronization.

### Returns

Download data for the given table name or null if not found.

### Parameters

- **table-name** Name of the table for which you want the download data.

## DownloadTableData interface

### Syntax

```
interface DownloadTableData
```

### Remarks

Encapsulates information for one download table for a synchronization. Use this interface to set the data operations that are downloaded to a synchronization client site.

### Example

For example, suppose you have the following table:

```
CREATE TABLE remoteOrders (  
    pk INT NOT NULL,  
    coll VARCHAR(200),  
    PRIMARY KEY (pk)  
);
```

The following example uses the `DownloadData.GetDownloadTableByName` method to return a `DownloadTableData` instance representing the `remoteOrders` table.

```
// The method used for the handle_DownloadData event  
public void HandleDownload() {  
    // _cc is a DBConnectionContext instance.  
  
    // Get the DownloadData for the current synchronization.  
    DownloadData my_dd = _cc.GetDownloadData();  
  
    // Get the DownloadTableData for the remoteOrders table.  
    DownloadTableData td = my_dd.GetDownloadTableByName("remoteOrders");  
  
    // User defined-methods to set download operations.  
    SetDownloadUpserts(td);  
    SetDownloadDeletes(td);  
  
    // ...  
}
```

In this example, the `SetDownloadInserts` method uses `DownloadTableData.GetUpsertCommand` to obtain a command for the rows you want to insert or update. The `IDbCommand` holds the parameters that you set to the values you want inserted on the remote database.

```
void SetDownloadInserts(DownloadTableData td) {  
    IDbCommand upsert_cmd = td.GetUpsertCommand();  
    IDataParameterCollection parameters = upsert_cmd.Parameters;  
  
    // The following method calls are the same as the following SQL  
    statement:  
    // INSERT INTO remoteOrders(pk, coll) values(2300, "truck");  
    ((IDataParameter) (parameters[0])).Value = (Int32) 2300;  
    ((IDataParameter) (parameters[1])).Value = (String) "truck";  
  
    if (upsert_cmd.ExecuteNonQuery() > 0) {  
        // Insert was not filtered.  
    }  
    else {  
        // Insert was filtered because it was uploaded  
        // in the same synchronization.  
    }  
}
```

The `SetDownloadDeletes` method uses the `DownloadTableData.GetDeleteCommand` to obtain a command for rows you want to delete.

```
void SetDownloadDeletes(DownloadTableData td) {  
    IDbCommand delete_cmd = t2_download_dd.GetDeleteCommand();  
  
    // The following method calls are the same as the following SQL  
    statement:  
    // DELETE FROM remoteOrders where pk = 2300;  
    IDataParameterCollection parameters = delete_cmd.Parameters;  
    ((IDataParameter) (parameters[0])).Value = (Int32) 2300;  
    delete_cmd.ExecuteNonQuery();  
}
```

## GetDeleteCommand method

### Syntax

```
IDbCommand GetDeleteCommand();
```

### Remarks

Gets a command which allows the user to add delete operations to the download data operations. The command returned has the same number of parameters as primary key columns in this table. For the delete to be included in the download, the column values for the primary key columns must be set and the statement executed with `ExecuteNonQuery()`.

**Note**

You must set all primary key values for download delete operations.

### Returns

A Command for deletes in the download.

### Example

See [“DownloadTableData interface” on page 623](#).

## GetLastDownloadTime method

### Syntax

```
DateTime GetLastDownloadTime();
```

### Remarks

Returns last download time for this table. This is the same last download time passed to several of the per table download events.

The last download time is useful for generating the table download data for a particular synchronization.

### Returns

The last download time for this download table.

## GetName method

### Syntax

```
string GetName();
```

### Remarks

Gets the table name of this instance. This is a utility function. The table name can also be accessed via the Schema for this instance.

### Returns

Table name of this instance.

## GetSchemaTable method

### Syntax

```
DataTable GetSchemaTable();
```

### Remarks

Gets a DataTable instance that describes the metadata for this download table.

If you want the DataTable to contain column name information, you must specify the client option to send column names.

### Returns

A DataTable that describes the column metadata.

### See also

- dbmlsync: “[SendColumnNames \(scn\) extended option](#)” [*MobiLink - Client Administration*]
- UltraLite: “[Send Column Names synchronization parameter](#)” [*UltraLite - Database Management and Reference*]

## GetUpsertCommand method

### Syntax

```
IDbCommand GetUpsertCommand();
```

### Remarks

Gets a command that allows you to add upsert (insert/update) operations to the direct download data operations. The command that is returned has the same number of parameters as columns in this table. The column values for the insert must be set and the statement executed with `ExecuteNonQuery()` for the insert/update to be included in the download. `ExecuteNonQuery()` on the command returns 0 if the insert/update operation was filtered and returns 1 if the insert/update was not filtered.

You cannot add or remove parameters to this command; you can only set their values.

### Returns

A Command for inserts/updates for the download.

**Example**

See [“DownloadTableData interface”](#) on page 623.

## LogCallback delegate

**Syntax**

```
delegate void LogCallback(  
    ServerContext sc  
    LogMessage message  
)  
Member of iAnywhere.MobiLink.Script
```

**Remarks**

Called when the MobiLink server prints a message.

## LogMessage class

**Syntax**

```
class LogMessage : iAnywhere.MobiLink.Script.LogMessage  
Member of iAnywhere.MobiLink.Script
```

**Remarks**

Contains information about a message printed to the log.

## MessageType enumeration

**Syntax**

```
enum MessageType  
Member of iAnywhere.MobiLink.Script.LogMessage
```

**Remarks**

Enumeration of the possible types of LogMessage.

## ERROR field

**Syntax**

```
ERROR
```

**Remarks**

The log message is an error.

## INFO field

### Syntax

INFO

### Remarks

The log info message.

## WARNING field

### Syntax

WARNING

### Remarks

The log message is a warning.

## Type property

### Syntax

LogMessage.MessageType **Type**

### Remarks

The type of the log message that this instance represents.

## User property

### Syntax

string **User**

### Remarks

The user for which this message is being logged. It may be null.

## Text property

### Syntax

string **Text**

### Remarks

The main text of the message.



## ServerContext interface

### Syntax

```
interface ServerContext  
Member of iAnywhere.MobiLink.Script
```

### Remarks

An instantiation of all the context that is present for the duration of the MobiLink server. This context can be held as static data and used in a background thread. It is valid for the duration of the .NET CLR invoked by MobiLink.

To access a ServerContext instance, use the DBConnectionContext.getServerContext method.

## GetStartClassInstances method

### Syntax

```
object[ ] GetStartClassInstances( )  
Member of iAnywhere.MobiLink.Script.ServerContext
```

### Remarks

Gets an array of the start classes that were constructed at server start time. The array length is zero if there are no start classes.

For more information about user-defined start classes, see [“User-defined start classes” on page 595](#).

The following is an example of getStartClassInstances():

```
void FindStartClass(ServerContext sc, string name) {  
    object[] startClasses = sc.GetStartClassInstances();  
    foreach (object obj in startClasses) {  
        if (obj is MyClass) {  
            // Execute some code.  
        }  
    }  
}
```

## LogCallback ErrorListener event

This event is triggered when the MobiLink server prints an error.

## LogCallback InfoListener event

This event is triggered when the MobiLink server prints info.

## LogCallback WarningListener event

This event is triggered when the MobiLink server prints a warning.

## MakeConnection method

### Syntax

**iAnywhere.MobiLink.Script.DBConnection makeConnection( )**  
Member of **iAnywhere.MobiLink.Script.ServerContext**

### Remarks

Creates a new server connection.

## ShutDown method

### Syntax

**void Shutdown( )**  
Member of **iAnywhere.MobiLink.Script.ServerContext**

### Remarks

Forces the server to shut down.

## ShutdownListener method

### Syntax

event **iAnywhere.MobiLink.Script.ShutdownCallback**  
**ShutdownListener(**  
**iAnywhere.MobiLink.Script.ServerContext sc)**  
Member of **iAnywhere.MobiLink.Script.ServerContext**

### Remarks

This event is triggered on shutdown. The following code is an example of how to use this event:

```
ShutdownCallback callback = new ShutdownCallback(shutdownHandler);
_sc.ShutdownListener += callback;

public void shutdownHandler(ServerContext sc) {
    _test_out_file.WriteLine("shutdownPerformed");
}
```

## getProperties method

### Syntax

```
NameValueCollection getProperties(  
    string component_name  
    string prop_set_name )
```

### Remarks

Returns the set of properties associated with the script version. These are stored in the MobiLink system table ml\_property.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

## getPropertiesByVersion method

### Syntax

```
NameValueCollection getPropertiesByVersion( string script_version )
```

### Remarks

Returns the set of properties associated with the script version. These are stored in the MobiLink system table ml\_property. The script version is stored in the property\_set\_name column when the component\_name is ScriptVersion.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

## getPropertySetNames method

### Syntax

```
StringCollection getPropertySetNames( string component_name )
```

### Remarks

Returns the list of property set names for a given component. These are stored in the MobiLink system table ml\_property.

### See also

- [“ml\\_property” on page 712](#)
- [“ml\\_add\\_property system procedure” on page 677](#)

## ServerException class

### Syntax

```
public class ServerException  
Member of iAnywhere.MobiLink.Script
```

### Remarks

Used to signal MobiLink that an error has occurred with the server and it should shut down immediately.

## ServerException constructors

### Syntax

```
public ServerException( )  
Member of iAnywhere.MobiLink.Script.ServerException
```

### Remarks

Constructs a ServerException with no detail message.

### Syntax

```
public ServerException( string message )  
Member of iAnywhere.MobiLink.Script.ServerException
```

### Remarks

Creates a new ServerException with the given message.

### Parameters

- **message** The message for this ServerException.

### Syntax

```
public ServerException( string message, SystemException ie )  
Member of iAnywhere.MobiLink.Script.ServerException
```

### Remarks

Creates a new ServerException with the given message and containing the given inner exception that caused this one.

### Parameters

- **message** The message for this ServerException.
- **ie** The exception that caused this ServerException.

## ShutdownCallback delegate

### Syntax

sealed delegate **ShutdownCallback** : **System.MulticastDelegate**  
Member of **iAnywhere.MobiLink.Script**

### Remarks

Called when the MobiLink server is shutting down. Implementations of this delegate can be registered with the `ServerContext.ShutdownListener` event to be called when the MobiLink server shuts down.

## SQLType enumeration

### Syntax

enum **SQLType**  
Member of **iAnywhere.MobiLink.Script**

### Remarks

Enumeration of all possible ODBC data types.

## SQL\_TYPE\_NULL field

### Syntax

**SQL\_TYPE\_NULL**

### Remarks

Null data type.

## SQL\_UNKNOWN\_TYPE field

### Syntax

**SQL\_UNKNOWN\_TYPE**

### Remarks

Unknown data type.

## SQL\_CHAR field

### Syntax

**SQL\_CHAR**

**Remarks**

Single byte string. Has .NET type String.

## SQL\_NUMERIC field

**Syntax**

SQL\_NUMERIC

**Remarks**

Numeric value of set size and precision. Has .NET type Decimal.

## SQL\_DECIMAL field

**Syntax**

SQL\_DECIMAL

**Remarks**

Decimal number of set size and precision. Has .NET type Decimal.

## SQL\_INTEGER field

**Syntax**

SQL\_INTEGER

**Remarks**

32-bit integer. Has .NET type Int32.

## SQL\_SMALLINT field

**Syntax**

SQL\_SMALLINT

**Remarks**

16-bit integer. Has .NET type Int16.

## SQL\_FLOAT field

### Syntax

SQL\_FLOAT

### Remarks

Floating point number with ODBC driver defined precision. Has .NET type Double.

## SQL\_REAL field

### Syntax

SQL\_REAL

### Remarks

Single precision floating-point number. Has .NET type Single.

## SQL\_DOUBLE field

### Syntax

SQL\_DOUBLE

### Remarks

Double precision floating point number. Has .NET type Double.

## SQL\_DATE field

### Syntax

SQL\_DATE

### Remarks

A date. Has .NET type DateTime.

## SQL\_DATETIME field

### Syntax

SQL\_DATETIME

### Remarks

A date and time. Has .NET type DateTime.

## SQL\_TIME field

### Syntax

SQL\_TIME

### Remarks

A time. Has .NET type DateTime.

## SQL\_INTERVAL field

### Syntax

SQL\_INTERVAL

### Remarks

An interval of time. Has .NET type TimeSpan.

## SQL\_TIMESTAMP field

### Syntax

SQL\_TIMESTAMP

### Remarks

A time stamp. Has .NET type DateTime.

## SQL\_VARCHAR field

### Syntax

SQL\_VARCHAR

### Remarks

Single byte string. Has .NET type String.

## SQL\_TYPE\_DATE field

### Syntax

SQL\_TYPE\_DATE

### Remarks

A date. Has .NET type DateTime.



## SQL\_TYPE\_TIME field

### Syntax

SQL\_TYPE\_TIME

### Remarks

A time. Has .NET type DateTime.

## SQL\_TYPE\_TIMESTAMP field

### Syntax

SQL\_TYPE\_TIMESTAMP

### Remarks

A timestamp. Has .NET type DateTime.

## SQL\_DEFAULT field

### Syntax

SQL\_DEFAULT

### Remarks

A default type. Has no type.

## SQL\_ARD\_TYPE field

### Syntax

SQL\_ARD\_TYPE

### Remarks

An ARD object. Has no type.

## SQL\_BIT field

### Syntax

SQL\_BIT

### Remarks

A single bit. Has .NET type Boolean.

## SQL\_TINYINT field

### Syntax

SQL\_TINYINT

### Remarks

An 8-bit integer. Has .NET type SByte.

## SQL\_BIGINT field

### Syntax

SQL\_BIGINT

### Remarks

A 64-bit integer. Has .NET type Int64.

## SQL\_LONGVARBINARY field

### Syntax

SQL\_LONGVARBINARY

### Remarks

Variable length binary data with a driver dependent maximum length. Has .NET type byte[ ].

## SQL\_VARBINARY field

### Syntax

SQL\_VARBINARY

### Remarks

Variable length binary data with a user specified maximum length. Has .NET type byte[ ].

## SQL\_BINARY field

### Syntax

SQL\_BINARY

### Remarks

Fixed length binary data. Has .NET type byte[ ].

## SQL\_LONGVARCHAR field

### Syntax

SQL\_LONGVARCHAR

### Remarks

Single byte string. Has .NET type String.

## SQL\_GUID field

### Syntax

SQL\_GUID

### Remarks

A Global Unique ID (also called a UUID). Has .NET type Guid.

## SQL\_WCHAR field

### Syntax

SQL\_WCHAR

### Remarks

Unicode character array of fixed size. Has .NET type String.

## SQL\_WVARCHAR field

### Syntax

SQL\_WVARCHAR

### Remarks

Null-terminated Unicode string of user-defined maximum length. Has .NET type String.

## SQL\_WLONGVARCHAR field

### Syntax

SQL\_WLONGVARCHAR

### Remarks

Null-terminated Unicode string of driver-dependent maximum length. Has .NET type String.

## SynchronizationException class

### Syntax

class **SynchronizationException**  
Member of **iAnywhere.MobiLink.Script**

### Remarks

Used to signal that a synchronization exception has occurred and that the current synchronization should be rolled back and restarted.

## SynchronizationException constructors

### Syntax

**SynchronizationException( )**  
Member of **iAnywhere.MobiLink.Script.SynchronizationException**

### Remarks

Constructs a SynchronizationException with no details.

### Syntax

public **SynchronizationException( string message )**  
Member of **iAnywhere.MobiLink.Script.SynchronizationException**

### Remarks

Creates a new SynchronizationException with the given message.

### Parameters

- **message** The message for this ServerException.

### Syntax

**SynchronizationException( string message, SystemException ie )**  
Member of **iAnywhere.MobiLink.Script.SynchronizationException**

### Remarks

Creates a new SynchronizationException with the given message and containing the given inner exception that caused this one.

### Parameters

- **message** The message for this ServerException.
- **ie** The exception that caused this ServerException.

## UploadData interface

### Syntax

```
public interface UploadData
```

### Remarks

Encapsulates upload operations for direct row handling. An upload transaction contains a set of tables containing row operations. An UploadData instance representing a single upload transaction is passed to the handle\_UploadData synchronization event.

#### Caution

You must handle direct row handling upload operations in the method registered for the handle\_UploadData event. The UploadData is destroyed after each call to the registered method. Do not create a new instance of UploadData to use in subsequent events.

Use the UploadData.GetUploadedTables or UploadData.GetUploadedTableByName methods to obtain UploadedTableData instances.

A synchronization has one UploadData unless the remote database is using transactional upload.

### See also

- [“Direct row handling” on page 649](#)
- [“UploadedTableData interface” on page 643](#)
- [“handle\\_UploadData connection event” on page 454](#)
- [“Handling direct uploads” on page 654](#)

### Example

See [“handle\\_UploadData connection event” on page 454](#).

## GetUploadedTableByName method

### Syntax

```
UploadedTableData GetUploadedTableByName(  
    string table-name);
```

### Remarks

Gets the named Uploaded table data in this uploaded transaction. Returns null if there is no table in this transaction with the given name.

### Parameters

- **table-name** Name of the table for which you want the uploaded data.

### Returns

Uploaded data for the given table name or null if not found.

## Example

Assume you use a method called `HandleUpload` for the `handle_UploadData` synchronization event. The following example uses the `GetUploadedTableByName` method to return an `UploadedTableData` instance for the `remoteOrders` table.

```
// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut) {
    UploadedTableData uploaded_t1 =
    ut.GetUploadedTableByName("remoteOrders");
    // ...
}
```

## GetUploadedTables method

### Syntax

```
UploadedTableData[] GetUploadedTables();
```

### Remarks

Gets an array of all the uploaded table data in this uploaded transaction. The order to the tables in the array is the same order that MobiLink uses for SQL row handling, and is the optimal order for preventing referential integrity violations. If your data source is a relational database, use this table order.

### Returns

An array of uploaded table data. The order of tables in the array is the same as the upload order of the client.

## Example

Assume you use a method called `HandleUpload` for the `handle_UploadData` synchronization event. The following example uses the `GetUploadedTables` method to return `UploadedTableData` instances for the current upload transaction.

```
// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ud) {
    UploadedTableData[] tables = ud.GetUploadedTables();
    //...
}
```

## UpdateDataReader interface

### Syntax

```
interface UpdateDataReader
```

### Remarks

Holds the update operations for one upload transaction for one table. New and old rows can both be accessed by changing the mode of the `DataReader` to `old` or `new`. Otherwise this can be used as a regular `DataReader`.

## SetNewRowValues method

### Syntax

```
void SetNewRowValues();
```

### Remarks

Sets the mode of this DataReader to return new column values (the post-update row). This is the default mode.

## SetOldRowValues method

### Syntax

```
void SetOldRowValues();
```

### Remarks

Sets the mode of this DataReader to return old column values (the pre-update row).

## UploadedTableData interface

### Syntax

```
public interface UploadedTableData
```

### Remarks

Encapsulates information for one uploaded table for a synchronization.

The insert, update and delete operations are all accessible via the standard ADO.NET IDataReader. The table metadata can be accessed via the GetSchemaTable() call or the insert and delete data readers. The delete data reader only includes the primary key columns of the table.

## GetDeletes method

### Syntax

```
IDataReader GetDeletes();
```

### Remarks

Gets a DataReader with the deletes for this uploaded table data. Each delete is represented by the primary key values needed to uniquely represent a row in this instances table.

Note: The index and order of the columns match the array for property DataTable.PrimaryKey for the schema of this table.

## Returns

A `DataReader` with primary key columns for deleted rows.

## Example

Assume your remote client contains a table called `sparse_pk`. The following example uses the `DownloadTableData.GetDeletes` method to obtain a data reader of deleted rows. In this case, the delete datareader includes two primary key columns. Note the index of each primary key column.

```
CREATE TABLE sparse_pk (
    pcol1 INT NOT NULL,
    col2 VARCHAR(200),
    pcol3 INT NOT NULL,
    PRIMARY KEY (pcol1, pcol3)
);

using iAnywhere.MobiLink.Script;
using System;
using System.IO;
using System.Data;
using System.Text;
...

// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut) {

    // Get an UploadedTableData for the sparse_pk table.
    UploadedTableData sparse_pk_table =
    ut.GetUploadedTableByName("sparse_pk");

    // Get deletes uploaded by the MobiLink client.
    IDataReader data_reader = sparse_pk_table.GetDeletes();

    while (data_reader.Read()) {
        StringBuilder row_str = new StringBuilder(" ( ");
        row_str.Append(data_reader.GetString(0)); // pcol1
        row_str.Append(", ");
        row_str.Append(data_reader.GetString(1)); // pcol3
        row_str.Append(" )");
        writer.WriteLine(row_str);
    }
    data_reader.Close();
}
```

## GetInserts method

### Syntax

```
IDataReader GetInserts();
```

### Remarks

Gets a `DataReader` with the inserts for this uploaded table data. Each Insert is represented by one row returned by the reader.

### Returns

A `DataReader` with inserts for this table data.



**Example**

```

CREATE TABLE sparse_pk (
    pcol1 INT NOT NULL,
    col2 VARCHAR(200),
    pcol3 INT NOT NULL,
    PRIMARY KEY (pcol1, pcol3)
);

using iAnywhere.MobiLink.Script;
using System;
using System.IO;
using System.Data;
using System.Text;
...

// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut) {

    // Get an UploadedTableData for the sparse_pk table.
    UploadedTableData sparse_pk_table =
    ut.GetUploadedTableByName("sparse_pk");

    // Get deletes uploaded by the MobiLink client.
    IDataReader data_reader = sparse_pk_table.GetInserts();

    while (data_reader.Read()) {
        StringBuilder row_str = new StringBuilder("( ");
        row_str.Append(data_reader.GetString(0)); // pcol1
        row_str.Append(", ");
        if (data_reader.IsDBNull(1)) {
            row_str.Append("<NULL>");
        }
        else {
            row_str.Append(data_reader.GetString(1)); // col2
        }
        row_str.Append(", ");
        row_str.Append(data_reader.GetString(2)); // pcol3
        row_str.Append(")");
        writer.WriteLine(row_str);
    }
    data_reader.Close();
}

```

**GetName method****Syntax**

```
string GetName();
```

**Remarks**

Gets the table name of this instance. This is a utility function. The table name can also be accessed via the Schema for this instance.

**Returns**

The table name of this instance.

## GetSchemaTable method

### Syntax

```
DataTable GetSchemaTable();
```

### Remarks

Gets a DataTable that describes the metadata for this download table.

If you want the DataTable to contain column name information, you must specify the client option to send column names.

### Returns

A DataTable that describes the column metadata.

### See also

- dbmlsync: “[SendColumnNames \(scn\) extended option](#)” [*MobiLink - Client Administration*]
- UltraLite: “[Send Column Names synchronization parameter](#)” [*UltraLite - Database Management and Reference*]

## GetUpdates method

### Syntax

```
UpdateDataReader GetUpdates();
```

### Remarks

Gets a DataReader with the updates for this uploaded table data. Each row in the result set represent one update. The mode of the result set can be flipped between new and old column values.

### Returns

A DataReader with updates for this table data.

### Example

The following example illustrates how to use the GetUpdates method.

```
CREATE TABLE sparse_pk (  
    pcol1 INT NOT NULL,  
    col2 VARCHAR(200),  
    pcol3 INT NOT NULL,  
    PRIMARY KEY (pcol1, pcol3)  
);  
  
using iAnywhere.MobiLink.Script;  
using System;  
using System.IO;  
using System.Data;  
using System.Text;  
...
```

```
// The method used for the handle_UploadData event.
public void HandleUpload(UploadData ut) {

    // Get an UploadedTableData for the sparse_pk table.
    UploadedTableData sparse_pk_table =
    ut.GetUploadedTableByName("sparse_pk");

    // Get deletes uploaded by the MobiLink client.
    UpdateDataReader data_reader = sparse_pk_table.GetInserts();

    while (data_reader.Read()) {
        data_reader.SetNewRowValues();
        StringBuilder row_str = new StringBuilder("New values ( ");
        row_str.Append(data_reader.GetString(0)); // pcol1
        row_str.Append(", ");
        if (data_reader.IsDBNull(1)) {
            row_str.Append("<NULL>");
        }
        else {
            row_str.Append(data_reader.GetString(1)); // col2
        }
        row_str.Append(", ");
        row_str.Append(data_reader.GetString(2)); // pcol3
        row_str.Append(")");
        data_reader.SetOldRowValues();
        row_str.Append(" Old Values (");
        row_str.Append(data_reader.GetString(0)); // pcol1
        row_str.Append(", ");
        if (data_reader.IsDBNull(1)) {
            row_str.Append("<NULL>");
        }
        else {
            row_str.Append(data_reader.GetString(1)); // col2
        }
        row_str.Append(", ");
        row_str.Append(data_reader.GetString(2)); // pcol3
        row_str.Append(")");
        writer.WriteLine(row_str);
    }
    data_reader.Close();
}
```

---

---

# Direct row handling

## Contents

Introduction to direct row handling .....	650
Handling direct uploads .....	654
Handling direct downloads .....	660

---

## Introduction to direct row handling

**Note**

Direct row handling is an advanced MobiLink feature. To use it, you must have a thorough understanding of how to create a MobiLink application and how to use the MobiLink APIs. See:

- [MobiLink - Getting Started](#)
- [MobiLink - Server Administration on page 1](#)
- [MobiLink - Client Administration](#)

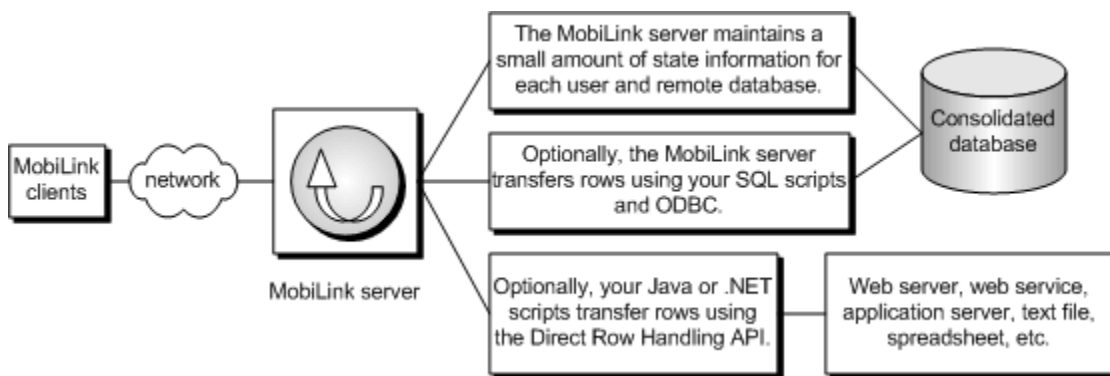
MobiLink supports two ways to handle rows: SQL and direct. You can use them separately or together.

- **SQL row handling** allows you to synchronize remote data to a supported consolidated database. SQL-based events provide a robust interface for conflict resolution and other synchronization tasks. You can use SQL directly or you can return SQL using the MobiLink server APIs for Java and .NET.
- **Direct row handling** allows you to synchronize remote data with any central data source. Direct row handling allows you to access raw synchronized data using special MobiLink events and the MobiLink server APIs for Java and .NET.

The data sources you can synchronize can be virtually anything, including an application, web server, web service, application server, text file, spreadsheet, non-relational database, or an RDBMS that cannot be used as a consolidated database. You still need a consolidated database to store your MobiLink system tables, and many implementations of direct row handling synchronizes to both the consolidated database and another data source.

To use direct row handling, you need familiarity with how to create a MobiLink consolidated database, add synchronization scripts, and create Mobilink remote users.

The following diagram shows the basic MobiLink architecture:



---

## The components of direct row handling

To implement direct row handling, you can use two synchronization events along with several interfaces and methods in the MobiLink server APIs for Java and .NET.

### Direct synchronization events

Direct row handling allows you to directly access the upload stream and download stream. You do this by writing Java or .NET methods for the `handle_UploadData` and `handle_DownloadData` synchronization events.

- **handle\_UploadData** accepts a single `UploadData` parameter that encapsulates operations uploaded by a MobiLink client for a single upload transaction. See:
  - [“Handling direct uploads” on page 654](#)
  - [“handle\\_UploadData connection event” on page 454](#)
- **handle\_DownloadData** allows you to set download operations using the `DownloadData` interface. See:
  - [“Handling direct downloads” on page 660](#)
  - [“handle\\_DownloadData connection event” on page 442](#)

### Components of the MobiLink server API for direct row handling

For the Java API:

- `DBConnectionContext` [“getDownloadData method” on page 544](#)
- [“DownloadData interface” on page 548](#)
- [“DownloadTableData interface” on page 550](#)
- [“UpdateResultSet” on page 578](#)
- [“UploadData interface” on page 579](#)
- [“UploadedTableData interface” on page 581](#)

For the .NET API:

- `DBConnectionContext` [“GetDownloadData method” on page 610](#)
- [“DownloadData interface” on page 622](#)
- [“DownloadTableData interface” on page 623](#)
- [“UpdateDataReader interface” on page 642](#)
- [“UploadedTableData interface” on page 643](#)
- [“UploadData interface” on page 641](#)

## Quick start

To use direct row handling, you need familiarity with how to create a MobiLink consolidated database, add synchronization scripts, and create Mobilink remote users.

To synchronize with a data source other than a consolidated database, complete the following steps.

## Overview of setting up direct row handling

1. Set up a consolidated database, if you do not already have one.

Whether you are synchronizing to a consolidated database, you need to have a consolidated database to hold MobiLink system tables.

See [“MobiLink consolidated databases” on page 3](#).

2. If you want to handle uploads, write a public method using the UploadData interface and register it for the handle\_UploadData connection event.

See [“Handling direct uploads” on page 654](#).

3. If you want to handle downloads, write a public method using the DownloadData interface and register it for the handle\_DownloadData connection event (or another event).

See [“Handling direct downloads” on page 660](#).

4. If you want to use the row handling API to refer to columns by name (rather than by index), specify in your client that column names should be sent with the upload. See:

- SQL Anywhere clients: [“SendColumnNames \(scn\) extended option” \[MobiLink - Client Administration\]](#)
- UltraLite: [“Send Column Names synchronization parameter” \[UltraLite - Database Management and Reference\]](#)

## Other resources for getting started

- [“Tutorial: Introduction to direct row handling” \[MobiLink - Getting Started\]](#)
- <http://www.sybase.com/detail?id=1058600#319>
- [“Setting up Java synchronization logic” on page 529](#)
- [“Setting up .NET synchronization logic” on page 591](#)

You can post questions on the MobiLink newsgroup: [sybase.public.sqlanywhere.mobilink](mailto:sybase.public.sqlanywhere.mobilink).

## Development tips for direct row handling

### Unique primary keys

For MobiLink synchronization, including direct row handling, your data source must have unique primary keys that are not updated. In a non-relational data source such as a spreadsheet or text file, this means that one column must contain unique, unchanging values that identify the row.

See [“Maintaining unique primary keys” on page 139](#).

### Column Names

When using direct row handling, the column names of tables are only available if the Mobilink client is configured to send column names. Alternatively, you can use column indexes to access row information.



To use column names, see:

- SQL Anywhere remotes: “[SendColumnNames \(scn\) extended option](#)” [*MobiLink - Client Administration*]
- UltraLite remotes: “[Send Column Names synchronization parameter](#)” [*UltraLite - Database Management and Reference*]

### Use the last download time for downloads

If possible, set up your direct row handling application like a timestamp-based SQL application; maintain a `last_modified` column and download data based on it. This method avoids unforeseen problems that could occur if you use a different download methodology.

See “[Timestamp-based downloads](#)” on page 129.

### Transaction management for uploads

You cannot commit transactions with the MobiLink consolidated database. However, you can commit transactions with your direct row handling data source. When setting up transaction management, keep the following tips in mind:

- **Commit the upload before MobiLink commits** When applying an upload, MobiLink commits the changes at the end of the `end_upload` event. You should make sure that all upload changes that you want to keep are committed before the end of your `end_upload` script. Otherwise, if there is an error or failure you may get into a state in which your application thinks that the upload is applied but MobiLink has not applied the data, which could result in lost data.
- **Handle redundant uploads** When an error or failure occurs after your application commits an uploaded row and before the MobiLink server commits it, the MobiLink server and your data source may get in an inconsistent state. You can solve this problem by allowing redundant uploads and having logic in place to make sure the redundant upload is applied properly. In particular, when your application sends the upload a second time, it should not be applied again.

### Handle errors

To handle errors, ensure you employ appropriate transaction management, as described above. In addition, your Java or .NET code that handles rows must send any exception that occurs to the MobiLink server. If an error occurs before the MobiLink server or your application has committed changes, MobiLink rolls back the transaction and maintains a consistent state with your application.

### Class instance

For direct row handling, MobiLink creates one class instance per database connection. The class instance is not destroyed at the end of a synchronization: it is destroyed when the database connection is closed. Class level variables retain values from previous synchronizations.

## Handling direct uploads

To handle direct uploads, complete the following steps:

### To handle direct uploads

1. Register a Java or .NET method for the [“handle\\_UploadData connection event” on page 454](#).
2. Write a method for the handle\_UploadData synchronization event. This event accepts one UploadData parameter. See:
  - Java server API: [“UploadData interface” on page 579](#)
  - .NET server API: [“UploadData interface” on page 641](#)

The handle\_UploadData event is usually called once per synchronization. However, for SQL Anywhere clients that use transaction-level uploads, there can be more than one upload per synchronization, in which case handle\_UploadData is called once per transaction.

For more information about dbmlsync transaction-level uploads, see [“-tu option” \[MobiLink - Client Administration\]](#).

For general information about writing Java or .NET synchronization scripts, see:

- [“Writing synchronization scripts in Java” on page 527](#)
- [“Writing synchronization scripts in .NET” on page 589](#)

For information about registering connection-level events, see:

- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)

### Classes for direct uploads

The MobiLink server APIs for Java and .NET provide the following interfaces for handling direct uploads:

- **UploadData** Encapsulates a single upload transaction. An upload transaction contains a set of tables containing row operations. See:
  - Java API: [“UploadData interface” on page 579](#)
  - .NET API: [“UploadData interface” on page 641](#)
- **UploadedTableData** Encapsulates a table's insert, update, and delete operations uploaded by a MobiLink client. For Java, UploadedTableData methods return an instance of an UpdateResultSet. For .NET, UploadedTableData methods return an instance of an UpdateDataReader interface. You traverse the result set IDataReaders to process the uploaded row operations. See:
  - Java API: [“UploadedTableData interface” on page 581](#)
  - .NET API: [“UploadedTableData interface” on page 643](#)
- **UpdateResultSet** For Java, this class represents an update result set returned by the UploadedTableData getUpdates method. It extends java.sql.ResultSet to include special methods for retrieving the new and old versions of an updated row.

See [“UpdateResultSet” on page 578](#).

For .NET, the `UpdateDataReader` interface represents a set of rows returned by the `UploadedTableData.GetUpdates` method. It extends `IDataReader` to include special methods for retrieving the new and old versions of an updated row.

See [“UpdateDataReader interface” on page 642](#).

### Example

See [“handle\\_UploadData connection event” on page 454](#).

## Handling conflicts for direct uploads

When a MobiLink client sends an updated row to the MobiLink server, it includes not only the updated values (the post-image or new row), but also a copy of the old row values (the pre-image or old row) obtained in the last synchronization with the MobiLink server. When the pre-image row does not match the current values in your central data source, a conflict is detected.

### SQL-based conflict resolution

For SQL-based uploads, the MobiLink consolidated database is your central data source and MobiLink provides special events for conflict detection and resolution.

See [“Handling conflicts” on page 146](#).

### Conflict resolution with direct row handling

For direct uploads, you can access new and old rows programmatically for conflict detection and resolution.

`UpdateResultSet` (returned by the `UploadedTableData.getUpdates` method) extends standard Java or .NET result sets to include special methods for handling conflicts. `setNewRowValues` sets `UpdateResultSet` to return new updated values from a remote client (the default mode). `setOldRowValues` sets `UpdateResultSet` to return old row values.

### Detecting conflicts with direct row handling

By using the `UpdateResultSet` method `.setOldRowValues`, you get the values of a row on the remote before it was changed. You compare the row values that are returned to the existing row values in your data source. If the rows you compare are not equal, then a conflict exists.

### Resolving conflicts with direct row handling

Once you have detected a conflict during an upload, you can use custom business logic to resolve the conflict. The resolution is handled by your Java or .NET code.

### Example

Suppose you track inventory in an XML document and want to use it as your central data source. User1 uses one of your remote databases called `Remote1`. User2 uses another remote database called `Remote2`.

Your XML document, User1, and User2 all start with an inventory of ten items. User1 sells three items and updates the `Remote1` inventory value to seven items. User2 sells four items and updates the `Remote2`

inventory to six items. When Remote1 synchronizes, the central database is updated to seven items. When Remote2 synchronizes, a conflict is detected because the value of the inventory is no longer ten items. To resolve this conflict programmatically, you need three row values:

- The current value in the central data source.
- The new row value that Remote2 uploaded.
- The old row value that Remote2 obtained during the last synchronization.

In this case, the business logic would use the following formula to calculate the new inventory value and resolve the conflict:

```
current data source - (old remote - new remote)
-> 7 - (10-6) = 3
```

The following procedures for Java and .NET demonstrate how you can resolve this conflict for direct uploads, using the following table as an example:

```
CREATE TABLE remoteOrders
(
    pk integer primary key not null,
    inventory integer not null
);
```

### To handle direct conflicts (Java)

1. Register a Java or .NET method for the handle\_UploadData connection event.

See [“handle\\_UploadData connection event” on page 454](#).

For example, the following stored procedure call registers a Java method called HandleUpload for the handle\_UploadData connection event when synchronizing the script version ver1. You run this stored procedure against your MobiLink consolidated database.

```
call ml_add_java_connection_script( 'ver1',
    'handle_UploadData',
    'OrderProcessor.HandleUpload' )
```

For more information about registering methods for synchronization events, see:

- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)

2. Obtain an UpdateResultSet for a table in the upload.

The OrderProcessor.HandleUpload method obtains an UpdateResultSet for the remoteOrders table:

```
// method for handle_UploadData event
public void HandleUpload( UploadData u_data )
{
    // Get UploadedTableData for the remoteOrders table.
    UploadedTableData u_table =
    u_data.getUploadedTableByName( "remoteOrders" );

    // Get an UpdateResultSet for the remoteOrders table.
    UpdateResultSet update_rs = u_table.getUpdates();
```

```
// (Continued...)
```

3. For each update, get the current values in your central data source.

In this example, the `UpdateResultSet` `getInt` method returns an integer value for the primary key column (the first column). You can implement and then use the `getMyCentralData` method to get data from your central data source.

```
while( update_rs.next() )
{
    // Get central data source values.

    // Get the primary key value.
    int pk_value = update_rs.getInt(1);

    // Get central data source values.
    int central_value = getMyCentralData(pk_value);

    // (Continued...)
```

4. For each update, get the old and new values uploaded by the MobiLink client.

The example uses the `UpdateResultSet` `setOldRowValues` and `UpdateResultSet` `setNewRowValues` for old and new values, respectively.

```
// Set mode for old row values.
update_rs.setOldRowValues();

// Get the _old_ stored value on the remote.
int old_value = update_rs.getInt(2);

// Set mode for new row values.
update_rs.setNewRowValues();

// Get the _new_ updated value on the remote.
int new_value = update_rs.getInt(2);

// (Continued...)
```

5. For each update, check for conflicts.

A conflict occurs when the old row value does not match the current value in the central data source. To resolve the conflict, a resolved value is calculated using business logic. If no conflict occurs, the central data source is updated with the new remote value. You can implement and then use the `setMyCentralData` method to perform the update.

```
// Check if there is a conflict.

if(old_value == central_value)
{
    // No conflict.
    setMyCentralData(pk_value, new_value);
}
else
{
    // Handle the conflict.
    int inventory = old_value - new_value;
    int resolved_value = central_value - inventory;

    setMyCentralData(pk_value, resolved_value);
```

```
}  
}
```

## To handle direct conflicts (.NET)

1. Register a method for the handle\_UploadData connection event.

For example, the following stored procedure call registers a .NET method called HandleUpload for the handle\_UploadData connection event when synchronizing the script version ver1. You run this stored procedure against your MobiLink consolidated database.

```
call ml_add_dnet_connection_script( 'ver1',  
    'handle_UploadData',  
    'MyScripts.OrderProcessor.HandleUpload' )
```

For more information about registering methods for synchronization events, see:

- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)

2. Obtain an UpdateDataReader for a table in the upload.

The MyScripts.OrderProcessor.HandleUpload method obtains an UpdateResultSet for the remoteOrders table:

```
// method for handle_UploadData event  
public void HandleUpload( UploadData u_data )  
{  
  
    // Get UploadedTableData for the remoteOrders table.  
    UploadedTableData u_table =  
    u_data.GetUploadedTableByName("remoteOrders");  
  
    // Get an UpdateDataReader for the remoteOrders table.  
    UpdateDataReader update_dr = u_table.GetUpdates();  
  
    // (Continued...)
```

3. For each update, get the current values in your central data source.

In this example, the UpdateDataReader GetInt32 method returns an integer value for the primary key column (the first column). You can implement and then use the getMyCentralData method to get data from your central data source.

```
while( update_dr.Read() )  
{  
    // Get central data source values.  
  
    // Get the primary key value.  
    int pk_value = update_dr.GetInt32(0);  
  
    // Get central data source values.  
    int central_value = getMyCentralData(pk_value);  
  
    // (Continued...)
```

4. For each update, get the old and new values uploaded by the MobiLink client.

The example uses the UpdateResultSet setOldRowValues and UpdateResultSet setNewRowValues for old and new values, respectively.

```
// Set mode for old row values.
update_dr.SetOldRowValues();

// Get an _old_ value.
int old_value = update_dr.GetInt32(1);

// Set mode for new row values.
update_dr.SetNewRowValues();

// Get the _new_ updated value.
int new_value = update_dr.GetInt32(1);

// (Continued...)
```

5. For each update, check for conflicts.

A conflict occurs when the old row value does not match the current value in the central data source. To resolve the conflict, a resolved value is calculated using business logic. If no conflict occurs, the central data source is updated with the new remote value. You can implement and then use the `setMyCentralData` method to perform the update.

```
// Check if there is a conflict.
if(old_value == central_value)
{
    // No conflict.
    setMyCentralData(pk_value, new_value);
}
else
{
    // Handle the conflict.
    int inventory = old_value - new_value;
    int resolved_value = central_value - inventory;

    setMyCentralData(pk_value, resolved_value);
}
}
```

## Handling direct downloads

To handle direct downloads, complete the following steps:

### To handle direct downloads

1. Register a Java or .NET method for the [“handle\\_DownloadData connection event” on page 442](#).
2. Write a method for the handle\_DownloadData synchronization event. In this event you use an instance of DBConnectionContext to get a DownloadData instance for the current synchronization. See:
  - Java: [“DBConnectionContext interface” on page 543](#)
  - Java: [“DownloadData interface” on page 548](#)
  - .NET: [“DBConnectionContext interface” on page 609](#)
  - .NET: [“DownloadData interface” on page 622](#)

You can create the entire direct download in the handle\_DownloadData synchronization event. Alternatively, you can use other synchronization events to set direct download operations. However, you must create a handle\_DownloadData script, even if its method does nothing. If you process the direct download in an event other than handle\_DownloadData, the event cannot be before begin\_synchronization and cannot be after end\_download.

For information about the order of events, see [“MobiLink complete event model” on page 346](#).

### Classes for direct downloads

The MobiLink server APIs for Java and .NET provide the following classes for creating direct downloads:

- **DownloadData** Encapsulates download tables containing operations to send down to a remote client during synchronization. See:
  - Java: [“DownloadData interface” on page 548](#)
  - .NET: [“DownloadData interface” on page 622](#)
- **DownloadTableData** Encapsulates upsert (update and insert) and delete operations to download to a MobiLink client.

For Java, DownloadTableData methods return an instance of a JDBC PreparedStatement. In Java, you add a row to the download by setting the prepared statement's column values and then executing the prepared statement.

For .NET, DownloadTableData methods return an instance of a .NET IDbCommand. In .NET, you add a row to the download by setting the command's column values and then executing the command.

See:

- Java: [“DownloadTableData interface” on page 550](#)
- .NET: [“DownloadTableData interface” on page 623](#)

### Example

See [“handle\\_DownloadData connection event” on page 442](#).



# MobiLink Reference

This section contains MobiLink reference material.

---

MobiLink server system procedures .....	663
MobiLink utilities .....	687
MobiLink server system tables .....	693
MobiLink data mappings between remote and consolidated databases .....	739
Character set considerations .....	789
iAnywhere Solutions ODBC drivers for MobiLink .....	793
Deploying MobiLink applications .....	799



---

# MobiLink server system procedures

## Contents

MobiLink system procedures ..... 664

---

## MobiLink system procedures

MobiLink provides the following stored procedures to help you create your applications.

### System procedures to add or delete scripts

You must add synchronization scripts to system tables in the consolidated database before you can use them. The following system procedures add or delete synchronization scripts in the consolidated database:

- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)

When you use the MobiLink server API for Java or .NET, you use these stored procedures to register a method as the script for an event, so that the method is run when the event occurs. You can also use them to unregister your methods.

When you add a script using a system procedure, the script is a string. Any strings within the script need to be escaped. For SQL Anywhere, each quotation mark (') needs to be doubled so as not to terminate the string.

You cannot use system procedures to add scripts longer than 255 bytes to Adaptive Server Enterprise 11.5 or earlier. Instead, use Sybase Central or direct insertion to define longer scripts.

DB2 mainframe version 8.1 supports a backward compatibility mode, where column names and other identifiers are limited to a maximum of 18 characters. To support this environment, all MobiLink system objects in DB2 mainframe have names of 18 characters or less. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

IBM DB2 LUW prior to version 6 only supports column names and other identifiers of 18 characters or less, and so the names are truncated. For example, ml\_add\_connection\_script is shortened to ml\_add\_connection\_.

### Other system procedures

- [“ml\\_add\\_property system procedure” on page 677](#)
- [“ml\\_delete\\_sync\\_state\\_before system procedure” on page 684](#)
- [“ml\\_reset\\_sync\\_state system procedure” on page 686](#)

## IBM DB2 mainframe system procedure name conversions

IBM DB2 mainframe consolidated databases only support column names and other identifiers of 18 characters or less. The following table identifies how system procedure names for DB2 mainframe consolidated databases are mapped to system procedure names for all other consolidated database types.

If the system procedure name does not appear in the table below, no conversion is required.

<b>System procedure name</b>	<b>System procedure name for DB2 mainframe consolidated databases</b>
“ml_add_connection_script system procedure” on page 667	ml_add_cs
“ml_add_dnet_connection_script system procedure” on page 668	ml_add_dcs
“ml_add_dnet_table_script system procedure” on page 669	ml_add_dts
“ml_add_java_connection_script system procedure” on page 671	ml_add_jcs
“ml_add_java_table_script system procedure” on page 672	ml_add_jts
“ml_add_lang_connection_script system procedure” on page 673	ml_add_lcs
“ml_add_lang_connection_script_chk system procedure” on page 673	ml_add_lcs_chk
“ml_add_lang_table_script system procedure” on page 673	ml_add_lts
“ml_add_lang_table_script_chk system procedure” on page 673	ml_add_lts_chk
“ml_add_passthrough system procedure” on page 673	ml_add_pt
“ml_add_passthrough_repair system procedure” on page 674	ml_add_pt_repair
“ml_add_passthrough_script system procedure” on page 676	ml_add_pt_script
“ml_add_table_script system procedure” on page 680	ml_add_ts
“ml_delete_passthrough system procedure” on page 681	ml_del_pt
“ml_delete_passthrough_repair system procedure” on page 682	ml_del_pt_repair

System procedure name	System procedure name for DB2 mainframe consolidated databases
<a href="#">“ml_delete_passthrough_script system procedure” on page 683</a>	ml_del_pt_script
<a href="#">“ml_delete_sync_state system procedure” on page 683</a>	ml_del_sstate
<a href="#">“ml_delete_sync_state_before system procedure” on page 684</a>	ml_del_sstate_b4
<a href="#">“ml_reset_sync_state system procedure” on page 686</a>	ml_reset_sstate

## ml\_add\_column system procedure

Registers information about columns on remote databases for use by named column parameters.

### Syntax

```
ml_add_column (
  'version',
  'table',
  'column',
  'type'
)
```

### Parameters

Syntax	Description
version	VARCHAR(128). The version name.
table	VARCHAR(128). The table name.
column	VARCHAR(128). The column name.
type	VARCHAR(128). Reserved for future use. Set to null.

### Remarks

This procedure populates the ml\_column MobiLink system table with information about the columns on the remote database. The information is used by named row parameters.

#### Caution

ml\_add\_column calls must be executed in the same order that the columns exist in the remote database table. Failing to do so may result in incorrect data.

You need to run this system procedure if both of the following are true:

- Your SQL scripts contain named parameters for columns (for example, o.column-name and r.column-name).
- You are not using the **Create Synchronization Model Wizard**.

Even if you are using the **Create Synchronization Model Wizard**, if you modify the remote schema outside Model mode, you need to use this stored procedure to send information about columns that are not registered in ml\_column.

To delete all entries for the table name in the given script version, set the column name to null.

### See also

- [“ml\\_column” on page 698](#)
- [“Script parameters” on page 320](#)

### Examples

The following stored procedure call populates the ml\_column MobiLink system table for col1 in MyTable for the script version Version1. This call allows you to use the named row parameters r.col1 and o.col1 in table scripts for MyTable1 in the Version1 script version.

```
CALL ml_add_column( 'Version1', 'MyTable1', 'col1', NULL )
```

The following stored procedure call deletes all entries in the ml\_column MobiLink system table for MyTable1 in script version Version1:

```
CALL ml_add_column( 'Version1', 'MyTable1', NULL, NULL )
```

## ml\_add\_connection\_script system procedure

Use this system procedure to add or delete SQL connection scripts in the consolidated database.

### Syntax

```
ml_add_connection_script (
  'version',
  'event',
  'script'
)
```

### Parameters

Syntax	Description
version	VARCHAR(128). The version name.
event	VARCHAR(128). The event name.

Syntax	Description
script	TEXT. The script contents. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

**Remarks**

To delete a connection script, set the script contents parameter to null.

When you add a script, the script is inserted into the ml\_script table and the appropriate references are defined to associate the script with the event and script version that you specify. If the version name is new, it is automatically inserted into the ml\_version table.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_cs. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“System procedures to add or delete scripts” on page 664](#)
- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)

**Example**

The following statement adds a connection script associated with the begin\_synchronization event to the script version custdb in a SQL Anywhere consolidated database. The script itself is the single statement that sets the @EmployeeID variable.

```
call ml_add_connection_script( 'custdb',
    'begin_synchronization',
    'set @EmployeeID = {ml s.username}' )
```

## ml\_add\_dnet\_connection\_script system procedure

Use this system procedure to register or unregister a .NET method as the script for a connection event.

**Syntax**

```
ml_add_dnet_connection_script (
    'version',
    'event',
    'script'
)
```



**Parameters**

Syntax	Description
version	VARCHAR(128). The version name.
event	VARCHAR(128). The event name.
script	TEXT. The script contents. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

**Remarks**

To unregister a method, set the script contents parameter to null.

The script contents value is a public method in a class in a .NET assembly (for example, MyClass.MyMethod).

When you call `ml_add_dnet_connection_script`, the method is associated with the event and script version that you specify. If the version name is new, it is automatically inserted into the `ml_version` table.

For DB2 mainframe consolidated database types, this procedure is called `ml_add_dcs`. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“System procedures to add or delete scripts” on page 664](#)
- [“Adding and deleting scripts” on page 327](#)
- [“IBM DB2 mainframe system procedure name conversions” on page 664](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)
- [“Methods” on page 595](#)
- [“Writing synchronization scripts in .NET” on page 589](#)

**Example**

The following example registers the `beginDownloadConnection` method of the `ExampleClass` class for the `begin_download` event.

```
call ml_add_dnet_connection_script( 'ver1',
  'begin_download',
  'ExamplePackage.ExampleClass.beginDownloadConnection' );
```

**ml\_add\_dnet\_table\_script system procedure**

Use this system procedure to register or unregister a .NET method as the script for a table event.

**Syntax**

```
ml_add_dnet_table_script (
  'version',
  'table',
  'event',
  'script'
)
```

**Parameters**

Syntax	Description
version	VARCHAR(128). The version name.
table	VARCHAR(128). The table name.
event	VARCHAR(128). The event name.
script	TEXT. The script contents. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

**Remarks**

To unregister a method, set the script contents parameter to null.

The script value is a public method in a class in a .NET assembly (for example, MyClass.MyMethod).

When you call ml\_add\_dnet\_table\_script, the method is associated with the table, event, and script version that you specify. If the version name is new, it is automatically inserted into the ml\_version table.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_dts. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“System procedures to add or delete scripts” on page 664](#)
- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“Methods” on page 595](#)
- [“Writing synchronization scripts in .NET” on page 589](#)

**Example**

The following example assigns the empDownloadCursor method of the EgClass class to the download\_cursor event for the table emp.

```
call ml_add_dnet_table_script( 'ver1', 'emp',
  'download_cursor', EgPackage.EgClass.empDownloadCursor' )
```

## ml\_add\_java\_connection\_script system procedure

Use this system procedure to register or unregister a Java method as the script for a connection event.

### Syntax

```
ml_add_java_connection_script (
  'version',
  'event',
  'script'
)
```

### Parameters

Syntax	Description
version	VARCHAR(128). The version name.
event	VARCHAR(128). The event name.
script	TEXT. The script contents. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

### Remarks

To unregister a method, set the script contents parameter to null.

The script value is a public method in a class in the MobiLink server classpath (for example, MyClass.MyMethod).

When you ml\_add\_java\_connection\_script, the method is associated with the event and script version that you specify. If the version name is new, it is automatically inserted into the ml\_version table.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_jcs. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

### See also

- [“System procedures to add or delete scripts” on page 664](#)
- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_java\\_table\\_script system procedure” on page 672](#)
- [“Methods” on page 533](#)
- [“Writing synchronization scripts in Java” on page 527](#)

### Example

The following example registers the endConnection method of the CustEmpScripts class for the end\_connection event.

```
call ml_add_java_connection_script( 'ver1',
    'end_connection',
    'CustEmpScripts.endConnection' )
```

## ml\_add\_java\_table\_script system procedure

Use this system procedure to register or unregister a Java method as the script for a table event.

### Syntax

```
ml_add_java_table_script (
    'version',
    'table',
    'event',
    'script'
)
```

### Parameters

Syntax	Description
version	VARCHAR(128). The version name.
table	VARCHAR(128). table name.
event	VARCHAR(128). The event name.
script	TEXT. The script content. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

### Remarks

To unregister a method, set the script content parameter to null.

The *script* value is a public method in a class in the MobiLink server classpath (for example, MyClass.MyMethod).

When you call ml\_add\_java\_table\_script, the method is associated with the table, event, and script version that you specify. If the version name is new, it is automatically inserted into the ml\_version table.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_jts. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“System procedures to add or delete scripts” on page 664](#)
- [“Adding and deleting scripts” on page 327](#)
- [“ml\\_add\\_connection\\_script system procedure” on page 667](#)
- [“ml\\_add\\_table\\_script system procedure” on page 680](#)
- [“ml\\_add\\_dnet\\_connection\\_script system procedure” on page 668](#)
- [“ml\\_add\\_dnet\\_table\\_script system procedure” on page 669](#)
- [“ml\\_add\\_java\\_connection\\_script system procedure” on page 671](#)
- [“Methods” on page 533](#)
- [“Writing synchronization scripts in Java” on page 527](#)

**Example**

The following example registers the empDownloadCursor method of the CustEmpScripts class for the download\_cursor event for the table emp.

```
call ml_add_java_table_script( 'ver1', 'emp',  
    'download_cursor', 'CustEmpScripts.empDownloadCursor' )
```

## ml\_add\_lang\_connection\_script system procedure

This procedure is for internal use only.

## ml\_add\_lang\_connection\_script\_chk system procedure

This procedure is for internal use only.

## ml\_add\_lang\_table\_script system procedure

This procedure is for internal use only.

## ml\_add\_lang\_table\_script\_chk system procedure

This procedure is for internal use only.

## ml\_add\_passthrough system procedure

Use this system procedure to identify remote databases that should execute a script. This procedure adds an entry to the ml\_passthrough system table. If an entry with the given remote\_id and run\_order already exists in the table, this procedure updates the entry.

### Syntax

```
ml_add_passthrough (
  'remote_id',
  'script_name',
  run_order
)
```

### Parameters

Syntax	Description
remote_id	VARCHAR(128). The remote ID of the database that should execute the script. This value can be a valid remote ID in the ml_database table to apply to a specific client, or null to apply to all the script clients listed in the ml_database table. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>Caution</b> Be very careful when applying a script to all, or even many, remotes. A poorly written script can leave most of even all of your remotes damaged or disabled.</p> </div>
script_name	VARCHAR(128). The name of the script being subscribed to. This value must be a valid script name defined in the ml_passthrough_script table.
run_order	INTEGER. The run_order parameter determines the order in which scripts are applied on the remote database. Scripts are always applied in order by run_order. Each remote stores the run_order of the last script that it attempted to apply and does not download or execute any script with a run_order less than this.  This value must be a non-negative integer or null.

### Remarks

If you define run\_order as null, the procedure assigns an integer based on the value of remote\_id. If remote\_id is null, the procedure assigns a value equal to the run\_order value in ml\_passthrough, plus 10. If remote\_id is not null, the procedure assigns the maximum value of the run\_order column for the remote\_id in the ml\_passthrough table plus 10.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_pt. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

### See also

- [“ml\\_database” on page 700](#)
- [“ml\\_passthrough” on page 707](#)
- [“ml\\_passthrough\\_script” on page 709](#)

## ml\_add\_passthrough\_repair system procedure

Use this system procedure to define rules for handling script errors. Each rule defines the action that a client should perform when a specific script generates a given error code. This procedure adds an entry to the

ml\_passthrough\_repair system table. If an entry with the given failed\_script\_name and error\_code already exists in the table, the procedure updates the entry.

### Syntax

```
ml_add_passthrough_repair (
  'failed_script_name',
  error_code,
  'new_script_name',
  'action'
)
```

### Parameters

Syntax	Description
failed_script_name	VARCHAR(128). The name of the failed script to which this rule applies. This value must be a valid script name in the ml_passthrough_script table.
error_code	INTEGER. The SQL Anywhere error code that this rule handles.
new_script_name	VARCHAR(128). The name of a script to replace the failed script when action is R. If action is S, P, or H, this value must be null. If action is R, this value must be a valid script name in the ml_passthrough_script table, and can be the same as failed_script_name.
action	CHAR(1). The action that a client should perform when error_code is generated for failed_script_name. This value must be one of the following: <ul style="list-style-type: none"> <li>• <b>R</b> (replace) Indicates that the failed script should be replaced with the one specified by <b>new script name</b> and an attempt should be made to run the new script. To rerun the failed script, choose <b>new script name</b> to be the same as <b>failed script name</b>.</li> <li>• <b>P</b> (purge) Indicates that the remote database should discard all the scripts that it has received and continue executing script normally after that.</li> <li>• <b>S</b> (skip) Indicates that the remote database should ignore the failed script and continue executing scripts as if the failed script had succeeded.</li> <li>• <b>H</b> (halt) Indicates that the remote should not execute any more scripts until it receives further instructions.</li> </ul>

### Remarks

You should make every effort to avoid failed SQL passthrough scripts by testing scripts thoroughly.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_pt\_repair. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“ml\\_passthrough\\_repair” on page 708](#)
- [“ml\\_passthrough\\_script” on page 709](#)

## ml\_add\_passthrough\_script system procedure

Use this system procedure to create a passthrough script. This procedure adds an entry to the ml\_passthrough\_script system table.

**Syntax**

```
ml_add_passthrough_script (
  'script_name',
  'flags',
  'affected_pubs',
  'script',
  'description'
)
```

**Parameters**

Syntax	Description
script_name	VARCHAR(128). The script name. This value must be unique.
flags	<p>VARCHAR(256). The value that tells clients how to run the script. This value can be null or contain a combination of the following keywords in a semicolon-delimited list:</p> <ul style="list-style-type: none"> <li>• <b>manual</b> Indicates that the script may only be run in manual execution mode. By default, all scripts can be run in either automatic or manual execution modes.</li> <li>• <b>exclusive</b> Indicates that the script may only be automatically executed at the end of a synchronization where exclusive locks were obtained on all tables being synchronized. This option is ignored if the affected_publications value lists no publications. This option is only meaningful to SQL Anywhere remotes.</li> <li>• <b>schema_diff</b> Indicates that the script should be run in schema-diffing mode. In this mode, the database schema is altered to match the schema described in the script. For example, a create statement for an existing table is treated as an alter statement. This flag only applies to scripts run on UltraLite remotes.</li> </ul> <p>For example:</p> <pre>'manual;exclusive;schema_diff'</pre>



Syntax	Description
affected_pubs	TEXT. A list of publications that must be synchronized before the script is run. An empty string or null indicates that no synchronization is required. This value is only meaningful for SQL Anywhere clients. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.
script	<p>TEXT. The contents of the passthrough script. This value cannot be null. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.</p> <p>The script content must be non-null. For UltraLite remotes, the <b>script</b> content should be a collection of SQL statements separated by the word <b>go</b>. Note that the word <b>go</b> must appear on a separate line. For SQL Anywhere remotes, the <b>script</b> content can be any collection of SQL statements that are valid when enclosed by a <b>begin... end</b> block.</p> <p>Example of <b>script</b> content on a SQL Anywhere remote:</p> <pre>DECLARE val INTEGER; SELECT c1 INTO val FROM t1 WHERE pk = 5; IF val &gt; 100 THEN     INSERT INTO t2 VALUES ('c1 is big'); ENDIF</pre> <p>Example of <b>script</b> content on an UltraLite remote:</p> <pre>CREATE TABLE myScript (c1 INT NOT NULL PRIMARY KEY) GO INSERT INTO myScript VALUES (1) GO</pre>
description	VARCHAR(2000). A comment or description of the script. This value may be null.

**Remarks**

This procedure generates an error if the specified script\_name already exists in ml\_passthrough\_script.

For DB2 mainframe consolidated database types, this procedure is called ml\_add\_pt\_script. See [“IBM DB2 mainframe system procedure name conversions”](#) on page 664.

**See also**

- [“ml\\_passthrough\\_script”](#) on page 709

## ml\_add\_property system procedure

Use this system procedure to add or delete MobiLink properties. This system procedure changes rows in the ml\_property system table.

**Syntax**

```
ml_add_property (
  'comp_name',
  'prop_set_name',
  'prop_name',
  'prop_value'
)
```

**Parameters**

Syntax	Description
comp_name	<p>VARCHAR(128). The component name. To save properties by script version, set to ScriptVersion. For MobiLink server properties, set to MLS. For server-initiated synchronization properties, set to SIS.</p>
prop_set_name	<p>VARCHAR(128). The property set name.</p> <p>If the component name is ScriptVersion, then this parameter is the name of the script version.</p> <p>If the component name is MLS, then this parameter can be either ml_user_log_verbosity to specify verbosity for a MobiLink user, or ml_remote_id_log_verbosity to specify verbosity for a remote ID.</p> <p>If the component name is SIS, then this parameter is the name of the Notifier, gateway, or carrier that you are setting a property for.</p>
prop_name	<p>VARCHAR(128). The property name.</p> <p>If the component name is ScriptVersion, then this parameter is a property that you define. You can reference these properties using DBConnectionContext: getVersion and getProperties, or ServerContext: getPropertiesByVersion, getProperties, and getPropertySetNames.</p> <p>If the component name is MLS, then this property is either a MobiLink user name or remote ID that you define.</p>
prop_value	<p>TEXT. The property value.</p> <p>If the prop_set_name is ml_user_log_verbosity or ml_remote_id_log_verbosity, this must be a valid mlsrv -v option.</p> <p>For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB. To delete a property, set to null.</p>

**Log verbosity for targeted MobiLink users and remote IDs**

The MobiLink server can be set to use different log verbosity for a targeted MobiLink user or remote ID. The MobiLink server checks the ml\_property table every five minutes and looks for verbose settings for a MobiLink user or remote ID. If verbose settings exist, then it uses the new setting to log output messages for the given MobiLink user or remote ID. This enables you to see the details for a specific user or remote

ID without the need for high verbosity settings that would negatively impact the server farm, and without requiring a restart of each server in the farm.

To set maximum verbosity for a targeted MobiLink user, for example *ml\_user1*, log into the consolidated database and issue the following SQL command:

```
call ml_add_property( 'MLS', 'ml_user_log_verbosity', 'ml_user1', '-v+' )
```

To set maximum verbosity for a targeted remote ID, for example *rid\_1*, log into the consolidated database and issue the following SQL command:

```
call ml_add_property( 'MLS', 'ml_remote_id_log_verbosity', 'rid_1', '-v+' )
```

Note that *verbose\_setting* must be a valid MobiLink server *-v* option. For example, to log row data and undefined table scripts, the *verbose\_setting* can be *-vru* or *vru*. The MobiLink server will use this verbose setting for *ml\_user1* or *rid\_1* after 5 minutes. See “[-v option](#)” on page 102.

To disable log verbosity for a MobiLink user, log into the consolidated database and issue the following SQL command:

```
call ml_add_property( 'MLS', 'ml_user_log_verbosity', 'ml_user', NULL )
```

To disable log verbosity for a MobiLink remote ID, log into the consolidated database and issue the following SQL command:

```
call ml_add_property( 'MLS', 'ml_remote_id_log_verbosity', 'rid_1', NULL )
```

The MobiLink server stops using the previous verbose setting for *ml\_user* or *rid\_1* after five minutes.

If **both** the **ml\_user\_log\_verbosity** and **ml\_remote\_id\_log\_verbosity** are set for a given MobiLink user and remote ID, and if the MobiLink user name and remote ID in a synchronization are identical to the given targeted MobiLink user and remote ID, the MobiLink server will use the **ml\_remote\_id\_log\_verbosity** setting to log output messages.

### Server-initiated synchronization

For server-initiated synchronization, the `ml_add_property` system procedure allows you to set properties for Notifiers, gateways, and carriers.

For example, to add the property `server=mailserver1` for an SMTP gateway called `x`:

```
ml_add_property( 'SIS', 'SMTP(x)', 'server', 'mailserver1' );
```

The verbosity property applies to all Notifiers and gateways, and so you cannot specify a particular property set name. To change the verbosity setting, leave the property set name blank:

```
ml_add_property( 'SIS', '', 'verbosity', 2 );
```

### Script Version

For regular MobiLink synchronization, you can use this system procedure to associate properties with a script version. In this case, set the `component_name` to `ScriptVersion`. You can specify any properties, and use Java and .NET classes to access them.

For example, to associate an LDAP server with a script version called `MyVersion`:

```
ml_add_property( 'ScriptVersion', 'MyVersion', 'ldap-server', 'MyServer' )
```

**See also**

- [“ml\\_property” on page 712](#)
- [“MobiLink server settings for server-initiated synchronization” \[\*MobiLink - Server-Initiated Synchronization\*\]](#)
- [“MobiLink server settings for server-initiated synchronization” \[\*MobiLink - Server-Initiated Synchronization\*\]](#)
- Java API DBConnectionContext: [“getProperties method” on page 545](#) and [“getVersion method” on page 547](#)
- .NET API DBConnectionContext: [“GetProperties method” on page 611](#) and [“GetVersion method” on page 612](#)
- Java API ServerContext: [“getPropertiesByVersion method” on page 569](#), [“getProperties method” on page 569](#), [“getPropertySetNames method” on page 570](#)
- .NET ServerContext: [“getPropertiesByVersion method” on page 631](#), [“getProperties method” on page 631](#), [“getPropertySetNames method” on page 631](#)

## ml\_add\_table\_script system procedure

Use this system procedure to add or delete SQL table scripts in the consolidated database.

**Syntax**

```
ml_add_table_script (
    'version',
    'table',
    'event',
    'script'
)
```

**Parameters**

Syntax	Description
version	VARCHAR(128). The version name.
table	VARCHAR(128). The table name.
event	VARCHAR(128). The event name.
script	TEXT. The script contents. For Adaptive Server Enterprise, this parameter is VARCHAR(16384). For DB2 LUW, this parameter is VARCHAR(4000). For Oracle, this parameter is CLOB.

**Remarks**

To delete a table script, set the script contents parameter to null.

When you add a script, the script is inserted into the ml\_script table and the appropriate references are defined to associate the script with the table, event and script version that you specify. If the version name is new, it is automatically inserted into the ml\_version table.

For DB2 mainframe consolidated database types, this procedure is called `ml_add_ts`. See [“IBM DB2 mainframe system procedure name conversions”](#) on page 664.

### See also

- [“System procedures to add or delete scripts”](#) on page 664
- [“Adding and deleting scripts”](#) on page 327
- [“ml\\_add\\_connection\\_script system procedure”](#) on page 667
- [“ml\\_add\\_dnet\\_connection\\_script system procedure”](#) on page 668
- [“ml\\_add\\_dnet\\_table\\_script system procedure”](#) on page 669
- [“ml\\_add\\_java\\_connection\\_script system procedure”](#) on page 671
- [“ml\\_add\\_java\\_table\\_script system procedure”](#) on page 672

### Example

The following command adds a table script associated with the `upload_insert` event on the Customer table.

```
call ml_add_table_script( 'default', 'Customer', 'upload_insert',
  'INSERT INTO Customer( cust_id, name, rep_id, active )
  VALUES ( {ml r.cust_id}, {ml r.name}, {ml r.rep_id}, 1 )' )
```

## ml\_add\_user system procedure

This procedure is for internal use only.

## ml\_delete\_passthrough system procedure

This stored procedure removes the row(s) in the `ml_passthrough` table that cause the specified script to be downloaded to the specified remote with the specified run order. If the script is downloaded to the remote before it is deleted then it is not deleted from the remote and executes as usual.

### Syntax

```
ml_delete_passthrough (
  'remote_id',
  'script_name',
  'run_order'
)
```

### Parameters

Syntax	Description
<code>remote_id</code>	VARCHAR(128). The remote ID. If <b>remote_id</b> is null then all rows in the <code>ml_passthrough</code> table for the specified script name and run order are removed.
<code>script_name</code>	VARCHAR(128). The script name.

Syntax	Description
run_order	INTEGER. The run order of the script applied on the remote database. If <b>run_order</b> is null then all rows for the specified <b>remote_id</b> and <b>script_name</b> are removed from the ml_passthrough table regardless of their run order.

**Remarks**

The MobiLink server does not automatically remove entries from the ml\_passthrough table. You must use this procedure to remove outdated passthrough scripts.

For DB2 mainframe consolidated database types, this procedure is called ml\_del\_pt. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“ml\\_passthrough” on page 707](#)

## ml\_delete\_passthrough\_repair system procedure

Use this system procedure to delete a repair rule from the ml\_passthrough\_repair system table.

**Syntax**

```
ml_delete_passthrough_repair (
    'failed_script_name',
    error_code
)
```

**Parameters**

Syntax	Description
failed_script_name	VARCHAR(128). The name of the script to which a rule applied.
error_code	INTEGER. The error code for which the rule applied.

**Remarks**

The MobiLink server does not automatically remove entries from the ml\_passthrough\_repair table. You must use this procedure to remove outdated passthrough repair scripts.

For DB2 mainframe consolidated database types, this procedure is called ml\_del\_pt\_repair. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

**See also**

- [“ml\\_passthrough\\_repair” on page 708](#)

## ml\_delete\_passthrough\_script system procedure

Use this system procedure to delete a passthrough script from the ml\_passthrough\_script system table.

### Syntax

```
ml_delete_passthrough_script (
  'script_name'
)
```

### Parameters

Syntax	Description
script_name	VARCHAR(128). The name of the script to remove.

### Remarks

Scripts can not be removed if they are referenced in the ml\_passthrough or ml\_passthrough\_repair system tables.

The MobiLink server does not automatically remove entries from the ml\_passthrough\_script table. You must use this procedure to remove outdated passthrough scripts.

For DB2 mainframe consolidated database types, this procedure is called ml\_del\_pt\_script. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

### See also

- [“ml\\_passthrough” on page 707](#)
- [“ml\\_passthrough\\_repair” on page 708](#)
- [“ml\\_passthrough\\_script” on page 709](#)

## ml\_delete\_sync\_state system procedure

Use this procedure to delete unused or unwanted synchronization states.

### Syntax

```
ml_delete_sync_state (
  'user',
  'remote_id'
)
```

### Parameters

Syntax	Description
user	VARCHAR(128). The MobiLink user name.
remote_id	VARCHAR(128). The remote ID.

## Remarks

These parameters can be null. If all the parameters are null, the procedure does nothing.

This stored procedure deletes all the rows from the `ml_subscription` table for the given MobiLink user name and remote ID. It also removes this remote ID from the `ml_database` table, if the remote ID is no longer referenced by any rows in the `ml_subscription` table.

If the remote ID is null and the MobiLink user name is not null, it removes all the rows that are referenced by the given MobiLink user name from the `ml_subscription` table and all the remote IDs from the `ml_database` table, if these remote IDs are no longer referenced by any rows in the `ml_subscription` table.

If the MobiLink user name is null and the remote ID is not null, this stored procedure removes all the rows from the `ml_subscription` table and the `ml_database` table for the given remote ID.

The MobiLink user is not removed by this stored procedure, even if all the remote IDs have been deleted from the `ml_database` table and this user is no longer referenced by any rows in the `ml_subscription` table. If this MobiLink user needs to be deleted, you may delete it by issuing a command such as

```
delete * from ml_user where name = 'user_name'
```

where `user_name` is the MobiLink user you want to delete.

Use this stored procedure with extreme caution because the MobiLink server automatically adds this remote ID in the `ml_database` and `ml_subscription` tables without checking its synchronization status the next time the MobiLink client requests synchronization for this remote ID. It may cause data inconsistency to delete synchronization states for a remote ID that did not have a successful synchronization in the last synchronization attempt.

For DB2 mainframe consolidated database types, this procedure is called `ml_del_sstate`. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

## See also

- [“ml\\_subscription” on page 733](#)
- [“ml\\_database” on page 700](#)

## Example

The following example cleans up MobiLink system table information about remote databases with the remote ID `remote_db_for_John` for the MobiLink user John:

```
CALL ml_delete_sync_state( 'John', 'remote_db_for_John' )
```

## **ml\_delete\_sync\_state\_before system procedure**

Use this procedure to clean up the MobiLink system tables when you have dropped remote databases.

## Syntax

```
ml_delete_sync_state_before (  
  'ts'  
)
```



## Parameters

Syntax	Description
ts	TIMESTAMP. The datetime must appear in exactly the order specified in the consolidated database. If the datetime format in the consolidated database is set to 'yyyy/mm/dd hh:mm:ss.ssss', then the timestamp must appear in the order year, month, day, hour, minute, second, fraction of second.

## Remarks

This stored procedure removes rows from MobiLink system tables that pertain to remote databases that are no longer being used. In particular, it does the following:

- Deletes all the rows from the ml\_subscription system table that have both the last\_upload\_time and last\_download\_time earlier than the given timestamp.
- Removes remote IDs from the ml\_database system table if the remote IDs are no longer referenced by any rows in the ml\_subscription table.

You should not use this system procedure for a time period that is so recent that it may delete rows for remote databases that have not actually been deleted. If you do, the deletion of the rows in ml\_subscription and ml\_database could cause problems for remote databases that are in an "unknown state" caused by an unsuccessful upload; in that unknown state, the remote relies on the MobiLink system tables to resend data.

The timestamp provided to this procedure must have a correct date-time format because the procedure does not validate the date-time format of the parameter.

For DB2 mainframe consolidated database types, this procedure is called ml\_del\_sstate\_b4. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

## See also

- [“ml\\_subscription” on page 733](#)
- [“ml\\_database” on page 700](#)

## Example

The following example cleans up MobiLink system table information about remote databases that have not synchronized since January 10, 2004. It works for a SQL Anywhere consolidated database where the date-time format in the consolidated database is yyyy/mm/dd hh:mm:ss.ssss.

```
CALL ml_delete_sync_state_before( '2004/01/10 00:00:00' )
```

## ml\_delete\_user system procedure

This procedure is for internal use only.

## ml\_reset\_sync\_state system procedure

Use this procedure to reset synchronization state information in MobiLink system tables.

### Syntax

```
ml_reset_sync_state (  
  'user',  
  'remote_id'  
)
```

### Parameters

Syntax	Description
user	VARCHAR(128). The MobiLink user name.
remote_id	VARCHAR(128). The remote ID.

### Remarks

The parameters can be null. If both parameters are null, this procedure does nothing.

This stored procedure sets the progress, last\_upload\_time, and last\_download\_time columns in the ml\_subscription table to their default values for the given user\_name and remote ID. The default value for the progress is 0. The default value for the last\_upload\_time and last\_download\_time columns is '1900/01/01 00:00:00'.

If the remote ID is null and the MobiLink user name is not null, this procedure sets those columns to the default values for the rows in the ml\_subscription table referenced by the given MobiLink user name. If the MobiLink user name is null and the remote ID is not null, it sets them to the default values for the rows in the ml\_subscription table with the given remote ID.

Use this stored procedure with extreme caution. The MobiLink server does not do any synchronization status checking for this remote ID the next time the MobiLink client requests synchronization for this remote ID. It may cause data inconsistency to reset a remote ID that did not have a successful synchronization in the last synchronization.

For DB2 mainframe consolidated database types, this procedure is called ml\_reset\_sstate. See [“IBM DB2 mainframe system procedure name conversions” on page 664](#).

## ml\_server\_delete system procedure

This procedure is for internal use only.

## ml\_server\_update system procedure

This procedure is for internal use only.

---

# MobiLink utilities

## Contents

Introduction to MobiLink utilities .....	688
MobiLink stop utility (mlstop) .....	689
MobiLink user authentication utility (mluser) .....	690

---

## Introduction to MobiLink utilities

There are two MobiLink server utilities:

- [“MobiLink stop utility \(mlstop\)” on page 689](#)
- [“MobiLink user authentication utility \(mluser\)” on page 690](#)

In addition, see:

- [MobiLink client utilities: “MobiLink client utilities” \[\*MobiLink - Client Administration\*\]](#)
- [UltraLite utilities: “UltraLite utilities” \[\*UltraLite - Database Management and Reference\*\]](#)
- [Utilities for using TLS certificates: “Certificate utilities” \[\*SQL Anywhere Server - Database Administration\*\]](#)
- [Other SQL Anywhere utilities: “Database administration utilities” \[\*SQL Anywhere Server - Database Administration\*\]](#)

## MobiLink stop utility (mlstop)

Stops the MobiLink server on the local computer.

### Syntax

**mlstop** [ *options* ] [ *name* ]

Option	Description
@ <i>data</i>	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used. See <a href="#">“Using configuration files” [SQL Anywhere Server - Database Administration]</a> .  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 <a href="#">“File Hiding utility (dbfhide)” [SQL Anywhere Server - Database Administration]</a> .
<b>-f</b>	Forced shutdown. Use if a hard shutdown does not work.
<b>-h</b>	Hard shutdown. MobiLink stops all synchronizations and exits. Some remotes may report an error.
<b>-q</b>	Quiet mode. This suppresses the banner.
<b>-t</b> <i>time</i>	Soft shutdown, with a hard shutdown after the specified time. <i>time</i> is a number followed by D, H, M, or S (for days, hours, minutes and seconds). For example, <code>-t 10m</code> specifies that the server should be shut down in 10 minutes or when current synchronizations complete, whichever is sooner. D, H, M, and S are not case sensitive.
<b>-w</b>	Waits for the MobiLink server to shut down before returning from the command.
<i>name</i>	If the MobiLink server is started using the <code>-zs</code> option, it must be shut down by specifying the same server name. See <a href="#">“-zs option” on page 119</a> .

### Description

By default (if none of `-f`, `-h` or `-t` are specified), `mlstop` does a soft shutdown.

- **Soft shutdown** means that the MobiLink server stops accepting new connections and exits when the current synchronizations are complete.
- **Hard shutdown** means that the MobiLink server stops all synchronizations and exits. Some remotes may report an error.

## MobiLink user authentication utility (mluser)

Registers MobiLink users at the consolidated database. For SQL Anywhere remotes, the users must have previously been created at the remote databases with the CREATE SYNCHRONIZATION USER statement.

### Syntax

```
mluser [ options ] -c "connection-string"
{ -f file | -u user [ -p password ] }
```

Option	Description
@data	Use this to read in options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used. See “Using configuration files” [ <a href="#">SQL Anywhere Server - Database Administration</a> ].  If you want to protect passwords or other information in the configuration file, you can use the File Hiding utility to obfuscate the contents of the configuration file. See “File Hiding utility (dbfhide)” [ <a href="#">SQL Anywhere Server - Database Administration</a> ].
-c "keyword=value;..."	Use this to supply database connection parameters. The connection string must give the utility permission to connect to the consolidated database using an ODBC data source. This parameter is required.
-d	Deletes the user name(s) specified by -f or -u.
-f filename	Reads the user names and passwords from the specified file. The file should be a text file containing one user name and password pair on each line, separated by white space. You must specify either -f or -u.
-fips	When set, mluser fails if FIPS support is not installed.
-o filename	Logs output messages to the specified file.
-ot filename	Truncate the log file and then append output messages to it. The default is to send output to the screen.
-p password	Password to associate with the user. This option can only be used with -u.

Option	Description
<b>-pc</b> <i>collation-id</i>	Supplies a database collation ID for character set conversion of the user name and password. This should be one of the SQL Anywhere collation labels such as those listed in <a href="#">“Supported and alternate collations” [SQL Anywhere Server - Database Administration]</a> . This option is required when user names and passwords are read from a file that is encoded in a different character set than the default character set determined by locale.
<b>-u</b> <i>ml_username</i>	Specify the user name to add (or delete, if used with -d). Only one user can be specified on a single command line. This option is used with -p if passwords are being used. You must specify either -f or -u.
<b>-v</b>	Specifies verbose logging.

### Remarks

Given a user/password pair, the mluser utility first attempts to add the user. If the user has already been added to the consolidated database, it attempts to update the password for that user.

There are alternative ways to register user names in the consolidated database:

- Use Sybase Central.
- Specify the -zu+ command line option with mlsrv11. In this case, any existing MobiLink users that have not been added to the consolidated database are added when they first synchronize.

The MobiLink user must already exist in a remote database. To add users at the remote, you have the following options:

- For SQL Anywhere remotes, set the name with CREATE SYNCHRONIZATION USER and synchronize with that user name.
- For UltraLite remotes, you can either use the user\_name field of the ul\_synch\_info structure; or in Java, use the SetUserName() method of the ULSynchInfo class before synchronizing.

### See also

- [“MobiLink users” \[MobiLink - Client Administration\]](#)
- [“-zu option” on page 121](#)
- [“CREATE SYNCHRONIZATION USER statement \[MobiLink\]” \[SQL Anywhere Server - SQL Reference\]](#)
- [“Transport-layer security” \[SQL Anywhere Server - Database Administration\]](#)

---



---

# MobiLink server system tables

## Contents

Introduction to MobiLink system tables .....	695
IBM DB2 mainframe system table name conversions .....	696
ml_active_remote_id .....	697
ml_column .....	698
ml_connection_script .....	699
ml_database .....	700
ml_device .....	701
ml_device_address .....	703
ml_listening .....	705
ml_passthrough .....	707
ml_passthrough_repair .....	708
ml_passthrough_script .....	709
ml_passthrough_status .....	711
ml_property .....	712
ml_qa_clients .....	713
ml_qa_delivery .....	714
ml_qa_delivery_archive .....	716
ml_qa_global_props .....	718
ml_qa_notifications .....	719
ml_qa_repository .....	720
ml_qa_repository_archive .....	721
ml_qa_repository_props .....	722
ml_qa_repository_props_archive .....	723
ml_qa_repository_staging .....	724
ml_qa_status_history .....	725
ml_qa_status_history_archive .....	726
ml_qa_status_staging .....	727
ml_script .....	728
ml_script_version .....	729
ml_scripts_modified .....	730
ml_server .....	731

ml_sis_sync_state .....	732
ml_subscription .....	733
ml_table .....	735
ml_table_script .....	736
ml_user .....	737

---

## Introduction to MobiLink system tables

MobiLink system tables store information about MobiLink users, subscriptions, tables, scripts, script versions, and other information. They are required for MobiLink synchronization. Unlike other system tables, you can modify the MobiLink system tables, although in most cases you do not need to.

MobiLink system tables are created when you run the MobiLink setup script for your consolidated database. They must be stored on your consolidated database. The database user who runs the setup script is the owner of the MobiLink system tables that are created by the script.

See [“Setting up a consolidated database” on page 6](#).

### Notes

- This chapter provides data types for the MobiLink system tables in SQL Anywhere consolidated databases. In some RDBMSs, the data types are slightly different.
- DB2 mainframe version 8.1 supports a backward compatibility mode, where column names and other identifiers are limited to a maximum of 18 characters. To support this environment, all MobiLink system objects in DB2 mainframe have names of 18 characters or less. See [“IBM DB2 mainframe system table name conversions” on page 696](#).
- IBM DB2 LUW version 5.2 only supports column names and other identifiers of 18 characters or less. In a DB2 LUW 5.2 consolidated database, MobiLink system table names are truncated where necessary.

## IBM DB2 mainframe system table name conversions

IBM DB2 mainframe consolidated databases only support column names and other identifiers of 18 characters or less. The following table identifies how system table names for DB2 mainframe consolidated databases are mapped to system table names for all other consolidated database types.

If the system table name does not appear in the table below, no conversion is required.

System table name	System table name for DB2 mainframe consolidated databases
<a href="#">“ml_active_remote_id” on page 697</a>	ml_active_rid
<a href="#">“ml_connection_script” on page 699</a>	ml_conn_script
<a href="#">“ml_passthrough” on page 707</a>	ml_pt
<a href="#">“ml_passthrough_repair” on page 708</a>	ml_pt_repair
<a href="#">“ml_passthrough_script” on page 709</a>	ml_pt_script
<a href="#">“ml_passthrough_status” on page 711</a>	ml_pt_status
<a href="#">“ml_scripts_modified” on page 730</a>	ml_script_modified

## ml\_active\_remote\_id

Stores the synchronization status for each remote database in the server farm.

Column	Description
remote_id	VARCHAR(128). The unique integer that identifies the ID of the remote database currently being synchronized.
server_id	INTEGER. The value that references the server_id in the ml_server table.
status	CHAR(1). The synchronization status for the remote database. The value O indicates an ongoing synchronization, and C indicates a canceled synchronization.

### Remarks

This table does not contain data unless a server farm is running.

For DB2 mainframe consolidated database types, this table is called ml\_active\_rid. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( remote\_id )

FOREIGN KEY( server\_id ) REFERENCES ml\_server( server\_id )

## ml\_column

Stores the names of columns for a specific table in a specific script version.

Column	Description
version_id	INTEGER. A number identifying the script version.
table_id	INTEGER. A number identifying the table.
idx	INTEGER. The index, origin 1, of this column in the table. The column order must be the order in which the columns were created in the remote database.
name	VARCHAR(128). The column name.
type	VARCHAR(128). Not currently used.

This table is only required when SQL scripts contain named parameters for columns (for example, o.column-name and r.column-name). (The exception is the column index, which is available even without this MobiLink system table being populated; for example, o.column-index and r.column-index.)

This table is populated by the **Create Synchronization Model Wizard** when you deploy a MobiLink model. If you did not use the **Create Synchronization Model Wizard**, or if you did use it but afterward changed the schema of synchronized columns on the remote database outside Sybase Central Model mode, you can use the ml\_add\_column stored procedure to populate the table.

Note: The dbmsync extended option SendColumnNames and UltraLite synchronization parameter Send Column Names are used by direct row handling, but are not used for named row parameters.

### Remarks

There is a system view, ml\_columns, that makes it easier to view the contents of this table.

### Constraints

PRIMARY KEY( idx, version\_id, table\_id )

UNIQUE( version\_id, table\_id, name )

FOREIGN KEY( version\_id ) REFERENCES ml\_script\_version( version\_id )

FOREIGN KEY( table\_id ) REFERENCES ml\_table( table\_id )

### See also

- [“ml\\_add\\_column system procedure” on page 666](#)
- [“Script parameters” on page 320](#)

## ml\_connection\_script

For a given script version, this table associates a script with a given event.

Column	Description
version_id	INTEGER. A number identifying the script version.
event	VARCHAR(128). The name of the event that triggers the connection script.
script_id	INTEGER. A number identifying the script. The text of the connection script is stored in the ml_script system table.

### Remarks

There is a system view, ml\_connection\_scripts, that makes it easier to view the contents of this table.

For DB2 mainframe consolidated database types, this table is called ml\_conn\_script. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( version\_id, event )

FOREIGN KEY( version\_id ) REFERENCES ml\_script\_version( version\_id )

FOREIGN KEY( script\_id ) REFERENCES ml\_script( script\_id )

## ml\_database

Stores a unique ID for each remote database that has synchronized.

**Caution**  
Do not alter this table.

Column	Description
rid	INTEGER. A unique integer that identifies the remote ID. This value is used internally.
remote_id	VARCHAR(128). The remote ID that uniquely identifies each remote database.
script_ldt	TIMESTAMP. The last time a passthrough script was downloaded.
description	VARCHAR(128). Reserved.

### Remarks

The remote ID is sent by the client in each synchronization. The MobiLink server tracks the state information for each remote database using this remote ID.

### Constraints

PRIMARY KEY( rid )



## ml\_device

This table is used only for server-initiated synchronization. It stores device names that are required by device tracking.

Column	Description
device_name	VARCHAR(255). The name given to the device. This name is extracted from the operating system unless you specify a name using the dblsn -e option.
listener_version	VARCHAR(128). Not null. The SQL Anywhere version number for the installed software on the device. Changing this value does not affect the operation of the software, but may be useful for diagnostic purposes.
listener_protocol	INTEGER. Not null. This column is 0, 1, or 2: <ul style="list-style-type: none"> <li>● <b>0</b> - for Listeners from SQL Anywhere prior to version 9.0.1</li> <li>● <b>1</b> - for post-9.0.0 Palm Listeners</li> <li>● <b>2</b> - for post-9.0.0 Windows Listeners</li> </ul>
info	VARCHAR(255). Not null. Operating system information about the listening device. This information can be overridden by providing information using the dblsn -f option.
ignore_tracking	VARCHAR(1). Not null. If this is <b>y</b> , tracking information is not written to the row. If it is <b>n</b> , tracking information is written to the row.
source	VARCHAR(255). Not null. This is <b>tracking</b> if the row was created by automatic device tracking. Otherwise, it is blank unless you change it using stored procedures to add information about where the data in this row came from. Unless it is set to <b>tracking</b> , the value in this column does not affect the operation of the software.

### Remarks

The MobiLink system tables ml\_device, ml\_device\_address, and ml\_listening contain information about devices for server-initiated synchronization. DeviceTracker gateways use this information to address target devices by MobiLink user name.

In most cases, you should not need to alter these tables. However, if your device does not support device tracking or if you want to override device tracking for troubleshooting purposes, you can add or delete rows to this system table using pre-defined stored procedures. See [“Adding support for device tracking” \[MobiLink - Server-Initiated Synchronization\]](#).

If you want to stop automatic tracking, set ignore\_tracking to **y**. In this case, it is also recommended that you use a source name other than **tracking**.

### Constraints

PRIMARY KEY( device\_name )

**See also**

- [“ml\\_set\\_device system procedure”](#) [*MobiLink - Server-Initiated Synchronization*]
- [“ml\\_delete\\_device system procedure”](#) [*MobiLink - Server-Initiated Synchronization*]

## ml\_device\_address

This table is used only for server-initiated synchronization. It stores addressing information that is required by device tracking.

Column	Description
device_name	VARCHAR(255). Not null. The name of the device. This name is extracted from the operating system unless you specify a name using the dblns -e option.
medium	VARCHAR(255). Not null. For UDP, this is <b>_UDP_</b> . Otherwise, it is the network provider ID.
address	VARCHAR(255). Not null. For UDP, this is <i>ip:port-number</i> , where <i>ip</i> is an IP address or host name. For SMS, this is the phone number.
active	VARCHAR(1). Not null. This is <b>y</b> for active, and <b>n</b> otherwise. A DeviceTracker gateway may deactivate a UDP channel if it is unresponsive and there is a fallback SMS delivery path.
last_modified	Timestamp. Not null. Default timestamp. The datetime when this row was last modified.
ignore_tracking	VARCHAR(1). Not null. If this is <b>y</b> , tracking information is not written to the row. If it is <b>n</b> , tracking information is written to the row.
source	VARCHAR(255). Not null. This is <b>tracking</b> if the row was created by automatic device tracking. Otherwise, it is blank unless you change it using stored procedures to add information about where the data in this row came from. Unless it is set to <b>tracking</b> , the value in this column does not affect the operation of the software.

### Remarks

The MobiLink system tables ml\_device, ml\_device\_address, and ml\_listening contain information about devices for server-initiated synchronization. DeviceTracker gateways use this information to address target devices by MobiLink user name.

In most cases, you should not need to alter these tables. However, if your device does not support device tracking or if you want to override device tracking for troubleshooting purposes, you can add or delete rows to this system table using pre-defined stored procedures. See [“Adding support for device tracking” \[MobiLink - Server-Initiated Synchronization\]](#).

If you want to stop automatic tracking, set ignore\_tracking to **y**. In this case, it is also recommended that you use a source name other than **tracking**.

### Constraints

PRIMARY KEY( device\_name, medium )

FOREIGN KEY( device\_name ) REFERENCES ml\_device( device\_name )

**See also**

- [“ml\\_set\\_device\\_address system procedure” \[MobiLink - Server-Initiated Synchronization\]](#)
- [“ml\\_delete\\_device\\_address system procedure” \[MobiLink - Server-Initiated Synchronization\]](#)

## ml\_listening

This table is used only for server-initiated synchronization. It maps a MobiLink user name to a device name for device tracking.

Column	Description
name	<p>VARCHAR(128). Not null. The name you use to address notification. This column can be populated in one of three ways:</p> <ul style="list-style-type: none"> <li>• If you use the <code>dblsn -t+</code> option, this is the alias that you have defined for the <code>ml_user</code>.</li> <li>• If you used the <code>dblsn -u</code> option, this is an <code>ml_user</code> name.</li> <li>• If you use neither <code>-t+</code> nor <code>-u</code>, the default name is <i>device-name-dblsn</i>, where <i>device-name</i> is the name of your device. You can find the device name in your Listener messages window. Optionally, you can set the device name using the <code>dblsn -e</code> option.</li> </ul> <p>See “<a href="#">Listener options for Windows</a>” [<i>MobiLink - Server-Initiated Synchronization</i>].</p>
device_name	<p>VARCHAR(255). Not null. The name given to the device. This name is extracted from the operating system unless you specify a name using the <code>dblsn -e</code> option.</p>
listening	<p>VARCHAR(1). Not null. This is <b>y</b> for an active Listener; otherwise it is <b>n</b>. This field is set when you use the <code>dblsn</code> option <code>-t</code>, or you can manually set it with stored procedures.</p>
ignore_tracking	<p>VARCHAR(1). Not null. If this is <b>y</b>, tracking information is not written to the row. If it is <b>n</b>, tracking information is written to the row.</p>
source	<p>VARCHAR(255). Not null. This is <b>tracking</b> if the row was created by automatic device tracking. Otherwise, it is blank unless you change it using stored procedures to add information about where the data in this row came from. The value in this column does not affect the operation of the software.</p>

### Remarks

The MobiLink system tables `ml_device`, `ml_device_address`, and `ml_listening` contain tracked information about devices for server-initiated synchronization. DeviceTracker gateways use this information to address target devices by MobiLink user name.

In most cases, you should not need to alter these tables. However, if your device does not support device tracking or if you want to override device tracking for troubleshooting purposes, you can add or delete rows in this table using pre-defined stored procedures. See “[Adding support for device tracking](#)” [*MobiLink - Server-Initiated Synchronization*].

If you want to stop automatic tracking, set `ignore_tracking` to **y**. In this case, it is also recommended that you use a source name other than **tracking**.

**Constraints**

PRIMARY KEY( name )

FOREIGN KEY( device\_name ) REFERENCES ml\_device( device\_name )

**See also**

- [“ml\\_set\\_listening system procedure”](#) [*MobiLink - Server-Initiated Synchronization*]
- [“ml\\_delete\\_listening system procedure”](#) [*MobiLink - Server-Initiated Synchronization*]

## ml\_passthrough

Stores rows that indicate which scripts should run on each remote database.

Column	Description
remote_id	VARCHAR(128). The remote ID of the client that should execute the script.
run_order	INTEGER. The run order of the script for this client. This value must be non-negative.
script_id	INTEGER. An integer that references the script_id in the ml_passthrough_script table. This value identifies the script to be executed.
last_modified	TIMESTAMP. The last time the passthrough information was added or modified.

### Remarks

You can use the system procedures ml\_add\_passthrough and ml\_delete\_passthrough to add, modify, and delete entries in this table.

For DB2 mainframe consolidated database types, this table is called ml\_pt. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( remote\_id, run\_order )

FOREIGN KEY( remote\_id ) REFERENCES ml\_database( remote\_id )

FOREIGN KEY( script\_id ) REFERENCES ml\_passthrough\_script( script\_id )

### See also

- [“ml\\_add\\_passthrough system procedure” on page 673](#)
- [“ml\\_delete\\_passthrough system procedure” on page 681](#)

## ml\_passthrough\_repair

Stores rules that define how clients should handle script errors.

Column	Description
failed_script_id	INTEGER. Identifies the script to which this rule applies. This value references the script_id in the ml_passthrough_script table.
error_code	INTEGER. The error code that this rule handles.
new_script_id	INTEGER. The script_id defined in the ml_passthrough_script table, representing the repair script. If the action is R, this value identifies the script with which to replace the failed script. The value refers to the script_id column of the ml_passthrough_script table. If the action is S, P, or H, the value is null.
action	<p>CHAR(1). The action to perform on the client when the script fails. This value must be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>R</b> (replace) Indicates that the failed script should be replaced with the one specified by <b>new script name</b> and an attempt should be made to run the new script. To rerun the failed script, choose <b>new script name</b> to be the same as <b>failed script name</b>.</li> <li>• <b>P</b> (purge) Indicates that the remote database should discard all the scripts that it has received and continue executing script normally after that.</li> <li>• <b>S</b> (skip) Indicates that the remote database should ignore the failed script and continue executing scripts as if the failed script had succeeded.</li> <li>• <b>H</b> (halt) Indicates that the remote should not execute any more scripts until it receives further instructions.</li> </ul>

### Remarks

You can use the system procedures ml\_add\_passthrough\_repair and ml\_delete\_passthrough\_repair to add, modify, and delete entries in this table.

For DB2 mainframe consolidated database types, this table is called ml\_pt\_repair. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( failed\_script\_id, error\_code )

FOREIGN KEY( failed\_script\_id ) REFERENCES ml\_passthrough\_script( script\_id )

### See also

- [“ml\\_add\\_passthrough\\_repair system procedure” on page 674](#)
- [“ml\\_delete\\_passthrough\\_repair system procedure” on page 682](#)



## ml\_passthrough\_script

Stores the names and contents for each passthrough script.

Column	Description
script_id	INTEGER. A unique integer that identifies the passthrough script name. This value is used internally.
script_name	VARCHAR(128). The unique passthrough script name.
flags	<p>VARCHAR(256). This value tells clients how to run the script. This value can be null or contain a combination of the following keywords in a semicolon-delimited list:</p> <ul style="list-style-type: none"> <li>• <b>manual</b> Indicates that the script may only be run in manual execution mode. By default, all scripts can be run in either automatic or manual execution modes.</li> <li>• <b>exclusive</b> Indicates that the script may only be automatically executed at the end of a synchronization where exclusive locks were obtained on all tables being synchronized. This option is ignored if the affected_ publications value lists no publications. This option is only meaningful to SQL Anywhere remotes.</li> <li>• <b>schema_diff</b> Indicates that the script should be run in schema-diffing mode. In this mode, the database schema is altered to match the schema described in the script. For example, a create statement for an existing table is treated as an alter statement. This flag only applies to scripts run on UltraLite remotes.</li> </ul> <p>For example:</p> <pre>'manual;exclusive;schema_diff'</pre>
affected_pubs	TEXT. A list of publications that must be synchronized before the script is run. An empty or null publication list indicates that no synchronization is required. This value is only meaningful with SQL Anywhere clients.
script	TEXT. The contents of the passthrough script.
description	VARCHAR(2000). A comment or description of the script.

### Remarks

Use the following procedures to add and delete entries in this table.

- [“ml\\_add\\_passthrough\\_script system procedure” on page 676](#)
- [“ml\\_delete\\_passthrough\\_script system procedure” on page 683](#)

It is recommended that you do not directly update the ml\_passthrough\_script table without using the ml\_add\_passthrough\_script and ml\_delete\_passthrough\_script system procedures. Clients do not receive

newly updated scripts if they have already downloaded the original passthrough scripts. Discrepancies between the various scripts on multiple clients makes managing script execution difficult or impossible.

For DB2 mainframe consolidated database types, this table is called ml\_pt\_script. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( script\_id )

### See also

- [“ml\\_add\\_passthrough\\_script system procedure” on page 676](#)
- [“ml\\_delete\\_passthrough\\_script system procedure” on page 683](#)

## ml\_passthrough\_status

Stores the status for each passthrough script after it is executed by a client.

Column	Description
status_id	INTEGER. A unique integer that identifies the row.
remote_id	VARCHAR(128). Identifies the remote database that executed the script. References the remote_id in the ml_database table.
run_order	INTEGER. The run order of the script on the client.
script_id	INTEGER. Identifies the script that was executed. References the script_id in the ml_passthrough_script table.
script_status	CHAR(1). The status of the script. Contains <b>S</b> for success or <b>E</b> for error.
error_code	INTEGER. The SQL code generated by the script on the remote database.
error_text	TEXT. The error text generated by the script on the remote database.
remote_run_time	TIMESTAMP. The time at the remote database when the script was executed.

### Remarks

The MobiLink server does not automatically remove entries from the ml\_passthrough\_status table. Deletions from this table must be performed manually.

For DB2 mainframe consolidated database types, this table is called ml\_pt\_status. See [“IBM DB2 mainframe system table name conversions” on page 696](#).

### Constraints

PRIMARY KEY( status\_id )

FOREIGN KEY( remote\_id ) REFERENCES ml\_database( remote\_id )

FOREIGN KEY( script\_id ) REFERENCES ml\_passthrough\_script( script\_id )

### See also

- [“upload\\_insert table event” on page 504](#)

## ml\_property

This table stores some MobiLink properties.

Column	Description
component_name	VARCHAR(128). For user-defined properties, this can be <b>ScriptVersion</b> or <b>SIS</b> .
property_set_name	VARCHAR(128). If the component_name is <b>ScriptVersion</b> , this is the name of the script version. If the component_name is <b>SIS</b> , this is the name of the Notifier, gateway, or carrier that you are setting a property for.
property_name	VARCHAR(128). The name of the property. If the component_name is <b>ScriptVersion</b> , this is a user-defined property. If the component_name is <b>SIS</b> , this is a property of the Notifier, gateway, or carrier. See <a href="#">“MobiLink server settings for server-initiated synchronization” [MobiLink - Server-Initiated Synchronization]</a> .
property_value	TEXT. The value of the property.

### Remarks

This table stores name-value pairs. Some of the properties in this table are used internally by MobiLink. In addition, you can use the stored procedure ml\_add\_property to add or delete rows in this table.

You can use the component\_name **ScriptVersion** to store information on a per script version basis that can be accessed by Java or .NET scripting logic.

### Constraints

PRIMARY KEY( component\_name, property\_set\_name, property\_name )

### See also

- [“ml\\_add\\_property system procedure” on page 677](#)

## ml\_qa\_clients

This table is used only for QAnywhere applications. It is a global temporary table that exists only on SQL Anywhere and Oracle consolidated databases.

**Caution**

Do not alter this table.

Column	Description
client	VARCHAR(128). Client targeted by uploaded messages.

## ml\_qa\_delivery

This table is used only for QAnywhere applications.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). Globally unique message identifier.
seqno	BIGINT. Used to give an ordering to the messages, which is necessary for true queuing.
address	VARCHAR(255). Address of the target recipient.
clientaddress	VARCHAR(128). Client part of the address.
client	VARCHAR(128). Client targeted by current client state.
originator	VARCHAR(128). The name of the originating client.
priority	INTEGER. A number from 0 to 9. Messages with a higher priority number are delivered before messages with a lower priority. The default is 4.
expires	TIMESTAMP. The expiry time after which the message might not be delivered.
kind	INTEGER. Indicates whether the message is binary (1) or text (2).
contentsize	BIGINT. The size of the message. For binary messages, this is the number of bytes. For text messages, this is the number of characters.
status	INTEGER. The status of the message. Can be 1 (pending), 10 (receiving), 30 (expired), 40 (canceled), 50 (unreceivable), or 60 (received).
statustime	TIMESTAMP. The time this status was achieved. The time is local to the client achieving the state.
syncstatus	INTEGER. The state of the synchronization between the client and server for this message. Can be 0 (not in sync), 1 (in sync), 2 (message should not be synchronized), or 3 (synchronizing).
receiverid	VARCHAR(128). An identifier set by the receiver that identifies the receiver of the message, if any.

**Constraints**

PRIMARY KEY( msgid, address )

FOREIGN KEY ( msgid ) REFERENCES ml\_qa\_repository( msgid )

## ml\_qa\_delivery\_archive

This table is used only for QAnywhere applications.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). Globally unique message identifier.
seqno	BIGINT. Used to give an ordering to the messages, which is necessary for true queuing.
address	VARCHAR(255). Address of the target recipient.
clientaddress	VARCHAR(128). Client part of the address.
client	VARCHAR(128). Client targeted by current client state.
originator	VARCHAR(128). The name of the originating client.
priority	INTEGER. A number from 0 to 9. Messages with a higher priority number are delivered before messages with a lower priority. The default is 4.
expires	TIMESTAMP. The expiry time after which the message might not be delivered.
kind	INTEGER. Indicates whether the message is binary (1) or text (2).
contentsize	BIGINT. The size of the message. For binary messages, this is the number of bytes. For text messages, this is the number of characters.
status	INTEGER. The status of the message. Can be 1 (pending), 10 (receiving), 30 (expired), 40 (canceled), 50 (unreceivable), or 60 (received).
statustime	TIMESTAMP. The time this status was achieved. The time is local to the client achieving the state.
syncstatus	INTEGER. The state of the synchronization between the client and server for this message. Can be 0 (not in sync), 1 (in sync), 2 (message should not be synchronized), or 3 (synchronizing).
receiverid	VARCHAR(128). An identifier set by the receiver that identifies the receiver of the message, if any.

**Constraints**

PRIMARY KEY( msgid, address )



FOREIGN KEY ( msgid ) REFERENCES ml\_qa\_repository( msgid )

## mi\_qa\_global\_props

This table is used only for QAnywhere applications. It contains global *name-value* pairs that are used in transmission rules.

**Caution**

Do not alter this table.

Column	Description
client	VARCHAR(128). The client associated with the property. A client value of 'ia-nywhere.server.defaultClient' indicates a property that is global to all clients.
name	VARCHAR(255). The name of the property.
modifiers	INTEGER. Bitfields used to further describe the property. Currently, only the first bit is used to indicate a property that should not be synchronized. All other bit fields are reserved for future use.
value	LONG VARCHAR. The value of the property.
last_modified	TIMESTAMP. The last time the value was changed. This is necessary to indicate when a property needs to be synchronized with the client.

**Constraints**

PRIMARY KEY ( client, name )

## ml\_qa\_notifications

This table is used only for QAnywhere applications. It is used by the Notifier to determine which QAnywhere clients to notify to initiate synchronization.

**Caution**

Do not alter this table.

Column	Description
user_id	INTEGER.
name	VARCHAR(128). The QAnywhere client name that uniquely identifies a client message store.

**Constraints**

PRIMARY KEY( name )

## ml\_qa\_repository

This table is used only for QAnywhere applications. It stores messages and their properties.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
props	LONG BINARY. An encoding of the message properties.
content	LONG BINARY. The content of the message. Text messages are encoded as UTF-8.

### Constraints

PRIMARY KEY( msgid )

## ml\_qa\_repository\_archive

This table is used only for QAnywhere applications. It archives messages and their properties.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
props	LONG BINARY. An encoding of the message properties.
content	LONG BINARY. The content of the message. Text messages are encoded as UTF-8.

**Constraints**

PRIMARY KEY( msgid )

## ml\_qa\_repository\_props

This table is used only for QAnywhere applications. This is an expansion of the props column in the ml\_qa\_repository table. Properties are only expanded as needed by the transmission rules engine. If there are no associated rules, a property is not expanded.

**Caution**  
Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
name	VARCHAR(128). The name of the property. If the property name was provided in Unicode, it is translated to the native character set of the database.
value	LONG VARCHAR. The value of the property.

### Constraints

PRIMARY KEY( msgid, name )

FOREIGN KEY( msgid ) REFERENCES ml\_qa\_repository( msgid )

## ml\_qa\_repository\_props\_archive

This table is used only for QAnywhere applications. This is an expansion of the props column in the ml\_qa\_repository table. Properties are only expanded as needed by the transmission rules engine. If there are no associated rules, a property is not expanded.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
name	VARCHAR(128). The name of the property. If the property name was provided in Unicode, it is translated to the native character set of the database.
value	LONG VARCHAR. The value of the property.

**Constraints**

PRIMARY KEY( msgid, name )

FOREIGN KEY( msgid ) REFERENCES ml\_qa\_repository( msgid )

## ml\_qa\_repository\_staging

This table is used only for QAnywhere applications. It contains messages that are to be sent to a QAnywhere client using SQL Anywhere version 9.0.1.

**Caution**

Do not alter this table.

Column	Description
seqno	BIGINT. Used to give a total ordering to the messages, this is necessary for true queuing.
msgid	VARCHAR(255). Globally unique message identifier.
destination	VARCHAR(128). The address of the message.
originator	VARCHAR(128). The name of the originating MobiLink user.
status	VARCHAR(128). The status of the message. Can be <b>pending</b> , <b>receiving</b> , <b>received</b> , <b>unreceivable</b> , <b>expired</b> , or <b>cancelled</b> . The default is <b>pending</b> .
statustime	TIMESTAMP. The last time the status was changed.
expires	TIMESTAMP. Expiry time after which the message are not delivered.
priority	INTEGER. A number from 0 to 9. Messages with a higher number are always delivered before messages with a lower number. The default is 4.
props	LONG BINARY. An encoding of the message properties.
kind	INTEGER. Indicates whether the message is binary (1) or text (2).
content	LONG BINARY. The content of the message. Text messages are encoded as UTF-8.
contentsize	BIGINT. The size of the message. For binary messages, this is the number of bytes. For text messages, this is the number of characters.
mluser	VARCHAR(128). The MobiLink user name that uniquely identifies a remote database.

### Constraints

PRIMARY KEY( msgid )



## ml\_qa\_status\_history

This table is used only for QAnywhere applications. It contains a history of message status changes.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
address	VARCHAR(255). The address of the target recipient.
status	INTEGER. The status of the message. Can be 1 (pending), 10 (receiving), 30 (expired), 40 (canceled), 50 (unreceivable), or 60 (received).
statustime	TIMESTAMP. The time this status was achieved. The time is local to the client achieving the state.
servertime	TIMESTAMP. The time the status change was received by the server.
details	VARCHAR(1000). The details of the status change, if any.
syncstatus	INTEGER. The state of the synchronization between the client and server for this message. Can be <b>0</b> (not in sync), <b>1</b> (in sync), <b>2</b> (message should not be synchronized), or <b>3</b> (synchronizing).

### Constraints

PRIMARY KEY( msgid )

FOREIGN KEY( msgid ) REFERENCES ml\_qa\_repository( msgid )

## ml\_qa\_status\_history\_archive

This table is used only for QAnywhere applications. It contains a history of message status changes.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
address	VARCHAR(255). The address of the target recipient.
status	INTEGER. The status of the message. Can be 1 (pending), 10 (receiving), 30 (expired), 40 (canceled), 50 (unreceivable), or 60 (received).
statustime	TIMESTAMP. The time this status was achieved. The time is local to the client achieving the state.
servertime	TIMESTAMP. The time the status change was received by the server.
details	VARCHAR(1000). The details of the status change, if any.
syncstatus	INTEGER. The state of the synchronization between the client and server for this message. Can be <b>0</b> (not in sync), <b>1</b> (in sync), <b>2</b> (message should not be synchronized), or <b>3</b> (synchronizing).

**Constraints**

PRIMARY KEY( msgid )

FOREIGN KEY( msgid ) REFERENCES ml\_qa\_repository( msgid )

## ml\_qa\_status\_staging

This table is used only for QAnywhere applications. It is a staging table that is used when synchronizing status changes with the originating client, where the originating client was using SQL Anywhere version 9.0.1.

**Caution**

Do not alter this table.

Column	Description
msgid	VARCHAR(128). A globally unique message identifier.
status	VARCHAR(255). The status of the message. Can be <b>pending</b> , <b>receiving</b> , <b>received</b> , <b>unreceivable</b> , <b>expired</b> , or <b>cancelled</b> . The default is <b>pending</b> .
statustime	TIMESTAMP. The last time the status was changed.
mluser	VARCHAR(128). The MobiLink user name that uniquely identifies a remote database.

**Constraints**

PRIMARY KEY( msgid )

## ml\_script

This table stores the content of all scripts.

Column	Description
script_id	INTEGER. A unique integer that identifies the script.
script	TEXT. The text of the script.
script_language	VARCHAR(128). The scripting language used for the script. The scripting language can be <b>sql</b> , <b>java</b> , or <b>dnet</b> .
checksum	VARCHAR(64). This column is used internally.

### Constraints

PRIMARY KEY( script\_id )

## ml\_script\_version

This table stores the name and description of scripts associated with each script version.

Column	Description
version_id	INTEGER. A unique integer that identifies the version.
name	VARCHAR(128). The name of the script version.
description	TEXT. The description given to the version. The description is not used by MobiLink, but is useful for application-specific comments. For example, you could describe the purpose of a given script version.

### Constraints

PRIMARY KEY( version\_id )

## ml\_scripts\_modified

This table stores the last time script tables were changed. The MobiLink server checks this table to determine if it must load new scripts.

Column	Description
last_modified	DATETIME. The last time when the ml_script, ml_table_script, or ml_connection_script system table was altered.

### Remarks

For DB2 mainframe consolidated database types, this table is called ml\_script\_modified. See [“IBM DB2 mainframe system table name conversions”](#) on page 696.

### Constraints

PRIMARY KEY( last\_modified )

## ml\_server

Stores the information for each MobiLink server running in the server farm.

Column	Description
server_id	INTEGER. A unique integer that identifies the server ID of the MobiLink server name. This value is used internally.
name	VARCHAR(128). The name of the MobiLink server actively running on the server farm. This value must be unique across the server farm.
version	VARCHAR(10). The current version of the MobiLink server.
connection_info	VARCHAR(2048). The connection information for the MobiLink server.
instance_key	VARCHAR(32). An assigned instance key. This value is used internally.
start_time	TIMESTAMP. The startup time of the MobiLink server.
liveness	TIMESTAMP. The last time the client responded to the server.

### Remarks

This table does not contain data unless a server farm is running.

### Constraints

PRIMARY KEY( server\_id )

## ml\_sis\_sync\_state

This table is used by Sybase Central to generate request cursors for server-initiated synchronization.

**Caution**

Do not alter this table.

Column	Description
remote_id	VARCHAR(128). The remote ID that uniquely identifies a database.
subscription_id	VARCHAR(128). The subscription_id is a number that is generated by a remote database. For SQL Anywhere clients, this value is the same as the sync_id in the SYS.ISYSSYNC system table.
publication_name	<p>VARCHAR(128). The user-defined name for the publication that is subscribed to by the subscription. In every synchronization, the client sends the publication name for each subscription_id.</p> <p>For UltraLite clients prior to version 10.0.0, this is always &lt;unknown&gt;; for UltraLite version 10 and later, it is the publication or the string ul_no_pub if there is no publication.</p>
user_name	VARCHAR(128). The MobiLink user name.
last_upload	TIMESTAMP. The last time an upload was applied to the consolidated database for a given remote ID and subscription_id. The default is January 1, 1900, 00:00:00.
last_download	TIMESTAMP. The last time a download was applied to the consolidated for a given user and subscription_id. The default is January 1, 1900, 00:00:00.

**Constraints**

PRIMARY KEY( remote\_id, subscription\_id )



## ml\_subscription

This table stores state information for each remote.

Column	Description
rid	INTEGER. A unique integer identifying the remote ID. This value is used internally.
subscription_id	<p>VARCHAR(128). The subscription_id is a number that is generated by a remote database. For SQL Anywhere clients, this value is the same as the sync_id in the SYS.ISYSSYNC system table.</p> <p>UltraLite clients do not use subscriptions, so for UltraLite clients, this value is the UltraLite publication ID for version 10.0.0 and later, and &lt;unknown&gt; for versions 8 and 9.</p>
user_id	INTEGER. The user who performed the last synchronization for the given rid and subscription_id. You can use the user_id column to find the MobiLink user who ran the last successful synchronization.
progress	NUMERIC(20,0). The synchronization progress, also called the offset, state, sequence number, or progress counter.
publication_name	<p>VARCHAR(128). The user-defined name for the publication that is subscribed to by the subscription. In every synchronization, the client sends the publication name for each subscription_id.</p> <p>For UltraLite clients prior to version 10.0.0, this is always &lt;unknown&gt;; for UltraLite version 10 and later, it is the publication or the string ul_no_pub if there is no publication.</p>
last_upload_time	TIMESTAMP. The last time an upload was applied to the consolidated database for a given remote ID and subscription_id. The default is January 1, 1900, 00:00:00.
last_download_time	TIMESTAMP. The last time a download was applied to the consolidated database for a given user and subscription_id. The default is January 1, 1900, 00:00:00. See <a href="#">“How download timestamps are generated and used”</a> on page 131.

### Remarks

In SQL Anywhere clients, the progress refers to a position in the transaction log of the remote database. It indicates the point to which all committed operations for the subscription have been uploaded from the database. The dbmlsync utility uses the offset to decide what data to upload. On the SQL Anywhere remote database, the offset is stored in the progress column of the SYS.ISYSSYNC system table.

See:

- “SYSSYNC system view” [*SQL Anywhere Server - SQL Reference*]
- “Progress offsets” [*MobiLink - Client Administration*]

In UltraLite clients, the progress is the synchronization sequence number or progress counter for the given publication. This counter indicates what rows have been synchronized. It is incremented every time the publication synchronizes. This number is used internally in the UltraLite database and cannot be accessed.

See “The progress counter” [*UltraLite - Database Management and Reference*].

### Constraints

PRIMARY KEY( rid, subscription\_id )

FOREIGN KEY( rid ) REFERENCES ml\_database( rid )

FOREIGN KEY( user\_id ) REFERENCES ml\_user( user\_id )

### See also

- “Remote IDs” [*MobiLink - Client Administration*]

---

## ml\_table

This table stores the names of remote tables. This list includes any table that is marked as a synchronized table.

Column	Description
table_id	INTEGER. A unique integer identifying the table.
name	VARCHAR(128). The name given to the table.

### Constraints

PRIMARY KEY( table\_id )

## ml\_table\_script

For a given script version, this table associates a table script with a given table and event.

Column	Description
version_id	INTEGER. A number identifying the script version.
table_id	INTEGER. A number identifying the table.
event	VARCHAR(128). The name of the event.
script_id	INTEGER. An number identifying the script. The script is stored in the ml_script table.

### Remarks

There is a system view, ml\_table\_scripts, that makes it easier to view the contents of the ml\_table\_script MobiLink system table.

### Constraints

FOREIGN KEY( version\_id ) REFERENCES ml\_script\_version( version\_id )

FOREIGN KEY( table\_id ) REFERENCES ml\_table( table\_id )

FOREIGN KEY( script\_id ) REFERENCES ml\_script( script\_id )

## ml\_user

Stores registered users and their hashed passwords.

Column	Description
user_id	INTEGER. A unique integer that identifies the user. This value is used internally.
name	VARCHAR(128). The registered user name.
hashed_password	BINARY(32). The password in obfuscated form. This value can be null however, it is recommended that you specify a password.

### Remarks

This table stores all registered users that are known by the MobiLink server. The user names are sent by clients in every synchronization and the clients may optionally send up the password for the user for authentication.

The MobiLink server uses its own algorithm to hash the user password.

Do not directly insert any user names with non-null passwords into this table. The user names can be added by using the mluser utility.

### Constraints

PRIMARY KEY( user\_id )

### See also

- [“MobiLink user authentication utility \(mluser\)” on page 690](#)

---

---

# MobiLink data mappings between remote and consolidated databases

## Contents

Adaptive Server Enterprise data mapping .....	740
IBM DB2 LUW data mapping .....	749
IBM DB2 mainframe data mapping .....	756
Microsoft SQL Server data mapping .....	767
MySQL data mapping .....	774
Oracle data mapping .....	779

---

## Adaptive Server Enterprise data mapping

### Mapping to Adaptive Server Enterprise consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to Adaptive Server Enterprise consolidated data types. For example, a column of type FLOAT on the remote database should be type REAL on the consolidated database.

Maximum column length (MCL) depends on the Adaptive Server Enterprise page size. If the page size is 2K the MCL is 1954; if the page size is 4K the MCL is 4002. For information about MCL, see the Adaptive Server Enterprise documentation.

SQL Anywhere or UltraLite data type	Adaptive Server Enterprise data type	Notes
BIGINT	NUMERIC(20) <sup>1</sup> or BIGINT <sup>2</sup>	
BIT	BIT	
BINARY( <i>n</i> <MCL)	BINARY( <i>n</i> )	
BINARY( <i>n</i> >MCL)	IMAGE	
CHAR( <i>n</i> <MCL)	VARCHAR( <i>n</i> )	
CHAR( <i>n</i> >MCL)	TEXT	On download, ensure the values are not too long.
DATE	DATE <sup>3</sup> or DATE-TIME <sup>4</sup>	For Adaptive Server Enterprise DATETIME, the year must be in the range 1753-9999. For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.



SQL Anywhere or UltraLite data type	Adaptive Server Enterprise data type	Notes
DATETIME	DATETIME	<p>The Adaptive Server Enterprise DATETIME values are accurate to 1/300 second. The last digit of the fractional second is always 0, 3, or 6. Other digits are rounded to one of these three digits, so 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10.</p> <p>For download, SQL Anywhere keeps the original values from Adaptive Server Enterprise, but for upload, the values may not be exactly the original values.</p> <p>If DATETIME is used for a primary key, conflict resolution may fail. To successfully synchronize DATETIME, you should round the fractional second to 10 milliseconds. Also, the year must be in the range 1753-9999.</p>
DECIMAL( $p < 39, s$ )	DECIMAL( $p, s$ )	The precision of the Adaptive Server Enterprise NUMERIC can be from 1 to 38 digits ( $p < 39$ ).
DECIMAL( $p \geq 39, s$ )		There is no corresponding data type in Adaptive Server Enterprise.
DOUBLE	DOUBLE PRECISION	
FLOAT( $p$ )	FLOAT( $p$ )	
IMAGE	IMAGE	
INTEGER	INTEGER	
LONG BINARY	IMAGE	
LONG NVARCHAR	UNITEXT	
LONG VARBIT	TEXT	
LONG VARCHAR	TEXT	
MONEY	MONEY	
NCHAR( $c \leq \text{MCL}$ )	UNIVARCHAR( $c/2$ )	

SQL Anywhere or UltraLite data type	Adaptive Server Enterprise data type	Notes
NCHAR( $c > \text{MCL}$ )	UNITEXT	On download, ensure the values are not too long.
NTEXT	UNITEXT	
NUMERIC( $p < 39, s$ )	NUMERIC( $p, s$ )	The precision of the Adaptive Server Enterprise decimal can be from 1 to 38 digits ( $p < 39$ ).
NUMERIC( $p \geq 39, s$ )		
NVARCHAR( $c = < \text{MCL}$ )	UNIVARCHAR( $c/2$ )	
NVARCHAR( $c > \text{MCL}$ )	UNITEXT	On download, ensure the values are not too long.
REAL	REAL	
SMALLDATETIME	DATETIME <sup>4</sup>	SQL Anywhere and UltraLite SMALLDATETIME is implemented as TIMESTAMP.  The Adaptive Server Enterprise DATETIME is accurate to the minute. 29.998 seconds or lower are rounded down to the nearest minute; values with 29.999 seconds or higher are rounded up to the nearest minute. SQL Anywhere or UltraLite SMALLDATETIME is accurate to the microsecond. To successfully synchronize, SQL Anywhere or UltraLite SMALLDATETIME must be rounded to the minute. Also, the year must be in the range 1753-9999.
SMALLINT	SMALLINT	
SMALLMONEY	SMALLMONEY	
TEXT	TEXT	

SQL Anywhere or UltraLite data type	Adaptive Server Enterprise data type	Notes
TIME	TIME <sup>3</sup> or DATE-TIME <sup>4</sup>	<p>The Adaptive Server Enterprise TIME values are accurate to 1/300 second. The last digit of the fractional second is always 0, 3, or 6. Other digits are rounded to one of these three digits, so 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10. For download, SQL Anywhere keeps the original values from Adaptive Server Enterprise, but for upload, the values may not be exactly the original values. If TIME is used for a primary key, conflict resolution may fail. To successfully synchronize TIME, you should round the fractional second to 10 milliseconds.</p>
TIMESTAMP	DATETIME	<p>The Adaptive Server Enterprise DATETIME values are accurate to 1/300 second. The last digit of the fractional second is always 0, 3, or 6. Other digits are rounded to one of these three digits, so, 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10.</p> <p>For download, SQL Anywhere keeps the original values from Adaptive Server Enterprise, but for upload, the values may not be exactly the original values.</p> <p>If DATETIME is used for a primary key, conflict resolution may fail. To successfully synchronize DATETIME, you should round the fractional second to 10 milliseconds. Also, the year must be in the range 1753-9999.</p>
TINYINT	TINYINT	
UNIQUEIDENTIFIER	CHAR(36)	
UNIQUEIDENTIFIERSTR	CHAR(36)	Do not use UNIQUEIDENTIFIERSTR. Use UNIQUEIDENTIFIER instead.
UNSIGNED BIGINT	NUMERIC(20) <sup>1</sup> or UNSIGNED BIGINT <sup>2</sup>	
UNSIGNED INTEGER	UNSIGNED INT	

SQL Anywhere or UltraLite data type	Adaptive Server Enterprise data type	Notes
UNSIGNED SMALLINT	UNSIGNED SMALLINT	
UNSIGNED TINYINT	TINYINT	
VARBINARY( $n \leq \text{MCL}$ )	VARBINARY	
VARBINARY( $n > \text{MCL}$ )	IMAGE	
VARBIT( $n \leq \text{MCL}$ )	VARCHAR( $n$ )	
VARBIT( $n > \text{MCL}$ )	TEXT	
VARCHAR( $n \leq \text{MCL}$ )	VARCHAR( $n$ )	
VARCHAR( $n > \text{MCL}$ )	TEXT	
XML	TEXT	

<sup>1</sup> Only applies to Adaptive Server Enterprise before version 15.0.

<sup>2</sup> Only applies to Adaptive Server Enterprise version 15.0 or later.

<sup>3</sup> Only applies to Adaptive Server Enterprise version 12.5.1 or later.

<sup>4</sup> Only applies to Adaptive Server Enterprise before version 12.5.1.

### Mapping to SQL Anywhere or UltraLite remote data types

The following table identifies how Adaptive Server Enterprise consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type DOUBLE PRECISION on the consolidated database should be type DOUBLE on the remote database.

Adaptive Server Enterprise data type	SQL Anywhere or UltraLite data type	Notes
BIGINT <sup>1</sup>	BIGINT	
BINARY( $n$ )	BINARY( $n$ )	
BIT	BIT	

Adaptive Server Enterprise data type	SQL Anywhere or UltraLite data type	Notes
CHAR( <i>n</i> )	VARCHAR( <i>n</i> )	There is no equivalence between SQL Anywhere CHAR/NCHAR and Adaptive Server Enterprise CHAR/NCHAR. SQL Anywhere CHAR/NCHAR is equivalent to VARCHAR/NVARCHAR. You should not use CHAR/NCHAR in a consolidated database column that is synchronized. If you must use non-SQL Anywhere CHAR/NCHAR, run the Mobi-Link server with the -b option.
DATE	DATE	For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.
DATETIME	DATETIME	<p>The Adaptive Server Enterprise DATETIME values are accurate to 1/300 second. The last digit of the fractional second is always one of 0, 3, or 6. Other digit numbers are rounded to one of these three digits, so 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10.</p> <p>For download, SQL Anywhere keeps the original values from Adaptive Server Enterprise, but for upload, the values may not be exactly the original values. Conflict resolution may fail. To successfully synchronize DATETIME, you should round the fractional second to 10 milliseconds. Also, the year must be in the range 1753-9999.</p>
DECIMAL( <i>p,s</i> )	DECIMAL( <i>p,s</i> )	
DOUBLE PRECISION	DOUBLE	
FLOAT( <i>p</i> )	FLOAT( <i>p</i> )	
IMAGE	LONG BINARY	
INT	INT	
MONEY	MONEY	

Adaptive Server Enterprise data type	SQL Anywhere or UltraLite data type	Notes
NCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	The Adaptive Server Enterprise NCHAR and NVARCHAR store multibyte national character strings, they are different from SQL Anywhere NCHAR and NVARCHAR. In a multibyte environment, use SQL Anywhere or UltraLite VARCHAR.
NUMERIC( <i>p,s</i> )	NUMERIC( <i>p,s</i> )	
NVARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	The Adaptive Server Enterprise NCHAR and NVARCHAR store multibyte national character strings, they are different from SQL Anywhere NCHAR and NVARCHAR. In a multibyte environment, use SQL Anywhere or UltraLite VARCHAR.
REAL	REAL	
SMALLDATETIME	SMALLDATETIME	SQL Anywhere and UltraLite SMALLDATETIME is implemented as TIMESTAMP.  The Adaptive Server Enterprise SMALLDATETIME is accurate to the minute. 29.998 seconds or lower are rounded down to the nearest minute; values with 29.999 seconds or higher are rounded up to the nearest minute. SQL Anywhere or UltraLite SMALLDATETIME is accurate to the microsecond. To successfully synchronize, SQL Anywhere or UltraLite SMALLDATETIME must be rounded to the minute. Also, the year must be in the range 1900-2078.
SMALLINT	SMALLINT	
SMALLMONEY	SMALLMONEY	
TEXT	LONG VARCHAR	

Adaptive Server Enterprise data type	SQL Anywhere or UltraLite data type	Notes
TIME	TIME	<p>The Adaptive Server Enterprise TIME values are accurate to 1/300 second. The last digit of the fractional second is always one of 0, 3, or 6. Other digit numbers are rounded to one of these three digits, so 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10.</p> <p>For download, SQL Anywhere keeps the original values from Adaptive Server Enterprise, but for upload, the values may not be exactly the original values. Conflict resolution may fail. To successfully synchronize TIME, it is recommended that you round the fractional second to 10 milliseconds.</p>
TIMESTAMP	VARBINARY(8)	<p>Within Adaptive Server Enterprise, TIME- STAMP is a binary counter that gets incremented with every change to a row. So, each table can only contain one TIMESTAMP column and it does not make sense to synchronize it. If it must be in a synchronization, map it to a VARBINARY(8) data type in SQL Anywhere or UltraLite.</p> <p>This TIMESTAMP column cannot be explicitly inserted or updated, because it is maintained by the server. Keep this in mind when you are implementing upload scripts for tables that contain such columns.</p>
TINYINT	TINYINT	
UNSIGNED BIGINT <sup>1</sup>	UNSIGNED BIGINT	
UNSIGNED INT <sup>1</sup>	UNSIGNED INT	
UNSIGNED SMALLINT <sup>1</sup>	UNSIGNED SMALLINT	
VARBINARY( <i>n</i> )	VARBINARY( <i>n</i> )	
VARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	
UNICHAR( <i>n</i> )	NVARCHAR( <i>n</i> )	Not available in UltraLite.

<b>Adaptive Server Enterprise data type</b>	<b>SQL Anywhere or UltraLite data type</b>	<b>Notes</b>
UNITEXT <sup>1</sup>	LONG NVARCHAR	Not available in UltraLite.
UNIVARCHAR( <i>n</i> )	NVARCHAR( <i>n</i> )	Not available in UltraLite.

<sup>1</sup> Only applies to Adaptive Server Enterprise before version 15.0.



## IBM DB2 LUW data mapping

### Mapping to IBM DB2 LUW consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to IBM DB2 LUW consolidated data types. For example, a column of type BIT on the remote database should be type SMALLINT on the consolidated database.

When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length (MRL) based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.

SQL Anywhere or UltraLite data type	IBM DB2 LUW data type	Notes
BIGINT	BIGINT	
BINARY( $n < \text{MRL}$ )	VARCHAR( $n$ ) FOR BIT DATA	
BINARY( $n \geq \text{MRL}$ )	BLOB( $n$ )	
BIT	SMALLINT	
CHAR( $n < \text{MRL}$ )	VARCHAR( $n$ )	
CHAR( $n \geq \text{MRL}$ )	CLOB( $n$ )	DB2 values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
DATE	DATE	For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.
DATETIME	TIMESTAMP	
DECIMAL( $p < 32, s$ )	DECIMAL( $p, s$ )	The precision of SQL Anywhere DECIMAL is between 1 and 127. The maximum precision of DB2 DECIMAL is 31.
DECIMAL( $p \geq 32, s$ )		Any data of SQL Anywhere DECIMAL precision greater than 31 cannot be synchronized to DB2.

SQL Anywhere or UltraLite data type	IBM DB2 LUW data type	Notes
DOUBLE	DOUBLE	DOUBLE is an imprecise numeric data type that is subject to rounding. When working with different types of computers, the underlying storage of DOUBLE is often different, resulting in different rounding. DOUBLE is a bad choice to use in a primary key because primary keys are looking for equality. This is especially true in a synchronization environment because the consolidated database often runs on different hardware from the remote database.
FLOAT(1-24)	REAL	FLOAT can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
FLOAT(25-53)	DOUBLE	FLOAT can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
IMAGE	BLOB( <i>n</i> )	
INTEGER	INTEGER	
LONG BINARY	BLOB( <i>n</i> )	
LONG NVARCHAR	CLOB( <i>n</i> )	There is no corresponding data type in DB2. If the DB2 character set is Unicode, SQL Anywhere LONG NVARCHAR can synchronize to DB2 CLOB. UltraLite doesn't have LONG NVARCHAR.
LONG VARBIT	CLOB( <i>n</i> )	
LONG VARCHAR	CLOB( <i>n</i> )	
MONEY	DECIMAL(19,4)	

SQL Anywhere or UltraLite data type	IBM DB2 LUW data type	Notes
NCHAR( <i>c</i> )	VARCHAR( <i>n</i> ) or CLOB( <i>n</i> )	There is no corresponding data type in DB2. If the DB2 character set is Unicode, NCHAR can synchronize to DB2 VARCHAR or CLOB. The size of SQL Anywhere NCHAR is characters and the size of DB2 VARCHAR is bytes. If you map to VARCHAR, the total bytes of NCHAR can not be bigger than MRL. Otherwise, NCHAR should map to CLOB. It is difficult to calculate the number of bytes in NCHAR( <i>c</i> ), but it is approximately $c=n/4$ . In general, if <i>c</i> is less than MRL/4, map to VARCHAR( <i>n</i> ), but if <i>c</i> is greater than or equal to MRL/4, map to CLOB( <i>n</i> ).
NUMERIC( <i>p</i> <32, <i>s</i> )	NUMERIC( <i>p</i> , <i>s</i> )	
NUMERIC( <i>p</i> >=32, <i>s</i> )		There is no corresponding data type in DB2.
NTEXT	CLOB( <i>n</i> )	There is no corresponding data type in DB2. If the DB2 character set is Unicode, NTEXT can synchronize to DB2 CLOB.
NVARCHAR( <i>c</i> )	VARCHAR( <i>n</i> ) or CLOB( <i>n</i> )	There is no corresponding data type in DB2. If the DB2 character set is Unicode, NVARCHAR can synchronize to DB2 VARCHAR or CLOB. The size of SQL Anywhere NVARCHAR is characters and the size of DB2 VARCHAR is bytes. If you map to VARCHAR, the total bytes of NVARCHAR can not be bigger than MRL. Otherwise, NVARCHAR should map to CLOB. It is difficult to calculate the number of bytes in NVARCHAR( <i>c</i> ), but it is approximately $c=n/4$ . In general, if <i>c</i> is less than MRL/4, map to VARCHAR( <i>n</i> ), but if <i>c</i> is greater than or equal to MRL/4, map to CLOB( <i>n</i> ).
REAL	REAL	REAL can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
SMALLDATETIME	TIMESTAMP	

SQL Anywhere or UltraLite data type	IBM DB2 LUW data type	Notes
SMALLINT	SMALLINT	
SMALLMONEY	DECIMAL(10,4)	
TEXT	CLOB( <i>n</i> )	
TIME	TIMESTAMP or TIME	SQL Anywhere and UltraLite TIME values with fractional seconds require DB2 TIME-STAMP. SQL Anywhere and UltraLite time values with fractional seconds that are always zero can use DB2 TIME.
TIMESTAMP	TIMESTAMP	
TINYINT	SMALLINT	For download, DB2 values must be non-negative.
UNIQUEIDENTIFIER	CHAR(36)	
UNIQUEIDENTIFIERSTR	CHAR(36)	UNIQUEIDENTIFIERSTR is not recommended for DB2. Use UNIQUEIDENTIFIER instead.
UNSIGNED BIGINT	DECIMAL(20)	For download, DB2 values must be non-negative.
UNSIGNED INTEGER	DECIMAL(11)	For download, DB2 values must be non-negative.
UNSIGNED SMALLINT	DECIMAL(5)	For download, DB2 values must be non-negative.
UNSIGNED TINYINT	SMALLINT	For download, DB2 values must be non-negative.
VARBINARY( <i>n</i> <MRL)	VARCHAR( <i>n</i> ) FOR BIT DATA	
VARBINARY( <i>n</i> >=MRL)	BLOB( <i>n</i> )	
VARBIT( <i>n</i> <MRL)	VARCHAR( <i>n</i> )	
VARBIT( <i>n</i> >=MRL)	CLOB( <i>n</i> )	
VARCHAR( <i>n</i> <MRL)	VARCHAR( <i>n</i> )	

SQL Anywhere or UltraLite data type	IBM DB2 LUW data type	Notes
VARCHAR( $n \geq \text{MRL}$ )	CLOB( $n$ )	DB2 values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
XML	CLOB( $n$ )	

### Mapping to SQL Anywhere or UltraLite remote data types

The following table identifies how IBM DB2 LUW consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type INT on the consolidated database should be type INTEGER on the remote database.

When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.

IBM DB2 LUW data type	SQL Anywhere or UltraLite data type	Notes
BLOB	LONG BINARY	
BIGINT	BIGINT	
CHAR( $n$ )	VARCHAR( $n$ )	There is no equivalent to DB2 CHAR in SQL Anywhere. You should not use CHAR in a consolidated database column that is synchronized. If you must synchronize DB2 CHAR columns, run MobiLink server with the -b option.
CHAR( $n$ ) FOR BIT DATA	BINARY( $n$ )	
CLOB( $n$ )	LONG VARCHAR	
DATE	DATE	For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.
DBCLOB( $n$ )	LONG VARCHAR	The data type DBCLOB( $n$ ) is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, DBCLOB( $n$ ) is equivalent to CLOB.
DECIMAL( $p,s$ )	DECIMAL( $p,s$ )	

IBM DB2 LUW data type	SQL Anywhere or UltraLite data type	Notes
DOUBLE	DOUBLE	DOUBLE is an imprecise numeric data type that is subject to rounding. When working with different types of computers, the underlying storage of DOUBLE is often different, resulting in different rounding. DOUBLE is a bad choice to use in a primary key because primary keys are looking for equality. This is especially true in a synchronization environment because the consolidated database often runs on different hardware from the remote database.
FLOAT	DOUBLE	FLOAT can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
GRAPHIC( <i>n</i> )	VARCHAR( <i>2n</i> )	DB2 GRAPHIC does blank-padding, but SQL Anywhere CHAR does not. We recommend that you do not use this data type.  The data type GRAPHIC is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, GRAPHIC is equivalent to CHAR.
INT	INTEGER	
LONG VARCHAR	VARCHAR(32700)	
LONG VARCHAR FOR BIT DATA	VARBINARY(32700)	
LONG VARGRAPHIC( <i>n</i> )	VARCHAR(32700)	The data type LONG VARGRAPHIC is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, LONG VARGRAPHIC is equivalent LONG VARCHAR.
NUMERIC( <i>p,s</i> )	NUMERIC( <i>p,s</i> )	

IBM DB2 LUW data type	SQL Anywhere or UltraLite data type	Notes
REAL	REAL	REAL can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
SMALLINT	SMALLINT	
TIME	TIME	The fractional seconds values from SQL Anywhere TIME values are truncated on download. To avoid problems, do not use fractional seconds.
TIMESTAMP	TIMESTAMP	
VARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	
VARCHAR( <i>n</i> ) FOR BIT DATA	VARBINARY( <i>n</i> )	
VARGRAPHIC( <i>n</i> )	VARCHAR( <i>2n</i> )	The data type VARGRAPHIC is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, VARGRAPHIC is equivalent to VARCHAR.

## IBM DB2 mainframe data mapping

### Mapping to DB2 mainframe consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to DB2 mainframe consolidated data types. For example, a column of type BIT on the remote database should be type SMALLINT on the consolidated database.

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
BIGINT	DECIMAL(20)	
BINARY( $n < \text{MRL}$ )	VARCHAR( $n$ ) FOR BIT DATA	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.
BINARY( $n \geq \text{MRL}$ )	BLOB( $n$ )	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.
BIT	SMALLINT	



SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
CHAR( $n < \text{MRL}$ )	VARCHAR( $n$ )	<p>MRL is DB2 maximum row length.</p> <p>DB2 VARCHAR can only hold up to 32672 bytes depending on page size.</p> <p>DB2 has a maximum row length (MRL) based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. For details, see your DB2 documentation.</p> <p>SQL Anywhere CHAR is identical to SQL Anywhere VARCHAR.</p>
CHAR( $n \geq \text{MRL}$ )	CLOB( $n$ )	<p>DB2 has a maximum row length (MRL) based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. For details, see your DB2 documentation.</p> <p>SQL Anywhere CHAR is identical to SQL Anywhere VARCHAR.</p> <p>DB2 CLOB values can be longer than SQL Anywhere or UltraLite values. You must make sure that downloaded values are not too big.</p>
DATE	DATE	For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.
DATETIME	TIMESTAMP	
DECIMAL( $p < 32, s$ )	DECIMAL( $p, s$ )	The precision of SQL Anywhere DECIMAL is between 1 and 127. The maximum precision of DB2 DECIMAL is 31. Any data of SQL Anywhere DECIMAL precision greater than 31 cannot be synchronized to DB2.
DECIMAL( $p \geq 32, s$ )		Any data of SQL Anywhere DECIMAL precision greater than 31 cannot be synchronized to DB2.

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
DOUBLE	DOUBLE	<p>DOUBLE can cause problems if the consolidated and remote databases do not allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>To successfully synchronize, SQL Anywhere/ UltraLite DOUBLE values must be within the DB2 mainframe DOUBLE value range.</p>
FLOAT(1-24)	REAL	<p>FLOAT can cause problems if the consolidated and remote databases do not allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>To successfully synchronize, DB2 mainframe REAL values must be within the SQL Anywhere/ UltraLite REAL value range.</p>
FLOAT(25-53)	DOUBLE	<p>FLOAT can cause problems if the consolidated and remote databases do not allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>To successfully synchronize, SQL Anywhere/ UltraLite DOUBLE values must be within the DB2 mainframe DOUBLE value range.</p>
IMAGE	BLOB( <i>n</i> )	
INTEGER	INTEGER	
LONG BINARY	BLOB( <i>n</i> )	
LONG NVARCHAR	CLOB( <i>n</i> )	<p>There is no corresponding data type in DB2. If the DB2 character set is Unicode, SQL Anywhere long NVARCHAR can synchronize to DB2 CLOB. UltraLite does not have LONG NVARCHAR.</p>
LONG VARBIT	BLOB( <i>n</i> )	

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
LONG VARCHAR	CLOB( <i>n</i> )	
MONEY	DECIMAL(19,4)	
NCHAR( <i>c</i> )	VARCHAR( <i>n</i> ) or CLOB( <i>n</i> )	<p>SQL Anywhere NCHAR is identical to SQL Anywhere NVARCHAR.</p> <p>There is no corresponding data type in DB2. If the DB2 character set is Unicode, NCHAR can synchronize to DB2 VARCHAR or CLOB. The size of SQL Anywhere NCHAR is characters and the size of DB2 VARCHAR is bytes. If you map to VARCHAR, the total bytes of NCHAR cannot be bigger than MRL. Otherwise, NCHAR should map to CLOB. It is difficult to calculate the many bytes in NCHAR(<i>c</i>), but it is approximately <math>c=n/4</math>. In general, if <i>c</i> is less than MRL/4, map to VARCHAR(<i>n</i>), but if <i>c</i> is greater than or equal to MRL/4, map to CLOB(<i>n</i>).</p>
NTEXT	CLOB( <i>n</i> )	<p>SQL Anywhere NTEXT is identical to SQL Anywhere LONG NVARCHAR.</p> <p>There is no corresponding data type in DB2. If the DB2 character set is Unicode, NCHAR can synchronize to DB2 VARCHAR or CLOB. The size of SQL Anywhere NCHAR is characters and the size of DB2 VARCHAR is bytes. If you map to VARCHAR, the total bytes of NCHAR cannot be bigger than MRL. Otherwise, NCHAR should map to CLOB. It is difficult to calculate the many bytes in NCHAR(<i>c</i>), but it is approximately <math>c=n/4</math>. In general, if <i>c</i> is less than MRL/4, map to VARCHAR(<i>n</i>), but if <i>c</i> is greater than or equal to MRL/4, map to CLOB(<i>n</i>).</p>
NUMERIC( <i>p</i> <32, <i>s</i> )	NUMERIC( <i>p</i> , <i>s</i> )	
NUMERIC( <i>p</i> >=32, <i>s</i> )		There is no corresponding data type in DB2.

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
NVARCHAR( <i>c</i> )	VARCHAR( <i>n</i> ) or CLOB( <i>n</i> )	There is no corresponding data type in DB2. If the DB2 character set is Unicode, NCHAR can synchronize to DB2 VARCHAR or CLOB. The size of SQL Anywhere NCHAR is characters and the size of DB2 VARCHAR is bytes. If you map to VARCHAR, the total bytes of NCHAR cannot be bigger than MRL. Otherwise, NCHAR should map to CLOB. It is difficult to calculate the many bytes in NCHAR( <i>c</i> ), but it is approximately $c=n/4$ . In general, if <i>c</i> is less than MRL/4, map to VARCHAR( <i>n</i> ), but if <i>c</i> is greater than or equal to MRL/4, map to CLOB( <i>n</i> ).
REAL	REAL	REAL can cause problems if the consolidated and remote databases do not allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.  To successfully synchronize, DB2 mainframe REAL values must be within the SQL Anywhere/UltraLite REAL value range.
SMALLDATETIME	TIMESTAMP	
SMALLINT	SMALLINT	
SMALLMONEY	DECIMAL(10,4)	
TEXT	CLOB( <i>n</i> )	SQL Anywhere TEXT is identical to SQL Anywhere LONG VARCHAR.
TIME	TIMESTAMP or TIME	SQL Anywhere and UltraLite TIME values with fractional seconds require DB2 TIMESTAMP. SQL Anywhere and UltraLite time values with fractional seconds that are always zero can use DB2 TIME.
TIMESTAMP	TIMESTAMP	
TINYINT	SMALLINT	
UNIQUEIDENTIFIER	CHAR(36)	

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
UNIQUEIDENTIFIERSTR	CHAR(36)	UNIQUEIDENTIFIERSTR is not recommended for DB2. Use UNIQUEIDENTIFIER instead.
UNSIGNED BIGINT	DECIMAL(20)	DB2 values must be non-negative.
UNSIGNED INTEGER	DECIMAL(11)	DB2 values must be non-negative.
UNSIGNED SMALLINT	DECIMAL(5)	DB2 values must be non-negative.
UNSIGNED TINYINT	SMALLINT	DB2 values must be non-negative.
VARBINARY( $n < \text{MRL}$ )	VARCHAR( $n$ ) FOR BIT DATA	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.
VARBINARY( $n \geq \text{MRL}$ )	BLOB( $n$ )	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
VARBIT( $n < \text{MRL}$ )	VARCHAR( $n$ )	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.
VARBIT( $n \geq \text{MRL}$ )	CLOB( $n$ )	When creating a DB2 table, you need to pay attention to the DB2 page size. DB2 has a maximum row length based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. The length of all columns in a table can't exceed the above limitation. If a table has a BLOB or CLOB column, you count row length using the LOB locator, not BLOB or CLOB data directly. For details, see the DB2 documentation.
VARCHAR( $n < \text{MRL}$ )	VARCHAR( $n$ )	<p>DB2 VARCHAR can only hold up to 32672 bytes, depending on page size.</p> <p>DB2 has a maximum row length (MRL) based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. For details, see your DB2 documentation.</p> <p>SQL Anywhere CHAR is identical to SQL Anywhere VARCHAR.</p>

SQL Anywhere or UltraLite data type	IBM DB2 mainframe data type	Notes
VARCHAR( $n \geq \text{MRL}$ )	CLOB( $n$ )	DB2 has a maximum row length (MRL) based on the page size: the MRL is 4005 when the page size is 4K, 8101 when 8K, 16293 when 16K and 32677 when 32K. For details, see your DB2 documentation.  SQL Anywhere CHAR is identical to SQL Anywhere VARCHAR.  DB2 CLOB values can be longer than SQL Anywhere or UltraLite values. You must make sure that downloaded values are not too big.
XML	CLOB( $n$ )	SQL Anywhere XML is identical to SQL Anywhere LONG VARCHAR.

### Mapping to SQL Anywhere or UltraLite remote data types

The following table identifies how DB2 mainframe consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type INT on the consolidated database should be type INTEGER on the remote database.

IBM DB2 mainframe data type	SQL Anywhere or UltraLite data type	Notes
BLOB	LONG BINARY	
CHAR( $n$ )	VARCHAR( $n$ )	There is no equivalent to DB2 CHAR in SQL Anywhere. You should not use CHAR in a consolidated database column that is synchronized. If you must synchronize DB2 CHAR columns, run MobiLink server with the -b option.
CHAR( $n$ ) FOR BIT DATA	BINARY( $n$ )	There is no equivalent to DB2 CHAR in SQL Anywhere. You should not use CHAR in a consolidated database column that is synchronized. If you must synchronize DB2 CHAR columns, run MobiLink server with the -b option.
CLOB( $n$ )	LONG VARCHAR	
DATE	DATE	For SQL Anywhere and UltraLite, the time value must in the format 00:00:00.

IBM DB2 mainframe data type	SQL Anywhere or UltraLite data type	Notes
DBCLOB( <i>n</i> )	LONG VARCHAR	The data type DBCLOB( <i>n</i> ) is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, DBCLOB( <i>n</i> ) is equivalent to CLOB.
DECIMAL( <i>p,s</i> )	DECIMAL( <i>p,s</i> )	
DOUBLE	DOUBLE	<p>DOUBLE can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>To successfully synchronize, SQL Anywhere/UltraLite DOUBLE values must be within the DB2 DOUBLE value range.</p>
FLOAT	DOUBLE	<p>FLOAT can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>To successfully synchronize, SQL Anywhere/UltraLite DOUBLE values must be within the DB2 DOUBLE value range.</p>
GRAPHIC( <i>n</i> )	VARCHAR( <i>2n</i> )	<p>DB2 GRAPHIC does blank-padding. There is no equivalent to DB2 GRAPHIC in SQL Anywhere. You should not use GRAPHIC in a consolidated database column that is synchronized.</p> <p>The data type GRAPHIC is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, GRAPHIC is equivalent to CHAR.</p>
INT	INTEGER	
NUMERIC( <i>p,s</i> )	NUMERIC( <i>p,s</i> )	



IBM DB2 mainframe data type	SQL Anywhere or Ultra-Lite data type	Notes
REAL	DOUBLE	<p>REAL can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.</p> <p>On DB2 mainframe, REAL, DOUBLE and FLOAT have the same range, between <math>-7.2E+75</math> and <math>7.2E+75</math>. The largest negative value is about <math>-5.4E-79</math>, and the smallest positive value is about <math>5.4E-79</math>. The range of SA/UL REAL is <math>-3.402823e+38</math> to <math>3.402823e+38</math>, and DOUBLE is <math>2.22507385850721e-308</math> to <math>1.79769313486231e+308</math>. To successfully synchronize, DB2 mainframe REAL must be in SA/UL REAL range and SA/UL DOUBLE must be in DB2 mainframe DOUBLE range.</p>
ROWID		There is no corresponding data type in SQL Anywhere or UltraLite. ROWID is maintained by DB2 server. This data type cannot be synchronized.
SMALLINT	SMALLINT	
TIME	TIME	The fractional seconds values from SQL Anywhere TIME values are truncated on download. To avoid problems, do not use fractional seconds.
TIMESTAMP	TIMESTAMP	
VARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	
VARCHAR( <i>n</i> ) FOR BIT DATA	VARBINARY( <i>n</i> )	

<b>IBM DB2 mainframe data type</b>	<b>SQL Anywhere or Ultra-Lite data type</b>	<b>Notes</b>
VARGRAPHIC( <i>n</i> )	VARCHAR( <i>2n</i> )	The data type VARGRAPHIC is only used for double-byte characters. SQL Anywhere does not have a corresponding data type. When the DB2 character set is Unicode, VARGRAPHIC is equivalent to VARCHAR.

## Microsoft SQL Server data mapping

### Mapping to Microsoft SQL Server consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to Microsoft SQL Server consolidated data types. For example, a column of type DATE on the remote database should be type DATETIME on the consolidated database.

SQL Anywhere or UltraLite data type	Microsoft SQL Server data type	Notes
BIGINT	BIGINT	
BINARY( $n \leq 8000$ )	VARBINARY( $n$ )	
BINARY( $n > 8000$ )	VARBINARY(MAX)	
BIT	BIT	
CHAR( $n \leq 8000$ )	VARCHAR( $n$ )	
CHAR( $n > 8000$ )	VARCHAR(MAX)	
DATE	DATE	
DATETIME	DATETIME2	SQL Server DATETIME2 and TIME values are accurate to 100 nanoseconds. However, TIMESTAMP and TIME values are only accurate to 1 microsecond. To successfully synchronize DATETIME2 and TIME, we suggest you round the fractional second to 1 microsecond.
DECIMAL( $p \leq 38, s$ )	DECIMAL( $p, s$ )	SQL Server DECIMAL/NUMERIC precision ranges from 1 to 38, so $p$ must be less than 39.
DECIMAL( $p > 38, s$ )		There is no corresponding data type in SQL Server.
DOUBLE	FLOAT(53)	
FLOAT( $p$ )	FLOAT( $p$ )	
IMAGE	VARBINARY(MAX)	
INTEGER	INT	

SQL Anywhere or UltraLite data type	Microsoft SQL Server data type	Notes
LONG BINARY	VARBINARY(MAX)	
LONG NVARCHAR	NVARCHAR(MAX)	
LONG VARBIT	VARCHAR(MAX)	
LONG VARCHAR	VARCHAR(MAX)	
MONEY	MONEY	
NCHAR( $n \leq 4000$ )	NVARCHAR( $c$ )	
NCHAR( $n > 4000$ )	NVARCHAR(MAX)	
NTEXT	NVARCHAR(MAX)	
NUMERIC( $p \leq 38, s$ )	NUMERIC( $p, s$ )	SQL Server DECIMAL/NUMERIC precision ranges from 1 to 38, so $p$ must be less than 39.
NUMERIC( $p > 38, s$ )		There is no corresponding data type in SQL Server.
NVARCHAR( $n \leq 4000$ )	NVARCHAR( $c$ )	
NVARCHAR( $n > 4000$ )	NVARCHAR(MAX)	
REAL	REAL	
SMALLDATETIME	SMALLDATETIME	SQL Anywhere and UltraLite SMALLDATETIME is implemented as TIMESTAMP. SQL Server SMALLDATETIME is accurate to the minute. 29.998 seconds or lower are rounded down to the nearest minute; values with 29.999 seconds or higher are rounded up to the nearest minute. SQL Anywhere or UltraLite SMALLDATETIME is accurate to the microsecond. To successfully synchronize, SQL Anywhere or UltraLite SMALLDATETIME must be rounded to the minute. The year must be in the range 1900-2078.
SMALLINT	SMALLINT	
SMALLMONEY	SMALLMONEY	

SQL Anywhere or UltraLite data type	Microsoft SQL Server data type	Notes
TEXT	VARCHAR(MAX)	
TIME	TIME	SQL Server DATETIME2 and TIME values are accurate to 100 nanoseconds. However, TIMESTAMP and TIME values are only accurate to 1 microsecond. To successfully synchronize DATETIME2 and TIME, we suggest you round the fractional second to 1 microsecond.
TIMESTAMP	DATETIME2	SQL Server DATETIME2 and TIME values are accurate to 100 nanoseconds. However, TIMESTAMP and TIME values are only accurate to 1 microsecond. To successfully synchronize DATETIME2 and TIME, we suggest you round the fractional second to 1 microsecond.
TINYINT	TINYINT	For download, values must be non-negative.
UNIQUEIDENTIFIER	UNIQUEIDENTIFIER	
UNIQUEIDENTIFIERSTR	UNIQUEIDENTIFIER	
UNSIGNED BIGINT	NUMERIC(20)	For download, values must be non-negative.
UNSIGNED INTEGER	NUMERIC(11)	For download, values must be non-negative.
UNSIGNED TINYINT	TINYINT	For download, values must be non-negative.
UNSIGNED SMALLINT	INT	For download, values must be non-negative.
VARBINARY( $n \leq 8000$ )	VARBINARY( $n$ )	
VARBINARY( $n > 8000$ )	VARBINARY(MAX)	
VARBIT( $n \leq 8000$ )	VARCHAR( $n$ )	
VARBIT( $n > 8000$ )	VARCHAR(MAX)	
VARCHAR( $n \leq 8000$ )	VARCHAR( $c$ )	
VARCHAR( $n > 8000$ )	VARCHAR(MAX)	

SQL Anywhere or UltraLite data type	Microsoft SQL Server data type	Notes
XML	XML or VARCHAR(MAX)	For SQL Server 2005, use XML. For other versions, use VARCHAR(MAX).

### Mapping to SQL Anywhere or UltraLite remote data types

The following table identifies how Microsoft SQL Server consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type TEXT on the remote database should be type LONG VARCHAR on the consolidated database.

Microsoft SQL Server data type	SQL Anywhere or UltraLite data type	Notes
BIGINT	BIGINT	
BINARY( <i>n</i> )	BINARY( <i>n</i> )	
BIT	BIT	
CHAR( <i>n</i> )	VARCHAR( <i>n</i> )	A Microsoft SQL Server CHAR column is blank padded. A SQL Anywhere CHAR column is not blank padded by default and is equivalent to a VARCHAR column. Therefore, try to avoid using the CHAR data type in the synchronization tables in Microsoft SQL Server. If you must use the CHAR data type in the Microsoft SQL Server consolidated database, please run the MobiLink server with the -b command line option to help resolve the differences between SQL Anywhere CHAR and non-SQL Anywhere CHAR.
DATE	DATE	

Microsoft SQL Server data type	SQL Anywhere or UltraLite data type	Notes
DATETIME	TIMESTAMP or DATETIME	SQL Server DATETIME values are accurate to 1/300 second. The last digit of the fractional second is always 0, 3, or 6. Other digits are rounded to one of these three digits, so, 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10. For download, SQL Anywhere keeps the original values from SQL Server, but for upload, the values may not be exactly the original values. If DATETIME is used for a primary key, conflict resolution may fail. To successfully synchronize DATETIME, you should round the fractional second to 10 milliseconds. The year must be in the range 1753-9999.
DATETIME2	TIMESTAMP	SQL Server DATETIME2 and TIME values are accurate to 100 nanoseconds. However, TIMESTAMP and TIME values are only accurate to 1 microsecond. To successfully synchronize DATETIME2 and TIME, we suggest you round the fractional second to 1 microsecond.
DECIMAL( <i>p,s</i> )	DECIMAL( <i>p,s</i> )	
FLOAT( <i>p</i> )	FLOAT( <i>p</i> )	
IMAGE	LONG BINARY	
INT	INT	
MONEY	MONEY	
NCHAR( <i>n</i> )	NVARCHAR( <i>c</i> )	Not available in UltraLite.  There is no equivalence between SQL Anywhere NCHAR and non-SQL Anywhere NCHAR. SQL Anywhere NCHAR is equivalent to NVARCHAR. You should not use NCHAR in a consolidated database column that is synchronized. If you must use non-SQL Anywhere NCHAR, run the MobiLink server with the -b option.
NTEXT	LONG NVARCHAR	Not available in UltraLite.

Microsoft SQL Server data type	SQL Anywhere or UltraLite data type	Notes
NVARCHAR( <i>c</i> )	NVARCHAR( <i>c</i> )	Not available in UltraLite.
NVARCHAR(MAX)	LONG NVARCHAR	
NUMERIC( <i>p,s</i> )	NUMERIC( <i>p,s</i> )	
REAL	REAL	REAL can cause problems if the consolidated and remote databases don't allow the exact same (imprecise) values. We do not test all possible values, so care must be taken. To avoid problems, do not use these types as part of a primary key.
SMALLDATETIME	SMALLDATETIME	SQL Anywhere and UltraLite SMALLDATETIME is implemented as TIMESTAMP. SQL Server SMALLDATETIME is accurate to the minute. 29.998 seconds or lower are rounded down to the nearest minute; values with 29.999 seconds or higher are rounded up to the nearest minute. SQL Anywhere or UltraLite SMALLDATETIME is accurate to the microsecond. To successfully synchronize, SQL Anywhere or UltraLite SMALLDATETIME must be rounded to the minute. The year must be in the range 1900-2078.
SMALLINT	SMALLINT	
SMALLMONEY	SMALLMONEY	
TEXT	LONG VARCHAR	
TIME	TIME	SQL Server DATETIME2 and TIME values are accurate to 100 nanoseconds. However, TIMESTAMP and TIME values are only accurate to 1 microsecond. To successfully synchronize DATETIME2 and TIME, we suggest you round the fractional second to 1 microsecond.



Microsoft SQL Server data type	SQL Anywhere or UltraLite data type	Notes
TIMESTAMP	VARBINARY(8)	<p>Within Microsoft SQL Server, <b>TIMESTAMP</b> is a binary counter that gets incremented with every change to a row. So, each table can only contain one <b>TIMESTAMP</b> column and it does not make sense to synchronize it. If it must be in a synchronization, map it to a <b>VARBINARY(8)</b> data type in SQL Anywhere or UltraLite.</p> <p>This timestamp column cannot be explicitly inserted or updated, because it is maintained by the server. Keep this in mind when you are implementing upload scripts for tables that contain such columns.</p>
TINYINT	TINYINT	
UNIQUEIDENTIFIER	UNIQUEIDENTIFIER	
VARBINARY( <i>n</i> )	VARBINARY( <i>n</i> )	
VARBINARY(MAX)	LONG BINARY	
VARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	
VARCHAR(MAX)	LONG VARCHAR	
XML	XML	

## MySQL data mapping

### Mapping to MySQL consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to MySQL consolidated data types. For example, a column of type TEXT on the remote database should be type LONGTEXT on the consolidated database.

SQL Anywhere or UltraLite data type	MySQL data type	Notes
BIGINT	BIGINT	
BINARY( $n \leq 255$ )	BINARY( $n$ )	
BINARY( $n > 255$ )	BLOB	
BIT	BIT	
CHAR( $n \leq 255$ )	CHAR( $n$ )	
CHAR( $n > 255$ )	TEXT( $n$ )	
DATE	DATE	The year must range from 1000 to 9999.
DATETIME	DATETIME	The MySQL DATETIME data type does not support fractional seconds. The year must range from 1000 to 9999.
DECIMAL( $p \leq 65, s \leq 30$ )	DECIMAL( $p, s$ )	
DECIMAL( $p > 65, s > 30$ )		There is no corresponding data type in MySQL if the precision is greater than 65 or if the scale is greater than 30.
DOUBLE	DOUBLE	
FLOAT	FLOAT	
IMAGE	LONGBLOB	
INTEGER	INTEGER	
LONG BINARY	LONGBLOB	
LONG NVARCHAR	LONGTEXT CHARACTER SET UTF8	
LONG VARBIT	LONGTEXT	

SQL Anywhere or UltraLite data type	MySQL data type	Notes
LONG VARCHAR	LONGTEXT	
MONEY	NUMERIC(19,4)	
NCHAR( $n \leq 255$ )	CHAR( $n$ ) CHARACTER SET UTF8	
NCHAR( $n > 255$ )	TEXT CHARACTER SET UTF8	
NTEXT	LONGTEXT CHARACTER SET UTF8	
NUMERIC( $p \leq 65, s \leq 30$ )	DECIMAL( $p, s$ )	
NUMERIC( $p > 65, s > 30$ )		There is no corresponding data type in MySQL.
NVARCHAR( $n$ )	VARCHAR( $n$ ) CHARACTER SET UTF8	
REAL	REAL	
SMALLDATETIME	DATETIME	The MySQL DATETIME data type does not support fractional seconds. The year must range from 1000 to 9999.
SMALLINT	SMALLINT	
SMALLMONEY	NUMERIC(10,4)	
TEXT	LONGTEXT	
TIME	TIME	The MySQL TIME data type does not support fractional seconds.
TIMESTAMP	DATETIME	The MySQL DATETIME data type does not support fractional seconds. The year must range from 1000 to 9999.
TINYINT	TINYINT UNSIGNED	TINYINT is always unsigned in SQL Anywhere and UltraLite.
UNIQUEIDENTIFIER	CHAR(36)	

SQL Anywhere or UltraLite data type	MySQL data type	Notes
UNIQUEIDENTIFIERSTR	CHAR(36)	
VARBINARY( <i>n</i> )	VARCHAR( <i>n</i> )	
VARBIT( <i>n</i> ≤8000)	VARCHAR( <i>n</i> )	
VARBIT( <i>n</i> >8000)	TEXT	
VARCHAR( <i>n</i> )	VARCHAR( <i>n</i> )	
XML	LONGTEXT	

### Mapping to SQL Anywhere or UltraLite remote data types

The following table identifies how MySQL consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type BOOL on the consolidated database should be type BIT on the remote database.

MySQL data type	SQL Anywhere or UltraLite data type	Notes
BIGINT	BIGINT	
BINARY( <i>n</i> )	BINARY( <i>n</i> )	
BIT(1)	BIT	
BIT( <i>n</i> >1)	UNSIGNED BIGINT	
BLOB( <i>n</i> ≤32767)	VARBINARY( <i>n</i> )	
BLOB( <i>n</i> >32767)	IMAGE	
BOOL	BIT	
CHAR( <i>n</i> )	CHAR( <i>n</i> )	
DATE	DATE	The year must range from 1000 to 9999.
DATETIME	DATETIME	The MySQL DATETIME data type does not support fractional seconds. The year must range from 1000 to 9999.
DOUBLE	DOUBLE	
DECIMAL	DECIMAL	

MySQL data type	SQL Anywhere or UltraLite data type	Notes
ENUM		There is no corresponding data type in SQL Anywhere or UltraLite.
GEOMETRY		There is no corresponding data type in SQL Anywhere or UltraLite.
INTEGER	INTEGER	
LINestring		There is no corresponding data type in SQL Anywhere or UltraLite.
LONGBLOB	IMAGE	
LONGTEXT	TEXT	
MEDIUMBLOB	IMAGE	
MEDIUMINT	INTEGER	
MEDIUMTEXT	TEXT	
MULTILINESTRING		There is no corresponding data type in SQL Anywhere or UltraLite.
MULTIPOINT		There is no corresponding data type in SQL Anywhere or UltraLite.
MULTIPOLYGON		There is no corresponding data type in SQL Anywhere or UltraLite.
NCHAR	NCHAR	Not available in UltraLite.
NUMERIC	NUMERIC	
NVARCHAR	NVARCHAR	Not available in UltraLite.
POINT		There is no corresponding data type in SQL Anywhere or UltraLite.
POLYGON		There is no corresponding data type in SQL Anywhere or UltraLite.
REAL	REAL	
SET		There is no corresponding data type in SQL Anywhere or UltraLite.

MySQL data type	SQL Anywhere or UltraLite data type	Notes
SMALLINT	SMALLINT	
TEXT( $n \leq 32767$ )	VARCHAR( $n$ )	
TEXT( $n > 32767$ )	TEXT	
TIME	TIME	The MySQL TIME data type does not support fractional seconds. The range of TIME in MySQL is '-838:59:59' to '838:59:59'. The range of TIME in SQL Anywhere or UltraLite is '00:00:00.000000' to '23:59:59.999999'.
TIMESTAMP	TIMESTAMP	The MySQL DATETIME data type does not support fractional seconds. The year must range from 1000 to 9999. Although MySQL offers automatic initialization and updating on TIMESTAMP columns, SQL Anywhere and UltraLite only offers automatic initialization.
TINYBLOB	VARBINARY	
TINYINT	SMALLINT	TINYINT is always unsigned in SQL Anywhere and UltraLite. Must be a positive value.
TINYINT UNSIGNED	TINYINT	TINYINT is always unsigned in SQL Anywhere and UltraLite.
TINYTEXT	VARCHAR	
VARBINARY( $n \leq 32767$ )	VARBINARY( $n$ )	
VARBINARY( $n > 32767$ )	IMAGE	
VARCHAR( $n \leq 32767$ )	VARCHAR( $n$ )	
VARCHAR( $n > 32767$ )	TEXT	
YEAR[(2 4)]	INTEGER	SQL Anywhere and UltraLite do not support the YEAR data type. YEAR needs to be mapped to INTEGER in a remote database. The INTEGER value must range from 1000 to 9999.

## Oracle data mapping

### Mapping to Oracle consolidated data types

The following table identifies how SQL Anywhere and UltraLite remote data types are mapped to Oracle consolidated data types. For example, a column of type BIT on the remote database should be type NUMBER on the consolidated database.

SQL Anywhere or UltraLite data type	Oracle data type	Notes
BIGINT	NUMBER(20)	
BINARY( $n \leq 2000$ )	RAW( $n$ )	
BINARY( $n > 2000$ )	BLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
BIT	NUMBER(1)	
CHAR( $n \leq 4000$ )	VARCHAR2( $n$ byte)	Oracle VARCHAR2 allows you to specify the maximum number of bytes or characters. The maximum length of VARCHAR2 data is 4000 bytes. If you specify the character number, make sure the maximum data length is not over 4000 bytes.
CHAR( $n > 4000$ )	CLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
DATE	DATE <sup>2</sup> or TIME-STAMP	SQL Anywhere or UltraLite fractional seconds cannot be preserved when using an Oracle DATE data type which has no fractional seconds. To avoid problems, do not use fractional seconds. The year must be in the range 1-9999.  When using the Interactive SQL utility, turn off the Return_date_time_as_string option before executing your SQL statement.

SQL Anywhere or UltraLite data type	Oracle data type	Notes
DATETIME	DATE <sup>2</sup> or TIME-STAMP	<p>SQL Anywhere or UltraLite fractional seconds cannot be preserved when using an Oracle DATE data type which has no fractional seconds. To avoid problems, do not use fractional seconds. The year must be in the range 1-9999.</p> <p>When using the Interactive SQL utility, turn off the Return_date_time_as_string option before executing your SQL statement.</p>
DECIMAL( $p \leq 38, s$ )	NUMBER( $p, 0 \leq s \leq 38$ )	<p>In SQL Anywhere DECIMAL, <math>p</math> is between 1 and 127, and <math>s</math> is always less than or equal to <math>p</math>. In Oracle NUMBER, <math>p</math> ranges from 1 to 38, and <math>s</math> ranges from -84 to 127. To synchronize, the Oracle NUMBER scale must be restricted to between 0 and 38.</p>
DECIMAL( $p > 38, s$ )		<p>There is no corresponding data type in Oracle.</p>
DOUBLE	DOUBLE PRECISION or BINARY_DOUBLE <sup>1</sup>	<p>The special values INF, -INF and NAN of Oracle 10g BINARY_FLOAT and BINARY_DOUBLE cannot be synchronized with SQL Anywhere or UltraLite.</p>
FLOAT( $p$ )	FLOAT( $p$ )	
IMAGE	BLOB	<p>Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.</p>
INTEGER	INT	
LONG BINARY	BLOB	<p>Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.</p>
LONG NVARCHAR	NCLOB	<p>Oracle CLOB and NCLOB can hold up to 4G of data. SQL Anywhere LONG VARCHAR and LONG NVARCHAR can only hold up to 2G.</p> <p>Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.</p>



SQL Anywhere or UltraLite data type	Oracle data type	Notes
LONG VARBIT	CLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
LONG VARCHAR	CLOB	Oracle CLOB and NCLOB can hold up to 4G of data. SQL Anywhere LONG VARCHAR and LONG NVARCHAR can only hold up to 2G.  Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
MONEY	NUMBER(19,4)	
NCHAR( <i>c</i> )	NVARCHAR2( <i>c</i> char) or NCLOB	The size of SQL Anywhere NCHAR and Oracle NVARCHAR2 indicates the maximum number of Unicode characters. The data length of Oracle NVARCHAR2 can't be over 4000 bytes. It is difficult to calculate the maximum byte length from character size. In general, if the size is over 1000, map to NCLOB, otherwise map to NVARCHAR2.
NTEXT	NCLOB	Oracle NCLOB can hold up to 4G of data. SQL Anywhere NTEXT (or LONG NVARCHAR) can only hold up to 2G.  Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
NUMERIC( <i>p</i> ≤38, <i>s</i> )	NUMBER( <i>p</i> , 0≤ <i>s</i> ≤38)	In SQL Anywhere NUMERIC, <i>p</i> is between 1 and 127, and <i>s</i> is always less than or equal to <i>p</i> . In Oracle NUMBER, <i>p</i> ranges from 1 to 38, and <i>s</i> ranges from -84 to 127. To synchronize, the Oracle NUMBER scale must be restricted to between 0 and 38.
NUMERIC( <i>p</i> >38, <i>s</i> )		There is no corresponding data type in Oracle.

SQL Anywhere or UltraLite data type	Oracle data type	Notes
NVARCHAR	NVARCHAR2( <i>c</i> char) or NCLOB	The size of SQL Anywhere NCHAR and Oracle NVARCHAR2 indicates the maximum number of Unicode characters. The data length of Oracle NVARCHAR2 can't be over 4000 bytes. It is difficult to calculate the maximum byte length from character size. In general, if the size is over 1000, map to NCLOB, otherwise map to NVARCHAR2.
REAL	REAL or BINARY_FLOAT <sup>1</sup>	The special values INF, -INF and NAN of Oracle 10g BINARY_FLOAT and BINARY_DOUBLE cannot be synchronized with SQL Anywhere or UltraLite.
SMALLDATETIME	DATE <sup>2</sup> or TIME-STAMP	SQL Anywhere or UltraLite fractional seconds cannot be preserved when using an Oracle DATE data type which has no fractional seconds. To avoid problems, do not use fractional seconds. The year must be in the range 1-9999.
SMALLINT	NUMBER(5)	
SMALLMONEY	NUMBER(10,4)	
TEXT	CLOB	<p>Oracle CLOB can hold up to 4G of data. SQL Anywhere TEXT (or LONG VARCHAR) can only hold up to 2G.</p> <p>Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.</p>
TIME	DATE <sup>2</sup> or TIME-STAMP	<p>SQL Anywhere or UltraLite fractional seconds cannot be preserved when using an Oracle DATE data type which has no fractional seconds. To avoid problems, do not use fractional seconds.</p> <p>When using the Interactive SQL utility, turn off the Return_date_time_as_string option before executing your SQL statement.</p>

SQL Anywhere or UltraLite data type	Oracle data type	Notes
TIMESTAMP	DATE <sup>2</sup> or TIME-STAMP	SQL Anywhere or UltraLite fractional seconds cannot be preserved when using an Oracle DATE data type which has no fractional seconds. To avoid problems, do not use fractional seconds. The year must be in the range 1-9999.  When using the Interactive SQL utility, turn off the Return_date_time_as_string option before executing your SQL statement.
TINYINT	NUMBER(3)	For download, Oracle values must be non-negative.
UNSIGNED BIGINT	NUMBER(20)	For download, Oracle values must be non-negative.
UNSIGNED INTEGER	NUMBER(11)	For download, Oracle values must be non-negative.
UNSIGNED SMALLINT	NUMBER(5)	For download, Oracle values must be non-negative.
UNSIGNED TINYINT	NUMBER(3)	For download, Oracle values must be non-negative.
UNIQUEIDENTIFIER	CHAR(36)	
UNIQUEIDENTIFIERSTR	CHAR(36)	UNIQUEIDENTIFIERSTR is not recommended to use for Oracle. Use UNIQUEIDENTIFIER instead.
VARBINARY( $n \leq 2000$ )	RAW( $n$ )	
VARBINARY( $n > 2000$ )	BLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
VARBIT( $n \leq 4000$ )	VARCHAR2( $n$ byte)	
VARBIT( $n > 4000$ )	CLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.

SQL Anywhere or UltraLite data type	Oracle data type	Notes
VARCHAR( $n \leq 4000$ )	VARCHAR2( $n$ byte)	Oracle VARCHAR2 allows you to specify the maximum number of bytes or characters. The maximum length of VARCHAR2 data is 4000 bytes. If you specify the character number, make sure the maximum data length is not over 4000 bytes.
VARCHAR( $n > 4000$ )	CLOB	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
XML	CLOB	Oracle CLOB and NCLOB can hold up to 4G of data. SQL Anywhere XML can only hold up to 2G.  Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.

<sup>1</sup> Only applies to Oracle version 10g or later.

<sup>2</sup> Only applies to Oracle version 8i.

**Note**

The LONG data types are deprecated in Oracle 8, 8i and 9i.

For Oracle LONG data types to synchronize properly, you must check the **Oracle Force Retrieval Of Long Columns** option in the **ODBC Data Source Configuration** window of the iAnywhere Solutions Oracle ODBC driver.

**Mapping to SQL Anywhere or UltraLite remote data types**

The following table identifies how Oracle consolidated data types are mapped to SQL Anywhere and UltraLite remote data types. For example, a column of type LONG on the consolidated database should be type LONG VARCHAR on the remote database.

Oracle data type	SQL Anywhere or UltraLite data type	Notes
BFILE	LONG BINARY	Download only.  Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.

Oracle data type	SQL Anywhere or UltraLite data type	Notes
BINARY_DOUBLE	DOUBLE	The special values INF, -INF and NAN of BINARY_FLOAT cannot be synchronized with SQL Anywhere or UltraLite. The precision of FLOAT and DOUBLE in Oracle is different from SQL Anywhere and UltraLite. The value of the data may change depending on the precision.
BINARY_FLOAT	REAL	The special values INF, -INF and NAN of BINARY_FLOAT cannot be synchronized with SQL Anywhere or UltraLite. The precision of FLOAT and DOUBLE in Oracle is different from SQL Anywhere and UltraLite. The value of the data may change depending on the precision.
BLOB	LONG BINARY	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
CHAR( <i>n</i> byte)	VARCHAR( <i>n</i> )	There is no equivalence between SQL Anywhere CHAR and Oracle CHAR. SQL Anywhere CHAR is equivalent to VARCHAR. You should not use CHAR/NCHAR in a consolidated database column that is synchronized. If you must use non-SQL Anywhere CHAR, run the MobiLink server with the -b option.  SQL Anywhere or UltraLite values can be longer than Oracle values, so make sure values are not too big when uploading.
CLOB	LONG VARCHAR	Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.
DATE	TIMESTAMP	The year must be in the range 1-9999.
INTERVAL YEAR( <i>year_precision</i> ) TO MONTH		There is no corresponding data type in SQL Anywhere or UltraLite.
INTERVAL DAY( <i>day_precision</i> ) TO SECOND( <i>p</i> )		There is no corresponding data type in SQL Anywhere or UltraLite.

Oracle data type	SQL Anywhere or UltraLite data type	Notes
LONG	LONG VARCHAR	
LONG RAW	LONG BINARY	
NCHAR( <i>c</i> char)	NVARCHAR( <i>c</i> )	<p>There is no equivalence between SQL Anywhere NCHAR and Oracle NCHAR. SQL Anywhere NCHAR is equivalent to NVARCHAR. You should not use NCHAR in a consolidated database column that is synchronized. If you must use non-SQL Anywhere NCHAR, run the MobiLink server with the -b option.</p> <p>SQL Anywhere or UltraLite values can be longer than Oracle values, so make sure values are not too big when uploading.</p>
NCLOB	LONG NVARCHAR	<p>Not available in UltraLite.</p> <p>Oracle values can be longer than SQL Anywhere or UltraLite values, so make sure values are not too big when downloading.</p>
NUMBER( <i>p,s</i> )	NUMBER( <i>p,s</i> )	<p>In SQL Anywhere NUMBER, <i>p</i> is between 1 and 127, and <i>s</i> is always less than or equal to <i>p</i>. In Oracle NUMBER, <i>p</i> ranges from 1 to 38, and <i>s</i> ranges from -84 to 127. To synchronize, the Oracle NUMBER scale must be between 0 and 38.</p>
NVARCHAR2( <i>c</i> char)	NVARCHAR( <i>c</i> )	<p>Not available in UltraLite.</p> <p>SQL Anywhere or UltraLite values can be longer than Oracle values, so make sure values are not too big when uploading.</p>
RAW	BINARY	<p>SQL Anywhere or UltraLite values can be longer than Oracle values, so make sure values are not too big when uploading.</p>
ROWID	VARCHAR(64)	<p>UROWID and ROWID are read-only and so are unlikely to be synchronized.</p>

Oracle data type	SQL Anywhere or UltraLite data type	Notes
TIMESTAMP( $p \leq 6$ )	TIMESTAMP	When $p < 6$ , you may need to ensure SQL Anywhere or UltraLite values have the same precision. Otherwise, conflict detection may fail and/or duplicate rows may result. The year must be in the range 1-9999.
TIMESTAMP( $p > 6$ )		There is no corresponding data type in SQL Anywhere or UltraLite.
TIMESTAMP( $p$ ) WITH LOCAL TIME ZONE		There is no corresponding data type in SQL Anywhere or UltraLite.
TIMESTAMP( $p$ ) WITH TIME ZONE		There is no corresponding data type in SQL Anywhere or UltraLite.
UROWID	VARCHAR(64)	UROWID and ROWID are read-only and so are unlikely to be synchronized.
VARCHAR2( $n$ byte)	VARCHAR( $n$ )	SQL Anywhere or UltraLite values can be longer than Oracle values, so make sure values are not too big when uploading.

---



---

# Character set considerations

## Contents

Character set considerations ..... 790

---

## Character set considerations

Each character of text is represented in one or more bytes. The mapping from characters to binary codes is called the **character set encoding**. Some character sets used for languages with small alphabets, such as European languages, use a single-byte representation. Others, such as Unicode, use a double-byte representation. Because they use twice the storage space for each character, double-byte character sets can represent a much larger number of characters.

Conversion errors can occur or data can be lost when text using one character set must be converted to another character set. Not all characters can be represented in all character sets. In particular, single-byte character sets can represent a much smaller number of characters than multibyte systems because of the limited number of codes available.

When the character set of your MobiLink remote database is the same as your consolidated database, character conversion issues are avoided.

Text often needs to be sorted to build indexes and to prepare ordered result sets, such as directory listings. The **sort order** identifies the order of the characters. For example, a sort order typically states that the letter "a" comes before the letter "b", which comes before the letter "c".

Each database has a **collation sequence**. You set the collation sequence when you create the database, although how you do so can differ between database systems. The collation sequence defines both the character set and the sort order for that database.

### Tip

Whenever possible, define the collation sequence of your remote database to be the same as that of your consolidated database. This arrangement reduces the chance of erroneous conversions.

### See also

- SQL Anywhere clients: “[International languages and character sets](#)” [*SQL Anywhere Server - Database Administration*]
- UltraLite clients: “[UltraLite character sets](#)” [*UltraLite - Database Management and Reference*]
- Information specific to your RDBMS: “[MobiLink consolidated databases](#)” on page 3

## Character set conversion during synchronization

During synchronization, characters may need to be converted from one character set to another. The following conversions occur as characters are passed between the remote application and the consolidated database.

### Character set conversion during upload

The MobiLink client sends data to the MobiLink server using the character set of the remote database.

1. The MobiLink server communicates with the consolidated database using the Unicode ODBC API. To do so, the MobiLink server converts all characters received from the remote database into Unicode and sends the Unicode to the ODBC driver.

2. If necessary, the ODBC driver for the consolidated database server converts the characters from Unicode into the character set of your consolidated database. This conversion is controlled solely by the ODBC driver for your consolidated database system. So, behavior can differ between two different database systems, particularly systems made by different manufacturers. MobiLink synchronization works with several database systems. Check the documentation of your particular consolidated server and ODBC driver for details.

### **Character set conversion during download**

1. The ODBC driver for the consolidated database system receives characters in the coding of the consolidated database. It converts these characters into Unicode to pass them through the Unicode API to the MobiLink server. This conversion is controlled solely by the ODBC driver for your consolidated database system. Check the documentation of your particular consolidated server and ODBC driver for details.
2. The MobiLink server receives characters through the Unicode ODBC API. If the remote database uses a different character set, the MobiLink server converts the characters before downloading them.

### **Examples**

- UltraLite applications on Windows Mobile devices use the Unicode character set.  
When you synchronize a Windows Mobile application, no character conversion occurs within the MobiLink server. The server finds that data arriving from the application is already in Unicode and passes it directly to the ODBC driver. Similarly, no character set conversion is necessary when downloading data.
- All SQL Anywhere databases and all UltraLite applications on platforms other than Windows Mobile use the character set determined by the collating sequence of the remote database.  
When you synchronize a remote database, the MobiLink server performs character set conversions between the character set of the remote database and Unicode.

## **Controlling ODBC driver character set conversion**

Because most consolidated databases are unlikely to use Unicode, it is important to understand how the ODBC driver for your consolidated database system converts data to and from Unicode. Some ODBC drivers use the language settings of the computer running MobiLink to determine what character set to use. In these cases, it is best if the language and code-page settings of the computer running the MobiLink server match those of the consolidated database.

Other ODBC drivers, such as the driver for Sybase Adaptive Server Enterprise, allow each connection to use a specific character set. To avoid conversion errors, the character set used by MobiLink should be set to match that of the consolidated database.

For a detailed description of how character set conversions take place in your consolidated database server's ODBC driver, consult that product's ODBC driver documentation.

---

---

# iAnywhere Solutions ODBC drivers for MobiLink

## Contents

ODBC drivers supported by MobiLink .....	794
iAnywhere Solutions Oracle driver .....	795

---

## ODBC drivers supported by MobiLink

The MobiLink server can work with a variety of consolidated databases and ODBC drivers, as shown in the table below. Some drivers, though compatible for use with MobiLink, may have functional restrictions associated with their use.

For more information about supported versions, see <http://www.sybase.com/detail?id=1002288>.

Database	ODBC Driver
SQL Anywhere 11	SQL Anywhere 11 <sup>1</sup>
Oracle 10g or 11g	iAnywhere Solutions 11 - Oracle <sup>1</sup>
Microsoft SQL Server	Microsoft SQL Server ODBC driver <sup>2</sup>
Sybase Adaptive Server Enterprise 12.5.1 or later	Sybase Adaptive Server Enterprise driver <sup>2</sup>
IBM DB2 LUW 8.1 or 8.2 for Windows, Linux and Unix	IBM DB2 8.2 CLI driver <sup>2</sup>
IBM DB2 LUW 9.x for Windows, Linux and Unix	IBM DB2 9.1 CLI driver <sup>2</sup>
IBM DB2 mainframe 8.1 for Z/OS	IBM DB2 8.2 CLI driver <sup>2</sup>
IBM DB2 mainframe 9.1 for Z/OS	IBM DB2 9.1 CLI driver <sup>2</sup>
MySQL 5.1	MySQL ODBC driver 5.1 <sup>2</sup>

<sup>1</sup> Provided with SQL Anywhere version 11. See [Recommended ODBC Drivers for MobiLink](#).

<sup>2</sup> Not provided with SQL Anywhere version 11. For installation and configuration instructions, see [Recommended ODBC Drivers for MobiLink](#).

## iAnywhere Solutions Oracle driver

The iAnywhere Solutions 11 - Oracle ODBC driver is custom-tailored for use with iAnywhere software. This driver does not work with third-party software.

If you use Oracle with MobiLink or OMNI, you must install an Oracle client on the same computer as this Oracle driver.

The Oracle driver can be configured using the ODBC Administrator, the `.odbc.ini` file (in Unix), or the `dbdsn` utility.

The following table provides the configuration options for the Oracle driver.

Windows ODBC Administrator	Configuration for dbdsn command line or <code>.odbc.ini</code> file	Description
Data source name	For <code>dbdsn</code> , use the <code>-w</code> option.	A name to identify your data source.
User ID	<b>UserID</b> In <code>dbdsn</code> , set this option in the connection string.	The default logon ID that the application uses to connect to your Oracle database. If you leave this field blank, you are prompted for the information when you connect.
Password	<b>Password</b> In <code>dbdsn</code> , set this option in the connection string.	The password that the application uses to connect to your Oracle database. If you leave this field blank, you are prompted for the information when you connect.
SID	<b>SID</b>	The TNS Service Name that is stored in <code>network/admin/tnsnames.ora</code> under your Oracle installation directory.
Enable Microsoft distributed transactions	For <code>dbdsn</code> , use the <b>enableMSDIC</b> option in the connection string. Not supported for <code>.odbc.ini</code> .	Select this checkbox if you want to enlist your transactions in the Microsoft Distributed Transaction Coordinator. When selected, the Oracle ODBC driver requires an Oracle binary file, <code>oramts.dll</code> for Oracle 9i clients or <code>oramts10.dll</code> for Oracle 10g clients.
Encrypt Password	For <code>dbdsn</code> , use the <code>-pe</code> option. Not supported for <code>.odbc.ini</code> .	Select this checkbox if you want the password to be stored in encrypted form in the data source.

Windows ODBC Administrator	Configuration for dbdsn command line or .odbc.ini file	Description
Procedures Return Results	<b>ProcResults</b> In dbdsn, set this option in the connection string.	Select this field if your stored procedures can return results. The default is that procedures do not return results (not selected). If your download_cursor or download_delete_cursor scripts are stored procedure invocations, set this to yes.
Array Size	<b>ArraySize</b> In dbdsn, set this option in the connection string.	The size, in bytes, of the byte array used to pre-fetch rows, on a per-statement basis. The default is 60000. Increasing this value can significantly improve fetch performance (such as during MobiLink server downloads) at the cost of extra memory allocation.

## Windows configuration

### To create a DSN for the Oracle driver in Windows

1. Open the ODBC Administrator:
  - Choose **Start » Programs » SQL Anywhere 11 » ODBC Administrator**.
 The **ODBC Data Source Administrator** appears.
2. Click **Add**.
3. Choose **iAnywhere Solutions 11 - Oracle**.
4. Specify the configuration options you need. The fields are explained above.
5. Click **Test Connection**, and then click **OK**.

## Unix configuration

On Unix, if you are setting up the driver in an ODBC system information file (typically called *.odbc.ini*), the section for this driver should appear as follows (with appropriate values entered for each field):

```
[sample_dsn_using_the_ias_odbc_driver_for_oracle]
Driver=full-path/libdboraodbc10_r.so
UserID=user-id
Password=password
SID=TNS-service-name
ProcResults=[yes|no]
ArraySize=bytes
```

For an explanation of each field, see above.



**DBDSN configuration**

To create an Oracle DSN with the dbdsn utility, use the following syntax:

**dbdsn -w** *data-source-name* **-or -c** *configuration-options*

The *configuration-options* are described above.

For example:

```
dbdsn -w MyOracleDSN -or -pe -c  
Userid=dba;Password=sql;SID=abcd;ArraySize=100000;ProcResults=y;enableMSDIC=n
```

See “Data Source utility (dbdsn)” [[SQL Anywhere Server - Database Administration](#)].

**See also**

- [Recommended ODBC Drivers for MobiLink](#)
- “Data Source utility (dbdsn)” [[SQL Anywhere Server - Database Administration](#)]

---

---

# Deploying MobiLink applications

## Contents

Introduction to MobiLink deployment .....	800
Deploying the MobiLink server .....	801
Deploying SQL Anywhere MobiLink clients .....	813
Deploying UltraLite MobiLink clients .....	815
Deploying QAnywhere applications .....	816

---

## Introduction to MobiLink deployment

Deploying MobiLink applications involves the following activities:

- Deploy the MobiLink server into a production setting.
- If required, deploy the Redirector.
- Deploy any SQL Anywhere MobiLink clients.
- Deploy any UltraLite MobiLink clients.

This chapter describes the files you need to include in your application's installation program for each of these items.

The **Deploy Synchronization Model Wizard** that can help with your deployment on Windows. See [“Using the Deployment Wizard”](#) [*SQL Anywhere Server - Programming*].

**Check your license agreement**

Redistribution of files is subject to your license agreement. No statements in this document override anything in your license agreement. Check your license agreement before considering deployment.

## Deploying the MobiLink server

The simplest way to deploy a MobiLink server into a production environment is to install a licensed copy of SQL Anywhere onto the production computer.

However, if you are redistributing a MobiLink server in a separate installation program, you may want to include only a subset of the files. In this case, you need to include the following files in your installation.

### Notes

- Test on a clean computer before redistributing.
- Files must be installed to the SQL Anywhere installation directory, with the exception of samples.
- The files should be in the same directory unless otherwise noted.
- When a location is given, the files must be copied into a directory of the same name.
- On Unix, environment variables must be set for the system to locate SQL Anywhere applications and libraries. It is recommended that you use the appropriate file for your shell, either *sa\_config.sh* or *sa\_config.csh* (located in the directory *install-dir/bin32* for 32-bit environments and *install-dir/bin64* for 64-bit environments) as a template for setting the required environment variables. Some of the environment variables set by the *sa\_config* files include `PATH`, `LD_LIBRARY_PATH`, `SQLANY11`, and `SQLANY11SAMP11`.
- On Windows, the `PATH` environment variable must be set for the system to locate SQL Anywhere applications and libraries. Check the `PATH` variable to ensure that it includes *install-dir/bin32* for 32-bit environments or *install-dir/bin64* for 64-bit environments. If both entries exist, remove the path that does not apply to your environment.
- To use Java synchronization logic, and to use the graphical administration tools (Sybase Central and the MobiLink Monitor), you must have JRE 1.6.0 installed.
- To deploy Sybase Central, see [“Deploying administration tools” \[SQL Anywhere Server - Programming\]](#).
- There is a deployment wizard for Windows. See [“Using the Deployment Wizard” \[SQL Anywhere Server - Programming\]](#).

### Windows 32-bit applications

All directories are relative to *install-dir*. For more details on the file structure of a 64-bit Windows environment, see [“Windows 64-bit applications” on page 804](#).

Description	Windows files
MobiLink server	<ul style="list-style-type: none"> <li>• <i>bin32\mlodbc11.dll</i></li> <li>• <i>bin32\mlsrv11.exe</i></li> <li>• <i>bin32\mlsrv11.lic</i></li> <li>• <i>bin32\mlsql11.dll</i></li> <li>• <i>bin32\dbicu11.dll</i></li> <li>• <i>bin32\dbicudt11.dll</i></li> </ul>

Description	Windows files
Language library	<ul style="list-style-type: none"> <li>● <i>bin32\dblgen11.dll</i><sup>1</sup></li> </ul>
Synchronization stream libraries (to support version 8 and 9 clients)	<ul style="list-style-type: none"> <li>● <i>bin32\mlhttp11.dll</i></li> <li>● <i>bin32\mlsock11.dll</i></li> </ul>
Java synchronization logic	<ul style="list-style-type: none"> <li>● <i>java\activation.jar</i><sup>2</sup></li> <li>● <i>java\imap.jar</i><sup>2</sup></li> <li>● <i>java\jodbc.jar</i></li> <li>● <i>java\log4j.jar</i><sup>2</sup></li> <li>● <i>java\mailapi.jar</i><sup>2</sup></li> <li>● <i>java\mlscript.jar</i></li> <li>● <i>java\mlsupport.jar</i></li> <li>● <i>java\pop3.jar</i><sup>2</sup></li> <li>● <i>java\smtp.jar</i><sup>2</sup></li> <li>● <i>bin32\mljava11.dll</i></li> <li>● <i>bin32\dbjodbc11.dll</i></li> <li>● <i>bin32\mljodbc11.dll</i></li> </ul>
.NET synchronization logic	<ul style="list-style-type: none"> <li>● <i>MobiLink\setup\dnet\mlDomConfig.xml</i></li> <li>● <i>bin32\mldnet11.dll</i></li> <li>● <i>bin32\dnetodbc11.dll</i></li> <li>● <i>Assembly\v2\iAnywhere.MobiLink.dll</i></li> <li>● <i>Assembly\v2\iAnywhere.MobiLink.Script.dll</i></li> <li>● <i>Assembly\v2\iAnywhere.MobiLink.Script.xml</i></li> <li>● <i>bin32\mlDomConfig.xsd</i></li> </ul>
Security option for version 10 and 11 clients (mlsrv11 -x) <sup>3</sup>	<ul style="list-style-type: none"> <li>● <i>bin32\mecc_tls11.dll</i></li> <li>● <i>bin32\mlrsa_tls11.dll</i></li> <li>● <i>bin32\mlrsa_tls_fips11.dll</i></li> <li>● <i>bin32\sbgse2.dll</i></li> </ul>
Security option <sup>3</sup> for version 8 and 9 clients (mlsrv11 -xo) <sup>5</sup>	<ul style="list-style-type: none"> <li>● <i>bin32\mlhttps11.dll</i></li> <li>● <i>bin32\mlhttpsfips11.dll</i></li> <li>● <i>bin32\mlrsafips11.dll</i></li> <li>● <i>bin32\mljrsa11.dll</i></li> <li>● <i>bin32\mljtls11.dll</i></li> <li>● <i>bin32\mlrsa11.dll</i></li> <li>● <i>bin32\mltls11.dll</i></li> <li>● <i>bin32\defaultmem.dll</i></li> <li>● <i>bin32\libs.b.dll</i></li> </ul>
Setup scripts (deploy the ones for your consolidated database)	<ul style="list-style-type: none"> <li>● <i>MobiLink\setup\</i></li> <li>● <i>MobiLink\upgrade\</i></li> </ul>

Description	Windows files
mluser utility	<ul style="list-style-type: none"> <li>• <i>bin32\mluser.exe</i></li> <li>• <i>bin32\mlodbc11.dll</i></li> <li>• <i>bin32\dbicu11.dll</i></li> <li>• <i>bin32\dbicudt11.dll</i></li> </ul>
mlstop utility	<ul style="list-style-type: none"> <li>• <i>bin32\mlstop.exe</i></li> <li>• <i>bin32\dbicu11.dll</i></li> </ul>
MobiLink Monitor	<ul style="list-style-type: none"> <li>• <i>java\mlmon.jar</i></li> <li>• <i>java\JComponents1100.jar</i></li> <li>• <i>java\mlstream.jar</i></li> <li>• <i>java\jsyblib600.jar</i></li> <li>• <i>Sun\JavaHelp-2_0\jh.jar</i></li> <li>• <i>Sun\jaxb1.0\</i></li> <li>• <i>bin32\jsyblib600.dll</i></li> <li>• <i>bin32\mlmon.exe</i></li> </ul> <p>For security with the MobiLink Monitor:<sup>3</sup></p> <ul style="list-style-type: none"> <li>• <i>bin32\mlcecc11.dll</i></li> <li>• <i>bin32\mlcrsa11.dll</i></li> <li>• <i>bin32\mlcrsafips11.dll</i></li> <li>• <i>bin32\mlczlib11.dll</i></li> </ul>
Online help for the MobiLink plug-in and Monitor	<ul style="list-style-type: none"> <li>• <i>\documentation\en\htmlhelp\dbadmin_en11.chm<sup>1</sup></i></li> <li>• <i>\documentation\en\htmlhelp\dbadmin_en11.map<sup>1</sup></i></li> </ul>
MobiLink Redirector	<ul style="list-style-type: none"> <li>• <i>MobiLink\redirector</i></li> </ul>
Notifier	<ul style="list-style-type: none"> <li>• <i>java\activation.jar<sup>2</sup></i></li> <li>• <i>java\jodbc.jar</i></li> <li>• <i>java\log4j.jar<sup>4</sup></i></li> <li>• <i>java\mailapi.jar<sup>2</sup></i></li> <li>• <i>java\mlnotif.jar</i></li> <li>• <i>java\mlscript.jar</i></li> <li>• <i>java\smtp.jar<sup>2</sup></i></li> <li>• <i>bin32\mljodbc11.dll</i></li> <li>• <i>bin32\mljstrm11.dll</i></li> </ul>

Description	Windows files
MobiLink server files required by QAnywhere	<ul style="list-style-type: none"> <li>● Notifier files</li> <li>● <i>java\commons-logging.jar</i></li> <li>● <i>java\commons-codec-1.3.jar</i></li> <li>● <i>java\commons-httpclient-3.0.jar</i></li> <li>● <i>java\jasyb600.jar</i></li> <li>● <i>java\log4j.jar<sup>4</sup></i></li> <li>● <i>java\mlscript.jar</i></li> <li>● <i>java\mlstream.jar</i></li> <li>● <i>java\qaconnector.jar</i></li> <li>● <i>bin32\jasyb600.dll</i></li> </ul>
Relay Server Outbound Enabler	<ul style="list-style-type: none"> <li>● <i>bin32\rsoe.exe</i></li> </ul> <p>For security with the Outbound Enabler:</p> <ul style="list-style-type: none"> <li>● <i>bin32\mlcrsa11.dll</i></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

<sup>2</sup> If you are redistributing an application, you must obtain these files directly from Sun.

<sup>3</sup> ECC and FIPS require that you obtain the separately-licensed SQL Anywhere security option and is subject to export regulations. RSA security is included with SQL Anywhere for version 10 and later. To order this component, see [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

<sup>4</sup> If you are redistributing an application, you must obtain this file directly from Apache.

<sup>5</sup> You must also create a registry key called *HKEY\_LOCAL\_MACHINE\SOFTWARE\Certicom\libs* and add a REG\_BINARY value named *expectedtag* with the data 5B0F4FA6E24AEF3B4407052EB04902711FD991B6.

## Windows 64-bit applications

All directories are relative to *install-dir*. For more details on the file structure of a 32-bit Windows environment, see [“Windows 32-bit applications” on page 801](#).

Description	Windows files
MobiLink server	<ul style="list-style-type: none"> <li>● <i>bin64\mlodbc11.dll</i></li> <li>● <i>bin64\mlsrv11.exe</i></li> <li>● <i>bin64\mlsrv11.lic</i></li> <li>● <i>bin64\mlsql11.dll</i></li> <li>● <i>bin64\dbicu11.dll</i></li> <li>● <i>bin64\dbicudt11.dll</i></li> </ul>
Language library	<ul style="list-style-type: none"> <li>● <i>bin64\dblgen11.dll<sup>1</sup></i></li> </ul>



Description	Windows files
Synchronization stream libraries (to support version 8 and 9 clients)	<ul style="list-style-type: none"> <li>• <i>bin64\mlhttp11.dll</i></li> <li>• <i>bin64\mlsock11.dll</i></li> </ul>
Java synchronization logic	<ul style="list-style-type: none"> <li>• <i>java\activation.jar<sup>2</sup></i></li> <li>• <i>java\imap.jar<sup>2</sup></i></li> <li>• <i>java\jdbc.jar</i></li> <li>• <i>java\log4j.jar<sup>2</sup></i></li> <li>• <i>java\mailapi.jar<sup>2</sup></i></li> <li>• <i>java\mlscript.jar</i></li> <li>• <i>java\mlsupport.jar</i></li> <li>• <i>java\pop3.jar<sup>2</sup></i></li> <li>• <i>java\smtp.jar<sup>2</sup></i></li> <li>• <i>bin64\mljava11.dll</i></li> <li>• <i>bin64\dbjdbc11.dll</i></li> <li>• <i>bin64\mljdbc11.dll</i></li> </ul>
.NET synchronization logic	<ul style="list-style-type: none"> <li>• <i>MobiLink\setup\dnet\mlDomConfig.xml</i></li> <li>• <i>bin64\mldnet11.dll</i></li> <li>• <i>bin64\dnetodbc11.dll</i></li> <li>• <i>Assembly\v2\iAnywhere.MobiLink.dll</i></li> <li>• <i>Assembly\v2\iAnywhere.MobiLink.Script.dll</i></li> <li>• <i>Assembly\v2\iAnywhere.MobiLink.Script.xml</i></li> <li>• <i>bin64\mlDomConfig.xsd</i></li> </ul>
Security option for version 10 and 11 clients (mlsrv11 -x) <sup>3</sup>	<ul style="list-style-type: none"> <li>• <i>bin64\mlecc_tls11.dll</i></li> <li>• <i>bin64\mlrsa_tls11.dll</i></li> <li>• <i>bin64\mlrsa_tls_fips11.dll</i></li> <li>• <i>bin64\sbgse2.dll</i></li> </ul>
Security option <sup>3</sup> for version 8 and 9 clients (mlsrv11 -xo) <sup>5</sup>	<ul style="list-style-type: none"> <li>• <i>bin64\mlhttps11.dll</i></li> <li>• <i>bin64\mlhttpsfips11.dll</i></li> <li>• <i>bin64\mlrsafips11.dll</i></li> <li>• <i>bin64\mljrsa11.dll</i></li> <li>• <i>bin64\mljtls11.dll</i></li> <li>• <i>bin64\mlrsa11.dll</i></li> <li>• <i>bin64\mtls11.dll</i></li> <li>• <i>bin64\defaultmem.dll</i></li> <li>• <i>bin64\libs.dll</i></li> </ul>
Setup scripts (deploy the ones for your consolidated database)	<ul style="list-style-type: none"> <li>• <i>MobiLink\setup\</i></li> <li>• <i>MobiLink\upgrade\</i></li> </ul>
mluser utility	<ul style="list-style-type: none"> <li>• <i>bin64\mluser.exe</i></li> <li>• <i>bin64\mlodbc11.dll</i></li> <li>• <i>bin64\dbicu11.dll</i></li> <li>• <i>bin64\dbicudt11.dll</i></li> </ul>

Description	Windows files
mlstop utility	<ul style="list-style-type: none"> <li>● <i>bin64\mlstop.exe</i></li> <li>● <i>bin64\dbicu11.dll</i></li> </ul>
MobiLink Monitor	<ul style="list-style-type: none"> <li>● <i>java\mlmon.jar</i></li> <li>● <i>java\JComponents1100.jar</i></li> <li>● <i>java\mlstream.jar</i></li> <li>● <i>java\jsyblib600.jar</i></li> <li>● <i>Sun\JavaHelp-2_0\jh.jar</i></li> <li>● <i>Sun\jaxb1.0\</i></li> <li>● <i>bin64\jsyblib600.dll</i></li> <li>● <i>bin64\mlmon.exe</i></li> </ul> <p>For security with the MobiLink Monitor:<sup>3</sup></p> <ul style="list-style-type: none"> <li>● <i>bin64\mlcecc11.dll</i></li> <li>● <i>bin64\mlcrsa11.dll</i></li> <li>● <i>bin64\mlcrsafips11.dll</i></li> <li>● <i>bin64\mlczlib11.dll</i></li> </ul>
Online help for the MobiLink plug-in and Monitor	<ul style="list-style-type: none"> <li>● <i>\documentation\en\htmlhelp\dbadmin_en11.chm<sup>1</sup></i></li> <li>● <i>\documentation\en\htmlhelp\dbadmin_en11.map<sup>1</sup></i></li> </ul>
MobiLink Redirector	<ul style="list-style-type: none"> <li>● <i>MobiLink\redirector</i></li> </ul>
Notifier	<ul style="list-style-type: none"> <li>● <i>java\activation.jar<sup>2</sup></i></li> <li>● <i>java\jodbc.jar</i></li> <li>● <i>java\log4j.jar<sup>4</sup></i></li> <li>● <i>java\mailapi.jar<sup>2</sup></i></li> <li>● <i>java\mlnotif.jar</i></li> <li>● <i>java\mlscript.jar</i></li> <li>● <i>java\smtp.jar<sup>2</sup></i></li> <li>● <i>bin64\mljodbc11.dll</i></li> <li>● <i>bin64\mljstrm11.dll</i></li> </ul>
MobiLink server files required by QAnywhere	<ul style="list-style-type: none"> <li>● Notifier files</li> <li>● <i>java\commons-logging.jar</i></li> <li>● <i>java\commons-codec-1.3.jar</i></li> <li>● <i>java\commons-httpclient-3.0.jar</i></li> <li>● <i>java\jsyblib600.jar</i></li> <li>● <i>java\log4j.jar<sup>4</sup></i></li> <li>● <i>java\mlscript.jar</i></li> <li>● <i>java\mlstream.jar</i></li> <li>● <i>java\qaconnector.jar</i></li> <li>● <i>bin64\jsyblib600.dll</i></li> </ul>

Description	Windows files
Relay Server Outbound Enabler	<ul style="list-style-type: none"> <li>• <i>bin64\rsoe.exe</i></li> </ul> <p>For security with the Outbound Enabler:</p> <ul style="list-style-type: none"> <li>• <i>bin64\mlcrsa11.dll</i></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

<sup>2</sup> If you are redistributing an application, you must obtain these files directly from Sun.

<sup>3</sup> ECC and FIPS require that you obtain the separately-licensed SQL Anywhere security option and is subject to export regulations. RSA security is included with SQL Anywhere for version 10 and later. To order this component, see [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

<sup>4</sup> If you are redistributing an application, you must obtain this file directly from Apache.

<sup>5</sup> You must also create a registry key called *HKEY\_LOCAL\_MACHINE\SOFTWARE\Certicom\libs* and add a REG\_BINARY value named *expectedtag* with the data 5B0F4FA6E24AEF3B4407052EB04902711FD991B6.

### Unix 32-bit applications on Unix, Linux, and Macintosh

All directories are relative to *install-dir*. For more details on the file structure of a 64-bit Unix environment, see [“Unix 64-bit applications on Unix and Linux” on page 810](#).

Description	Unix files
MobiLink server	<ul style="list-style-type: none"> <li>• <i>bin32/mlsrv11</i></li> <li>• <i>bin32/mlsrv11.lic</i></li> <li>• <i>lib32/libdbodm11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libmlodbc11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libmlsql11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libdbtasks11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libdbicu11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libdbicudt11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libdbodbcinst11_r.so<sup>3</sup></i></li> </ul>
Language library	<ul style="list-style-type: none"> <li>• <i>res/dblgen11.res<sup>1</sup></i></li> </ul>
Synchronization stream libraries for version 8 and 9 clients (deploy the ones you use)	<ul style="list-style-type: none"> <li>• <i>lib32/libmlhttp11_r.so<sup>3</sup></i></li> <li>• <i>lib32/libmlsock11_r.so<sup>3</sup></i></li> </ul>

Description	Unix files
Java synchronization logic	<ul style="list-style-type: none"> <li>● <i>java/activation.jar</i><sup>2</sup></li> <li>● <i>java/imap.jar</i><sup>2</sup></li> <li>● <i>java/jodbc.jar</i></li> <li>● <i>java/log4j.jar</i><sup>5</sup></li> <li>● <i>java/mailapi.jar</i><sup>2</sup></li> <li>● <i>java/mlscript.jar</i></li> <li>● <i>java/mlsupport.jar</i></li> <li>● <i>java/pop3.jar</i><sup>2</sup></li> <li>● <i>java/smtp.jar</i><sup>2</sup></li> <li>● <i>lib32/libmljava11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmljodbc11.so</i><sup>3</sup></li> </ul>
.NET synchronization logic	<ul style="list-style-type: none"> <li>● Not applicable</li> </ul>
Security option for version 10 and 11 clients (mlsrv11 -x) <sup>4</sup>	<ul style="list-style-type: none"> <li>● <i>lib32/libmlecc_tls11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmlrsa_tls11_r.so</i><sup>3</sup></li> </ul>
Security option for version 8 and 9 clients (mlsrv11 -xo) <sup>4</sup>	<ul style="list-style-type: none"> <li>● <i>lib32/libmlhttps11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmljrsa11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmljtls11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmlrsa11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libmltls11_r.so</i><sup>3</sup></li> </ul>
Setup scripts (deploy the ones for your consolidated database)	<ul style="list-style-type: none"> <li>● <i>MobiLink/setup</i></li> <li>● <i>MobiLink/upgrade</i></li> </ul>
mluser utility	<ul style="list-style-type: none"> <li>● <i>bin32/mluser</i></li> <li>● <i>lib32/libmlodbc11_r.so</i><sup>3</sup></li> <li>● <i>lib32/libdbicu11.so</i><sup>3</sup></li> <li>● <i>lib32/libdbicudt11.so</i><sup>3</sup></li> </ul>
mlstop utility	<ul style="list-style-type: none"> <li>● <i>bin32/mlstop</i></li> <li>● <i>lib32/libdbicu11.so</i><sup>3</sup></li> </ul>
MobiLink Monitor	<ul style="list-style-type: none"> <li>● <i>bin32/mlmon</i></li> <li>● <i>java/mlmon.jar</i></li> <li>● <i>java/mlstream.jar</i></li> <li>● <i>lib32/libjsyblib600_r.so</i><sup>3</sup></li> <li>● <i>sun/JavaHelp-2_0/jh.jar</i></li> <li>● <i>sun/jaxb1.0/</i></li> <li>● <i>java/JComponents1100.jar</i></li> <li>● <i>java/jsyblib600.jar</i></li> </ul>

Description	Unix files
MobiLink Redirector	<ul style="list-style-type: none"> <li>● <i>Mobilink/redirector/redirector.config</i></li> <li>● <i>MobiLink/redirector/apache/</i></li> <li>● <i>MobiLink/redirector/java/</i></li> <li>● <i>MobiLink/redirector/MBusinessAnywhere/</i></li> <li>● <i>MobiLink/redirector/nsapi/</i></li> </ul>
Online help for the MobiLink plug-in and MobiLink Monitor	<ul style="list-style-type: none"> <li>● <i>java/sqlanywhere_en11.jar</i><sup>1</sup></li> </ul>
Notifier	<ul style="list-style-type: none"> <li>● <i>java/activation.jar</i><sup>2</sup></li> <li>● <i>java/jodbc.jar</i></li> <li>● <i>java/log4j.jar</i><sup>2</sup></li> <li>● <i>java/mailapi.jar</i><sup>2</sup></li> <li>● <i>java/mlnotif.jar</i></li> <li>● <i>java/mlscript.jar</i></li> <li>● <i>java/smtplib.jar</i><sup>2</sup></li> <li>● <i>lib32/libmljstrm11_r.so</i><sup>3</sup></li> </ul>
MobiLink server files required by QAnywhere	<ul style="list-style-type: none"> <li>● Notifier files</li> <li>● <i>java/commons-codec-1.3.jar</i></li> <li>● <i>java/commons-httpclient-3.0.jar</i></li> <li>● <i>java/commons-logging.jar</i></li> <li>● <i>java/jsyblib600.jar</i></li> <li>● <i>java/log4j.jar</i><sup>5</sup></li> <li>● <i>java/mlscript.jar</i></li> <li>● <i>java/mlstream.jar</i></li> <li>● <i>java/qaconnector.jar</i></li> <li>● <i>lib32/libjsyblib600_r.so</i><sup>3</sup></li> </ul>
Relay Server Outbound Enabler	<ul style="list-style-type: none"> <li>● <i>bin32/rsoe</i></li> </ul> <p>For security with the Outbound Enabler:</p> <ul style="list-style-type: none"> <li>● <i>lib32/libmlcrsa11_r.so</i><sup>3</sup></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

<sup>2</sup> If you are redistributing an application, you must obtain these files directly from Sun.

<sup>3</sup> For Linux, the file extension is *.so*. For Macintosh, the file extension is *.dylib*.

<sup>4</sup> Transport-layer security requires that you obtain the separately-licensed SQL Anywhere security option and is subject to export regulations. To order this component, see [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

<sup>5</sup> If you are redistributing an application, you must obtain these files directly from Apache.

### Unix 64-bit applications on Unix and Linux

All directories are relative to *install-dir*. For more details on the file structure of a 32-bit Unix environment, see [“Unix 32-bit applications on Unix, Linux, and Macintosh” on page 807](#).

Description	Unix files
MobiLink server	<ul style="list-style-type: none"> <li>● <i>bin64/mlsrv11</i></li> <li>● <i>bin64/mlsrv11.lic</i></li> <li>● <i>lib64/libdbodm11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmlodbc11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmlsql11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libdbtasks11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libdbicu11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libdbicudt11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libdbodbcinst11_r.so<sup>3</sup></i></li> </ul>
Language library	<ul style="list-style-type: none"> <li>● <i>res/dblgen11.res<sup>1</sup></i></li> </ul>
Synchronization stream libraries for version 8 and 9 clients (deploy the ones you use)	<ul style="list-style-type: none"> <li>● <i>lib64/libmlhttp11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmlsock11_r.so<sup>3</sup></i></li> </ul>
Java synchronization logic	<ul style="list-style-type: none"> <li>● <i>java/activation.jar<sup>2</sup></i></li> <li>● <i>java/imap.jar<sup>2</sup></i></li> <li>● <i>java/jodbc.jar</i></li> <li>● <i>java/log4j.jar<sup>5</sup></i></li> <li>● <i>java/mailapi.jar<sup>2</sup></i></li> <li>● <i>java/mlscript.jar</i></li> <li>● <i>java/mlsupport.jar</i></li> <li>● <i>java/pop3.jar<sup>2</sup></i></li> <li>● <i>java/smtp.jar<sup>2</sup></i></li> <li>● <i>lib64/libmljava11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmljodbc11.so<sup>3</sup></i></li> </ul>
.NET synchronization logic	<ul style="list-style-type: none"> <li>● Not applicable</li> </ul>
Security option for version 10 and 11 clients (mlsrv11 -x) <sup>4</sup>	<ul style="list-style-type: none"> <li>● <i>lib64/libmlecc_tls11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmlrsa_tls11_r.so<sup>3</sup></i></li> </ul>
Security option for version 8 and 9 clients (mlsrv11 -xo) <sup>4</sup>	<ul style="list-style-type: none"> <li>● <i>lib64/libmlhttps11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmljrsa11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmljtls11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmlrsa11_r.so<sup>3</sup></i></li> <li>● <i>lib64/libmltls11_r.so<sup>3</sup></i></li> </ul>
Setup scripts (deploy the ones for your consolidated database)	<ul style="list-style-type: none"> <li>● <i>MobiLink/setup</i></li> <li>● <i>MobiLink/upgrade</i></li> </ul>

Description	Unix files
mluser utility	<ul style="list-style-type: none"> <li>• <i>bin64/mluser</i></li> <li>• <i>lib64/libmlodbc11_r.so<sup>3</sup></i></li> <li>• <i>lib64/libdbicu11.so<sup>3</sup></i></li> <li>• <i>lib64/libdbicudt11.so<sup>3</sup></i></li> </ul>
mlstop utility	<ul style="list-style-type: none"> <li>• <i>bin64/mlstop</i></li> <li>• <i>lib64/libdbicu11.so<sup>3</sup></i></li> </ul>
MobiLink Monitor	<ul style="list-style-type: none"> <li>• <i>bin64/mlmon</i></li> <li>• <i>java/mlmon.jar</i></li> <li>• <i>java/mlstream.jar</i></li> <li>• <i>lib64/libjsyblib600_r.so<sup>3</sup></i></li> <li>• <i>sun/JavaHelp-2_0/jh.jar</i></li> <li>• <i>sun/jaxb1.0</i></li> <li>• <i>java/JComponents1100.jar</i></li> <li>• <i>java/jsyblib600.jar</i></li> </ul>
MobiLink Redirector	<ul style="list-style-type: none"> <li>• <i>Mobilink/redirector/redirector.config</i></li> <li>• <i>MobiLink/redirector/apache/</i></li> <li>• <i>MobiLink/redirector/java/</i></li> <li>• <i>MobiLink/redirector/MBusinessAnywhere/</i></li> <li>• <i>MobiLink/redirector/nsapi/</i></li> </ul>
Online help for the MobiLink plug-in and MobiLink Monitor	<ul style="list-style-type: none"> <li>• <i>java/sqlanywhere_en11.jar<sup>1</sup></i></li> </ul>
Notifier	<ul style="list-style-type: none"> <li>• <i>java/activation.jar<sup>2</sup></i></li> <li>• <i>java/jodbc.jar</i></li> <li>• <i>java/log4j.jar<sup>2</sup></i></li> <li>• <i>java/mailapi.jar<sup>2</sup></i></li> <li>• <i>java/mlnotif.jar</i></li> <li>• <i>java/mlscript.jar</i></li> <li>• <i>java/sntp.jar<sup>2</sup></i></li> </ul>
MobiLink server files required by QAnywhere	<ul style="list-style-type: none"> <li>• Notifier files</li> <li>• <i>java/commons-codec-1.3.jar</i></li> <li>• <i>java/commons-httpclient-3.0.jar</i></li> <li>• <i>java/commons-logging.jar</i></li> <li>• <i>java/jsyblib600.jar</i></li> <li>• <i>java/log4j.jar<sup>5</sup></i></li> <li>• <i>java/mlscript.jar</i></li> <li>• <i>java/mlstream.jar</i></li> <li>• <i>java/qaconnector.jar</i></li> <li>• <i>lib64/libjsyblib600_r.so<sup>3</sup></i></li> </ul>

Description	Unix files
Relay Server Outbound Enabler	<ul style="list-style-type: none"><li>● <i>bin64/rsoe</i></li></ul> For security with the Outbound Enabler: <ul style="list-style-type: none"><li>● <i>lib64/libmlcrsa11_r.so</i><sup>3</sup></li></ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

<sup>2</sup> If you are redistributing an application, you must obtain these files directly from Sun.

<sup>3</sup> For Solaris SPARC and Linux, the file extension is *.so*. For AIX, the file extension is *.a*.

<sup>4</sup> Transport-layer security requires that you obtain the separately-licensed SQL Anywhere security option and is subject to export regulations. To order this component, see [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

<sup>5</sup> If you are redistributing an application, you must obtain these files directly from Apache.



# Deploying SQL Anywhere MobiLink clients

## Notes

- For SQL Anywhere clients, you need to deploy a SQL Anywhere database server and the MobiLink client.  
See “[Deploying databases and applications](#)” [[SQL Anywhere Server - Programming](#)].
- If you are redistributing MobiLink synchronization clients you need to include the following files in your installation, in addition to those required for the SQL Anywhere database.
- When deploying the files below, place them in the same directory structure unless otherwise noted.
- To deploy Sybase Central, see “[Deploying administration tools](#)” [[SQL Anywhere Server - Programming](#)].
- There is a deployment wizard for Windows. See “[Using the Deployment Wizard](#)” [[SQL Anywhere Server - Programming](#)].
- For Windows Mobile deployment, files that are listed below in the *bin32* directories are located in the *ce\arm.50* directory. .NET assemblies are placed in the *ce\Assembly\v2* directory.

## Windows applications

All directories are relative to *install-dir*.

Description	Windows files
MobiLink synchronization client (dbmlsync)	<ul style="list-style-type: none"> <li>• <i>bin32\dbcon11.dll</i><sup>2</sup></li> <li>• <i>bin32\dbcu11.dll</i><sup>3</sup></li> <li>• <i>bin32\dblgen11.dll</i><sup>1</sup></li> <li>• <i>bin32\dblib11.dll</i></li> <li>• <i>bin32\dbmlsync.exe</i></li> <li>• <i>bin32\dbtool11.dll</i><sup>2</sup></li> </ul>
Dbmlsync integration component (deprecated)	<ul style="list-style-type: none"> <li>• MobiLink synchronization client files</li> <li>• Visual component: <i>bin32\dbmlsynccomg.dll</i></li> <li>• Non-visual component: <i>bin32\dbmlsynccom.dll</i></li> </ul>
Security option <sup>2</sup>	<ul style="list-style-type: none"> <li>• <i>bin32\mlcecc11.dll</i></li> <li>• <i>bin32\mlcrsa11.dll</i></li> <li>• <i>bin32\mlcrsafips11.dll</i></li> <li>• <i>bin32\sbgse2.dll</i></li> </ul>
ActiveSync and HotSync utilities	<ul style="list-style-type: none"> <li>• <i>bin32\mlasinst.exe</i></li> <li>• <i>bin32\mlasdesk.dll</i></li> <li>• <i>bin32\dbcon11.exe</i></li> <li>• <i>ce\chip\mlasdev.dll</i> (where <i>chip</i> can be any supported chip, such as arm.50)</li> </ul>

Description	Windows files
Listener	<ul style="list-style-type: none"> <li>• <i>bin32\dblgen11.dll</i><sup>1</sup></li> <li>• <i>bin32\dblsn.exe</i></li> <li>• <i>bin32\lsn_udp.dll</i></li> <li>• <i>bin32\lsn_swi510.dll</i></li> <li>• <i>bin32\maac555.dll</i></li> <li>• <i>bin32\maac750.dll</i></li> <li>• <i>bin32\maac750r3.dll</i></li> <li>• <i>bin32\mabridge.dll</i></li> </ul>

<sup>1</sup> For German, Japanese, and Chinese editions, substitute en with de, ja, and zh, respectively.

<sup>2</sup> Not required on Windows Mobile unless you use the dbtools interface.

<sup>3</sup> Not required if the database is initialized with dbinit -zn UTF8BIN. See [“Initialization utility \(dbinit\)” \[SQL Anywhere Server - Database Administration\]](#).

### Unix applications on Unix, Linux, and Macintosh

All directories are relative to *install-dir*.

Description	Unix files
MobiLink synchronization client	<ul style="list-style-type: none"> <li>• <i>bin32/dbmlsync</i></li> <li>• <i>res/dblgen11.res</i></li> <li>• <i>lib32/libdbcon11_r.so</i><sup>1</sup></li> <li>• <i>lib32/libdbicu11_r.so</i><sup>1</sup></li> <li>• <i>lib32/libdblib11_r.so</i><sup>1</sup></li> <li>• <i>lib32/libdbtool11_r.so</i><sup>1</sup></li> </ul>
Security option <sup>2</sup>	<ul style="list-style-type: none"> <li>• <i>lib32/libmlcecc11_r.so</i><sup>1</sup></li> <li>• <i>lib32/libmlcrsa11_r.so</i><sup>1</sup></li> </ul>

<sup>1</sup>For Linux, the file extension is *.so*. For the Macintosh, the file extension is *.dylib*.

<sup>2</sup> Transport-layer security requires that you obtain the separately-licensed SQL Anywhere security option and is subject to export regulations. To order this component, see [“Separately licensed components” \[SQL Anywhere 11 - Introduction\]](#).

## Deploying UltraLite MobiLink clients

For UltraLite clients, the UltraLite runtime library or the UltraLite component includes the required synchronization stream functions. The UltraLite runtime library is compiled into your application. Deployment is subject to your license agreement.

### See also

- “Deploying UltraLite to devices” [*UltraLite - Database Management and Reference*]
- C/C++: “Deploying Palm applications” [*UltraLite - C and C++ Programming*] and “Deploying Windows Mobile applications” [*UltraLite - C and C++ Programming*]
- M-Business Anywhere: “Deploying UltraLite for M-Business Anywhere applications” [*UltraLite - M-Business Anywhere Programming*]
- .NET: “Lesson 5: Build and deploy application” [*UltraLite - .NET Programming*]

## Deploying QAnywhere applications

QAnywhere provides C++, Java, and .NET API support for SQL Anywhere message stores. The Java and .NET APIs also support UltraLite message stores. The files required for deploying QAnywhere applications are based on your Windows environment, message store type, and API selection. Additional files are required if you are developing Mobile Web Service applications.

In addition to the files listed below, a QAnywhere application requires:

- All files listed in the MobiLink synchronization client, Listener, and optionally the Security sections of [“Deploying SQL Anywhere MobiLink clients” on page 813](#). The Listener files are required only if you are using push notifications, which is the default.
- dbeng11 or dbsrv11 files from [“Deploying database servers” \[SQL Anywhere Server - Programming\]](#).

To deploy Sybase Central, see [“Deploying administration tools” \[SQL Anywhere Server - Programming\]](#).

### Windows applications

All directories are relative to *install-dir*.

For more details on the file structure of a Windows Mobile environment, see [“Windows Mobile applications” on page 818](#).

The following is a list of files required to set up a SQL Anywhere message store.

Client API	Windows files
C++	<ul style="list-style-type: none"> <li>• <i>bin32\qany11.dll</i></li> <li>• <i>bin32\qaagent.exe</i></li> <li>• <i>bin32\qastop.exe</i></li> </ul>
Java	<ul style="list-style-type: none"> <li>• <i>bin32\qaagent.exe</i></li> <li>• <i>bin32\qastop.exe</i></li> <li>• <i>java\qaclient.jar</i></li> <li>• <i>java\jodbc.jar</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>• <i>java\iawsrt.jar</i></li> <li>• <i>java\jaxrpc.jar</i></li> </ul>

Client API	Windows files
.NET	<ul style="list-style-type: none"> <li>● <i>bin32\qazlib.dll</i></li> <li>● <i>bin32\qaagent.exe</i></li> <li>● <i>bin32\qastop.exe</i></li> <li>● <i>assembly\v2\iAnywhere.QAnywhere.Client.dll</i></li> <li>● <i>assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>● <i>assembly\v2\iAnywhere.Data.SQLAnywhere.dll</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>● <i>Assembly\v2\iAnywhere.QAnywhere.WS.dll</i></li> </ul>

The following is a list of files required to set up an UltraLite message store with deployments using QAnywhere Agent.

Client API	Windows files
Java	<ul style="list-style-type: none"> <li>● <i>bin32\qauagent.exe</i></li> <li>● <i>bin32\qastop.exe</i></li> <li>● <i>bin32\qadbiuljni.dll</i></li> <li>● <i>java\qaclient.jar</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>● <i>java\iawsrt.jar</i></li> <li>● <i>java\jaxrpc.jar</i></li> </ul>
.NET	<ul style="list-style-type: none"> <li>● <i>bin32\qazlib.dll</i></li> <li>● <i>bin32\qauagent.exe</i></li> <li>● <i>bin32\qastop.exe</i></li> <li>● <i>assembly\v2\iAnywhere.QAnywhere.Client.dll</i></li> <li>● <i>assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>● <i>ultralite\ultralite.NET\assembly\v2\iAnywhere.Data.UltraLite.dll</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>● <i>Assembly\v2\iAnywhere.QAnywhere.WS.dll</i></li> </ul>

When creating an UltraLite message store, you must create a udb database file using the UltraLite Create Database utility, then initialize the database using the QAnywhere UltraLite Agent's `-si` option. See [“UltraLite Create Database utility \(ulcreate\)”](#) [*UltraLite - Database Management and Reference*] and [“qauagent utility”](#) [*QAnywhere*].

The following is a list of files required to set up a deployment with the QAnywhere standalone client.

Client API	Windows files
Java	<ul style="list-style-type: none"> <li>• <i>java\qastandaloneclient.jar</i></li> <li>• <i>bin32\qadbiulsjni.dll</i></li> </ul>
.NET	<ul style="list-style-type: none"> <li>• <i>assembly\v2\iAnywhere.QAnywhere.StandAloneClient.dll</i></li> <li>• <i>assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>• <i>ultralite\ultralite.NET\assembly\v2\iAnywhere.Data.UltraLite.dll</i></li> </ul>

### Windows Mobile applications

All directories are relative to *install-dir*.

For more details on the file structure of a Windows environment, see [“Windows applications” on page 816](#).

The following is a list of files required to set up a SQL Anywhere message store.

Client API	Windows Mobile files
C++	<ul style="list-style-type: none"> <li>• <i>ce\arm.50\qany11.dll</i></li> <li>• <i>ce\arm.50\qaagent.exe</i></li> <li>• <i>ce\arm.50\qastop.exe</i></li> </ul>
Java	<ul style="list-style-type: none"> <li>• <i>ce\arm.50\qaagent.exe</i></li> <li>• <i>ce\arm.50\qastop.exe</i></li> <li>• <i>java\qaclient.jar</i></li> <li>• <i>java\jodbc.jar</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>• <i>java\iawsrt.jar</i></li> <li>• <i>java\jaxrpc.jar</i></li> </ul>
.NET	<ul style="list-style-type: none"> <li>• <i>ce\arm.50\qazlib.dll</i></li> <li>• <i>ce\arm.50\qaagent.exe</i></li> <li>• <i>ce\arm.50\qastop.exe</i></li> <li>• <i>ce\assembly\v2\iAnywhere.QAnywhere.Client.dll</i></li> <li>• <i>ce\assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>• <i>ce\assembly\v2\iAnywhere.Data.SQLAnywhere.dll</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>• <i>ce\Assembly\v2\iAnywhere.QAnywhere.WS.dll</i></li> </ul>

The following is a list of files required to set up an UltraLite message store with deployments using QAnywhere Agent.

Client API	Windows Mobile files
Java	<ul style="list-style-type: none"> <li>● <i>ce\arm.50\qauagent.exe</i></li> <li>● <i>ce\arm.50\qastop.exe</i></li> <li>● <i>ce\arm.50\qadbiuljni.dll</i></li> <li>● <i>java\qaclient.jar</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>● <i>java\iawsrt.jar</i></li> <li>● <i>java\jaxrpc.jar</i></li> </ul>
.NET	<ul style="list-style-type: none"> <li>● <i>ce\arm.50\qazlib.dll</i></li> <li>● <i>ce\arm.50\qauagent.exe</i></li> <li>● <i>ce\arm.50\qastop.exe</i></li> <li>● <i>ce\assembly\v2\iAnywhere.QAnywhere.Client.dll</i></li> <li>● <i>ce\assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>● <i>ultralite\ultralite.NET\ce\assembly\v2\iAnywhere.Data.UltraLite.dll</i></li> </ul> <p>For Mobile Web Service applications, you also need the following:</p> <ul style="list-style-type: none"> <li>● <i>ce\Assembly\v2\iAnywhere.QAnywhere.WS.dll</i></li> </ul>

When creating an UltraLite message store, you must create a database file using the UltraLite Create Database utility, then initialize the database using the `-si` option for the QAnywhere UltraLite Agent. See [“UltraLite Create Database utility \(ulcreate\)” \[UltraLite - Database Management and Reference\]](#) and [“qauagent utility” \[QAnywhere\]](#).

The following is a list of files required to set up a deployment with the QAnywhere standalone client.

Client API	Windows Mobile files
Java	<ul style="list-style-type: none"> <li>● <i>java\qastandaloneclient.jar</i></li> <li>● <i>ce\arm.50\qadbiulsjni.dll</i></li> </ul>
.NET	<ul style="list-style-type: none"> <li>● <i>ce\assembly\v2\iAnywhere.QAnywhere.StandAloneClient.dll</i></li> <li>● <i>ce\assembly\v2\iAnywhere.QAnywhere.Resources.dll</i></li> <li>● <i>ultralite\ultralite.NET\ce\assembly\v2\iAnywhere.Data.UltraLite.dll</i></li> </ul>

### Registering the QAnywhere .NET API DLL

The QAnywhere .NET API DLL (*Assembly\v2\iAnywhere.QAnywhere.Client.dll*) needs to be registered in the Global Assembly Cache on Windows (except on Windows Mobile). The Global Assembly Cache lists all the registered programs on your computer. When you install SQL Anywhere, the installation program registers it. In Windows Mobile you do not need to register the DLL.

If you are deploying QAnywhere, you must register the QAnywhere .NET API DLL (*Assembly\v2\iAnywhere.QAnywhere.Client.dll*) using the `gacutil` utility that is included with the .NET Framework.

---



# Glossary

---

Glossary ..... 823



---

# 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 847](#).

## **agent ID**

See also: [“client message store ID” on page 825](#).

## **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 845](#)
- [“publication” on page 842](#)

## **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 849](#)
- [“view” on page 851](#)

### **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 827](#)
- [“user-defined data type” on page 851](#)

### **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 846](#).

### **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 825](#)
- [“encoding” on page 831](#)
- [“collation” on page 825](#)

### **check constraint**

A restriction that enforces specified conditions on a column or set of columns.

See also:

- [“constraint” on page 827](#)
- [“foreign key constraint” on page 832](#)
- [“primary key constraint” on page 842](#)
- [“unique constraint” on page 850](#)

### **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 824](#)
- [“encoding” on page 831](#)
- [“collation” on page 825](#)

## 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 824](#)
- [“code page” on page 825](#)
- [“encoding” on page 831](#)

## 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 849](#)
- [“isolation level” on page 835](#)

### **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 846](#).

### **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 849](#)
- [“replication” on page 845](#)

---

## **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 824](#)
- [“foreign key constraint” on page 832](#)
- [“primary key constraint” on page 842](#)
- [“unique constraint” on page 850](#)

## **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 827](#)
- [“cursor position” on page 827](#)

## **cursor position**

A pointer to one row within the cursor result set.

See also:

- [“cursor” on page 827](#)
- [“cursor result set” on page 827](#)

## **cursor result set**

The set of rows resulting from a query that is associated with a cursor.

See also:

- [“cursor” on page 827](#)
- [“cursor position” on page 827](#)

### **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 830](#).

### **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 832](#)
- [“primary key” on page 842](#)
- [“database management system \(DBMS\)” on page 829](#)
- [“relational database management system \(RDBMS\)” on page 844](#)

### **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 830](#).

## 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 844](#).

## 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 829](#).

## 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 828](#)
- [“SYS” on page 849](#)

## 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 828](#).

## **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 829](#).

## **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 846](#).

## **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 826](#)
- [“SQL-based synchronization” on page 847](#)

## **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 828](#).

## **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 824](#)
- [“code page” on page 825](#)
- [“collation” on page 825](#)

## **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 845](#).

## **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 845](#).

### **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 832](#).

### **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 842](#)
- [“foreign table” on page 832](#)

### **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 827](#)
- [“check constraint” on page 824](#)
- [“primary key constraint” on page 842](#)
- [“unique constraint” on page 850](#)

### **foreign table**

The table containing the foreign key.

See also: [“foreign key” on page 832](#).

### **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 834](#).

### **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 846](#).

---

## 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 836](#)
- [“join condition” on page 836](#)

## 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 832](#).

## 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 849](#)
- [“local temporary table” on page 836](#)

## 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 835](#)
- [“jConnect” on page 835](#)

### **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 849](#).

### **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 836](#)
- [“outer join” on page 840](#)

### **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 844](#).

## 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 847](#).

## 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 835](#)
- [“iAnywhere JDBC driver” on page 834](#)

## 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 835](#)
- [“iAnywhere JDBC driver” on page 834](#)

### **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 836](#)
- [“generated join condition” on page 833](#)

### **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 836](#).

### **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 846](#).

### **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 846](#).

### **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 849](#)
- [“global temporary table” on page 833](#)



---

## 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 835](#)
- [“concurrency” on page 826](#)
- [“integrity” on page 835](#)

## 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 849](#)
- [“transaction log mirror” on page 850](#)
- [“full backup” on page 832](#)

## 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 845](#).

## 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 823](#)
- [“view” on page 851](#)

### **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 825](#)
- [“server message store” on page 847](#)

### **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 845](#)
- [“FILE” on page 831](#)

### **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 845](#)
- [“consolidated database” on page 826](#)

### **metadata**

Data about data. Metadata describes the nature and content of other data.

See also: [“schema” on page 846](#).

### **mirror log**

See also: [“transaction log mirror” on page 850](#).

---

## **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 826](#)
- [“synchronization” on page 849](#)
- [“UltraLite” on page 850](#)

## **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 841](#).

## **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 846](#)
- [“Listener” on page 836](#)

## 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 848](#).

## 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 836](#)
- [“inner join” on page 834](#)

## 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 848](#).

## **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 846](#).

## **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 832](#).

### **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 827](#)
- [“check constraint” on page 824](#)
- [“foreign key constraint” on page 832](#)
- [“unique constraint” on page 850](#)
- [“integrity” on page 835](#)

### **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 838](#).

### **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 845](#)
- [“article” on page 823](#)
- [“publication update” on page 842](#)

### **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 845](#)
- [“publication” on page 842](#)

---

## **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 845](#).

## **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 843](#)
- [“server-initiated synchronization” on page 846](#)

## **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 846](#).

## **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 847](#).

## **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 832](#).

### **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 842](#).

### **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 842](#)
- [“foreign key” on page 832](#)

### **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 829](#).

### **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 849](#)
- [“consolidated database” on page 826](#)

### **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 829](#).



---

## 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 837](#).

## 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 845](#).

## 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 845](#)
- [“publication update” on page 842](#)

## 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 837](#).

## 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 832](#).

### **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 849](#).

### **row-level trigger**

A trigger that executes once for each row that is changed.

See also:

- [“trigger” on page 850](#)
- [“statement-level trigger” on page 848](#)

### **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 831](#).

### **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 835](#).

## 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 846](#)
- [“SQL” on page 847](#)
- [“database management system \(DBMS\)” on page 829](#)

### **statement-level trigger**

A trigger that executes after the entire triggering statement is completed.

See also:

- [“trigger” on page 850](#)
- [“row-level trigger” on page 846](#)

### **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 842](#)
- [“MobiLink user” on page 839](#)

### **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 839](#)
- [“SQL Remote” on page 847](#)

## **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 836](#)
- [“global temporary table” on page 833](#)

## **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 849](#).

### **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 846](#)
- [“statement-level trigger” on page 848](#)
- [“integrity” on page 835](#)

### **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 832](#)
- [“primary key constraint” on page 842](#)
- [“constraint” on page 827](#)

---

**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 830](#).

**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 851](#).

**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

- a option
  - MobiLink server (mlsrv11), 51
- b option
  - MobiLink server (mlsrv11), 52
- bn option
  - MobiLink server (mlsrv11), 53
- c option
  - MobiLink server (mlsrv11), 54
  - MobiLink user authentication (mluser), 690
- classic option
  - MobiLink server (mlsrv11) -sl java, 92
- classpath option
  - MobiLink server (mlsrv11) -sl java, 92
- clrConGC option
  - MobiLink server (mlsrv11) -sl dnet, 90
- clrFlavor option
  - MobiLink server (mlsrv11) -sl dnet, 90
- clrVersion option
  - MobiLink server (mlsrv11) -sl dnet, 90
- cm option
  - MobiLink server (mlsrv11), 55
- cn option
  - MobiLink server (mlsrv11), 56
- cp option
  - MobiLink server (mlsrv11) -sl java, 92
- cr option
  - MobiLink server (mlsrv11), 57
- cs option
  - MobiLink server (mlsrv11), 58
- ct option
  - MobiLink server (mlsrv11), 59
- d option
  - MobiLink server (mlsrv11) -sl java, 92
  - MobiLink user authentication (mluser), 690
- dl option
  - MobiLink server (mlsrv11), 60
  - MobiLink user authentication (mluser), 690
- DMLStartClasses
  - Java user-defined start classes, 536
  - MobiLink server (mlsrv11) -sl java, 92
- dr option
  - MobiLink server (mlsrv11), 61
- ds option
  - MobiLink server (mlsrv11), 62
- dsd option
  - MobiLink server (mlsrv11), 63
- dt option
  - MobiLink server (mlsrv11), 64
- e option
  - MobiLink server (mlsrv11), 65
- esu option
  - MobiLink server (mlsrv11), 66
- et option
  - MobiLink server (mlsrv11), 67
- f option
  - MobiLink server (mlsrv11), 68
  - MobiLink stop utility (mlstop), 689
  - MobiLink user authentication (mluser), 690
- fr option
  - MobiLink server (mlsrv11), 70
- ftr option
  - MobiLink server (mlsrv11), 71
- h option
  - MobiLink stop utility (mlstop), 689
- hotspot option
  - MobiLink server (mlsrv11) -sl java, 92
- jrepath option
  - MobiLink server (mlsrv11) -sl java, 92
- lsc option
  - MobiLink server (mlsrv11), 72
- m option
  - MobiLink server (mlsrv11), 73
  - QAnywhere starting MobiLink server (mlsrv11), 73
- MLAutoLoadPath option
  - about, 601
  - MobiLink server (mlsrv11) -sl dnet, 90
- MLDomConfigFile option
  - about, 601
  - MobiLink server (mlsrv11) -sl dnet, 90
- MLStartClasses
  - .NET user-defined start classes, 595
  - MobiLink server (mlsrv11) -sl dnet, 90
- nba option
  - MobiLink server (mlsrv11), 74
- nc option
  - MobiLink server (mlsrv11), 75
- notifier option
  - MobiLink server (mlsrv11), 76
- o option
  - MobiLink server (mlsrv11), 77

- MobiLink user authentication (mluser), 690
- on option
  - MobiLink server (mlsrv11), 78
- oq option
  - MobiLink server (mlsrv11), 79
- os option
  - MobiLink server (mlsrv11), 80
  - MobiLink user authentication (mluser), 690
- ot option
  - MobiLink server (mlsrv11), 81
  - MobiLink user authentication (mluser), 690
- p option
  - MobiLink user authentication (mluser), 690
- pc option
  - MobiLink user authentication (mluser), 690
- ppv option
  - MobiLink server (mlsrv11), 82
- q option
  - MobiLink server (mlsrv11), 86
  - MobiLink stop utility (mlstop), 689
- r option
  - MobiLink server (mlsrv11), 87
- rd option
  - MobiLink server (mlsrv11), 88
- s option
  - MobiLink server (mlsrv11), 89
- server option
  - MobiLink server (mlsrv11) -sl java, 92
- sl dnet option
  - MobiLink server (mlsrv11), 90
  - user-defined start classes, 595
  - using -MLAutoLoadPath, 592
  - using -MLDomConfigFile, 601
- sl java option
  - MobiLink server (mlsrv11), 92
  - user-defined start classes, 536
- sm option
  - MobiLink server (mlsrv11), 94
- t option
  - MobiLink stop utility (mlstop), 689
- tc option
  - MobiLink server (mlsrv11), 96
- tf option
  - MobiLink server (mlsrv11), 97
- tx option
  - MobiLink server (mlsrv11), 98
- u option
  - MobiLink user authentication (mluser), 690
- ud option
  - MobiLink server (mlsrv11), 99
- ui option
  - MobiLink server (mlsrv11), 100
- urc option
  - MobiLink performance benefits, 173
- ux option
  - MobiLink server (mlsrv11), 101
- v option
  - MobiLink [dbmlsync] performance, 172
  - MobiLink server (mlsrv11), 102
- v+ option
  - MobiLink server (mlsrv11), 102
- vc option
  - MobiLink server (mlsrv11), 102
- ve option
  - MobiLink server (mlsrv11), 102
- verbose option
  - MobiLink server (mlsrv11) -sl java, 92
- vf option
  - MobiLink server (mlsrv11), 102
- vh option
  - MobiLink server (mlsrv11), 102
- vm option
  - MobiLink server (mlsrv11), 102
- vn option
  - MobiLink server (mlsrv11), 102
- vp option
  - MobiLink server (mlsrv11), 102
- vr option
  - MobiLink server (mlsrv11), 102
- vs option
  - MobiLink server (mlsrv11), 102
- vt option
  - MobiLink server (mlsrv11), 102
- vu option
  - MobiLink server (mlsrv11), 102
- w option
  - MobiLink server (mlsrv11), 105
  - MobiLink stop utility (mlstop), 689
- wu option
  - MobiLink server (mlsrv11), 106
- x option
  - MobiLink server (mlsrv11), 107
  - MobiLink server (mlsrv11) -sl java, 92
- xo option
  - MobiLink server (mlsrv11), 113
- zp option

- MobiLink server (mlsrv11), 118
  - zs option
    - MobiLink server (mlsrv11), 119
    - shared server state, 119
  - zt option
    - MobiLink server (mlsrv11), 120
  - zu option
    - MobiLink server (mlsrv11), 121
  - zus option
    - MobiLink server (mlsrv11), 122
  - zw option
    - MobiLink server (mlsrv11), 123
  - zwd option
    - MobiLink server (mlsrv11), 124
  - zwe option
    - MobiLink server (mlsrv11), 125
  - .NET
    - MobiLink data types, 594
    - MobiLink object-based data flow, 649
    - MobiLink server API reference, 606
    - MobiLink synchronization scripts, 589
  - .NET classes
    - instantiation for .NET synchronization logic, 593
  - .NET CLR
    - MobiLink options, 90
  - .NET MobiLink server API (*see* MobiLink server API for .NET)
  - .NET synchronization example
    - MobiLink .NET synchronization logic, 604
  - .NET synchronization logic
    - .NET class instantiations, 593
    - DBCommand, 606
    - DBConnection interface, 608
    - DBConnectionContext, 609
    - DBParameter class, 612
    - DBParameterCollection class, 615
    - DBRowReader interface, 620
    - debugging, 599
    - deploying on 32-bit Unix, 807
    - deploying on 32-bit Windows, 801
    - deploying on 64-bit Unix, 810
    - deploying on 64-bit Windows, 804
    - LogCallback delegate, 627
    - LogMessage class, 627
    - MessageType enumeration, 627
    - methods, 595
    - MobiLink performance, 172
    - MobiLink server API, 606
    - sample, 604
    - ServerContext interface, 629
    - setup, 591
    - ShutdownCallback delegate, 633
    - SQLType enumeration, 633
    - supported languages, 590
    - SynchronizationException class, 640
  - .NET synchronization techniques
    - about, 600
  - @data option
    - MobiLink server (mlsrv11), 50
    - MobiLink stop utility (mlstop), 689
    - MobiLink user authentication (mluser), 690
  - @EmployeeID variable
    - using with MobiLink primary key pools, 143
- ## A
- a.
    - MobiLink named parameter prefix, 320
    - MobiLink user-defined parameter prefix, 323
  - active property
    - MobiLink Monitor synchronization statistics, 195
  - ActiveSync
    - MobiLink client deployment on Windows, 813
  - Adaptive Server Enterprise
    - begin\_connection\_autocommit event, 368
    - MobiLink consolidated database, 10
    - MobiLink data mapping, 740
    - MobiLink isolation levels, 165
    - MobiLink synchronization, 10
    - StaticCursorLongColBufLen, 10
    - using DDL in MobiLink, 368
  - add table script wizard
    - using, 327
  - Add( object value ) method [ML .NET]
    - DBParameterCollection class syntax, 616
  - addErrorListener method [ML Java]
    - ServerContext syntax, 567
  - addInfoListener method [ML Java]
    - ServerContext syntax, 566
  - adding
    - MobiLink .NET connection scripts, 668
    - MobiLink .NET table scripts, 669
    - MobiLink Java connection scripts, 671
    - MobiLink Java table scripts, 672
    - MobiLink properties, 677
    - MobiLink SQL connection scripts, 667

- MobiLink SQL table scripts, 680
  - Monitor users, 228
  - synchronization scripts with Sybase Central, 327
  - user names in MobiLink, 690
  - adding or deleting scripts
    - MobiLink, 327
  - adding script versions
    - MobiLink, 325
  - adding scripts
    - MobiLink about, 327
  - adding synchronization scripts
    - using stored procedures, 328
  - addShutdownListener method [ML Java]
    - ServerContext syntax, 568
  - addWarningListener method [ML Java]
    - ServerContext syntax, 568
  - admin user
    - Monitor about, 228
  - administrators
    - Monitor users, 228
  - AdventureWorks
    - synchronization issues, 21
  - agent IDs
    - glossary definition, 823
  - alerts
    - Monitor, 232
    - Monitor email notification, 233
    - Monitor suppressing, 235
  - antialiasing
    - MobiLink Monitor option, 188
  - Apache
    - configuring for the Apache Redirector, 287
    - configuring servlet Redirector for MobiLink, 284
  - Apache on Linux
    - deploying the Relay Server, 256
  - Apache Redirector
    - configuring, 287
  - Apache Tomcat
    - servlet Redirector, 284
  - Apache web servers
    - configuring the Apache Redirector, 287
  - APIs
    - MobiLink server API for .NET, 606
    - MobiLink server API for Java, 543
  - application pool
    - creating, 254
  - application servers
    - synchronizing with MobiLink, 649
  - applications
    - deploying MobiLink, 799
  - array size
    - Oracle driver option, 795
  - articles
    - glossary definition, 823
  - ASE
    - (*see also* Adaptive Server Enterprise)
  - assemblies
    - implementing in MobiLink, 601
    - locating in MobiLink .NET synchronization logic, 591
  - atomic transactions
    - glossary definition, 823
  - authenticate\_file\_transfer
    - connection event, 353
  - authenticate\_parameters
    - connection event, 355
  - authenticate\_user
    - connection event, 358
  - authenticate\_user property
    - MobiLink Monitor synchronization statistics, 195
  - authenticate\_user\_hashed
    - connection event, 363
  - authentication
    - MobiLink mluser utility, 690
  - authentication parameters
    - MobiLink, 323
    - MobiLink scripts, 320
  - authentication\_status synchronization parameter
    - about, 358
  - autoincrement methods
    - Oracle MobiLink consolidated databases, 25
  - automatic validation
    - MobiLink file-based download, 298
  - AvantGo (*see* M-Business Anywhere)
- ## B
- backend farm section
    - Relay Server configuration file, 244
  - backend server section
    - Relay Server configuration file, 245
  - base tables
    - glossary definition, 823
  - begin\_connection
    - connection event, 367
  - begin\_connection\_autocommit

- connection event, 368
- begin\_download
  - connection event, 369
  - table event, 371
- begin\_download\_deletes
  - table event, 374
- begin\_download\_rows
  - table event, 377
- begin\_publication
  - connection event, 380
- begin\_sync property
  - MobiLink Monitor synchronization statistics, 195
- begin\_synchronization
  - connection event, 383
  - table event, 385
- begin\_upload
  - connection event, 387
  - table event, 389
- begin\_upload\_deletes
  - table event, 391
- begin\_upload\_rows
  - table event, 394
- bi-directional synchronization
  - about, 138
  - required scripts, 326
- bit arrays
  - glossary definition, 824
- blackouts
  - about, 226
- BLOBs
  - downloaded from ASE, 10
- blocking download acknowledgement
  - about, 161
- bottlenecks
  - MobiLink performance, 174
- broadcast download
  - MobiLink file-based download, 293
- buffer\_size protocol option
  - MobiLink server (mlsrv11) -x option for HTTP, 109
  - MobiLink server (mlsrv11) -x option for HTTPS, 110
- bugs
  - providing feedback, xix
- business rules
  - glossary definition, 824

## C

- C# programming language
  - MobiLink .NET support, 590
  - MobiLink options, 90
  - MobiLink synchronization scripts, 589
- C++ programming language
  - MobiLink .NET support, 590
- carriers
  - glossary definition, 824
- central databases
  - MobiLink consolidated databases, 3
- changing the last download time
  - MobiLink, 131
- CHAR columns
  - ASE MobiLink consolidated databases, 10
  - DB2 MobiLink consolidated databases, 13
  - MobiLink issues, 8
  - MobiLink server (mlsrv11) -b option, 52
  - Oracle MobiLink consolidated databases, 25
  - SQL Server MobiLink consolidated databases, 20
- CHAR data type
  - MobiLink and other DBMSs, 8
- character set considerations
  - MobiLink, 790
- character set conversion
  - by ODBC drivers, 791
  - during MobiLink synchronization, 790
- character sets
  - glossary definition, 824
  - MobiLink synchronization, 790
- chart pane
  - MobiLink Monitor, 188
- CHECK constraints
  - glossary definition, 824
- checkpoints
  - glossary definition, 824
- checksums
  - glossary definition, 825
- class instances
  - Java synchronization logic, 531
  - MobiLink .NET synchronization logic, 593
- CLASSPATH environment variable
  - MobiLink Java synchronization logic, 529
- Clear method [ML .NET]
  - DBParameterCollection class syntax, 616
- client
  - connecting to the Relay Server farm, 262

- client event-hook procedures
  - (*see also* event hooks)
- client message store IDs
  - glossary definition, 825
- client message stores
  - glossary definition, 825
- client/server
  - glossary definition, 825
- Close method [ML .NET]
  - DBCommand syntax, 608
  - DBConnection syntax, 609
  - DBRowReader interface syntax, 621
- CLR
  - MobiLink options, 90
- code pages
  - glossary definition, 825
- collation sequences
  - MobiLink synchronization, 790
- collations
  - glossary definition, 825
- collisions
  - MobiLink conflict resolution, 146
- column sizes
  - ASE MobiLink consolidated databases, 10
- ColumnNames property [ML .NET]
  - DBRowReader interface syntax, 621
- ColumnTypes property [ML .NET]
  - DBRowReader interface syntax, 622
- command files
  - glossary definition, 825
- command line
  - starting mlsrv11, 45
- command line utilities
  - MobiLink stop utility (mlstop), 689
  - MobiLink synchronization, 687
  - MobiLink user authentication (mluser), 690
- command prompts
  - conventions, xvii
  - curly braces, xvii
  - environment variables, xvii
  - parentheses, xvii
  - quotes, xvii
- command shells
  - conventions, xvii
  - curly braces, xvii
  - environment variables, xvii
  - parentheses, xvii
  - quotes, xvii
- CommandText property [ML .NET]
  - DBCommand syntax, 608
- Commit method [ML .NET]
  - DBConnection syntax, 609
- common language runtime
  - MobiLink options, 90
- communication streams
  - glossary definition, 826
- communications
  - MobiLink mlsrv11 -c option, 54
  - MobiLink server -x option, 107
- complete event model
  - MobiLink, 343
  - MobiLink pseudocode, 346
- completed property
  - MobiLink Monitor synchronization statistics, 195
- composite keys
  - MobiLink unique primary keys, 139
- concurrency
  - glossary definition, 826
  - MobiLink performance, 170
- configuring
  - Apache web servers, 287
  - M-Business Anywhere, 289
  - Microsoft web servers, 282
  - MobiLink consolidated databases, 6
  - MobiLink Redirector about, 268
  - NSAPI web servers on Unix, 280
  - NSAPI web servers on Windows, 277
  - servlet Redirector for Apache web servers, 284
  - Tomcat, 284
- configuring an Apache Redirector for Apache web servers
  - about, 287
- configuring an ISAPI Redirector for Microsoft web servers
  - about, 282
- configuring an NSAPI Redirector for Netscape/Sun web servers
  - Unix, 280
  - Windows, 277
- configuring Redirector properties
  - Redirectors that don't support server groups, 275
  - Redirectors that support server groups, 273
- configuring the servlet Redirector
  - Apache web servers, 284
- conflict detection
  - MobiLink, 147

- 
- MobiLink statement-based uploads, 147
  - conflict resolution
    - forcing in MobiLink, 154
    - glossary definition, 826
    - MobiLink, 146
    - MobiLink conflict detection, 147
    - MobiLink default behavior, 146
    - MobiLink detection, 147
    - resolve\_conflict script, 149
    - upload\_update script, 151
  - conflicted\_deletes property
    - MobiLink Monitor synchronization statistics, 195
  - conflicted\_inserts property
    - MobiLink Monitor synchronization statistics, 195
  - conflicted\_updates property
    - MobiLink Monitor synchronization statistics, 195
  - conflicts
    - MobiLink, 146
    - MobiLink default behavior, 146
    - MobiLink detection, 147
    - MobiLink direct row handling, 655
    - MobiLink forced, 154
  - connecting
    - MobiLink client-server prior to version 10, 113
    - MobiLink mlsrv11 -c option, 54
    - MobiLink server -x option, 107
  - connection IDs
    - glossary definition, 826
  - connection parameters
    - MobiLink server -x option, 107
  - connection profiles
    - glossary definition, 826
  - connection properties
    - MobiLink server -x option, 107
  - connection scripts
    - about, 318
    - adding .NET scripts, 668
    - adding Java scripts, 671
    - adding SQL scripts, 667
    - adding with Sybase Central, 327
    - alphabetic list of MobiLink scripts, 342
    - defined, 318
    - deleting .NET scripts, 668
    - deleting Java scripts, 671
    - deleting SQL scripts, 667
    - ml\_global, 325
  - connection strings
    - MobiLink mlsrv11, 54
  - connection-initiated synchronization
    - glossary definition, 826
  - connection-level scripts
    - defined, 318
  - connection\_retries property
    - MobiLink Monitor synchronization statistics, 195
  - connections
    - MobiLink mlsrv11 -c option, 54
    - MobiLink server -x option, 107
  - consolidated databases
    - about, 3
    - adding synchronization scripts to, 327
    - ASE as MobiLink, 10
    - creating MobiLink, 6
    - databases other than SQL Anywhere, 8
    - DBMS dependencies, 8
    - glossary definition, 826
    - IBM DB2 LUW as MobiLink, 12
    - IBM DB2 mainframe as MobiLink, 15
    - mapping of data types in MobiLink, 739
    - MobiLink isolation levels, 165
    - MobiLink system tables, 694
    - MySQL as MobiLink, 22
    - Oracle as MobiLink, 25
    - relating tables to MobiLink remote tables, 5
    - SQL Anywhere as MobiLink, 28
    - SQL Server as MobiLink, 20
  - constraint errors (*see* conflicts)
  - constraints
    - glossary definition, 827
  - constructors
    - MobiLink .NET synchronization logic, 594
    - MobiLink Java synchronization logic, 532
  - Contains( object value ) method [ML .NET]
    - DBParameterCollection class syntax, 617
  - Contains( string parameterName ) method [ML .NET]
    - DBParameterCollection class syntax, 615
  - contd\_timeout protocol option
    - MobiLink Redirector, 269
  - contention
    - glossary definition, 827
    - MobiLink performance, 170
    - MobiLink performance explanation, 174
  - conventions
    - command prompts, xvii
    - command shells, xvii
    - documentation, xvi
    - file names in documentation, xvi

- conversion
  - character set by ODBC drivers, 791
- conversion between character sets
  - MobiLink synchronization, 790
- CopyTo(Array array, int index) method [ML .NET]
  - DBParameterCollection class syntax, 618
- correlation names
  - glossary definition, 827
- Count property [ML .NET]
  - DBParameterCollection class syntax, 619
- create connection script wizard
  - using, 327
- create script version wizard
  - using, 325
- create service wizard
  - MobiLink, 36
- create your java synchronization script
  - MobiLink Java synchronization logic example, 538
- CreateCommand method [ML .NET]
  - DBConnection syntax, 609
- creating
  - download file for MobiLink file-based download, 296
  - file-definition database, 295
  - MobiLink consolidated databases, 6
- creating consolidated databases
  - MobiLink about, 6
- creating databases
  - consolidated, 6
- creating download files
  - MobiLink file-based download, 296
- creating file-definition databases
  - MobiLink, 295
- creator ID
  - glossary definition, 827
- cursor positions
  - glossary definition, 827
- cursor result sets
  - glossary definition, 827
- cursor scripts
  - defined, 318
- cursors
  - glossary definition, 827
- custom validation
  - MobiLink file-based download, 300
- customizing your statistics
  - MobiLink Monitor, 194
- D**
  - daemon
    - running MobiLink as a, 35
  - data cube
    - glossary definition, 828
  - data entry
    - MobiLink, 155
  - data exchange (*see* synchronization)
  - data flow (MobiLink) (*see* direct row handling)
  - data inconsistency
    - MobiLink conflict-handling, 146
  - data manipulation language
    - glossary definition, 828
  - data mappings
    - about, 739
  - data source name
    - Oracle driver option, 795
  - data type mapping
    - MobiLink consolidated databases, 739
  - data types
    - ASE in MobiLink, 740
    - glossary definition, 828
    - IBM DB2 LUW in MobiLink, 749
    - IBM DB2 mainframe in MobiLink, 756
    - Microsoft SQL Server in MobiLink, 767
    - MobiLink .NET and SQL, 594
    - MobiLink consolidated database mappings, 739
    - MobiLink Java and SQL, 532
    - MySQL in MobiLink, 774
    - Oracle in MobiLink, 779
  - database administrator
    - glossary definition, 828
  - database connections
    - glossary definition, 829
    - MobiLink performance, 176
    - MobiLink performance setting maximum, 171
  - database files
    - glossary definition, 829
  - database names
    - glossary definition, 829
  - database objects
    - glossary definition, 829
  - database owner
    - glossary definition, 829
  - database schemas
    - relating consolidated tables to MobiLink remote tables, 5



---

- database servers
  - glossary definition, 829
- database worker threads
  - MobiLink, 174
  - MobiLink performance, 170
- databases
  - glossary definition, 828
  - MobiLink consolidated databases, 3
- daylight savings time
  - MobiLink, 132
- DB2
  - maximum identifier length in IBM, 664, 696
  - MobiLink data mapping for LUW, 749
  - MobiLink data mapping for Mainframe, 756
  - MobiLink isolation levels, 165
- DB2 LUW
  - MobiLink consolidated database, 12
- DB2 mainframe
  - maximum identifier length in IBM, 695
  - MobiLink consolidated database, 15
- DBA authority
  - glossary definition, 829
- DBCommand interface [ML .NET]
  - syntax, 606
- DBConnection interface [ML .NET]
  - syntax, 608
- DBConnectionContext
  - constructors, 594
- DBConnectionContext interface [ML .NET]
  - syntax, 609
- DBConnectionContext interface [ML Java]
  - syntax, 543
- dbmsync integration component (deprecated)
  - deploying on Windows, 813
- dbmsync utility
  - deploying, 813
  - deploying on Unix, 814
  - deploying on Windows, 813
- DBMS
  - glossary definition, 829
- DBParameter class [ML .NET]
  - syntax, 612
- DBParameterCollection class [ML .NET]
  - syntax, 615
- DBParameterCollection method [ML .NET]
  - DBParameterCollection class syntax, 615
- DBRowReader interface [ML .NET]
  - syntax, 620
- dbspaces
  - glossary definition, 830
- DbType property [ML .NET]
  - DBParameter syntax, 613
- DCX
  - about, xiv
- DDL
  - glossary definition, 828
- deadlocks
  - glossary definition, 830
- debugging
  - .NET synchronization logic, 599
  - MobiLink connections, 41
  - MobiLink server log, 33
  - MobiLink synchronization using Java classes, 533
- debugging .NET synchronization logic
  - about, 599
- debugging Java classes
  - MobiLink Java synchronization logic, 533
- default global autoincrement
  - MobiLink declaring, 141
- default isolation levels
  - MobiLink, 165
- deletes
  - MobiLink downloads, 335
  - stopping upload of for SQL Anywhere clients, 156
- deleting
  - MobiLink .NET connection scripts, 668
  - MobiLink .NET table scripts, 669
  - MobiLink Java connection scripts, 671
  - MobiLink Java table scripts, 672
  - MobiLink properties, 677
  - MobiLink SQL connection scripts, 667
  - MobiLink SQL table scripts, 680
  - rows in remote MobiLink databases, 335
- deleting all the rows in a table
  - MobiLink, 336
- deleting rows
  - MobiLink remote databases, 335
  - MobiLink techniques, 156
- deleting rows with the download\_delete\_cursor script
  - MobiLink, 335
- deleting scripts
  - MobiLink about, 327
- deploying
  - MobiLink applications, 799
  - MobiLink applications and databases, 799
  - MobiLink performance, 169

- MobiLink server, 801
- overview of MobiLink, 800
- QAnywhere applications, 816
- SQL Anywhere MobiLink clients, 813
- UltraLite MobiLink clients, 815
- deploying MobiLink applications
  - about, 799
- deploying QAnywhere clients
  - about, 816
- deploying remote databases
  - about, 799
- deploying SQL Anywhere MobiLink clients
  - about, 813
- deploying the MobiLink server
  - about, 801
- deploying UltraLite MobiLink clients
  - about, 815
- deployment (*see* deploying)
- deployment overview
  - MobiLink, 800
- details table pane
  - MobiLink Monitor, 184
- detecting conflicts
  - MobiLink, 147
  - MobiLink with upload\_fetch scripts, 147
  - MobiLink with upload\_update scripts, 148
- developer community
  - newsgroups, xiv
- development tips
  - Mobilink direct row handling, 652
  - MobiLink synchronization, 128
- device tracking
  - glossary definition, 830
- direct inserts of scripts
  - MobiLink, 328
- direct row handling
  - about, 649
  - development tips, 652
  - DownloadData interface [ML Java], 548
  - downloads, 660
  - DownloadTableData interface [ML Java], 550
  - glossary definition, 830
  - handle\_DownloadData connection event, 442
  - handle\_UploadData connection event, 454
  - quick start, 651
  - SendColumnNames, 652
  - UpdateResultSet interface, 578
  - UploadData interface [ML Java], 579
  - UploadedTableData interface [ML Java], 581
  - uploads, 654
- direct synchronization events
  - about, 651
- Direction property [ML .NET]
  - DBParameter class syntax, 613
- disjoint partitioning
  - defined, 135
  - MobiLink, 135
- distributable download
  - MobiLink file-based download, 293
- DML
  - glossary definition, 828
- DocCommentXchange (DCX)
  - about, xiv
- documentation
  - conventions, xvi
  - SQL Anywhere, xiv
- domain configuration files
  - MobiLink, 602
- domains
  - glossary definition, 830
- download acknowledgement
  - about, 161
  - MobiLink performance, 171
- download buffer
  - MobiLink performance, 171
- download events
  - MobiLink synchronization, 351
- download failure
  - MobiLink restartable downloads, 158
- download file
  - creating for MobiLink file-based download, 296
- download property
  - MobiLink Monitor synchronization statistics, 195
- download timestamp
  - about MobiLink, 130
  - MobiLink generation of, 131
- download transaction
  - MobiLink, 345
- download-only synchronization
  - about, 138
  - required scripts, 326
- download\_bytes property
  - MobiLink Monitor synchronization statistics, 195
- download\_cursor
  - about, 334
  - example using a stored procedure call, 162

- MobiLink disjoint partitioning, 135
    - partitioning child tables, 137
    - partitioning with overlaps, 136
    - performance, 173
    - table event, 396
    - timestamp-based synchronization, 130
      - using a stored procedure call, 162
      - writing scripts to download rows, 333
  - download\_delete\_cursor
    - about, 335
    - disjoint partitioning, 135
    - example using a stored procedure call, 162
    - partitioning child tables, 137
    - partitioning with overlaps, 136
    - performance, 173
    - table event, 400
    - timestamp-based synchronization, 129
      - using a stored procedure call, 162
      - writing scripts to download rows, 333
  - download\_deleted\_rows property
    - MobiLink Monitor synchronization statistics, 195
  - download\_errors property
    - MobiLink Monitor synchronization statistics, 195
  - download\_fetched\_rows property
    - MobiLink Monitor synchronization statistics, 195
  - download\_filtered\_rows property
    - MobiLink Monitor synchronization statistics, 195
  - download\_statistics
    - connection event, 403
    - table event, 406
  - download\_timestamp
    - MobiLink generation of, 131
  - download\_warnings property
    - MobiLink Monitor synchronization statistics, 195
  - DownloadData interface [ML .NET]
    - syntax, 622
  - DownloadData interface [ML Java]
    - syntax, 548
  - downloading a result set from a stored procedure call
    - synchronization techniques, 162
  - downloading data
    - file-based download in MobiLink, 293
  - downloading deletes
    - MobiLink download\_delete\_cursor scripts, 335
  - downloading rows
    - synchronization scripts, 333
  - downloads
    - file-based MobiLink, 293
      - glossary definition, 830
    - MobiLink failed downloads, 158
    - MobiLink performance, 173
    - MobiLink scripts to download rows, 333
    - MobiLink transaction, 345
      - timestamp-based, 129
  - DownloadTableData interface [ML .NET]
    - syntax, 623
  - DownloadTableData interface [ML Java]
    - syntax, 550
  - drivers
    - supported by MobiLink, 794
  - duration property
    - MobiLink Monitor synchronization statistics, 195
  - dynamic SQL
    - glossary definition, 830
- ## E
- EBFs
    - glossary definition, 830
  - ECC protocol option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
    - MobiLink server (mlsrv11) -x option for TCP/IP, 108
  - emailing
    - Monitor alert notification, 233
    - Monitor users, 229
  - embedded SQL
    - glossary definition, 831
  - empty strings
    - Oracle MobiLink consolidated databases, 25
    - Oracle not supported, 25
  - enabling Microsoft distributed transactions
    - Oracle driver option, 795
  - encoding
    - glossary definition, 831
  - encrypting passwords
    - Oracle driver option, 795
  - end\_connection
    - connection event, 409
  - end\_download
    - connection event, 411
    - table event, 414
  - end\_download\_deletes
    - table event, 417
  - end\_download\_rows

- table event, 420
- end\_publication
  - connection event, 423
- end\_sync property
  - MobiLink Monitor synchronization statistics, 195
- end\_synchronization
  - connection event, 426
  - table event, 428
- end\_upload
  - connection event, 431
  - table event, 433
- end\_upload\_deletes
  - table event, 436
- end\_upload\_rows
  - table event, 439
- ending
  - MobiLink server, 32
- enterprise databases
  - synchronizing with MobiLink, 649
- environment variables
  - command prompts, xvii
  - command shells, xvii
- ERROR [ML Java]
  - Java LogMessage interface, 563
- ERROR field [ML .NET]
  - MessageType enumeration syntax, 627
- error handling
  - during MobiLink synchronization, 338
- error logs
  - MobiLink server (mlsrv11), 65
- errors
  - handling during MobiLink synchronization, 338
  - MobiLink modify\_error\_message connection event, 460
  - recording, 338
- event model
  - glossary definition, 831
  - MobiLink pseudocode, 346
- events
  - about MobiLink, 313
  - about MobiLink events, 315
  - about MobiLink synchronization, 343
  - MobiLink, 342
  - MobiLink direct row handling, 651
- events during download
  - about, 351
  - writing scripts to download rows, 333
- events during upload
  - about, 349
  - writing scripts to upload rows, 330
- examples
  - MobiLink file-based download, 301
- Excel
  - synchronizing with MobiLink, 649
- ExecuteNonQuery method [ML .NET]
  - DBCommand syntax, 607
- ExecuteReader method [ML .NET]
  - DBCommand syntax, 607
- external logins
  - glossary definition, 831
- extraction
  - glossary definition, 831
- F**
- failed downloads
  - MobiLink, 158
  - synchronization techniques, 158
- failover
  - glossary definition, 831
  - MobiLink Redirector, 266
  - MobiLink server farm, 40
- feedback
  - documentation, xix
  - providing, xix
  - reporting an error, xix
  - requesting an update, xix
- FILE
  - glossary definition, 831
- FILE message type
  - glossary definition, 831
- file-based downloads
  - about, 293
  - examples, 301
  - glossary definition, 832
- file-definition database
  - about, 295
  - creating, 295
  - glossary definition, 832
- file\_authentication\_code
  - authenticate\_file\_transfer parameter, 353
- files
  - MobiLink file-based download, 293
- finding out more and requesting technical assistance
  - technical support, xix
- FIPS

- mlsrv11 using HTTPS, 110
  - MobiLink server -x option, 107
  - FIPS option
    - MobiLink server (mlsrv11), 69
    - MobiLink user authentication (mluser), 690
  - FIPS protocol option
    - mlsrv11 -x option using TCP/IP, 108
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
  - firewalls
    - configuring MobiLink clients, 269
    - configuring MobiLink server, 269
    - MobiLink server, 269
    - routing MobiLink requests, 266
  - forced conflicts
    - MobiLink, 154
  - forcing conflicts
    - MobiLink, 154
  - foreign key constraints
    - glossary definition, 832
  - foreign keys
    - glossary definition, 832
  - foreign tables
    - glossary definition, 832
  - fragmentation
    - (*see also* partitioning)
  - FTP
    - MobiLink file-based download, 293
  - full backups
    - glossary definition, 832
  - fundamental rules
    - MobiLink, 128
- G**
- gateways
    - glossary definition, 832
  - generated join conditions
    - glossary definition, 833
  - generation numbers
    - glossary definition, 833
    - MobiLink file-based download, 299
  - GetConnection method [ML .NET]
    - DBConnectionContext syntax, 610
  - getConnection method [ML Java]
    - DBConnectionContext syntax, 544
  - GetDeleteCommand method [ML .NET]
    - DownloadTableData interface syntax, 625
  - getDeletePreparedStatement method [ML Java]
    - DownloadTableData syntax, 552
  - GetDeletes method [ML .NET]
    - UploadedTableData interface syntax, 643
  - getDeletes method [ML Java]
    - UploadedTableData syntax, 582
  - GetDownloadData method [ML .NET]
    - DBConnectionContext syntax, 610
  - getDownloadData method [ML Java]
    - DBConnectionContext syntax, 544
  - GetDownloadTableByName method [ML .NET]
    - DownloadData interface syntax, 623
  - getDownloadTableByName method [ML Java]
    - DownloadData syntax, 549
  - GetDownloadTables method [ML .NET]
    - DownloadData interface syntax, 622
  - getDownloadTables method [ML Java]
    - DownloadData syntax, 550
  - GetEnumerator method [ML .NET]
    - DBParameterCollection class syntax, 618
  - GetInserts method [ML .NET]
    - UploadedTableData interface syntax, 644
  - getInserts method [ML Java]
    - UploadedTableData syntax, 583
  - GetLastDownloadTime method [ML .NET]
    - DownloadTableData interface syntax, 625
  - getLastDownloadTime method [ML Java]
    - DownloadTableData syntax, 556
  - getMetaData method [ML Java]
    - DownloadTableData syntax, 555
    - UploadedTableData syntax, 586
  - GetName method [ML .NET]
    - DownloadTableData interface syntax, 625
    - UploadedTableData interface syntax, 645
  - getName method [ML Java]
    - DownloadTableData syntax, 554
    - UploadedTableData syntax, 585
  - GetProperties method [ML .NET]
    - DBConnectionContext syntax, 611
  - getProperties method [ML .NET]
    - ServerContext interface syntax, 631
  - getProperties method [ML Java]
    - DBConnectionContext syntax, 545
    - ServerContext syntax, 569
  - getPropertiesByVersion method [ML .NET]
    - ServerContext interface syntax, 631
  - getPropertiesByVersion method [ML Java]
    - ServerContext syntax, 569

- getPropertySetNames method [ML .NET]
    - ServerContext interface syntax, 631
  - getPropertySetNames method [ML Java]
    - ServerContext syntax, 570
  - GetRemoteID method [ML .NET]
    - DBConnectionContext syntax, 611
  - getRemoteID method [ML Java]
    - DBConnectionContext syntax, 546
  - GetSchemaTable method [ML .NET]
    - DownloadTableData interface syntax, 626
    - UploadedTableData interface syntax, 646
  - GetServerContext method [ML .NET]
    - DBConnectionContext syntax, 611
  - getServerContext method [ML Java]
    - DBConnectionContext syntax, 546
  - GetStartClassInstances method [ML .NET]
    - ServerContext interface syntax, 629
  - getStartClassInstances method [ML Java]
    - ServerContext syntax, 571
  - getText method [ML Java]
    - LogMessage syntax, 565
  - getting help
    - technical support, xix
  - getType method [ML Java]
    - LogMessage syntax, 564
  - GetUpdates method [ML .NET]
    - UploadedTableData interface syntax, 646
  - getUpdates method [ML Java]
    - UploadedTableData syntax, 584
  - GetUploadedTableByName method [ML .NET]
    - UploadData interface syntax, 641
  - getUploadedTableByName method [ML Java]
    - UploadData syntax, 580
  - GetUploadedTables method [ML .NET]
    - UploadData interface syntax, 642
  - getUploadedTables method [ML Java]
    - UploadData syntax, 581
  - GetUpsertCommand method [ML .NET]
    - DownloadTableData interface syntax, 626
  - getUpsertPreparedStatement method [ML Java]
    - DownloadTableData syntax, 553
  - getUser method [ML Java]
    - LogMessage syntax, 565
  - getValue method [ML Java]
    - InOutInteger syntax, 558
    - InOutString syntax, 559
  - GetVersion method [ML .NET]
    - DBConnectionContext syntax, 612
  - getVersion method [ML Java]
    - DBConnectionContext syntax, 547
  - global
    - script versions in MobiLink, 325
  - global assembly cache
    - implementing in MobiLink, 601
  - global autoincrement
    - algorithm, 141
    - MobiLink declaring, 141
    - MobiLink unique primary keys, 140
    - setting global\_database\_id for MobiLink, 141
  - global script versions
    - MobiLink, 325
  - global temporary tables
    - glossary definition, 833
  - global\_database\_id option
    - setting in MobiLink, 141
  - glossary
    - list of SQL Anywhere terminology, 823
  - grant options
    - glossary definition, 833
  - graph pane
    - MobiLink Monitor, 186
  - GUIDs
    - (*see also* UUIDs)
- ## H
- handle\_DownloadData
    - connection event, 442
  - handle\_error
    - connection event, 446
    - synchronization scripts, 338
  - handle\_odbc\_error
    - connection event, 450
  - handle\_UploadData
    - connection event, 454
  - handling conflicts
    - MobiLink, 146
    - MobiLink direct row handling, 655
  - handling conflicts for direct uploads
    - MobiLink direct row handling, 655
  - handling deletes
    - MobiLink, 156
  - handling direct downloads
    - MobiLink direct row handling, 660
  - Handling direct uploads
    - MobiLink direct row handling, 654

---

handling errors  
     MobiLink server, 446

handling MobiLink server errors in Java  
     MobiLink Java synchronization logic, 534

handling MobiLink server errors with .NET  
     MobiLink .NET synchronization logic, 598

handling multiple errors in a single SQL statement  
     MobiLink, 339

hard shutdown  
     MobiLink stop utility (mlstop), 689

hash  
     glossary definition, 833

health and statistics  
     Monitor, 201

help  
     technical support, xix

high availability  
     MobiLink Redirector, 266

histograms  
     glossary definition, 833

hooks  
     (*see also* event hooks)

host protocol option  
     MobiLink Redirector, 269  
     MobiLink server (mlsrv11) -x option for HTTP, 109  
     MobiLink server (mlsrv11) -x option for HTTPS, 110  
     MobiLink server (mlsrv11) -x option for TCP/IP, 107  
     MobiLink server (mlsrv11) -x option for TLS over TCP/IP, 108

hosted Relay Server  
     about, 260  
     adding a server farm, 260  
     logging in, 260  
     subscribing, 260

HotSync  
     MobiLink client deployment on Windows, 813

how conflicts are detected  
     MobiLink, 147

how default values are chosen  
     MobiLink global autoincrement, 141

HTTP  
     mlsrv11 -x option, 109  
     MobiLink server -x option, 107

HTTP load balancer  
     Relay Server, 242

httpd.conf  
     Apache native Redirector, 287

HTTPS  
     mlsrv11 -x option, 110  
     MobiLink server -x option, 107

**I**

iAnywhere developer community  
     newsgroups, xix

iAnywhere JDBC driver  
     glossary definition, 834

iAnywhere Solutions ODBC drivers  
     support, 793

iAnywhere Solutions Oracle driver  
     about, 795

iaredirect.dll  
     configuring the ISAPI Redirector, 282  
     configuring the NSAPI Redirector on Unix, 280  
     configuring the NSAPI Redirector on Windows, 277

iaredirect.so  
     configuring the NSAPI Redirector on Unix, 280  
     configuring the NSAPI Redirector on Windows, 277

IBM DB2  
     DB2 LUW as MobiLink consolidated database, 12  
     maximum identifier length in, 664, 696  
     MobiLink data mapping for LUW, 749  
     MobiLink data mapping for Mainframe, 756

IBM DB2 LUW  
     MobiLink consolidated database, 12

IBM DB2 LUW consolidated database  
     MobiLink, 12

IBM DB2 mainframe  
     maximum identifier length in, 695  
     MobiLink consolidated database, 15

IBM DB2 mainframe consolidated database  
     MobiLink, 15

IBM DB2 mainframe system table name conversions  
     ml\_active\_rid, 696, 697  
     ml\_conn\_script, 696, 699  
     ml\_pt, 696, 707, 730  
     ml\_pt\_repair, 696, 708  
     ml\_pt\_script, 696, 709  
     ml\_pt\_status, 696, 711  
     ml\_script\_modified, 696, 730

icons

- used in this Help, xviii
  - identifiers
    - glossary definition, 834
    - maximum length in IBM DB2 LUW, 695
    - maximum length in IBM DB2 mainframe, 664, 695, 696
  - identity option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
  - identity protocol option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
  - identity\_password protocol option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
  - ignore protocol option
    - MobiLink server (mlsrv11) -x option for TCP/IP, 107
    - MobiLink server (mlsrv11) -x option for TLS over TCP/IP, 108
  - ignored\_deletes property
    - MobiLink Monitor synchronization statistics, 195
  - ignored\_inserts property
    - MobiLink Monitor synchronization statistics, 195
  - ignored\_updates property
    - MobiLink Monitor synchronization statistics, 195
  - IIS
    - configuring for ISAPI, 282
  - IIS on Windows
    - deploying the Relay Server, 253
    - Relay Server, performance tips, 255
  - inconsistency
    - MobiLink conflict-handling, 146
  - incremental backups
    - glossary definition, 834
  - indexes
    - glossary definition, 834
    - MobiLink performance, 173
  - IndexOf( object value ) method [ML .NET]
    - DBParameterCollection class syntax, 617
  - IndexOf( string parameterName ) method [ML .NET]
    - DBParameterCollection class syntax, 616
  - INFO [ML Java]
    - Java LogMessage interface, 564
  - INFO field [ML .NET]
    - MessageType enumeration syntax, 628
  - InfoMaker
    - glossary definition, 834
  - inner joins
    - glossary definition, 834
  - InOutInteger interface [ML Java]
    - syntax, 557
  - InOutString interface [ML Java]
    - syntax, 559
  - Insert( int index, object value ) method [ML .NET]
    - DBParameterCollection class syntax, 617
  - inserting
    - scripts in MobiLink, 328
  - install-dir
    - documentation usage, xvi
  - installing
    - Monitor on a separate computer, 236
  - integrated logins
    - glossary definition, 834
  - integrity
    - glossary definition, 835
  - Interactive SQL
    - glossary definition, 835
  - iPlanet
    - configuring for the NSAPI Redirector on Unix, 280
    - configuring for the NSAPI Redirector on Windows, 277
  - ISAPI Redirector
    - calling, 282
    - configuring, 282
  - IsFixedSize property [ML .NET]
    - DBParameterCollection class syntax, 619
  - IsNullable property [ML .NET]
    - DBParameter class syntax, 613
  - isolation levels
    - glossary definition, 835
    - MobiLink, 165
  - IsReadOnly property [ML .NET]
    - DBParameterCollection class syntax, 619
  - IsSynchronized property [ML .NET]
    - DBParameterCollection class syntax, 619
- ## J
- JAR files
    - glossary definition, 835
  - Java
    - MobiLink data types, 532
    - MobiLink object-based data flow, 649
    - MobiLink server API reference, 543
    - MobiLink synchronization scripts, 527



---

Java classes  
    glossary definition, 835  
    instantiation for Java synchronization logic, 531

Java MobiLink server API (*see* MobiLink server API for Java)

Java synchronization  
    MobiLink Java synchronization logic, 538

Java synchronization logic  
    deploying on 32-bit Unix, 807  
    deploying on 32-bit Windows, 801  
    deploying on 64-bit Unix, 810  
    deploying on 64-bit Windows, 804  
    Java class instantiations, 531  
    methods, 533  
    MobiLink performance, 172  
    MobiLink server API, 543  
    sample, 538  
    setup, 529  
    specifying in MobiLink server command line, 529

Java VM  
    MobiLink options, 92

Java vs. SQL synchronization logic  
    MobiLink performance, 172

Javadoc  
    MobiLink, 543

jConnect  
    glossary definition, 835

JDBC  
    glossary definition, 835

join conditions  
    glossary definition, 836

join types  
    glossary definition, 836

joins  
    glossary definition, 836

## K

keep partial download synchronization parameter  
    restartable downloads, 159

key joins  
    glossary definition, 833

key pools  
    MobiLink synchronization application, 143

killing  
    MobiLink server, 32

## L

language libraries  
    MobiLink server deployment on 32-bit Unix, 807  
    MobiLink server deployment on 32-bit Windows, 801  
    MobiLink server deployment on 64-bit Unix, 810  
    MobiLink server deployment on 64-bit Windows, 804

last download time  
    about MobiLink, 130

last download timestamp  
    about MobiLink, 130  
    MobiLink generation of, 131  
    modify\_last\_download\_timestamp connection event, 463  
    modify\_next\_last\_download\_timestamp connection event, 466

last modified column  
    MobiLink, 129

last\_download  
    MobiLink named parameter, 130  
    modify\_last\_download\_timestamp connection event, 463

last\_download\_timestamp  
    MobiLink generation of, 131  
    MobiLink named parameter, 130

last\_table\_download  
    MobiLink named parameter, 130  
    modify\_last\_download\_timestamp connection event, 463

limitations  
    Monitor, 203

Listener utility  
    MobiLink client deployment on Windows, 813

Listeners  
    glossary definition, 836

load balancer  
    HTTP, 242

load balancing  
    MobiLink Redirector, 266  
    MobiLink server farm, 40  
    Redirector example (for Redirectors that don't support server groups), 276  
    Redirector example (for Redirectors that support server groups), 274

loading assemblies  
    MobiLink .NET synchronization logic, 601

- local temporary tables
  - glossary definition, 836
- locks
  - glossary definition, 837
- log file viewer
  - MobiLink server logs, 34
- log files
  - glossary definition, 837
  - MobiLink server, 33
  - MobiLink server viewing, 34
- LOG\_LEVEL
  - Redirector property (for Redirectors that don't support server groups), 275
  - Redirector property (for Redirectors that support server groups), 273
- LogCallback delegate [ML .NET]
  - DBRowReader interface syntax, 627
- LogCallback ErrorListener event [ML .NET]
  - ServerContext interface syntax, 629
- LogCallback InfoListener event [ML .NET]
  - ServerContext interface syntax, 629
- LogCallback WarningListener event [ML .NET]
  - ServerContext interface, 630
- logging
  - MobiLink performance, 172
  - MobiLink server actions, 33
- logging MobiLink server actions
  - about, 33
- logical deletes
  - writing download\_delete\_cursor scripts, 335
- logical indexes
  - glossary definition, 837
- LogListener interface [ML Java]
  - syntax, 560
- LogMessage class [ML .NET]
  - syntax, 627
- LogMessage class [ML Java]
  - syntax, 561
- logs
  - (*see also* log files)
- LONG data type
  - Oracle synchronization, 784
- LTM
  - glossary definition, 837
- LUW
  - DB2 LUW as MobiLink consolidated database, 12

## M

- M-Business Anywhere
  - configuring for synchronization, 289
  - Redirector, 289
- M-Business Anywhere Redirector
  - configuring, 289
- magnus.conf
  - configuring for the NSAPI Redirector on Unix, 280
  - configuring for the NSAPI Redirector on Windows, 277
- Mainframe
  - DB2 as MobiLink consolidated database, 15
- maintaining unique primary keys
  - about, 139
  - composite keys, 139
  - global autoincrement, 140
  - primary key pools, 143
  - UUIDs, 139
- maintenance releases
  - glossary definition, 837
- MakeConnection method [ML .NET]
  - ServerContext interface syntax, 630
- makeConnection method [ML Java]
  - ServerContext syntax, 571
- Manage Anywhere
  - MobiLink file-based download, 293
- many-to-many relationships
  - partitioning, 136
  - synchronization, 136
- mapping
  - MobiLink consolidated database data types, 739
- marquee tool
  - MobiLink Monitor overview pane, 189
- materialized views
  - glossary definition, 837
- message log
  - glossary definition, 838
- message properties files
  - deprecated in version 10.0.0, 73
- message stores
  - glossary definition, 838
- message systems
  - glossary definition, 838
- message types
  - glossary definition, 838
- messageLogged method [ML Java]
  - LogListener syntax, 561

---

MessageType enumeration [ML .NET]  
     syntax, 627

messaging  
     MobiLink QAnywhere system tables, 694

metadata  
     glossary definition, 838

methods  
     MobiLink .NET synchronization logic, 595  
     MobiLink Java synchronization logic, 533

metrics  
     editing collection intervals, 223  
     Monitor, 216

Microsoft Distributed Transaction Coordinator  
     Oracle driver option, 795

Microsoft Excel  
     synchronizing with MobiLink, 649

Microsoft SQL Server  
     as MobiLink consolidated database, 20  
     MobiLink data mapping, 767  
     MobiLink isolation levels, 165

Microsoft SQL Server consolidated database  
     MobiLink, 20

mirror logs  
     glossary definition, 838

ML  
     Redirector property (for Redirectors that don't support server groups), 275  
     Redirector property (for Redirectors that support server groups), 273

ml\_active\_remote\_id  
     MobiLink system table, 697

ml\_active\_rid  
     IBM DB2 mainframe system table name conversions, 696, 697

ml\_add\_column system procedure  
     syntax, 666

ml\_add\_connection\_script system procedure  
     syntax, 667

ml\_add\_cs system procedure  
     IBM DB2 conversion, 664  
     IBM DB2 mainframe system procedure name conversion, 667

ml\_add\_dcs system procedure  
     IBM DB2 conversion, 664  
     IBM DB2 mainframe system procedure name conversion, 668

ml\_add\_dnet\_connection\_script system procedure  
     syntax, 668

ml\_add\_dnet\_table\_script system procedure  
     syntax, 669

ml\_add\_dts system procedure  
     IBM DB2 conversion, 664  
     IBM DB2 mainframe system procedure name conversion, 669

ml\_add\_java\_connection\_script system procedure  
     syntax, 671

ml\_add\_java\_table\_script system procedure  
     syntax, 672

ml\_add\_jcs system procedure  
     IBM DB2 conversion, 664  
     IBM DB2 mainframe system procedure name conversion, 671

ml\_add\_jts system procedure  
     IBM DB2 conversion, 664  
     IBM DB2 mainframe system procedure name conversion, 672

ml\_add\_lang\_connection\_script system procedure  
     syntax, 673

ml\_add\_lang\_connection\_script\_chk system procedure  
     syntax, 673

ml\_add\_lang\_table\_script system procedure  
     syntax, 673

ml\_add\_lang\_table\_script\_chk system procedure  
     syntax, 673

ml\_add\_lcs system procedure  
     IBM DB2 conversion, 664  
     syntax, 673

ml\_add\_lcs\_chk system procedure  
     IBM DB2 conversion, 664  
     syntax, 673

ml\_add\_lts system procedure  
     IBM DB2 conversion, 664  
     syntax, 673

ml\_add\_lts\_chk system procedure  
     IBM DB2 conversion, 664  
     syntax, 673

ml\_add\_passthrough system procedure  
     syntax, 673

ml\_add\_passthrough\_repair system procedure  
     syntax, 674

ml\_add\_passthrough\_script system procedure  
     syntax, 676

ml\_add\_property system procedure  
     syntax, 677

ml\_add\_pt system procedure  
     IBM DB2 conversion, 664

- IBM DB2 mainframe system procedure name conversion, 673
- ml\_add\_pt\_repair system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 674
- ml\_add\_pt\_script system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 676
- ml\_add\_table\_script system procedure
  - syntax, 680
- ml\_add\_ts system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 680
- ml\_add\_user system procedure
  - syntax, 681
- ML\_CLIENT\_TIMEOUT
  - Redirector property (for Redirectors that don't support server groups), 275
  - Redirector property (for Redirectors that support server groups), 273
- ml\_column
  - MobiLink system table, 698
- ml\_conn\_script
  - IBM DB2 mainframe system table name conversions, 696, 699
- ml\_connection\_script
  - MobiLink system table, 699
- ml\_database
  - MobiLink system table, 700
- ml\_del\_pt system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 681
- ml\_del\_pt\_repair system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 682
- ml\_del\_pt\_script system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 683
- ml\_del\_sstate system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 683
- ml\_del\_sstate\_b4 system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 684
- ml\_delete\_passthrough system procedure
  - syntax, 681
- ml\_delete\_passthrough\_repair system procedure
  - syntax, 682
- ml\_delete\_passthrough\_script system procedure
  - syntax, 683
- ml\_delete\_sync\_state system procedure
  - syntax, 683
- ml\_delete\_sync\_state\_before system procedure
  - syntax, 684
- ml\_delete\_user system procedure
  - syntax, 685
- ml\_device
  - MobiLink system table, 701
- ml\_device\_address
  - MobiLink system table, 703
- ml\_global script version
  - about, 325
- ml\_listening
  - MobiLink system table, 705
- ml\_passthrough
  - MobiLink system table, 707
- ml\_passthrough\_repair
  - MobiLink system table, 708
- ml\_passthrough\_script
  - MobiLink system table, 709
- ml\_passthrough\_status
  - MobiLink system table, 711
- ml\_property
  - MobiLink system table, 712
- ml\_pt
  - IBM DB2 mainframe system table name conversions, 696, 707
- ml\_pt\_repair
  - IBM DB2 mainframe system table name conversions, 696, 708
- ml\_pt\_script
  - IBM DB2 mainframe system table name conversions, 696, 709
- ml\_pt\_status
  - IBM DB2 mainframe system table name conversions, 696, 711
- ml\_qa\_clients
  - QAnywhere client system table, 713

---

- ml\_qa\_delivery
  - QAnywhere client system table, 714
- ml\_qa\_delivery\_archive
  - QAnywhere client system table, 716
- ml\_qa\_global\_props
  - QAnywhere client system table, 718
- ml\_qa\_notifications
  - QAnywhere client system table, 719
- ml\_qa\_repository
  - QAnywhere client system table, 720
- ml\_qa\_repository\_archive
  - QAnywhere client system table, 721
- ml\_qa\_repository\_props
  - QAnywhere client system table, 722
- ml\_qa\_repository\_props\_archive
  - QAnywhere client system table, 723
- ml\_qa\_repository\_staging
  - QAnywhere client system table, 724
- ml\_qa\_status\_history
  - QAnywhere client system table, 725
- ml\_qa\_status\_history\_archive
  - QAnywhere client system table, 726
- ml\_qa\_status\_staging
  - QAnywhere client system table, 727
- ml\_reset\_sstate system procedure
  - IBM DB2 conversion, 664
  - IBM DB2 mainframe system procedure name conversion, 686
- ml\_reset\_sync\_state system procedure
  - syntax, 686
- ml\_script
  - MobiLink system table, 728
- ml\_script\_modified
  - IBM DB2 mainframe system table name conversions, 696, 730
- ml\_script\_version
  - MobiLink system table, 729
- ml\_scripts\_modified
  - MobiLink system table, 730
- ml\_server
  - MobiLink system table, 731
- ml\_server\_delete system procedure
  - syntax, 686
- ml\_server\_update system procedure
  - syntax, 686
- ml\_set\_sis\_state system procedure
  - IBM DB2 conversion, 664
- ml\_sis\_sync\_state
  - MobiLink system table, 732
- ml\_subscription
  - MobiLink system table, 733
- ml\_table
  - MobiLink system table, 735
- ml\_table\_script
  - MobiLink system table, 736
- ml\_user
  - MobiLink system table, 737
  - MobiLink user authentication (mluser), 690
- mlDomConfig.xml
  - about, 602
- mlmon
  - about MobiLink Monitor, 179
  - starting, 181
- mlMonitorSettings
  - MobiLink Monitor settings, 190
- mlscript.jar
  - MobiLink Java synchronization logic, 529
- mlsrv11
  - (*see also* MobiLink server)
  - nc option, 75
  - connection string, 54
  - logging, 33
  - Notifier, 76
  - options, 45
  - QAnywhere, 73
  - reports error context in message log, 77
  - starting, 30
  - stopping, 32
  - syntax, 45
- mlsrv11 options
  - alphabetical list, 45
- mlsrv11 syntax
  - about, 45
- mlstop utility
  - deploying on 32-bit Unix, 807
  - deploying on 32-bit Windows, 801
  - deploying on 64-bit Unix, 810
  - deploying on 64-bit Windows, 804
  - methods for stopping MobiLink server, 32
  - options, 689
  - syntax, 689
- mluser utility
  - deploying on 32-bit Unix, 807
  - deploying on 32-bit Windows, 801
  - deploying on 64-bit Unix, 810
  - deploying on 64-bit Windows, 804

- options, 690
- syntax, 690
- mobile device
  - connecting to the Relay Server farm, 262
- MobiLink
  - .NET synchronization logic, 589
  - alphabetic list of events, 342
  - character set considerations, 789
  - connection parameters for mlsrv11, 107
  - connection parameters for Monitor, 181
  - consolidated databases, 3
  - data types, 739
  - deploying applications, 799
  - development tips, 128
  - event overview, 343
  - file-based download, 293
  - glossary definition, 839
  - handling conflicts, 146
  - Java synchronization logic, 527
  - mlsrv11 options, 44
  - Monitor, 179
  - multiple synchronization servers, 267
  - ODBC driver support, 794
  - performance, 169
  - Redirector, 265
  - running outside the current session, 35
  - running the synchronization server, 29
  - scripts, 313
  - starting, 30
  - stopping the MobiLink server, 32
  - synchronization techniques, 127
  - system procedures, 663
  - system tables, 694
  - web server configuration, 265
- MobiLink clients
  - deploying, 813
  - glossary definition, 839
- MobiLink connections
  - debugging, 41
- MobiLink consolidated databases
  - about, 3
  - ASE, 10
  - IBM DB2 LUW as, 12
  - IBM DB2 mainframe as, 15
  - MySQL as, 22
  - Oracle as, 25
  - SQL Anywhere as, 28
  - SQL Server as, 20
- MobiLink data mappings
  - about, 739
- MobiLink data mappings between remote and consolidated databases
  - about, 739
- MobiLink data types
  - .NET and SQL, 594
  - Java and SQL, 532
- MobiLink events
  - listed, 342
- MobiLink file transfer utility (mlfiletransfer)
  - mlsrv11 -ftr option, 71
- MobiLink generation numbers
  - file-based download, 299
- MobiLink log file viewer
  - MobiLink server logs, 34
- MobiLink Monitor
  - about, 179
  - chart pane, 188
  - deploying on 32-bit Unix, 807
  - deploying on 32-bit Windows, 801
  - deploying on 64-bit Unix, 810
  - deploying on 64-bit Windows, 804
  - details table pane, 184
  - glossary definition, 839
  - graph pane, 186
  - marquee tool, 189
  - options, 190
  - overview pane, 189
  - restoring defaults, 190
  - sample properties, 190
  - saving data, 193
  - session properties, 190
  - specifying watches, 194
  - starting, 181
  - statistical properties, 195
  - user interface, 184
  - using, 184
  - viewing in MS Excel, 193
  - Watch Manager, 194
  - zooming, 188
- MobiLink object-based data flow for Java and .NET
  - about, 649
- MobiLink performance
  - about, 169
  - key factors, 174
  - monitoring, 178
- MobiLink scripts

---

- listed, 342
- MobiLink server
  - (*see also* mlsrv11)
  - deploying, 801
  - glossary definition, 839
  - monitoring, 201
  - options, 45
  - running, 29
  - starting, 30
  - stop utility, 689
  - syntax, 45
- MobiLink server API for .NET
  - API reference, 606
  - ml\_property system table, 712
- MobiLink server API for Java
  - ml\_property system table, 712
- MobiLink server farm
  - lsc option, 72
  - failover, 40
  - load balancing, 40
- MobiLink server groups
  - about, 271
- MobiLink server log file viewer
  - MobiLink server logs, 34
- MobiLink server monitoring
  - about, 201
- MobiLink server options
  - about, 44
- MobiLink server shared state
  - server farm, 40
- MobiLink statistical properties
  - MobiLink Monitor, 195
- MobiLink stop utility (mlstop)
  - syntax, 689
- MobiLink stored procedures (*see* MobiLink system procedures)
- MobiLink synchronization
  - .NET synchronization logic, 589
  - consolidated databases, 3
  - file-based download, 293
  - Java synchronization logic, 527
  - overview of events, 343
  - performance, 169
  - restartable downloads, 158
  - web server configuration, 265
  - writing .NET classes, 595
  - writing Java classes, 533
- MobiLink synchronization logic
  - .NET, 589
  - alphabetic list of scripts, 342
  - data types for .NET and SQL, 594
  - data types for Java and SQL, 532
  - Java, 527
  - synchronization techniques, 127
  - writing scripts, 313
- MobiLink synchronization scripts
  - about, 313
  - alphabetic list of scripts, 342
  - constructing .NET classes, 594
  - constructing Java classes, 532
  - database transactions and .NET classes, 594
  - database transactions and Java classes, 531
  - debugging Java classes, 533
  - preserving database transactions in .NET, 594
  - preserving database transactions in Java, 531
  - writing .NET classes, 595
  - writing Java classes, 533
- MobiLink synchronization server (*see* MobiLink server)
- MobiLink system database
  - about, 7
- MobiLink system procedures
  - about, 663
- MobiLink system tables
  - about, 694
  - creating in consolidated database, 6
  - glossary definition, 839
  - ml\_active\_remote\_id, 697
  - ml\_column, 698
  - ml\_connection\_script, 699
  - ml\_database, 700
  - ml\_device, 701
  - ml\_device\_address, 703
  - ml\_listening, 705
  - ml\_passthrough, 707
  - ml\_passthrough\_repair, 708
  - ml\_passthrough\_script, 709
  - ml\_passthrough\_status, 711
  - ml\_property, 712
  - ml\_script, 728
  - ml\_script\_version, 729
  - ml\_scripts\_modified, 730
  - ml\_server, 731
  - ml\_sis\_sync\_state, 732
  - ml\_subscription, 733
  - ml\_table, 735

- ml\_table\_script, 736
- ml\_user, 737
- MobiLink user authentication utility (mluser)
  - syntax, 690
- MobiLink users
  - glossary definition, 839
  - ml\_user system table, 737
  - MobiLink user authentication (mluser), 690
  - registering with the mluser utility, 690
- MobiLink utilities
  - about, 687
  - MobiLink stop utility (mlstop), 689
  - MobiLink user authentication (mluser), 690
  - server, 687
- mod\_iaredirect.dll
  - configuring the Apache Redirector, 287
  - configuring the M-Business Anywhere Redirector, 289
- mod\_iaredirect.so
  - configuring the M-Business Anywhere Redirector, 289
- modify\_error\_message
  - connection event, 460
- modify\_last\_download\_timestamp
  - connection event, 463
- modify\_next\_last\_download\_timestamp
  - connection event, 466
- modify\_user
  - connection event, 469
- Monitor
  - about, 201
  - adding resources, 222
  - admin user, 228
  - alerts, 232
  - alerts error reports, 235
  - associating users with resources, 229
  - blackouts, 226
  - collection intervals, 223
  - connecting, 213
  - creating users, 228
  - deleting alerts, 233
  - deleting metrics, 220
  - deleting users, 230
  - disconnecting, 214
  - editing users, 229
  - email notification, 233
  - enabling email notification, 234
  - error reports, 235
  - exiting, 212
  - installing on a separate computer, 236
  - limitations, 203
  - logins required, 230
  - metric tabs, 217
  - metrics, 216
  - MobiLink Monitor, 179
  - monitoring resources, 222
  - network configuration, 203
  - quick start, 205
  - removing resources, 227
  - requirements, 202
  - resources, 215, 222
  - running in a production environment, 202
  - security, 230
  - start monitoring resources, 222
  - starting locally, 211
  - starting on a separate computer, 211
  - state, 215
  - status, 215
  - stop monitoring manually, 226
  - stop monitoring resources, 225
  - stop monitoring using blackouts, 226
  - stopping the Monitor, 212
  - tabs, 217
  - troubleshooting, 237
  - tutorial, 206
  - user types, 228
  - viewing alerts, 232
- Monitor metrics
  - alerts, 224
  - alerts tab, 218
  - collection intervals, 223
  - consolidated database tab, 220
  - deleting, 220
  - machine resources tab, 220
  - metric tabs, 217
  - server tab, 218
  - specifying metrics to collect , 223
  - suppress alerts for the same condition that occur within a specified time, 224
  - synchronization tab, 219
- monitoring
  - MobiLink performance, 178
  - MobiLink servers, 201
  - synchronizations in MobiLink, 179
- monitoring MobiLink performance
  - overview, 178



---

## MySQL

- MobiLink consolidated database, 22

- MobiLink data mapping for , 774

## MySQL consolidated database

- MobiLink, 22

## N

### named parameters

- about MobiLink, 320

- last\_download, 130

- last\_table\_download, 130

### named row parameters

- about MobiLink scripts, 320

- adding column information to the consolidated database, 666

### named script parameters

- about MobiLink, 320

- ml\_add\_column system procedure, 666

### natural joins

- glossary definition, 833

### Netscape web servers

- configuring the NSAPI Redirector on Unix, 280

- configuring the NSAPI Redirector on Windows, 277

### network parameters

- MobiLink server -x option, 107

### network protocols

- glossary definition, 839

- mlsrv11 -x option using TCP/IP, 107

- mlsrv11 -x option using TLS over TCP/IP, 108

- mlsrv11 using HTTP, 109

- mlsrv11 using HTTPS, 110

- MobiLink server, 107

### network server

- glossary definition, 839

### newsgroups

- technical support, xix

### NextRow method [ML .NET]

- DBRowReader interface syntax, 621

### non-blocking download acknowledgement

- MobiLink server (mlsrv11) -nba option, 74

- nonblocking\_download\_ack connection event, 471

- publication\_nonblocking\_download\_ack

- connection event, 475

### non-blockingdownload acknowledgement

- about, 161

### non-relational databases

- synchronizing with MobiLink, 649

### nonblocking\_download\_ack

- connection event, 471

### normalization

- glossary definition, 839

### Notifiers

- deploying on 32-bit Unix, 807

- deploying on 32-bit Windows, 801

- deploying on 64-bit Unix, 810

- deploying on 64-bit Windows, 804

- glossary definition, 840

### NSAPI Redirector

- configuring on Unix, 280

- configuring on Windows, 277

## O

### o.

- MobiLink named parameter prefix, 320

### obj.conf

- configuring for the NSAPI Redirector on Unix, 280

- configuring for the NSAPI Redirector on Windows, 277

### object trees

- glossary definition, 840

### object-based data flow (*see* direct row handling)

### objects

- MobiLink server API for .NET, 606

- MobiLink server API for Java, 543

### ODBC

- glossary definition, 840

- MobiLink drivers, 794

- multiple errors in MobiLink, 339

- Oracle driver, 795

### ODBC Administrator

- glossary definition, 840

### ODBC data sources

- glossary definition, 840

### ODBC drivers

- MobiLink character set conversion by, 791

- Oracle, 795

- supported by MobiLink, 794

### ODBC drivers supported by MobiLink

- about, 794

### offsets

- progress column of ml\_subscription table, 733

### old row parameters

- MobiLink scripts, 320

- online books
    - PDF, xiv
  - operators
    - Monitor users, 228
  - options
    - mlsrv11, 45
    - MobiLink server (mlsrv11), 45
    - MobiLink stop utility (mlstop), 689
    - MobiLink user authentication (mluser), 690
  - options section
    - Relay Server configuration file, 245
  - options window
    - MobiLink Monitor, 190
  - Oracle
    - as MobiLink consolidated database, 25
    - MobiLink data mapping, 779
    - MobiLink isolation levels, 165
    - ODBC driver, 795
    - sequences in MobiLink synchronization, 25
    - synchronizing LONG data, 784
  - Oracle consolidated database
    - MobiLink, 25
  - Oracle driver
    - encrypting passwords, 795
    - ODBC, 795
  - Oracle varray
    - example, 26
    - restrictions, 27
    - using in stored procedures, 26
  - Outbound Enabler
    - about, 247
    - deploying, 247
    - Relay Server, 242
    - syntax, 247
  - outer joins
    - glossary definition, 840
  - overlaps
    - partitioning, 135
  - overview pane
    - MobiLink Monitor, 189
  - P**
  - packaged download
    - MobiLink file-based download, 293
  - packages
    - glossary definition, 840
  - ParameterName property [ML .NET]
    - DBParameter class syntax, 613
  - parameters
    - synchronization scripts, 320
  - Parameters property [ML .NET]
    - DBCommand syntax, 608
  - parse trees
    - glossary definition, 840
  - partial download retained synchronization parameter
    - restartable downloads, 159
  - partitioning
    - about MobiLink, 135
    - defined, 135
    - MobiLink disjoint, 135
  - partitioning child tables
    - MobiLink, 137
  - partitioning rows
    - MobiLink among remote databases, 135
  - partitioning tables
    - example, 135
  - partitioning with overlaps
    - MobiLink, 136
- passwords
    - encrypting in Oracle driver, 795
    - MobiLink mluser utility, 690
    - Monitor users, 228
  - PDB
    - glossary definition, 840
  - PDF
    - documentation, xiv
  - performance
    - downloads, 173
    - MobiLink, 169
    - MobiLink -sm option, 94
    - MobiLink forced conflicts, 154
  - performance statistics
    - glossary definition, 841
  - permissions
    - MobiLink server, 30
  - personal server
    - glossary definition, 841
  - physical indexes
    - glossary definition, 841
  - plug-in modules
    - glossary definition, 841
  - policies
    - glossary definition, 841
  - polling
    - glossary definition, 841

---

port protocol option

- MobiLink Redirector, 269
- MobiLink server (mlsrv11) -x option for HTTP, 109
- MobiLink server (mlsrv11) -x option for HTTPS, 110
- MobiLink server (mlsrv11) -x option for TCP/IP, 107
- MobiLink server (mlsrv11) -x option for TLS over TCP/IP, 108

PowerDesigner

- glossary definition, 841

PowerJ

- glossary definition, 841

Precision property [ML .NET]

- DBParameter class syntax, 614

predicates

- glossary definition, 841

prefixes

- MobiLink named parameters, 320

Prepare method [ML .NET]

- DBCommand syntax, 607

prepare\_for\_download

- connection event, 473

prepare\_for\_download property

- MobiLink Monitor synchronization statistics, 195

primary key constraints

- glossary definition, 842

primary key pools

- generating unique values using default global autoincrement for MobiLink, 140
- MobiLink example, 143
- MobiLink unique primary keys, 143

primary keys

- glossary definition, 842
- MobiLink about, 128
- MobiLink uniqueness techniques, 139
- Oracle sequences, 25

primary tables

- glossary definition, 842

printing information from .NET

- MobiLink .NET synchronization logic, 597

priority synchronization

- MobiLink performance, 172

private assemblies

- implementing in MobiLink, 601

procedure calls

- SQL Server MobiLink consolidated databases, 20

procedures

- MobiLink, 663

procedures return results

- Oracle driver option, 795

ProcResults

- Oracle driver option, 795

progress

- progress column of ml\_subscription table, 733

progress counter

- progress column of ml\_subscription table, 733

progress offsets

- progress column of ml\_subscription table, 733

properties

- QAnywhere server, 73

property

- DBParameter class syntax, 614

protocols

- mlsrv11 -x option using TCP/IP, 107
- mlsrv11 -x option using TLS over TCP/IP, 108
- mlsrv11 using HTTP, 109
- mlsrv11 using HTTPS, 110
- MobiLink server, 107

proxy tables

- glossary definition, 842

proxy web servers

- MobiLink, 266

pseudocode

- MobiLink events, 343

publication updates

- glossary definition, 842

publication\_nonblocking\_download\_ack

- connection event, 475

publications

- glossary definition, 842

publisher

- glossary definition, 843

push notifications

- glossary definition, 843

push requests

- glossary definition, 843

**Q**

QAnywhere

- deploying, 816
- glossary definition, 843
- MobiLink system tables, 694
- properties, 73



---

- referenced object
  - glossary definition, 844
- referencing object
  - glossary definition, 844
- referential integrity
  - glossary definition, 844
- registering
  - methods as MobiLink scripts, 664
- registering methods
  - MobiLink server API for .NET, 595
  - MobiLink server API for Java, 533
- registering MobiLink users
  - mluser utility, 690
- regular expressions
  - glossary definition, 844
- Relay Server
  - about, 240
  - architecture, 240
  - back-end server farm, 242
  - configuration file, 243
  - deployment, 253
  - hosted service, 260
  - HTTP load balancer, 242
  - Outbound Enabler, 242, 247
  - Outbound Enabler deployment, 247
  - Outbound Enabler syntax, 247
  - Relay Server farm, 241
  - rshost command line syntax, 251
  - rsoe syntax, 247
  - State Manager, 250
  - synchronizing through a web server, 239
  - updating configuration, 258
  - using MobiLink, 262
- Relay Server configuration file
  - about, 243
  - backend farm section, 244
  - backend server section, 245
  - format, 246
  - options section, 245
  - Relay Server section, 243
- Relay Server deployment
  - Apache on Linux, 256
  - application pool, creating, 254
  - files for Linux, 256
  - files for Windows, 253
  - IIS on Windows, 253
  - State Manager, 254, 256
  - web server extensions, 254, 256
- Relay Server farm
  - connecting a client, 262
  - connecting a mobile device, 262
- Relay Server farm configuration
  - updating, 258
- Relay Server for MobiLink
  - sample scenario, 263
- Relay Server hosting service
  - subscribing, 260
- Relay Server section
  - Relay Server configuration file, 243
- Relay Server State Manager
  - about, 250
  - command line syntax, 251
  - starting as a Windows service, 250
  - starting automatically, 251
  - starting automatically with custom options, 251
- Relay Server web extensions
  - about, 239
  - enabling, 254
- remote databases
  - glossary definition, 844
  - mapping of data types in MobiLink, 739
  - relating consolidated tables to MobiLink remote tables, 5
- REMOTE DBA authority
  - glossary definition, 844
- remote IDs
  - getRemoteID method in MobiLink Java API, 546
  - glossary definition, 845
  - ml\_database system table, 700
- remote tables
  - deleting rows in MobiLink, 335
- Remove( object value ) method [ML .NET]
  - DBParameterCollection class syntax, 618
- RemoveAt( int index ) method [ML .NET]
  - DBParameterCollection class syntax, 618
- RemoveAt( string parameterName ) method [ML .NET]
  - DBParameterCollection class syntax, 616
- removeErrorListener method [ML Java]
  - ServerContext syntax, 572
- removeInfoListener method [ML Java]
  - ServerContext syntax, 572
- removeShutdownListener method [ML Java]
  - ServerContext syntax, 572
- removeWarningListener method [ML Java]
  - ServerContext syntax, 573

- replication
    - glossary definition, 845
  - Replication Agent
    - glossary definition, 845
  - replication frequency
    - glossary definition, 845
  - replication messages
    - glossary definition, 845
  - Replication Server
    - glossary definition, 845
  - report\_error
    - connection event, 477
    - syntax, 338
  - report\_odbc\_error
    - connection event, 480
  - reporting errors
    - MobiLink synchronization, 338
  - requests
    - routing in MobiLink, 265
  - required scripts
    - MobiLink, 326
  - resetting
    - MobiLink last download time, 131
  - resolution
    - MobiLink conflict resolution, 146
  - resolve\_conflict
    - table event, 483
    - using, 149
  - resolving
    - MobiLink conflicts, 146
  - resolving conflicts
    - MobiLink overview, 149
    - MobiLink with resolve\_conflict scripts, 149
    - MobiLink with upload\_update scripts, 151
    - resolve\_conflict script, 149
    - upload\_update script, 151
  - resources
    - Monitor, 215
  - restartable downloads
    - MobiLink, 158
  - resume partial download synchronization parameter
    - restartable downloads, 159
  - resuming failed downloads
    - MobiLink, 158
  - return values
    - .NET synchronization, 595
    - Java synchronization, 533
  - reverse proxy
    - defined, 266
  - role names
    - glossary definition, 845
  - roles
    - glossary definition, 845
  - rollback logs
    - glossary definition, 846
  - Rollback method [ML .NET]
    - DBConnection syntax, 609
  - routing requests
    - MobiLink synchronization, 265
  - row handling in MobiLink (*see* direct row handling)
  - row parameters
    - MobiLink scripts, 320
  - row-level triggers
    - glossary definition, 846
  - rows
    - deleting on remote MobiLink databases, 335
    - partitioning, 135
  - rsa protocol option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
    - MobiLink server (mlsrv11) -x option for TCP/IP, 108
  - rshost (*see* Relay Server State Manager)
  - running
    - MobiLink server, 29
  - running .NET synchronization logic
    - about, 591
  - running Java synchronization logic
    - about, 529
  - running MobiLink
    - about, 29
    - as a daemon, 35
    - outside the current session, 35
  - running MobiLink server
    - as a service, 35
  - running the MobiLink server
    - about, 30
    - in a server farm, 40
- S**
- s.
    - MobiLink named parameter prefix, 320
  - sample domain configuration file
    - MobiLink, 602
  - sample properties

---

- MobiLink Monitor, 190
- samples
  - .NET synchronization logic, 604
  - Java synchronization logic, 538
- samples-dir
  - documentation usage, xvi
- saving Monitor data
  - MobiLink Monitor, 193
- Scale property [ML .NET]
  - DBParameter class syntax, 614
- schemas
  - glossary definition, 846
  - relating consolidated tables to MobiLink remote tables, 5
- script parameters
  - about MobiLink, 320
  - last\_download, 130
  - last\_table\_download, 130
- script types
  - MobiLink, 318
- script versions
  - about MobiLink synchronization, 324
  - adding, 325
  - global, 325
  - glossary definition, 846
  - reserved names, 324
- script-based uploads
  - glossary definition, 846
- scripts
  - about MobiLink, 313
  - adding and deleting .NET connection scripts, 668
  - adding and deleting .NET table scripts, 669
  - adding and deleting Java connection scripts, 671
  - adding and deleting Java table scripts, 672
  - adding and deleting SQL connection scripts, 667
  - adding and deleting SQL table scripts, 680
  - adding to the consolidated database in MobiLink, 327
  - connection scripts, 318
  - global script versions, 325
  - glossary definition, 846
  - MobiLink event overview, 343
  - MobiLink events, 342
  - MobiLink ml\_active\_remote\_id system table, 697
  - MobiLink ml\_column system table, 698
  - MobiLink ml\_connection\_script system table, 699
  - MobiLink ml\_database system table, 700
  - MobiLink ml\_device system table, 701
  - MobiLink ml\_device\_address system table, 703
  - MobiLink ml\_listening system table, 705
  - MobiLink ml\_passthrough system table, 707
  - MobiLink ml\_passthrough\_repair system table, 708
  - MobiLink ml\_passthrough\_script system table, 709
  - MobiLink ml\_passthrough\_status system table, 711
  - MobiLink ml\_property system table, 712
  - MobiLink ml\_script system table, 728
  - MobiLink ml\_script\_version system table, 729
  - MobiLink ml\_scripts\_modified system table, 730
  - MobiLink ml\_server system table, 731
  - MobiLink ml\_sis\_sync\_state system table, 732
  - MobiLink ml\_subscription system table, 733
  - MobiLink ml\_table\_script system table, 736
  - MobiLink ml\_user system table, 737
  - required by MobiLink, 326
  - supported DBMS scripting strategies, 8
  - table scripts, 318
  - versions, 324
  - writing scripts to download rows, 333
  - writing scripts to upload rows, 330
- secured features
  - glossary definition, 846
- security
  - MobiLink user authentication (mluser) utility, 690
  - Monitor users, 230
- selective sharing (*see* partitioning)
- self-referencing tables
  - MobiLink, 164
- sending
  - Monitor alert emails, 233
- sequence number
  - progress column of ml\_subscription table, 733
- sequence of MobiLink events
  - pseudocode, 346
- sequences
  - primary key uniqueness in MobiLink synchronization, 25
- server farm
  - zs option, 119
  - back-end server, 242
  - load balancing, 40
  - MobiLink, 40
  - Relay Server, 241
- server groups
  - MobiLink, 271
- server management requests

- glossary definition, 846
- server message stores
  - glossary definition, 847
- server monitoring
  - about, 201
- server system procedures
  - MobiLink, 663
- server-initiated synchronization
  - glossary definition, 846
- ServerContext [ML Java]
  - syntax, 565
- ServerContext interface [ML .NET]
  - syntax, 629
- ServerException class [ML .NET]
  - syntax, 632
- ServerException class [ML Java]
  - syntax, 574
- ServerException constructors [ML .NET]
  - syntax, 632
- ServerException constructors [ML Java]
  - syntax, 575
- servers
  - MobiLink synchronization [mlsrv11], 30
- service dependencies
  - MobiLink, 38
- services
  - configuring, 36
  - deleting, 36
  - dependencies, 38
  - glossary definition, 847
  - MobiLink, 36
  - MobiLink server, 35
  - running MobiLink as a service, 35
  - running multiple, 38
- servlet Redirector
  - Apache Tomcat, 284
  - Apache web servers, 284
- servlets
  - installing the Redirector for Apache web servers, 284
- session properties
  - MobiLink Monitor, 190
- session-based synchronization
  - glossary definition, 847
- session-wide variables
  - DB2 MobiLink consolidated databases, 13
  - Oracle MobiLink consolidated databases, 25
- session\_key
  - MobiLink server (mlsrv11) -xo option for HTTP, 115
- session\_key protocol option
  - MobiLink server (mlsrv11) -xo option for HTTPS, 116
- SET NOCOUNT
  - SQL Server MobiLink consolidated databases, 20
- SetNewRowValues method [ML .NET]
  - UpdateDataReader interface syntax, 643
- setNewRowValues method [ML Java]
  - SynchronizationException syntax, 578
- SetOldRowValues method [ML .NET]
  - UpdateDataReader interface syntax, 643
- setOldRowValues method [ML Java]
  - MobiLink server API for Java
  - SynchronizationException class, 579
- setting up
  - MobiLink consolidated databases, 3
  - MobiLink file-based downloads, 295
  - MobiLink Java synchronization logic, 529
  - MobiLink Redirector about, 268
- setting up .NET synchronization logic
  - about, 591
- setting up a MySQL consolidated database
  - MobiLink, 22
- setting up a SQL Anywhere consolidated database
  - MobiLink, 28
- setting up a Sybase ASE consolidated database
  - MobiLink, 10
- setting up an IBM DB2 LUW consolidated database
  - MobiLink, 12
- setting up an IBM DB2 mainframe consolidated database
  - MobiLink, 15
- setting up an Oracle consolidated database
  - MobiLink, 25
- setting up direct row handling
  - about, 651
- setting up the Redirector
  - about, 268
- setup
  - MobiLink .NET synchronization logic, 591
  - MobiLink consolidated databases, 6
  - MobiLink Java synchronization logic, 529
- setup scripts
  - MobiLink consolidated databases, 6
  - MobiLink system database, 7
- setValue method [ML Java]



---

- InOutInteger syntax, 558
- InOutString syntax, 560
- shadow tables
  - writing download\_delete\_cursor scripts, 335
- shared assemblies
  - implementing in MobiLink, 601
- shared server state
  - zs option, 119
- shared stated
  - MobiLink server farm, 40
- sharing rules (*see* partitioning)
- ShutDown method [ML .NET]
  - ServerContext interface syntax, 630
- shutdown method [ML Java]
  - ServerContext syntax, 573
- ShutdownCallback delegate [ML .NET]
  - syntax, 633
- ShutdownListener interface [ML Java]
  - syntax, 575
- ShutdownListener method [ML .NET]
  - ServerContext interface syntax, 630
- shutdownPerformed method [ML Java]
  - ShutdownListener syntax, 576
- shutting down
  - MobiLink server, 32
  - MobiLink stop utility (mlstop), 689
- SID
  - Oracle driver option, 795
- SLEEP
  - Redirector property (for Redirectors that don't support server groups), 275
  - Redirector property (for Redirectors that support server groups), 273
- snapshot isolation
  - glossary definition, 847
  - MobiLink, 165
  - MobiLink -dsd option to disable, 63
  - MobiLink -dt option for SQL Server, 64
  - MobiLink -esu option to enable for uploads, 66
- snapshot synchronization
  - about, 133
- soft shutdown
  - MobiLink stop utility (mlstop), 689
- sort order
  - characters and MobiLink synchronization, 790
- spreadsheets
  - synchronizing with MobiLink, 649
- SQL
  - glossary definition, 847
  - SQL Anywhere
    - as MobiLink consolidated database, 28
    - documentation, xiv
    - glossary definition, 847
    - MobiLink isolation levels, 165
  - SQL Anywhere consolidated database
    - MobiLink, 28
  - SQL Remote
    - glossary definition, 847
  - SQL Server
    - (*see also* Microsoft SQL Server)
    - as MobiLink consolidated database, 20
    - MobiLink data mapping, 767
  - SQL statements
    - glossary definition, 847
  - SQL synchronization logic
    - MobiLink, 313
    - MobiLink performance, 172
  - SQL syntax
    - MobiLink server (mlsrv11), 45
  - SQL-.NET data types
    - MobiLink .NET synchronization logic, 594
  - SQL-based synchronization
    - glossary definition, 847
  - SQL-java data types
    - about, 532
  - SQL\_ARD\_TYPE field [ML .NET]
    - SQLType enumeration syntax, 637
  - SQL\_BIGINT field [ML .NET]
    - SQLType enumeration, 638
  - SQL\_BINARY field [ML .NET]
    - SQLType enumeration syntax, 638
  - SQL\_BIT field [ML .NET]
    - SQLType enumeration syntax, 637
  - SQL\_CHAR field [ML .NET]
    - SQLType enumeration syntax, 633
  - SQL\_DATE field [ML .NET]
    - SQLType enumeration syntax, 635
  - SQL\_DATETIME field [ML .NET]
    - SQLType enumeration syntax, 635
  - SQL\_DECIMAL field [ML .NET]
    - SQLType enumeration syntax, 634
  - SQL\_DEFAULT field [ML .NET]
    - SQLType enumeration syntax, 637
  - SQL\_DOUBLE field [ML .NET]
    - SQLType enumeration syntax, 635
  - SQL\_FLOAT field [ML .NET]

- SQLType enumeration syntax, 635
- SQL\_GUID field [ML .NET]
  - SQLType enumeration syntax, 639
- SQL\_INTEGER field [ML .NET]
  - SQLType enumeration syntax, 634
- SQL\_INTERVAL field [ML .NET]
  - SQLType enumeration syntax, 636
- SQL\_LONGVARBINARY field [ML .NET]
  - SQLType enumeration syntax, 638
- SQL\_LONGVARCHAR field [ML .NET]
  - SQLType enumeration syntax, 639
- SQL\_NUMERIC field [ML .NET]
  - SQLType enumeration syntax, 634
- SQL\_REAL field [ML .NET]
  - SQLType enumeration syntax, 635
- SQL\_SMALLINT field [ML .NET]
  - SQLType enumeration syntax, 634
- SQL\_TIME field [ML .NET]
  - SQLType enumeration syntax, 636
- SQL\_TIMESTAMP field [ML .NET]
  - SQLType enumeration syntax, 636
- SQL\_TINYINT field [ML .NET]
  - SQLType enumeration syntax, 638
- SQL\_TXN\_READ\_COMMITTED
  - MobiLink isolation levels, 165
- SQL\_TYPE\_DATE field [ML .NET]
  - SQLType enumeration syntax, 636
- SQL\_TYPE\_NULL field [ML .NET]
  - SQLType enumeration syntax, 633
- SQL\_TYPE\_TIME field [ML .NET]
  - SQLType enumeration syntax, 637
- SQL\_TYPE\_TIMESTAMP field [ML .NET]
  - SQLType enumeration syntax, 637
- SQL\_UNKNOWN\_TYPE field [ML .NET]
  - SQLType enumeration syntax, 633
- SQL\_VARBINARY field [ML .NET]
  - SQLType enumeration syntax, 638
- SQL\_VARCHAR field [ML .NET]
  - SQLType enumeration syntax, 636
- SQL\_WCHAR field [ML .NET]
  - SQLType enumeration, 639
- SQL\_WLONGVARCHAR field [ML .NET]
  - SQLType enumeration syntax, 639
- SQL\_WVARCHAR field [ML .NET]
  - SQLType enumeration syntax, 639
- SQLType enumeration [ML .NET]
  - ServerException class syntax, 633
- start classes
  - DMLStartClasses option for Java, 92
  - MLStartClasses option for .NET, 90
  - MobiLink .NET synchronization logic, 595
  - MobiLink Java synchronization logic, 536
- start\_time property
  - MobiLink Monitor synchronization statistics, 195
- starting
  - MobiLink Monitor (mlmon), 181
  - MobiLink server, 30
  - Notifiers, 76
- starting the MobiLink Monitor
  - about, 181
- state
  - Monitor, 215
  - progress column of ml\_subscription table, 733
- state management
  - Relay Server, 250
- State Manager
  - command line syntax, 251
- statement-based scripts
  - uploading rows, 330
- statement-based uploads
  - conflict detection, 147
- statement-level triggers
  - glossary definition, 848
- StaticCursorLongColBuffLen
  - ASE, 10
- statistical properties
  - MobiLink, 195
- statistics
  - MobiLink, 195
- status
  - Monitor, 215
- STOP SYNCHRONIZATION DELETE statement
  - SQL Anywhere clients, 156
  - usage, 335
- stop utility (mlstop)
  - syntax, 689
- stopping
  - MobiLink server, 32
  - MobiLink stop utility (mlstop), 689
  - upload of deletes for SQL Anywhere clients, 156
- stored procedures
  - glossary definition, 848
  - MobiLink, 663
  - MobiLink stored procedure source code, 328
  - using to add or delete synchronization scripts, 328
  - using to download data, 162

---

- string literal
  - glossary definition, 848
- subqueries
  - glossary definition, 848
- subscriptions
  - glossary definition, 848
  - ml\_subscription system table, 733
- subsets
  - downloading subsets of data to remotes, 135
- Sun Java System web servers
  - configuring the NSAPI Redirector on Unix, 280
  - configuring the NSAPI Redirector on Windows, 277
- Sun One
  - configuring for the NSAPI Redirector on Unix, 280
  - configuring for the NSAPI Redirector on Windows, 277
- Sun web servers
  - configuring the NSAPI Redirector on Unix, 280
  - configuring the NSAPI Redirector on Windows, 277
- support
  - newsgroups, xix
- supporting old and new clients
  - MobiLink, 271
- suppress unsubmitted error reports
  - Monitor, 235
- switches
  - MobiLink server (mlsrv11), 45
  - MobiLink user authentication (mluser), 690
- Sybase Adaptive Server Enterprise (*see* Adaptive Server Enterprise)
- Sybase Central
  - glossary definition, 848
  - MobiLink server deployment on 32-bit Unix, 807
  - MobiLink server deployment on 64-bit Unix, 810
- sync property
  - MobiLink Monitor synchronization statistics, 195
- sync.conf
  - M-Business Anywhere Redirector, 289
- sync\_deadlocks property
  - MobiLink Monitor synchronization statistics, 195
- sync\_errors property
  - MobiLink Monitor synchronization statistics, 195
- sync\_request property
  - MobiLink Monitor synchronization statistics, 195
- sync\_tables property
  - MobiLink Monitor synchronization statistics, 195
- sync\_warnings property
  - MobiLink Monitor synchronization statistics, 195
- syncase.sql
  - about, 10
- syncdb2long.sql
  - about, 12
- synchronization
  - alphabetic list of scripts, 342
  - ASE data types in MobiLink, 740
  - conflict resolution, 146
  - connection parameters for Monitor, 181
  - consolidated databases, 3
  - data type mappings in MobiLink, 739
  - deleting rows, 335
  - downloading rows, 333
  - event overview, 343
  - glossary definition, 849
  - IBM DB2 LUW data types in MobiLink, 749
  - IBM DB2 mainframe data types in MobiLink, 756
  - many-to-many relationships, 136
  - Microsoft SQL Server data types in MobiLink, 767
  - MobiLink character set conversion, 790
  - MobiLink character sets, 790
  - MobiLink system procedures, 663
  - MobiLink system tables, 694
  - MobiLink utilities, 687
  - MySQL data types in MobiLink, 774
  - Oracle data types in MobiLink, 779
  - performance tips, 169
  - process, 317
  - protocol options for mlsrv11, 107
  - restartable downloads, 158
  - running the MobiLink server, 29
  - snapshot, 133
  - techniques, 127
  - web server configuration, 265
  - writing MobiLink scripts in .NET, 589
  - writing MobiLink scripts in Java, 527
  - writing scripts, 313
- synchronization errors
  - handling MobiLink, 338
  - troubleshooting, 65
- synchronization events
  - about, 342
  - about MobiLink synchronization, 343
  - alphabetic list of event scripts, 342
  - ASE begin\_connection\_autocommit connection event, 368

- MobiLink download, 351
- MobiLink upload, 349
- synchronization logic
  - MobiLink, 313
- synchronization parameters
  - HTTP synchronization, 107
  - HTTPS synchronization, 107
  - TCP/IP synchronization, 107
- synchronization properties
  - MobiLink Monitor, 191
- synchronization scripts
  - .NET, 589
  - .NET methods, 595
  - about, 313
  - adding and deleting, 327
  - adding or deleting with stored procedures, 328
  - adding with Sybase Central, 327
  - connection scripts, 318
  - DBMS dependencies, 8
  - download\_cursor, 334
  - example, 316
  - execution during, 317
  - handle\_error event, 338
  - implementing for .NET, 591
  - implementing for Java, 529
  - Java, 527
  - Java methods, 533
  - MobiLink events, 342
  - parameters, 320
  - report\_error, 338
  - script versions, 324
  - supported DBMS scripting strategies, 8
  - table scripts, 318
  - types, 318
  - writing scripts to download rows, 333
  - writing scripts to upload rows, 330
- synchronization sequence number
  - progress column of ml\_subscription table, 733
- synchronization server (*see* MobiLink server)
- synchronization stream libraries
  - MobiLink server deployment on 32-bit Unix, 807
  - MobiLink server deployment on 32-bit Windows, 801
  - MobiLink server deployment on 64-bit Unix, 810
  - MobiLink server deployment on 64-bit Windows, 804
- synchronization subscriptions
  - (*see also* subscriptions)
- synchronization tables
  - MobiLink ml\_table system table, 735
- synchronization techniques
  - about, 127
  - data entry, 155
  - deleting rows, 156
  - failed downloads, 158
  - partitioning, 135
  - primary key pools, 143
  - snapshot-based synchronization, 133
  - stored procedures to download, 162
  - timestamp-based synchronization, 129
  - uploading rows, 330
- synchronization users
  - MobiLink user authentication (mluser), 690
- synchronization\_statistics
  - connection event, 486
  - table event, 489
- SynchronizationException class [ML .NET]
  - syntax, 640
- SynchronizationException class [ML Java]
  - syntax, 576
- SynchronizationException constructors [ML .NET]
  - SynchronizationException class syntax, 640
- SynchronizationException constructors [ML Java]
  - SynchronizationException syntax, 577
- synchronizing data sources other than consolidated databases
  - about, 649
- synchronizing new remotes
  - MobiLink file-based download, 296
- synchronizing self-referencing tables
  - MobiLink, 164
- synchronizing through a web server
  - Redirector (deprecated), 265
  - Relay Server, 239
- syncmss.sql
  - about, 20
- syncora.sql
  - about, 25
- SyncRoot property [ML .NET]
  - DBParameterCollection class syntax, 620
- syncsa.sql
  - about, 28
- syntax
  - MobiLink scripts, 342
  - MobiLink server (mlsrv11), 45
  - MobiLink stop utility (mlstop), 689

- 
- MobiLink synchronization utilities, 687
  - MobiLink system procedures, 663
  - MobiLink user authentication (mluser), 690
  - SYS
    - glossary definition, 849
  - system database
    - MobiLink, 7
  - system objects
    - glossary definition, 849
  - system parameters
    - MobiLink scripts, 320
  - system procedures
    - alphabetical list of MobiLink system procedures, 664
    - ml\_add\_cs, 667
    - ml\_add\_dcs, 668
    - ml\_add\_dts, 669
    - ml\_add\_jcs, 671
    - ml\_add\_lcs, 673
    - ml\_add\_lcs\_chk, 673
    - ml\_add\_lts, 673
    - MobiLink , 663
    - MobiLink IBM DB2 mainframe system procedure name conversions, 664
  - system procedures to add or delete properties
    - MobiLink server, 664
  - system procedures to add or delete scripts
    - MobiLink server, 664
  - system tables
    - creating in MobiLink consolidated database, 6
    - glossary definition, 849
    - MobiLink synchronization, 694
  - system views
    - glossary definition, 849
  - T**
  - table scripts
    - about, 318
    - adding .NET scripts, 669
    - adding Java scripts, 672
    - adding SQL scripts, 680
    - adding with Sybase Central, 327
    - alphabetic list of MobiLink scripts, 342
    - defined, 315, 318
    - deleting .NET scripts, 669
    - deleting Java scripts, 672
    - deleting SQL scripts, 680
  - table-level scripts
    - defined, 318
  - tables
    - MobiLink ml\_table system table, 735
    - partitioning, 135
    - relating consolidated tables to MobiLink remote tables, 5
  - tablespace capacity
    - DB2 MobiLink consolidated databases, 13
  - TCP/IP
    - MobiLink server -x option, 107
  - technical support
    - newsgroups, xix
  - temporary tables
    - glossary definition, 849
  - text files
    - synchronizing with MobiLink, 649
  - Text property [ML .NET]
    - DBRowReader interface syntax, 628
  - this[ int index ] property [ML .NET]
    - DBParameterCollection class syntax, 620
  - this[ string parameterName ] property [ML .NET]
    - DBParameterCollection class syntax, 620
  - threading
    - (*see also* threads)
    - MobiLink performance, 170
  - threads
    - MobiLink worker threads and performance, 170
  - time changes
    - MobiLink, 132
  - time\_statistics
    - connection event, 492
    - table event, 495
  - timestamp-based downloads
    - about, 129
  - timestamp-based synchronization
    - about, 129
    - download\_cursor script, 130
    - download\_delete\_cursor script, 129
  - timestamps
    - MobiLink download, 131
  - tips
    - performance of MobiLink, 169
    - synchronization techniques, 128
  - TLS
    - (*see also* transport-layer security)
    - MobiLink client deployment on Unix, 814
    - MobiLink client deployment on Windows, 813

- MobiLink server -x option, 107
  - MobiLink server deployment on 32-bit Unix, 807
  - MobiLink server deployment on 32-bit Windows, 801
  - MobiLink server deployment on 64-bit Unix, 810
  - MobiLink server deployment on 64-bit Windows, 804
  - tls\_type protocol option
    - MobiLink server (mlsrv11) -x option for HTTPS, 110
    - MobiLink server (mlsrv11) -x option for TCP/IP, 108
  - Tomcat
    - configuring the servlet Redirector, 284
    - Redirector supported versions, 284
  - tools
    - MobiLink Monitor marquee tool, 189
  - topics
    - graphic icons, xviii
  - transaction log
    - glossary definition, 849
  - transaction log mirror
    - glossary definition, 850
  - transactional integrity
    - glossary definition, 850
  - transactions
    - glossary definition, 849
    - MobiLink, 343
    - MobiLink .NET synchronization logic, 594
    - MobiLink Java synchronization logic, 531
  - translation between character sets
    - MobiLink synchronization, 790
  - transmission rules
    - glossary definition, 850
  - triggers
    - glossary definition, 850
  - troubleshooting
    - handling failed downloads, 158
    - MobiLink remote data loss, 131
    - MobiLink restartable downloads, 158
    - MobiLink server log, 33
    - MobiLink server startup, 41
    - Monitor, 237
    - newsgroups, xix
    - synchronization errors, 65
  - tuning performance
    - MobiLink , 174
  - tutorials
    - Monitor, 206
  - Type property [ML .NET]
    - DBRowReader interface syntax, 628
- ## U
- u.
    - MobiLink user-defined parameter prefix, 322
  - ULRollbackPartialDownload function
    - restartable downloads, 159
  - UltraLite
    - deploying, 815
    - glossary definition, 850
  - UltraLite runtime
    - glossary definition, 850
  - uni-directional synchronization
    - about, 138
  - unique
    - primary keys in MobiLink, 139
  - unique constraints
    - glossary definition, 850
  - unique keys
    - MobiLink, 139
  - unique primary keys
    - generating for MobiLink using composite keys, 139
    - generating for MobiLink using UUIDs, 139
    - generating using key pools for MobiLink, 143
    - global autoincrement for MobiLink, 140
    - MobiLink, 139
  - Unix
    - MobiLink server as a daemon, 35
  - unknown\_timeout protocol option
    - synchronizing across firewalls, 269
  - unload
    - glossary definition, 851
  - unsubmitted error reports
    - Monitor , 218
    - Monitor alerts, 235
  - UPDATE conflicts
    - MobiLink, 146
  - UpdateData interface [ML .NET]
    - syntax, 641
  - UpdateDataReader interface [ML .NET]
    - syntax, 642
  - UpdateResultSet [ML Java]
    - SynchronizationException syntax, 578
  - upgrading applications

- using multiple MobiLink script versions, 324
- upload events
  - about, 330
  - MobiLink synchronization, 349
- upload property
  - MobiLink Monitor synchronization statistics, 195
- upload transaction
  - MobiLink, 344
- upload-only and download-only synchronizations
  - about, 138
- upload-only synchronization
  - about, 138
  - required scripts, 326
- upload\_bytes property
  - MobiLink Monitor synchronization statistics, 195
- upload\_deadlocks property
  - MobiLink Monitor synchronization statistics, 195
- upload\_delete
  - table event, 498
- upload\_deleted\_rows property
  - MobiLink Monitor synchronization statistics, 195
- upload\_errors property
  - MobiLink Monitor synchronization statistics, 195
- upload\_fetch
  - detecting conflicts, 147
  - overview of conflict detection, 147
  - table event, 500
- upload\_fetch\_column\_conflict
  - detecting conflicts, 147
  - overview of conflict detection, 147
  - table event, 502
- upload\_insert
  - table event, 504
- upload\_inserted\_rows property
  - MobiLink Monitor synchronization statistics, 195
- upload\_new\_row\_insert
  - table event, 506
- upload\_old\_row\_insert
  - table event, 509
- upload\_statistics
  - connection event, 512
  - table event, 517
- upload\_update
  - detecting conflicts, 148
  - overview of conflict detection, 147
  - table event, 522
  - using, 151
- upload\_updated\_rows property
  - MobiLink Monitor synchronization statistics, 195
- upload\_warnings property
  - MobiLink Monitor synchronization statistics, 195
- UploadData interface [ML Java]
  - syntax, 579
- UploadedTableData interface [ML .NET]
  - syntax, 643
- UploadedTableData interface [ML Java]
  - syntax, 581
- uploading data from self-referencing tables
  - about, 164
- uploading rows
  - .NET synchronization techniques, 600
  - MobiLink performance, 173
  - writing scripts, 330
- uploads
  - glossary definition, 851
  - MobiLink scripts to upload rows, 330
  - MobiLink temporarily stopping, 156
  - MobiLink transaction, 344
- url\_suffix protocol option
  - MobiLink Redirector, 269
- user authentication utility (mluser)
  - syntax, 690
- user names
  - MobiLink user authentication utility (mluser), 690
- user parameters
  - MobiLink, 322
- user property
  - MobiLink Monitor synchronization statistics, 195
- User property [ML .NET]
  - DBRowReader interface, 628
- user-defined data types
  - glossary definition, 851
- user-defined parameters
  - MobiLink, 322
- user-defined procedures
  - DB2 MobiLink consolidated databases, 13
- user-defined start classes
  - MobiLink .NET synchronization logic, 595
  - MobiLink Java synchronization logic, 536
- users
  - Monitor admin user, 228
  - Monitor administrators, 228
  - Monitor creating, 228
  - Monitor default user, 228
  - Monitor deleting, 230
  - Monitor emailing, 229

- Monitor operators, 228
- Monitor read-only users, 228
- Monitor security, 230
- Monitor types, 228
- utilities
  - MobiLink, 687
  - MobiLink Redirector, 265
  - MobiLink server (mlsrv11), 45
  - MobiLink stop utility (mlstop), 689
  - MobiLink user authentication (mluser), 690
- utilization graph pane
  - MobiLink Monitor, 186
- UUIDs
  - MobiLink synchronization application, 139
- V**
- validate
  - glossary definition, 851
- validating
  - MobiLink automatically, 298
  - MobiLink custom , 300
  - MobiLink file-based download, 298
- validation checks
  - MobiLink file-based download, 298
- Value property [ML .NET]
  - DBParameter class syntax, 615
- VARBIT data type
  - restrictions in ASE MobiLink consolidated databases, 10
- VARCHAR data type
  - MobiLink and other DBMSs, 8
- varray (Oracle)
  - example, 26
  - restrictions, 27
  - using in stored procedures, 26
- verbosity
  - MobiLink performance, 172
  - MobiLink server (mlsrv11) -v option, 102
- version property
  - MobiLink Monitor synchronization statistics, 195
- version protocol option
  - MobiLink server (mlsrv11) -x option for HTTP, 109
  - MobiLink server (mlsrv11) -x option for HTTPS, 110
- versions
  - about MobiLink synchronization scripts, 324

- adding script versions, 325
- viewing MobiLink logs
  - about, 34
- views
  - glossary definition, 851
- Visual Basic
  - MobiLink synchronization scripts, 589
  - support in MobiLink .NET, 590
- Visual Studio
  - MobiLink synchronization scripts, 589

**W**

- WARNING [ML Java]
  - Java LogMessage interface, 564
- WARNING field [ML .NET]
  - MessageType enumeration syntax, 628
- web extensions
  - Relay Server, 239
- web servers
  - configuration options for MobiLink, 267
  - configuring Apache for synchronization, 287
  - configuring for MobiLink (for Redirectors that don't support server groups), 275
  - configuring for MobiLink (for Redirectors that support server groups), 273
  - configuring ISAPI Microsoft for synchronization, 282
  - configuring M-Business Anywhere for synchronization, 289
  - configuring NSAPI for synchronization on Unix, 280
  - configuring NSAPI for synchronization on Windows, 277
  - MobiLink clients, 269
  - MobiLink Redirector, 265
  - synchronizing with MobiLink, 649
- web services
  - synchronizing with MobiLink, 649
- WebLogic
  - MobiLink and, 649
- window (OLAP)
  - glossary definition, 851
- Windows
  - glossary definition, 851
- Windows Mobile
  - glossary definition, 851
- work tables



---

- glossary definition, 851
- worker threads
  - MobiLink, 174
  - MobiLink performance, 170
- writing
  - .NET synchronization logic, 589
  - Java synchronization logic, 527
- writing .NET synchronization logic
  - about, 593
- writing download\_cursor scripts
  - MobiLink, 334
- writing download\_delete\_cursor scripts
  - MobiLink, 335
- writing Java synchronization logic
  - about, 531
- writing scripts to download rows
  - MobiLink, 333
- writing scripts to handle errors
  - MobiLink, 338
- writing scripts to upload rows
  - MobiLink, 330
- writing synchronization scripts
  - SQL, 313
  - supported DBMS scripting strategies, 8
- writing synchronization scripts in .NET
  - about, 589
- writing synchronization scripts in Java
  - about, 527
- writing upload\_delete scripts
  - MobiLink, 331
- writing upload\_fetch scripts
  - MobiLink, 332
- writing upload\_insert scripts
  - MobiLink, 330
- writing upload\_update scripts
  - MobiLink, 331

## **X**

- Xusage.txt
  - location, 92

---